


Implementing SAP R/3 on OS/400

Learn how to use SAP R/3 with the IBM
eServer iSeries server

Explore and implement the best
practices, tips, and hints

Run R/3 on the iSeries faster,
more smoothly, and easily



Aco Vidovic
Monti Abrahams
Ali Ayub
Lisa Gargulak
Michael Power
Marco Wermann



International Technical Support Organization

SG24-4672-03

Implementing SAP R/3 on OS/400

August 2001

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special notices" on page 561.

Fourth Edition (August 2001)

This edition applies to Version 4.6C of SAP R/3 for use with OS/400 Version 4 Release 5.

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, 2001. 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
Part 1. Understanding the solution	1
Chapter 1. Introduction to the iSeries server	3
1.1 iSeries success factors	3
1.1.1 IBM server positioning	6
1.2 iSeries architecture	7
1.2.1 Technology independence	8
1.2.2 64-bit RISC technology	9
1.2.3 Hierarchy of microprocessors	10
1.2.4 Object-based architecture	11
1.2.5 Storage	11
1.2.6 Integrated relational database (DB2 UDB for iSeries)	14
1.2.7 Security, user profiles, and authority management	15
1.2.8 Integrated file system (IFS)	15
1.2.9 System management	17
1.2.10 Logical partitioning (LPAR)	21
1.2.11 Work management	23
Chapter 2. Introduction to SAP R/3	25
2.1 Client/server	25
2.1.1 Client/server computing	25
2.1.2 Client process	26
2.1.3 Server process	26
2.1.4 Client/server architectures	26
2.1.5 SAP R/3 client/server architecture	27
2.2 SAP R/3 system	28
2.3 SAP R/3 instance	29
2.4 SAP R/3 work processes, dispatcher, sessions, and modes	30
2.5 SAP R/3 servers	30
2.6 SAP R/3 landscapes	31
2.7 Solution architecture	32
2.7.1 SAP R/3 jobs	32
2.7.2 National language support for SAP R/3	34
2.8 iSeries support for multiple SAP R/3 systems	35
2.8.1 Memory	35
2.8.2 Disk (auxiliary storage)	36
2.8.3 OS/400 new version and release support	36
2.8.4 Coexistence of multiple SAP R/3 systems	37
2.8.5 iSeries server shared facilities	37
2.8.6 Experience to date	38
2.9 Differences between iSeries and UNIX implementation	38
2.9.1 Memory management	38
2.9.2 Database	39
2.9.3 System level	41
2.9.4 Authority	41
2.9.5 Character sets	42
Chapter 3. SAP R/3 system landscapes on iSeries	43
3.1 Two-tier landscape	43

3.2	Three-tier landscape	43
3.3	Multi-tier landscape with mySAP.com and ITS	44
3.4	SAP R/3 landscapes with logical partitioning (LPAR)	47
3.4.1	LPAR benefits	48
3.4.2	R/3 development, test, and production system.	48
3.4.3	e-business applications with SAP R/3	49
3.4.4	Other applications with SAP R/3	49
3.4.5	Backup system with test system	50
3.4.6	Multiple SAP R/3 application servers	50
3.4.7	Other LPAR scenarios	53
Chapter 4.	Sizing and capacity planning	55
4.1	Sizing terminology	56
4.2	SAP R/3 benchmarks	57
4.3	Sizing and capacity planning approaches	58
4.4	Sizing materials, processes, and contacts	60
4.4.1	IBM sizing support.	60
4.4.2	IBM SAP Capacity Planning Service Offering.	60
4.4.3	IBM Performance and Testing Support/Analysis Services	61
4.4.4	SAP Quick Sizer	61
4.4.5	Sizing the IBM solution	62
4.4.6	Suggestions for disk sizing	63
Part 2. Implementation		65
Chapter 5.	Installation and configuration.	67
5.1	Hardware and software requirements	67
5.1.1	iSeries hardware requirements	67
5.1.2	iSeries software requirements	67
5.1.3	Front-end requirements	68
5.2	TCP/IP configuration	68
5.2.1	TCP/IP configuration on the iSeries server	69
5.3	SAP R/3 installation concepts	74
5.3.1	SAP R/3 directory structures	74
5.3.2	Sample configuration	76
5.3.3	Objects created during the R/3 installation.	78
5.4	Installation overview	80
5.4.1	Preload option	80
5.4.2	Installation tasks	81
5.4.3	SAP R/3 online help installation.	85
5.4.4	National language support (NLS) considerations	86
5.4.5	System date and time settings	87
5.5	SAP R/3 menus	88
5.6	Software upgrades	89
5.6.1	OS/400 upgrade	90
5.6.2	SAP R/3 kernel upgrade	90
5.6.3	SAP R/3 database upgrade	90
5.7	Configuring SAPNet	90
5.7.1	X.25 connection.	91
5.7.2	Frame relay connection to SAPNet	107
Chapter 6.	National language support	113
6.1	Single-byte character set (SBCS) languages	113

6.1.1	Installing a non-Latin-1 SBCS language	114
6.1.2	Importing the language	116
6.1.3	Language possibilities with EBCDIC SAP	116
6.2	Double-byte character set (DBCS) languages	116
6.2.1	ASCII application support on the iSeries server	117
6.2.2	Installing an SAP R/3 DBCS system	118
6.2.3	Importing the language	121
6.3	Multiple language support (MDMP) in one SAP system	122
Chapter 7	Connectivity	123
7.1	Introduction to OptiConnect for iSeries	123
7.1.1	OptiMover for AS/400 (5799-FWQ)	123
7.2	Gigabit Ethernet support	128
7.3	TCP/IP	129
7.3.1	Performance tips	129
7.3.2	TCP/IP improvements	130
7.3.3	Integrated virtual private network (VPN)	134
Chapter 8	Data porting	135
8.1	Concept of data porting to SAP R/3	135
8.2	Writing programs for data porting to SAP R/3	136
8.2.1	Data transfer program	136
8.2.2	Batch input program	137
8.3	Data porting services	139
8.4	Data porting tools	139
8.4.1	Legacy System Migration Workbench (LSMW)	139
8.4.2	MIDAS400	147
8.4.3	R/3-KOMPAKT	150
Chapter 9	R/3 system printing	151
9.1	Introduction to R/3 system printing	151
9.2	TemSe data	153
9.3	The spool work process	153
9.3.1	Placing the spool work processes on the database server	154
9.3.2	Placing the spool work processes on an application server	155
9.4	Spool administration	156
9.4.1	Components of a printer definition	156
9.4.2	Output devices	157
9.4.3	Device types	159
9.4.4	Format types	162
9.4.5	Print controls	164
9.4.6	SAP characters and character sets	165
9.5	Example of configuring a new device to the R/3 system	166
9.6	Special definitions for barcode printing	173
9.7	Spool request versus spooled files	179
9.8	Managing R/3 spooled and output requests	181
9.8.1	Spool request actions	182
9.9	iSeries printer commands	182
9.10	Advanced Function Presentation (AFP) printing with R/3	183
9.10.1	Concept	183
9.10.2	Requirements	185
9.10.3	AFP resources	186
9.10.4	How the AFP driver for R/3 works	186
9.10.5	Access method and device type for AFP printers	188

9.10.6	Using multiple overlays with CVTPRTDTA	191
9.10.7	Setup to support the new OTF user fonts	199
9.10.8	Setup to support a new OTF user barcode	201
9.11	LAN-attached printers	202
9.11.1	LAN-attached IPDS printers	203
9.11.2	LAN-attached ASCII printers	207
9.11.3	Creating a remote output queue	210
9.12	Resolution tips for printing problems	212
9.12.1	General tips	212
9.12.2	Access methods using CVTPRTDTA	212
Chapter 10.	iSeries system administration using GUI	215
10.1	Operations Navigator	215
10.1.1	Database management	216
10.2	Management Central	217
10.2.1	Performance monitoring and collection	219
10.2.2	Examples of Management Central usage with SAP R/3	220
Chapter 11.	Availability, backup, and recovery	221
11.1	Availability considerations	221
11.1.1	Terminology	221
11.1.2	Scheduled outages	222
11.1.3	Unscheduled outages	222
11.1.4	Availability solutions for unscheduled outages	223
11.1.5	Availability solutions for unscheduled outages in R/3 environment	227
11.2	Save and restore commands and menu options	227
11.2.1	OS/400 save and restore commands	230
11.3	Save strategies	232
11.3.1	SAP R/3 libraries	232
11.3.2	SAP R/3 IFS structure	232
11.3.3	What needs to be saved	233
11.3.4	Saving logical partitions	234
11.4	Backup considerations	236
11.4.1	Backup requirements	236
11.4.2	Backup media options	236
11.4.3	Before you begin	237
11.5	Backup	239
11.5.1	Backup methods	239
11.5.2	Initializing the tape	240
11.5.3	Offline backup	240
11.5.4	Online backup with disconnect from the database	250
11.5.5	Online backup without disconnect	255
11.5.6	Saving journal receivers	260
11.6	Recovery	263
11.6.1	User ASP overflow recovery	263
11.6.2	Restoring the entire system	263
11.6.3	Restoring the SAP R/3 environment	264
11.6.4	Recovering the R/3 database	265
11.7	High availability	272
11.7.1	Solution discussion	272
11.7.2	Business partner solutions	273
11.7.3	Concepts	273
11.7.4	Minimizing backup and recovery windows	279

Part 3. Advanced topics	281
Chapter 12. Coexistence with non-SAP R/3 applications	283
12.1 Considerations	283
12.2 Example programs	285
12.2.1 RPG/400 example	285
12.2.2 ABAP example	286
12.3 Accessing SAP R/3 data using an RPG/400 program	288
12.4 Accessing non-SAP R/3 data using ABAP programs	289
12.4.1 Accessing non-SAP R/3 data through the ABAP Dictionary	290
12.4.2 Accessing iSeries stream files using ABAP programs	292
12.5 SAP R/3 processing OS/400 commands	293
12.5.1 ABAP system command interface	294
12.5.2 ABAP program processing CL streams	295
12.5.3 Working with OS/400 commands	297
12.5.4 External command interface for SAP R/3 background jobs	299
12.6 iSeries jobs starting SAP R/3 applications	302
12.6.1 The SAPEVT command	302
12.6.2 The STRREPORT command	304
12.7 Interactive program communication	304
12.7.1 Using the CPI-C connection	305
12.7.2 Using RFC connection	312
Chapter 13. MQSeries link for R/3	323
13.1 Introduction and scope	323
13.2 Customer requirements and situations	324
13.3 MQSeries overview	326
13.3.1 How messaging and queuing works	326
13.3.2 An example of a distributed application using MQSeries	327
13.3.3 Data integrity and resource protection	328
13.3.4 Message attributes	328
13.3.5 Triggering	329
13.3.6 Message sizes	329
13.3.7 MQSeries products	329
13.4 ALE overview	330
13.4.1 Concepts	330
13.4.2 Data distribution	330
13.4.3 Outbound processing	330
13.4.4 Communications	331
13.4.5 Inbound processing	331
13.5 How the MQSeries link for R/3 works	331
13.5.1 Components of the MQSeries link for R/3	332
13.6 Installing MQSeries link for R/3	337
13.6.1 Installing the MQSeries link for R/3 on the iSeries server	337
13.6.2 Three stages of configuration	338
13.7 Ordering and installing MQSeries link for R/3	339
13.8 Benefits of using MQSeries link for R/3	340
Chapter 14. Migration to another platform	343
14.1 System copy	343
14.2 Migration tools	343
14.3 Sap Migration Service	343
14.4 Requirements	344

14.5	Preparation	345
14.6	System migration testing	345
14.7	Final system migration.....	346
14.8	SAP R/3 license.....	347
Chapter 15.	Change and transport system	349
15.1	SAP R/3 in a three-system landscape	349
15.2	Global transport directory	352
15.2.1	Setting up the Transport Management System.....	354
15.2.2	Defining the transport domain and transport routes	355
15.3	Customizing and development process in an R/3 system.....	357
15.3.1	Client attributes	357
15.3.2	Tasks and change requests	358
15.3.3	Transporting change requests	358
15.4	Example of creating a repository object change request	359
15.4.1	Transport Organizer.....	360
15.4.2	Strategies for importing requests	360
15.4.3	Importing single requests using STMS.....	360
15.4.4	Importing single requests using TP	361
15.5	Transporting objects between SAP systems on different host systems	362
15.5.1	iSeries-only environments	362
15.5.2	Heterogeneous environments (NFS-connected machines).....	362
15.5.3	Heterogeneous environments (no NFS available)	363
15.6	Testing the Transport Management System.....	364
Chapter 16.	Performance management	365
16.1	Introduction to performance.....	365
16.1.1	Performance concepts.....	365
16.1.2	iSeries work management concepts.....	370
16.2	SAP memory management concepts	371
16.2.1	Early implementations	372
16.2.2	Later enhancements	374
16.2.3	SAP memory management (iSeries)	376
16.3	A approach to performance analysis	378
16.4	iSeries system monitoring	380
16.4.1	iSeries performance indicator guidelines	380
16.4.2	OS/400 commands versus SAP R/3 transactions.....	382
16.4.3	OS/400 system values.....	386
16.4.4	OS/400 system status and disk status	388
16.4.5	SAP R/3 system on iSeries: Snapshot information.....	392
16.4.6	SAP R/3 system on iSeries: Statistics from previous hours	395
16.4.7	SAP R/3 system on: Performance database.....	398
16.4.8	Active jobs on the iSeries server	399
16.4.9	OS/400 system log.....	404
16.4.10	SAP R/3 system on iSeries: Additional functions	406
16.4.11	Collecting iSeries performance data.....	408
16.4.12	iSeries Performance Tools.....	409
16.4.13	Insight SAP performance reports	412
16.5	iSeries system tuning	412
16.5.1	Initial iSeries tuning	412
16.5.2	Expert cache	429
16.5.3	Maximum active threads per memory pool.....	430
16.5.4	Run priority	431

16.5.5	Separate memory pool for an SAP R/3 instance	433
16.5.6	Network communication	434
16.5.7	Installing additional SAP R/3 instances	435
16.5.8	Temporary storage	436
16.5.9	Logon load balancing	437
16.6	SAP R/3 application performance	437
16.6.1	SAP R/3 monitoring	438
16.6.2	Database monitor	459
16.6.3	ST04 Database performance	461
16.6.4	ST05 SQL trace (ABAP level)	472
16.7	SAP R/3 tuning	472
16.7.1	SAP R/3 buffers	473
16.7.2	SAP work processes	473
16.7.3	File reorganizing	473
16.7.4	Build required indexes	474
16.7.5	Table locks	474
16.7.6	Long running job	474
16.7.7	Resource-intensive functions	474
16.7.8	Deleting SQL packages	474
16.7.9	Short dumps	475
16.7.10	Reporting performance problems	475
16.8	LPAR performance	475
16.8.1	Virtual OptiConnect and DMA	476
16.8.2	Performance tests	477
Chapter 17.	Access Builder for SAP R/3	479
17.1	Overview	479
17.2	Using SAP business objects and BAPIs	480
17.3	Accessing the SAP system	480
17.4	Logging on to an SAP system	481
17.5	Business Object Repository (BOR)	481
Chapter 18.	mySAP.com	483
18.1	SAP Workplace	483
18.2	SAP R/3	485
18.3	SAP Advanced Planner and Optimizer (APO)	485
18.4	SAP Business Information Warehouse (BW)	486
18.5	SAP Business-to-business Procurement (BBP)	486
18.6	SAP Corporate Finance Management (CFM)	487
18.7	SAP Customer Relationship Management (CRM)	487
18.8	SAP Environment, Health, and Safety (EH&S)	488
18.9	SAP Knowledge Management (KM)	488
18.10	SAP Strategic Enterprise Management (SEM)	488
18.11	Internet Business Framework Architecture	489
18.12	SAP Internet Transaction Server (ITS)	489
18.13	SAP Business Connector Interface (BC)	489
Chapter 19.	SAP R/3 and Domino connection	491
19.1	Lotus Domino Connector for SAP R/3	491
19.2	Integration of OS/400, Lotus Domino, and SAP R/3	492
19.3	Architecture	492
19.4	Benefits of Lotus Connector Lotus Script Extensions (LC LSX)	492
19.5	Sample scenarios for Lotus Connector Lotus Script Extensions	493
19.5.1	Scenario 1	493

19.5.2 Scenario 2	493
19.6 Domino Access to SAP R/3 business workflow	495
19.7 Domino Mail Transfer Agent (MTA) for SAP R/3 R1.5	495
19.8 Further information	496
Chapter 20. Using an IBM Network Station as an SAP front end	497
20.1 IBM Network Station: Basic description	497
20.1.1 Introducing various scenarios	498
20.1.2 WTS with separate PC Server	499
20.1.3 Windows-Based Terminal Server on the Integrated xSeries Server	500
20.1.4 Java SAP GUI	500
20.2 Further information	501
Chapter 21. Problem determination	503
21.1 Implementation of SAP R/3 on the iSeries server	503
21.1.1 Job structure	503
21.2 Working with job logs	505
21.2.1 Changing the job attributes	505
21.2.2 Work process overview	505
21.2.3 The WRKPID command	507
21.2.4 The database lock monitor (DB01)	508
21.2.5 The WRKACTJOB command	509
21.2.6 Printing and locating the job log	509
21.3 R/3 profiles	510
21.4 Trace and log files	510
21.4.1 System log (SM21)	511
21.4.2 Short dumps (ST22)	511
21.4.3 Developer traces (ST11)	511
21.4.4 SQL trace (ST05)	511
21.4.5 Startup log	511
21.4.6 Upgrade logs	511
21.4.7 Transport logs	512
21.4.8 XDA trace	512
21.5 Problem analysis	515
21.5.1 Where to look first	515
21.5.2 Database error	517
21.5.3 Performance problems	518
21.5.4 IFS problems	520
21.5.5 Printing problems	520
21.5.6 OptiMover/400 or OptiConnect problems	520
21.5.7 Unable to start R/3	521
21.6 Reporting the problem	522
21.6.1 R/3 environment	523
21.6.2 OS/400 environment	523
21.6.3 Saving spooled files	524
21.6.4 Reporting the problem to SAP	524
21.7 Additional information	524
Part 4. Appendices	525
Appendix A. OS/400 user tools	527
A.1 Edit File (EDTF)	528
A.2 Display File (DSPF)	530

A.3	STRASPBAL, ENDASPBAL, and TRCASPBAL	532
A.4	SAVDLTRCV and SAVDLTRCVE	532
A.5	CRTSAPUSR	538
A.6	CPYTOSTMF and CPYFRMSTMF	538
A.7	SAVR3SYS	540
A.8	STARTSAP and STOPSAP	540
A.9	SQL Utility (SQLUTIL)	542
Appendix B. Apply Journalled Changes Extended		545
B.1	Example of recovering a database with APYJRNCHGX PRPQ	545
B.1.1	Original save	545
B.1.2	Second save	546
B.2	Planning the recovery	547
B.2.1	Finding the starting receiver	547
B.2.2	Finding the ending receiver	549
B.2.3	Finding the starting journal entry	550
B.2.4	Finding the ending journal entry	550
B.3	Looking inside the journal	551
B.3.1	Counting the record level changes	551
B.4	The APYJRNCHGX command	552
B.5	Performance hints	552
Appendix C. Support for SAP R/3		553
C.1	Marketing and technical support	553
C.1.1	Defect support	553
C.2	Competence centers	553
C.2.1	European Competence Center	553
C.3	North American Competence Center contact	554
C.4	Latin America contacts	554
C.5	Asia Pacific Competence Center contacts	554
C.6	Africa and Middle East Competence Center contacts	555
C.7	IBM SAP International Competence Center (ISICC) support	556
C.8	Regional support	556
C.8.1	ISICC InfoService (Walldorf)	557
C.8.2	Philadelphia IBM SAP Competency Center outline	557
C.8.3	SAP regional support	557
C.9	Information access	558
C.9.1	SAP Service Marketplace	558
C.9.2	SAPNet - R/3 Frontend	558
C.9.3	iSeries Informational APARs	558
C.9.4	iSeries Technology Solutions Center	558
C.9.5	SAP Support Services	559
Appendix D. Special notices		561
Appendix E. Related publications		565
E.1	International Technical Support Organization publications	565
E.2	Redbooks on CD-ROMs	565
E.3	Other publications	566
E.4	SAP documentation	568
E.5	Web sites	568

How to get IBM Redbooks	571
List of abbreviations	573
IBM Redbooks review.....	585

Preface

This IBM Redbook features a collection of knowledge gained from IBM and SAP solution experts who work with customers that use SAP R/3 on the IBM @server iSeries server. It was written to assist R/3 basis consultants and other IT professionals in implementing a total business solution consisting of iSeries and AS/400 servers, OS/400 Version 4 Release 5 (V4R5), SAP R/3 Release 4.6C, DB2 UDB for iSeries database, and complementary solution products.

The primary content of this redbook is divided into three parts:

- Part 1, “Understanding the solution”, presents the concepts and other basic knowledge necessary to understand the structure, features, and functions of the SAP R/3 solution on the iSeries server.
- Part 2, “Implementation”, describes the implementation techniques necessary to install and properly set up R/3 in the iSeries environment. It contains detailed guidance and explanations of the specific tasks associated with the implementation. Professionals involved in implementing R/3 on OS/400 may, at some stage, face all the topics covered in this part.
- Part 3, “Advanced topics”, covers topics that will be of interest to those who want to enhance their SAP R/3 installation by improving performance and adding additional functionality.

Note

Throughout this redbook, the name “iSeries” refers to both the iSeries and AS/400 servers.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Rochester Center.

Aco Vidovic is an SAP Certified Basis Consultant and Senior I/T Specialist in IBM ITSO Rochester Center on assignment from IBM Slovenia. Before joining the ITSO in 1999, he worked in the AS/400 area for 10 years in a variety of roles from second-level technical support to marketing. In the ITSO, Aco teaches extensively and writes Redbooks about ERP solution implementation on the iSeries server.

Monti Abrahams is an iSeries IT Specialist and SAP R/3 Basis consultant at IBM South Africa. He has 10 years of experience with AS/400 system development, two years of experience with SAP R/3 production planning (PP)-module configuration and implementation (on the AS/400), and three years of experience with SAP R/3 Basis on the AS/400 server. Monti currently supports SAP R/3 iSeries installations in South Africa and Namibia.

Ali Ayub is the Mid-Range Solutions Manager with ea consulting, Inc., an IBM, Lotus and SAP Business Partner based in Folsom, California. Prior to joining ea consulting, Inc, Ali worked for IBM, where he was responsible for system and product management for the IBM Mid-Range Market. He has more than 10 years of IT experience with many IBM iSeries-based projects.

Lisa Gargulak is an iSeries IT Specialist and a SAP-Certified Basis Consultant in IBM Rochester, Minnesota. She has been with IBM for five years in the Rochester Support Center. She has spent the last three years as a member of the ERP team focusing SAP customer support on the AS/400 and iSeries servers.

Michael Power is an SAP-Certified Basis Consultant with the IBM Business Partner, Data#3 Limited, in Australia. He has eight years of experience in working with the AS/400 and iSeries servers. Currently he provides technical consulting and support for SAP customers on the iSeries platform in Australia.

Marco Wermann is an IT Specialist and SAP-Certified Basis Consultant for IBM. He is currently on assignment at SAP in Walldorf, Germany. He joined IBM in 1996 and performs technical support on the iSeries platform. In 1999, Marco became involved in SAP R/3 support for customers in EMEA.

The authors of the first three editions of this redbook are:

Jaejin Ahn
Frank Benson
Sally Broughton
Manfred Engelbart
Randy Grimm
Willy Guenther
Jacques Hofstetter
Yessong Johng
Walter Lang
Latief Mazema
Peter Mullner
Lloyd Perera
Ivo van Proosdij
Adhie Rachman
Gottfried Schimunek
Frank Zimmer

Thanks to the following people for their invaluable contributions to this project.

Christian Bartels
Lilo Bucknell
Dave Christenson
Brian Clark
Mark Diez
Mike Faunce
Dave Hubka
Dan Moravec
Osamu Nakayama
Ron Peterson
Chris Place
Tom Poturica
Nitin Raut
Dan Sanders
Ron Schmerbauch
Mike Snyder
Chang Wang
IBM Rochester

Melanie Phares
IBM Boulder

Frank Zimmer
IBM Germany

Carmen Coll Ibañez
Christian Goldbach
Volker Gueldenpfennig
Bernhard Wolf
SAP Walldorf

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 “IBM Redbooks review” on page 585 to the fax number shown on the form.
- Use the online evaluation form found at ibm.com/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Part 1. Understanding the solution

This part contains concepts and other basic information about R/3 and the iSeries server that are required during the pre-installation phase. You should be familiar with the topics covered in this part before you start installing SAP R/3. This part includes the following chapters:

- Chapter 1, “Introduction to the iSeries server” on page 3, provides an overview of the iSeries and AS/400 platform. It includes key concepts, architecture, technologies, and functions. Read this chapter if you are new to the iSeries and AS/400 area.
- Chapter 2, “Introduction to SAP R/3” on page 25, presents an overview of some basic SAP R/3 concepts that are necessary for you to understand the iSeries implementation, as well as topics that are discussed in later chapters. If you are new to the R/3 on iSeries arena, be sure to read this chapter.
- Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43, discusses R/3 landscapes in relation to the iSeries environment.
- Chapter 4, “Sizing and capacity planning” on page 55, offers tips, techniques, and contacts that you can use for proper system sizing.

Chapter 1. Introduction to the iSeries server

Note

Throughout this redbook, the name “iSeries” refers to both the iSeries and AS/400 servers.

The IBM @server iSeries server is a follow on system to the AS/400 server. While there are differences in hardware features supported by these two systems, they both run the same operating system, Operating System/400 (OS/400), which distinguish them from other server platforms.

One of the key factors that differentiates the iSeries server from other systems is the level of hardware and software integration. For example, in a UNIX environment, the customer is required to select software components from different vendors (operating system, relational database, security, system management software, and so on) and *integrate* them in order to build a robust working environment for their business applications.

The iSeries server, on the other hand, is an integrated system that offers an out-of-the-box, ready-to-run approach. OS/400 includes a complete range of “licensed programs” (middleware) that offer the highest level of integration. By reducing the extent of integration required to be performed during implementation, the iSeries server approach minimizes implementation costs and increases reliability. The iSeries server also provides a customer the highest level of “ease-of-use” in today's market. The initial ease of implementation and the on-going ease of use, combined with its reliable integration, make the iSeries server a high-performing, high availability, and low-cost business solution.

1.1 iSeries success factors

The iSeries server has a long and successful history worldwide. By the end of last millennium, more than 700,000 iSeries servers had been shipped to over 150 countries. There are also more than 30,000 business applications worldwide, including over 2,700 for e-business. The reason for this success is founded in six basic factors:

- **Architecture:** The iSeries server has a layered architecture that is divided into the actual user interface (OS/400) and a Technology Independent Machine Interface (TIMI). This architectural concept has allowed the iSeries server to undergo several fundamental technology changes, while protecting the customer's investment in information technology. The *transparent* migration of user applications to 64-bit RISC-based hardware and middleware done in 1995 is an example of the benefit of the iSeries architecture. Refer to 1.2, “iSeries architecture” on page 7, for more information.
- **High level of integration:** The iSeries server offers the highest integration of both its hardware and software components. Hardware, microcode, the operating system, and IBM middleware are tightly interlaced allowing maximum exploitation of all available computing resources. Integration of input/output processors (IOPs) and direct access storage devices (DASD) yields valuable benefits, such as an extremely high level of reliability and availability. Other benefits from this kind of integration are proactive. Some of

these benefits include remote hardware and software maintenance, or detailed and informative diagnostic aids for problem determination and monitoring and future projection of system performance.

Some of the features that have been integrated into the iSeries server are:

- System availability:
 - Battery backup unit (BBU)
 - Continuously powered mainstorage (CPM)
 - Uninterruptible power supply (UPS)
 - Protection against system failures
 - Backup of database servers and systems
 - Mirroring and RAID-5 (both at minimal performance cost)
 - Menu-driven backup and recovery
 - Journaling
 - Commitment control
 - Auxiliary storage pools (ASPs)
 - Save while active function
 - Clustering support
 - Logical partitioning
- Standard ease-of-use functions for:
 - System customization
 - Automatic procedures, startup programs, and so on
 - System values
 - System tuning (managing memory, disk)
 - Graphical user interface (GUI) to many system functions
- Database management system integrated with the operating system:
 - No additional cost for database software
 - Integrated database administration tools and automated self-managing database administration functions
 - Excellent performance through microcode-embedding, fine tuned with hardware and OS/400
 - Support for parallel database operations, symmetric multiprocessing (SMP), and parallel I/O processing
 - DB2 Universal Database (DB2 UDB) for iSeries supports the storing, managing, and indexing of all forms of information, including binary objects such as spreadsheets, word processing documents, and multimedia objects
- Easy system management through:
 - Easy hardware configuration and reconfiguration
 - Autoconfiguration of devices
 - Integrated file system accessible through a standard iSeries interface (including the PC and UNIX file system)
 - Minimal database installation, management, and operations activity
 - Simple menu-driven functions
 - Balancing data across disk units
 - Native performance optimization of DB2 UDB for iSeries

- **Interoperability:** The iSeries server offers a wide range of communication capabilities and functions that enable the iSeries server to communicate with IBM and non-IBM systems.

The wide range of communications protocols and networks supported by the iSeries server include:

- TCP/IP
- SNA (APPC, APPN, and HPR)
- OSI
- NetBIOS
- IPX/SPX

Note

You can't use IPX/SPX with SAP, even if the iSeries supports it. SAP R/3 must be connected through a TCP/IP protocol.

- Asynchronous Transfer Mode (ATM) LANs
- ISDN Data Link Control (IDLC)
- Ethernet Version 2 or IEEE 802.3
- IBM Token-Ring Network (IEEE 802.5 and 802.2)
- T1/E1/J1 and Fractional T1 networks (high bandwidth)
- FDDI LANs
- Synchronous Data Link Control (SDLC)
- X.25
- X.21
- Asynchronous
- Binary synchronous

- **Client/server capability:** The iSeries server can operate with almost any client in any communication environment. The following *clients* can work with the iSeries server:

- Microsoft Windows 3.1, 3.11, 95, 98, NT, and Windows 2000
- OS/2
- MacIntosh
- UNIX (IBM-AIX, HP-UX, SUN-SPARC)
- Thin clients with Web browsers

The iSeries server can accommodate the Integrated xSeries Server for iSeries (or simply the *Integrated xSeries Server*), an integrated card that fits into the iSeries server box, which is used to run the following *server* environments:

- Microsoft Windows NT and Windows 2000 Server
- IBM Firewall for AS/400
- IBM Warp Server
- Novell NetWare

- **Scalability:** The iSeries server product line covers a wide range of performance capacities. Server Models 2xx, 7xx, and 8xx provide many processor options to choose from based on the performance requirements. At the high end of these servers, the iSeries server 840, with 24 central processors, provides 330 times the performance boost over the smallest model 250. Figure 1 shows the performance span of the newest iSeries Models 250, 270, 820, 830, and 840, measured in Commercial Processing Workload (CPW).




Figure 1. iSeries scalability

For the SAP R/3 environment, this means the iSeries server can support from a few users to thousands of users. Many of these iSeries server models are field-upgradable to more powerful models. This degree of customer investment protection is exceptional. It is one of the contributing factors to the iSeries server's low cost of operation in the long term.

- **Price and performance:** Many independent analysts have confirmed that the iSeries server represents a cost-effective platform in the long run. The iSeries server's extensive integration yields significant cost advantages, high availability, easy system management, and significant investment protection. This has been the basis for its ten-year success in the dynamic world of information technology. The iSeries server delivers high computing performance at a low cost of ownership and, therefore, scores high in customer satisfaction.

1.1.1 IBM server positioning

The iSeries is one of the four IBM @server product lines on which SAP R/3 can be implemented. The other three servers are the zSeries, pSeries, and xSeries servers. The positioning of the IBM @server family is shown in Figure 2.




Figure 2. IBM @server group positioning

The iSeries is the premier integrated business server that's ideal for small-to-medium-size businesses that want to enter today's sophisticated world of changing information technology, without having to manage the complexity the new world can bring.

With the iSeries, customers are never locked into only one operating system environment. The iSeries advanced architecture can run any combination of UNIX, Windows NT or 2000, Java, Lotus Domino, and OS/400 applications concurrently, exploiting PowerPC and the onboard Intel chip or chips. This gives iSeries customers the unique ability to choose from a large array of leading industry applications, without having to purchase a separate server for each application operating environment. No other server can give you this flexibility.

The high-end solution for large customers is clearly provided by the IBM @server zSeries database server with a DB2 UDB database. The iSeries server, along with the IBM @server pSeries servers (AIX system), covers the medium-to-large customer range. These products also play an important role in the small business area. xSeries servers cover the low-end capacity range for small to medium size customers.

1.2 iSeries architecture

One of the key factors contributing to the commercial success of the iSeries server is its integrated architecture. This section discusses the architectural aspects of the iSeries server.

The iSeries server represents an integrated computing environment. This integration includes the hardware, operating system, and middleware as shown in Figure 3.




Figure 3. iSeries server: An integrated system

IBM engineers designed and built the iSeries PowerPC hardware to integrate with OS/400. OS/400 includes a wide array of middleware components including a relational database, system management facilities, built-in facilities for change management, backup and recovery, and spool or print handling. It also includes a Web server, Web application server, POP3 mail server, support for UNIX applications (PASE (Portable Application Solution Environment)), and client support (Windows clients, MacIntosh, OS/2, and some UNIX clients such as IBM AIX). As a part of OS/400, these middleware products share the same packaging, competitive pricing, user interface, installation, maintenance, security, and management facilities.

The integration of these components by one team (the IBM Rochester lab) ensures that the iSeries server does not need a skilled systems integrator at the customer site. The iSeries server arrives ready to start. This integration is the reason why the iSeries server delivers a low cost of ownership. Studies by such consulting groups as International Data Corporation (IDC) and Emerging Technologies Group reveal that the more integrated a computer system is, the lower the total cost of the solution is.

1.2.1 Technology independence

Technology Independent Machine Interface (TIMI) is a software layer that separates the hardware and System Licensed Internal Code (SLIC) from the operating system (Figure 4).




Figure 4. iSeries advanced application architecture

This permits the instructions used by the compilers of high-level languages to be generic and machine (hardware) independent. These instructions are translated to a specific hardware instruction set as part of the back end of the compilation process. Hardware dependencies are absorbed by SLIC (internal microcode).

The ability of the iSeries architecture to accommodate a seamless transition from CISC-based processors to RISC-based processors, made in 1995, is one example of technology independence. Once the hardware technology was changed, all applications could automatically take *full* advantage of these new hardware capabilities.

1.2.2 64-bit RISC technology

iSeries PowerPC RISC hardware includes 64-bit processors, 64-bit data paths, and 64-bit registers.

OS/400 is a 64-bit operating system, and its instructions perform 64-bit operations. Functional units can sort, compare, load, add, and subtract values that are 64-bit. Storage addressing uses 64-bits. Data paths move data from one location to another, 64-bits at a time (for example, from the cache to the processor). OS/400 middleware is also 64-bit, including the database.

All iSeries server applications exploit the new 64-bit architecture. Older iSeries server applications are automatically migrated to run under 64-bit RISC technology. They do not have to be recompiled or rewritten to support 64-bit addressing, 64-bit data paths, or 64-bit registers.

Complete 64-bit support is a key enabler for delivering performance to demanding applications. Enterprise Resource Planning (ERP) software, such as SAP R/3, requires the movement and analysis of a tremendous amount of information.

Processor technology

The newest iSeries Models 270, 820, 830, and 840 use Pulsar and iStar processors with copper-chip technology. This is the sixth generation of AS/400e PowerPC 64-bit RISC processors. Northstar processors used in earlier AS/400 systems deploy aluminum for on-chip wiring.

Copper's better conductivity permits thinner wires to be used, which enables the transistors to be packed closer together. This new denser technology allows for additional micro architecture methods to improve performance. Denser processor technology also permits more on-chip cache.

iStar processors used in iSeries Model 830s, 840s, and some Model 820s are the first processors in the industry that use the Silicon on Insulator (SOI) processor technology. The addition of SOI alone can increase performance up to 20 to 30 percent beyond the use of copper by protecting the millions of transistors on a chip with a blanket of insulation. This reduces harmful electrical leakage that wastes power.

The transistors are built within and on top of a thin layer of silicon that is on top of an insulating layer. The insulating layer is fabricated by implanting a thin layer of oxide beneath the primary silicon surface of the wafer. This allows the high-end iSeries Model 840 to be 3.6 times faster than the previous high-end AS/400e Model 740.

1.2.3 Hierarchy of microprocessors

In addition to its central system processor (CPU), the iSeries server has a range of other processors, each dedicated to a particular input/output (I/O) device type. A single large iSeries configuration can have well over 200 processors.

When the central system processor encounters a request for data to be read to or written from any I/O device (an operation whose duration is measured in milliseconds (10^{-3} second), rather than nanoseconds (10^{-9} second), which is the unit of time used to measure main storage access times), that request for data is delegated to the particular microprocessor dedicated to that I/O device.

Meanwhile, the CPU continues running another application program. The CPU can actually comprise many separate PowerPC processors. At the time when this redbook was written, the CPU in the high-end iSeries server could have up to 24 separate processors.

This design provides the iSeries server with its outstanding performance in the commercial, transaction-based, environment. The iSeries server is designed for business computing. One of the main characteristics of that environment is that it is I/O-intensive, rather than compute-intensive.

In addition to the benefit of outstanding performance in the business environment, this design gives the iSeries server an elegant method of integrating diverse environments into a single, harmonious customer solution. The microprocessors that look after a particular I/O device are accommodated on I/O cards that fit into slots on the iSeries server bus. One of these cards can be the Integrated xSeries Server. This is a PC on a card and enables the iSeries server to run, for example, Windows NT server. The iSeries server's Internet firewall capability also exploits the Integrated xSeries Server.

1.2.4 Object-based architecture

Objects are the means through which information is stored and retrieved on the iSeries server. This concept is different from typical byte-string and file manipulation in many other systems. Object orientation is part of the architecture and affects both the operating system implementation and high-level language interaction with the system.

As previously mentioned, the TIMI is a boundary (set of instructions) that separates the hardware and LIC from the operating system. The iSeries server instructions have an operation code and operands. Unlike many other systems, the operands are objects, and the instructions act on objects.

Objects have operational characteristics and a defined set of operations that can be performed on them. Objects are addressed through 16-byte pointers (eight bytes are used for a machine address. The other eight bytes are used for information about the object pointed to and for reserved space). In addition to providing addressability to the object, pointers provide access to the associated storage and are used to ensure data integrity and security.

Below the TIMI, LIC provides a tag bit for each *quadword* (16 bytes that must be aligned on a 16-byte boundary) within main storage. This bit is not addressable by the TIMI instructions used to address storage (that is, programs above the TIMI have no direct access to the tag bit). The bit identifies quadwords in storage containing TIMI pointers. The tag bit is turned on by LIC, when a pointer is set, and turned off by the hardware anytime the quadword is modified. This procedure allows the system to detect invalid pointers and prevent illegal use of a pointer. An attempt to subsequently use this data as a pointer results in an exception, and the instruction is not completed. It is not possible to counterfeit a pointer or to modify a pointer in an invalid way. This provides for increased integrity of the machine against intentional or unintentional software actions.

1.2.5 Storage

iSeries storage (memory) design and storage management is another iSeries unique feature. While OS/400 application programs are unaware of underlying hardware characteristics because of the TIMI, they are also unaware of any storage devices characteristics, because of the single-level storage.

1.2.5.1 Single-level storage

In the similar manner as TIMI, the concept of single-level storage delegates the knowledge of the characteristics of underlying hardware devices (in this case, the storage devices – main storage and disk storage) to the SLIC. Programs work with objects. The objects are accessed by name, and never by address. With single-level storage, there is a single, very large address space for *all* storage, including main storage (main memory) and disk storage. Storage is addressed using a 64-bit (8-byte) addressing structure. The iSeries server can address the number of bytes that 64 bits allows it to address, which is 18.4 quintillion bytes.

When an object is created, it is defined in a unique virtual address space. There is a single page table (sometimes referred to as a *page directory*) that maps all virtual addresses to corresponding physical addresses. The benefits of single-level storage are:

- All iSeries applications use one address space. Switching processes (for example, from the database to the Internet) requires little time. There is no need to flush out one address space and bring in another one.
- With single-level storage, there is no need for swap or paging space. When data needs to be moved out of memory, it goes directly to its permanent location on disk. It is not written first to a swap file and then its permanent location. iSeries disk space is not wasted for swap files.
- With single-level storage, all information is in one large address space, regardless of it being in memory or on disk. All objects appear to be in memory. When a program wants to reference data, it simply references it by its address. There is no need for the program to translate addresses or go through a lengthy process to determine where the file exists on disk.

With disk and memory managed through a common address space, the system can take advantage of increased memory without needing any complex memory management activity.

- Single-level storage also holds tremendous promise for object-oriented programs such as Java. With the ability to have shared persistent objects, the iSeries server can eliminate the processing required on other systems to hydrate and dehydrate objects.

Single-level storage is a key reason why it is easy to develop applications for the iSeries server. It is also why the iSeries can support multiple applications and is scaled to support a large number of users. Single-level storage, combined with the automated storage management capabilities, makes the iSeries server an easier system to manage.

1.2.5.2 Storage management

One of the key resources on today's computer system is the disk space that stores the data, applications, and programs. The management of this storage space is critical to having an efficient and high performance server installation.

Storage management on the iSeries server is automated. The iSeries server takes care of selecting the physical disk drive (DASD) to store the data, spreads the data across the disk arms or disk units, and continues to add records to files until specified threshold levels are reached.

The iSeries customer manages the iSeries disk space as a single entity. The disk space can be used by all of the applications. The system determines when and where to extend individual files and spreads the data across the disk arms to maximize performance. The system keeps track of disk usage and warns the iSeries system administrator when the disk space in the system auxiliary storage pool (ASP) reaches a specified threshold value. iSeries server installations can deploy other disk management options to increase control of the space management process.

OS/400 has sophisticated space allocation routines to create and store information in the right-size spaces. These routines look for the right-sized block, as opposed to selecting the first available block, regardless of size. The objective of these routines is to effectively use the disk space by reducing file and free space fragmentation. The iSeries server automatically consolidates disk space fragmentation where possible and also has a separate disk fragmentation utility.

1.2.5.3 Object persistence

Single level storage enables another extremely important iSeries benefit, *object persistence*. Object persistence means that the object continues to exist in the memory system forever. An ordinary machine requires that information be stored in a separate file system if the information is to be shared or if it is to be retained for a long time. The persistence of objects is extremely important for future support of object-oriented databases. Objects need to continue to exist even after their creator goes away. The iSeries server is uniquely positioned to exploit this characteristic of object persistence. Ordinary systems use a less elegant mechanism that requires them to store their persistent objects in a separate file system with all the attendant performance implications.

1.2.5.4 Teraspace storage

The iSeries provides what is known as a single-level storage (SLS) model. All processes share a single, 64-bit address space, which is partitioned into 16 MB (where 1 MB equals 1,048,576 bytes) segments. The segment is the basic unit of storage allocation throughout the system. While the SLS model provides many benefits to the OS/400 in terms of efficiency, scalability, and integrity, the 16 MB segment size may in certain situations represent a limitation. In such cases, larger contiguous storage can be emulated with the program to support main storage intensive applications that are scaled to support more users or larger systems.

Teraspace, introduced in OS/400 V4R4, is a new temporary storage area, which provides up to 1 TB (where 1 TB equals 1024 GB) of private storage to a single process. Teraspace storage overcomes the 16 MB limitation of segmented storage and, therefore, provides intensive applications with large contiguous storage. Teraspace storage also supports memory mapping, used to efficiently access single-level storage objects, such as files, via memory accesses.

OS/400 V4R4 and later provides enhanced C runtime routines for allocating teraspace heap storage of up to 2 GB in a single allocation. The POSIX shared memory manager is also enhanced to provide shared memory objects of up to 4 GB in size.

Teraspace storage is being employed within OS/400 for tasks as diverse as manipulating large performance data collections to efficiently handling large journal entries.

In SAP R/3 releases prior to 4.6A, memory type (ROLL memory, extended memory, HEAP memory, paging memory) was based on SLS segments in OS/400. Starting with release 4.6A, extended memory, HEAP memory, paging memory, and the R/3 buffers are in the teraspace. Only the ROLL memory is implemented directly with SLS.

Some background information

Due to the single-level storage concept on OS/400, each address used has a physical representation in an ASP. When a segment is allocated, temporary storage is assigned to it in the system ASP. When an address is accessed, OS/400 only loads the page (4 KB block) that contains the address and not the entire segment (neither in the SLS nor in the teraspace).

OS/400 uses the 16-byte tagged-pointer concept for the physical representation of pointers. A pointer allocates 16 bytes, plus a tag bit that specifies whether a valid pointer is contained in a certain 16-byte area.

1.2.6 Integrated relational database (DB2 UDB for iSeries)

The integrated relational database system (DB2 UDB for iSeries) is one of the most significant features of OS/400. Relying on a database manager integrated into the operating system means that almost all of the user data on the iSeries server is stored in a relational database. Access to the database is implemented by the operating system itself. Moreover, some of the database functions are implemented below OS/400, in SLIC and hardware to improve performance.

DB2 UDB for iSeries provides stability and compatibility with previous releases of the iSeries database. It also complies with standards coupled with advanced functions, distributed capabilities, and performance. DB2 UDB for iSeries provides the following features:

- Structured Query Language (SQL) standards conformance, which supplies the industry standard database access language conforming to the IBM SQL Version 1, ANSI X3.135.1992, ISO 9075-1992, and FIPS 127-2 standards. Support is provided for embedded static, dynamic, and extended dynamic SQL, IBM Distributed Relational Database Architecture (DRDA), Java Database Connectivity (JDBC), Microsoft Open Database Connectivity (ODBC), and Apple Data Access Language (DAL). A Call Level Interface (CLI) server mode is also provided that allows developers to write applications that do database serving for multiple users.
- Encoded Vector Indexes (EVI) can be created using SQL. EVIs can in many cases improve query performance.
- Declarative referential integrity, which prevents from entering the conflicting data in the database.
- Stored procedures that allow the distribution of application workloads between a client and an application server.
- Triggers that cause automatic program execution before and after database modifications.
- Two-phase commit transaction management to allow access to multiple heterogeneous databases simultaneously.
- Automatic data replication in a distributed DB2 UDB family environment.
- System-wide database catalog allowing applications to query information concerning all objects on a system using a single system catalog.
- Multiple-level concurrency control providing read stability, cursor stability, uncommitted read, and no commit isolation levels.

- National language support to store data in a preferred language, character set (single and double byte), and a sort sequence.

1.2.7 Security, user profiles, and authority management

The iSeries server has an outstanding reputation as a mission-critical commercial server. Part of this reputation is due to the security of the iSeries server. OS/400 has a single integrated security model. This security model is shared by the operating system, database, mail, systems management, the Internet, and communications functions. You can add a user once to OS/400, and they can have access to all the appropriate components. These include Windows NT, OS/2 Warp Server, and Novell NetWare on the Integrated xSeries Server or Lotus Domino running natively on the iSeries server.

OS/400 has a single directory, which is the system distribution directory. This directory is used by OS/400 e-mail and Domino for iSeries.

The iSeries server has a C2 security certification from the U.S. Department of Defense and is the only computing system to have received certification for both a database and operating system as a unit. The iSeries server configuration that was certified was a functional configuration that included non-programmable terminals.

The iSeries server has an object-based system architecture (see 1.2.4, “Object-based architecture” on page 11). Everything is an object (including data, programs, files, and so on) and is subject to object rules. These rules describe what can be read, changed, deleted, or added, and by whom or what. The iSeries server checks those rules before it touches any object. This basic and underlying capability provides iSeries server customers with a powerful security capability within their system. In addition, program objects have special rules that do not allow them to run unless they are properly encapsulated by the system. This makes the iSeries server extremely virus resistant. iSeries server objects, including programs, files, and databases, are protected by the OS/400 security mechanisms. When correctly implemented, this protection cannot be compromised by someone who merely has access to the machine.

System authorization management is based on user profiles that are also objects. All objects created on the system are owned by a specific user. Each operation or access to an object must be verified by the system to ensure the user's authority. The owner or appropriately authorized user profiles may delegate to other user profiles various types of authorities to operate on an object. Authority checking is provided uniformly to all types of objects within a file system.

The object authorization mechanism provides various levels of control. A user's authority may be limited to exactly what is needed.

1.2.8 Integrated file system (IFS)

To enable the iSeries server as a *server-of-choice* for all the major clients, it was necessary to implement each client's file system on the iSeries server as natively as possible.

The integrated file system (IFS) is part of the iSeries server that supports input, output, and storage management similar to the PC and UNIX operating systems.

At the same time, it provides an integrated structure over all information stored in the iSeries server. The key features of the IFS are:

- Support for information stored in stream files
- A hierarchical directory structure
- An integrated interface to access stream files, database files, documents, and other objects stored in the iSeries server

A diagram of the IFS is shown in Figure 5.




Figure 5. Integrated file systems

The integrated file system on the OS/400 can be accessed via native OS/400 commands and also from the graphical user interface of Operations Navigator. Figure 6 shows the Operations Navigator IFS access interface and functions.




Figure 6. Operations Navigator IFS interface

For more information on IFS, refer to *OS/400 Integrated File System Introduction Guide*, SC41-5711.

1.2.9 System management

A number of system management functions are integrated with OS/400. They are key to lowering the cost of ownership of the iSeries server solution. These functions include:

- **Job management:** A job is a unit of work in OS/400. OS/400 includes an environment to manage jobs for users, operators, and programmers. An operator can look at all of the jobs running on the system, or select only those associated with a specific queue or user. An operator can sort them by CPU usage, view their properties, change their priority, stop them, or delete them within the limitations of authority granted to the operator.
- **Printer management:** The iSeries server has a print output environment. An operator can select what output to see all of it by printer, user, or queue. With the printer management facilities, an operator can route the output to a different printer, change the number of copies, print only selected pages, or print duplex copies. iSeries server print jobs can also be sent or mailed to other iSeries servers and users. Plain text print output can be viewed on the iSeries server before being printed. When using Advanced Function Printing (AFP), a user can see the actual iSeries printed output (text, forms, or graphics) before it is printed. For more information, refer to Chapter 9, "R/3 system printing" on page 151.
- **Backup and recovery:** OS/400 includes a backup scheduling application that allows an administrator to easily create a daily, weekly, and monthly backup plan along with other scheduled tasks. The administrator selects what to back up, where to back it up, and when to back it up. For more information, refer to Chapter 11, "Availability, backup, and recovery" on page 221.

- **Managed availability:** OS/400 V4R4 introduced iSeries Logical Partitioning (LPAR) to enhance the role of the iSeries server as a consolidated server. LPAR provides the ability to run multiple independent OS/400 instances or partitions (each with its own processor, memory and disks, in n-way symmetric multi-processing iSeries). With LPAR, companies have both the power and flexibility to address multiple system requirements in a single machine. LPAR is of value to customers for server consolidation, business unit consolidation, mixed production and test environments, as well as integrated clusters. iSeries clustering is taking a major step forward with the introduction of Cluster Resource Services as part of OS/400 V4R4 (APIs). The complexity of managing systems in a cluster and keeping track of data and applications is now handled by OS/400 V4R4. Protecting your business from unplanned and planned outages, as well as site loss disasters, is easier than ever before. Cluster management and enhanced data resilience applications, both provided by high-availability business partners, complete the total solution.
- **User management:** When a new user is added to OS/400, besides entering the user ID and password, an administrator can define the user's environment. This includes the maximum amount of disk space the user can use to store information, the maximum priority at which any of their jobs run, accounting information for tracking iSeries server resource consumption, the output queue and message queue for the user, and the language and sort order for the user. They can also integrate the user ID and password with Windows NT, Novell NetWare, or OS/2 Warp Server running on the Integrated xSeries Server or with Lotus Domino.
- **National language support (NLS):** The iSeries server has national language support for over 50 languages and can support multiple languages, sort sequences, and character sets per system. With this support, one user can interact with the system, for example, in the German language, while another user converses in French. Unicode (UCS-2) support is also available. For more information on national language support, refer to Chapter 6, "National language support" on page 113.
- **Problem management:** A complete problem management application is included with OS/400. This support allows system administrators to track, analyze the causes of, prioritize, and report to IBM the problems associated with the computing environment. The iSeries server keeps track of the installed products, release levels, and program temporary fixes (PTFs) that were applied. For more information, refer to Chapter 21, "Problem determination" on page 503.
- **Software fixes:** Individual software fixes or PTFs can be applied to LIC, OS/400, and licensed programs either permanently or temporarily. The benefit of applying a fix temporarily is that it can be applied, tested, and removed if it is appropriate to do so.

With advanced, built-in management features, the iSeries server is a reliable, secure, and virtually self-sufficient computing system. And the iSeries server comes with all of the services and support for which IBM is traditionally known.

Electronic Customer Support (ECS) is a direct electronic link to IBM marketing, administration, technical, and service operations. ECS gives users and technical staff online advice and provides configuration management assistance, problem determination, and other service needs. ECS simplifies the PTF ordering procedure by enabling the speedy receipt and application of PTFs and quicker

problem resolution. The Copy-Screen facility permits a more accurate user-to-specialist problem description and direct electronic contact with iSeries server specialists. ECS support is provided for software problems (LIC, operating system, licensed programs, and third-party applications), as well as for hardware problems.

In addition to ECS, iSeries server customers under warranty or an IBM maintenance agreement may receive Service Director support. Service Director monitors the iSeries server in real time. When an error is entered into the problem log, it is immediately and automatically analyzed and reported to IBM through the ECS function on the iSeries server. The problem is fed into the IBM remote technical assistance information network, and IBM resources are used to evaluate the situation. If a PTF is available that corrects the problem, it can be downloaded to the iSeries server. This can significantly reduce overall downtime.

In addition to these system management facilities, additional IBM and third-party products are available to manage larger and distributed iSeries server environments. These products provide support for software distribution, client management, performance analysis, capacity planning, and more.

1.2.9.1 Management Central

A suite of system management functions known as Management Central, which is part of Operations Navigator, includes system management functions such as Collection services, object packaging, PTF management for distributed environment, and inventory on multiple systems. These functions are briefly described here:

- **Collection services:** This tool for collecting and managing performance data replaces the traditional performance monitor function with a low-overhead, automated, and on-going data collector. Data is captured with reduced system impact. Processing occurs only if and when needed. In addition, collection services lets you control what data is collected and how that data is managed.

Each type of data supported by collection services can be controlled individually without data loss or affecting the collection of other data.

A compatible performance monitor database is created based on the data in the management collection object. However, you can defer the creation of the database until a later time.

Collecting performance sample data is enhanced by:

- Reducing the impact of collecting performance data, especially on large systems
- Allowing flexibility in the data collected
- Simplifying the management of performance data
- Promoting automated, continuous data collection

- **Object packaging:** The object packaging and distribution graphical interface provides an easy way to send objects from any file system to one or more iSeries servers in a network. You can also restore objects, take snapshots of the objects, version packages of objects, and post execution of commands. All of these functions can be performed on a group or network of iSeries servers and be scheduled to occur at a time that is most convenient for your staff.

- **PTF management for a distributed environment:** If managing PTFs among several iSeries servers is too complicated, the new PTF management wizards are for you. The easy-to-use wizards walk you through comparing the PTF levels of multiple iSeries servers to a model system that has a proven set of PTFs already installed. You then distribute and install any missing PTFs on the remote iSeries servers by simply identifying the system or group of systems to be updated. You can run iSeries commands as part of completing PTF installations or as part of normal day-to-day operations.
- **Inventory for multiple systems:** With the new graphical interface, you can schedule regular inventory collections of hardware, software, and PTF information for a group or network of iSeries servers. From the data collected, you can search for a specific piece of information, export the information to a PC application for analysis, or just compare information in multiple systems.

Refer to Chapter 10, “iSeries system administration using GUI” on page 215.

1.2.9.2 System management facilities

A variety of tools and functions that provide system availability and management are available in the iSeries server. Some are discussed here:

- **System Managed Access Path Protection (SMAPP):** SMAPP supports and automates the process of selecting which access paths should be protected. The system uses the EDTRCYAP value to estimate the amount of journaling to perform. The shorter the time is in this value, the more journaling takes place. This impedes system performance, but leads to shorter IPLs. The longer the value is, the longer IPLs are. However, the impact of journaling on CPU and DASD utilization is less.
- **Expert cache:** Expert cache provides a disk cache tuner option, which allows the iSeries server to take advantage of available main storage capacity. It dynamically responds to system jobs to cache pages of data in main storage to reduce the time to process disk I/O.
- **Integrated hardware disk compression:** Beginning with OS/400 V4R3, the compression of data on disk is supported by OS/400. Data is dynamically compressed and uncompressed by the iSeries DASD controller as data is written to and read from disk. Disk compression does not affect the main CPU utilization since this function is performed by the DASD IOP.
- **Hierarchical Storage Management (HSM):** OS/400 includes HSM APIs that are used by Backup and Recovery Media Services (BRMS), 5769-BR1, to provide HSM functions. These APIs can also be used to develop custom HSM applications. The APIs are documented in *Complementing AS/400 Storage Management Using Hierarchical Storage Management APIs*, SG24-4450. Refer to the following Web site for more information on BRMS and HSM:
<http://www.as400.ibm.com/hsmcomp>
- **PTFs available on Internet:** Beginning with OS/400 V4R3, iSeries customers can download PTFs over the Internet. The client hardware needed is a PC with Windows 95 or Windows NT, a TCP connection to the iSeries server over a LAN, and access to the Internet. The various configurations and setup information are documented on IBM @server iSeries and AS/400 Technical Support site at: <http://as400service.rochester.ibm.com>

Except for the medium of transport (Internet), the functionality is the same as the ECS method of transport. The user selects the PTFs and options using a Web browser and submits the order.

At the IBM @server iSeries and AS/400 Technical Support site, the user can also search on PTF cover letters and read them before the order is even placed. The same entitlement rules that apply on the ECS connection are enforced. In other words, if a user can acquire PTFs electronically over the ECS, they can acquire PTFs over the Internet.

1.2.10 Logical partitioning (LPAR)

Logical partitioning is a means used to create multiple independent systems within the same computer, by splitting the hardware resources of a single iSeries server. The resulting structure consists of a *primary partition* and one or more *secondary partitions*.

Primary partition

When a new system is installed, it is always configured with the Primary Partition only. This partition owns all the resources available on the machine, such as processors, main storage and system buses. It functions as one of the logical systems and also provides management functions for the configuration of the secondary partitions, such as:

- Power management (includes concurrent maintenance)
- Virtual Operations Panel (controls power up or power down, IPL, virtual key lock)
- Logical partition definition (for configuring logical partitions)

The primary partition can also function as a hub for external communications or Electronic Customer Support for the whole system.

Secondary partition

Secondary partitions are created and managed from the primary partition, but function as independent systems within the same physical box. They have their own resources such as processors, main storage, and system buses. And they also have their own primary language, system values, time-of-day, user profiles, OS/400 SLIC, IBM Licensed Programs (LPP), database files, and applications such as SAP R/3. Furthermore, they can be independently powered down and up without affecting the primary or other secondary partitions. However, they retain some dependencies on the primary partition. For example, performing a system restart or power down on the primary partition will power down all secondary partitions. This means that, before these actions are undertaken, all secondary partitions must be powered down to avoid possible abnormal secondary partitions termination.

In OS/400 V4R5 and earlier, each partition has at least one CPU. It also has its own bus infrastructure, memory, disk environment, load source unit, alternate IPL units (either CD-ROM or tape), system console, and IOPs such as LAN, communication, or tape adapters. Figure 7 shows an example of a 12-way iSeries server partitioned into three LPARs: SYSTEMA, SYSTEMB, and SYSTEMC.




Figure 7. 12-way iSeries Logical Partitioning example

SYSTEMA has three processors, owns bus 1 and also uses bus 2. SYSTEMB has five processors and owns two buses, bus 2 and bus 3. The bus 2 is owned and shared while the bus 3 is owned as dedicated. SYSTEMC has four processors, owns bus 4, and also uses bus 2.

The CD-ROM and 3570 tape drive attached to I/O processor IOPxyz can be switched between SYSTEMB and SYSTEMC.

Each logical partition runs its own OS/400, microcode (SLIC), license programs, and applications independently of all other partitions. This can be done with different releases, PTF levels, and system settings (for example, different time zones, languages, and so on).

A system must have main memory. Otherwise, it cannot execute any programs. Main memory is assigned to a partition in 1 MB increments. The minimum main memory required is 256 MB for the primary partition and 64 MB for the secondary partition. The memory assigned cannot be touched by other partition. It is where you run your own SLIC, LPPs, and application programs.

The end objective of LPAR, released with OS/400 V4R4, is to provide users with the ability to split a single iSeries server into several independent systems capable of running applications in multiple, independent environments simultaneously. For example, logical partitioning makes it possible to run an application, such as SAP R/3, on a single system partitioned to provide a 3-tier environment.

For more information on LPAR, refer to *Slicing the AS/400 with Logical Partitioning: A How to Guide*, SG24-5439.

1.2.11 Work management

In the iSeries server operational environment, work is managed by one or more OS/400 *subsystems* and SLIC. All system and user jobs run under the control of a subsystem, while SLIC performs the actual task management according to user-specified and internal priority schemes.

Subsystems are provided to permit job grouping for easier control of CPU and main memory. The assignment of jobs to either a separate or a shared main memory pools can be accomplished within a single subsystem or managed by separate subsystems. User jobs may be assigned to IBM-supplied or user-defined subsystems.

For more information on iSeries work management, refer to *OS/400 Work Management Guide*, SC41-5306.

Chapter 2. Introduction to SAP R/3

Note

This chapter gives an overview of some basics about SAP R/3 to help you understand how to implement it on the iSeries. Some of this information will also help you understand topics that are presented in later chapters. Be sure to read this chapter if you are new to working with R/3 on the iSeries.

SAP is one of the world's largest software vendors. The SAP R/3 system is well suited for companies of all sizes and industries. R/3 products offer a strong combination of functionality, flexibility, and technology allowing companies to respond quickly to change. SAP R/3 applications cover accounting and controlling, production and materials management, quality management and plant maintenance, sales and distribution, human resource management, and project management. Although R/3 is designed as an integrated system, modules can also be used individually.

R/3 provides multi-currency capabilities, dual currency functionality, and exchange rate support. The SAP product suite continues to evolve by incorporating technologies such as Object-oriented (OO), the Internet, and Business Intelligence. For more information on these products, refer to Chapter 18, "mySAP.com" on page 483.

R/3 was first made available on the AS/400 system in the summer of 1996. It runs on all iSeries and AS/400 models, except for the 150 model. Today there are more than 1,000 R/3 installations on the iSeries server.

2.1 Client/server

This section offers a simple overview of the SAP R/3 client/server architecture and its client/server implementation on the iSeries server. Before we examine the SAP R/3 client/server architecture, let's briefly look at the client/server architecture in general.

2.1.1 Client/server computing

Client/server computing is the logical extension of *modular programming*. Modular programming has as its fundamental assumption that separation of a large piece of software into its constituent parts (modules) creates the possibility for easier development and better maintainability. Client/server computing takes this a step further by recognizing that those modules do not need to be run within the same memory space. With this architecture, the calling module becomes the "client" (which requests a service), and the called module becomes the "server" (which provides the service).

The logical extension of this is to have clients and servers running on the appropriate hardware and software platforms for their functions. For example, database management system servers running on platforms specially designed and configured to perform queries or file servers running on platforms with special elements for managing files.

2.1.2 Client process

The client is a process (program) that sends a message to a server process (program), requesting that the server perform a task (service). Client programs usually manage the user-interface portion of the application, validate data entered by the user, dispatch requests to server programs, and sometimes execute business logic. The client-based process is the front end of the application that the user sees and with which it interacts. The client process contains solution-specific logic and provides the interface between the user and the rest of the application system. The client process also manages the local resources that the user interacts with such as the monitor, keyboard, workstation CPU, and peripherals. One of the key elements of a client workstation is the graphical user interface. Normally a part of operating system, such as the window manager, detects user actions, manages the windows on the display, and displays the data in the windows.

2.1.3 Server process

A server process (program) fulfills the client request by performing the task requested. Server programs generally receive requests from client programs, execute database retrieval and updates, manage data integrity, and dispatch responses to client requests. Sometimes server programs execute common or complex business logic. The server-based process "may" run on another machine on the network. This server could be the host operating system or network file server. The server is then provided both file system services and application services. Or in some cases, another machine provides the application services. The server process acts as a software engine that manages shared resources such as databases, printers, communication links, or high powered-processors. The server process performs the back-end tasks that are common to similar applications.

2.1.4 Client/server architectures

The basic characteristics of client/server architectures are:

- Combination of a client or front-end portion that interacts with the user, and a server or back-end portion that interacts with the shared resource. The client process contains solution-specific logic and provides the interface between the user and the rest of the application system. The server process acts as a software engine that manages shared resources such as databases, printers, modems, or high powered processors.
- The front-end task and back-end task have fundamentally different requirements for computing resources such as processor speeds, memory, disk speeds and capacities, and input/output devices.
- The environment is typically heterogeneous. The hardware platform and operating system of the client and server are not usually the same. Client and server processes communicate through a well-defined set of standard application program interfaces (APIs) and remote procedure calls (RPCs).
- An important characteristic of client/server systems is scalability. They can be scaled horizontally or vertically. *Horizontal scaling* means adding or removing client workstations with only a slight performance impact. *Vertical scaling* means migrating to a larger and faster server machine or multiple servers.

2.1.5 SAP R/3 client/server architecture

To manage the complex business needs of today's competitive information technology infrastructure, SAP R/3 offers leading client/server technology. SAP address a wide range of customer requirements and provides a multi-tier client/server architecture.

From a software viewpoint, excluding Internet applications, the architecture of the R/3 system consists of three main layers:

- Presentation
- Application
- Database

From the hardware perspective, these three layers can run separately on different computers or all together on the same computer. R/3 also allows the distribution of the presentation and application layers over multiple computers.

Presentation layer

Usually SAP's graphical user interfaces (SAP GUIs) and other front-end applications at the presentation level run on PCs, one per user. This guarantees that resources on this level are at the disposal of each user, and no single user can be affected by individual behavior of another. The SAP GUI accepts the user input and passes it to the next layer, which is the application layer, for further processing. The SAP GUI accepts data from the application layer and presents this data to the user.

Application layer

User requests are passed from the presentation layer to the R/3 application layer. This is where the actual calculations and evaluations are performed. Data required to run these tasks is requested from the database layer. Incoming data from the presentation layer is processed here and passed to the database layer.

Database layer

The most important part of the database layer is its relational database management system (RDBMS). An SQL interface provides the data exchange between the application processes and the RDBMS.

2.1.5.1 SAP R/3 components

The SAP R/3 system consists of the following components:

- SAP R/3 applications such as financial accounting, sales and distribution, asset management, and so on
- SAP R/3 Basis, which is a middleware that runs below the applications and on top of an operating system such as OS/400
- Hardware and system software provided by other vendors such as IBM

The relationship between these components is shown in Figure 8.




Figure 8. Basic SAP R/3 overview

The hardware and system software platforms supported by SAP R/3 are shown in Figure 9.




Figure 9. SAP R/3 technology environment

Note

This redbook covers topics related to SAP R/3 *Basis* and iSeries server hardware and software.

2.2 SAP R/3 system

An SAP R/3 system includes a kernel and database. The kernel contains the executable code, while the database includes the SAP database and the Advanced Business Application Programming (ABAP) program modules. The R/3

system is serviced by one or more instances. Access to SAP R/3 services is provided through at least one central instance and, optionally, one or more dialog instances.

Multiple SAP R/3 systems may be installed on a *single iSeries server* using a unique three-character system identifier (<SID>), for example, T01, P01, and D01. An SAP R/3 system has one and only one database, which is on the iSeries system and referred to as *SQL collection*. The name of this SQL collection is R3<SID>DATA, where <SID> refers to the system identifier.

A *client* is a way to organize and isolate data in an R/3 system. For example, you may develop and customize your code in one client and have a different client for productive use. SAP ships the database with client 000, 001, and 066.

The R/3 system consists of client-dependent and client-independent data. Client-dependent data applies only to a specific client. An example of client-dependent data is application data. Client-independent data applies to all clients. An example of client-independent data is transaction codes.

In an SAP R/3 system that is just installed, customers normally copy client 001 and make changes in the copied client. Additional client copies may be made to provide for refreshing a client in case of application errors or usage corrupts the data in a client. This is recommended in development and testing, rather than production environment.

2.3 SAP R/3 instance

An SAP R/3 instance is a collection of processes and resources within a single SAP R/3 system that provides SAP resources to end-user requests. An instance is connected to only one R/3 database that is specified when the instance is defined. Multiple instances can be defined for a single SAP R/3 system. The characteristics of an instance are specified by its *instance profile*.

On the iSeries server, an instance is implemented as a *subsystem* with a name *R3_nn*, where *nn* corresponds to the instance number. The instance number is assigned automatically by the system when an instance is created and cannot be altered. An instance becomes active after it is started using the Start SAP (STARTSAP) command. The STARTSAP command has parameters that require you to specify the SAP R/3 system and the instance number to be started. An instance is stopped after the completion of the corresponding Stop SAP (STOPSAP) command.

A *central instance* is an instance that is defined, through its profile, to have a *message server* and an *enqueue server*. A central instance does not depend on another instance to be operational. A central instance has complete information of all SAP R/3 system resources and their status. Each SAP R/3 system has only one central instance.

A *dialog instance* is an instance that is defined without a message server or enqueue server. A dialog instance depends on the central instance and uses the message server and enqueue server under that central instance to be fully operational. A dialog instance does not have complete information of all the resources of an R/3 system, nor the status of those resources.

Often, an SAP R/3 system (SID) has only one instance on each iSeries server. However, *multiple* instances may be configured on each server.

The instance to which the presentation services on the workstation attach is specified in the folder with the SAP icon or on the SAPLOGON menu.

2.4 SAP R/3 work processes, dispatcher, sessions, and modes

SAP R/3 work processes on the iSeries server are jobs within the R3_nn subsystem that actually perform the work.

Each work process is assigned a primary role by the dispatcher. The dispatcher is one of the jobs in the subsystem that controls the type of work that is performed by that work process. End-user requests and units of work are assigned by the dispatcher to the work processes of the instance for processing.

For example, a work process that is designated as a dialog process can handle end-user requests from the presentation services. Each instance has a single dispatcher that manages the workload of the instance. Presentation services interact with the dispatcher. The number of work processes and the types that exist for an instance are determined by the instance profile and can be controlled from within the SAP R/3 system by Computing Center Management System (CCMS).

The *gateway job* is responsible for communications between an instance and jobs outside the instance. For example, the gateway is used to communicate between a work process and a program that is running external to the SAP R/3 environment.

The *enqueue work process* is a special job that is responsible for handling the special locking requirements of SAP R/3. The enqueue process handles logical locks. There can be more than one enqueue work process if necessary (SAP note 127773). All of them are in the central instance.

The *message server* is a special job within a central instance. It facilitates communications between instances within the same SAP R/3 system and identifies services in the SAP R/3 system to external jobs. There is only one message server for each SAP R/3 system.

Sessions are initiated to support specific user transactions. A session corresponds to a primary window on the end-user's display.

Modes corresponds to the nesting of displays within a window.

2.5 SAP R/3 servers

An SAP R/3 in iSeries environment consists of the following servers and services:

- **Presentation services:** The part of SAP R/3 that interacts with the end user. It provides the graphical user interface, usually running on the PC. Sometimes it is referred to as *SAP GUI*. The presentation services interact with an appropriate server instance of SAP R/3 to carry out end-user requests.
- **Database server:** An iSeries server on which the SAP R/3 database is stored. It contains the support required to access that database through an instance.

- **Application server:** An iSeries server with an SAP R/3 instance that is ready to service requests from the presentation services. It works in conjunction with the database server to provide services to the end user.

An SAP R/3 system has one database server and one or more application servers.

2.6 SAP R/3 landscapes

Several SAP R/3 servers and clients that access these servers make up an SAP R/3 landscape. There are three types of SAP R/3 landscapes:

- Two-tier
- Three-tier and
- Multi-tier landscape

The two-tier and three-tier landscapes are shown in Figure 10.




Figure 10. Two-tier and three-tier configurations

In a *two-tier* landscape, a single iSeries server functions as the application server and database server. From an R/3 viewpoint, this is a stand-alone computer that does not need to interact with another computer to service the R/3 end-user requests coming from the presentation services layer. This configuration is sometimes referred to as *centralized implementation*.

In a *three-tier* landscape, there is a network of two or more iSeries servers. One server in the network is a database server, while all the other servers in the network are application servers. Optionally, the database system can also provide application server functions. This configuration is sometimes referred to as a *client/server implementation*. With the introduction of logical partitions (LPAR), the three-tier configuration can run on a single iSeries server, partitioned into two or more logical partitions, each accommodating a separate iSeries server. For additional information on LPARs, refer to Chapter 3, "SAP R/3 system landscapes on iSeries" on page 43.

The multi-tier landscape brings SAP R/3 to the Internet. In addition to the database, application, and presentation layers, an Internet communication and transaction layer is added in a multi-tier landscape. This landscape has one database server, at least one application server (can be multiple servers), multiple presentation servers, multiple Internet Transaction Servers, and multiple Web servers.




Figure 11. SAP R/3 multi-tier architecture

For more information on landscapes, refer to Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43.

2.7 Solution architecture

SAP R/3 runs on the iSeries server in *native* mode. The solution involves *no emulation* whatsoever. The iSeries platform is compliant with the commercial aspects of the Single UNIX Specifications (SUS, formerly SPECS 1170) and provides all the functions required by the SAP R/3 system.

Less than 10% of the SAP R/3 code (the kernel) is written in C. The rest is written in ABAP, a programming language developed by SAP. ABAP generates code that runs in an interpretive mode on SAP R/3 application servers, regardless of the underlying implementation platform, which makes the SAP R/3 application functions platform independent.

2.7.1 SAP R/3 jobs

The iSeries server allocates SAP R/3 resources to an iSeries subsystem that runs the following SAP R/3 processes:

- Message server
- Dispatcher
- Multiple work processes
- Gateway

These processes are *batch immediate* jobs (BCI) on the iSeries server. A batch immediate job is started (spawned) by another job and may be functionally dependent on that parent job. BCI jobs bypass the job queue and run in the same subsystem as the parent job and inherit most attributes from the parent job.

A *batch job* (BCH), on the other hand, is submitted for execution by another job, but runs as an independent job on the iSeries server.

A user request from a workstation is received by the dispatcher who assigns the request to a work process. Work processes manage the request and run the tasks. Depending on the activity of each workstation user, a single dialog process may support multiple SAP R/3 users (Figure 12).




Figure 12. SAP R/3 job management on the iSeries server

Figure 13 shows an example of a possible implementation of multiple SAP R/3 systems on a single iSeries server.




Figure 13. SAP R/3 systems on an iSeries server

In this example, there are three SAP R/3 systems (D01, T01, and P01), each with its own separate database. The database is referred to as an *OS/400 SQL collection* and is held in an iSeries library. Each database has multiple SAP R/3 clients (000, 001, and so on). The application functions are available through four SAP R/3 instances (00, 01, 03, and 05), which are represented with OS/400 subsystems. Each subsystem has the name R3_<nn>, where <nn> is the instance number that must be unique within an iSeries server. The R/3 system P01 has multiple instances (03 and 05).

When an instance is started, the initiating job is a batch job, running in the R3_nn subsystem, while each SAP process is started as a batch immediate job. Each instance has its own dispatcher process and several SAP R/3 work processes, depending on the specifications made in the instance profile.

Each SAP logged on user can have up to six sessions. Within each session, a user can have up to nine operation modes.

2.7.2 National language support for SAP R/3

SAP provides single-byte character set (SBCS), such as Latin-1 and Latin-2, and double-byte character set (DBCS) language support on the iSeries server. The iSeries default language support is the Latin-1 language group. This is represented by SAP code page 0120.

If you want to work with a non-Latin-1 code page, there are several items of which you should be aware. For more information on national language support, refer to Chapter 6, “National language support” on page 113, and SAP Notes 99792 and 69337.

2.8 iSeries support for multiple SAP R/3 systems

For any computer-based information processing solution, it is prudent to isolate the application development, system testing, and production environments as much as possible. This is true for the SAP R/3 implementation as well. The degree of isolation can range from running each environment on a separate machine or having all the environments on a single machine. A company should evaluate the cost compared to the benefit of the various degrees of isolation. Your company may be unique in the following areas:

- Extent of development activity
- Change management procedures
- Business dependence on information systems
- Budgetary considerations

Some companies have installed individual iSeries servers for development, testing, and production. Similarly, there are organizations that have installed an iSeries server to support all of these environments on a single machines. In a customer scenario with major development activity, well-developed change management processes, a high degree of business dependence on the information systems, and a sufficient budget to support information technology, the organization may choose to install separate iSeries machines for:

- Application development
- Pre-installation testing
- Staff training
- Production
- Failover backup

At the other end of the spectrum, a customer may decide to solve their IT requests with a single machine. They can take advantage of the facilities available on the iSeries server to implement development, testing, and production environments on a single machine, with a degree of isolation that meets their requirements. They can do this at a cost they are comfortable with, while accepting the risks associated with running the diverse environments on a single iSeries machine.

Another option is to consider combining the development and testing environments on one machine, while dedicating a second machine exclusively for production work.

With the introduction of logical partitions in OS/400 V4R4, a large iSeries server can be partitioning appropriately for hosting multiple SAP R/3 environments in independently defined system partitions. For examples, refer to Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43.

2.8.1 Memory

Through the use of iSeries subsystems and the allocation of memory pools, it is possible to prevent memory contention between users of separate environments. In an SAP R/3 implementation, each SAP R/3 system (development, test, production), with its associated database, can have multiple instances defined. Each instance runs in a separate iSeries subsystem.

The subsystem definition allows you to define and allocate individual memory pools to the subsystem that is not accessible by users from any other SAP R/3 instance. In addition, by using the iSeries *shared memory pools*, you can choose selected memory pools shared across SAP R/3 instances.

An iSeries *class* object contains the run attributes that control the run-time environment of a job. Within the class, it is possible to automatically manage run-time priorities across different subsystems. In an SAP R/3 environment, this means that R/3 work processes within an instance can have their execution priority determined *automatically* by the class. For example, you can give to your production instances higher priority than to test instances. This is in addition to the capability to further modify the execution priorities of work processes within an instance.

2.8.2 Disk (auxiliary storage)

iSeries storage management allows for segmentation of the total disk capacity into separate auxiliary storage pools (ASPs) for each SAP R/3 system (for example, development, test, or production). This is used by allocating specific disk drives to each of the user ASPs that you want to create.

Information written to disk is spread over all of the available disk drives within each ASP. This helps balance the workload of the disk arms within each ASP. By defining separate user ASPs for each of the SAP R/3 systems installed, it is possible to minimize the impact of disk activity from one SAP R/3 environment on other environments.

The isolation of disk drives can be optionally extended to the disk IOPs and even to the bus to which the disks are attached.

The disk hardware protection mechanisms of mirroring or RAID-5 apply to all of the disks attached to the iSeries server. The protection that is chosen can be different for each ASP.

It is important to note that the unused disk space may vary greatly depending when R/3 systems are started, stopped, or upgraded. Additional disk space must be sized when customers run multiple systems for the additional database and customer data of another R/3 system, and for the temporary disk space required to start another R/3 system.

2.8.3 OS/400 new version and release support

IBM ensures that a new version and release of the OS/400 can support all of the functions of the preceding system software. This has been a major benefit of the iSeries server that protects a customer's investment in software.

A new version and release of OS/400 can be expected to support SAP R/3 software that ran effectively on a previous version and release of OS/400. However, it may be necessary to upgrade the SAP kernel to support the new OS/400 release. If it is necessary to install a new version of SAP R/3 software that requires functions included in a new version/release of OS/400, install the new OS/400 first. This allows the old and new versions of SAP R/3 to coexist on the same iSeries server but in separate libraries.

2.8.4 Coexistence of multiple SAP R/3 systems

It is possible to install multiple SAP R/3 systems on a single iSeries server. Naturally, this is subject to the availability of adequate disk, memory, and processing capacity to support the additional SAP R/3 systems. When submitting input for sizing to the IBM Competence Centers, it is important to identify how many SAP systems you plan to run on a single machine so they can account for the additional resource requirements (memory, disk, and processor) in their recommended configuration.

You can opt to install these SAP R/3 systems with shared or separate kernel libraries. If you install a second production system, you may choose to share the kernel libraries to save space. If you install multiple SAP R/3 systems, where you may want different versions of SAP R/3, use separate SAP R/3 kernel libraries.

2.8.5 iSeries server shared facilities

When using iSeries Logical Partitioning (LPAR), it is possible to run *multiple iSeries servers* in separate logical partitions of a *single iSeries machine*. With LPAR, each iSeries server uses its own system resources and facilities independently from other iSeries servers, even though they all run on the same iSeries machine. This section discusses the sharing of resources of one iSeries server, that runs either in one of the logical partitions or is the only system on the iSeries machine.

Multiple SAP R/3 systems running on a single iSeries server (in the same Logical Partition) share two key resources: the central processor (CPU) and the operating system. The result of this share are the three considerations that are discussed in this section. These are key factors in deciding whether to run multiple R/3 systems on a *single* iSeries server or *multiple* iSeries servers.

2.8.5.1 Sharing a CPU

One of the common resources that you have to consider is the CPU. Regardless of the number of CPUs, the iSeries server manages all its CPUs as a single entity. The iSeries server dispatches tasks to its CPUs while maintaining a balance in their usage.

Therefore, if you encounter a “runaway” program or some other resource-intensive task during application development and testing, it can impact production activity if SAP development test and production systems are running on the same iSeries server. You can minimize the potential impact of this by running SAP development and test systems at a lower priority than the production system. Also, SAP R/3 includes numerous parameters that can be used in the development and test systems to control “runaway” programs. It is prudent for the developers to be aware that they can impact production activity and be more cautious in their activity.

2.8.5.2 Sharing a copy of the operating system

SAP completes a certification process to confirm that specific releases of SAP R/3 are supported on each new release of OS/400. Similarly, IBM goes through stringent quality reviews prior to each release of OS/400 and follow-on PTFs.

By having SAP production, development, and testing all in the same copy of OS/400, you cannot test a new PTF, cumulative package, or OS/400 upgrade

independently from your SAP production system. You *can* do this in separate LPARs.

This is a risk that must be evaluated when determining the costs versus the benefits of such an approach. For example, assume you want to run production in the same copy of OS/400 as development and test. If you want to upgrade your test system, you may first have to upgrade OS/400 and install a new R/3 kernel. Consequently, your production system is also impacted since you upgraded OS/400. SAP recommends that you upgrade R/3 first on an R/3 test system while running an R/3 production system on another hardware system. Once an upgrade is performed with sufficient time for quality assurance testing, the production system is upgraded.

2.8.5.3 Using LPAR in an SAP R/3 environment

Running R/3 in separate logical partitions overcomes the limitations described in 2.8.5.2, “Sharing a copy of the operating system” on page 37. For examples of using LPAR with SAP R/3, refer to 3.4, “SAP R/3 landscapes with logical partitioning (LPAR)” on page 47.

2.8.6 Experience to date

There are many customers today running multiple R/3 systems on a single iSeries machine. This works extremely well for multiple test, development, QA, and training system needs. There are also iSeries customers running their R/3 production and development systems on a single machine.

In both of these cases, it is imperative that the iSeries server is configured properly for the necessary additional resources required to run more than one R/3 system concurrently. And, in the case of running production and development R/3 systems concurrently on a single iSeries machine, it is important to understand the potential risks to the production environment outlined previously to determine whether this option satisfies your company's specific requirements.

2.9 Differences between iSeries and UNIX implementation

The iSeries is an integrated platform, which includes hardware and database management software to support SAP R/3. In most other platforms, the customer has to select and integrate the various components of the platform. Ease of use and low costs (both implementation and on-going costs) are associated with the integrated solution provided by the iSeries server.

The following sections outline some of the differences between the iSeries server and UNIX-based platforms.

2.9.1 Memory management

The UNIX System V kernel divides the virtual address space of a process into logical *regions*. A region is a contiguous area of the virtual address space that can be treated as a distinct space to be shared (with other processes) or protected. UNIX address spaces are unique within a process. The same address in different processes refer to different spaces.

The iSeries address space is unique across the machine. The same address always points to the same space. Consequently, the way addresses are

translated and the way memory is managed within the iSeries architecture is fundamentally different to the architecture typically associated with UNIX systems. These differences are summarized in Table 1.

Table 1. iSeries and UNIX storage management differences

OS/400	UNIX
Single level storage	Process address storage
Absolute addresses	Relative addresses
Single process (job)	Multiple processes
Full job structure	Lightweight process

Every effort has been made to make the implementation of SAP R/3 on the iSeries server as simple as possible (this includes a factory preload option of the SAP R/3 software).

2.9.2 Database

Some of the major iSeries database management specifics are:

- You do not have to be concerned about compatibility issues of different release levels of the database and operating system. When there is a new release or version of the iSeries server software, IBM tests the integration of all components involved to ensure compatibility.
- The full integration of DB2 UDB for iSeries with the operating system eliminates the need for a separate database software installation and configuration procedure.
- The automatic load balancing of the disk drives by the iSeries server eliminates the need for complex manual database requirements analysis and installation. It also minimizes the need for on-going database management activity.

If you already have an iSeries-based application, you can easily integrate or exchange data with your existing applications because all iSeries applications use the same database management system. Your developers do not have to learn about a new database or operating system. You can also use your existing backup facilities.

- There is no separate starting or stopping of the database. When the iSeries server is running, the database is also active. There is no SAPDBA utility or equivalent necessary. The SAPDBA tool provides the database administrator functions in a UNIX-Oracle environment. On the iSeries, these functions are integrated in the operating system.

Note

There is a STARTSAP *DB command that is needed for three-tier implementation. The STOPSAP command only stops SAP on the machine on which you issue the command. The collector (SAPOSCOL) continues to run, but may be stopped either from the R3MAIN manu or by using the CALL SAPOSCOL '-k' command.

- The tablespace management process, such as creating or extending existing tablespaces, is not required on the iSeries server. There are no “table spaces” in the iSeries implementation. DB2 objects are stored in a large address space using 64-bit addressing. These are mapped to the installed disk and memory by the iSeries server. When additional space for a table is required, the database management system on the iSeries server can extend the table without any user involvement. If the space usage reaches a specified threshold, the system issues a message. Then, you need to take corrective action by deleting unwanted data or by adding more disk drives.

Any fragmentation of disk space is automatically retrieved by the iSeries server if possible.

- Space within a table resulting from row deletion is available for reuse by new rows added to the table. In general, you do not need to reorganize the database files to reclaim space occupied by deleted records. If there is a significant deletion of rows (such as in a client deletion), a database reorganization can return the unused space to the system. Use the RGZPFM command with the table name and the library name (R3<SID>DATA) to perform a reorganization. The table must not be in use for the command to run.
- The Export/Import function to and from UNIX databases is performed on the iSeries server using a simple CPYF command that copies a database file. When you use the CPYF command in an SAP environment, keep in mind that this command always generates a physical file and not a table. SAP R/3 recognizes physical files during runtime. However, it does not recognize them during transports and upgrades. For this reason, we recommend that you use CPYF together with the CRTDUPOBJ DATA (*NO) command in an SAP environment.
- Redo logs in UNIX are represented in the iSeries server by journal receivers. We recommend you define them in a separate auxiliary storage pool located on a separate disk unit. You can protect the logs using mirroring or RAID-5, the same as the other disks.
- The iSeries equivalent of archive mode in UNIX is achieved by the journal receivers being saved to tape.
- On UNIX-Oracle systems, *database table buffers* have to be set up and regularly managed to optimize access to the databases. It is important to have good buffer quality (an indication of how many database requests are satisfied directly from the buffer instead of going to the disks). On the iSeries server, this function is integrated in the system.
- DB2 for the iSeries server automatically creates temporary indexes when necessary (for table access and sorting). The system makes recommendations on which indexes should be made permanent in the system for best performance if the DB monitor is collecting data at the time that the index is built. Oracle does not use temporary indexes or advise adding indexes.
- Expert mode in UNIX enables database recovery to be protected against unauthorized execution through the use of a password. After entering the correct password, the database administrator can perform a recovery with SAPDBA for an Oracle database. On the iSeries server, all user functions are controlled by an integrated security management system.

2.9.3 System level

At the system level, you will notice some differences in the information shown by certain SAP transactions. For example, in the ST06 transaction, the following fields do not apply on the iSeries server:

- System calls
- Interrupts
- Context switches
- Free physical memory
- Swap space

2.9.4 Authority

Files stored in the QOpenSys and / (Root) directories are authorized in the same manner as UNIX files. Figure 14 shows the relationship between UNIX permissions and security used on the iSeries database files. This screen shows the results of running the WRKAUT command.

```
Work with Authority

Object. . . . . : /sapmnt/R01/profile/DEFAULT.PFL
Owner . . . . . : MONTI
Primary group . . . . . : R01GROUP
Authorization list . . . . . : *NONE

Type options, press Enter.
  1=Add user   2=Change user authority   4=Remove user

      Data
Opt  User      Authority  Objopr  Read  Add  Update  Delete  Execute
-----Data Authorities-----
      *PUBLIC   *EXCLUDE
      MONTI     *RWX      X      X      X      X      X      X
      R01GROUP  *RWX      X      X      X      X      X      X
```

Figure 14. Mapping UNIX permissions to iSeries security

You can change these authorities on this screen or through the Operations Navigator GUI as shown in Figure 15.




Figure 15. Object authority management from Operations Navigator

2.9.5 Character sets

Most UNIX applications run on hardware that uses ASCII character encoding, while the iSeries server normally uses EBCDIC encoding. The collating sequence (ordering of characters) is also different. This architectural difference is not a concern for high-level language applications that do not depend on (or make an assumption about) the character set. With SAP R/3 using the native iSeries database and being platform independent, there are no dependencies on the character set code.

Chapter 3. SAP R/3 system landscapes on iSeries

SAP R/3 implementation on the iSeries supports two-tier, three-tier, and multi-tier landscapes. In a two-tier configuration only one iSeries server is required. In a three-tier configuration, two or more iSeries servers make up the landscape. To enable these two- and three-tier configurations for the Internet, at least one Internet Transaction Server is required.

3.1 Two-tier landscape

In a *two-tier* configuration (Figure 16), a single iSeries server functions as the application server and database server. From the R/3 viewpoint, this is a stand-alone machine that does not need to interact with another machine to service the R/3 end-user requests coming from the presentation services layer. This configuration is sometimes referred to as a *centralized implementation*.




Figure 16. iSeries two-tier configuration

3.2 Three-tier landscape

In a *three-tier* landscape (Figure 17), there is a network of two or more iSeries servers. One system in the network is a database server, while all other systems in the network are application servers. Optionally, the database system can also provide application server functions. This configuration is sometimes referred to as a client/server implementation, where the application servers are the *clients* and the database server is the *server*.




Figure 17. iSeries three-tier configuration

3.3 Multi-tier landscape with mySAP.com and ITS

SAP Internet Transaction Server (ITS) is a tool that enhances the three-tiered SAP architecture for use in the Internet. It combines the existing Internet technology with SAP technology and provides reliable access to SAP functions from the Internet or intranets.

ITS is a component of mySAP.com. With its functionality, ITS is a part of the mySAP Workplace middleware, together with a Web server and the Workplace Engine. It mainly controls the generation of HTML pages and enables end users to interact with the different SAP systems through a Web browser. ITS also serves as the server component for the SAP GUI for HTML. Figure 18 shows a sample scenario of a multi-tier landscape.




Figure 18. SAP R/3 Internet and intranet solution using the ITS

In a multi-tier landscape, the presentation layer runs on a Web browser, which sends requests through a firewall to the Web server. The ITS is an intermedior between the Web server and R/3 system. The ITS consists of two software components (Figure 19):

- A-gate process (Application Gate)
- W-gate process (Web Gate)




Figure 19. ITS components

The A-gate process establishes the connection to an R/3 server. The W-gate process establishes the connection to the Web server. The communication between the two processes (A-gate and W-gate) is via TCP/IP. Each Internet Transaction Server has exactly one A-gate. The W-gate passes the request to the A-gate. The A-gate can be connected to *several* W-gates, which allows the support for multiple Web servers for load balancing and other landscape distribution purposes.

To manage the load at the ITS layer, several Internet Transactions Servers can be connected to the same R/3 system. Similarly one ITS can cater to several

application servers of one R/3 system, either via load balancing or separate selection of a specific application server.

The ITS provides the SAP R/3 architecture with a Web-efficient, browser-based interface that supports a large number of concurrent users. The multi-tier architecture of SAP R/3 with ITS offers maximum flexibility in terms of scalability.

The SAP Internet Transaction Server runs on a Microsoft NT Server, which can run either on a stand-alone PC Server or on the Integrated xSeries Server. Figure 20 and Figure 21 show mySAP.com landscapes for two- and three-tier iSeries configurations. The ITS can be installed on a stand-alone Windows NT server or it can also be installed on the Integrated xSeries Server.




Figure 20. ITS implementation for SAP R/3 iSeries two-tier landscape




Figure 21. ITS implementation for SAP R/3 iSeries three-tier landscape

3.4 SAP R/3 landscapes with logical partitioning (LPAR)

The objective of LPAR, released with OS/400 V4R4, is to provide users with the ability to split a single iSeries machine into several independent systems capable of running applications in multiple, independent environments simultaneously. For example, logical partitioning makes it possible for a three-tier SAP R/3 configuration (that would normally require multiple iSeries servers) to run on a single iSeries machine, as if it was running on separate physical iSeries machines. You can learn more about the concepts of LPAR in 1.2.10, “Logical partitioning (LPAR)” on page 21.

Figure 22 shows how a three-tier SAP R/3 configuration can be implemented on a single machine using LPARs. This particular example has two application servers and one database server, all running on a single iSeries server using logical partitioning.




Figure 22. SAP R/3 three-tier configuration using LPAR

A 4-way iSeries server in this example is partitioned into three logical partitions. The communication between the partitions is provided by Virtual OptiConnect, which is an iSeries feature that provides the high-performance link between partitions that can be used by SAP R/3. Multiple application servers in different partitions can access a single database server in the particular partition.

The primary partition controls all other partitions through a layer of kernel code called the *Hypervisor*. When the primary partition (LPAR 0) is shut down, this automatically shuts down all secondary partitions (LPAR 1, LPAR 2, and so on).

Recommendation

Always run the R/3 production system in LPAR 0. Running the production system in any other LPAR will shut down your production system whenever LPAR 0 is shut down. In case you run two production systems on a single machine, consider using LPAR 0 only for the Hypervisor and running the production systems in LPAR 1 and LPAR 2.

3.4.1 LPAR benefits

iSeries Logical Partitioning, described in 1.2.10, “Logical partitioning (LPAR)” on page 21, is a way to have several independent iSeries *servers* on one machine. With LPAR you can achieve the following benefits:

- Flexible system configurations with the ability to shift hardware resources between logical partitions. For example, if you have an R/3 development system, test system, and production system on one iSeries box, you can re-allocate CPUs, memory, disks, and other hardware resources between these systems according to changes in your workload.
- The ability to run different software releases (including IBM system software and application software), different PTF or patch levels, and different system settings (for example, different time zones, languages, and so on) on one iSeries machine.
- The ability to run SAP R/3 together with other applications on one hardware platform.
- The opportunity to reduce costs due to fewer hardware systems and fewer software licenses (which affect support and maintenance costs as well), less space occupied by IT equipment, and lower consumption of AC power.
- The ability to switch high performance, high capacity, and expensive external tape drives among different partitions, therefore, reducing backup and recovery time.
- Improved performance due to high-speed connection between logical partitions. Very fast internal communication between different partitions opens the possibility for scenarios such as:
 - High availability solutions implemented between the partitions to enhance the system availability for planned outages
 - Data replication, for example, for data warehouse applications, such as SAP BW
 - Data interchange between different applications or between development, test, and production systems

The following sections describe landscape scenarios with LPAR. In these scenarios, each partition is represented by a picture that resembles a stand-alone iSeries machine. This is because logical partitions are really separate, autonomous systems. Except for the control of the Hypervisor, they are functionally independent.

3.4.2 R/3 development, test, and production system

Probably the most typical use of R/3 with LPAR will be the installation of the R/3 development system, test system, and production system into separate partitions on one machine. In this scenario, the R/3 landscape is defined the same as if using multiple iSeries machines, but installed on one machine, resulting in a three-tier environment running on one single physical box. The benefit of this solution is that you can reallocate CPUs, memory, disks, and other hardware resources between these systems according to changes in your workload. This scenario is illustrated in Figure 23.




Figure 23. Different SAP systems on logical partitions

3.4.3 e-business applications with SAP R/3

Following the trends toward Internet-based e-business solutions in the latest releases of R/3, using one partition for the Web server (HTTP server) with a firewall provides a high degree of protection and integration to the R/3 database and applications. Figure 24 shows an example of using a partition for e-business and firewall with R/3.




Figure 24. Partition configured for firewall protection and R/3

3.4.4 Other applications with SAP R/3

Another scenario may run R/3 software in one partition and another vendor's software, or even other SAP software (for example SAP Business Warehouse), in another partition on the same machine. Figure 25 shows an example of the scenario where R/3 and other applications are installed on separate LPARs.




Figure 25. SAP R/3 and BW systems on logical partitions

3.4.5 Backup system with test system

Replication between LPARs and iSeries servers to permit functions, such as high availability, save and restore, query, and reporting from the replicated LPAR, is another scenario. The R/3 production system can be replicated to another machine for high availability reasons. Since the replicated machine does not need the same resources (memory, disks, CPUs, and so on) as the production machine all the time, it can be configured into two LPARs, with another LPAR used for testing, for example. Performance improvements may be realized if certain batch functions (queries, customer created reports, backups) are performed on the replicated system instead of the production system.

Figure 26 shows an example of R/3 production system replicated to an LPAR on another iSeries machine.




Figure 26. R/3 production system replicated to LPAR 0 on a different iSeries machine

3.4.6 Multiple SAP R/3 application servers

Another scenario would be to use partitions for individual server functions. An example may be to have the application servers installed in one or more partitions, and the central database server installed on a different partition.




Figure 27. SAP application and database servers on logical partitions

The drawback of this scenario is that the system resources are split across the partitions and are not used efficiently. For this reason, such a scenario is generally not recommended unless your installation requires multiple application servers, each with different settings, such as multiple language support, different time zones, system values, and so on.

3.4.6.1 Implementing an archiving solution

In this scenario, the customer has a very large database containing a mix of live and historical data. Most of the interactive and batch jobs use the live records on the database. However, the customer needs to access the historical data both interactively and in batch mode from time to time, with reasonable response times. Furthermore, they also need to reduce their backup time without jeopardizing data security. Assuming that the application can distinguish between live and historical data, one possible solution is to create a three-partitioned configuration as shown in Figure 28.




Figure 28. Setting up an archiving solution

3.4.6.2 Minimizing backup and recovery windows

The imperatives driving today's businesses demand total machine availability, 7 days-a-week, 24 hours-a-day (7 x 24). On the other hand, data security demands that a foolproof backup strategy be in place to meet unforeseen contingencies. A third ingredient in this situation is the relentless growth of databases as businesses become more complex. Logical partitioning can provide one solution to balance these conflicting needs.

Let's look at a scenario where backup is becoming so time consuming that the production window is decreasing fast. Figure 29 illustrates a possible solution to minimize backup of the production database. Incidentally, this scenario also facilitates recovery.




Figure 29. Minimizing the backup window

In this scenario, the production server is running on partition 0, and all updates are replicated on partition 1, using remote journaling. At preset intervals, the partition 1 update is suspended, and the partition database is backed up. After the backup, the partition 1 database is re-synchronized with partition 0, by applying all accumulated journal entries from partition 0 to partition 1.

This scenario provides for recovery of the production database onto the same partition or a different physical machine, with minimum inconvenience.

Using one of the High Availability Business Partner (HABP) applications, it is possible to replicate data from one partition to another on the same machine. Although this option does not offer any higher level of availability with unplanned outages, it offers a higher level of availability with planned outages. See Figure 30.




Figure 30. Backup solution using High Availability Business Partner Applications

The CPU required for this partition will have to be adequate so that data changes from other partitions can be received and applied and, when required, the save tasks can be run. A save task is typically I/O intensive so this in itself will require very little CPU cycles.

Enough DASD needs to be installed in this partition to support the amount of data it will be required to store. This is likely to be the largest partition on the system.

When a save is required, the apply process in the backup partition is stopped and the save is taken from this copy of the data. Work continues to run in the original partition and journal entries generated by changes to the data are still sent to the backup partition but not applied.

3.4.7 Other LPAR scenarios

Some real-life examples of other scenarios using LPAR are:

- A company outsourcing the iSeries server to numerous other companies, while using separate logical partitions to maintain system integrity
- Using one logical partition to perform and test OS/400 or R/3 upgrades and apply patches
- Testing iSeries server values (such as QSECURITY, QDATE, QTIME, QPFRADJ, and so on) in a logical partition

Chapter 4. Sizing and capacity planning

Sizing means determining the hardware requirements of an SAP system such as:

- Network bandwidth
- Physical memory
- CPU power
- I/O capacity

Determining the correct size of a platform for an SAP R/3 installation in terms of processor speed, memory, disk space, and configuration design is not a trivial task. Every configuration has to be analyzed based on:

- Individual customer requirements
- Number of users
- Transactions per time period
- SAP R/3 modules that are being used
- Customer workloads
- Choice of hardware platform

The IBM/SAP sizing methodology is based on SAP R/3 benchmarks, information from SAP, and actual customer experiences. The IBM Sizing Center uses sizing tools and customer input to approximate the system resource requirements; however, actual customer results may vary. Figure 31 shows the sizing factors.




Figure 31. Sizing factors

Prior to ordering the final hardware configuration, all customers should consult with a trained SAP Basis specialist to develop a detailed plan for the SAP R/3 implementation.

This chapter explains sizing methodology, terminology, considerations, and basic information that will help you to understand the sizing estimate results.

4.1 Sizing terminology

The phrase “R/3 users” can have diverse meanings. It is important to understand which type of user is being discussed at any particular time, as noted in the following explanations.

Named user

A named user is equivalent to a user master record in the SAP system (for example, someone who has the right to log on to the system). This is a metric used by SAP as a base for calculating the SAP R/3 license fee. The number of named users is not relevant for bench-marking or sizing a platform.

Information user

An information user is a user that can only perform queries, but not updates, of the system.

Logged-on user

A logged-on user is logged on to the SAP system but may or may not be actively working. This is typically the kind of a user that customers refer to when being asked for the number of users on their R/3 system.

Active user

Active users are logged on to the system and are actively working, which means that they press the Enter key to initiate R/3 dialog steps (displays). Some assumptions include:

- A user’s key think time is 25 seconds (which calculates to 2.4 dialog steps per minute).
- The number of active users is 50% to 60% of the number of logged-on users.

Benchmark user

A benchmark user only has a think time of 10 seconds between dialog steps. Therefore, one benchmark user is usually seen as being equivalent to three active users.

Dialog step

This is an SAP term used to measure a unit of work. In an interactive environment, a dialog step is equivalent to a screen change. SAP also uses dialog steps as a measure of throughput in batch and update processing workloads. As shown in Figure 32, one dialog step encompasses all of the system activity that takes place as a user moves from one R/3 application screen to the next. Each R/3 module (for example, FI for financial accounting, AA for asset accounting, and PA for payroll accounting) is made up of several transactions. Each transaction has multiple dialog steps.




Figure 32. Dialog step

Transactions

SAP transactions (not to be confused with other types of transactions such as IMS or CICS transactions) consist of one to 12 dialog steps depending on the business transactions you look at. For example, posting a bill requires fewer dialog steps than running an online MRP for material. A general assumption is that one transaction equals four dialog steps.

SAPS

In discussions about performance, sometimes the term SAP Application Benchmark Performance Standard (SAPS) is used. It was introduced by SAP as a way to measure the maximum number of dialog steps on a processor. It is a measurement of maximum throughput independent of response time. IBM does not consider SAPS as a useful basis for performance evaluation. For more information, refer to 4.4.5.2, “SAP Application Benchmark Performance Standard (SAPS)” on page 62.

The size of the hardware and database is influenced by both business aspects and technological aspects. This means that the number of users using the various application components and the data load they put on the network must be taken into account. Therefore, a system should also be configured in a balanced way so that all relevant system components are sized in such a way that they work with an adequate average utilization without creating a bottleneck. Such a configuration ensures that the system can handle peak loads and has a predictable behavior.

4.2 SAP R/3 benchmarks

Since SAP R/3 Release 1.1H, SAP has provided all of its hardware partners with benchmark samples (application and data) as a means to measure the performance of their respective platforms. Right now these samples are available for all major R/3 modules. These are so-called “real benchmarks” based on SAP transactions. The SAP transactions in these benchmarks were chosen based on the analysis of typical SAP R/3 customer environments.

Along with the benchmark scripts, SAP delivers additional tools, benchmark clients, and a performance monitor. This represents the standard benchmark environment for all hardware vendors. Running benchmarks is the responsibility of SAP's hardware partners, such as IBM.

While benchmark results are only an important factor used to determine the suitability of a hardware platform, for the following reasons, they shouldn't be the only factor:

- Important system characteristics such as availability, recovery backup, or flexibility are not assessed by benchmarks.
- Standard benchmarks use only a small part of the possible SAP transactions per application.
- The benchmark test environment is artificial.
- SAP standard benchmarks do not reflect the specific implementation (setting of customization parameters) at an individual customer site.
- SAP standard benchmarks do not reflect the real application environment at the individual customer site.

4.3 Sizing and capacity planning approaches

There are several approaches to sizing and capacity planning:

- Questionnaire-based Input Analysis
- Pre-production Performance Evaluation
- Measured Customer Data Analysis

Which approach and when to use it depends on the phase of the implementation project. There are three distinct phases in the SAP R/3 implementation project (presented in order):

- **Pre-sales phase:** Preparation of initial sizing (or sizings) configurations for a specific customer proposal. The recommended approach in this phase is *Questionnaire Based Input Analysis*.
- **Pre-production phase:** Verification of planned capacity requirements is done by running a pre-defined representative subset of users and modules with the customer's own data and application customization. The recommended approach in this phase is *Pre-production Performance Evaluation*.
- **Post-production phase:** Analysis of future capacity requirements is done based on a collection of production performance statistics coupled with future business planning information. The recommended approach in this phase is *Measured Customer Data Analysis*.

Questionnaire based input analysis is designed to collect information necessary for developing an IBM hardware configuration proposal for your R/3 implementation. The guidelines for sizing an R/3 system come from a number of sources. The sources include SAP, actual SAP R/3 benchmarks, and customer feedback. Based on information from these sources, the sizing methodology is used to analyze the customer requirements to arrive at a recommended configuration.

Two types of sizings can be performed:

- **User-based sizing:** Based on the number of concurrent users for each of the SAP R/3 modules to be used
- **Transaction-based sizing:** Based on the transaction information provided for the SAP R/3 modules to be used

The user and transaction information is translated into a number of normalized FI (financial) dialog steps per hour by applying different weighting factors for each module.

Pre-production Performance Evaluation is a validation of the actual customer environment where tests are run using customer configured transactions against the customer's data. A representable workload is defined (appropriate transactions within the modules) along with a representable application mix (correct proportion of users working in various modules). Only a subset of the actual users is needed to run this test, and the results are then extrapolated to project production level workload. It is helpful to have defined/repeatable scripts prepared to reproduce any problematic transactions that are identified during this test. This test should be defined as closely as possible to the customer's typical production environment. Therefore, if batch workload or complementary products, such as EDI or Fax (typical), this workload should also be included in this pre-production test.

Experience to date in running this pre-production evaluation has been extremely positive. Customers not only validate their hardware configuration prior to production but also have the opportunity to identify problematic transactions prior to going live that may need focus through additional customization changes or SAP Notes. It is obviously the best practice to identify and address such changes prior to going into production to reduce/eliminate the impact to end users. The confidence of going live without performance problems is much higher for customers when they have included this pre-production performance evaluation in their rollout plans.

Measured Customer Data Analysis is the best of all approaches. First of all, it overcomes a psychological barrier for customers who do not believe in something they do not see (someone is doing a calculation for their configuration somewhere else with unknown rules and data). Measured data modeling is done on an iSeries server that assures the customer even more that the modelling is more trustful. The customers see what has been measured (they actually should be involved in selecting the right transactions, the data fed into their system, and even the execution of the transactions).

Modeling the measured data is a matter of running the resulting data through various queuing calculations specific to the iSeries architecture and specific to a selected hardware and software configuration. It also provides an easy mechanism to combine SAP R/3 capacity planning data with other workloads the customer might have today or expect to have in the future (for example, FAX support, telephony integration, and so on). For additional information on capacity planning with measured SAP data, please consult *AS/400 Server Capacity Planning*, SG24-2159, which devotes an entire chapter to SAP R/3 on AS/400 capacity planning.

4.4 Sizing materials, processes, and contacts

Many people around the world have been trained to assist with SAP R/3 sizings on IBM platforms. We strongly recommend you to seek assistance from trained personnel when proceeding with a customer opportunity. The best place to start is to contact the nearest SAP/IBM Competence Center. Refer to Appendix C, "Support for SAP R/3" on page 553, for Competence Center contacts.

Look for the following *key documents* to assist you in understanding the sizing and capacity planning approaches available today:

- IBM SAP R/3 Sizing and Planning Questionnaire
- IBM SAP R/3 Sizing Guidelines
- IBM SAP R/3 Comprehensive Testing Results

4.4.1 IBM sizing support

IBM offers hardware sizing support for new and installed SAP R/3 implementations via the IBM Americas Techline Sizing Center, located in Wayne, Pennsylvania, outside of Philadelphia. For six years, this team of 10 professionals has been sizing IBM hardware systems for ERP software. The team members average over 10 years of technical sales support experience with the full range of IBM hardware platforms.

Contact the Sizing Center by e-mail at ibmerp@us.ibm.com or by telephone at 1-800-426-0222 (within the US).

IBM Insight for SAP R/3

IBM Insight for SAP R/3 is an IBM offering that provides high level R/3 workload analyses for SAP installed, in-production systems. IBM Insight for SAP R/3 is available free of charge to existing SAP customers. This offering includes the Insight utility program, which gathers performance data on a customer's R/3 system, and formal analysis of the data by IBM technical specialists. The Insight results are delivered to the customer in a custom workload report. The report includes active user counts, machine utilization, user and module load distribution, dialog steps, batch and reporting usage, and other system and database information.

Insight analysis provides valuable planning information for customers who are migrating from one release of R/3 to another and for customers who are planning to add users to an existing R/3 system.

For more information about IBM Insight for SAP R/3, and to download the utility program, go to: <http://www.ibm.com/erp/sap/insight>

4.4.2 IBM SAP Capacity Planning Service Offering

In this service offering, skilled SAP R/3 consultants analyze the customer's current system. Also, capacity planning specialists build a performance model representing future requirements of all of the capacity elements in the client server application. The performance model is built from the customer's production R/3 system accounting data plus operating system data captured through a monitoring tool. The IBM SNAP/SHOT discrete simulation modeling tool is used to simulate all aspects of capacity consumption and resulting

performance, from the database server to the desktop. The project concludes with a workshop in which various “what if” scenarios are modeled and analyzed.

For more information on this service offering, contact IBM Global Services at 1-919-301-4162.

4.4.3 IBM Performance and Testing Support/Analysis Services

The IBM Solutions Center, part of the Americas Advanced Technical Support organization, provides free and fee-based performance analysis and diagnostic services for customers who encounter performance problems in their ERP or SCM implementations. This team of specialists has a rich history and broad experience supporting response time, batch window, and overall capacity and testing challenges in ERP and SCM environments.

For more information, call 1-650-286-6832.

4.4.4 SAP Quick Sizer

The Quick Sizer is a tool jointly developed by SAP and their hardware partners to help give customers an idea about initial sizing. It is free of charge. Customers may provide sizing input to IBM through either the IBM/SAP Sizing and Planning Questionnaire or SAP's Internet-based Quick Sizer. If the IBM/SAP Sizing and Planning Questionnaire is used, the data is entered into the SAP Quick Sizer.

The SAP Quick Sizer is an Internet application that guides the user through a structured sizing questionnaire. The Quick Sizer gathers information about an organization's business requirements and translates this data into generic system requirements, that is platform independent specifications for processor, memory and disk. The IBM Sizing Center uses the generic SAP results to develop the IBM hardware recommendation. The Quick Sizer offers two different sizing models: user-based sizings and quantity-structure-based sizings.

4.4.4.1 User-based sizings

The user-based model asks the user to count the number of *active* R/3 users by module. SAP considers this model to be limited in its ability to estimate the R/3 resource requirements because it does not consider important sizing factors such as user behavior, peak versus average workload, the amount of batch processing, reporting, and user customization. SAP recommends the user-based sizing model for small businesses only.

4.4.4.2 Quantity-structure-based sizings

The quantity-structure-based model is more thorough than the user-based model because it considers actual or expected R/3 workload throughput. Besides the number of R/3 users, this model gathers detailed information about the business processes and objects used, including the number of R/3 dialogs, workload profiles, peak usage times, retention periods for business objects, and background and reporting processes. SAP recommends the quantity-structure-based model for medium and large businesses.

Customers may complete the user-based sizing questions, quantity-structure-based sizing questions, or both. When both models are used, the Quick Sizer provides R/3 workload estimates for both models. However, the IBM sizing specialist will develop the IBM hardware recommendation from the larger of the two workload estimates.

4.4.5 Sizing the IBM solution

The IBM Sizing Center uses information from the SAP Quick Sizer as an input to the sizing process. The Quick Sizer provides processor, memory, and disk requirements. This section explains how information from the Quick Sizer tool is used to develop the IBM hardware recommendation.

4.4.5.1 Target CPU utilization

Based on the sizing inputs, the Quick Sizer approximates processor, memory, and disk consumption. For user-based sizings, SAP's sizing results are calculated to meet an average CPU utilization of 33% for the *online interactive workload*. Note that this percentage differs from the traditional IBM sizing methodology in which the target CPU utilization is 65%. There are several reasons for the differences in the target CPU utilization. The IBM sizing methodology uses a higher CPU target because it takes into consideration peak-hour workload and batch and reporting requirements. The Quick Sizer uses the lower CPU target because it does not consider peak processing requirements. The lower target is also used to leave sufficient processor resources available for batch, reporting, printing, and software interfaces.

The Quick Sizer quantity-structure-based model, like the IBM sizing tool, uses a target CPU utilization of 65% for most hardware platforms.

4.4.5.2 SAP Application Benchmark Performance Standard (SAPS)

The SAP Quick Sizer calculates a customer's potential R/3 workload in terms of SAP Application Benchmark Performance Standard (SAPS). SAP has defined SAPS as the standard measurement of system throughput.

One hundred SAPSs are equal to 2,000 *fully business-processed order line items per hour in the standard SD application benchmark*. "Fully business processed... in the SD standard benchmark" refers to the entire business workflow of an order line item, including creating the order and delivery note for the order, displaying the order, changing the delivery, posting a goods issue, listing the orders, and creating an invoice for the order. This throughput is achieved by processing 6,000 dialog steps and 2,000 postings per hour in the SD benchmark, or by processing 2,400 SAP transactions.

For sizing purposes, IBM rates the capacity of each IBM server in terms of SAPS.

4.4.5.3 Resource categories

For CPU and disk, the Quick Sizer results are expressed as resource categories, which are SAP's hardware independent resource specifications.

For CPU sizing, each resource category represents a range of SAPS. Table 2 shows an example of the CPU resource categories and SAPS requirements.

Table 2. CPU resource categories

Category	SAPS
1	Up to 125
2	Up to 250
3	Up to 500

For disk sizing, each resource category represents a range of disk space. Table 3 shows an example of the disk resource categories and disk space requirements.

Table 3. Disk resource categories

Category	Disk range
1	16-25 GB
2	26-35 GB
3	36-50 GB

4.4.5.4 Steps in the sizing process

The Sizing Center performs the following steps for every Quick Sizer sizing request:

1. Determine the customer's potential R/3 workload requirements.

The SAP Quick Sizer analyzes the customer input and calculates a generic sizing result. The Quick Sizer provides the memory requirement and resource categories for CPU and disk. If the customer completed both the user and quantity-structure-based sizings, the larger of the two workloads to determine the IBM hardware requirements is used.

2. Select either a two-tier or three-tier hardware configuration.

Based on the potential customer workload, we select either a two-tier or three-tier hardware configuration. In a two-tier configuration, one server provides both the database and application server functions. In a three-tier configuration, there is one database server and one or more separate application servers. If one server can handle both the database and application server workloads, a two-tier configuration is selected, which is typically appropriate for customers with smaller R/3 workloads. If the customer workload will not fit in a single server, a three-tier configuration is selected.

Refer to Chapter 3, "SAP R/3 system landscapes on iSeries" on page 43, for configuring a three-tier landscape on the iSeries server using logical partitioning.

3. Determine the IBM server requirements.

Based on the SAPS capacity rating for each IBM processor, the processor or processors that can best support the potential R/3 workloads are selected.

4. Communicate the sizing estimate results.

Finally, the summarized results of the sizing estimate in a customized report are forwarded to the requesting IBM Representative or Business Partner. The IBM Representative or Business Partner reviews the sizing results with the customer. Together they determine the hardware configuration that best meets the customer's needs. In this decision process, they consider the customer's growth plans, upgrade options, financial issues, etc.

4.4.6 Suggestions for disk sizing

Sizing the disk arms is an important, yet very difficult, task. With new SAP R/3 releases, the CPU requirements (dramatically), the requirements for memory, and the amount of required disk storage have increased.

The number of disk arms are calculated based on the CPW of the machine (the IBM performance metric for the iSeries). When calculating the number of disk arms, the following factors are taken into account:

- CPW
- Disk I/O workload (light, medium, heavy)
- Disk technology
- IOP technology
- Type of disk protection (RAID-5 is assumed in most cases)

The critical factor for SAP R/3 is recognizing that for a central server load, the disk I/O workload would be “light”, while the disk workload for just the database server in a three-tier environment would be “heavy”. This is because on a central server, most of the work is an SAP application server workload, which is CPU and memory intensive, but requires very little disk I/O.

As higher capacity disk (DASD) devices (8 GB and 18 GB) for the iSeries become available, fewer arms are needed to satisfy the capacity requirements. This can lead to configuring too few disk arms (actuators) to meet the workload placed on them. A lack of disk arms can bottleneck the processor's performance. To avoid such a bottleneck, a minimum number of disk devices is needed for optimum performance on each iSeries processor level. This number is independent of the quantity of drives needed to meet the desired storage capacity. For detailed disk sizing reference and considerations, go to:

<http://www.as400.ibm.com/developer/performance/dasdmnu>

Part 2. Implementation

This part describes the implementation tasks and techniques that are necessary to install and properly set up R/3 on the iSeries server. During implementation, all R/3 iSeries customers will experience the topics in the following chapters in this part:

- Chapter 5, “Installation and configuration” on page 67, briefly describes the installation and configuration procedures for a standard SAP R/3 Release 4.6C system on the iSeries server.
- Chapter 6, “National language support” on page 113, discusses issues and solutions related to DBCS and non-Latin-1 national languages.
- Chapter 7, “Connectivity” on page 123, describes the different connectivity options you have when connecting iSeries servers in an R/3 environment. It also addresses optical connections, Virtual OptiConnect, TCP/IP, and 1Gb Ethernet.
- Chapter 8, “Data porting” on page 135, covers data porting from legacy applications into an SAP R/3 environment. It takes you through the steps that are necessary to perform data porting and provides examples of using the data porting tools.
- Chapter 9, “R/3 system printing” on page 151, describes the fundamentals of SAP R/3 system printing and provides several definition examples.
- Chapter 10, “iSeries system administration using GUI” on page 215, covers the GUI-based system administration and system management tools for the iSeries server.
- Chapter 11, “Availability, backup, and recovery” on page 221, outlines the standard iSeries server backup, recovery, and availability facilities that are available. It also provides an example of the backup and recovery facilities and procedures in the SAP R/3 environment.

Chapter 5. Installation and configuration

This chapter briefly describes the installation and configuration procedures for a standard SAP R/3 Release 4.6C system on the iSeries server. You can find a detailed description of the installation process in the SAP document *Installing R/3 on IBM AS/400*, which is delivered with your SAP R/3 software. Use this redbook as a complementary resource, rather than as a replacement to the installation guide. You can usually find special instructions on upgrading to a new release of the SAP R/3 software in your upgrade kit.

If your iSeries server uses logical partitions (LPAR), you need to configure LPAR first.

This chapter does not include the steps and procedures required to install SAP add-ons, industry solutions, or plug-ins.

Note

The software distribution CDs for an initial installation are different from the CD for an upgrade.

5.1 Hardware and software requirements

This section describes the system requirements to successfully install an SAP R/3 system on the iSeries server.

5.1.1 iSeries hardware requirements

SAP R/3 requires the iSeries server to have PowerPC processors. Exact hardware configuration requirements depend on the recommended network solution. A fast LAN adapter (either Token-Ring or Ethernet) is required for attaching frontend clients through TCP/IP. In a three-tier environment, you also need a fast connection between the database server and the application servers. The fast connection can be:

- 1 Gigabit Ethernet on the iSeries 8xx models

Refer to 7.2, “Gigabit Ethernet support” on page 128, for information on 1 Gb Ethernet.

- Fiber-optic link
- Virtual OptiConnect on iSeries servers with LPAR

5.1.2 iSeries software requirements

You must have an OS/400 release that is certified and supported by SAP.

Note

You can find details about the required OS/400 release for SAP R/3 4.6C in SAP Note 156557. Refer to SAPNet, Quick link dbosplatforms for OS/400 release planning and SAP Note 78541 for SAP R/3 release strategy.

When connecting two or more iSeries servers (such as application servers to the database server in a three-tier environment) using a fibre-optic link, use OptiMover PRPQ 5799-FWQ. Refer to 7.1.1, “OptiMover for AS/400 (5799-FWQ)” on page 123, for more information about OptiMover.

5.1.3 Front-end requirements

The installation of the front-end SAP GUI software is independent of the SAP R/3 server platform. The following front-end operating systems are supported by SAP:

- Windows 32-bit versions
- Windows 16-bit versions
- OS/2 Warp
- Java GUI

SAP R/3 requires a TCP/IP connection from the application server to the front-end SAP GUI. The SAP document *SAP-Supported Network Products* contains details of the supported network software for front-end presentation.

For detailed information about the front-end requirements and the installation process, refer to the following SAP documents:

- *Check list - Installation requirements: Frontends*
- *SAP Front-end Installation Guide*

For information on how to use an IBM Network Station as an SAP R/3 front end, refer to Chapter 20, “Using an IBM Network Station as an SAP front end” on page 497.

Note

IBM Client Access software is *not* a prerequisite for the SAP R/3 front end.

5.2 TCP/IP configuration

Before you start installing SAP R/3, TCP/IP must be configured on the iSeries server. This section describes how to do this. If your TCP/IP is already configured, skip this section and go to 5.3, “SAP R/3 installation concepts” on page 74.

The TCP/IP communications protocol function, along with the related administration, is packaged with OS/400. This includes the following features:

- Packet Internet Groper (PING)
- Network Status (NETSTAT)
- Domain Name System (DNS) Server
- Dynamic Host Configuration Protocol (DHCP) Server
- Point-to-Point (PPP)
- Routing Information Protocol (RIPv2)
- Simple Network Management Protocol (SNMP)
- IP over Twinax

5.2.1 TCP/IP configuration on the iSeries server

Prior to configuration, you must know and have:

- IP addresses and a subnet mask of your iSeries servers
- Router
- Gateway
- Line description of your token-ring or Ethernet line
- Your local domain name
- Host name of the iSeries server

The following steps show you how to configure a basic TCP/IP connection. Your network administrator should check the connection if applicable. You can find detailed coverage of this topic in *TCP/IP Fastpath Setup*, SC41-5430.

1. Sign on to the iSeries server. Go to the Configure TCP/IP menu by typing on the command line:

```
CFGTCP
```

2. Select option 1 (Work with TCP/IP interfaces).

You need two entries, one for the loopback entry and one for the IP address of your iSeries server (Figure 33).

System: AS23

Work with TCP/IP Interfaces

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display 9=Start 10=End

Opt	Internet Address	Subnet Mask	Line Description	Line Type
	199.5.83.158	255.255.255.128	TRNLINE	*TRIAN
	127.0.0.1	255.0.0.0	*LOOPBACK	*NONE

F3=Exit F5=Refresh F6=Print list F11=Display interface status

F12=Cancel F17=Top F18=Bottom

Figure 33. CFGTCP: Work with TCP/IP Interfaces

Note that the loopback entry always has the IP address of 127.0.0.1, subnet mask of 255.0.0.0, and line description of *LOOPBACK. To add an entry, type 1 and press Enter. The screen in Figure 34 appears.

Add TCP/IP Interface (ADDTCPIFC)

Type choices, press Enter.

Internet address	> ' '	
Line description		Name, *LOOPBACK...
Subnet mask		
Associated local interface . . .	*NONE	
Type of service	*NORMAL	*MINDELAY, *MAXTHRPUT...
Maximum transmission unit . . .	*LIND	576-16388, *LIND
Autostart	*YES	*YES, *NO
PVC logical channel identifier		001-FFF
+ for more values		
X.25 idle circuit timeout . . .	60	1-600
X.25 maximum virtual circuits .	64	0-64
X.25 DDN interface	*NO	*YES, *NO
TRLAN bit sequencing	*MSB	*MSB, *LSB

Bottom

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

Figure 34. Add TCP/IP Interface (ADDTCPIFC)

You need to add entries for the first three fields and leave the other fields as the default.

- If the route to the remote host (in this case, the PC workstation) is through a gateway, or the remote host resides in a different network or subnetwork to the local host, it is necessary to configure a route. Select option 2 (Work with TCP/IP routes) on the Configure TCP/IP menu. Add an entry that is similar to the entry in Figure 35.

Work with TCP/IP Routes

System: AS23

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display

Route	Subnet	Next	Preferred
Opt Destination	Mask	Hop	Interface
*DFTRROUTE	*NONE	199.5.83.29	*NONE

Bottom

F3=Exit F5=Refresh F6=Print list F11=Display type of service
F12=Cancel F17=Top F18=Bottom

Figure 35. CFGTCP: Work with TCP/IP Routes

If a route has not been added, this list is empty.

- The local host table on the iSeries server contains a list of the Internet addresses and associated host names for this network.

Note

If you *do not have a name server*, you must add an entry for each system to which this iSeries server connects. To add a host table entry to the local host table on your iSeries server, return to the Configure TCP/IP menu and select option 10 (Work with TCP/IP Host Table Entries). The screen shown in Figure 36 appears.

```

                                Work with TCP/IP Host Table Entries
                                System:  AS23
Type options, press Enter.
  1=Add  2=Change  4=Remove  5=Display  7=Rename

      Internet      Host
Opt  Address      Name
  1
    127.0.0.1      LOOPBACK
                   LOCALHOST

                                Bottom
F3=Exit  F5=Refresh  F6=Print list  F12=Cancel  F17=Position to
```

Figure 36. CFGTCP: Work with Host Table Entries before adding names

Enter **1** the Opt field to see the Add TCP/IP Host Table Entry display (Figure 37). In our example, we use a name server and only need to set up the entry as shown in Figure 37.

```

                                Add TCP/IP Host Table Entry (ADDTCPHTE)
Type choices, press Enter.

Internet address . . . . . > '199.5.83.158 '
Host names:
  Name . . . . . 'as23.ibm.com'

                                + for more values
Text 'description' . . . . . as23

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

Figure 37. ADDTCPHTE: Add TCP/IP Host Table Entry

Go back to the Configure TCP/IP menu, and select option **10** (Work with TCP/IP Host Table Entries (Figure 38)).

Work with TCP/IP Host Table Entries

System: AS23

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display 7=Rename

Opt	Internet Address	Host Name
5	127.0.0.1	LOOPBACK LOCALHOST

F3=Exit F5=Refresh F6=Print list F12=Cancel F17=Position to

Bottom

Figure 38. CFGTCP: Work with Host TCP/IP Table Entries

You do not need to create a loopback entry and add an additional host name under it called “LOCALHOST”. Your Work with TCP/IP Host Table Entries display should look exactly the same as the example in Figure 38.

Note

You have to be careful to match TCP/IP system names between what is defined in the SAP R/3 system and what is defined in the TCP/IP system. R/3 uses the names in the /usr/sap/<SID>/sys/profile/DEFAULT.PFL file as the TCP/IP system names (where <SID> is your SAP system ID). The names are in this format:

<system>_<SID>_<INSTANCE_NUMBER>

The system portion of this name must match what is in your TCP/IP configuration (CFGTCP option 12 and option 10) or in your name server. Note that these names are also case-sensitive. To see the names that the SAP system uses or if your SAP system fails to start, you can look in the /usr/sap/<SID>/DVEBMGS<INSTANCE_NUMBER>/work/dev_w0 file. You can also, look at the job logs for **SAP<number>** and **sapstart.log**.

5. Select option 12 (Change TCP/IP domain information). Be sure that you have entries under “Domain Name” and “Host Name”. If you have a remote name server (or servers), you need to define the IP address for it here. See Figure 39.

```

Change TCP/IP Domain (CHGTCPDMN)

Type choices, press Enter.

Host name . . . . . 'as23'

Domain name . . . . . 'ibm.com'


Host name search priority . . . *LOCAL      *REMOTE, *LOCAL, *SAME
Domain name server:
  Internet address . . . . . '199.4.191.76'
                          '199.4.191.75'


F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys
Bottom

```

Figure 39. Change TCP/IP Domain (CHGTCPDMN)

For the sake of convenience, we assume that the TCP/IP host name is equal to the SNA system name. However, the TCP/IP host name can be different from the SNA system name. The “Current system name” is on your Network Attributes display. To access this display, type:

DSPNETA

See Figure 40.

```

Display Network Attributes

Current system name . . . . . : AS23
Pending system name . . . . . :
Local network ID . . . . . : ASAA
Local control point name . . . . . : AS23
Default local location . . . . . : AS23
Default mode . . . . . : BLANK
APPN node type . . . . . : *NETNODE
Data compression . . . . . : *NONE
Intermediate data compression . . . . . : *NONE
Maximum number of intermediate sessions . . . . : 200
Route addition resistance . . . . . : 128
Server network ID/control point name . . . . . : *LCLNETID *ANY

More...
Press Enter to continue.

F3=Exit  F12=Cancel

```

Figure 40. Display Network Attributes

6. To start the entire TCP/IP stack on the iSeries server, type:

STRTCP

You can ensure that TCP/IP is started after each IPL by inserting the STRTCP command in the QSTRUP program, which is called each time you IPL the iSeries server. The Command Language Program (CLP) QSTRUP source for this can be retrieved if you want to use the system-supplied startup CLP as a model. If you want to retrieve CL source, use the RTVCLSRC command.

5.3 SAP R/3 installation concepts

This section introduces some key concepts and features that may help you better understand what happens during the R/3 installation process.

5.3.1 SAP R/3 directory structures

The SAP R/3 application suite used the iSeries Integrated File System (IFS). The IFS provides stream file support that is required by SAP R/3. The SAP R/3 database uses native iSeries database files located in an OS/400 library (or SQL collection).

Figure 41 shows a part of the SAP R/3 directory structure on the iSeries server. As you can see, the directory structure starts with the root (/) file system.




Figure 41. SAP directory structure

IFS consists of several file systems. Three of them are shown in Figure 41:

- File system QSYS.LIB
- File system QFileSvr.400
- Root file system

The QSYS.LIB file system is shown on the right-hand side of the directory tree. It contains the OS/400 library structure, which is used by the SAP R/3 database, the executable SAP R/3 kernel programs, work management objects, and so on.

The QFileSvr.400 file system, shown in the middle of Figure 41, allows access to information on *remote* iSeries server. The /QFileSvr.400 directory tree contains subdirectories that represent all *other* iSeries host systems that can be accessed by the local iSeries server. For more information on the functions of the QFileSvr.400 file system, refer to *OS/400 Integrated File System Introduction*, SC41-5711.

The dotted lines between some of the directories are symbolic links (or soft links). They point to another directory. Symbolic links contain only pointers to the referenced data, not the data itself. For example, the /usr/sap/<SID>/SYS/profile directory is a symbolic link that is pointing to /QFileSvr.400/<hostname>/sapmnt/<SID>/profile. The latter is where the information is actually stored. /QFileSvr.400/<hostname> identifies the iSeries server where the information is located. This <hostname> can be the name of any iSeries server in the network.

Performance tip

In many situations, you may gain a performance benefit if you change some of the default soft links. One such case is documented in SAP Note 68487.

The /sapmnt/trans directory (shown in the shaded box in Figure 41) is used by the SAP R/3 Transport Management System (TMS) to transport the SAP R/3 objects between different systems. For more information on the Transport Management System, see Chapter 15, “Change and transport system” on page 349.

If you lose the connection with the system where this directory resides, your SAP R/3 system will continue running. However, you will not be able to use the Transport Management System because it needs the /QFileSvr.400/<hostname>/sapmnt/trans directory. In an SAP R/3 landscape with multiple iSeries servers, the sapmnt/trans directory is used on only *one* of the iSeries servers (although each iSeries server contains its own /sapmnt/trans directory). All iSeries servers in the SAP R/3 landscape should use the /QFileSvr.400 directory to reference the /sapmnt/trans directory of the designated iSeries server.

The /usr directory, shown on the left-hand side of Figure 41, contains all stream files used by SAP R/3, including the profiles and SAP R/3 log files.

Other directories that are of interest are:

- **/usr/sap/<SID>/sys/profile** (which is normally linked to /QFileSvr.400/<hostname>/sapmnt/<SID>/profile): Contains all the profiles that are necessary to control your SAP R/3 system.
- **/usr/sap/<SID>/SYS/exe/run**: Contains references to SAP R/3 kernel executables in the QSYS.LIB file system.
- **/usr/sap/<SID>/<instance_name>** (for each of the defined instances): This directory contains the data, log, and work subdirectories that are used to store SAP R/3 job logs and logs of work processes for that instance. You can also find the /sec subdirectory here. This subdirectory contains the SAP Personal Security Environment files, which are used for running digital signatures.

- **/usr/sap/trans/config/<SID>**: Contains the central configuration files: `<hostname>_<system number>.cfg.` and `system.cfg.`

These configuration files contain information about all the SAP R/3 systems and the configured instances. They are used to manage the installation of the SAP R/3 systems. Every iSeries server in your landscape must be linked to the same configuration files. You need these configuration files if you want to make changes to your installation configuration.

Recommendation

We strongly recommend that you make alterations to the configuration files only through the R3MENU or by using ADD, CHG, or RMV commands from the SAP R/3 menu interface.

To view the details of the directory structures on the iSeries server, use the OS/400 WRKLNK command. The WRKLNK command contains parameters that allow you to specify the level of detail you require to be displayed from the IFS.

The WRKLNSAP command, contained in the SAP R/3 kernel library, also offers options for changing, copying, deleting, displaying, renaming, moving, printing, and working with stream files using numbered options.

5.3.1.1 SAP R/3 profiles

SAP R/3 uses several profiles to store the information about the configuration of the system and its instances.

The SAP R/3 installation program automatically creates a *start profile* and an *instance profile*. The start profile contains information about which SAP R/3 services to start when the instance is started.

The instance profile contains additional configuration parameters pertaining only to one specific instance and complements the values that were set by the default profile. For example, the SAP R/3 buffer sizes for the instance, as well as the number of work processes allowed in the instance are defined in this file.

The installation program also creates a *default profile* (DEFAULT.PFL), if this is the first instance of the SAP R/3 system. Otherwise, the existing default profile is updated with the new information. The default profile contains general configuration values for all instances in the SAP R/3 system, for example, the SAP R/3 system name (SID), the name of the database system, and the system name on which the message server is running.

All of these profiles are created in the `/usr/sap/<SID>/sys/profile` directory (normally linked to the `/QFileSvr.400/<hostname>/sapmnt/<SID>/profile` directory).

5.3.2 Sample configuration

The following example shows a configuration that consists of three SAP R/3 systems: development, testing, and production. The R/3 production system is installed across three iSeries servers as a three-tier solution. Table 4 shows the names that we assigned in this sample configuration.

Table 4. SAP R/3 system example configuration

SAP R/3 system			iSeries servers
System	<SID>	Instance	Host name
Development	DEV	00	DEVSYS
Test	TST	01	TSTSYS
Production	PRD	02 03 04	PRODDB (Database Server) PRODAPP1 (Application Server 1) PRODAPP2 (Application Server 2)

The system landscape for this example is shown in Figure 42. The landscape has three SAP R/3 systems installed on a total of five iSeries servers.




Figure 42. SAP R/3 system example configuration

The shaded area depicts the iSeries application servers that are *not* present in a two-tier SAP R/3 system landscape. Consequently, the corresponding start profiles, instance profiles, and configuration files for these application server instances would also *not* be created in the case of a two-tier system landscape.

There are several things you need to decide before starting the implementation. First, you need to decide on which iSeries server to store the /usr/sap/trans/config directory, which will contain the configuration files for each SAP R/3 system in the landscape. The /usr/sap/trans directory is also used to support transporting SAP R/3 objects between different SAP R/3 systems through the Transport Management System.

In our example (Figure 42), we decided to use the /sapmnt/trans directory on the iSeries server with the host name DEVSYS (development system). The /usr/sap/trans directories of all the SAP R/3 systems link through the /QFileSvr.400 file system to this directory on DEVSYS. They also share the configuration information in the /usr/sap/trans/config directory on DEVSYS. Accordingly, the development system is installed first.

The database server of the production system (PRD) contains the /QFileSvr.400/PRODDB/sapmnt/PRD/profile directory, as well as the profiles for the entire Production SAP R/3 system. In Figure 42, both application servers have soft links from their respective profile directories to this directory (dotted lines).

The first instance for the production system to be installed is the central instance (02) on the database server. This is followed by the instances (03 and 04) for the two application servers of the production system.

Note

The installation of multiple SAP R/3 systems (SIDs) on one iSeries machine is fully supported. Adjustments have to be made in the hardware sizing exercise to accommodate multiple SAP R/3 systems. This is a strategic decision that has to take into account that operating system and database software upgrades, as well as PTFs, will affect all SAP R/3 systems installed on the same iSeries server.

Each installation of SAP R/3 on different iSeries servers has its own set of executable code regardless of whether they are part of a single SAP R/3 system (sharing a single database) or separate SAP R/3 systems belonging to a single SAP R/3 transport domain. There is no sharing of SAP R/3 executable code across separate servers. They can (and mostly do) share profiles and transport directories. Although multiple SAP R/3 systems on the same iSeries server may share the same set of executable code, the relatively small amount of additional disk space required more than compensates for the increased flexibility in having separate executable kernels.

5.3.3 Objects created during the R/3 installation

This section gives an overview of the objects that the SAP R/3 installation creates on the iSeries server. It also looks at the functions of these objects.

For each SAP R/3 system (SID), the following iSeries objects are created:

- **Library R3400** contains objects that are common to all the SAP R/3 systems running on a single iSeries server. For example, the objects may include user spaces for storing data and indexes.
- **Library R3<rel>OPT** (by default, <rel> is the SAP R/3 kernel release level, but you may select any name) is created on each iSeries server. This library contains the SAP R/3 kernel executables that are required for running the SAP R/3 system.
- **Library R3<SID>400** is created for each SAP R/3 system. All jobs for an SAP R/3 instance run in their own subsystem environment. This library contains all work management objects required for establishing this environment.

- **Subsystem R3_<nn>**. For each instance within a R/3 system, there is an OS/400 subsystem created with an associated job description, job queue, and class (all called R3_<nn>, where <nn> is the instance number). In other words, every R/3 instance is implemented in its own iSeries subsystem.
- The following libraries are created on each database server:
 - **R3<SID>DATA** (database library): This library contains the database, including the tables, views, indexes, journals, and cross-reference files. All the application programs (ABAP code) also reside in the database. The security and availability of this library are very important because all the data resides here. In a three-tier environment, this library is located only on the database server.
 - **R3<SID>JRN** (journal receiver library): This library contains the journal receivers where all database changes are recorded. You should install this library in a separate user ASP. SAP recommends that this user ASP is mirrored at a hardware level.
- Other libraries with names starting with **R3<SID>**. These automatically created libraries contain SQL packages.
- The following user profiles with special authorities are created:
 - **R3OWNER**: Owns the executable kernel objects
 - **R3GROUP**
 - **<SID>GROUP**: Primary Group for SAP R/3 IFS objects
 - **<SID>OFR**: System officer for the <SID>
 - **<SID>OPR**: System operator for the <SID>
 - **<SID>OPRGRP**: Group profile for <SID>OPR
 - **<SID>OWNER**: Owns the database
 - **<SID><nn>**: Here *nn* is the instance number. The <SID> work processes run under this profile

The initial passwords for the user profiles are set as follows:

- User profile <SID>OFR: SAPOFR
- User profile <SID>OPR: SAPOPR
- User profile <SID><nn>: SAPnnPWD (*nn* is the instance number profile)

You should ensure that the <SID>OFR and <SID><nn> passwords are the same on all systems in a three-tier environment.

After the installation, change these default passwords because they are well-known and are a security exposure to your system.

Note

The user profiles <SID>OFR and <SID>OPR are the only ones that should be used to sign on to an SAP R/3 system. You should ensure that the <SID>OFR and <SID><nn> passwords are the same in a three-tier environment. Refer to the SAP documentation, *Installing R/3 on IBM iSeries*, for more details.

5.4 Installation overview

This section briefly describes the installation of SAP R/3 on the iSeries server. Installation of the front-end presentation is not described here.

IBM customers may choose to install the SAP R/3 themselves or use the IBM Preload on iSeries offering.

5.4.1 Preload option

There is a preload option available for the iSeries server that allows you to obtain an iSeries server directly from IBM with SAP R/3 already installed and partly customized. The following requirements must be satisfied to obtain a preload for SAP R/3 on the iSeries server:

- You must have a signed SAP contract *before* you can order your iSeries server with the SAP R/3 preload option. Before preloading, IBM verifies with SAP AG that you have a valid SAP customer number and installation number before shipping the preloaded system.
- You must submit a completed preload form together with your order for the iSeries server.

At the present time, the following preload activities are carried out:

- OS/400 is installed.
- The required PTFs are installed.
- A user ASP is configured for SAP R/3 journal receivers.
- iSeries server values related to the SAP R/3 environment are set.
- Basic TCP/IP settings on the iSeries are configured.
- The iSeries user tools for the SAP R/3 environment are installed.
- SAP R/3 systems and instances are installed.
- The Remote Function Call (RFC) kit is installed.
- The CPI-C kit is installed.
- The SAP R/3 license key is installed.
- The SAP R/3 system start is tested.
- The instance and configuration profiles are verified.
- The SAP R/3 performance profiles, based on the iSeries model, are tuned.
- The correction and transport system is initialized and checked in a "stand-alone environment".
- Additional language import (optional) is done.
- An entire system backup is done.

The following activities must be performed at the customer site when the system arrives:

1. Install the SAP R/3 Fontend on the workstation.
2. Configure and test the connection to the SAPNet - R/3 Frontend system.
3. Configure SAP R/3 users, printers, and central system log.
4. Execute client copies (optional).
5. Import supplemental languages if required.
6. Request a developer license key from the SAPNet - R/3 Frontend.
7. Configure the Transport Management System to conform to your SAP R/3 system landscape.

5.4.2 Installation tasks

This section contains a general description of the tasks required to install R/3 4.6C. Since the installation process is described in detail in the SAP publications shipped with the software, we do not go into great detail here. Be sure to obtain the latest SAP Notes specified in *SAP System Installation: IBM AS/400*.

Note: The following items that are marked with an asterisk (*) indicate actions that were already performed on the iSeries servers that are preloaded with SAP R/3 (see 5.4.1, “Preload option” on page 80). If you do not have a preloaded system, you must perform these tasks manually.

1. Before you start to install SAP R/3, complete these steps:
 - a. * Check the hardware and software requirements for the iSeries server and front-end PCs. See 5.1, “Hardware and software requirements” on page 67, for further information.
 - b. * Ensure that the necessary OS/400 PTFs that are required for the SAP R/3 installation are loaded and applied. The list of Informational APARs includes:
 - APAR II11832 for OS/400 V4R4
 - APAR II12399 for OS/400 V4R5
 - APAR II12833 for OS/400 V5R1

You can obtain the APARs through the IBM ECS link with the following command:

```
SNDPTFORD PTFID((IIInnnnn))
```

Here, *nnnnn* is the APAR number.

The APARs list is also available through the Internet at:

<http://www.as400.ibm.com/service/bms/support.htm>

Select your relevant OS/400 release from the list and search for *SAP400*. Also refer to SAP Note 83292.

Attention

Do not install SAP R/3 without first installing the required PTFs. The installation will fail because the SAP R/3 installation program checks for the presence of certain PTFs.

- c. * Set the required iSeries server values for the SAP R/3 environment. Refer to the SAP guide, *Installing R/3 on IBM AS/400*, for more information.
- d. * Set up and configure a separate user ASP on the iSeries database server for storing journal receivers. For a detailed description, see the SAP publication, *Installing R/3 on IBM AS/400*, and the IBM document *Backup and Recovery*, SC41-5304.
- e. Ensure that Electronic Customer Support (ECS) is set up on the iSeries server. We recommend that you set up the ECS connection with IBM. ECS provides an integrated set of service and support functions and provides online and remote technical support.
- f. * Configure and test the TCP/IP configuration on all iSeries servers in the system landscape. If you do not have a TCP/IP network configured, contact your network administrator about proper planning for your TCP/IP network.

For a detailed description about TCP/IP, refer to *TCP/IP Configuration and Reference V4R3*, SC41-5420.

Note

The host name and domain name entries are case sensitive. We recommend that you consistently specify these entries in lowercase and enclose them in single quotation marks ('). If the quotation marks are omitted, these entries will be converted to uppercase. This could result in the system being unable to resolve the host name and domain name.

- g. Adjust the iSeries default startup program QSTRUP to automatically start the TCP/IP services and other required subsystem and service jobs. We recommend that you first copy the default QSYS/QSTRUP program, which is shipped with your iSeries server, to a user library. Then, do the necessary modifications in this copy of QSTRUP program. Note that any changes to the QSYS/QSTRUP program may be overwritten by future upgrades to OS/400. Remember to change the system value QSTRUPPGM to point at the QSTRUP program in your user library.
- h. Review the relevant SAP Notes for installation-related updates and take action if required:
 - Note 86061 for SAP R/3 Rel 4.0A
 - Note 97815 for SAP R/3 Rel 4.0B
 - Note 108376 for SAP R/3 Rel 4.5A
 - Note 139942 for SAP R/3 Rel 4.5B
 - Note 135511 for SAP R/3 Rel 4.6A
 - Note 192231 for SAP R/3 Rel 4.6B
 - Note 300097 for SAP R/3 Rel 4.6C
 - Note 324968 for SAP R/3 Rel 4.6C SR1
 - Note 319471 for SAP R/3 Rel 4.6D
- i. Review the SAP Note *AS400 400 Latest News*, and perform the described actions if required:
 - Note 77864 for SAP R/3 Rel 4.0A
 - Note 99814 for SAP R/3 Rel 4.0B
 - Note 103805 for SAP R/3 Rel 4.5A
 - Note 139924 for SAP R/3 Rel 4.5B
 - Note 146836 for SAP R/3 Rel 4.6A
 - Note 179665 for SAP R/3 Rel 4.6B
 - Note 300105 for SAP R/3 Rel 4.6C
 - Note 324971 for SAP R/3 Rel 4.6C SR1
 - Note 319741 for SAP R/3 Rel 4.6D
- j. In a three-tier environment, connect your database server and application server with a fast communications link (fiber-optic link or a 1 GB Ethernet in case of iSeries 8xx models). You can find more information about fiber-optic connection in Chapter 7, "Connectivity" on page 123, and in *OptiConnect for OS/400*, SC41-4414.
- k. * Make a full iSeries server backup before you install the SAP R/3 code. Have a copy of this backup and the OS/400 installation CD available. Refer to Chapter 11, "Availability, backup, and recovery" on page 221, for the full backup procedure. You should also refer to the manual *Backup and*

Recovery, SC41-5304, for general information about backup and recovery on the iSeries server.

Note: After SAP R/3 is installed, perform another backup of user data.

2. Verify that the iSeries optical media drive is varied on and accessible. To verify this, insert SAP R/3 Kernel CD into the optical drive, and issue the command:

```
wrklnk '/QOPT'
```

Ensure that the contents of the CD can be displayed.

3. Install the R3SETUP program on the iSeries, using the SAP R/3 Kernel CD. The R3SETUP program is used to install the SAP R/3 system on the iSeries server and provides automatic restart options in the event of a failure during the installation process. Refer to the SAP publication *Installing R/3 on IBM AS/400* for details.
4. Copy the R3SETUP command files from SAP R/3 installation CD and install the SAP R/3 system on the iSeries server by running the appropriate R3SETUP command file (centrdb.r3s, dialog.r3s, and so on), depending on the type of installation required. The R3SETUP program will prompt you for information relevant to the installation you are performing. The information that is entered is stored in the R3SETUP command file that you are executing. This happens after the command file is automatically executed (similar to a script file), using the information that was entered.

The following installation steps are described in detail in the SAP publication *Installing R/3 on IBM AS/400*. Always use the most current documentation that accompanies your SAP R/3 software.

Three-tier environment

In a three-tier environment, the user profile of the user installing the SAP R/3 system on the application server must also exist on the database server. The passwords must be *the same* on both systems. Sign on to both systems with the same user name and password to verify this.

- a. * Configure and install the SAP R/3 system, database instance, and central instance.
 - b. * Configure and install the application server instances.
 - c. * Configure TCP/IP on the PCs, and test the connection to the iSeries server. Consult your PC Network guide for more information.
 - d. Install the front-end SAP GUI software.
5. Carry out the post-installation tasks:
 - a. * Register the SAP license key. After the installation, you must request a permanent license key for your SAP R/3 software. For the installation, you use a temporary license key that is only valid for four weeks. You can find information on how to obtain a permanent license key in the SAP documentation *Installing R/3 on IBM AS/400*.
 - b. * Install the optional components such as RFC SDK and CPI-C SDK. Remote Function Call (RFC) and Common Programming Interface - Communication (CPI-C) are both communication interfaces you can use to establish a program-to-program connection between ABAP programs and

other applications running outside of SAP R/3 on the same system or a system connected to it.

- c. Change the password for the standard SAP R/3 user profiles.
If you have a three-tier environment, you must set the password for <SID>OFR the same on all systems.
- d. * Check the initial system and instance tuning parameters for shared memory and paging as outlined in 16.4.1, “iSeries performance indicator guidelines” on page 380.
- e. * Check the initial pool sizes, activity levels, and system values on the iSeries server as outlined in 16.4.1, “iSeries performance indicator guidelines” on page 380, to tune the system.
- f. * Start the SAP R/3 system.
- g. Configure the Transport Management System. These tasks require that SAP R/3 is operational and you have a connection to a SAP GUI front end. This configuration must correspond to the SAP R/3 system landscape designed for your organization. You must perform the following steps:
 - i. Initialize TMS.
 - ii. Verify that the system global setting is set to “modifiable”.
 - iii. Verify the correctness of the TMS tables on all systems.
 - iv. Configure and test the Transport Management System.
 - v. Perform a client copy.
- h. * Verify the instance (first server) configuration.
- i. Verify that the system language is specified correctly in the instance profile parameter `zcsa/system_language`.
- j. * Verify the installation of additional languages.
- k. Verify the existence and configuration of ISDN/X.25/frame relay gateway.
- l. SAP provides an online service called SAPNet - R/3 Frontend that each customer is required to access in order to report and monitor problems. Developer keys are also available via SAPNet - R/3 Frontend. This enables a developer to define or change objects in the SAP R/3 system. The SAPNet - R/3 Frontend requires an established connection to SAP through X.25, ISDN, or frame relay. SAP must be notified about the relevant network data. Certain support functions are also available through SAPNet at: <http://www.sapnet.sap.com>

Refer to 5.7, “Configuring SAPNet” on page 90, for the detailed iSeries configuration steps for the SAPNet - R/3 Frontend. Read the *SAP Remote Connection to the Online Services* document, and follow the instructions before calling for assistance.
- m. After you connect to SAPNet - R/3 Frontend, review the SAP Note about SAP R/3 installation for the iSeries server, using BC-INS-AS4 in your search criteria. You can further specify the desired SAP R/3 release level.
- n. Configure SAP R/3 printing. SAP provides its own printing subsystem, which is independent from the underlying operating system. You must define your iSeries printer and network printers in the SAP R/3 printer subsystem in order to print from your SAP R/3 application. For a detailed description, refer to 9.5, “Example of configuring a new device to the R/3 system” on page 166.

- o. Import and apply all relevant SAP R/3 support packages, including the latest SPAM/SAINT updates. SAP R/3 support packages include Basis support packages (component SAP_BASIS), ABA support packages (component SAP_ABA), R/3 support packages (component SAP_APPL), and R/3 HR support packages (component SAP_HR). Refer to Chapter 15, “Change and transport system” on page 349, for more information on importing and applying support packages.
- p. Test your installed SAP R/3 system:
 - i. Run an online ABAP program.
 - ii. Run a batch ABAP program.
 - iii. Install and test SAP online documentation.
 - iv. Verify database consistency.
 - v. Check the SAP R/3 system log for errors, warnings and system dumps.
 - vi. Perform initial tuning tasks on the iSeries server.
- q. * Initiate a complete backup, including the installed SAP R/3 system. Refer to Chapter 11, “Availability, backup, and recovery” on page 221, for the full backup procedure. Also, refer to *Backup and Recovery*, SC41-5304, for general information about backup and recovery on the iSeries server.

5.4.3 SAP R/3 online help installation

SAP R/3 online help provides assistance and information to assist you in using your SAP R/3 system more effectively. The Application Help, Glossary, Implementation Guide (IMG), and Release Notes are delivered in HTML format. You need a Web browser on the front end to view this online help documentation. This section offers a brief overview of the installation of the SAP R/3 HTML-based online help.

Profile parameter settings

The online documentation settings are no longer stored as profile parameters. These settings are now contained in the SHLPADM1 table as variants. The *eu/iwb/...* profile parameters are, therefore, no longer required. Refer to the SAP documentation *Installing the SAP Library* for information regarding upgrading to SAP R/3 Release 4.6C.

Refer to SAP Note 203256 and the information files contained on the SAP Documentation Installation CD before you start the installation.

SAP provides three different types of online documentation:

- **HtmlHelpFile:** This form of online documentation is available for PCs running Windows 95, Windows 98, and Windows NT. It consists of compiled HTML files that you can access from a file server or directly from the CD. The online documentation is displayed with an HTML-Help viewer and requires Microsoft Internet Explorer.
- **PlainHtmlHttp:** This online documentation can be used from all supported front-end PC platforms. It uses standard HTML files in a compressed format. You can only access the online documentation from a Web server and can be viewed using a standard Web browser. Direct access from the CD is not supported.

- **PlainHtmlFile:** This type of online documentation can be used from all supported front-end PC platforms, except Windows 3.1x. It consists of standard HTML files in a compressed format. You can access the online documentation from a file server and view it using a standard Web browser. Direct access from the CD is not supported.

Customized documentation

You can access both the SAP R/3 online documentation and your own customized documentation using the SAP Knowledge Warehouse. SAP refers to this type of online documentation as DynamicHelp.

Installation

Install the online documentation on the file server (or Web server) according to the instructions in *SAP System Installation: IBM AS/400*.

To display the online documentation from within the SAP R/3 system, you need to specify the variants that will be used to display the online documentation. These settings are maintained in the SAP R/3 IMG, under **General Settings-> Variants to Adjust Help (SAP Library)**.

Select the type of online documentation that you require (HtmlHelpFile, PlainHtmlHttp, or PlainHtmlFile) by selecting the appropriate tab and create a variant to be used. For each variant, you can specify:

- Variant name
- Server name, path name, and file name where the online documentation is located
- Language version of the online documentation
- Front-end PC platform that will be used to display the online documentation
- Help area to which the online documentation refers

These settings will be used when the Application Help, SAP Library, and Glossary options are selected from the SAP R/3 Help menu.

Documentation file names

Do not use special characters and spaces in the path name and file name of the online documentation. Also, limit these names to 64 characters. The system will not be able to locate the online documentation if you do not follow these rules.

Refer to the SAP documentation *Installing the SAP Library* for complete details on installing the SAP online help documentation.

5.4.4 National language support (NLS) considerations

West European languages, such as English or German, belong to the Latin-1 group of languages. This section discusses the installation considerations for Latin-1 languages.

Note

Refer to Chapter 6, “National language support” on page 113, for details on installing the *non-Latin-1* languages, such as Latin-2 (Central and East European), Cyrillic, or double-byte character set (DBCS).

To ensure correct and consistent language support for SAP R/3, you must perform the following tasks:

1. Install the proper front-end software and set the correct code page for the front end (that is, the wincode page in the system.ini file for Windows).
2. Change the SAP R/3 `zcsa/installed_languages` profile parameter to include the language that you plan to add to the system. This parameter lists all languages that are installed in SAP R/3. For example, if German and English are installed, this parameter should read `zcsa/installed_languages =DE`.
3. Import the language.
4. Change the SAP R/3 `zcsa/system_language` profile parameter to specify the main language used by the system (that is, the language in which the logon screen appears).
5. Verify that the tables TCP0C, TCP09, TCPDB, and TCP0D contain the correct entries (see SAP Note 69337 for SAP R/3 Release 3.x and 99792 for SAP R/3 Release 4.x).
6. Update all the instance profiles to add the language key. Alternatively, add the language key to the DEFAULT.PFL profile of the SAP R/3 system.

Here is an example of profile parameters for an installation in Italy:

```
zcsa/installed_languages      = DEI
zcsa/system_language         = I
install/codepage/db/transp   = 0120
install/codepage/db/non_transp = 0120
install/codepage/appl_server = 0120
abap/locale_ctype            = IT_IT_ISO1
```

5.4.5 System date and time settings

It is important that the OS/400 system values QDATE (current date), QTIME (current Time of day), and QUTCOFFSET (Coordinated Universal Time Offset) are set correctly. Incorrect or inconsistent specification of these system values could produce incorrect date and timestamps on spooled output, IFS contents, and journal entries.

The QTIME system value should be set to the current time at your specific geographic location. Set the system value QUTCOFFSET to coincide with the time difference between your geographic area and Coordinated Universal Time (GMT).

In areas where daylight saving time (summer in the northern hemisphere and winter in the southern hemisphere) is observed, both the QTIME and QUTCOFFSET system values have to be manually adjusted to coincide with these change overs. Changes to these system values are effective immediately (no IPL of the iSeries server is required for these changes to take effect).

However, we strongly recommend that you do not adjust any of the above system values (and other related system values) while your SAP R/3 system is running. The SAP R/3 system performs certain date and timestamp checks and could “lock up” if it detects inconsistent sequences in these checks. This is especially relevant when switching from daylight saving time back to standard time. Special consideration should also be given to scheduled jobs when changing any of these system values.

These settings should also coincide with the SAP R/3 customizing settings that are maintained via the SAP R/3 IMG.

Three-tier environment

In the case of three-tier environments, the QDATE, QTIME, and QUTCOFFSET system values of the database server and all applications servers should be set the same. Inconsistencies between these system values could result in inconsistent data being written to the database.

5.5 SAP R/3 menus

There are several SAP R/3 menus that simplify the management and operation of the SAP R/3 system at the OS/400 level. These menus are used for tasks that must be performed outside of the SAP R/3 system, including the R/3 installation. Note that an SAP R/3 system has its own system management, operation, and monitoring tools to control the SAP R/3 system itself. These menus are complementary to the SAP R/3 system tools. All of these menus are contained within the SAP R/3 kernel library.

During the initial installation process, you use options from the SAP R/3 Installation menu Figure 43.

R3SETUP	R/3 Installation	System: AS23
Select one of the following:		
1. Copy Installation Files from CD 2. Install R/3 Systems and Instances 3. Work with SAP license information 4. Change the location of the /usr/sap/trans directory 5. Load RFC SDK library 6. Load CPI-C SDK library 8. Display R/3 System configuration 9. Create R/3 System objects 10. Delete R/3 System objects 11. Create R/3 Instance objects 12. Delete R/3 Instance objects		
Selection or command		
====>		
F3=Exit F4=Prompt F9=Retrieve F12=Cancel		
(C) Copyright SAP AG 1997. All rights reserved.		

Figure 43. SAP R/3 Installation menu

The *Create R/3 System objects* and *Create R/3 Instance objects* menu options allow you to make changes *prior* to commencement of the actual SAP R/3 kernel and database installation. These changes are reflected in the configuration files in the /usr/sap/trans/config/<SID> directory that control the installation process. Once the installation is complete, changing certain elements in your SAP R/3 system become more complex. This includes the following elements:

- The name of the SAP R/3 kernel library
- The ASP for the SAP R/3 database
- The ASP for journal receivers

5.6 Software upgrades

There are three types of upgrades to the software in your SAP R/3 installation. These include version or release upgrades to:

- OS/400
- SAP R/3 kernel
- SAP R/3 database

One or more of these upgrades may have to be implemented at the same time. Any upgrade has the potential of impacting the availability of the SAP R/3 system to the users. Therefore, careful planning of such an exercise is extremely important. We suggest that you first test any upgrade on your test or development system as far as possible before performing the upgrade on the production system. A well-documented upgrade strategy and review process should be an integral part of your installation.

5.6.1 OS/400 upgrade

This upgrade may involve replacing a release or version of the OS/400 operating system or replacing a cumulative PTF package. When planning the upgrade, refer to 5.4.2, “Installation tasks” on page 81, for the correct APAR that identifies the cumulative PTF package level and the additional PTFs required for running SAP R/3 on the current release of the OS/400 operating system. Make sure that you also have the latest version of the database Fix Pack available at:

<http://as400service.ibm.com/as400/startfr.html?supporthome.nsf/document/17403848>

If you run your database server and application servers on different releases of OS/400, it is required that the SAP R/3 kernel that corresponds to the earliest release of OS/400 is loaded on all systems in the SAP R/3 system landscape. For example, if you run your database server on OS/400 V4R5 and you run your application servers on V4R4, the SAP R/3 kernel corresponding to OS/400 V4R4 must be installed on the database server and the application servers.

However, we *do not* recommend that you mix different releases of OS/400 in your SAP R/3 system.

5.6.2 SAP R/3 kernel upgrade

This upgrade involves downloading the latest kernel patch from sapservX or loading the new kernel from the CD. It updates the current patch level of the SAP R/3 kernel. Each patch resolves problems that occurred since the last release of the kernel. Enhancements to the kernel are also added in this way. We recommend that you always install the latest patch level of the kernel. SAP ensures downward compatibility of SAP R/3 kernels within certain release ranges. You can find information on operating systems and kernels that are supported by SAP on the Internet at: <http://service.sap.com/dbosplatforms/>

5.6.3 SAP R/3 database upgrade

This is a very complex upgrade and should only be performed by an experienced and skilled SAP Basis consultant. There is a great deal of planning involved in upgrading an SAP R/3 release. A SAP R/3 release upgrade involves updating the SAP R/3 database, ABAP repository code, data dictionary content, and many other activities. You should not underestimate the complexity and time required for this type of upgrade.

Note

Collect all the relevant SAP Notes pertaining to the upgrade. Then integrate these into a single, clear plan of activities.

5.7 Configuring SAPNet

SAP has built up a worldwide service network that you can use to obtain services for supporting the implementation and operation of your SAP R/3 systems. This service network is called SAPNet. There are two parts to SAPNet:

- **SAPNet Web Frontend:** Provides access to various interactive services as well as knowledge databases and communication options.

- **SAPNet - R/3 Frontend** (formerly known as OSS): The primary medium for problem management.

This section discusses the configuration of SAPNet - R/3 Frontend. In the text, SAPNet - R/3 Frontend is referred to simply as “SAPNet”.

Before you start to configure your iSeries server for SAPNet, you should read and understand the online information for SAPNet provided by SAP. In particular, you should understand how to use a SAProuter and the security provisions afforded by using it. After you connect your system to SAPNet, it is on a public network. Be sure that you have taken the necessary precautions to prevent unauthorized access to your system.

The following sections explain how to setup a connection over an X.25 line and Frame relay. ISDN is another connection alternative that is not discussed here.

5.7.1 X.25 connection

To establish a remote connection via X.25, complete the following steps:

1. Obtain the necessary hardware to connect your iSeries server to an X.25 network.
2. Obtain the services of an X.25 network provider.
3. Obtain a unique IP address for your iSeries server.
4. Obtain from SAP the host name, IP address, and X.25 address of the SAP X.25 server. Also obtain from SAP the host name, IP address, and subnet mask of the SAPNet server you are using.
5. Configure the X.25 line on the iSeries server.
6. Configure TCP/IP over the X.25 line.
7. Fill out and send to SAP the *Remote Connection Data Sheet* that specifies the necessary information for SAP to enable their servers for your use.

Each of the preceding steps is discussed in more detail in the following sections.

5.7.1.1 X.25 hardware requirements

The connection to the X.25 network requires a communications controller and a communications adapter card on the iSeries server. There are a number of configurations that can be used. Table 5 lists the currently available iSeries controllers and adapters. Additional information is available in the *iSeries and AS/400e System Builder Version 5 Release 1*, SG24-2155.

Table 5. Communications controllers and adapters

Part number	Part description
MFIOP	Base communication controller. Supports two communication lines.
2623	Controller that supports up to six communication lines.
2609	EIA 232/V.24 two-line adapter. Connects to 2623.
2610	X.21 two-line adapter. Connects to 2623.
2612	EIA 232/V.24 one-line adapter. Connects to MFIOP or 2623.
2613	V.35 one-line adapter. Connects to MFIOP or 2623.

Part number	Part description
2614	X.21 one-line adapter. Connects to MF1OP or 2623.
9612	Standard EIA 232/V.24 one-line adapter. Connects to MF1OP.

The adapter you need is determined by your network provider. Choose the adapter card based on the connection requirements of your X.25 network.

An alternative to a communications adapter is a router that is attached to the same LAN as your iSeries server. In this case, the iSeries server communicates with the router over the LAN. The router itself provides the connection to the X.25 network. The supplier of the router can provide the necessary information.

Note

The use of routers is not addressed in this book.

5.7.1.2 X.25 network provider

The X.25 network provider can be your local telephone company or some other private company. X.25 networks are available worldwide. An SAP or IBM representative can help you find a suitable provider. Alternatively, use the SAP Note for your region (Table 6) for an updated list of network service providers of SAPNet connectivity. The provider you choose provides X.25 network that allows your iSeries server to remotely connect to the SAP X.25 server machine.

Table 6. SAP Notes for local network providers

Region	SAP Notes listing network providers
Asia (except Japan)	SAP Note 37946
Australia	SAP Note 102414. Please request a list of network providers in your region from the Australian Local Support
Europe/Africa/Middle East	SAP Note 33953
Japan	SAP Note 39894
America/Canada	SAP Note 40739

Your network provider assigns an X.25 network address to you. This is the number by which your system is recognized on the network. You need to provide SAP with this number to establish a connection between the SAP X.25 server and your iSeries server.

The subscription you have with your network provider determines whether you use permanent, switched, or both types of circuits, and the number you have of each. You need to know this information to configure X.25 and TCP/IP on your iSeries server. In the example shown in Figure 44, we assume that the logical channel is a switched virtual circuit for both incoming and outgoing calls.

5.7.1.3 IP address

In addition to an X.25 network address, you also need an IP address. You may already have an IP address for your iSeries server associated with your LAN communication line. This IP address is necessary for you to run your R/3 application. Because IP addresses defined on your iSeries server must be

unique, you should assign a separate IP address to your X.25 connection. A subnet mask is associated with your IP address. You need this as well as your IP address to configure TCP/IP.

SAP only connects customers with officially registered IP addresses. If you cannot obtain an official IP address, the SAP network hotline allocates an official IP address for the SAProuter subnet system on request. The addresses for the machines behind the SAProuter system can continue to use the unofficial IP addresses.

5.7.1.4 SAPNet configuration example

This section shows an example of setting up a SAPNet connection (Figure 44).




Figure 44. SAPNet configuration

In this example, the following values are assigned to the iSeries server (source host):

- iSeries server = AS23
- Network address = 0000499 (one switched circuit)
- X.25 IP address = 199.5.83.6
- Subnet mask = 255.255.255.3

SAP has provided the following information on the SAPNet and X.25 servers:

- SAPNet server = SAPSERV
- Network address = 0000222
- X.25 IP address = 199.8.2.5
- Subnet mask = 255.255.255.0
- X.25 Server = X25SERV
- X.25 IP address = 199.8.38.1

There are several other variations possible for establishing a connection to SAPNet. For example, you can establish a X.25 connection using a pSeries

server and SAProuter. If the pSeries is running router software, requests directed to the iSeries server can be forwarded to it by the pSeries server. Likewise, requests from the iSeries server can be routed by the pSeries server to the SAP X.25 server. Another variation is that you can use ISDN or frame relay in place of X.25.

5.7.1.5 Configuring the source system

This section outlines the steps to configure the components that represent the source system. Configuration of the X.25 connections includes these steps:

1. Creating a line description
2. Creating a controller description
3. Adding TCP/IP information

Creating an X.25 line description

To use the X.25 network, you need to configure an X.25 line description on the iSeries server. This is accomplished using the CRTLINX25 command. The values you need to enter for the various parameters depend on your X.25 network subscription that you have arranged with your communications service provider.

The screens of the CRTLINX25 command are shown in the following four figures. In this example, we use the name SAPNETLN for the line description that we created.

Create Line Desc (X.25) (CRTLINX25)

Type choices, press Enter.

Line description	> SAPNETLN	Name
Resource name	> LIN012	Name, *NWID
Logical channel entries:		
Logical channel identifier . .	> 001	001-FFF, *PROMPT
Logical channel type	> *SVCBOTH	*PVC, *SVCIN, *SVCBOTH...
PVC controller		Name
	+ for more values	
Local network address	> 0000499	
Connection initiation	> *LOCAL	*LOCAL, *REMOTE, *WAIT...
Online at IPL	> *YES	*YES, *NO
Physical interface	> *X21BISV24	*X21BISV24, *X21BISV35...
Connection type	> *NONSWTPP	*NONSWTPP, *SWTPP...
Vary on wait	> *NOWAIT	*NOWAIT, 15-180 seconds
Line speed	> 9600	*CALC, 600, 1200, 2400...
Exchange identifier	> *SYSGEN	05600000-056FFFFFF, *SYSGEN
Information transfer type . . .	> *UNRESTRICTED	
		More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 45. CRTLINX25: Create Line Description (Part 1 of 4)

Some of the parameters for the screens are explained in the following list. The values depend on the network provider. For any parameter, you may position the cursor on the parameter and press F1 for help.

- **Logical channel entries**

This is the logical channel configuration that the network provider has specified for your use. You may need to enter more than one channel specification here. In the example, we specified *SVCBOTH for the channel

type for the single circuit we are configuring. This means the circuit is a switched line for both input and output.

Note that if your network provider gives you a permanent virtual circuit (*PVC type) to the SAPNet X.25 server, you need to use that channel number later when you are configuring TCP/IP.

- **Local network address**

This is your machine network address on the X.25 network. This is supplied by your network provider.

- **Physical interface**

Specify the physical interface on the adapter card. The adapter card you are using was determined by the physical characteristics of the network to which you are connecting.

- **Line speed**

Specify the speed of the line.

- **Default packet size**

Specify the default packet size for both transmit and receive that your network supports.

Create Line Desc (X.25) (CRTLINX25)

Type choices, press Enter.

Extended network addressing . .	*NO	*YES, *NO
Maximum frame size	1024	1024, 2048, 4096
Default packet size:		
Transmit value	128	64, 128, 256, 512, 1024...
Receive value	*TRANSMIT	*TRANSMIT, 64, 128, 256...
Maximum packet size:		
Transmit value	*DFTPKTSIZE	*DFTPKTSIZE, 64, 128, 256...
Receive value	*TRANSMIT	*DFTPKTSIZE, *TRANSMIT, 64...
Modulus	8	8, 128
Default window size:		
Transmit value	2	1-15
Receive value	*TRANSMIT	1-15, *TRANSMIT
Insert net address in packets .	*YES	*YES, *NO
Text 'description'	*BLANK	

More...

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

Figure 46. CRTLINX25: Create Line Description (Part 2 of 4)

- **Maximum packet size**

This is the maximum packet size that the network supports for both transmit and receive.

- **Modulus**

Specify whether extended sequence numbers are used on your network. The value 8 is specified if they are not, and the value 128 is specified if they are.

Create Line Desc (X.25) (CRTLINX25)

Type choices, press Enter.

Additional Parameters

X.25 DCE support	*NO	*NO, *YES, *NEG
Network controller		Name
Switched controller list	*NONE	Name, *NONE, *ALL
+ for more values		
Idle timer	600	3-600 in 0.1 second intervals
Frame retry	7	0-64
Error threshold level	*OFF	*OFF, *MIN, *MED, *MAX
Modem type supported	*NORMAL	*NORMAL, *V54, *IBMWRAP
Modem data rate select	*FULL	*FULL, *HALF
Clear To Send timer	25	10-60 seconds
Link speed	*INTERFACE	*INTERFACE, *MIN, 1200...
Cost/connect time	128	0-255
Cost/byte	128	0-255
More...		

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

Figure 47. CRTLINX25: Create Line Description (Part 3 of 4)

Create Line Desc (X.25) (CRTLINX25)

Type choices, press Enter.

Security for line	*PKTSWINET	*NONSECURE, *PKTSWINET...
Propagation delay	*PKTSWINET	*MIN, *LAN, *TELEPHONE...
User-defined 1	128	0-255
User-defined 2	128	0-255
User-defined 3	128	0-255
Recovery limits:		
Count limit	2	0-99, *SYSVAL
Time interval	5	0-120 minutes
Message queue	*SYSVAL	Name, *SYSVAL, *SYSOPR
Library		Name
Authority	*LIBCRTAUT	Name, *LIBCRTAUT, *CHANGE...

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

Bottom

Figure 48. CRTLINX25: Create Line Description (Part 4 of 4)

Creating a controller description

After you create the X.25 line description, you must create a controller for that line. To do this, use the Create Controller Network (CRTCTLNET) command. Prompt on the CRTCTLNET command and the screen in Figure 49 shown. In this example, we name the controller SAPNETCTL, and we create it for line SAPNETLN.

```

                                Create Ctl Desc (Network) (CRTCTLNET)

Type choices, press Enter.

Controller description . . . . . SAPNETCTL      Name
Online at IPL . . . . . *YES                  *YES, *NO
Attached line . . . . . SAPNETLN              Name
Connection response timer . . . 170             1-3600 (seconds)
Text 'description' . . . . . X.25 TCP-IP Controller

                                                                Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys

```

Figure 49. Create Controller Description (CRTCTLNET)

Adding TCP/IP interface information

You can now add the TCP/IP interface information for your iSeries server. From the CFGTCP main menu, select option 1. This brings up the Work with TCP/IP Interfaces screen shown in Figure 50.

```

                                Work with TCP/IP Interfaces
                                                                System:  AS23

Type options, press Enter.
    1=Add   2=Change   4=Remove   5=Display   9=Start   10=End

    Internet      Subnet      Line      Line
    Opt  Address  Mask      Description  Type
    1_
    ___ 127.0.0.1  255.0.0.0   *LOOPBACK   *NONE

                                                                Bottom
F3=Exit      F5=Refresh   F6=Print list   F11=Display interface status
F12=Cancel   F17=Top        F18=Bottom

```

Figure 50. Work with TCP/IP Interfaces

Enter option 1 on this display. Then you see the Add TCP/IP Interface screen shown in Figure 51.

Add TCP/IP Interface (ADDTCPIFC)

Type choices, press Enter.

Internet address	199.5.83.6		
Line description	SAPNETLN	Name, *LOOPBACK...	
Subnet mask	255.255.255.3		
Associated local interface . . .	*NONE		
Type of service	*NORMAL	*MINDELAY, *MAXTHRPUT...	
Maximum transmission unit . . .	*LIND	576-16388, *LIND	
Autostart	*YES	*YES, *NO	
PVC logical channel identifier		001-FFF	
+ for more values			
X.25 idle circuit timeout . . .	60	1-600	
X.25 maximum virtual circuits .	1	0-64	
X.25 DDN interface	*NO	*YES, *NO	
TRLAN bit sequencing	*MSB	*MSB, *LSB	

Bottom

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

Figure 51. ADDTCPIFC: ADD TCP/IP Interface

On this display, enter the IP address for your X.25 adapter (199.5.83.6 in this example), the name of the X.25 line description that was previously created (SAPNETLN in this example), and the subnet mask that is associated with your IP address (255.255.255.3 in this example).

You may want to also fill in the X.25 maximum virtual circuits parameter to something other than the default value of 64. This parameter controls how many switched virtual circuits TCP/IP can use. A value of 64 means to use them all.

The number specified here must be equal to or less than the number of switched circuits you specified when the line description was created. In our example, we only created one switched circuit on the CRTLINX25 command so we can only specify one here. However, if you create more than one, you need to decide how many to allocate for TCP/IP use and enter that number on this parameter.

If you want TCP/IP to use permanent circuits in place of, or in addition to, switched circuits, enter the channel identifier number (or numbers) in the PVC logical channel identifier. The number (or numbers) you specify here must also be specified when the line description is created.

5.7.1.6 Configuring destination (SAP) systems

First you need to identify the TCP/IP configuration information for the remote systems used by the SAPNet link. Before you configure TCP/IP, you need to know some information about the SAPNet servers. Obtain the following information from SAP:

- The name and IP address for the SAPNet server you are using:
In the following displays, we use the name *SAPSERV* with the IP address 199.8.2.5. You also need the subnet mask corresponding to this IP address. We used 255.255.255.0 in this example.

- The host name and IP address for the SAP X.25 server:

In the following displays, we use the name *X25SERV* with IP address 199.8.38.1.

- The X.25 address of the X.25 server or channel number for the *PVC type connection:

In the example, we use switched lines so we need the X.25 address. We used address 0000222 as an example.

Configuring TCP/IP

Use the CFGTCP command to configure TCP/IP. Perform the following steps:

1. Add the host name and IP address of the *SAPNet server* you are using to the host name table.
2. Add the host name and IP address of the *X.25 server* to the host name table.

Note

This step assumes you are using your local host name table. If you are using a remote name server, you need to ensure that the X.25 and SAPNet servers are already in the host name table on that server. Check option 13 on the CFGTCP display to see if you are using the local iSeries host table or a remote name server.

3. Add a new TCP/IP route entry to route a request from your host to the X.25 server and then to the SAPNet server you are using.
4. Add the TCP/IP interface information for the X.25 line description that was created previously.
5. Add a new TCP/IP remote system entry for the SAP X.25 server.

Adding TCP/IP host table name entries

Figure 52 shows the main menu for the CFGTCP command menu.

CFGTCP	Configure TCP/IP	System: AS23

Select one of the following:

1. Work with TCP/IP interfaces
2. Work with TCP/IP routes
3. Change TCP/IP attributes
4. Work with TCP/IP port restrictions
5. Work with TCP/IP remote system information

10. Work with TCP/IP host table entries
11. Merge TCP/IP host table
12. Change TCP/IP domain information

20. Configure TCP/IP applications
21. Configure related tables
22. Configure point-to-point TCP/IP

Selection or command
 ===> 10

F3=Exit F4=Prompt F9=Retrieve F12=Cancel

Figure 52. Configure TCP/IP

From this display, enter option 10. Then you see the Work with TCP/IP Host Table Entries (Figure 53).

Work with TCP/IP Host Table Entries		System: AS23

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display 7=Rename

	Internet	Host	
Opt	Address	Name	
<u>1</u>	<u>127.0.0.1</u>	LOOPBACK	
		LOCALHOST	

Bottom

F3=Exit F5=Refresh F6=Printlist F12=Cancel F17=Positionto

Figure 53. Work with TCP/IP Host Table Entries

Now select option 1 and press Enter. This brings up the display shown in Figure 54. On this screen, enter the IP address for the SAPNet server, the host name of the SAPNet server machine, and a text description that describes what the new entry is for.

```

Add TCP/IP Host Table Entry (ADDTCPHTE)

Type choices, press Enter.

Internet address . . . . . '199.8.2.5'
Host names:
  Name . . . . . SAPSERV

+ for more values
Text 'description' . . . . . SAP SAPNET Server Machine

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

```

Figure 54. Add TCP/IP Host Table Entry (ADDTCPHTE)

Repeat this process for the X.25 server shown in Figure 55.

```

Add TCP/IP Host Table Entry (ADDTCPHTE)

Type choices, press Enter.

Internet address . . . . . '199.8.38.1'
Host names:
  Name . . . . . X25SERV

+ for more values
Text 'description' . . . . . SAP X25 Server Machine

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

```

Figure 55. Add TCP/IP Host Table Entry (ADDTCPHTE)

Note

The IP address for the iSeries host must be explicitly defined in the host name table.

Adding TCP/IP route information

After adding the SAPNet server and X.25 server to the host name table, you need to set up the route to go from your machine to the SAPNet server by using the X.25 server as an intermediate hop. To do this, use the CFGTCP command again. From the main menu, select option 2 (Work with TCP/IP Routes). This brings up the screen shown in Figure 56.

Work with TCP/IP Routes				System: AS23
Type options, press Enter.				
1=Add 2=Change 4=Remove 5=Display				
Opt	Route Destination	Subnet Mask	Next Hop	Preferred Interface
<u>1</u>	<u>*DFTROUTE</u>	<u>*NONE</u>	<u>10.8.90.1</u>	<u>*NONE</u>
				Bottom
F3=Exit F5=Refresh F6=Print list F11=Display type of service				
F12=Cancel F17=Top F18=Bottom				

Figure 56. Work with TCP/IP Routes

From this display, enter option 1, which brings the screen shown in Figure 57.

Add TCP/IP Route (ADDTCPRTE)			
Type choices, press Enter.			
Route destination		>	199.8.2.5
Subnet mask		>	*HOST
Type of service			*NORMAL *MINDELAY, *MAXTHPUT...
Next hop		>	199.8.38.1
Preferred binding interface	. .		*NONE
Maximum transmission unit	. . .		576 576-16388, *IFC
Route metric			1 1-16
Route redistribution			*NO *YES
Duplicate route priority			5 1-10
			Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display			
F24=More keys			

Figure 57. Add TCP/IP Route (ADDTCPRTE)

On this display, enter the SAPNet server IP address as the route destination, the value *HOST for the subnet mask, *NORMAL for the type of service, and the X.25 server IP address for the next hop. Set the smallest packet size in any gateway router between the system and the X.25 line as the maximum transmission unit.

Press Enter to save your settings.

Note

Any variation from this example configuration will require additional routes to be configured. For example, if the SAProuter is running on the host AS24 and the X.25 line is on AS23 as shown in Figure 44, you have to add a route on host AS24 (Next hop parameter), similar to the example shown in Figure 58.

Add TCP/IP Route (ADDTCPRTE)

Type choices, press Enter.

```

Route destination . . . . . > 199.8.2.5
Subnet mask . . . . . > *HOST
Type of service . . . . . *NORMAL      *MINDELAY, *MAXTHRPUT...
Next hop . . . . . > 199.5.83.5
Preferred binding interface . . *NONE
Maximum transmission unit . . . 576      576-16388, *IFC
Route metric . . . . . 1      1-16
Route redistribution . . . . . *NO      *NO, *YES
Duplicate route priority . . . . 5      1-10

```

Bottom

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

Figure 58. Add TCP/IP Route (ADDTCPRTE)

Adding TCP/IP remote system

After you add the TCP/IP interface information, you need to add a new TCP/IP remote system entry. From the main menu for CFGTCP, type option 5. This brings up the display shown in Figure 59.

Work with TCP/IP Remote System Information

System: AS23

Type options, press Enter.

1=Add 4=Remove 5=Display

Opt	Internet Address	Network Address	PVC	Reverse Charges
1				

(No remote system information)

Bottom

F3=Exit F5=Refresh F6=Print list F12=Cancel F17=Top F18=Bottom

Figure 59. Work with TCP/IP Remote System Information

Enter option 1. This brings up the display shown in Figure 60.

Add TCP/IP Remote System (ADDTCPRSI)

Type choices, press Enter.

Internet address > 199.8.38.1

Network address 0000222

PVC logical channel identifier 001-FFF

X.25 reverse charge *NONE *NONE, *REQUEST, *ACCEPT...

Bottom

F3=Exit
F4=Prompt
F5=Refresh
F10=Additional parameters
F12=Cancel

F13=How to use this display
F24=More keys

Figure 60. Add TCP/IP Remote System (ADDTCPRSI)

On this display, enter the following values:

- For Internet address, enter the IP address of the SAPNet X.25 server.
- If you are using switched circuits (as we are in this example), for Network address, enter the X.25 address of the SAPNet X.25 server.
- If you are using permanent circuits instead of switched circuits, enter the logical channel identifier for the channel to be used when connecting to the X.25 server.

5.7.1.7 Verifying the connection

Once you configure TCP/IP, you can verify the connection by using the PING command with the SAPNet server name. It is sometimes necessary to increase the WAITTIME on the PING command if the response time is greater than the default of 1 second.

Note

Note: Verify that the interface status is active by using the CFGTCP command and pressing F11 (Display interface status) on the Work with TCP/IP Interfaces display. When you work with the SAPNet site that you are using, the SAPNet site can also verify the connection from their end.

5.7.1.8 Remote Connection Data Sheet

In addition to the steps that were just discussed, you need to provide SAP with some additional information before a SAPNet connection can finally be made. SAP requires that you submit a *Remote Connection Data Sheet*. You can print copies of the Remote Connection Data Sheet from the Online Documentation CD. On this sheet, fill in the pertinent information about your company and personnel within the company who are using SAPNet. In addition, provide SAP with the necessary information about the remote connection. This includes your X.25 network number and your IP address.

SAP requests that you provide the IP address of your machine that has an X.25 connection and the IP address of your machine that is running the SAProuter. In

our examples, the machine running the SAProuter is the same machine that has the X.25 connection. Therefore, there is only one IP address for both. Once SAP has processed the remote connection information you provided, SAPNet is ready to be used.

When the IP network is different between the customer system (for example, 192.45.33.7) and the SAPNet X.25 server (for example, 199.8.38.1), you need to set the IP datagram forwarding parameter to *YES (the default is *NO). This is found in CFGTCP option 3 (Change TCP/IP Attributes) shown in Figure 61.




Figure 61. Change TCP/IP Attributes (CHGTCPA)

5.7.1.9 SAPROUTTAB file

A routing table is mandatory for the SAProuter program since SAP R/3 release 3.0E. Before the SAProuter can be used, you need to create an access (route) permission table in the /usr/sap/ directory. This file can be created using the OS/400 EDTF command.

Note
SAP Note 79411 recommends you place saproustab in the /usr/sap directory.

The file has five fields separated by one or more blanks:

- **Route permission code:** Specify P (=permit) or D (=deny). Permission S (=secure) is only permitted for the SAP protocol.
- **Source host:** Specify the IP address of the local host.
- **Destination host:** Specify the IP address of the remote host.
- **Destination service:** Specify the port number to be used.
- **Password:** Specify a password (optional).

Enter the following two entries in the saproustab file if you do not require an authorization check:

```
P * * *  
P * * 23
```

Note

An extra entry for port 23 is needed to allow a Telnet connection.

For more information on saproustab, please refer to SAP Note 30289 and the SAP online documentation CD.

5.7.1.10 Starting SAProuter

You have to start the SAProuter on the iSeries server by typing the following command:

```
SEMJOB CMD(CALL PGM(SAPROUTER) PARM('-r' '-R' '/usr/sap/saproustab'))  
JOBQ(QSYSNOMAX)
```

Alternatively, you can start and stop SAProuter from the R3MAIN menu. Select option 1 for General Task and option 9 to Start SAProuter.

For a more detailed description of setting up the Route Permission table, refer to the online documentation on SAPNet from SAP and SAP Note 79411.

5.7.1.11 Configuring for SAPNet - R/3 Frontend

Follow these steps:

1. To start your SAPNet connection, start SAP GUI on your PC.
2. Sign on to an R/3 system.
3. Start transaction **oss1** or choose **System-> Services-> SAP Service** to reach the SAPNet logon window.
4. Choose **Parameters** and click **Technical settings**. Click the **Change** push button and add parameters as shown in Figure 62.




Figure 62. Sample SAPNet route setup window

5. Be aware that your iSeries server now has two IP addresses. One is associated with your X.25 line, and the other one is associated with your LAN. You should add the LAN IP address to the SAProuter 1 section of the window. The Instance no. parameter does not refer to your R/3 system instance. Instead, it refers to the TCP/IP *service no.* used by the SAProuter running on your *iSeries* server to establish a connection with the SAP site. The default value for this parameter is 99. If the SAProuter on your *iSeries* server uses a service other than sapdp99 (3299), change this parameter accordingly (use the CFGTCP command and option 21 (Configure related tables) to check).
6. Click the **Save** push button to save your settings.
7. Log on to SAPNet. Once you are logged on, the first step is to define your new installation and service connection within SAPNet. Refer to SAP Note 86251.

Note

See SAP Note 60856 for more details on how the SAPNet - R/3 Frontend works on the iSeries server.

5.7.2 Frame relay connection to SAPNet

In this configuration example, only the configuration of TCP/IP is required on the source iSeries server.

It is assumed that the network infrastructure for the frame relay connection is provided by a network service provider and that the customer will use TCP/IP to

connect to SAPNet. Refer to Table 6 on page 92 for a list of network suppliers for SAPNet connection. In this section, we only concentrate on configuring TCP/IP on the iSeries server.

5.7.2.1 Prerequisites

To establish a remote connection, complete the following steps:

1. Obtain the services of a network provider to establish the frame relay connection.
2. Provide an official IP address for the router that will connect to your LAN. Refer to SAP Note 50430 for more information on the official IP addresses for SAP access.
3. Decide to which SAP location you will connect.
4. Configure TCP/IP on the iSeries server.
5. Configure saproustab and start SAProuter.
6. Complete and fax the Remote Data Connection Sheet to SAP.
7. Verify the connection.
8. Configure the router data for SAPNet logon.

5.7.2.2 Frame relay configuration example

In this section, the service provider has established the route from the customer site to the SAP. This network diagram illustrates the network configuration described in the following section.

The iSeries server (source host) has the following values assigned to it:

- iSeries server = AS23
- LAN IP address = 192.41.246.5
- Customer Router = 192.41.246.95

SAP has provided the following information on the SAPNet server and SAPNetrouter:

- SAPNet server = SAPSERV
- SAPNet server IP = 199.8.2.5
- SAPNet router = SAPROUTE
- SAPNet router IP = 199.8.31.1

5.7.2.3 Configuring TCP/IP for a frame relay connection

This step assumes that you are using the local host table and that SAP has verified the route from SAPNet server to the customer router. Normally the verification is done when processing the remote connection information is completed.

Add host name table entries

Figure 63 shows the main menu for the CFGTCP command.

CFGTCP	Configure TCP/IP	System: AS23

Select one of the following:

1. Work with TCP/IP interfaces
2. Work with TCP/IP routes
3. Change TCP/IP attributes
4. Work with TCP/IP port restrictions
5. Work with TCP/IP remote system information

10. Work with TCP/IP host table entries
11. Merge TCP/IP host table
12. Change TCP/IP domain information

20. Configure TCP/IP applications
21. Configure related tables
22. Configure point-to-point TCP/IP

Selection or command
 ===> 10

F3=Exit F4=Prompt F9=Retrieve F12=Cancel

Figure 63. Configure TCP/IP

On this screen, enter option 10. This brings up the Work with TCP/IP Host Table Entries screen shown in Figure 64.

Work with TCP/IP Host Table Entries		System: AS23

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display 7=Rename

	Internet	Host
Opt	Address	Name
1	_____	
_____	127.0.0.1	LOOPBACK LOCALHOST

Bottom

F3=Exit F5=Refresh F6=Print list F12=Cancel F17=Position to

Figure 64. Work with TCP/IP Host Table Entries

Now enter option 1 and press Enter. This brings up the Add TCP/IP Host Table Entry screen shown in Figure 65.

```

Add TCP/IP Host Table Entry (ADDTCPHTE)

Type choices, press Enter.

Internet address . . . . . '199.8.2.5'
Host names:
  Name . . . . . SAPSERV

+ for more values
Text 'description' . . . . . SAP SAPNET Server Machine

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

```

Figure 65. Add TCP/IP Host Table Entry (ADDTCPHTE)

On this screen, you can enter the IP address for the SAPNet server, the host name of the SAPNet server machine, and a text description that describes what the new entry is for.

5.7.2.4 Adding TCP/IP route information

After you add the SAPNet server to the host name table, you need to set up the route to go from your machine to the SAPNet server, using the Customer Site Router (IP = 192.41.246.95) as an intermediate hop. To do this, use the CFGTCP command again. From the main menu, select option 2 (Work with TCP/IP Routes). This brings up the screen shown in Figure 66.

```

Work with TCP/IP Routes

System: AS23

Type options, press Enter.
  1=Add  2=Change  4=Remove  5=Display

Route      Subnet      Next      Preferred
Opt  Destination  Mask      Hop      Interface
  1
  *DFTROUTE      *NONE      10.8.90.1  *NONE

Bottom
F3=Exit      F5=Refresh  F6=Print list  F11=Display type of service
F12=Cancel   F17=Top     F18=Bottom

```

Figure 66. Work with TCP/IP Routes

From this display, enter 1. Then you see the screen shown in Figure 67.

Add TCP/IP Route (ADDTCPRTE)

Type choices, press Enter.

Route destination	> 199.8.2.5	
Subnet mask	> *HOST	
Type of service	*NORMAL	*MINDELAY, *MAXTHRPUT...
Next hop	> 199.41.246.95	
Preferred binding interface . .	*NONE	
Maximum transmission unit . . .	*IFC	576-16388, *IFC
Route metric	1	1-16
Route redistribution	*NO	*NO, *YES
Duplicate route priority	5	1-10

Bottom

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

Figure 67. Add TCP/IP Route (ADDTCPRTE)

Enter the SAPNet server IP address as the route destination, the value *HOST for subnet mask, *NORMAL for the type of service, and the IP address associated with Customer Site Router for next hop. Set the maximum transmission unit to the minimum of any router or gateway along the route. Press Enter to save your settings.

With TCP/IP configured, you can verify the connection by using the PING command with the SAPNet server name. If the PING command was successful, SAProuter starts, and the SAPNet technical settings are configured, you can log on to SAPNet.

Chapter 6. National language support

There are nearly 30 national languages supported by SAP. As you know, different languages have different characters. For example, French has accents, German has “umlauts”, and Chinese uses symbols. As long as a set of languages can be represented by 256 separate positions without any overlap, they can share a common code page where each character is represented with *one byte* of data. Therefore, the name single-byte character set (SBCS) language. Languages that use the symbols (such as Chinese, Japanese, or Korean) require far more than 256 code points and must be represented by more than one byte. These languages are known as double-byte character set (DBCS) languages.

This chapter discusses code pages and coded character set identification (CCSID). Code pages are the specification of code points (hexadecimal value) for each character in a character set. Each code page has its own identification number within the manufacturers. For example, the IBM code page for Japanese is 0290, where the SAP code page for Japanese is 8000.

CCSID is an integrated concept that includes an encoding scheme, character set, and code page. Each CCSID has its own identification number and IBM defines CCSIDs for EBCDIC, ASCII, and Unicode. The EBCDIC CCSID for Japanese is 5026, the ASCII CCSID for Japanese is 943, and the Unicode CCSID is 13488.

Note

Unicode has only one CCSID because it is not language specific.

All objects within OS/400 are defined by CCSIDs. Therefore, this chapter mentions CCSID numbers when referring to OS/400 objects.

6.1 Single-byte character set (SBCS) languages

Table 7 shows the SBCS languages and code pages in the standard SAP installations in an EBCDIC environment.

Table 7. Standard SAP installations for EBCDIC based SBCS languages

Code page name and IBM CCSID	SAP code page number	Language supported by the code page
Latin-1 Installation (0500)	0120	Danish, Dutch, English, Finnish, French, German, Italian, Norwegian, Portuguese, Spanish, Swedish
Latin-2 Installation (0870)	0410	Czech, German, English, Hungarian, Polish, Slovak, Romanian, Slovenian, Croatian
Russian Installation (1025)	0500	English, Russian
Turkish Installation (1026)	0610	German, English, Turkish
Hebrew Installation (0424)	0800	English, Hebrew
Thai Installation (839)	0860	English, Thai

Code page name and IBM CCSID	SAP code page number	Language supported by the code page
Greek Installation (0875)	0700	English, Greek

German and English are part of the Latin-1 code page and are provided with the SAP software. This section explains the steps needed to install a SBCS language that is not in the Latin-1 code page. In our example, we will use Russian.

SAP Notes 693377(R/3 3.X releases) and 99792 (4.X releases) show further details on code page and locale relationships.

As you will see, code page numbers differ by manufacturer. Four different code pages are introduced during this section. For the Russian example, the IBM code page is 1025, the SAP code page for the application server is 500, the SAP code page for the front end is 1504, and the Microsoft code page for the PC operating system is 1251.

6.1.1 Installing a non-Latin-1 SBCS language

The installation prerequisites in the example are:

- OS/400 installed with Russian (Cyrillic) and English. If Russian is installed as the primary language, then English should be installed as the secondary language.
- The standard installation procedure for R/3 on the iSeries server must be completed (see Chapter 5, "Installation and configuration" on page 67).

6.1.1.1 Installing a front end

Perform following steps:

1. Install the Russian version of the Windows operating system of choice.
2. Set the Windows code page to 1251.
3. Install the SAP presentation server using the documentation *Installing SAP Frontend Software for PCs*.
4. Create a new front-end session, and select 1504 as its SAP code page. This is done when using the SAP Logon to create your new session. Click **New-> Advanced**. Then deselect **default code page**. Select the pull-down menu by **language**, and select the language of your choice.
5. In the autoexec.bat file, type the following line:

```
SET PATH_TO_CODEPAGE = <SAPGUI _directory>
```

For example: <SAPGUI_directory> = C:\SAPGUI
6. The appropriate conversion files must be in the SAP GUI directory. In our example, the SAP application server will use SAP code page 500 (Cyrillic EBCDIC, not to be confused with IBM 500, which is EBCDIC Latin 1). The Windows PC will use MS Windows code page 1252, which is SAP code page 1504. The names of the conversion files you have to create are 15040500.CDP and 05001504.CDP. You create them with transaction SM59. Choose **RFC-> Generate conversion tables**, and download them to the iSeries server. Send these *.CDP files by FTP in binary mode to the SAP

Frontend PC and store them in the directory given in the autoexec.bat parameter PATH_TO_CODEPAGE.

6.1.1.2 Adjusting profile parameters

You have to adjust the profiles of the instances. The path to the profiles is /usr/sap<SID>/SYS/profile. The parameters to be adjusted are listed in the following example text, along with their values as they would appear in the Russian example.

- `zcsa/installed_languages = ER`
The installed languages parameter should list all languages that are installed on your system. The original value of this parameter is DE, which you should change to ER.
- `zcsa/system_language = R`
The system languages parameter should list the main language used by this instance.

Note: This is also where you control the language that the GUI session will display for the logon screen, and the default language to be used for the users session if one is not specified on the logon screen.
- `install/codepage/db/transp = 0500`
`install/codepage/db/non_transp = 0500`
`install/codepage/appl_server = 0500`
The codepage parameters should list the code page used in this R/3 system.
- `abap/locale_ctype = RU_RU_ISO5`
The locale type parameter should list the locale type of your system. The Russian installation has a code page 0500. If you take this parameter and look in table TCP0C, you will find that the locale for a Russian application server is RU_RU_ISO5.

6.1.1.3 Verifying tables TCP0C, TCP09, TCPDB, and TCP0D

Tables TCP0C, TCP09, TCPDB, and TCP0D are updated as a result of activating the language with the RMCPINST report during the language installation procedure.

Table 8 shows the data that should be included in the TCP0C table for the Russian example.

Table 8. TCP0C settings for Russian

Platform	Language	Country	Language and country
OS/400	E	RU	RU_RU_ISO5
OS/400	R		RU_RU_ISO5
OS/400	R	RU	RU_RU_ISO5

Table 9 shows the data that should be included in the TCP09 table for the Russian example.

Table 9. TCP09 settings for Russian

Language	Code Page
E	0500
R	0500

Follow the instructions in SAP Note 15023 to have the TCPDB table contain the proper data. Table 10 shows the TCPDB settings for the Russian example.

Table 10. TCPDB settings for Russian

Character set	Character set
0500	0500

Follow the instructions in SAP Note 42305 to have the TCP0D table contain the proper data. Table 11 shows the settings for the Russian example.

Table 11. TCP0D settings for Russian

Country	DB Language
RU	R

6.1.2 Importing the language

The language import procedure is described in the guide *R/3 Language Transport*. It is the same for all platforms. After the language transport, you should once again verify the tables TCP0C, TCP09, TCPDB, and TCP0D as described in 6.1.1.3, “Verifying tables TCP0C, TCP09, TCPDB, and TCP0D” on page 115.

6.1.3 Language possibilities with EBCDIC SAP

You can have multiple <SIDs> with different SAP code pages on one iSeries server, but you cannot have one <SID> with different SAP code pages. For example, you can have an R/3 system with Russian, and you can have an R/3 system with Greek, but an *EBCDIC* R/3 system with both Russian and Greek is not supported. Support for multiple SAP code pages in one SAP system is provided in ASCII, as described 6.3, “Multiple language support (MDMP) in one SAP system” on page 122. Configuring and EBCDIC system to run with multiple SAP code pages will result in a database that cannot be migrated to an ASCII or Unicode database on the iSeries or any other SAP-supported database platform in the future.

6.2 Double-byte character set (DBCS) languages

The iSeries server has long had the ability to handle DBCS languages based on EBCDIC code pages. However, this capability was not available for SAP on the iSeries server because the current SAP DBCS implementation is based on ASCII code pages.

With the ASCII/Unicode solution available with OS/400 V4R5 and SAP R/3 Release 4.6D, it is possible to implement four different DBCS languages:

- Japanese
- Simplified Chinese

- Traditional Chinese or Taiwanese
- Korean

As a first step to complete the multiple language support for SAP on the iSeries server, an ASCII Application Server was created for the iSeries server to support the full range of languages currently available with SAP. Since this ASCII Application Server implicitly uses all DBCS resources coded within mySap.com, it automatically supports DBCS languages. The general SAP DBCS solution based on ASCII code pages has had excellent stability and reliability over the past years making this a proven approach for Asian language support.

6.2.1 ASCII application support on the iSeries server

The ability to run ASCII applications on the iSeries server (without first having to convert them to EBCDIC) has been available for several years. Lotus Domino and Lotus Notes are perhaps the most visible examples of this. A library of C “shell” functions that provide an ASCII/EBCDIC translation layer between the system and the application is available to application providers to use for this purpose.

For the SAP DBCS support, these functions were enhanced to support DBCS ASCII and now make a separate product. This product is called *ASCII runtime*. It is available as PRPQ 5799AAS. ASCII runtime is a set of C functions that provide:

- The ASCII equivalent version of a C-runtime function
- A simple translation layer before invoking the native EBCDIC runtime function
- High performance translation functions for converting data between ASCII, Unicode, and EBCDIC

This allows the application to run and manipulate ASCII data transparently while the underlying operating system and job run in an EBCDIC code page. With this approach, the SAP kernel runs in ASCII on the iSeries server.

ASCII runtime is designed to work with SAP and other business partners’ products. After installing ASCII runtime, SAP users are required to perform a few steps that are unique to the SAP environment. These steps are described in 6.2.2.1, “Installation prerequisites” on page 119.

6.2.1.1 ASCII to Unicode conversion

Data in this solution is stored in a Unicode database. To be more exact, the character data is stored in columns of the data type GRAPHIC, which is used to store Unicode data in the iSeries server database. Data represented in the ASCII application servers is converted to Latin 1 Unicode, when it is stored in the database, and converted back to ASCII, when it is loaded into the application server to be used by application programs.

When writing into the database, character data is converted from the ASCII code page of the application server to the Unicode code page of the database (UCS-2). When reading from the database, character data is converted from UCS-2 to the ASCII code page of the application server.

The conversion is done by the SAP database interface that services all types of database requests from the SAP application server. To the SAP application programs only, ASCII data is visible. See Figure 68 for further illustration.




Figure 68. ASCII to Unicode conversion

Numeric data and binary data (for example, pool and cluster table data) do not require any conversion and are stored using the appropriate data types as with the EBCDIC version.

6.2.2 Installing an SAP R/3 DBCS system

In this example of how to install an SAP R/3 DBCS system, we use Japanese. The IBM code page for Japanese is 0290, where the SAP code page for Japanese is 8000. The EBCDIC CCSID for Japanese is 5026, the ASCII CCSID for Japanese is 943, and the Unicode CCSID is 13488.

Table 12 shows the different DBCS languages you can use with R/3 on the iSeries server and their corresponding IBM and SAP code page numbers.

Table 12. Standard SAP installations for DBCS languages

Code page name	IBM code page	IBM EBCDIC CCSID	IBM ASCII CCSID	SAP code page number	Languages supported by the code page
Shift JLS (943)	01027	5035	943	8000	English, Japanese
GB2312-80 (1386)				8400	English, Simplified Chinese, Mainland China
Big 5 (950)				8300	English, Taiwanese, Traditional Chinese
KSC5601 (1363)				8500	English, Korean

Figure 69 shows the relationship between the code pages and CCSIDs in our Japanese example. Note that the SAP code page numbers used are the SAP ASCII code page numbers. See SAP Note 103935 for a complete table of the

SAP code pages and R/3 languages. All of the standard ASCII R/3 code pages are supported with the ASCII application server on the iSeries.




Figure 69. Code page and CCSID relationships

We discuss these code pages and set the CCSID in later sections.

6.2.2.1 Installation prerequisites

To install the SAP R/3 Japanese version, you need the following prerequisites:

- If OS/400 is installed with Japanese as the primary language, then English must be installed as the secondary language. This sets the OS/400 EBCDIC CCSID to 5035.
- Install ASCII runtime (PRPQ 5799AAS) on the iSeries server.
- Complete the ASCII and Unicode R/3 installation procedure on the iSeries server.

6.2.2.2 Installing ASCII runtime

ASCII runtime is installed by performing the following steps:

1. Run the RSTLICPGM command:

```
RSTLICPGM LICPGM(5799AAS) DEV(OPTXX) RSTOBJ(*PGM)
```

Here *OPTXX* is the optical drive with the PRPQ CD-ROM. The installation creates a QADRT library on the system. There is a README file inside of the QADRT library that contains further installation instructions.

Note

The default on the RSTLICPGM command for the RSTOBJ parameter is *ALL. However, because this product works on all national language versions, you need to specify *PGM for the RSTOBJ parameter to avoid an installation error message.

2. Apply all required PTFs for 5799AAS.

3. Run the following command:

```
CALL PGM(QADRT/QADRTCINF) PARM ('update')
```

You need this step to update the conversion tables and other necessary information for the SAP environments.

4. Add the QADRT library (contains all 5799AAS functions) to the library list by placing it at the bottom of the system library list (OS/400 system value QSYSLIBL) or at the top of the user library list (system value QUSRLIBL). Failing to perform this step will prevent the activation of ASCII Runtime functions. You must perform this step prior to installing SAP R/3 ASCII or Unicode. This step won't affect other applications or system programs.

6.2.2.3 SAP R/3 ASCII/Unicode installation procedure

The ASCII or Unicode R/3 installation procedure differs from the EBCDIC installation procedure. During the installation procedure, you are prompted for the SAP code page that you want to use. There will be a table in the installation instructions to help you decide what this value should be, based on the language you want to use. In our example, this would be 8000. The installation procedure saves the SAP code page you specified, along with the ASCII CCSID and EBCDIC CCSID that correspond to it. These will be used by later steps to create locales, sort tables, and other language specific objects. If the iSeries primary language is a DBCS language, then there will be an additional prompt for the library name of the English secondary language that is to be used by the SAP system. This will be either QSYS2984 or QSYS2938 and must be entered in uppercase. The library must be present on the system. This library will be specified as the SYSLIBL for the SAP subsystem so that any system messages or text retrieved by the SAP system jobs will be in English.

6.2.2.4 Installing a front end

To install the front end, install the SAP presentation server using the documentation *Installing SAP Frontend Software for PCs*.

6.2.2.5 Verifying profile parameters

The instance profile should be set correctly based on the SAP code page entered during the install prompting. The path to the profiles is /usr/sap<SID>/SYS/profile. Check the following parameters, which also show their values as they would appear in the Japanese example:

- zcsa/installed_languages = EJ
The installed languages parameter should list all languages that are installed on your system.
- zcsa/system_language = J
The system languages parameter should list the main language used by this instance.

Note

This is also where you control what language the GUI session will display for the log on screen.

- install/codepage/db/transp = 8000
install/codepage/db/non_transp = 8000
install/codepage/appl_server = 8000
The codepage parameters should list the code page used in this R/3 system.

- abap/locale_ctype = JA_JP_SJIS

The locale type parameter should list the locale type of your system. The Japanese installation has a locale of JA_JP_SJIS.

6.2.2.6 Verifying the TCP0C, TCP09, TCPDB, and TCP0D tables

Tables TCP0C, TCP09, TCPDB, and TCP0D should have been updated by running the report RSCPINST using transaction SE38 (see SAP Note 42305).

Table 13 shows the data that should be included in the TCP0C table for the Japanese example.

Table 13. TCP0C settings for Japanese

Platform	Language	Country	Language and country
OS/400A	E	JP	JA_JP_SJIS
OS/400A	J		JA_JP_SJIS
OS/400A	J	JP	JA_JP_SJIS

Table 14 shows the data that should be included in the TCP09 table for the Japanese example.

Table 14. TCP09 settings for Japanese

Language	Code Page
E	8000
J	8000

Follow the instructions in SAP Note 15023 to have table TCPDB contain the proper data. Table 15 shows the settings for the Japanese example.

Table 15. TCPDB settings for Japanese

Character set	Character set
8000	8000

Follow the instructions in SAP Note 42305 to have table TCP0D contain the proper data. Table 16 shows the settings for the Japanese example.

Table 16. TCP0D settings for Japanese

Country	DB language
JP	J

6.2.3 Importing the language

The language import procedure is described in the guide *R/3 Language Transport*. It is the same for all platforms. After the language transport, you should once again verify the tables TCP0C, TCP09, TCPDB, and TCP0D as described in 6.1.1.3, “Verifying tables TCP0C, TCP09, TCPDB, and TCP0D” on page 115.

6.3 Multiple language support (MDMP) in one SAP system

The ability to support more than one SAP code page in the same SAP system is known as Multi-Display Multi_Process (MDMP) and is described in the SAP document *R/3 Language Combination*. See the SAP CSN Note 73606 for access to this document. MDMP support is provided on iSeries with the ASCII/Unicode version. SAP blended code pages are currently not supported on the iSeries server.

Installation of an MDMP system (or the conversion of an existing single SAP code page system to MDMP) on the iSeries server is accomplished in essentially the same manner as on the other SAP platforms. The installation process is documented in CSN Note 73606. At a high level, the steps to do this are:

1. Install an SAP system in the normal manner, specifying one of the planned SAP code pages during the installation process. Or, if an existing single code page ASCII system is to be converted to MDMP, proceed to step 2.
2. Modify the instance profile or profiles to configure an additional language or languages. For example, to add Korean to a Japanese system, update the *zcsa/installed_languages* parameter as follows:

```
zcsa/installed_languages = EJK
```

3. Delete the following parameters from the profile. They are not needed for an MDMP system:

```
install/codepage/db/transp = 8000  
install/codepage/db/non_transp = 8000  
install/codepage/appl_server = 8000  
abap/locale_ctype = JA_JP_SJIS
```

4. Create the required locales. Sign on as <sid>OFR and use the CRTSAPL CL command as follows:

```
CRTSAPLCL SID(<sid>) CODEPAGE(8500)
```

Run the command for each SAP code page to be added to the system.

5. Use the RSCPINST report in transaction SE38 to activate each of the additional languages (see CSN Note 42305).
6. Use the SMLT transaction to import each of the languages as described in the *R/3 Language Transport* guide.

Chapter 7. Connectivity

This chapter describes the different connectivity options when you connect iSeries servers in an R/3 environment. It talks about optical connections, Virtual OptiConnect, TCP/IP, and 1Gb Ethernet connectivity options.

7.1 Introduction to OptiConnect for iSeries

The concept of a two-tier and three-tier iSeries configuration is presented in Figure 16 on page 43 and Figure 17 on page 44. The three-tier implementation of SAP R/3 is also referred to as a *client/server* configuration, where the application servers are the *clients* and the database server is the *server*.

OptiConnect is a fiber-optic-based solution used to efficiently implement a three-tier solution for your SAP R/3 applications. OptiConnect for iSeries consists of a set of:

- Hardware that allows multiple iSeries servers to connect to each other through fiber-optic cables
- Software that is used by R/3 to move data quickly across that connection

OptiConnect hardware and software was specifically designed for high speed transfer of data from one iSeries server to another.

The connection between any two iSeries servers is not a typical communications type connection such as LAN, ISDN, or FDDI. Rather, each iSeries server in an OptiConnect network is connected to a dedicated, shared I/O bus. Each system in this network is connected to every other system in the same network through this shared I/O bus. The connection may be best viewed as a device connection where one system is accessing another system as if it was an attached device. The data transfer rates provided by the fiber-optic link using the shared I/O bus is in the 1 Gbps range. The software used to move data across is lightweight with a normal path length much shorter than other communications methods.

Note

Information-only RPQ number 843871 describes OptiConnect in greater detail.

7.1.1 OptiMover for AS/400 (5799-FWQ)

Note

OptiMover is available on OS/400 V4R5 or earlier releases. With the release of OS/400 V5R1, it has been withdrawn from marketing.

OptiMover for AS/400 is used to move data across the fiber-optic link. This program offering contains a set of system APIs that gives the using program (SAP R/3 for example) access to the facilities of the OptiConnect hardware. OptiMover does not use DDM files as the OptiConnect program offering does. The protocol across the fiber-optics connection is a private protocol similar to a device protocol.

Through the OptiMover APIs, an agent job is started on the database for each work process. The agent job runs the R/3 remote database access code and handles requests only for the work process that established it. For all remote database transactions, the job initiates action against the database by directing API requests to its agent, which then carries out the transaction. The agent returns data to its work process by also using an API. Therefore, the OptiMover APIs allow for inter-process communications between two jobs that are running on different systems. There is no attempt to make a remote file appear local, as is the case with DDM files.

This work process-to-agent relationship stays in effect until the work process ends. At that time, the agent job is also automatically ended by the OptiMover support.

Figure 70 illustrates the use of the OptiMover APIs by the R/3 system.




Figure 70. OptiMover/400 API usage

The OptiMover program provides a good price per performance solution for SAP R/3 three-tier implementations on the AS/400 system. You can find more information about OptiMover in *OptiMover for the AS/400*, SC41-0626.

7.1.1.1 Three-tier configurations using OptiMover for AS/400

A network of iSeries servers for SAP R/3 consists of two or more machines connected together by fiber-optic cables. One machine in this network is designated as the *hub machine*. The others are referred to as *satellite machines*. Figure 71 illustrates a network consisting of one hub machine and two satellite machines.




Figure 71. A single path network

The hub machine provides the communication route for all machines on the network. This includes communications between the hub and a satellite as well as communications between two satellites. There is no processing overhead of any significance on the hub machine. If the hub machine fails for any reason, all communications routed through this hub also fail, and consequently, the network can fail.

A dual path configuration is also available. This introduces a redundant hub into the network so that if one hub fails, the redundant hub can take over, and the network remains operational. Figure 72 illustrates a network consisting of two hub machines and two satellite machines.




Figure 72. Dual path network

Virtual OptiConnect

When an iSeries server is configured into logical partitions (LPAR), OptiConnect can be configured for communication between these partitions. This is referred to as *Virtual OptiConnect* since it uses the internal hardware path to connect two logical partitions in a single iSeries server. Refer to 3.4, “SAP R/3 landscapes with logical partitioning (LPAR)” on page 47, for more information on LPAR.

7.1.1.2 SAP R/3 server location

OptiMover support does not control or enforce any restrictions on where the SAP R/3 database or application servers should be located. However, given the importance of the hub machine in the communications network and the critical nature of the SAP R/3 database server, the OptiMover hub and SAP R/3 database server are best located on the same machine. Meanwhile the SAP R/3 application servers are located on the satellite machines.

If a satellite machine fails, only the corresponding application server is impacted. The failure of the hub results in both communications and the SAP R/3 database becoming unavailable. This arrangement reduces the exposure of the total system to a single point of failure only and maximizes the availability. The system has the potential to be available as long as the hub machine is up.

On the other hand, if the database server is placed on a satellite machine, the availability of the system is impacted if the satellite (and the associated database server) or the OptiMover/400 hub machine fails. This arrangement of database and application servers is illustrated in Figure 73.




Figure 73. SAP R/3 server placement

7.1.1.3 Configuring and ordering OptiMover

OptiMover for AS/400 (5799-FWQ) can be ordered at anytime. A copy of the OptiMover program is required for each machine in the network. Some of the hardware components are *I-listed RPQs*, indicating that they can be ordered only through an IBM Representative after approval has been obtained from the OptiConnect Service Center at IBM Rochester.

Table 17 presents two examples of the components required to implement OptiMover/400.

Table 17. OptiMover/400 and RPQs

Feature	RPQ	Description	Single hub two satellites		Dual hubs two satellites	
			Hub	Each satellite	Each hub	Satellite
#5072		System unit expansion tower	1	-	1	-
#2685	843873	Bus receiver	2	-	3	-
-	843877	20m cables	4	-	5	-
#2688	843875	1063 Mb optical link card	-	1	1	1
5799-FWQ	-	OptiMover/400	1	1	1	1

To obtain the requirements needed to connect systems in an SAP R/3 three-tier landscape, you must:

1. Determine the required capacity of the iSeries servers and the network.
2. Advise the OptiConnect Services Center of the iSeries configurations and the SAP R/3 server placement.

On receiving the configuration details, the OptiConnect Services Center provides the IBM Representative with:

- Detailed product ordering information
- Connection diagrams
- Installation instruction sheets

Based on this information, the IBM Representative can determine the price and place an order for the products.

7.1.1.4 Installing OptiMover

The OptiMover software is installed from the distribution media using the RSTLICPGM command. Enter 5799FWQ as the identifier of the licensed program on this command. Adding the OptiMover software to the system creates a new library, QSOC, which contains all of the OptiMover objects.

Once you install the hardware and the software, you must start the QSOC subsystem by using the STRSBS QSOC/QSOC command.

Then, use the DSPOPCLNK command to verify the connections between the hubs and satellites. This command presents a display that shows the system name followed by the OptiMover resources for that system. If the status of the resources is *Active*, or *Varied on*, the support is installed correctly.

All of the hardware and software involved in an OptiMover/400 network are supported through normal IBM hardware and software support channels.

7.2 Gigabit Ethernet support

With V4R5 and the iSeries 270 and 8xx models, a Gigabit Ethernet connection between a database server and application server is an alternative to OptiConnect.

To use the Gigabit Ethernet, the R/3 database server should run on iSeries Model 270 or 8xx, while the R/3 application server runs on iSeries Model SB2, SB3, or 270.

The Gigabit Ethernet, which is a viable alternative for OptiConnect, uses TCP/IP to move the data over the dedicated Gigabit Ethernet connection. Consider the following points when choosing between a Gigabit Ethernet connection and an OptiConnect connection:

- The performance that Gigabit Ethernet offers matches that of OptiConnect.
- The price for Gigabit Ethernet is lower than for OptiConnect.
- Industry standard switches can be used with Gigabit Ethernet as well as PCI-based cards.
- No extra software is needed for a Gigabit Ethernet.
- Gigabit Ethernet runs only on iSeries 8xx and 270 models.

As with OptiConnect, the communication between the database server and the application server should be on a dedicated network, for best performance (Figure 74).




Figure 74. Three-tier configuration with 1 Gb Ethernet

An additional network card is used to connect the two servers to the rest of network.

7.3 TCP/IP

This section covers the latest iSeries TCP/IP improvements and tips.

7.3.1 Performance tips

The following tips will enable the iSeries server to improve the TCP/IP handling.

7.3.1.1 Domain name server

If you use a domain name server, change your configuration to search the local name server (host table) first. To do this, use the `CHGTCP` command, option 12 (see Figure 75).

Change TCP/IP Domain (CHGTCPDMN)

Type choices, press Enter.

Host name 'ASM23'

Domain name 'ASAA.IBM.COM'

Host name search priority . . . *LOCAL *REMOTE, *LOCAL, *SAME

Domain name server:

Internet address '199.4.191.76'

'199.4.191.75'

Bottom

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel

F13=How to use this display F24=More keys

Figure 75. Change TCP/IP Domain (CHGTCPDMN)

The Host name search priority parameter defines the order that a name server is searched. Maybe your value was to `*REMOTE`, meaning that a remote name server is always searched first.

Change this to search the local name server first. If a match is not found, the remote name server is searched. This has the effect in R/3 of saving a trip to and from your remote name server every time you look up the local system name, since the local system is defined in the local name server.

7.3.1.2 Send and receive buffer size

The TCP/IP send and receive buffer sizes are set by typing the command Change TCP/IP Attributes (CHGTCPA). The TCP/IP send buffer size specifies the maximum amount of data that is not sent, which can be queued for a given application (TCP/IP connection). If you make the send buffer too small, the application will have to be blocked (put to sleep) if it tries to send data after the send buffer size has been reached. On the other hand, making the send buffer too large, may unnecessarily consume a large amount of storage, resulting in excessive page faulting.

The TCP/IP receive buffer size specifies the maximum amount of received data that can be queued waiting for the application to process it. Once the receive buffer is full, TCP/IP will stop sending window updates, forcing the sending application to stop sending. Making the receive buffer too large may unnecessarily consume a large amount of storage, resulting in excessive page faulting.

In general, you should set the send and receive buffer to the maximum amount of data that the application streams in one direction before waiting for data to be received from the remote application. Also, the line speed should be taken into consideration. Having a large send and receive buffer on a fast line can significantly improve performance. However, increasing the buffer size on a slow line could actually degrade throughput with overall system performance.

Note

Optimal send and receive buffer size is customer specific. It depends on the application you're using with R/3, the network, and the size of the iSeries. The default send and receive buffer size works best for most customers.

7.3.2 TCP/IP improvements

Recent enhancements to Transmission Control Protocol/Internet Protocol (TCP/IP) relevant to the R/3 environment pertain to:

- TCP/IP over OptiConnect and Virtual OptiConnect
- Improved scalability and performance
- Integrated virtual private network (VPN)

7.3.2.1 TCP/IP over OptiConnect

With OS/400 V4R4 and later, TCP/IP can be configured for OptiConnect, the high speed communication link. OptiConnect allows data transfer functions, such as R/3 remote client copy or FTP, to use the high-speed optical link at the rate of up to 1063 Mbps.

When connecting to two or more iSeries servers using TCP/IP over OptiConnect, you can create the connection as an emulated LAN or as a *point-to-point unnumbered connection* (subnet mask 255.255.255.255). Although either of these methods can be used with SAP R/3 implementation, the point-to-point unnumbered configuration is typically the common choice because:

- It is simple to implement.
- It allows protection of the optical bus from user activity.

This section provides details and examples of the unnumbered connection.

Configuration example

Figure 76 shows a configuration with three iSeries server connections with optical link and attached to a local area network (LAN). The IP addresses shown are used only as an example.




Figure 76. TCP/IP over OptiConnect point-to-point

By using an IP addressing scheme for the OptiConnect interface different than the one used for the LAN connection (10.1.1.x vs. 10.2.4.x), the optical bus is effectively isolated from direct access by any users from the LAN. Instead of using the unnumbered mask of 255.255.255.255 for the OptiConnect interface, you may also configure a mask definition that would restrict the last octet of the IP address to a finite, small group of values, possibly 252 to 255. To prevent serious network problems, only qualified network specialists should be responsible for these configurations.

In a point-to-point unnumbered configuration, a TCP/IP interface must be created for each pair of OptiConnect hosts along with the actual LAN resource. In the R/3 environment, a host name must also be added into the TCP/IP host table. In this example, the Add TCP/IP Interface (ADDTCPIFC) and Add TCP/IP Host Table Entry (ADDTDPHTE) commands are used as shown here:

1. On the AS1 system, add a TCP/IP interface for LAN (Figure 77).

Add TCP/IP Interface (ADDTCPIFC)

Type choices, press Enter.

Internet address	10.2.4.164	
Line description	TRNLINE	Name, *LOOPBACK...
Subnet mask	255.255.255.0	
Associated local interface . . .	*NONE	
Type of service	*NORMAL	*MINDELAY, *MAXTHRPUT...
Maximum transmission unit . . .	*LIND	576-16388, *LIND
Autostart	*YES	*YES, *NO
PVC logical channel identifier		001-FFF
+ for more values		
X.25 idle circuit timeout . . .	60	1-600
X.25 maximum virtual circuits .	64	0-64
X.25 DDN interface	*NO	*YES, *NO
TRLAN bit sequencing	*MSB	*MSB, *LSB

Bottom

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

Figure 77. Add TCP/IP Interface (ADDTCPIFC) for LAN

2. On the AS1 system, add a TCP/IP interface for the OptiConnect bus using unnumbered subnet masking (Figure 78).

```

                                Add TCP/IP Interface (ADDTCPIFC)

Type choices, press Enter.

Internet address . . . . . 10.1.1.11
Line description . . . . . *OPC          Name, *LOOPBACK...
Subnet mask . . . . . 255.255.255.255
Associated local interface . . . 10.2.4.164
Type of service . . . . . *NORMAL      *MINDELAY, *MAXTHRPUT...
Maximum transmission unit . . . *LIND    576-16388, *LIND
Autostart . . . . . *YES              *YES, *NO
PVC logical channel identifier
+ for more values
X.25 idle circuit timeout . . . 60      1-600
X.25 maximum virtual circuits . 64      0-64
X.25 DDN interface . . . . . *NO       *YES, *NO
TRLAN bit sequencing . . . . . *MSB     *MSB, *LSB

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

```

Figure 78. Add TCP/IP Interface (ADDTCPIFC) for the OptiConnect bus

3. On the AS1 system, add an Internet address and its associated host name used for the *LAN connection* to the local host table (Figure 79).

```

                                Add TCP/IP Host Table Entry (ADDTCPHTE)

Type choices, press Enter.

Internet address . . . . . > '10.2.4.165 '
Host names:
  Name . . . . . 'AS1'

+ for more values
Text 'description' . . . . .

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

```

Figure 79. Add TCP/IP Host Table Entry (ADDTCPHTE) for LAN

4. On the AS1 system, add an Internet address and its associated host name used for the OptiConnect connection to the local host table (Figure 80).

Add TCP/IP Host Table Entry (ADDTCPHTE)

Type choices, press Enter.

Internet address > '10.1.1.11'

Host names:

Name 'AS1_OPTI'

+ for more values

Text 'description'

Bottom

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

Figure 80. Add TCP/IP Host Table Entry (ADDTCPHTE) for OptiConnect

To do the same activities on the AS2 system, enter the following command statements in the order given:

```
ADDTCPIFC INTNETADR(10.2.4.165) LIND(TRNLINE) SUBNETMASK(255.255.255.0)
ADDTCPIFC INETADR(10.1.1.12) LIND(*OPC) SUBNETMASK(255.255.255.255)
ADDTCPHTE INTNETADR(10.2.4.165) HOSTNAME(AS2)
ADDTCPHTE INTNETADR(10.1.1.12) HOSTNAME(AS2_OPTI)
```

To do the same activities on the AS3 system, enter the following command statements in the order shown:

```
ADDTCPIFC INTNETADR(10.2.4.166) LIND(TRNLINE) SUBNETMASK(255.255.255.0)
ADDTCPIFC INTNETADR(10.1.1.13) LIND(*OPC) SUBNETMASK(255.255.255.255)
ADDTCPHTE INTNETADR(10.2.4.166) HOSTNAME(AS3)
ADDTCPHTE INTNETADR(10.1.1.13) HOSTNAME(AS3_OPTI)
```

7.3.2.2 TCP/IP over Virtual OptiConnect

When the iSeries server is configured into logical partitions, TCP/IP must be configured for communication between these partitions. Because the logical partitions are independent processing units, the configuration of TCP/IP for LPARs is similar to the configuration of TCP/IP for multiple iSeries servers over OptiConnect.

The most common method of configuring TCP/IP over Virtual OptiConnect is the point-to-point unnumbered configuration. As with OptiConnect, use the Add TCP/IP Interface (ADDTCPIFC) command to configure the IP address interfaces for Virtual OptiConnect. Each logical partition must have an IP address assigned to the internal optical bus in this manner (Figure 81).




Figure 81. Virtual OptiConnect: TCP/IP implementation

7.3.2.3 Scalable TCP/IP performance improvements

In OS/400 V4R4 and higher, the TCP/IP performance has been improved by redesigning the protocol stack and by the support of two new TCP/IP industry standards, which are specified in two Request for Comments (RFCs, which are not to be confused with Remote Function Calls):

- **RFC 1191: Path Maximum Transmission Unit (MTU) Discovery.** This RFC specifies a technique for dynamically discovering the largest datagram size of an arbitrary Internet path.
- **RFC 1323: TCP extensions for high performance.** This RFC specifies a set of TCP extensions to improve performance over high-speed communication links such as fiber optic link.

For additional information on these and other RFCs, please visit the Web site at:
<ftp://ftp.isi.edu/in-notes/>

7.3.3 Integrated virtual private network (VPN)

The VPN protocol provides more secure inter-site networking through the Internet by using built-in security protocols Layer 2 Tunneling Protocol, Internet Key Exchange (IKE), and Internet Protocol Security (IPSEC).

In V4R4 and above, IBM Cryptographic Access Provider licensed program is a prerequisite for integrated VPN. For more information on VPN, please refer to the ITSO Web site at: <http://www.redbooks.ibm.com>

Chapter 8. Data porting

This chapter covers data porting from legacy applications into an SAP R/3 environment. It discusses the steps necessary to perform data porting and provides examples of using the data porting tools.

8.1 Concept of data porting to SAP R/3

If you implement a new application software such as SAP R/3, one of the major tasks is to migrate the data from existing applications into the new environment. This activity normally is done only once. But in some cases, you still need to run some of the existing applications in addition to the R/3 application. Therefore, data exchange between the two environments may be performed regularly. All of these activities are called *data porting*.

The implementation of SAP R/3 normally begins with the installation of the new software on a development or test system. This development or test system is the basis for developing new business processes, organizational changes, and data transition. It is also used to educate the users who work with the new system. In the meantime, the development or test system is customized and configured according to the requirement, and sometimes an additional specific function must be developed. After the test is completed, the test system with all of its parameters, additional specific functions, and its file is moved to the new production environment.

SAP R/3 provides a facility to transfer the existing application data into the SAP R/3 database using a function called *batch input*. This function reads the data from a sequential file and stores it into the SAP database using normal R/3 interactive transactions. The sequential file that contains the data to be imported should have the following characteristics:

- All data must be in character format.
- Data must have the logical structure expected by the batch input program.

This sequential file is used as a transfer medium between the existing application and the SAP R/3 database. You are responsible for producing this sequential file by using a data transfer program that reads the existing data files, checks and converts it, and exports it into the sequential file. You can either write your own data transfer program or use data porting tools for this purpose. The data porting concept is shown in Figure 82.




Figure 82. Concept of data porting to SAP R/3

8.2 Writing programs for data porting to SAP R/3

To perform data porting to an SAP R/3 environment, data transfer programs and batch input programs are required.

8.2.1 Data transfer program

The objective of a data transfer program is to produce the sequential file required by the batch input program from an existing application database. In other words, the resulting sequential file must have a structure that is expected by the batch input program.

Basically, the data transfer program performs the following tasks:

- Checks to see if the data records need to be exported.
- Converts the data records if necessary. For example, if the data type or length is not the same as expected in the target file, the conversion is done.
- Exports the data to the sequential (target) file.

In addition to these basic tasks, the data transfer program initializes the SAP format (data structure) in the sequential file with the special *no-data character*. If a batch input program finds the no-data character in the field, the program allows the field to default to its standard value in the SAP transaction that contains the field. By default, the no-data character is “/”.

To write the data transfer program for batch input, use the following procedure:

1. Analyze the data that is required by the batch input program. SAP provides a facility to generate a data structure for SAP tables in the COBOL, PL/1, or C languages. You can incorporate this data structure into your data transfer program. To generate the data structure, select **Environment** and select the Generate table description in the ABAP Dictionary.

Note

The term ABAP/4 is used in R/3 3.1H and earlier releases. After this release, the term ABAP/4 is replaced by ABAP.

2. Analyze the SAP format (data structure) for the existing application and determine which fields to transfer and map to which field in the SAP system.
3. Determine the conversion rules. The batch input program requires that:
 - The data field must be in character format.
 - No data field is longer than those in the SAP format.
 - If the field is shorter, left-justify it and fill in the right end with blanks.
4. Write the data transfer program. It can be developed in ABAP or other external languages such as COBOL, RPG, and so on.

8.2.2 Batch input program

Batch input is a technique of transferring data into the SAP system from other SAP systems or non-SAP systems by carrying out normal SAP transactions just as a user does. When you start batch input, the system runs the transaction automatically. It is, therefore, suitable for entering a large amount of data that is already available in electronic form.

This technique offers the following advantages:

- No manual interaction is required during the data transfer.
- Batch input enters data into an SAP system using the same transaction that interactive users do. It goes through the checks and controls that apply to data entered by a normal interactive means so that it ensures data integrity.

The batch input program is written in ABAP and basically performs the following tasks:

- Reads the data to be imported to the SAP database from a sequential file.
- Converts the data record if necessary.
- Simulates an SAP transaction to enter the data into an SAP database.

There are three methods of batch input processing:

- The first method:

The program reads the external data to be entered into the SAP system and stores it in the “batch input session”. To generate the session, the program uses the BDC_OPEN_GROUP, BDC_INSERT, and BDC_CLOSE_GROUP function modules. Then you can either explicitly start and monitor the session or have the session run in the background. To do this, on the R/3 GUI, choose **System-> Services-> Batch input**.

- The second method:

The program uses the CALL TRANSACTION USING statement to run the SAP transaction.

- The third method:

The program uses the CALL DIALOG statement. This is not recommended because it is outdated and more complex than the other methods.

All of the preceding methods use the SAP format (data structure) called BDCDATA in the ABAP Dictionary for holding the data to be entered into the SAP system. They also use the actions necessary to process the data.

The following statement is used to declare the SAP format (data structure) in the ABAP program:

```
DATA: BEGIN OF <bdc-table-name> OCCURS <occurs-parameter>.  
      INCLUDE STRUCTURE BDCDATA  
DATA: END OF <bdc-table-name>
```

Figure 83 shows the batch input technique.




Figure 83. Batch input technique

SAP provides a ready-to-run standard batch input program for most of the SAP applications. If necessary, you can also create your own batch input program. To do this, analyze the SAP transaction as shown in the following steps:

1. Go through the application function just as you are doing a normal transaction.
2. Note the program names and display (dynpro) number. Do this by selecting **System-> Status**.
3. Note any field names, check boxes, or buttons that require input. You can do this by pressing F1 (Help) on the field, check box, or button, and choose **Technical Info**.
4. Note the display (dynpro) sequence and function codes.
5. Create a BDC table structure.

The structure of BDCDATA is described as follows:

Field name	Type	Length	Description
PROGRAM	CHAR	8	BDC Module Pool
DYNPRO	NUMC	4	BDC Dynpro Number
DYNBEGIN	CHAR	1	BDC Dynpro Start
FNAM	CHAR	35	BDC Field Name
FVAL	CHAR	80	BDC Field Value

Every display that is processed in the transaction is identified with a BDCDATA record that uses the following fields: *Program*, *Dynpro*, and *Dynbegin*. Then it is followed by other BDCDATA records that use the *Fnam* and *Fval* fields. These records are used to enter values such as:

- Data that is entered into the display field
- Function code that is executed in the transaction, such as *Save* (using *Fnam* BDC_OKCODE)
- Cursor position (using *Fnam* BDC_CURSOR)

In this chapter, you can find an example of a batch input program that shows how the batch input data structure may appear.

8.3 Data porting services

IBM can assist you in connecting to data porting services through iSeries Porting Team in PartnerWorld for Developers. Their Web site is at:

<http://www.iseries.ibm.com/developer/porting/>

Information about non-IBM porting services and tools from iSeries Business Partners is available on the Web at: <http://www.ibm.com/software/>

8.4 Data porting tools

When you start implementing an SAP R/3 system in a legacy environment, specific one-time-use programs may be developed for data porting purposes. However, you may save time and money by using ready-to-use data porting tools, especially if you are performing data porting from a complex application environment with large databases. Using such a tool will, in many cases, result in more efficient and faster data porting compared to developing many one-time-use programs.

This section examines some of the data porting tools available on the market today.

8.4.1 Legacy System Migration Workbench (LSMW)

The LSM Workbench is an R/3 based tool that supports the one-time or periodic transfer of data from non-SAP or legacy systems to SAP R/3. LSMW uses standard R/3 interfaces.

The LSM Workbench helps in organizing data migration project and provides guidance through the process by using a clear sequence of steps. The most common conversion and migration rules are predefined with LSM Workbench.

Reusable conversion rules assure consistent data conversion for different data objects. Figure 84 shows the steps in the data migration.




Figure 84. Data migration in three steps

The LSM Workbench covers the tasks (Figure 85):

- Reads the legacy data from one or several files (such as spreadsheet tables and sequential files).
- Converts the data from source format to target format.
- Imports the data using standard interfaces (Batch Input, Direct Input, BAPI, and IDoc).




Figure 85. How LSMW works

Some benefits for LSM Workbench (LSMW) users are:

- LSMW is free of charge for all SAP customers and business partners.
- It provides maximum data consistency.
- LSMW is supported by SAP via Online Service System: Component XX-LSM.
- It leads you smoothly through all the steps of data migration.

- It can be downloaded easily and quickly from SAPNet.
- LSMW meets your requirements because it is a highly flexible tool.
- It is independent from R/3 releases, platforms, and the kind of data to be migrated.
- It can be used without deep knowledge of ABAP.
- LSMW comes with different control functions.
- It allows reusability of data mapping and conversion rules.
- It can be used in conjunction with DX-Workbench.

8.4.1.1 Installing LSM Workbench

By installing LSMW, no objects are imported that are also part of the standard version delivered for your R/3 system. Installing LSMW, therefore, does not affect the functions of the R/3 system in any way.

The prerequisites are:

- The SAP R/3 system must be on one of the following maintenance levels: 4.0A, 4.0B, 4.5A, 4.5B, 4.6A, 4.6B, and 4.6C.
- The R/3 correction and transport system has been set up correctly and tested.
- The following background jobs are released (see System A Services A Jobs A Job overview (be sure to enter '*' in the Start after event field)):
 - RDDIMPDP in client 000
 - RDDIMPDP_CLIENT_<client> in the remaining clients

Otherwise, start report RDDNEWPP as user DDIC in all clients.

To install LSMW, perform following steps:

1. Copy the archive LSMW170.CAR to an arbitrary directory <sourcedir>.
2. Switch to the /usr/sap/trans directory:

```
cd /usr/sap/trans
```

3. Unpack the archive named LSMW170.CAR:

```
CAR -xvf <sourcedir>/LSMW170.CAR
```

4. Switch to the /usr/sap/trans/bin directory:

```
cd bin
```

5. Import the transport request named U46K900170:

```
tp addtobuffer U46K900170 <SID>
tp import U46K900170 <SID>
```

Distribution of authorization profiles

Logon as user <SID>OFR and create the file LSMW.CMD with the following contents:

```
import
client cascade = yes
file = '/usr/sap/trans/data/R900170.U46'
including 'R3TRTABU'
including 'R3TRTDAT'
```

Logon as user <sid>adm resp. <sid>ofr and run the following command:

```
R3trans '-u 1 LSMW.CMD'
```

Check the file trans.log in the current directory. The maximum admissible return code (see end of trans.log) is 4.

Resetting the buffers

Reset the buffer using the following command in the OK field:

```
/ $SYNC
```

For LSMW latest version updates and additional information, see:

<http://service.sap.com/LSMW>

8.4.1.2 Basic LSMW steps

This section provides an introduction to LSMW and shows simple LSMW screens, with steps and functions.

To start working with the LSM Workbench, use transaction LSMW (Figure 86).




Figure 86. Defining the project, subproject, and object

On the initial screen, you can create a new project, corresponding subprojects, and objects by clicking **Edit-> Create new entry**.

On the initial screen shown in Figure 86, *All objects* provides a list of all projects created already. *My objects* displays a list of all objects you created personally. *All objects of the project* displays all objects of the selected project as tree structure. *Project documentation* displays any documentation written for the individual pop-ups and processing steps. You can print the project documentation out, send it, and save it in various file formats.

Figure 87 shows an example for a project with several subprojects and objects. This representation is displayed by clicking the *All objects of the project* button.




Figure 87. Project structure

After you select an object, press ENTER or CONTINUE to go to the interactive process guide. Here you are taken through the individual steps of data migration.




Figure 88. Selecting a migration step

In the screen shown in Figure 89, the object type and import technique are selected.




Figure 89. Maintaining the attributes

On this display, you can perform the following steps:

1. Name your object. By entering data into field Owner, add the project to the list of all projects you created. You can display it afterwards in the initial screen under My objects.

2. Choose whether data transfer is one-time or periodic. In the case of periodic transfer, files cannot be read from the PC. This processes the step Frame program for the periodic data transfer.
3. Flag whether the file names are system dependant (this gives you the chance to enter file names later per system ID).
4. Select the object type and import technique. Here, F4 help is available for the input field. This help displays the relevant lists from which you can select the objects.

In the display shown in Figure 90, the data structures are defined for the legacy system, so that they can be mapped in the R/3 system. You define the structures of the object with a name, description and the hierarchical relationships.




Figure 90. Entering the structure for the legacy system in R/3

In the pop-up window, click **Change**. You can now define, change, relink, or remove structures. All these functions are available via pushbuttons.

When you define more than one structure, a pop-up display appears that queries the relationship between the structures same level/subordinated.

Figure 91 shows the rules definitions for the source fields to target fields. It also defines how the field contents will be converted.




Figure 91. Defining conversion rules and field mapping

The display in Figure 92 shows the generated conversion code in ABAP.




Figure 92. Generated conversion program

Data from PC applications or data already converted onto a PC-based file can also be converted to SAP R/3 using LSMW. The display in Figure 93 shows the interface for PC data.




Figure 93. Interface for PC data

For additional help and detailed examples of how to use LSMW, refer to:
<http://service.sap.com/LSMW>

8.4.2 MIDAS400

One of the tools for porting data from iSeries databases to SAP R/3 is MIDAS400, developed by IT-Services and Solutions. MIDAS400 can migrate data from any OS/400 application into the SAP R/3 system. The data migration includes the following processes:

- Defining field relationships
- Exporting data
- Importing data
- Permanent data interface

8.4.2.1 Defining field relationships

For each field, a relationship between the OS/400 application and SAP R/3 is created. The relevant SAP R/3 data fields are predefined in a MIDAS400-supplied field table. The user specifies the corresponding OS/400 data fields. MIDAS400 checks the OS/400 fields and shows the field type and length. If the type and length are the same in both environments, the user may choose to specify a migration rule that controls the migration process. If the type and length differ, such a migration rule must be specified. In addition, fixed values, default values, or translation tables can be used.

The process of defining field relations requires applicable knowledge of both the OS/400 application and SAP R/3. It is the basis for the whole migration process. By defining field relations within a table, a documentation of the migration process is created and updated automatically with every change that is made.

For data migration with MIDAS400, you can create different projects. Before you create a new project, you need to copy all field information from the SAP R/3 system. This is realized by ABAP programs that scan the required data for field

information of the relevant SAP R/3 version and save this information to files. If SAP R/3 is installed on a different platform, all necessary files can be copied to the iSeries platform using file transfer.

After you select one of the SAP R/3 data types, you need to specify the corresponding primary file of the existing OS/400 application. For each OS/400 data item, one record is created. In case there are additional files that need to be accessed to extract migration data, these files must be declared as secondary files. Secondary files are linked to their primary file through keys. The secondary file can be linked to the primary file through a maximum of seven key fields. MIDAS400 checks both the linked files and the key fields.

Cyclic fields

Some SAP R/3 data types contain so-called *cyclic fields*. These fields can be defined more than once per cycle. This allows, for example, procurement or sales order or material master items to have different amounts of quantity units assigned to them. If this information is not provided by the primary file, secondary files can be used as well. However, single records are not addressed and sequentially processed through key assignment. It is rather the particular group of records that is dealt with in this way.

Cyclic fields (loop fields) are found in several SAP R/3 dynpros. Therefore, for a special data record, more than one value per field can be entered here. In this case, two types of fields have to be clearly distinguished for data migration:

- **Cyclic fields in the primary file:** All values for these fields stem from the primary file.
- **Cyclic fields in the secondary file:** All values for these fields stem from a secondary file that is linked to the primary file through a key.

In both cases, each record containing cyclic fields is written out several times until all cycles are completed.

Fixed value overlay

In some cases, fields for data migration must be modified before transferring them into the SAP R/3 system. Besides migration rules, MIDAS400 also provides the feature to overlay an OS/400 field with a fixed value.

Field transfer depending on other fields

In data migration, OS/400 fields can be transferred under certain conditions. If the condition is true, field contents for this data record are transferred. If the condition is false, the field is marked empty.

8.4.2.2 Data export

For data transfer and test purposes, the OS/400 data files must be exported. Export parameters control processing of the data within the SAP R/3 system. For data export, the selected master and transaction records are read. Then, the data fields are converted according to the field relation table and are finally written to a sequential output file.

To start data export, select a data type from all available data types. Assign a primary file and possibly linked secondary files in use to the selected data type. The data type needs to have at least one view edited by the user. Views must contain entries for all required fields. To prevent errors, views should be checked with the field check option.

For data export, several export parameters are required. You can choose which records of the primary file to export. This enables you to sort out records not to be exported to SAP R/3, or to split up large datasets by processing data export several times using different select criteria each time. With this, you can export different views.

For selection, you can enter select criteria manually on the screen or use predefined select criteria. For batch processed data export, only predefined select criteria are allowed. During dialog processing, the number of selected records are shown on the screen. After this selection, the relevant records of the primary file are processed sequentially. Data fields are read referring to the defined field maps, converted according to the migration rules and written to the output file. In case of errors during the process, error records are written to the error file.

8.4.2.3 Data import

For data import, a “batch input session” is created to which the sequential output file containing the exported data is written. The session is processed in much of the same way as a set of on-line transactions. The programs use the same logic tests and checks for the data as during dialog processing. Through this technique, data consistency is assured for imported batch data in the same way as for data manually entered by the user.

If data import takes place on a different platform than data export, the exported file needs to be transferred to the SAP R/3 system environment using data transfer.

8.4.2.4 Permanent data interfaces

MIDAS400 also provides the facilities for setting up permanent data interfaces between OS/400 applications and SAP R/3. These interfaces are defined in the same way as they are for data migration. Data transfer using an interface is activated by the sender when time or event controlled data export is started. The receiver waits until the sequential data file is sent by the data export process of the sender. After data processing completes successfully, the receiver verifies the correct data transfer to the sender. At this point, the sent data file can be archived or deleted.

Like a onetime migration, this is done in three steps:

1. Define field relations
2. Export data
3. Import data

The first step is performed only once, where steps two and three are used for permanent data transfer. As a general rule for permanent data transfer, two opposite directions must be distinguished:

- **iSeries → SAP R/3**

The interface transfers data from the OS/400 application to SAP R/3.

- **SAP R/3 → iSeries**

The interface transfers data from SAP R/3 to the OS/400 application.

The desired direction can be selected along with a project call. After a selection is made, data types that are available for the chosen direction are shown.

For more information about MIDAS400, refer to the MIDAS400 documentation or contact IT-Services and Solutions at: <http://www.itsas.de>

8.4.3 R/3-KOMPAKT

R/3-KOMPAKT Switch from Plaut enables users to transfer transaction data and master data from any iSeries application into the R/3 system. Data relationships between the OS/400 application files are defined in control tables. The completed migration file is available for transfer as soon as the export program is executed.

With the help of suitable transfer options, the program system can be used as a permanent interface between subsystems and R/3. For more information on R/3-KOMPAKT, refer to the Plaut home page at: <http://www.plaut.com>

Chapter 9. R/3 system printing

This chapter describes the fundamentals of SAP R/3 system printing and provides several definition examples.

9.1 Introduction to R/3 system printing

To ensure a consistent printing interface across diverse platforms, SAP R/3 provides its own internal spool support. To-be-printed data that is produced by an application running under R/3 goes into a R/3 spooled file. When the data is in the internal R/3 spooled file, it is not in a suitable form for printing yet. It is ready for printing only after R/3 transforms it into its final form and takes one of two actions:

- Sends the to-be-printed data to the iSeries server where it is stored in a spooled file and managed and printed using the standard iSeries printing facilities.
- Sends the to-be-printed data to a remote print server where it is printed. The remote print server may or may not be another iSeries server. In this case, the data is sent to a remote print server, using the standard Berkeley LPD protocol or the special SAP LPD protocol. The special protocol is used for those systems that do not support the standard Berkeley protocol.

Figure 94 illustrates the flow of data from its origin within an R/3 application to its arrival on an iSeries server output queue.




Figure 94. Flow of R/3 output data on an iSeries server

In R/3, there are several ways to produce output that is to be printed. The most common ways are:

- Clicking a print button on a SAP GUI display, which causes a simple list to be printed
- Requesting the SAP programming editor to print an ABAP source file
- Running an ABAP report
- Producing a document from SAPscript

Therefore, printed output from R/3 can consist of a simple list or complex documents that require advanced printer functions. In any case, R/3 does not actually print the data. R/3 always hands the final output data stream over to the operating system that is responsible for its final printing.

The R/3 spool system is separate from the iSeries server spooled support. As a result, the R/3 spool system requires a definition of print resources and capabilities. Making a printer available to an R/3 application requires a two-step process:

1. Configure the printer on the iSeries server using the normal iSeries configuration support. This makes the printer known to the iSeries server, but not to the R/3 system yet.

You can find some iSeries server configuration examples in 9.11, “LAN-attached printers” on page 202.

2. Add the printer to R/3 using the support provided by the R/3 system. This adds a new printer to the set that is known by the R/3 system and makes it available for use by R/3 applications.

Through the SAP GUI at the workstation, interfaces are provided to define and add new printers to the R/3 system (using transaction SPAD).

All printer methods supported by the iSeries server can be made available to an R/3 application, including SCS, AFPDS, and ASCII using local, LAN, and remote attachment. For a complete description of the different printer models and attachment methods supported by the iSeries server, refer to the redbook *IBM AS/400 Printing V*, SG24-2160. A detailed description of how printers are defined to R/3 is described in 9.5, “Example of configuring a new device to the R/3 system” on page 166.

The R/3 spool system does not recognize or use DDS, externally described data, and printer files. To define the proper format of print data, the R/3 spool system provides its own mechanisms for you to use. These mechanisms are described in more detail in the following sections.

Output data, while it is in an internal R/3 spooled file, is encoded in a character set that is not necessarily the same as the one supported by the printer. R/3 allows for the specification of a printer's character set and automatically converts the data as the final printer data stream is prepared. Character sets are discussed in more detail in 9.4.6, “SAP characters and character sets” on page 165.

R/3 spooled files are managed using the R/3 facilities, not of the iSeries facilities. Through the R/3 SAP GUI at the workstation, interfaces are provided so you can:

- Display the list of outstanding spooled requests.
- Display the content of a spooled file.

- Delete spooled requests.
- Print spooled requests.

Printing a spooled file causes the R/3 spool system to transform the data and give it to the iSeries server for actual printing. Once the data is given to the iSeries server, it enters an output queue. At that point, the printer actually prints the data.

Note

The system printer file QSYSPRT is used by SAP R/3 if the operating system spooler is involved. R/3 overwrites the printer file to match the printer customization within R/3.

Managing the R/3 spooled files is discussed in 9.8, “Managing R/3 spooled and output requests” on page 181.

For more information on R/3 printing, refer to:

- SAP R/3 online help for more information:
 - *BC-CCM SAP Printing Guide*: This publication contains information about all aspects of printing from SAP R/3.
 - *BC-ABA ABAP Programming*: Refer to this publication for information on how to print through ABAP programs.
 - *BC-SRV-SCR SAPscript*: Refer to this publication for information about producing documents through SAPscript.
- SAP R/3 Basis Administration Training 4.6 - BC370
- Web site: <http://sapnet.sap.com/output>

9.2 TemSe data

The set of spooled data held internal to R/3 is referred to collectively by the term *TemSe data*. This term reflects the nature of this data, which is that the data is temporary (Tem) and only accessed sequentially (Se).

TemSe data can be stored in one of two ways in the iSeries server: in database or in the IFS. The *rspo/store_location* profile parameter controls the method by which the TemSe data is stored. If this parameter has a value of *db*, the TemSe data is stored in the database. If it has a value *G*, the TemSe data is stored in the integrated file system. The default value for this parameter is *db*. For more information, refer to SAP Note 20176.

Depending on which system the spool data is stored, you may want to run the spool work process on the same machine.

9.3 The spool work process

A *spool work process* is a special work process that converts spooled data from its internal form into the final output data stream. It sends the final results to a iSeries server spooled file on the appropriate output queue or an external print server. Whether there is a spool work process is controlled by a profile value in

the profile file for the instance. The profile parameter that controls this is *rdisp/wp_no_spo*. This parameter specifies if there are any spool work processes (there can be multiple processes) defined for this instance (use transaction SM51 to see which instances have spool work processes). If a spool work process is not running, data cannot be printed, but R/3 applications can still place data into the TemSe data.

Note

Beginning with SAP R/3 Release 4.0A, it is possible to have more than one spool work process within an R/3 instance. For more information, refer to SAP Note 107899.

In a two-tier configuration, the spool work process always runs on the same machine where the TemSe data is stored. Therefore, the spool work process has local access to the data it needs to process.

In a three-tier configuration, there are two basic options for placing the spool work process:

- On the database server
- On the application server

9.3.1 Placing the spool work processes on the database server

The advantage in this case is that spool work processes have local access to the TemSe data. Because it is local, the access to the TemSe data is very fast due to the reduced traffic on the network.

The main disadvantage is that it places an additional load on the database server's CPU. The additional load comes from the spool work processes and the iSeries server spool support when the data is finally printed. Moreover, if the printed material is needed at the application server, the iSeries server spool support may need to use the network anyway. However, you can control loads by holding output queues and releasing them only during periods of low line traffic.

Figure 95 illustrates the placement of the spool work processes on the database server.




Figure 95. Spool work process on the database server

9.3.2 Placing the spool work processes on an application server

Using this approach, the processing load of the TemSe data is removed from the database server and placed on an application server. Therefore, the database server benefits because the additional resources can be used for database workload. If there are multiple application servers running different instances of R/3, a spool work process can run on each server. The main advantage of this approach is the reduction of the CPU load placed on the database server.

The main disadvantage of this option is that it increases the network traffic between an application server and the system containing the spooled data. If the TemSe is stored in the database, then the system containing the spooled data is the database server. If the TemSe is in the IFS, it is the system that contains `'/usr/sap/<SID>/global'`. If the TemSe data must flow two ways, the R/3 applications initially send the data to the TemSe database. Later, the spool work process must bring the data back for processing.

If the TemSe data is kept in the database, the data flow between an application server and database server is either over fiber optic cables or Gigabit Ethernet. Because of the high transfer rates on those types of connection, there should be no performance problem due to network traffic. If the TemSe data is kept in the IFS file, the data flow is across a TCP/IP connection where performance might be a consideration (unless a high speed connection is used).

Figure 96 illustrates the placement of the spool work processes on the application server. In this example, there are multiple application servers, each running the same instance.




Figure 96. Spool work process on the application server

9.4 Spool administration

All definitions in the R/3 spool system are done at the workstation using the SAP GUI interfaces. This section explains how you navigate to the correct window so you can create a new spool administration task.

To define something to the R/3 spool system, use the SPAD transaction. This brings you immediately to the spool administration window.

9.4.1 Components of a printer definition

Figure 97 illustrates the components that make up a printer's definition.




Figure 97. Components of a printer definition

A printer definition is linked to a device type. A device type basically provides information about the capabilities of any printer of that type. To specify those capabilities, a device type itself is linked to three other components, which include:

- **Print driver:** This function is used to format the final printer output stream. Print drivers are predefined by R/3. You can only choose from those that are available.
- **Character set:** The character set is supported by the printer.
- **Print controls:** This is a mapping of the generic printer controls of R/3 to the specific control characters of the printer.

A device type and its components do not provide any information about how data is to be formatted for the printer. This information is provided by a format (paper type). A format is linked to another component, a *page format*, which specifies the physical dimensions of a paper page and the orientation of the printing on the page.

A format is associated with a device type through a device type format. In addition to linking a format to a device type, this component specifies printer initialization information. Since a device type is capable of processing more than one data format, there may be multiple device type formats associated with a specific device type. As a result, multiple data formats can be associated with a single device type. Conversely, a specific format can be associated with multiple device types by creating unique device type formats for each combination.

9.4.2 Output devices

To define a new printer to the R/3 spool system, complete the following steps:

1. Select the **Output Devices** category from the SPAD transaction.

2. The next window that appears lists all of the current output devices that are already defined to R/3. To define a new one, click the **Create** or the **Create using template** button.

3. On the next window, fill in the following parameters:

- **Output device:** This is the logical name of the output device. The name is case-sensitive and can be up to 30 characters long.
- **Short name:** This is the name R/3 uses to access the printer.
- **Device type:** This parameter is the R/3 method of identifying the model of the device. The R/3 spool system provides a number of predefined device types from which you may choose. A device type defines certain model specific features of any device for that model. For example, it identifies the character set that the device supports. To view the list of device types already defined, click in the input field and the list box that appears. On this list, there are a number of predefined device types for various IBM printers and printers from other manufacturers.

If you cannot find a suitable predefined device type, you may create a new one. We discuss defining device types in 9.4.3, “Device types” on page 159.

- **Spool server:** This identifies the location of the spool work process that processes any data that is directed to this printer. To determine which spool work processes are available, click in the input field and then in the list box that appears. A list of names appears that identifies where the various spool work processes are located. Each name is made up of three parts that are connected by the underscore character. The first part is the host name, the second is the R/3 system ID, and the third is the instance number (this is the same name that is displayed on transaction SM51).
- **Device class:** Choose **Standard printer** from the list box to indicate that the device is a printer.
- **Model, Location and Message:** These are optional fields where you can enter descriptive data about the device.
- **Lock printer in SAP system:** Select the check box if the device is to be logically locked within R/3. When locked, R/3 generated output requests to the device are held until the device is logically unlocked. Note that this does not lock the device at the system level. While it is locked to R/3, the device can still be used outside of that environment.
- **Host spool access method:** This parameter is important because it tells the spool work process what to do with the final output data stream. To send the final output to an iSeries server spooled file, choose one of the following possible values for the iSeries server:
 - **C:** This option allows to use output queues. R/3 sends the final output data stream to a spooled file. This is mainly used for SCS data stream (R/3 device type ZIBMSCS) or ASCII data sent to a remote printer via remote writer (data stream in the spooled file *USERASCII). You can use either locally attached printers or the concept of remote output queues.
 - **F:** Front-end printing allows you to print on locally attached front-end PCs. The iSeries server spool system is not involved at all in this case.

- **Z:** The output is placed in a IFS stream file, and then the CVTPRTDTA command (belonging to the PrintSuite, option 4) is run. This is used to print AFPDS. The access method Z is no longer supported, but can still be used. Access method L supersedes Z.
- **L:** This method is more flexible than access method Z. L uses the CVTPRTDTA command as well.
- **U:** The output is sent via TCP/IP directly to the printer (without going through an output queue) using the Berkeley LPD protocol. This method is not recommended because problems with the printer could lock up the spool work process for quite a while. The iSeries server spool system is not involved at all in this case.
- **S:** The S value also causes R/3 to send the final output to another host rather than putting it into a spooled file. In this case, the protocol being used is SAP's own private protocol. For this to work, the receiving host must have SAP's transfer program running. This method sends print data to a host in a network using TCP/IP that does not support the Berkeley LPD protocol. The iSeries server spool system is not involved at all in this case.
- **X:** The X access method causes the final output stream to go to the SAPcomm support. This method is used, for example, to send a document through the SAP mail system. The iSeries server spool system is not involved at all in this case.

- **Host printer:** The name of the printer at operating system level. If you are using an iSeries server output queue to receive the data that is produced by the spool work process, enter the name of the output queue. R/3 does not create output queues. Therefore, you must ensure the output queue exists prior to any attempt by R/3 to use it.

For front-end printing (access method F), specify `__DEFAULT` to access the default printer on the front-end PC.

If you are using SAPLPD, enter the name of the printer on the remote host that is LPT1.

- **Destination host:** If you used the U or S access method, this field appears after pressing Enter. If you press the Save button before you press Enter, a message appears that requests you to enter “switching computer”.

Enter the name of the server to which the printer is attached.

- **Do not query host spooler for output status:** Select this check box if you do not want R/3 to query the output status from the host spooler (valid for the Z and L methods).

9.4.3 Device types

To create a new device type, follow these steps:

1. Enter the SPAD transaction.
2. Click the **Full administration** button.
3. Click the **Device Types** tab.
4. Click the **Device types** button.

5. The next display that appears shows a list of existing device types and allows several actions to occur. Click the **Change** button.
6. You can use an existing device type as a starting point by selecting one and clicking the **Create using template** button. This uses the selected device as source and lets you enter the name of a new device type as destination.

You can modify an existing R/3 standard device type instead of creating one, but you may lose these changes at the next release upgrade if you do this. As a result, we strongly recommend that you follow the example of this book and create a new device type using the template of an existing device type. Since SAP R/3 Release 4.6B, a reference to the standard device type is kept in the new device type. The benefit is that changes made to the standard device types (most likely by SAP) will immediately affect all device types referencing to that one.

Table 18 shows the currently supported IBM printers and device types. For more information, refer to SAP Note 8928.

Table 18. IBM printers supported by SAP

Printer	Description
IBM239X	Device type for the IBM 238X/239X Plus line printer series from Lexmark. This includes the types IBM 2380 Plus, IBM 2381 Plus, IBM 2390 Plus, IBM 2391 Plus. The IBM emulation and character set IBM 850 are used.
IBM4226	Device type for the IBM 4226 line printers from Lexmark. The IBM emulation and the character set IBM 850 are used.
IBM4232	Device type for the IBM 4232-302 line printer from IBM. The 4202 emulation and the character set IBM 2 are used. OCR-A and OCR-B are included and supported. Barcode printing is not supported.
IBM6408	Device type for the IBM 6408-A00 line printer from IBM. The PropPrinter III XL emulation and the character set IBM 850/IBM 2 are used. OCR-A and OCR-B are included and supported. Barcode printing is not supported.
IBMAFP	Device type for the IBM external CVTPRTDTA converter. The R/3 spool output is converted to AFPDS format and passed to IBM AFP software. IBMAFP can only be used in conjunction with spool-exit (access method Z or L when defining the device type). Selection of printers directly connected to R/3 is not possible. IBMAFP must be used for 240 pel AFP printers. Character set is ISO 8859-1 (Latin 1). OCR-A and OCR-B are supported, as well as barcodes. IBMAFP is for use on R/3 UNIX and Windows NT platforms.
IBMAFP3	IBMAFP3 must be used for 300 pel AFP printers. For more information, refer to the IBMAFP printer.
IBMEFP	IBMEFP is the proper device type for iSeries server R/3 Platforms and uses the EBCDIC character set. For more information, refer to the IBMAFP printer.
IBMEFP3	IBMEFP3 must be used for 300 pel AFP printers. For more information, refer to the IBMEFP printer.
IBMSCS	Device Type IBMSCS supports the "basic IBMSNA Character String" data stream for printers that are connected under IBM OS/400. IBMSCS is only supported for use under SAP R/3 on the IBM iSeries server hardware platforms.

Printer	Description
IBMNP	Device type for laser printer IBM Infoprint 20 as well as the IBM Network Printer 12, 17, 24, IBM Infoprint 32, Infoprint 40. OCR- and MICR fonts are supported by the device type. The printer needs an additional module for these fonts. "Barcode-, MICR and OCR A+B SIMM for IBM Network Printers" is required. Read SAP Note 133660 also. Barcode printing from R/3 is not supported. Important note: IBMNP uses new 4.0A features (scalable font under PCL-5) and can only be used in maintenance level 4.0A and higher.
Laser printer IBM 3912/IBM 3916/IBM 3130	These IBM laser printers can be operated in the PCL-5 mode with device types HPLJ4/HPLJIIID. If the HP font card (as it is called for HPLJ4/HPJIIID) is used, OCR-A and OCR-B can also be used. Barcode printing from R/3 is not supported.
Line printer IBM 4230	According to IBM, this line printer (IBM 4230-4I3, IBM 4230-4S3, IBM 4230-5I3, IBM 4230-5S3) is compatible with IBM 4232 and can, therefore, be operated with definition IBM4232.
Line printer IBM 4247	According to IBM, this line printer (IBM 4247-A00, IBM 4247-001, IBM 4247-002) is compatible with IBM 4232 and can, therefore, be operated with definition IBM4232.
Line printer IBM 6400	According to IBM, this line printer (IBM 6400-005, IBM 6400-009, IBM 6400-014) is compatible with IBM 6408 and can, therefore, be operated with definition IBM6408.
Line printer IBM 6412	According to information from IBM, this line printer is identical to the BM6408 printer and can, therefore, be operated with the device type definition of IBM6408.
Laser printers IBM Network Printer 12, 17, 24 and 24PS	These IBM laser printers can be used in PCL-5 mode with device type HPLJ4/HPLJ5 or IBMNP. OCR-A and OCR-B fonts are supported when using a special IBM Flash-SIMM (see Note 133660). Barcode printing from R/3 is not supported.
Laser printers IBM Infoprint 20, 32, 40	You can operate this laser printer in the PCL-5 mode with device types HPLJ4/HPLJ5 or IBMNP. OCR-A and OCR-B are supported if a special SIMM module is used. Read Note 133660. Barcode printing is not supported.
SAPWIN	Generic device type for printers linked (or also fax devices) to PCs running under MS Windows 3.1, Windows 95, or Windows NT by means of the R/3 program SAPLPD. Windows printer drivers are used, and the character set is ISO 8859-1. Barcode printing from R/3 is possible with the additional installation of a third-party DLL (see Note 25344), but is not supported in the standard system. OCR-A and OCR-B fonts are possible with an appropriate TrueType font (see also Note 48803). As of Release 3.0E, you can use SAPWIN to: • Print proportional fonts and lines or boxes in SAPscript • Print black and white and color TIFF graphics (with the 32-bit SAPlpd on Windows 95 and Windows NT only) See Note 39031.
SAPGOF	SAP Generic Output Format. This data format is used to supply external conversion programs with R/3 printouts. The SAPGOF data format can be used to describe both ABAP list format and SAPscript documents. ASCII version.

Printer	Description
SAPGOF_E	EBCDIC version of SAPGOF. See 9.10.5, "Access method and device type for AFP printers" on page 188, for more information.

When creating a new device type, fill in the following input fields:

- **Device type:** The name by which the new device type is identified. We recommend that this name start with the character Y or Z to avoid a collision with the names of the predefined device types.
- **Name:** A field that allows you to enter descriptive information about the device type. For example, you can enter the name and model of the related printer here.
- **Version:** A field that allows you to associate a version number with the device type definition.
- **Base device type:** You can select the base device type from the list box.
- **SAPscript driver:** The driver to use in this device type to convert the output text format (OTF) data stream generated by SAPscript into the final data stream for the printer. For example, drivers are provided to convert to PostScript, Prescribe, and line printer formats. A special driver is provided to convert from OTF to AFP. To determine which print drivers are available, click in the list box. Select an entry from this list.
- **List printer driver:** Select the printer driver for converting the ABAP list format into the final data stream for the printer.
- **Printer character set:** Enter three SAP ID numbers for three character sets that are compatible with the printer model for which you are defining the device type. Note that these are SAP character set IDs and not a manufacturer's code page number.

The first ID is used to convert all output data sent to a printer of this type to the correct representation for the printer. The second and third character sets are not currently used, but must be filled in. You can use the same ID in all three fields.

A number of character sets are predefined by SAP. To see the list of predefined character sets, click in the field and the list symbol that appears. Each entry in the list shows the SAP ID for the character set, an indicator of what manufacturer the set is intended to support, and a brief description of the character set. In most cases, this is enough information for you to choose the correct set. If you cannot find the correct set, examine each character set and each character in the set to ensure that its hexadecimal representation in the set is the same as on the printer. If a character is in the set but not on the printer, you may produce unprintable characters. If the character does not have the same hexadecimal value, you may print a different character than intended.

If no set is found that meets your needs, create your own character set. This is discussed later in 9.4.6, "SAP characters and character sets" on page 165.

9.4.4 Format types

A format (paper type) groups formatting information for a printed page. The formatting information is not given directly in the format, but by a page format that is linked to the format.

Note

In SAP R/3 Release 3.1H, the term “paper type” is renamed to “format”.

To get a list of the existing formats, follow these steps:

1. Enter the SPAD transaction.
2. Click the **Full administration** button.
3. Click the **Device Types** tab.
4. Click the **Format types** button.

To create a format, follow these steps:

1. Click the **Change** button
2. Click the **Create** button.
3. To use an existing format as a model, select one from the list shown and click the **Create using template** button. This fills in the values from the existing format so you need to type in only the values that need to change.
4. To create a new format, fill in the following fields:
 - **Format type:** Provide the name by which the new format is known. Start your name with the letter Y or Z to distinguish the ones you create from those provided by R/3.
 - **Type:** Select the format type that is most likely the format type for ABAP lists or SAPscript.
 - **Page format:** Provide the name of the page format that is to be associated with this format. If you use the value ANY, any page format is valid for this format.
 - **Orientation:** Use this field to specify the orientation of the printed page. Since this is for documentation only, you may select both check boxes for portrait and landscape.
 - **Comment:** Enter any descriptive comment that you want to document the format.

9.4.4.1 Page formats

A page format specifies the width and length of a page and whether the orientation is portrait or landscape. To create a new page format, follow these steps:

1. Click the **Page formats** button and then the **Change** button. This shows a list of existing page formats.
2. To create a new one, click the **Create** button.
3. To create a new one whose characteristics match that of the paper loaded in the printer, select a page format from the list and click the **Create using template** button.
4. Enter a value into the following fields:
 - **Page format:** Enter the name by which the new format is to be known.
 - **Orientation:** Click the appropriate orientation for the printed data (either portrait or landscape).

- **Paper size:** Enter the physical dimensions of the page width and length. You must indicate what the units for the dimensions are. To see the possible values for units, click in the field and the list symbol that appears.

9.4.4.2 Device type formats

A device type format links a format to a device type and provides additional information about how the printer is to be initialized when the format is formatted for the printer. A device type format is not identified by its own name. Rather, it is identified by providing the combination of the device type name and the format name.

A device type format is created after the device type and format are created. To create a new device type format, follow these steps:

1. Click the **Device types** button and then click the **Change** button.
2. Select the device type and click the **List of implemented formats** button.
3. Double-click the format and the window appears where you can edit the control characters for all kinds of actions that can be taken at different points during the printing of a page.
4. Not all actions that are listed need apply for a specific format. You need to choose those that do. Your choices are influenced by what the print driver is for the device type. Some drivers used for SAPscript do not use what is specified here while others do. To understand what needs to be specified, refer to the *SAP Printing Guide*.
5. To edit the control characters for an action, double-click the button of the particular action. This shows a window where you can enter one or more lines of information. What you need to enter is the exact sequence of escape characters that are needed to accomplish the desired effect at the printer. The control character sequences needed can be found in the printer manual for the specific printer being used. The SAP programming editor is processing this window. As a result, there is a language syntax that you use for entering the information. This syntax is described in the *SAP Printing Guide*.
6. After you enter the information for each action, you need to enter the attributes for the device type format. From the window that lists the actions, click the **Attributes** tab.
7. Most fields on the this window are for documentation purposes only. Select the **PostScript format active** check box if the format is for PostScript output. Otherwise, deselect the check box.

9.4.5 Print controls

Actual print controls are not placed in the data until the final data stream is being formatted. Until that time, print controls are represented in the data stream by symbolic names.

To see the list of standard print controls, follow these steps:

1. Click the **Full administration** button
2. Click the **Device Types** tab
3. Click the **Print Controls** button.

You see, for example, that to set printing at 8 CPI, the symbolic name CI008 is used. A symbolic name is really a place holder that must be assigned a value in the final data stream to the printer.

The value to assign a symbolic name is determined by the print control table associated with the printer's device type. To see the print control tables that are available, click **Utilities-> For device type -> Print PrCtrls**. A window appears that allows you to select either all print controls for every device type or you can specify a print control and device type. If you leave the defaults, a list is shown where each print control table is identified by the name of the associated device type. A print control table does not exist independently of its associated device type, nor do two device types share the same print control table.

To see the print controls that are currently assigned to your device type, complete these tasks:

1. Click the **Full administration** button
2. Click the **Device Types** tab
3. Click the **Device type** button.
4. Click the **Change** button, and select the device type from the list.
5. Click **List of implemented print controls** to access the print control table for that specific device type.

Each entry consists of a couple of fields. The most important field is the one that holds the Control character sequence. Please note the following points:

- Not all print control names have a substitution value. If a printer does not support a particular print control function, there is no value to fill in.
- If the radio button in the Hexadecimal column is active, the value in the Control character sequence field is in hexadecimal representation. Each hexadecimal character is a code point in the encoding scheme of the printer (ASCII code points for ASCII printers and EBCDIC code points for EBCDIC printers). The values entered here are the control characters needed to activate a specific print behavior. These are found in the technical manual for the specific printer.

The only time you need to define a new print control table is when you create a new device type. The easiest way to create a new print control table is to copy an existing one and modify it.

Note

Any print control name that is specified in a print control table must also be specified in the Standard Print Control table. Therefore, if you add a new print control name, you need to add it to the standard table and the device type specific table first. The standard table is edited in a manner similar to a device type table.

9.4.6 SAP characters and character sets

R/3 maintains a master list of characters that it supports. Each character in this master list is assigned a sequential number referred to as the *SAP identification number*. This number is used throughout the spooled system to identify the

character. Every character that is to be processed by the spooled system must be in this master list.

To see the current content of the master list, follow these steps:

1. Click the **Full administration** button.
2. Click the **Char. sets** tab.
3. Click the **SAP characters** button.

As you can see from the master list that is displayed, there is no encoding scheme associated with a character. That is, the master list only specifies the SAP identification number and a description of what the associated character is.

To specify the exact encoding for printing a character at a printer, a character set is used. A character set provides mapping from the SAP identification number to the exact byte encoding for the character. Therefore, by associating a character set with a device type, the spooled system knows how a character must be encoded for the final output data stream.

A character set is identified by a four-digit number. To see a list of character sets already defined, click the **Character sets** button. You can look at the mapping for a character set by selecting one from the list followed by clicking the **Edit character set** button.

It is unlikely that you need to define your own characters and character sets. The character sets predefined by R/3 are extensive. Even if you need to define a new device type, you can use one of the predefined character sets. However, R/3 has provisions for you to:

- Add new characters to the master list.
- Create a new character set.

There are a number of rules that you need to follow when doing this. To understand fully how to work with characters and character sets, refer to the Maintaining Character Sets chapter in the *BC - SAP Printing Guide*.

9.5 Example of configuring a new device to the R/3 system

Before you can add a device type or printer device in SAP R/3, you have to configure your iSeries server printer. Some configuration examples are provided in 9.11, "LAN-attached printers" on page 202. More detailed information about the attachment method and configuration is available in the redbook *AS/400 Printing V*, SG24-2160.

Sometimes it is necessary to add a new device type to your R/3 system because the existing printer device drivers do not match the function of your printer.

Refer to SAP Note 17054 for more information about how to copy or create a device type.

To add a new device type, complete the following steps:

1. Enter the SPAD (Spool administration) transaction.

When the Spool Administration: Initial screen window (Figure 98) appears, click the **Full administration** button, click the **Device Types** tab, and click the **Device types** button.




Figure 98. Spool Administration: Initial Screen window

2. Click the **Change** button and try to find an existing device type on the list where the control sequences resemble your new printer. This is your template device.

Note

For example, most of today's impact printers have the ability to emulate an EPSON type or an IBM PROPRINTER III XL type printer. This means, you can use one of those printers as a template if your new printer is an impact printer. The same analogy can be applied to laser printers as well.

Many IBM SCS printers have no escape control sequences. Because of this, these printers must be controlled by iSeries server printer files or directly by their own settings from the control panel.

Select the printer and click the **Create using template** to create your new device type. The new name should start with a Y or Z according to the SAP naming convention for customer objects as shown in Figure 99.




Figure 99. Copy Device Type window

3. Save the new the device type by clicking the **Save** button.

Go back to the spool administration window and click the **Print controls** button. This shows a list of existing standard print controls as shown in Figure 100.




Figure 100. List of Standard Print Controls

- Click the **Create using template** button to create new standard print controls based on an existing standard print control. Then, the Spool Administration: Copy Standard PrCtrls window (Figure 101) appears. Enter the name of your new standard print control, and enter a description in the Comment field. Choose which of the three print types this new print control can be used with in the Print control status box.

Note

If your new functionality falls into the categories of characters per inch or lines per inch, use the standard form Cxxxx and Lxxxx as the name of the new standard print control, where xxx is CPI or LPI. If your new functionality does not fit into one of the standard categories, type a five-character code that makes sense to you.




Figure 101. Creating a standard print control

5. Save your new print control by clicking the **Save** button and go back to the spool administration window. Click the **Device types** button. Select the device type you just created (in this example ZIBM6408), and click the **List of implemented print controls** button. A list of print controls is shown that was copied from your template device and. These print controls are currently assigned to your device type.

To add a new print control, click the **Standard print controls**. This shows all standard print controls that are available. Double-click the print control you just created (in this example CI018). Then the Spool Administration: Edit Print Controls window (Figure 102) appears.




Figure 102. Adding a print control to the device type

6. Enter the needed control character sequence according to the printer manual. Click the **Save** button to save the device type.
7. Initialize the new device type for more page formats. To do this, go back to the spool administration window, and click the **Device types** button. Then, select the device type, and click the **List of implemented formats** button. The Spool Administration: Format for Device Type window (Figure 103) shows a list of the implemented formats for the selected device type.




Figure 103. Format for Device Type

- Double-click the format to modify the printer initialization. Then, the Maintain Format window (Figure 104) appears.




Figure 104. Maintain Format for device type

9. The most frequently modified format action is printer initialization. Double-click **Printer initialization**. Then the Printer init. window (Figure 105) appears. Enter the commands here to be sent to the printer whenever this particular page format is selected for printout. This includes page size definition (lines per page), initial characters per inch, and lines per inch. The definitions are entered as described in 9.4.4.2, “Device type formats” on page 164.




Figure 105. Printer initialization

10. It may be helpful to add some or all of the printer initialization control characters to the *Start of page* format as well, specially for very large output requests. This will increase the time needed for the print output a little bit. The benefit is that this initializes the printer at the very beginning of every page. This helps usually to prevent from displaced print output.
11. If any special formats and page format are needed, follow the instructions in 9.4.4, “Format types” on page 162.

9.6 Special definitions for barcode printing

If your printer supports barcode printing, you need to define the control sequence for two print controls for each barcode type. SBPxx (barcode prefix) is the control sequence that precedes the variable barcode information in your printout. SBS (barcode suffix) is the control sequence following the variable barcode information.

1. Call the SPAD transaction.

2. Click the **Full administration** button. Click the **Device Types** tab and then click the **Print Controls** button.
3. Select the **SBP..** entry from the list. Click the **Create using template** button to create the barcode prefix as shown in Figure 106.




Figure 106. Create standard print control for barcode printing

4. Click the **Save** button to save the data.
5. Select the **SBS..** entry from the list, and use **Create using template** to create the barcode suffix.

If you defined barcode print controls for your new device type, associate these with an existing system barcode or define a new system barcode to associate them with.

To do this, complete the following steps:

1. Call the SE73 (SAPscript Font Maintenance) transaction.
2. If you need to define a system barcode, select **System barcode** as shown in Figure 107.




Figure 107. SAPscript Font Maintenance: Initial Screen

3. Click the **Change** button. On the next window, click the **Create** button and fill in the needed information as shown in Figure 108.




Figure 108. Create/change system barcode

4. To associate your new device type print control for barcodes with a standard barcode or with one that was created by you, return to the SAPscript Font Maintenance: Initial Screen. Select **Printer barcodes** and click the **Change** button. On the next window, double-click your new device type. If the device type you use as a template has defined barcodes, these are listed as shown in Figure 109.




Figure 109. List of existing barcodes associated with a device type

- To add either standard barcodes or your own barcode previously defined, click the **Create** button and select a desired barcode. Enter the associated print controls SBPxx (barcode prefix) and SBSxx (barcode suffix) as shown in Figure 110.




Figure 110. Create/change printer barcode

- Save the entry and go back to the list of printer barcodes.
- Select your device type. Then, click the **Check print controls** button to see the listing shown in Figure 111.




Figure 111. List of associated printer barcodes

8. The print controls for barcode printing are now associated to the device type.

To create the output device, perform the following steps:

1. Call the **SPAD** transaction.
2. Click the **Output devices** button and the Spool Administration: List of Output Devices window (Figure 112) appears.




Figure 112. Selecting the create output device

3. Click the **Change** button and then click the **Create** button.
4. As shown in Figure 113, specify the output device and the short name. Select the device type and the spool server from the list box.

The screenshot shows the 'Spool Administration: Create Output Device' window in SAP R01 (1) (300). The window has a menu bar (Output device, Edit, Goto, Extras, Utilities, System, Help) and a toolbar. The main area is divided into tabs: DeviceAttributes, HostSpoolAccMethod, Output attributes, and Tray info. The 'DeviceAttributes' tab is active, showing the following fields:

- Output device: IBM4226
- Short name: ZPRT
- Device type: IBM4226 : IBM 4226 IBM emulation
- Spool server: AS23_R01_01
- Device class: Standard printer
- Authorization group: (empty)
- Model: (empty)
- Location: My office
- Message: (empty)
- ☐ Lock printer in SAP System

Figure 113. Sample Create Output Device window (Part 1 of 2)

5. Specify the Host spool access method (in this case C), and enter the Host printer, which is the name of the output queue on the iSeries server. The Spool Administration: Create Output Device window in Figure 114 shows an example.




Figure 114. Sample create output device window (Part 2 of 2)

9.7 Spool request versus spooled files

Spool requests provide the mechanism that the spool system uses to manage and process the TemSe data. A spooled request is a master record that provides detailed information about a related spooled file. Figure 115 shows the print flow in an R/3 system.




Figure 115. R/3 spool request and output request

When you print a document, R/3 generates a spool request and places the spool data stored in the *TemSe*. The spool request and the spool data make up an output request. You can generate several output requests out of one spool request. To combine the two-step process into one, you can specify “print immediately”.

Note

The relationship between the output request and operating system spooled file is one to one. That means every output request results in a spooled file on the iSeries server if the operating system spooler is involved.

A spool request has two parts:

- Administrative information (origin, date, author name, logical printer) stored in the R/3 database.
- The data to be printed is stored in a repository called *TemSe data*. The R/3 spool system uses generic representations of printer formatting commands and R/3 internal character set to represent the characters to be printed.

In most cases, a spool request is used to manage a single spooled file. However, if multiple spooled files are generated that have the exact same attributes, what once were multiple spool requests are now merged into one that is used to manage and process the multiple spooled files.

The separation of the spool request record from the actual output data allows for an easier way to manage the spooled data and for changing attributes after the output data is generated. For example, you can redirect the output to a different printer.

9.8 Managing R/3 spooled and output requests

The following chapter discusses how to manage R/3 spooled files and the actions you can take in this process. In effect, you do not manage the spooled files directly, but rather through the spool requests. By working with a spool request, you indirectly manage the TemSe data.

Managing R/3 spool and output requests requires that you navigate to the Output controller window. To do this from the SAP R/3 main menu, call the SP01 transaction.

Once the window appears, you can display a list of spool or output requests that are in the R/3 spool system. To list the spool request, click the **Spool requests** tab and enter the criteria for identifying the spooled requests in which you are interested.

There are different criteria that you can use to identify only those spooled requests in which you are interested. As criteria, you can enter one or more of the following values:

- **Spool request number:** The number assigned by R/3 to the spool request.
- **Created by:** The name of the R/3 user that created the spool request.
- **Date created:** The creation date of the spooled request. This includes all spool requests created on or after the specified date.
- **Client:** The name of the client that was used to create the spool request.
- **Authorization:** R/3 spool authorization.
- **Output device:** Spool requests for a specified output device. The device is identified by the R/3 device name.
- **Title:** The title of the printed output as it appears on the standard SAP cover sheet. The title can be used only for spool requests where the standard cover sheet is included.
- **Recipient:** The name of the person to receive the document as it appears on the standard SAP cover sheet.
- **Department:** The department of the recipient as it appears on the standard SAP cover sheet.
- **System name:** The R/3 system name. It has to be a valid RFC destination. If you leave the field empty, all the spool requests for all the systems in table ALCONSEG (can be maintained use transaction RZ20) will be listed.

Once you enter the criteria, a list appears of those spool requests that meet the criteria. For each entry selected, important attributes about that entry are shown, including spool request number, output status, size of the data to be printed, and title.

To list the output requests, click the **Output requests** tab and enter criteria for identifying the output requests. The selection criteria is basically the same as the above list.

Note

You can display spool or output requests from other R/3 systems beginning with SAP R/3 Release 4.6A. Enter a valid RFC destination in the System name parameter to do so.

9.8.1 Spool request actions

From the list of spool requests, you can take various actions on one or more of those requests. This section summarizes these actions. For details about any one action, refer to the SAP publication *R/3 Administration*.

1. Display the content of the spool request in character or hexadecimal format.
2. Change attributes of a spool request. You can change the following attributes among other things:
 - The title, the name and the user name
 - The output device to print the data
 - The format to use to format the data
 - The recipient and department to appear on the cover page
 - The delete date after which the spool system can automatically delete the spool request and file if it still exists
 - The status of the cover page, controlling whether it is printed
 - The number of copies to print
 - The priority at which the spool request should be processed for printing
 - The status of the spool request after printing which controls whether it is deleted after successfully printing
3. Delete a spool request or a block of spool requests before they are printed. To delete old spool requests automatically, schedule ABAP program RSP01041 (refer to SAP Note 130978).

Note

It is important to keep the size consumed by the spool requests to a minimum.

9.9 iSeries printer commands

The following list contains important commands that are used on the iSeries server for printing. A brief description for each command is given. For more complete description of these commands, refer to the *Printer Device Programming*, SC41-5713, and to the *iSeries Information Center* for the CL Reference.

- **Create Printer Device Description (CRTDEVPRT)**

This creates a device description of the printer. You can indicate how the printer is attached, the type of printer you are attaching, the model number,

whether the device is varied on at IPL, and depending upon the specific models chosen, other optional parameters.

- **Create Output Queue (CRTOUTQ)**

An output queue is where spooled files are placed prior to actually being printed. You can order the spooled file entries on an output queue, you can identify that any spooled files that come to this output queue should be sent to a remote system for actual printing. Depending on what is specified, some other selections are optionally presented.

- **Start Print Writer (STRPRTWTR)**

This command associates an output queue with a specific printer device. The spooled files in that output queue are printed on the specified device.

- **Start Remote Writer (STRRMTWRT)**

This command associates a remote writer with an output queue. The remote writer sends the files in the specified output queue to an output queue on the remote host that was specified as part of the output queue definition.

- **Work with Writers (WRKWTR)**

This command allows you to work with printer writers, all spooling writers, or a specific writer. The default is to work with just printing writers. If you do not see your writer, try using the command:

```
WRKWTR WTR (*ALL)
```

This command allows you to see:

- Which writers are currently active
- Which output queue is currently associated with an active writer
- Which forms types can currently be printed
- Any messages associated with the writer
- A number of other options

- **Work with Output Queue (WRKOUTQ)**

This command allows you to work with output queues. You can see the number of spooled files, any active writer for that queue, and the current status of the output queue.

You must have the necessary authority to actually use the commands.

9.10 Advanced Function Presentation (AFP) printing with R/3

This section explains how to use Advanced Function Presentation (AFP) printers for volume printing AFP data from SAP R/3. For more information, refer to *SAP R/3 AFP Printing*, S544-5412.

9.10.1 Concept

The SAP R/3 AFP PrintSuite feature provides OS/400 users the Convert Print Data (CVTPRTDTA) command to enable them to use AFP high-speed printers. The CVTPRTDTA command provides a direct transform of SAP R/3 print data into the AFP native MO:DCA-P data stream. This MO:DCA-P data stream contains text records that you can print using fonts.

The SAP R/3 AFP PrintSuite feature includes two types of files that you require:

- Some AFP resources, the CVTPRTDTA command and command processor, the message file, and product information in the QPRTTOOL library
- Configuration files that are installed in the /QIBM/ProdData/PrintSuite directory when the PrintSuite feature is installed

To print output from the CVTPRTDTA command, you must install the appropriate fonts on the system from which you will print. The fonts that you need to install are found in the IBM AFP Font Collection.

The AFP driver can be used like all other print drivers under SAP R/3 for ABAP report and SAPscript printing.

The Advanced Function Presentation architecture provides one of the strongest solutions in this area. IBM provides an SAP driver using this architecture and builds an integrated environment. Figure 116 presents an overview of the printing elements with the AFP driver for the SAP R/3 system.




Figure 116. AFP environment

Each of the numbers in Figure 116 corresponds to the numbered explanation presented here:

1. The application program is invoked by the user to print data in the SAP application and produce a spooled request in the SAP spooling. The print data is placed in storage management, and the device information is placed in device management. The user has to produce a print request to generate a spooled file. This invokes the spool work process.

2. The spool work process generates the formatted data according the device type information. When the device is an AFP device, the AFP print driver is invoked and produces the AFPDS spooled file following the SAPscript instructions, for example. This AFPDS spooled file is placed in the iSeries server spooled environment.
3. The spooled file contains the data from the program with the appropriate formatting instructions. External resources, such as fonts or overlays (the formdef invokes the overlay in R3; see 9.10.6, “Using multiple overlays with CVTPRTDTA” on page 191), are not embedded in the spooled file. Only the references to them are embedded in this file.
4. PSF/400 passes the AFP resources referenced in the spooled file to the printer (the merge occurs at print time). PSF recognizes if the most current copy of a resource is already available in the printer and optimizes the data stream. AFP resources are sent only one time unless they changed on the iSeries server since they were sent. They temporarily reside in the printer until the connection with the system is terminated (using the ENDWTR command).
5. PSF/400 sends the print data and the resources to the printer.
6. PSF/400 manages all of the printer tasks such as printer characteristics, resources management, and error recovery.
7. Only IPDS printers communicate with the system to provide information about the printer and the status of the print job.

9.10.2 Requirements

Table 19 shows the products that are required prior to using SAP R/3 AFP printing.

Table 19. Required products

Product name	Product number	Option	OS/400 library	Current release
Print Services Facility (PSF/400) ¹	5769SS1	36, 37, 38	QAFPLIB1, QAFPLIB2, QAFPLIB3	V4R5M0
AFP Compatibility Fonts ²	5769SS1	8	QFNTCPL	V4R5M0
AFP Font Collection ³ - Optional	5648B45	-	in example: QFNT300CPL, QFNT300LA1	-
AFP PrintSuite	5798AF3	*BASE	QAF3	V3R7M1
SAP R/3 AFP Print for AS/400	5798AF3	4	QPRTTOOL	V3R7M1
1. PSF/400 is the AFP system software for iSeries server printers using the IPDS printer data stream. PSF is required to print with AFP support for SAP R/3. Install either option 36, 37, or 38 based on the speed of your printer. The IPDS printer must be configured with the AFP(*YES) option. For more information, refer to the redbook <i>IBM AS/400 Printing V</i>, SG24-2160. 2. The AFP Compatibility Fonts are free of charge but contain only fonts in 240-pel resolution. 3. If the target printer is a 300-pel resolution, the AFP Font Collection is required as this product includes font family in 240 and 300-pel resolution. This product contains libraries for code pages and fonts. The AFP Font Collection is not visible in the DSPSFWRSC output.				

9.10.3 AFP resources

AFP resources are objects that PSF can use at print time. The resources are referenced in the spool, but not included in the spooled file themselves. The following resources are part of the AFP architecture:

- **Overlays:** A collection of predefined data such as lines, text, boxes, barcodes, images, or graphics. All of these elements build an electronic form that can be merged with the application data at print time. Some elements of the overlay such as images (in this case, page segments) and graphics are not in the overlay, but are an external resource of the overlay.
- **Page segments:** Objects that contain images or text information. Page segments can be referenced in an overlay or directly from an application (but not from SAP yet). Page segments and all other AFP resources are compatible across system platforms with AFP support.
- **Fonts:** A set of graphic characters of a given size and style. There are different types of font objects on the iSeries server. Most applications can use fonts with the iSeries server as printer-resident fonts (Font ID), a code page and character set, or as a coded font.
- **Form definitions:** AFP resources that specify how the printer controls the processing of a sheet of paper. For more information about form definitions and SAP, see 9.10.6, "Using multiple overlays with CVTPRTDTA" on page 191.
- **Page definitions:** AFP resources that contain a set of formatting controls to specify how you want data positioned on the page. This includes controls for the number of lines per printed sheet, font selection, print direction, and mapping fields in the data to positions on the paper. A page definition can be specified at the printer file level.

The design of the resource (overlays, logo) implementation, maintenance, or compatibility includes elements that can dramatically affect your project. This section does not describe how AFP resources can be produced. The redbook *IBM AS/400 Printing V*, SG24-2160, gives a good overview and examples on producing AFP resources.

IBM also has services available to create AFP resources. For more information, contact your IBM Printing Systems Representative or call the Application Solutions Group at 1-303-924-6700.

9.10.4 How the AFP driver for R/3 works

This section describes how the AFP driver for SAP R/3 is invoked and what operation or elements are involved. How and which elements are used in the printer driver process depends on whether your output is an ABAP report or an OTF output.

Figure 117 shows the elements of the AFP print process with SAP R/3 using the AFP driver.




Figure 117. AFP driver for SAP R/3

Each of the numbers in Figure 117 corresponds to the numbered explanation presented here:

1. The AFP driver can be invoked in two different ways:
 - The SAP spool work process invokes the AFP driver using the CVTPTRDTA command automatically when the access method defined in the SAP output device is Z or L.
 - The AFP driver can be invoked manually using the CVTPRTDTA command directly.
2. If your output is an ABAP or an OTF report, the AFP driver works differently:
 - The ABAP report is converted into line data. The AFP formdef and pagedef information from the pagedef.tab are referenced in the line data. A line data spooled file is placed in the iSeries server output queue.
 - The OTF report is converted in an AFPDS data stream. The AFP driver takes all needed information from the following table file:
 - barcode.tab with barcode information
 - defcp.tab for the default code page
 - xxxxyyyy.tab to map ASCII code page to EBCDIC code page

- `fonts.tab` for font information
- `pagedef.tab` contains AFP formdef and pagedef information

3. PSF prints the two different types of spooled files.

- The line data is formatted using the information from the AFP formdef and pagedef and includes the other elements such as fonts or overlays.
- PSF prints the AFPDS formatted data and includes the other elements such as fonts and overlays.

9.10.5 Access method and device type for AFP printers

This section contains information about the proper access method and device type (data stream) for AFP printing.

9.10.5.1 Access method L

The access method for AFP printer is either Z or L (refer to 9.4.2, “Output devices” on page 157). The access method Z is available but no longer supported. The access method Z is replaced by method L that is the access method for AFP printers.

Note

The way how to define the output device for AFP printer within R/3 has changed since R/3 release 4.0A. Access method L has superseded access method Z.

In any case, we use the Convert Printer Data (CVTPRTDTA) command to create the spooled file on the iSeries server. Neither the configuration on operating system level nor the format of the output has changed.

9.10.5.2 Device type SAPGOF_E

The relatively new device type SAPGOF (SAP Generic Output Format) provided by SAP generates generic ASCII or EBCDIC output format.

The SAPGOF data format is used to supply external conversion programs (such as CVTPRTDTA) with R/3 printouts. These external software components can also contain printer drivers for printer protocols that are not directly supported in R/3, such as IBM AFPDS. The SAPGOF data format can be used to describe both ABAP print lists and SAPscript (OTF) documents. SAPGOF is available in a BETA version from Release 3.1G and in a formally released version from Release 4.0A. A SAPGOF data output is produced by setting up an output device with transaction SPAD, which has the device type SAPGOF (or SAPGOF_E).

There are two versions of the SAPGOF data format, an ASCII version (device type SAPGOF) and an EBCDIC version (device type SAPGOF_E).

The SAPGOF data format replaces the so-called “Spool Exit” that was available in an earlier release (access method Z). All product suppliers for the spool exit that are known to SAP have changed their products to the SAPGOF data format. The spool exit (access method Z) is no longer supported by SAP after release 4.0A.

9.10.5.3 Creating an output device

Follow these steps to create a proper output device:

1. Call the SPAD transaction.
2. Click the **Change** button, and then click the **Output devices** button. Click the **Create** or **Create using template** button on the next window. The Spool Administration: Create Output Device window (Figure 118) appears.

Output device: ZIBM3916 Short name: ZPRT

Device type: SAPGOF_E : Generic output format EBCDIC

Spool server: AS23

Device class: Standard printer

Authorization group:

Model:

Location: IBM3916 with access method L and device type SAPGOF_E

Message:

☐ Lock printer in SAP System

Make an entry in all required fields

Figure 118. Access method L and SAPGOF_E (Part 1 of 3)

3. Enter the output device and the short name. Select the device type of **SAPGOF_E**.
4. Click the **HostSpoolAccMethod** tab to access method L. The window shown in Figure 119 appears.




Figure 119. Access method L and SAPGOF_E (Part 2 of 3)

5. Enter the output queue on the iSeries server in the Host printer field. Select **Edit- > Command set** to access the Command record ID field. Type something (like A) and double-click. The window shown in Figure 120 appears.




Figure 120. Access method L and SAPGOF_E (Part 3 of 3)

6. Fill in a description and the Command to transfer print data.
7. Save your settings.

9.10.6 Using multiple overlays with CVTPRTDTA

The Convert Print Data (CVTPRTDTA) command allows the use of electronic overlays, media orientation, and bin selection through the use of form definitions (formdefs). A formdef can consist of multiple copy groups (media maps). A media map is what actually formats the output. However, media maps are often referred to as *formdefs*. A formdef is always used for IPDS printing. CVTPRTDTA defines the formdef to use for printing through the use of formats. In the pagedef.tab configuration file, SAP formats are mapped to AFP formdefs.

If a job's pages all have the same media style, you can simply map a format that fits your needs to a formdef in pagedef.tab. In this case, the first media map within the formdef is used to define the media style. If there is a need to switch media styles on a page-by-page basis (for different overlays on different pages), SAPscript data (OTF) can invoke a media map on a page-by-page basis by defining paper resources for pages in a custom form (formerly known as *layout set*). Each paper resource name passed in the OTF page command defines a media map. The AFP produced contains the Invoke Media Map (IMM) commands, which call out a media map in a formdef. A media mapping remains in effect until another media mapping is invoked.

This section describes:

- Solutions with a unique media style
- Solutions with varying media style

Note

OS/400 comes with an assortment of formdefs. One of the existing formdefs may fit your needs without requiring any modification. For more information on existing OS/400 formdefs, refer to *Printer Device Programming*, SC41-5713.

9.10.6.1 Solutions for jobs with the same media style

One solution is to use F1SAP (if it meets your needs) and have the AFP generated call out the media map in F1SAP that you want to use. To do this, go to Form Painter: Request (transaction SE71) and copy or create a new custom form. In the page resource name field under print attributes, define the media map name for your start page. In your SAPscript, be sure to use this custom form. Now whenever this new custom form is used, it calls out the media map specified by the resource name for the start page. This remains in effect for the entire document.

Another possible solution for a job with the same media style for all pages is to build a formdef that has the specific style you desire. For example, you may choose the same formdef as media map found in F1SAP.PPFA, such as S200000L, which is simplex, pull from bin 2, and print landscape). Make sure that the new formdef has only one copy group in it (the one that defines the behavior you desire). Create a new format (such as ZS2L). Edit pagedef.tab to include this new entry. In the formdef field for your new entry, enter the formdef you made. Now, the document using this new format prints with this formdef's media style.

9.10.6.2 Support for automated formdef selection

A solution for a job that requires media styles on a page-by-page basis is a little more complex. First, create a new format (such as ZCHKS). Next go to Form

Painter: Request (transaction SE71) and create a new form. In the page resource name field under print attributes, define the media map name you want to use for that page. This can be done on a page-by-page basis. Now set the page format in your new custom form to your new format (for example, select **Utilities-> Change Page Format** and then enter ZCHKS). In your SAPscript, be sure to use this custom form. Next, go into pagedef.tab and add a line for your new format (ZCHKS). In the formdef field, specify the formdef that contains *all* of the media maps (copy groups) that the custom form defined for use.

9.10.6.3 Defining a new format

CVTPRTDTA uses formats to map to AFP formdefs, and for ABAP, other relevant processing information. This step creates a blank format in the R/3 system.

For ABAP list printing

Follow these steps to create a format for an ABAP list printing:

1. Call the SPAD transaction.
2. Click the **Full administration** button. Click the **Device Types** tab and then click the **Format types** button.
3. Click the **Change** button. Select a format to copy, and click the **Create using template** button. The window shown in Figure 121 appears.

Note

All ABAP system formats follow the naming convention *X_<number-of-rows>_<number-of-columns>*, and all other system formats are SAPscript formats (for example, Letter). Be sure that if you define a new ABAP format, follow the naming convention *Z_<number-of-rows>_<number-of-columns>_<description>*.

The screenshot shows the SAP SPAD transaction window. The title bar is 'R01 (1) (300)'. The menu bar includes 'Format', 'Edit', 'Goto', 'Utilities', 'System', and 'Help'. The toolbar contains various icons for file operations. The main window title is 'Spool Administration: Copy Format from X_65_255'. Below this, there are icons for 'Format type', 'Attributes', and 'Comment'. The 'Format type' section shows 'Format type' as 'Z_65_255_AFP' and 'Type' as 'Format type for ABAP lists'. The 'Attributes' section shows 'Page format' as 'ANY', 'Number of lines' as '65', 'Number of columns' as '255', and 'List driver' checked. The 'Comment' section contains the text 'ABAP list: At least 65 rows with a maximum number of columns'. At the bottom, a status bar indicates 'Make an entry in all required fields'.

Figure 121. Copy format for ABAP list printing

4. Click the **Save** button to save the format type.

For OTF (SAPscript) printing

OTF formats require a page format with the exact same name first created. Follow these steps to create a page format and format for OTF printing:

1. Call the SPAD transaction.
2. Click the **Full administration** button. Click the **Device Types** tab, and click the **Page format** button.
3. Click the **Change** button. Select a page format to copy, and click the **Create using template** button. The window shown in Figure 122 appears.




Figure 122. Copy page format for OTF printing

4. Click the **Save** button and return to the SPAD window.
5. Click the **Formats types** button and click the **Create** button. The window shown in Figure 123 appears.




Figure 123. Create format for OTF printing

6. Click the **Save** button.

9.10.6.4 Connecting the new format to an existing device type

This section describes how to connect the format previously generated to an existing device type:

1. Go back to the spool administration window.
2. Click the **Device types** button. Select the device type that you want to connect your format to, and click the **List of implemented formats** button.
3. Click the **Create** button and select the format as shown in Figure 124.




Figure 124. Connecting the format to the device type

4. Click the **Execute** button. Then the Maintain Format window (Figure 125) appears.




Figure 125. Maintain Format

5. Notice that the new format is *uninitialized*. Click the **Copy format** button. Enter a source device type and format to initialize the new format as shown in Figure 126.




Figure 126. Copy Format

6. Click the **Copy reference** button. Notice that the format is now *initialized*. Click the **Save** button.

9.10.6.5 Setup to support page-by-page media map selection

This redbook does not present the details of creating a custom form (formerly know as *layout set*). Instead, it shows an example of defining a media name for each defined page in an existing form. You should use a custom form when making these changes, or you may lose your changes when you upgrade your R/3 release level.

Consider the example of defining a media name for each defined page in an existing form. Follow these steps to do this:

1. Call the SE71 (Form Painter) transaction.
2. Enter the existing form and activate the **Pages** radio button as shown in Figure 127.




Figure 127. Form Painter: Request

3. Click the **Change** button. The Change Header window (Figure 128) appears.

R01 (1) (300)

Form Edit Goto Attributes Utilities Settings System Help

Form: Change Header: ZLSETAFP

Pages Windows Page windows Paragraph formats Character formats

Administrative data **Basic settings**

Administration information

Form: **ZLSETAFP**

Description: Test Form Painter

Status: New - Not saved

Classification:

Development class:

Client number: 300

Created on: 00:00:00 by: Release:

Changed on: 00:00:00 by: Release:

Language attributes

Language key: EN

Original language: EN

Translate:

☒ to all languages

☐ to individual languages

☐ do not translate

Figure 128. Change Header (Part 1 of 3)

4. Click the **Basic settings** button to access the window in Figure 129.

The screenshot shows the 'Form: Change Header: ZLSETAFP' dialog box in SAP. The 'Basic settings' tab is selected. Under 'Set up page', 'Page format' is set to 'p1NA4', 'Orientation' is 'Portrait format', 'Lines/inch' is 6.00, 'Characters/inch' is 10.00, and 'First page' is empty. Under 'Default values for text formatting', 'Default paragr.' is empty, 'Tab stop' is 1.00 CM, 'Font family' is COURIER, 'Font size' is 12.0 pt, and 'Bold', 'Italic', and 'Underlined' are unchecked. There are also input fields for 'Spacing', 'Thickness', and 'Intensity'.

Figure 129. Change Header (Part 2 of 3)

5. Click the **Pages** button, and enter the setting like in the example shown in Figure 130.

You must define at least one page for every form. And you must designate a “first” page in the form header. Otherwise, text formatting is not possible. In addition, you should inform the system which page is to be used after reaching the end of the first page. If you do not specify a next page, the output of your text ends at the end of the current page.

With Resource name, you can specify that the paper for this page should be taken from a particular paper tray at the printer. In Resource name, enter the print control that has been defined for switching to the paper tray you want to use.

Page	Meaning	NextPage	Mode	No.type
FIRST	First page	SEC	INC	ARABIC
SEC	2 pages		INC	ARABIC

Page 1 of 2

Standard attributes

Page: FIRST Description: First page

Next page: SEC

Page counter

Mode: INC

Numbering type: ARABIC

Output length:

☐ Uppercase

Print attributes

Resource name: S1D0CAFF

Print mode:

Figure 130. Change Header (Part 3 of 3)

6. Click the **Save** button. Remember that media mapping remains in effect until another media map is encountered.

9.10.7 Setup to support the new OTF user fonts

To set up to print with a new font, perform the following steps:

1. Call the SE73 (SAPscript Font Maintenance) transaction to create a new font family. The SAPscript Font Maintenance window (Figure 131) appears.




Figure 131. Font Maintenance

2. Select the **Font families** radio button. Click the **Change** button and then click the **Create** button on the next window. A display like the example in Figure 132 appears.




Figure 132. Creating the font family

3. Click the **Continue** button and save your settings.
4. Go back to the SAPscript Font Maintenance: Initial Screen. Select the **System fonts** radio button, and click the **Change** button.
5. Click the **Create** button on the next window. A window like the example in Figure 133 appears.




Figure 133. Creating the system font

6. Click the **Continue** button and save your settings.
7. Go back to the SAPscript Font Maintenance: Initial Screen. Select the **Printer fonts** radio button, and click the **Change** button.
8. Double-click your device type (for example, ZDOCAFP), and click the **Create** button. The window shown in Figure 134 appears.



Figure 134. Creating the printer font

9. Click the **Continue** button and save your settings.
10. To use the new OTF font, you must define a form (refer to 9.10.6.5, “Setup to support page-by-page media map selection” on page 195) to use your device type.
11. As a final step, you must add an entry in
`'/QIBM/UserData/PrintSuite/fonts.tab'` for the new font.

Make sure the new resources are on the machine where the physical printer resides.

9.10.8 Setup to support a new OTF user barcode

To set up R/3 system to support a new OTF user barcode, perform these tasks:

1. Call the SE73 (SAPscript Font Maintenance) transaction.
2. Select the **System barcodes** radio button, and click the **Change** button. Click the **Create** button on the next window. The display shown in Figure 135 appears.




Figure 135. Creating the system barcode

3. Click the **Continue** button and save your settings.
4. Go back to the SAPscript Font Maintenance: Initial Screen. Select the **Printer barcodes** radio button, and click the **Change** button and. Double-click your device type on the next window.
5. Click the **Create** button on the next window. Then, the window shown in Figure 136 appears.




Figure 136. Creating the printer barcode

6. Click the **Continue** button.
7. To use the new OTF barcode, you have to define a format (refer to 9.10.6.5, "Setup to support page-by-page media map selection" on page 195) to use your device type (for example, ZDOCAFP).
8. Modify the barcode.tab file to map the barcode name to an actual barcode type. In the barcode.tab file, the format is:

```
BarCode=ZDOCBAR Type=017 Mode=002 Flag=128
```

Also make sure that the new barcode resources are in your resource path and are available on the same machine as your physical printer.

9.11 LAN-attached printers

Several printer attachment methods are available on the iSeries server. This chapter provides information on how to configure the iSeries server LAN-attached Intelligent Printer Data Stream (IPDS) or ASCII printers.

For considerations on all attachment methods, LAN-attached IPDS printers, and LAN-attached ASCII printers, as well as additional configuration examples, refer to the redbook *AS400 Printing V*, SG24-2160.

9.11.1 LAN-attached IPDS printers

Until version 3 of OS/400, printing to a system-attached twinax printer and printing to a network-attached printer were two very different propositions. With system-attached printers printing SCS and IPDS applications, you had printer file functionality, interactive print process management, and complete error recovery. With network-attached printers, much of this control and function was missing. Standard TCP/IP printing is done through a simple file transfer called *line printer requestor* (LPR). Most of the iSeries server printer file function and all of the print management control is missing. The spooled file is simply sent to an IP address. An IPDS connection applies an interactive, bi-directional print protocol to this environment. Except for auto-configuration, LAN-attached IPDS printers function the same as system-attached printers. Figure 137 illustrates the conceptual flow from the iSeries server (using PSF/400) to the LAN-attached IPDS printer.




Figure 137. LAN-attached IPDS printer

The following IBM IPDS printers can be LAN-attached to the iSeries server:

- Any IPDS printers with an IBM Advanced Function Common Control Unit (AFCCU) such as IBM 3130, 3160, Infoprint 60, Infoprint 62, 3900, and Infoprint 4000
- IBM Network Printers NP12 (4312), NP17 (4317), NP24 (4324), or IBM Infoprint 20 with the appropriate LAN card (token-ring or Ethernet)
- The following IPDS printers; IBM 3812, 3816, 3912, 3916, 3112, 3116, 4028, 4230, and 6400 using the I-DATA 7913 printer LAN attachment box (token-ring or Ethernet) 3900, and Infoprint 4000.

Before you begin to configure the printer, you should consider these points:

- Your TCP/IP network must already be set up on your iSeries server.
- Check that Print Services Facility/400 (PSF/400) is installed on your system (Display the list of installed licensed programs by using the `GO LICPGM` command and search for **5769SS1**, option 36, 37, or 38).
- To avoid any problem, install the latest cumulative PTFs on your system.

9.11.1.1 Creating the device description

To create the device description for your printer, perform these steps:

1. Type the Create Device description Printer (CRTDEVPRT) command on any command line and press F4 (Prompt).
2. Specify values for the parameters similar to the ones shown in Figure 138.

```

                                Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Device description . . . . . > PRT01      Name
Device class . . . . . > *LAN             *LCL, *RMT, *VRT, *SNPT, *LAN
Device type . . . . . > *IPDS             3287, 3812, 4019, 4201...
Device model . . . . . > 0                0, 1, 2, 3, 4, 10, 13, 301...
LAN attachment . . . . . *IP              *LEXLINK, *IP, *USRDFN
Advanced function printing . . . *YES      *NO, *YES
Port number . . . . . 0                  0-65535
Online at IPL . . . . . *YES              *YES, *NO
Font:
  Identifier . . . . . > 011              3, 5, 11, 12, 13, 18, 19...
  Point size . . . . . *NONE              000.1-999.9, *NONE
Form feed . . . . . > *AUTOCUT            *TYPE, *CONT, *CONT2, *CUT...
Separator drawer . . . . . *FILE          1-255, *FILE
Separator program . . . . . *NONE         Name, *NONE
Library . . . . . *LIBL                  Name, *LIBL, *CURLIB
Printer error message . . . . . *INQ      *INQ, *INFO
More...

F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys
```

Figure 138. Create Device Description (Part 1 of 3)

3. Page down to continue, and the screen shown in Figure 139 appears.

```

                                Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Message queue . . . . . *CTLD             Name, *CTLD, *SYSOPR, QSYSOPR
Library . . . . . *LIBL                  Name, *LIBL, *CURLIB
Activation timer . . . . . > *NOMAX        1-2550, *NOMAX
Image configuration . . . . . *NONE        *NONE, *IMGA01, *IMGA02...
Maximum pending requests . . . . 6         1-31
Print while converting . . . . . *YES      *NO, *YES
Form definition . . . . . F1C10110        Name
Library . . . . . *LIBL                  Name, *LIBL, *CURLIB
Remote location:
  Name or address . . . . . > '9.9.99.99'

User-defined options . . . . . *NONE       Character value, *NONE
+ for more values

More...

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

Figure 139. Create Device Description (Part 2 of 3)

4. If only one iSeries server uses the printer, use the default value for Activation timer (170 seconds). If more than one system shares the printer, set the value to *NOMAX, which causes PSF/400 to wait to establish a connection.
5. Page down to continue. The screen shown in Figure 140 appears.

Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

User-defined object:

Object	> NP17	Name, *NONE
Library	> QGPL	Name, *LIBL, *CURLIB
Object type	> *PSFCFG	*DTAARA, *DTAQ, *FILE...
Data transform program	*NONE	Name, *NONE
Library		Name, *LIBL, *CURLIB
User-defined driver program	*NONE	Name, *NONE
Library		Name, *LIBL, *CURLIB
Text 'description' DEVD for IBM Network Printer 17		

Bottom

F3=Exit
F4=Prompt
F5=Refresh
F10=Additional parameters
F12=Cancel
F13=How to use this display
F24=More keys

Figure 140. Create Device Description (Part 3 of 3)

6. Press Enter to save the settings.

9.11.1.2 Creating the PSF configuration object

To create the PSF configuration support, follow these steps:

1. Enter the Create PSF configuration (CRTPSFCFG) command on any command line and press F4 (Prompt). The screen shown in Figure 141 is shown.

Create PSF Configuration (CRTPSFCFG)

Type choices, press Enter.

PSF configuration	> NP17	Name
Library	> QGPL	Name, *CURLIB
User resource library list . . .	*JOBLIBL	*JOBLIBL, *CURLIB, *NONE
Device resource library list . .	*DFT	Name, *DFT
+ for more values		
IPDS pass through	> *YES	*NO, *YES
Activate release timer	*NORDYF	*NORDYF, *IMMED...
Release timer	> *SEC15	1-1440, *NOMAX, *SEC15...
Restart timer	*IMMED	1-1440, *IMMED
APPC and TCP/IP retry count . .	> 10	1-99, *NOMAX
Delay between APPC retries . .	> 10	0-999
Automatic session recovery . . .	*NO	*NO, *YES
Acknowledgment frequency . . .	100	1-32767
Text 'description'	PSF conf. object for IBM Network Printer 17	

Bottom

F3=Exit	F4=Prompt	F5=Refresh	F10=Additional parameters	F12=Cancel
F13=How to use this display		F24=More keys		

Figure 141. Create PSF Configuration object (Part 1 of 3)

- The name of the PSF configuration must correspond to the name specified in screen in Figure 140. The same PSF configuration object can be used for more than one printer.

IPDS pass through reduces the PSF/400 conversion time for some *SCS or *IPDS spooled files.

If only one system is using the printer, specify *NOMAX for Release timer because there is no need to release the printer for another system.

- To continue, press F10 (additional parameters) and page down. The screen shown in Figure 142 appears.

```

                                Create PSF Configuration (CRTPSF CFG)

Type choices, press Enter.

                                Additional Parameters

Blank page . . . . . > *NO          *YES, *NO
Page size control . . . . . *NO      *NO, *YES
Resident fonts . . . . . *YES       *YES, *NO
Resource retention . . . . . *YES    *YES, *NO
Edge orient . . . . . *NO          *YES, *NO
Use outline fonts . . . . . > *YES    *YES, *NO
PSF defined option . . . . . *NONE
    + for more values
Font substitution messages . . . *YES    *YES, *NO
Capture host fonts at printer . *NO     *NO, *YES
Font resolution for formatting  *SEARCH  *SEARCH, 240, 300
Font mapping table . . . . . *NONE    Name, *NONE
    Library . . . . .                Name, *LIBL, *CURLIB
                                           More...

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

```

Figure 142. Create PSF Configuration object (Part 2 of 3)

- Page down to continue. The screen shown in Figure 143 appears.

```

                                Create PSF Configuration (CRTPSF CFG)

Type choices, press Enter.

Cut sheet emulation mode . . . . *NONE    *NONE, *CHKFIRST, *CHKALL
Replace . . . . . *YES           *YES, *NO
Authority . . . . . *LIBCRTAUT    Name, *LIBCRTAUT, *CHANGE...

                                           Bottom

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

```

Figure 143. Create PSF Configuration object (Part 3 of 3)

- Leave the default parameter values as they are, and press Enter to create the PSF configuration object.

9.11.2 LAN-attached ASCII printers

The iSeries server can route print to LAN ASCII printers using TCP/IP. This includes the IBM Network Printer 12, 17, and 24 and the IBM 3130, as well as

non-IBM ASCII printers with appropriate network attachments. There are two methods to direct print to these printers:

- **Send TCP/IP spooled file:** This is the iSeries server implementation of the standard TCP/IP print file transfer, called *line printer requestor* (LPR). The spooled file is sent to an IP address. The printer must be capable of receiving an LPR transmission (this capability is called *line printer daemon* (LPD)). With the Send TCP/IP spooled file, the print file is simply sent. There is no print management by the iSeries server. The transmission has the following characteristics:
 - The entire file is sent.
 - Some printers ignore the control file that is sent along, losing job control.
 - This is a one-way transmission – no control, status, or error recovery.
 - The entire spooled file is resent on transmission errors.
 - Most printer file parameters on the iSeries server are inoperable.
 - Cancel printing will yield unpredictable results.
- **PJL driver:** The PJL driver (available with OS/400 V3R7 and later releases) increases the network printing function beyond what is provided by the Send TCP/IP spooled file. With the PJL driver, a printer device description is created (unlike the case with the Send TCP/IP spooled file). This facilitates a dialog between the iSeries server and the printer, although this dialog is limited in comparison with IPDS. The PJL driver supports the copies and page range parameters of the printer file. Some status information is returned from the printer.

Before you begin to configure the printer, you should consider these points:

- Your TCP/IP network must already be set up on your iSeries server.
- To avoid any problem, install the latest cumulative PTFs on your system.

9.11.2.1 Creating the device description

To be LAN-attached with the PJL driver, your printer must support Printer Job Language (PJM) and PCL. Complete these steps to configure LAN-attached ASCII printers with PJL drivers:

1. To create the device description for your printer, type the Create Device Description Printer (CRTDEVPRN) command on any command line and press F4 (Prompt). A screen appears like the example in Figure 135.

```

                                Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Device description . . . . . > NPLAN          Name
Device class . . . . . > *LAN                *LCL, *RMT, *VRT, *SNPT, *LAN
Device type . . . . . > 3812                 3287, 3812, 4019, 4201...
Device model . . . . . > 1                   0, 1, 2, 3, 4, 10, 13, 301...
LAN attachment . . . . . > *IP                *LEXLINK, *IP, *USRDEFN
Port number . . . . . > 2501                 0-65535
Online at IPL . . . . . *YES                 *YES, *NO
Font:
  Identifier . . . . . > 11                  3, 5, 11, 12, 13, 18, 19...
  Point size . . . . . *NONE                 000.1-999.9, *NONE
Form feed . . . . . > *AUTOCUT               *TYPE, *CONT, *CONT2, *CUT...
Separator drawer . . . . . *FILE             1-255, *FILE
Separator program . . . . . *NONE            Name, *NONE
Library . . . . . *LIBL                     Name, *LIBL, *CURLIB
Printer error message . . . . . *INQ         *INQ, *INFO

More...

F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys

```

Figure 144. CRTDEVPRT for LAN-attached ASCII printer using the PJI driver (Part 1 of 3)

2. Page down to continue. Then you see the display shown in Figure 145.

```

                                Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Message queue . . . . . > *SYSOPR          Name, *CTLD, *SYSOPR, QSYSOPR
Library . . . . . *LIBL                    Name, *LIBL, *CURLIB
Activation timer . . . . . 170              1-2550, *NOMAX
Inactivity timer . . . . . *ATTACH          1-30, *ATTACH, *NOMAX...
Host print transform . . . . . > *YES        *NO, *YES
Manufacturer type and model . . > *IBM4317
Paper source 1 . . . . . > *A4              *MFRTYPMDL, *LETTER...
Paper source 2 . . . . . > *A4              *MFRTYPMDL, *LETTER...
Envelope source . . . . . > *C5             *MFRTYPMDL, *MONARCH...
ASCII code page 899 support . . *NO         *NO, *YES
Image configuration . . . . . > *IMGA02     *NONE, *IMGA01, *IMGA02...
Character identifier:
  Graphic character set . . . . . *SYSVAL    1-32767, *SYSVAL
  Code page . . . . . *SYSVAL               1-32767

More...

F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys

```

Figure 145. CRTDEVPRT for LAN-attached ASCII printer using the PJI driver (Part 2 of 3)

3. Specify *YES for Host print transform because the spooled files from the iSeries server must be transformed from EBCDIC to ASCII. Page down to continue. Then you see the screen shown in Figure 146.

```

                                Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Remote location:
  Name or address . . . . . > '9.9.99.99'

User-defined options . . . . . *NONE      Character value, *NONE
      + for more values
User-defined object:
  Object . . . . . *NONE      Name, *NONE
  Library . . . . .          Name, *LIBL, *CURLIB
  Object type . . . . .          *DTAARA, *DTAQ, *FILE...
  Data transform program . . . . . *NONE      Name, *NONE
  Library . . . . .          Name, *LIBL, *CURLIB
  System driver program . . . . . > *IBMPJLDRV
  Text 'description' . . . . . DEVD for NPLAN

                                                                    Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys

```

Figure 146. CRTDEVPRT for LAN-attached ASCII printer using the PJI driver (Part 3 of 3)

4. Press Enter to create the device description.

9.11.3 Creating a remote output queue

To create a remote output queue for your printer, complete the following tasks:

1. Type the Create Output Queue (CRTOUTQ) command on any command line, and press the F4 (Prompt) function key. The screen shown in Figure 147 appears.

```

                                Create Output Queue (CRTOUTQ)

Type choices, press Enter.

Output queue . . . . . > RMT      Name
  Library . . . . . > MYLIB      Name, *CURLIB
Maximum spooled file size:
  Number of pages . . . . . *NONE      Number, *NONE
  Starting time . . . . .          Time
  Ending time . . . . .          Time
      + for more values
Order of files on queue . . . . . *FIFO      *FIFO, *JOBNBR
Remote system . . . . . > *INETADR

Remote printer queue . . . . . 'PASS'

                                                                    More...
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys

```

Figure 147. Create Output Queue (Part 1 of 3)

2. Page down to continue. Then you see the screen shown in Figure 148.

Create Output Queue (CRTOUTQ)

Type choices, press Enter.

Writers to autostart	> 1	1-10, *NONE
Queue for writer messages . . .	QSYSOPR	Name
Library	*LIBL	Name, *LIBL, *CURLIB
Connection type	> *IP	*SNA, *IP, *IPX, *USRDFN
Destination type	> *OTHER	*OS400, *OS400V2, *PSF2...
Host print transform	*YES	*YES, *NO
Manufacturer type and model . .	*IBM4317	
Workstation customizing object	*NONE	Name, *NONE
Library		Name, *LIBL, *CURLIB
Image configuration	*NONE	*NONE, *IMGA01, *IMGA02...
Internet address	'9.9.99.99'	
Destination options	*NONE	
Print separator page	*YES	*YES, *NO
F3=Exit F4=Prompt F5=Refresh		More...
F13=How to use this display		F10=Additional parameters F12=Cancel
		F24=More keys

Figure 148. Create Output Queue (Part 2 of 3)

3. Page down to continue. Next you see the screen shown in Figure 149.

Create Output Queue (CRTOUTQ)

Type choices, press Enter.

User defined option	*NONE	Option, *NONE
+ for more values		
User defined object:		
Object	*NONE	Name, *NONE
Library		Name, *LIBL, *CURLIB
Object type		*DTAARA, *DTAQ, *FILE...
User driver program	*NONE	Name, *NONE
Library		Name, *LIBL, *CURLIB
Spooled file ASP	*SYSTEM	*SYSTEM, *OUTQASP
Text 'description'	Remote output queue for IBM4317	
F3=Exit F4=Prompt F5=Refresh		Bottom
F13=How to use this display		F10=Additional parameters F12=Cancel
		F24=More keys

Figure 149. Create Output Queue (Part 3 of 3)

4. Press Enter to save your changes.

9.12 Resolution tips for printing problems

Sometimes your spooled file does not print and it may be difficult to determine why. This section points out some common causes as to why your file did not print if the iSeries server spool system was involved.

9.12.1 General tips

Here are some general tips for specific printing related problems.

9.12.1.1 Spooled file is not generated

In this case, try the following tasks:

- Search for the output request using the SP01 transaction. Make sure the status is "Compl." (completed).
- Look at the corresponding job log on the iSeries server. To do so, you have to find out the process ID of the spool work process using transaction SM50/SM51. Use the `WRKPID` command to display the job log of the spool work process. Make sure you select the proper spool work process in the correct instance.
- It may have happened that the output queue was not in the library list of the spool work process. In this case, the system may have placed it into the default output queue. Search for the spooled file using the command:

```
WRKSPLF SELECT(<SID>OWNER)
```

Here, *nn* is the instance number of the spool work process job.

9.12.1.2 Spool work process or the spool server is not active

If you do not have spool work processes in your instance or you defined a spool server on another system, make sure that the spool work process can be accessed by the application.

9.12.1.3 A spooled file is in the output queue, but it is not printing

In this case, try these tasks:

- Check whether the output queue is released using the `WRKOUTQ` command. If not, release the output queue using the `RLSOUTQ` command.
- Verify that the printer device has been started using the `WRKWTR` command.
- Check whether the printer, attached to the output queue, can handle the spooled file type.
- It is possible that the iSeries server spool system tried to print the spooled file, but found some type of error. Depending where the messages for this printer go, you may find a spooled error message indicating that the iSeries server spooled support encountered a problem. There may be a problem with the iSeries server spooled support, but check the SAPNet - R/3 Frontend (formerly known as OSS) if this is a known problem.

9.12.2 Access methods using CVTPRTDTA

This section lists some additional problem resolution tips for access methods using the CVTPRTDTA command. The access method can be Z or L.

9.12.2.1 CVTPRTDTA command is not found by spool work process

If you see a call to the CVTPRTDTA command in the job log of the spool work process, but the command is not found, try these actions:

- Note the library in which the command is being searched for and copy the command to that library.
- If the command is being looked for in *LIBL but is not found, include the library (usually QPRTTOOL) to the R/3 startup program R3STRUP.

9.12.2.2 CVTPRTDTA command has errors in the job log

If you see the CVTPRTDTA command was called successfully, but there were errors (display detailed messages to see errors), address the errors that were generated by using the CVTPRTDTA command.

9.12.2.3 The spooled file gets held when I try to print it

Check the printer writer job log for messages using the command WRKWTR. Then select option 5 (Work with) and then F17 (Writer job). Note that you have a WTR and a PDJ (for AFPDS) job. Therefore, check both job logs for the following messages:

- AFP resources (fonts, pagedef, formdef, overlays, segments) not found
 1. Delete the spooled file
 2. Find the resources on your system, and note the libraries where you found them. Then, add the libraries to the R/3 startup program R3STRUP.
- Barcode specification checks

Barcodes generated within R/3 must conform to BCOCA specifications (can't fall out of area, be too small, etc.)
- Other PSF/400 errors

Call IBM support.

9.12.2.4 The device type is incorrect

ABAP list format data should be generated into a iSeries server spooled file with a device type of *LINE. SAPscript output should be generated as an *AFPDS spooled file. If this is not the case, verify that your R/3 output device definition is correct.

9.12.2.5 The form definition is not correct

An iSeries server spooled file is generated but the formdef is not what was expected for the SAPscript output:

1. Verify that the device type of your spooled file is *AFPDS and note what the form definition of the spooled file is.
2. If the formdef is *DEV D, the R/3 format used to spool the file does not start with Z.

9.12.2.6 The page definition and form definition is wrong

The form definition and page definition are not what you were expecting for your ABAP list format output.

The format (paper type) you are using to print the ABAP output has an entry in either the '/QIBM/UserData/PrintSuite/pagedef.tab' or

'/QIBM/ProdData/PrintSuite/pagedef.tab' file. The entry is used to pick up the name of the page definition and form definition used to spool the output.

Edit the '/QIBM/UserData/PrintSuite/pagedef.tab' file and change the entry for the format your spooled output is using. It is necessary that user SIDnn is allowed to read this file. Remember, all AFP output you print with this format will be changed, so you may want to create another format, add this entry to the pagedef.tab file, and print with that format instead.

9.12.2.7 The format is wrong

The paper format you've created that starts with Z in R/3 is not available for spooling the data. Refer to 9.10.6.4, "Connecting the new format to an existing device type" on page 194.

Chapter 10. iSeries system administration using GUI

This chapter covers the GUI-based system administration and system management tools for iSeries server.

10.1 Operations Navigator

The graphical user interface to OS/400 functions is provided through Operations Navigator. Operations Navigator is software that runs on a Windows workstation and provides system administrators and operators with a Windows Explorer-like view of the iSeries resources. It is a standard feature of the base operating system, installed together with Client Access Express. SAP users may find that administering the iSeries server through this graphical user interface is easier and more intuitive than using the character-based interface, especially if they are not already familiar with iSeries operations. The functions of the Operations Navigator include:

- **Management Central:** The system management tool that manages multiple iSeries servers, as explained in 10.2, “Management Central” on page 217.
- **Basic operations:** Work with messages, printer output, and printer management.
- **Job management:** Displays and manages jobs on the system, including server jobs.
- **System configuration and services:** An inventory of hardware and software fixes, as well as a definition and collection of system performance data.
- **Network service configuration:** Configuration of interfaces, protocols, network security features, and so on.
- **Security and user profile management:** Policies, users, groups, authorizations.
- **Database and file system management:** The creation and management of SQL collections, tables, libraries, files, data sources, SQL scripts, and SQL performance monitors.
- **Multimedia:** The ability to store multimedia data, such as audio and video.

Figure 151 shows a typical Operations Navigator panel. It shows connections to two iSeries servers (named System1 and System2), the function groups for one of the servers (System2), and the functions available in some of the function groups (Basic Operations, Job management, Configuration and Service, Network, Security, Users and Groups, Database). It is possible to perform TCP/IP network configuration, for example, using this panel.




Figure 150. Operations Navigator panel

10.1.1 Database management

There are some new database management features in the V4R5 version of Operations Navigator that may be useful to R/3 users. One of these is Visual Explain.

Visual Explain provides a graphical way of identifying and analyzing database performance. It provides a graphical representation of the optimizer implementation of a query and the query result. It can identify the tables used by the query and the indexes considered for the query. It shows the lowest cost index, and therefore, the index that will be used. Where there is no index, it provides guidance as to whether an index could be beneficial and if so, the criteria to use for index creation. Visual Explain can help you understand where the greatest cost is being incurred.

Visual Explain shows the job run environment details and the levels of database parallelism that were used to process the query. It also shows the access plan in a diagrammatic form. This allows you to zoom to any part of the diagram for further details.

If query performance is an issue, the information provided by Visual Explain can help you to determine whether to:

- Rewrite or alter the SQL statement.
- Change the query attributes or environment settings.
- Create the recommended indexes.

Best of all, you do not even have to run the query to find this information. Visual Explain has a modeling option that allows you to explain the query without running it. That means you could try any of the changes suggested and see how they are likely to work, before you decide whether to implement them.

Visual Explain is an advanced tool that assists you with the task of enhancing query performance. It does not enhance the performance itself. You still need to understand the process of query optimization and the different access plans that can be implemented.

For more information on Visual Explain, refer to following documents, which are available at: <http://www.redbooks.ibm.com>

- *DB2 UDB for AS/400 Visual Explain: Illustrating the Secrets*, REDP0505
- *Using AS/400 Database Monitor and Visual Explain to Identify and Tune SQL Queries*, REDP0502
- *DB2 UDB for AS/400: Database Administration with Operations Navigator V4R5*, SG24-5993 (Redpiece)

For more information on Operations Navigator, refer to *AS/400 Client Access Express for Windows: Implementing V4R4M0*, SG24-5191.

10.2 Management Central

Management Central is a GUI-based, easy-to-use tool for managing multiple iSeries systems. It contains a suite of systems management functions that were introduced in V4R3 of the OS/400 as part of Operations Navigator. The functions supplied by Management Central are very useful in SAP R/3 environments with two or more R/3 systems (development, test, and production) running on separate iSeries servers or LPARs.

Management Central provides iSeries system administrators and operators with the ability to manage multiple iSeries servers that are interconnected across a TCP/IP network. Management Central expands the Operations Navigator's system management capabilities with:

- Additional features, such as software fix management and performance collection services
- The ability to manage several iSeries servers in a network

There are two types of systems in the Management Central environment:

- Central system (there is always only one central system in the network)
- One or more endpoint systems

The central system can be any iSeries V4R4 (or higher) server that you use to manage the other iSeries servers in your network. The other systems that are managed by the central system are called *endpoint systems*. Figure 151 shows the Management Central topology.




Figure 151. Management Central topology

The system administrator works with Management Central through Operations Navigator, which runs on a PC client attached to the central iSeries server. The central system broadcasts requests, collects data, receives response information, and provides the central data store for persistent management information. Depending on the business needs, different systems at various times can assume different roles. Although Management Central appears to be hierarchical in nature, it is actually implemented using a peer-to-peer model. For example, an endpoint system can directly send the data to another endpoint system without moving the information to the central system first. To use the full functionality of Management Central, all of the systems must be run using OS/400 V4R4 or higher.

Figure 152 shows an example of the administrator's display with Management Central.




Figure 152. Operations Navigator Management Central

Management Central provides real-time performance monitoring and has a set of integrated graphical tools to manage iSeries servers. These tools include:

- **Inventory collection for AS/400 hardware, software and software fixes:** In native AS/400 terminology, the fixes are called PTFs.
- **Software fix management:** To install, uninstall, compare, and send fixes among selected systems.
- **Integrated simple schedule:** To set when a certain task or action should occur.
- **Running commands:** To define a CL command into a task and then send it to multiple systems at once.
- **Package and object distribution:** To group QSYS objects or IFS files into a package and then send it to multiple systems.
- **Performance collection services:** To collect performance data for multiple systems.

Commands and package definitions can be reused later, which is not the case with some other similar tools such as FTP.

10.2.1 Performance monitoring and collection

Management Central can collect, monitor, and display real-time performance data for multiple iSeries servers.

10.2.1.1 Performance monitoring

Detailed graphs help you to visualize what is going on with your systems. You choose which iSeries servers and performance measurements to display. Then, use Management Central functions to create and start the appropriate monitors.

You can have multiple monitors measuring different system resources from different systems active at the same time.

In addition, you can establish thresholds for selected system resource collected by each monitor and automate the triggering of warning messages or other actions when the measurements exceed these thresholds. Management Central will continue to monitor your systems and perform any threshold commands or actions that you specified, even if your PC is inoperative.

10.2.1.2 Performance collection

Management Central's Collection Services can collect performance data for future analysis by the iSeries licensed program Performance Tools (5769-PT1) or other performance report applications. You can use Collection Services instead of the OS/400 performance monitor function, which is run by the Start Performance Monitor (STRPFRMON) command. As opposed to the OS/400 performance monitor, Collection Services gathers the performance data for each collection in a single collection object. This means a lower system overhead when collecting performance data.

10.2.2 Examples of Management Central usage with SAP R/3

Some examples of using Management Central in the R/3 environment are:

- With inventory collection and software fix management, system administrators can load and test a new software fix on a test system first, and then transfer and install it to other systems at once.
- Running iSeries commands with integrated scheduler can be used to:
 - Start and stop R/3 systems on multiple iSeries servers
 - Make backups and transports
 - Create an iSeries user profile on multiple systems or system groups
 - Set iSeries system values or network attributes on multiple systems or system groups
 - Change an iSeries user profile password on multiple systems or system groups
 - Set up your own help desk or operations run book to handle customer and system needs
- Object distribution can be used to distribute and apply SAP patches to all R/3 systems, including OS/400 commands and programs that are started after the distribution is complete. Objects can be grouped into packages with configuration data, Java applications, HTML Web pages, or software programs, for example.
- To observe real-time generic performance data for several systems at a time, consider using Management Central. However, to have more detailed application level performance information, use the SAP GUI interface with SAP transactions.
- You can send a command to remote systems connected over the Internet that creates a user profile on the remote iSeries server (with the initial password in it!). Since Management Central supports Secure Sockets Layer (SSL) communications, the password cannot be tapped.

Chapter 11. Availability, backup, and recovery

The integrity and consistency of the SAP R/3 database is protected by the R/3 transaction concept, which handles both dialog (“interactive”) and dialog-free (“background”) transactions as logical units of work. However, there are many potential causes for a breakdown in the information processing system. These causes include human error, hardware or software malfunction, power failures, natural disasters, fire, and so on.

An effective backup and recovery strategy is necessary to safeguard an organization from loss of information caused by a system failure and to minimize the time taken to recover from the failure. The impact of the failure depends on the duration of interruption as well as the importance of the applications impacted by the failure.

This chapter outlines the standard iSeries server backup, recovery, and availability facilities available. It also provides an example of backup and recovery facilities and procedures in the SAP R/3 environment. You can find full, detailed coverage of the subjects in:

- *Backup and Recovery*, SC41-5304
- R/3 online documentation
- *BC R/3 Database Guide: DB2/400*
- iSeries Information Center at
<http://publib.boulder.ibm.com/pubs/html/as400/infocenter.htm>

We strongly recommend that you read these documents to plan backup, recovery, and system availability to protect and maintain crucial business functions.

11.1 Availability considerations

This section describes the techniques you can use to provide the appropriate availability for your R/3 application running on an iSeries server.

11.1.1 Terminology

Before you determine the best solution for availability, it is important to understand some availability terminology:

- **Outage:** A period when the system is not available to users. During a *scheduled outage*, you deliberately make your system unavailable to users. You might use a scheduled outage to run batch work, save your system, or apply PTFs. An *unscheduled outage* is usually caused by a failure of some type.
- **High availability:** The system has no unscheduled outages.
- **Continuous operations:** The system has no scheduled outages.
- **Continuous availability:** The system has no scheduled or unscheduled outages.
- **Disaster recovery:** Plans are in place and ready to go to recover off site from a disaster.

Furthermore, you have to understand which activities or events (planned or unplanned) can influence the availability of your system.

11.1.2 Scheduled outages

Scheduled outages are required to maintain your system. During this time, SAP R/3 is not available on this system. To completely avoid scheduled outages and provide continuous operations, you must implement a backup system.

Activities that cause scheduled outages are:

- Hardware maintenance
 - Processor
 - DASD ¹
 - Network ¹
- Operating system maintenance
 - OS/400 release upgrade
 - Applying PTFs ¹
 - Save system
- R/3 system maintenance
 - Installing a new R/3 database release or R/3 kernel release
 - Installing patches for the R/3 system ¹
 - Saving the R/3 kernel ¹
- Offline backup, only needed for:
 - Saving an entire system
 - Saving *all* R/3 stream files located in the IFS
- Reorganize the R/3 database tables using the OS/400 RGZPFM command (multiple files at a time). ^{1 2}

Notes

1. May not require you to shut down the R/3 system.
2. Is only needed in special cases such as tables containing a lot of deleted rows that cannot be reused and causing performance or space problems.

11.1.3 Unscheduled outages

The following unscheduled outages influence system availability. These items are quite different. For each of them, you need a specific solution that covers your demands:

- **Complete system loss:** A fire, flood, or other natural disaster could destroy your entire system. To rebuild your entire system, you should have a complete set of save tapes and documentation stored off site at a secure, accessible location.
- **System failure:** A system failure means that some part of your system hardware, other than the disk units subsystems, fails. Some system failures, such as processor problems, cause your system to stop without warning. This is called an *abnormal end*. When your system ends abnormally, the following problems can occur:
 - Files may be partially updated.
 - Access paths for files may be incomplete.

- Objects that are in use may be damaged.
- Relationships between files may be partially validated.

When you restart (IPL) your system after the failed component is repaired, the system analyzes the possible damage, rebuilds or recovers access paths, tries to verify file relationships, and attempts to synchronize files to transaction boundaries. The first IPL after the system ends abnormally can take a long time.

- **Power failure:** Loss of power also causes your system to end abnormally. You may experience the same types of problems that occur with a system failure. You should either apply the feature called system power control network (SPCN) or use an uninterruptible power supply (UPS).
- **Disk failure:** If a disk unit on your system fails, in most cases, the data on that disk unit is destroyed. This requires recovering all the data in the auxiliary storage pool (ASP) that contains the failed unit. The single-level storage architecture makes the iSeries server a very productive system to program and to manage. However, the architecture makes recovering from a disk failure more difficult. The system spreads information across all the disk units in an ASP to provide good performance and storage management. If a unit in an ASP is lost, you cannot determine what data was on that unit because objects are spread across the ASP. You must recover all the data in the ASP. The disk protection tools, mirrored protection and device parity protection, are designed to reduce the recovery time if a disk unit fails or, in some cases, to eliminate the need for the recovery of data.
- **Program failure or human error:** Sometimes programs are not adequately tested before they are put into production. Or a condition occurs that was not anticipated by the software developers. A program error can cause incorrect information in some of your data files. People using the system can make mistakes, too. An operator might run a month-end program twice. A data entry person might enter the same batch of orders twice. A system manager might delete a file by mistake. When these types of errors occur, you need to correct or restore the data that has been damaged.

11.1.4 Availability solutions for unscheduled outages

Availability options come in many different types. It is important to evaluate each alternative based on the degree, to which it impacts overall availability time, overall system performance, and the cost to implement.

If a failure does occur, the impact to the business must be minimized. That is why the iSeries server places continuous emphasis on the improvement of recovery times.

11.1.4.1 Hardware solutions

Table 20 shows a comparison of availability solutions. For more information, refer to the *Backup and Recovery Guide*, SC41-5304.

- **System Power Control Network:** This feature provides a function called *Continuously Powered Main Store*. If your system has this feature, a battery provides sufficient power to shut down the system and maintain the contents of memory for up to two days after a power loss. In many cases, this can significantly reduce the amount of time the system requires to perform an initial program load (IPL) after a power loss.

- **Uninterruptible power supply:** It provides auxiliary power to the processing unit, disk units, system console, and other devices that you choose to protect. You can continue operations during brief power interruptions, provide normal ending of operations to reduce the recovery time, and protect the system from voltage peaks.
- **Device parity protection:** This is intended to prevent data from being lost if a disk failure occurs. In many cases, device parity protection can also prevent your system from stopping when a disk unit fails. Device parity protection provides:
 - Technology similar to the RAID-5 (redundant array of independent disks) technique
 - Redundant power
 - Concurrent maintenance for single disk failures
 - Concurrent maintenance for power supply failures for the 9337 disk array subsystem

Note

This should be the minimum protection for the ASP on which the R/3 database library is located (usually the system ASP).

Do not include the disks for the journal receiver user ASP to the same parity set used for the database library.

- **Mirrored protection:** If a disk failure occurs, mirrored protection is intended to prevent data from being lost. Mirrored protection is a software function that uses duplicates of disk-related hardware components to keep your system available if one of the components fails. It can be used on any model of the iSeries server and is a part of the Licensed Internal Code (LIC). Different levels of mirrored protection are possible, depending on what hardware is duplicated. You can duplicate:
 - Disk units (including the load source unit): The lowest relative level of availability
 - Disk controllers
 - Disk I/O processors
 - A bus: The highest relative level of availability

Note

Mirrored protection is recommended for the user ASP that contains the journal receivers. The benefits are maximum disk unit protection and better disk I/O performance than device parity protection can accomplish.

- **Dual systems:** Installations with very high availability requirements use a dual-systems approach. Some or all data is maintained on two systems. The secondary system can take over critical application programs if the primary system fails. The most common method for maintaining data on the secondary system is through the use of journaling. Journal entries from the primary system are transmitted to the secondary system. A user-written program on the secondary system receives the journal entries and uses them to update files and other journaled objects. Several software packages are available

from business partners to support dual systems on the iSeries server. For more information, refer to 11.7, “High availability” on page 272.

Table 20. Availability solutions

Impact	Availability	Performance	Cost
High Low	Dual systems	Mirrored protection	Device parity protection
	Mirrored protection	Dual systems	Mirrored protection
	Device parity protection	Device parity protection	Dual systems

11.1.4.2 Software solutions

This section lists the available software solutions. For more information, refer to the *Backup and Recovery Guide*, SC41-5304.

- **Auxiliary storage pools (ASPs):** An ASP is a group of units that are defined from all the disk units that make up auxiliary storage. It is a software definition of how the disk units are arranged. Since V4R5M0, you can use disk management wizards in Operations Navigator to help you manage your auxiliary storage pools.

An ASP does not necessarily correspond to the physical arrangement of disks. ASPs allow you to isolate objects on one or more specific disk units. This may reduce the loss of data due to a disk media failure. In most cases, only data that is stored on disk units in the affected ASP is lost. However, when a disk unit fails, the entire system is unusable until the disk unit is repaired, unless the failed unit is protected by device parity protection or mirrored protection.

Note

We strongly recommend you create a user ASP to provide dedicated resources for journal receiver objects (in the R3<SID>JRN library) for recovery reasons and to improve the performance.

- **Access path protection:** An access path describes the order in which records in a database file are processed. A file can have multiple access paths, if different programs need to see the records in different sequences. If your system ends abnormally when access paths are in use, the system may have to rebuild the access paths before you can use the files again. This is a time-consuming process. To perform an IPL on a large, busy iSeries server that has ended abnormally can take many hours. System-managed access-path protection (SMAPP) provides a simple method to reduce your recovery time after an abnormal system end. SMAPP manages the required environment for you, and it is turned on by default. SMAPP settings can be controlled with command EDTRCYAP. You do not need to use any type of journal management to use SMAPP.
- **Journal management:** You can use journal management to recover the changes to database files or other objects that have occurred since your last complete save. You use a journal to define what files and access paths you want to protect with journal management. This is often referred to as “journaling a file or an access path”. A journal receiver contains the entries (called *journal entries*) that the system adds when events occur that are

journalled, such as changes to database files, changes to other journalled objects, or security-relevant events.

Commitment control is an extension of journal management that assists you in keeping your database files synchronized. A single transaction on your system may update more than one database file. Journaling uses two object types: *journals* (*JRN) and *journal receivers* (*JRNRCV). The journal acts like a funnel. All database table adds, changes, and deletes are received by the journal. The journal writes them to the journal receivers.

SAP R/3 journals the database files using commitment control automatically (SQL collection) but it does not journal the access paths. The journal object itself is located in the R/3 database library (R3<SID>DATA), where the journal receiver objects are stored in the journal receiver library (R3<SID>JRN).

11.1.4.3 Backup solutions

This section shows the available backup solutions for iSeries server:

- **Backup Recovery Media Services (BRMS):** You can use BRMS to simplify and automate your backups and to manage your tape library. BRMS keeps track of what you have saved, when you saved it, and where it is saved. When you need to perform a recovery, BRMS helps ensure that the correct information is restored from the correct tapes in the correct sequence. The BRMS licensed program 5769-BR1 offers a set of functions for defining and performing these tasks: backup, recovery, archiving, retrieval, and media management.

BRMS can be used in conjunction with the Job Scheduler (5769-JS1) to provide a very robust and flexible unattended automated backup strategy. A save strategy using multiple tape drives operating concurrently is recommended when performing unattended backups on systems with 300 or more gigabytes of disk. A 3494 Tape Library can further automate this approach if 3590 or 3490 tape drives are used.

- Policy driven full function tape media management software for a single, or multiple networked, iSeries server
- Auto archive and recall Hierarchical Storage Management (HSM) Dynamic Retrieval

See 11.5.4.2, “BRMS using DSCR3SYS and RCNR3SYS” on page 252, to learn how to use BRMS to save the SAP R/3 system. For more information, refer to *Backup Recovery and Media Services*, SC41-5345.

- **Tivoli Storage Manager:** You can use Tivoli Storage Manager for iSeries server (formerly known as AD5M/400) to protect data on your workstations and LAN file servers. The Tivoli Storage Manager can automatically back up critical LAN and workstation data and archive files that are used infrequently. It provides a disaster recovery solution for LANs and workstations.

You can administer the Tivoli Storage Manager from a client workstation that is attached to an iSeries server. It can back up data from a variety of workstation platforms. You can use the BRMS program to back up iSeries server user data to any Tivoli Storage Manager when the iSeries server is in a client/server environment. You can use BRMS to manage the data that you save on the Tivoli Storage Manager and to manage the backup of the system data to local media.

- Client/Server Storage Management Tool (iSeries server code only)

- Save/Restore Functions for Networked Devices Across Multiple Platforms
- Can use BRMS to manage media

For more information, refer to:

http://www.tivoli.com/support/storage_mgr/ad4serv.htm

- **OnDemand:** You can use OnDemand (5969-RD1), formally known as R/DARS, to create search criteria and a search index for large volumes of data, such as history reports or files, that you want to archive. You can archive the data either to a folder or to tape or to optical media. You can retrieve specific archived data that meets your search criteria.
 - Stores and retrieves data on DASD, optical or tape
 - Hierarchical Storage Management with record level retrieval
 - Spool file archive
 - Can use BRMS to manage media

For more information, refer to *IBM Content Manager for OnDemand for AS/400 User's Guide*, SC41-5325.

11.1.5 Availability solutions for unscheduled outages in R/3 environment

In the R/3 system environment, PTFs that must be applied to the system and OS/400 release upgrades have an impact on system availability. Some customers cannot afford to wait for hours until their systems are available again after such processes. You can minimize this situation with either logical partitioning or an effective dual systems approach. The dual systems approach is the most secure way to satisfy the requirement of continuous operations.

The software environment is mirrored in either a separate off-site location or locally on the same network with a redundant hardware and software configuration. This approach uses the journaling capabilities of the iSeries server, along with independent software options to reproduce the duplicate system over a communications link. If an outage is planned, the users are switched to the duplicate system. They continue processing until the scheduled outage is completed on the master system and journaled transactions are applied.

To use R/3 in dual systems configuration, it is necessary to obtain the necessary license keys from SAP for both the production and backup machines. Provide SAP with the following information for each machine:

- The machine serial number
- The R/3 system identifier

Once this information is provided, SAP returns the license keys. Until the license keys are obtained for a machine, the R/3 system cannot be used on that machine.

For more information, refer to 11.7, “High availability” on page 272.

11.2 Save and restore commands and menu options

Figure 153 shows the save and restore commands for the integrated file system (IFS).

Restore Commands	File System	Save Commands
RST	Root (/)	SAV, SAVR3SYS
SC41-5304 Chapters 12 & 13 RSTUSRPRF, RSTAUT, RSTCFG, RSTLIB, RSTOBJ, RST	QSYS.LIB (Library)	SAVSYS, SAVCFG, SAVSECDTA, SAVLIB, SAVOBJ, SAVCHGOBJ, SAV, SAVR3SYS
RSTDLO RST	QDLS (Document Library Services)	SAVDLO SAV
RST	QLANSrv (OS/2 Warp Server)	SAV
RST	QOpenSys (Open Systems)	SAV
RST	QNetWare (Novell Netware)	SAV
RST	User-Defined File System (/dev/QASPxx/)	SAV
RST	(Other File Systems)	SAV

Figure 153. Save and restore commands for the IFS

Note

The following file systems cannot be saved:

- NFS (Network)
- QFileSvr.400 (File server)
- QOPT (Optical)
- QNTC (Windows NT server)

Figure 154 shows the commands and menu options you can use to save parts of the system or the entire system. If you choose to use a simple save strategy, review this figure to see which parts of your system are saved when you use the options from the save menu.




Figure 154. Save commands and save menu options

Figure 155 shows the commands and menu options you can use to restore parts of the system or the entire system.




Figure 155. Restore commands and restore menu options

11.2.1 OS/400 save and restore commands

All iSeries server save and restore commands are provided in the base OS/400 operating system. If you need backup and save media management, iSeries server Backup and Recovery Media Services (BRMS) can accomplish this.

Save and restore functions can be accessed several different ways: menus via the `GO` command, indirectly through other commands, or the commands themselves. The common save and go menu commands are:

- **SAVDLO:** Save Document Library Objects
- **SAVLIB:** Save Library
- **SAVOBJ:** Save Object(s)

- **SAVSECDTA**: Save Security Data
- **SAVSTG**: Save Storage
- **SAVCFG**: Save Configuration
- **SAVSYS**: Save System
- **SAVCHGOBJ**: Save Changed Objects
- **SAV**: Save Integrated File System Objects
- **GO CMDSAV**: Displays a menu of save commands
- **GO SAVE**: Displays a menu of save options
- **GO BACKUP**: Displays a menu of backup tasks

The common restore and go menu commands are:

- **RSTDLO**: Restore Document Library Objects
- **RSTLIB**: Restore Library
- **RSTOBJ**: Restore Object(s)
- **RSTAUT**: Restore User Authorities to Objects
- **RSTCFG**: Restore System Configuration
- **RSTUSRPRF**: Restore User Profiles and Authority Tables
- **RST**: Restore Integrated File System Objects
- **GO CMDRST**: Displays a menu or panel group of restore commands
- **GO RESTORE**: Displays a menu or panel group of restore tasks

Note that there is no Restore System command since this is done as a dedicated service tools function. Nor is there a Restore Changed Object command. There are additional commands, but they are not listed or described here.

This chapter does not intend to explain each command in detail. However, you can find detailed information on the save and restore commands in *Backup and Recovery*, SC41-5304, or *CL Reference*, SC41-5722.

11.2.1.1 Save commands requiring a dedicated system

A dedicated system (restricted state) on the iSeries server only has the console and system jobs running. No users or other applications can use the system.

We list the commands that require a dedicated system. If you wrote your own save or backup programs using save-while-active and the system is in a dedicated mode, the save-while-active processing is ignored.

The commands are:

- SAVSYS
- SAVSTG
- SAVLIB LIB(*NONSYS)

11.2.1.2 Save commands not requiring a dedicated system

The following commands do not require a dedicated system (restricted state). However, use caution to prevent an interruption to production users due to the iSeries server locking their objects during the execution of these saves. The save commands are:

- SAV
- SAVCFG
- SAVCHGOBJ OBJ(*ALL) LIB(<library-name>)
- SAVCHGOBJ OBJ(*ALL) LIB(*ALLUSR)
- SAVDLO
- SAVLIB LIB(*IBM)

- SAVLIB LIB(*ALLUSR)
- SAVOBJ OBJ(<object-name/library-name>)
- SAVSECDTA

11.3 Save strategies

This section provides an overview of the objects that you need to save to make sure that you can restore your data to a consistent state in case of system failure.

11.3.1 SAP R/3 libraries

Table 21 shows the iSeries libraries that are part of the SAP R/3 product (based on the SAP R/3 database 4.6C and kernel release 4.6D).

Table 21. iSeries server libraries included in the R/3 product

Library	Description	Initial number of objects	Initial size (approx.)
R3<REL>OPT	Library for optimized executables	542	525 MB
R3<SID>DATA	Database library	29491	16.7 GB ¹
R3<SID>JRN	Journal receiver library	1	² —
R3<SID>400	Library for work management objects	26	1.4 MB
R3400	Library for not <SID> specific objects.	15	58.5 MB ³
R3WRKnn	Internal R/3 library	0	0
R3<SID>xxxxx	SQL package library		⁴ —
R3<REL>RFC	Library for RFC SDK (optional)		
R3<REL>CPIC	Library for CPIC SDK (optional)		
1. This value does not include customer data. 2. We recommend that you allow approximately 4 GB for a productive system. This can be more or less depending on the activity on the system and the frequency with which the journal receivers are backed up. 3. The R3400 library contains objects that are not <SID> specific, including the OS collector data. 4. The R3<SID><xxxxx> libraries that contain SQL packages are dynamically created when you run the R/3 application so the number of libraries and their size vary.			

In this table, consider the following facts:

- <SID> is the SAP system ID (for example "P01")
- <REL> is the R/3 kernel release (for example, "46D")
- nn is the instance number.

11.3.2 SAP R/3 IFS structure

You can see the IFS structure of SAP R/3 on the iSeries server in Figure 41 on page 74. The IFS structure that is shown only contains the important symbolic links.

11.3.3 What needs to be saved

We recommend you save the R/3 objects as shown in Table 22.

Table 22. When to save what objects

Object	Frequency
R3<SID>DATA	Daily
R3<SID>JRN	At least daily. Make sure the user ASP does not overflow!
R3<REL>OPT	At least weekly (and before and after each change)
R3<SID>400	Weekly (and before and after each change)
R3400	Weekly (and before and after each change)
R3<REL>CPIC	Weekly (and before and after each change)
R3<REL>RFC	Weekly (and before and after each change)
IFS objects (refer to Table 23)	Daily

You do not have to save the following libraries:

- R3SETUP (Installation library)
- R3<SID>xxxxx (SQL packages get created automatically)

Table 23. What to save in the IFS

Path name	Type of data	Save action
/usr/sap/<SID>/	system-specific data	Save this data from each server where an instance resides.
/sapmnt/<SID>	system-specific data	Save the system-specific data from the database server.
/sapmnt/trans	transport data	Save the data from the system where the global transport directory resides.

You also need to consider the objects and paths that you may have entered into table SPTH.

Note

Save the entire system in offline mode after the installation or upgrade of SAP R/3 is complete! Refer to 11.5.3.1, “Saving the entire system” on page 241, for the backup instructions.

Table 24 shows the recommended backup methods.

Table 24. The recommended backup methods

Object	Weekly backup	Daily backup
IFS objects	Offline (Save entire system) Section 11.5.3.1, "Saving the entire system" on page 241	Online with disconnect from database using the command SAVR3SYS R3STS(*DSCDB) or SAVR3SYS R3STS(*RUN) or BRMS (with or without DSCR3SYS and RCNR3SYS) Refer to 11.5.4, "Online backup with disconnect from the database" on page 250
R3<SID>DATA		
R3<SID>JRN		Online (SAVDLTRCV) Refer to 11.5.6, "Saving journal receivers" on page 260
R3<REL>OPT R3<SID>400 R3400 R3<REL>RFC (opt.) R3<REL>CPIC (opt.) R3SYS		-

Refer to 11.5.1, "Backup methods" on page 239, for the available backup methods.

11.3.4 Saving logical partitions

Backing up logical partitions follows *the same* procedures that you would use for a system without logical partitions. You must perform backups individually for each logical partition. However, it is possible to perform multiple saves or restores in different partitions at the same time.

The system automatically maintains the configuration data for your logical partitions. Each logical partition load source contains the LPAR configuration data. This data is *not* saved to removable media. This allows entire partition data to be restored to a system regardless of whether it has logical partitions.

You must perform each backup from the console or a workstation that is attached to that logical partition. Some backup commands require the system to be in a restricted state and the operation run from a console.

Note

Changing the operation mode to a restricted state on a primary partition will not affect other partitions.

11.3.4.1 ObjectConnect

One of the methods of saving data is using the ObjectConnect licensed program, which is included as a part of OS/400. ObjectConnect provides the ability to transfer entire objects between systems over a communications link.

ObjectConnect is a set of CL commands for sending objects between iSeries servers *simply* and efficiently. When you use an ObjectConnect command, the system sends the object directly to the target system. ObjectConnect provides better performance than other methods for transferring objects between systems. And, ObjectConnect does not require additional disk space to store an intermediate copy of the object that is being sent. ObjectConnect cannot run in a restricted state.

ObjectConnect commands are named SAVRSTxxx and are closely related to the SAVxxx and RSTxxx commands. The ObjectConnect commands mostly support the same parameters.

To use the ObjectConnect functions, you must have ObjectConnect installed on both the source and target systems. The systems must be connected and configured with one of the following methods:

- Local area network or remote communications line with APPC and APPN
- Local area network or remote communications line with TCP/IP and AnyNet support (to run APPC and APPN over TCP/IP)
- OptiConnect/400 with TCP/IP and AnyNet or with APPC and APPN

For information about how to set up AnyNet, refer to *AS/400 AnyNet Scenarios*, SG24-2531.

After the initial setup is done, you can include commands that start the ObjectConnect environment in your startup procedures. After that, sending objects is easy. You simply need to type the appropriate SAVRSTxxx command and specify the Remote Location Name (RMTLOCNAME) where the objects are to be restored. Table 25 lists the available ObjectConnect commands.

Table 25. List of available ObjectConnect commands

Command name	Short description
Save/Restore (SAVRST)	It supports the same options as the Save (SAV) command.
Save/Restore Object (SAVRSTOBJ)	It supports the same options as the Save Object (SAVOBJ) command.
Save/Restore Library (SAVRSTLIB)	It supports the same options as the Save Library (SAVLIB) command. You may also use generic values for the *LIB parameter.
Save/Restore Changed Objects (SAVRSTCHG)	It supports most of the same options as the Save Changed Objects (SAVCHGOBJ) command. You may use the OMITLIB parameter, and you may also specify generic values for the LIB parameter.
Save/Restore Document Library Object (SAVRSTDLO)	It supports the same options as the Save Document Library Object (SAVDLO) command.
Save/Restore Configuration (SAVRSTCFG)	It supports most of the same options as the Save Configuration (SAVCFG) and Restore Configuration (RSTCFG) commands.

11.4 Backup considerations

Be sure to read this section before you save anything.

11.4.1 Backup requirements

The backup and recovery requirements depend on the availability goals for a specific installation. Consider the following factors:

- The cost to the business resulting from a loss of availability during the failure
- The probability of the failure occurring
- The cost of the backup strategy such as operator time, unavailability during backup, media cost, storage costs, and so on

The cost and benefits vary depending not only on the location of the installation, but also on the company policies. They also depend on the availability of the required skills to perform the backup and recovery operations.

The strategy should cover the loss of the database, as well as the suite of application code. The strategy should consider recovery from a disaster resulting in the loss of the computer site. It must include the following components:

- **System:** Including the operating system (OS/400), user profiles, authorities, configurations, system, and network values.
- **Application software:** Required for normal operations of the business, including compilers, utilities libraries, application and general purpose libraries (QGPL), IBM licensed program libraries, job descriptions, output queues, data areas, message queues, and other application dependent objects.
- **Databases:** Containing the organization's information.

11.4.2 Backup media options

The iSeries server offers a range of IBM tape devices and offline storage media to hold the backed up information. The selection of the device depends on the volume of information to be backed up and the backup "window" of time available to complete the backup.

The available tape drives provide a range of capacities per tape volume as well as a choice of backup speeds. The actual elapsed time for backup of a user system depends not only on the tape drive selected, but also on the sizes and number of objects to be saved. The capacity of the CPU and other jobs running during the backup can also affect the backup times.

An alternative approach to backup is to perform it in an unattended environment by backing up the information to special "save files" on disk after business hours. The information can then be copied on the offline media during business hours interleaved with the operator's normal work such as report distribution and so on.

Table 26 gives an *indication only* of the maximum capacity per cartridge and the maximum media data rates for the most important tape drives.

Table 26. iSeries server tape drives

Type/ model	Description	Maximum capacity per cartridge		Maximum media data rate	
		Uncompressed	Compressed	Uncompressed	Compressed
3490E-Fxx	10 cartridges (½ inch)	800 MB	2.4 GB	3 MB/s	9 MB/s
7208-342	1 cartridge (8 mm)	20 GB	40 GB	3 MB/s	6 MB/s
3570-Bxx	20 cartridges (0.31 inch)	5 GB	15 GB	2.2 MB/s	6.6 MB/s
3570-Cxx	20 cartridges (0.31 inch)	7 GB	21 GB	7 MB/s	15 MB/s
3575-Lxx	60-324 cartridges (0.31 inch)	7 GB	21 GB	7 MB/s	15 MB/s
3590-Bxx	10 cartridges (½ inch)	20 GB	60 GB	9 MB/s	27 MB/s
3590-Exx	10 cartridges (½ inch)	40 GB	120 GB	14 MB/s	42 MB/s
3494-xxx	210-6240 cartridges (½ inch)	40 GB	120 GB	14 MB/s	42 MB/s

The actual save and restore performance can vary significantly based on the iSeries server processing speed, the type of data, the number and speed of the disk units, and the operating characteristics of the tape unit.

11.4.3 Before you begin

Please read the following recommendations before you begin to back up your system:

- **Backup performance:** If you have a large system or a very short window of time in which to perform a backup, we recommend you use the 3590 or 3570 tape drives. They are high speed, high capacity, and very reliable tape drives. The 3590 tape drive is recommended on systems with more than 100 GB of disk space.
- **Parallel backup and restore support:** This enhancement can significantly speed up the backup and recovery process of very large libraries and objects. For a save function, this support enables the system to spread portions of the same object onto multiple tapes concurrently. The system records the information about objects saved in parallel on each tape. Objects saved in parallel should be restored in parallel. If necessary, they can be restored with fewer tape drives. When using this support, the save commands are limited to use a single library per command.

An object type called the *Media Definition Object* is used to specify the tape devices and media used for the parallel save or restore. An authorized user uses the DEV(*MEDDFN) parameter on the appropriate save or restore command to initiate the parallel save and restore. The *MEDDFN objects can be created, modified, and deleted in one of two ways:

- With a program that uses system APIs. An authorized user program can use the APIs to create and modify the media definition object. The devices that you specify in a media definition must be compatible stand-alone tape devices or tape libraries. The tape volumes that you specify must have

compatible media formats. Please refer to the *System API Reference*, SC41-5801, for more information about these APIs.

- Using the Backup and Recovery Media Services (BRMS) licensed program. BRMS automatically creates and manages media definition objects and also keeps track of which tapes contain the saved object portions. Therefore, we recommend that you use BRMS for all parallel save and restore related tasks.
- **Separate backup media:** Use separate backup media for the R/3 database and journal receivers to provide a better protection.
- **Saving spooled files:** Please note that spooled output cannot be backed up using any of the standard options available with OS/400, but it is possible using BRMS, OnDemand, or a vendor-supplied product.
- **Stop SAP R/3 before offline backup:** Make sure you've ended the R/3 system (central and secondary instances) before you perform an offline backup.
- **Symbolic links:** Saving symbolic links in the IFS does not save the object to which it points. Saving the symbolic link itself is all it does. Therefore, do not save, for example the '/usr/sap/trans' directory. Instead, you should save the '/sapmnt/trans' directory on the system that holds the global transport directory.
- **Download SAVR3SYS from sapservX:** The SAVR3SYS, DSCR3SYS, and RCNR3SYS commands are only up-to-date on sapservX. The versions that have been shipped with the kernel CD-ROM may be out-of-date. Therefore, you have to replace it as described in SAP Note 86305.
- **Save access paths:** Saving the access paths can significantly reduce the recovery time because the system does not have to rebuild the access paths. You should consider that it takes more time, and it consumes more storage on your backup media. We recommend you save the access paths and specify ACCPTH(*YES) wherever possible for save commands.
- **Precheck option:** You should use the precheck parameter when you save objects to ensure that all of the objects you intend to save can be successfully saved, especially when using online backup methods. When you specify PRECHK(*YES), all of the objects you are saving in a library must meet the conditions. If they do not, no objects in the library are saved. Do not use this option for saving IFS objects online because the backup will always fail.
- **Update history option:** We recommend you update the save history information for objects that you are going to save. To do so, specify UPDHST(*YES) for the SAV and SAVxxx commands.
- **Saving IFS objects online:** It is currently not possible to save the all R/3 IFS objects online because, for example, developer trace files ('/usr/sap/R01/DVEBMGS01/work/dev_<xxx>') are permanently in use and do not have the system attribute QP0L_ATTR_ALWCKPWRT set. This attribute allows the SAV command to save the objects with the SAVACT(*YES) parameter and SAVACTOPT(*ALWCKPWRT). However, this is not possible in the current SAP R/3 kernel release (4.6D).

You have to either save the R/3 IFS objects offline or save them online knowingly that some objects cannot be saved. Those stream files are usually not absolutely necessary like the developer trace files, statistic trace files, and

so on. But you have to read the job log to find out if something important could not be saved.

- **Job logs and save listings:** Always check the spooled output of the save procedure and the job log to be absolutely sure that the save operation completed successfully. Furthermore, you should retain the information for disaster recovery.

11.5 Backup

Backup and recovery options for SAP R/3 on an iSeries server are managed using *native iSeries server facilities*. BRMS can provide automated tape backup and archive operations, recovery services, and tape media management to treat tape as an extension of DASD from the application point of view.

The OS/400 provides many backup facilities that can be selected from menu options, depending on the backup requirements. These facilities are an integral part of the operating system, as is the database management system, security, and so on. We do not provide a full save strategy and procedure in this section. Therefore, refer you to *Backup and Recovery*, SC41-5304, for full, detailed coverage of the subject. We include this section for your general guidance for developing a suitable save strategy.

11.5.1 Backup methods

There are a three methods of saving the SAP R/3 database on iSeries server:

- **Offline backup:** This means saving the objects while either R/3 is down or the iSeries server is in restricted state (which implies that R/3 is down).
- **Online backup with disconnect from database:** This method disconnects each of the SAP work processes from the database, effectively suspending them. This allows them to reach a commitment boundary within 5 minutes, performs the save, and allows the work processes to continue once a save-while-active checkpoint has been reached. If they do not reach the commit within 5 minutes, they are rolled back.
- **Online backup without disconnect:** This function allows you to continue using the R/3 while the database is being backed up. You have to make sure that there no long running R/3 background jobs are active at backup time, because background jobs do not commit their work very often. Therefore, it doesn't make sense to wait until they reach the commitment boundary.

Table 27 lists the advantages and disadvantages of each method.

Table 27. *Methods of saving the database*

Method	Advantages	Disadvantages
Offline backup	You do not have to take care of synchronization at all.	R/3 needs to be down for a relatively long time.
	The fastest way to save the data.	The content of the R/3 buffers is lost.

Method	Advantages	Disadvantages
Online backup with disconnect from database	Checkpoint processing usually faster than without disconnect. You do not lose the contents of the R/3 buffers.	The work processes must be disconnected from database for a while. Rollback will be forced for jobs that do not commit their work within five minutes.
Online backup without disconnect	No disconnect from database necessary. You do not lose the contents of the R/3 buffers.	Save operation may not be successful if jobs do not commit their work within the specified amount of time.

11.5.2 Initializing the tape

Before you begin with the backup operation, perform these tasks:

1. Make sure you have enough tapes.
2. Clean the read and write head of your tape unit.
3. Insert a blank tape into the tape drive.
4. Initialize the tape using the `INZTAP` command as shown in Figure 156. If your media density does not comply with your tape drive, you may have to specify a different density.

Initialize Tape (INZTAP)

Type choices, press Enter.

Device	> TAP01	Name
New volume identifier	> SAPR3	Character value, *NONE...
New owner identifier	*BLANK	
Volume identifier	*MOUNTED	Character value, *MOUNTED
Check for active files	> *NO	*YES, *NO, *FIRST
Tape density	*DEVTYPE	*DEVTYPE, *CTGTTYPE, *QIC120...
Code	*EBCDIC	*EBCDIC, *ASCII
End of tape option	*REWIND	*REWIND, *UNLOAD
Clear	*NO	*NO, *YES

Bottom

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

Figure 156. Initialize Tape (INZTAP)

11.5.3 Offline backup

The backup options can either be selected from the save menu (GO SAVE) or a command line entry.

11.5.3.1 Saving the entire system

This task saves the entire iSeries server, including all SAP R/3 objects. This must be done after the initial installation of SAP R/3 is complete. Periodic entire system saves are recommended on a weekly basis, as well as before or after major system changes, such as:

- Hardware addition/removals
- Extensive system configuration changes
- Operating system upgrades
- PTF installations
- Application software additions, removals, and upgrades

Note

This option requires a dedicated system or *restricted state*, which means that no subsystem, except the controlling subsystem (for example QCTL), is available, and only to support the console. Upon completion, the controlling subsystem is restarted, and the system is made ready for production. If the STARTSAP command is not in the IPL startup program, it has to be issued manually to start SAP R/3.

To perform the backup operation, follow these steps:

1. Stop SAP R/3.
2. Sign on from the console as QSECOFR.
3. Go to the save menu by using the GO SAVE command and page down. The screen shown in Figure 157 appears.

The screenshot shows a terminal window with the following content:

```
SAVE                                Save                                System:  AS23

Select one of the following:

  Save System and User Data
    20. Define save system and user data defaults
    21. Entire system
    22. System data only
    23. All user data

  Save Document Library Objects
    30. All documents, folders, and mail
    31. New and changed documents, new folders, all mail
    32. Documents and folders
    33. Mail only
    34. Calendars

More...

Selection or command
====> 21

F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel  F13=Information Assistant
F16=AS/400 Main menu
```

Figure 157. GO SAVE: Option 21 (Entire system)

4. Select option 21 to save the entire system. Press Enter to view what this option will do as shown in Figure 158.

Save the Entire System

What this option will do

- o End all subsystems
- o Save the Licensed Internal Code
- o Save the operating system
- o Save the security data
- o Save the device configuration objects
- o Save all user libraries (including libraries for licensed programs)
- o Save all documents and folders
- o Save all distribution and mail objects
- o Save all directories
- o Start the controlling subsystem

What this option will not do

- o Save the contents of job queues, output queues, or data queues that exist on the system

Press Enter to continue.

F3=Exit F12=Cancel

Figure 158. GO SAVE: Option 21 (Save the entire system) (Part 1 of 3)

Figure 159 shows the recommended save options. For more information about these options, refer to *Backup and Recovery*, SC41-5304.

Specify Command Defaults

Type choices, press Enter.

Devices > TAP01	Names
Prompt for commands > N	Y=Yes, N=No
Check for active files > N	Y=Yes, N=No
Message queue delivery > *NOTIFY	*BREAK, *NOTIFY
Start time *CURRENT	*CURRENT, time
Vary off network servers > *ALL	*NONE, *ALL, *BASE, *WINDOWSNT
Unmount file systems > Y	Y=Yes, N=No
More...	
F3=Exit F12=Cancel	

Figure 159. GO SAVE: Option 21 (Save the entire system) (Part 2 of 3)

5. You can use the system reply list to perform an unattended backup operation as shown in Figure 160. See the command help for requirements.

Specify Command Defaults

Type choices, press Enter.

Print system information . . . > Y	Y=Yes, N=No
Use system reply list > Y	Y=Yes, N=No

F3=Exit F12=Cancel
Bottom

Figure 160. GO SAVE: Option 21 (Save the entire system) (Part 3 of 3)

6. Check the spooled output of the save procedure and the job log to be absolutely sure that the save operation completed successfully. You should retain the information for disaster recovery.

This backup or save, when completed, produces an offline backup of the entire system to tape. These tapes can be used for recovery purposes of the entire system or individual objects.

11.5.3.2 Saving all user data

Option 23 from the save menu saves all SAP R/3 objects and the following data:

- Security data (SAVSECDTA)
- Configuration data (SAVCFG)
- All user libraries (SAVLIB LIB(*ALLUSR))
- All folders (SAVDLO)
- Integrated File System (SAV OBJ('/*')) (but omits /QSYS.LIB and /QDLS)

That means this option does not save the Licensed Internal Code or the operating system, but it saves all other data that is included in option 21 (save the entire system).

We recommend that you *do not* use this option because you'll reduce the backup time very much compared to option 21. And you cannot restore an older version of the Licensed Internal Code or the operating system which, might be helpful in case of a software failure that came along with a PTF.

11.5.3.3 Saving changed objects

Select option 6 from the Save menu to save only the objects that have been changed since a particular referenced date and time, or since the last time the library was saved with the SAVLIB command.

In a situation where a single record of the database objects has changed, this option may not be advantageous, since the *entire* object is saved if it was changed since the referenced time.

If only a subset of the R/3 database files change, then you could save time during the backup process. You need to balance saving time during the backup against the longer and more complicated recovery process with the SAVCHGOBJ command.

Note

Specify OBJJRN(*YES) when you use the Save Changed Objects (SAVCHGOBJ) command.

11.5.3.4 Saving SAP R/3 relevant data manually

If you apply patches for the R/3 system, or you create a new system or instance, perform a backup of the R/3 system before and after the changes. This section shows how to save the R/3 database and IFS objects.

Refer to 11.3, “Save strategies” on page 232, for what needs to be saved. Section 11.5.6, “Saving journal receivers” on page 260, tells you how to save journal receivers.

To perform the backup operation, follow these steps:

1. Stop SAP R/3.
2. Sign on as QSECOFR or <SID>OFR.
3. Go to the Save menu using the GO SAVE command. Select option 11 (Save object in directories). Or you can prompt the SAV command, and enter the device name and '+' for more values in the Objects parameter as shown in Figure 161.

Save Object (SAV)

Type choices, press Enter.

Device '/qsys.lib/tap01.devd'

+ for more values

Objects: +

Name '*'

Include or omit *INCLUDE *INCLUDE, *OMIT

+ for more values

Directory subtree *ALL *ALL, *DIR, *NONE, *OBJ

Save active *NO *NO, *YES, *SYNC

Bottom

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys

Figure 161. SAV: Save objects in directory (Part 1 of 4)

4. Enter the IFS paths as shown in Figure 162.

```

                                Specify More Values for Parameter OBJ

Type choices, press Enter.

Objects:
  Name . . . . . > '/usr/sap/R01/DVEBMGS01'

  Include or omit . . . . . *INCLUDE      *INCLUDE, *OMIT

  Name . . . . . > '/sapmnt/R01'

  Include or omit . . . . . *INCLUDE      *INCLUDE, *OMIT

  Name . . . . . > '/sapmnt/trans'

  Include or omit . . . . . *INCLUDE      *INCLUDE, *OMIT

  Name . . . . . '*'

  Include or omit . . . . . *INCLUDE      *INCLUDE, *OMIT

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

Figure 162. SAV: Save objects in directory (Part 2 of 4)

5. Press F9 and F10 to see the additional parameters as shown in Figure 163. Specify *PRINT for the Output and *LEAVE for End of media option (to save rewind time).

```

                                Save Object (SAV)

Type choices, press Enter.

Output . . . . . > *PRINT

Volume identifier . . . . . *MOUNTED
+ for more values

Label . . . . . *GEN
Optical file . . . . . '*'

Sequence number . . . . . *END      1-16777215, *END
File expiration date . . . . . *PERM   Date, *PERM
End of media option . . . . . > *LEAVE  *REWIND, *LEAVE, *UNLOAD
Use optimum block . . . . . *YES      *YES, *NO
Save active message queue . . . *NONE

Type of output information . . . *ALL      *ALL, *ERR, *SUMMARY

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

Figure 163. SAV: Save objects in directory (Part 3 of 4)

6. Specify *YES for Object pre-check and *YES for Update history as shown in Figure 164.

```

                                Save Object (SAV)

Type choices, press Enter.

                                Additional Parameters

System . . . . . *LCL          *ALL, *LCL, *RMT
Time period for last change:
  Start date . . . . . *ALL      Date, *ALL, *LASTSAVE
  Start time . . . . . *ALL      Time, *ALL
  End date . . . . . *ALL       Date, *ALL
  End time . . . . . *ALL       Time, *ALL
Object pre-check . . . . . > *YES *NO, *YES
Target release . . . . . *CURRENT *CURRENT, *PRV, V3R2M0...
Update history . . . . . > *YES  *NO, *YES, *SYS, *PC

Clear . . . . . *NONE          *NONE, *ALL, *AFTER, *REPLACE
Data compression . . . . . *DEV  *DEV, *NO, *YES
Data compaction . . . . . *DEV  *DEV, *NO

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

```

Figure 164. SAV: Save objects in directory (Part 4 of 4)

7. Check the job log to make sure that all IFS objects could be saved. You should retain the information for disaster recovery.
8. Save the R/3 database by selecting option 2 (Save libraries) from the save menu, or prompt the SAVLIB command as shown in Figure 165.

```

                                Save Library (SAVLIB)

Type choices, press Enter.

Library . . . . . > R3R01DATA    Name, generic*, *NONSYS...
+ for more values
Device . . . . . > TAP01        Name, *SAVF, *MEDDFN
+ for more values
Volume identifier . . . . . *MOUNTED
+ for more values
Sequence number . . . . . *END    1-16777215, *END
Label . . . . . *LIB
File expiration date . . . . . *PERM Date, *PERM
End of media option . . . . . *REWIND *REWIND, *LEAVE, *UNLOAD
Use optimum block . . . . . *YES  *YES, *NO

                                Additional Parameters

Target release . . . . . *CURRENT *CURRENT, *PRV, V3R2M0...
Update history . . . . . *YES     *YES, *NO

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

```

Figure 165. SAVLIB: Save the R/3 database (Part 1 of 3)

9. Specify *YES for Object pre-check and *YES for Save access paths as shown in Figure 166.

Save Library (SAVLIB)		
Type choices, press Enter.		
Clear	*NONE	*NONE, *ALL, *AFTER, *REPLACE
Object pre-check >	*YES	*NO, *YES
Save active	*NO	*NO, *LIB, *SYNCLIB, *SYSDFN
Save active wait time	120	0-99999, *NOMAX
Save active message queue	*NONE	Name, *NONE, *WRKSTN
Library	*LIBL	Name, *LIBL, *CURLIB
Save access paths >	*YES	*NO, *YES
Save file data	*YES	*YES, *NO
Storage	*KEEP	*KEEP, *FREE
Data compression	*DEV	*DEV, *NO, *YES
Data compaction	*DEV	*DEV, *NO
Libraries to omit	*NONE	Name, generic*, *NONE
+ for more values		
		More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 166. SAVLIB: Save all R/3 libraries (Part 2 of 3)

10.Specify *PRINT for the Output parameter as shown in Figure 167.

Save Library (SAVLIB)		
Type choices, press Enter.		
Objects to omit:		
Object	*NONE	Name, generic*, *NONE, *ALL
Library	*ALL	Name, generic*, *ALL
Object type	*ALL	*ALL, *ALRTBL, *BNDDIR...
+ for more values		
Output >	*PRINT	*NONE, *PRINT, *OUTFILE
File to receive output		Name
Library	*LIBL	Name, *LIBL, *CURLIB
Output member options:		
Member to receive output . . .	*FIRST	Name, *FIRST
Replace or add records . . .	*REPLACE	*REPLACE, *ADD
Type of output information . . .	*OBJ	*OBJ, *LIB, *MBR, *ERR
		Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 167. SAVLIB: Save all R/3 libraries (Part 3 of 3)

11.Check the spooled output of the SAVLIB procedure and the job log to be absolutely sure that the save operation completed successfully. You should retain the information for disaster recovery.

11.5.3.5 Saving R/3 database and IFS objects

Save the R/3 database and IFS objects with command SAVR3SYS R3STS(*END). Refer to 11.3, “Save strategies” on page 232, for what needs to be

saved. Refer to 11.5.4.1, “Saving the R/3 database and IFS objects” on page 250, for a detailed description of the SAVR3SYS command. You can also refer to 11.5.6, “Saving journal receivers” on page 260, which tells you how to save journal receivers.

You should use *SPTH for IFS files to save rather than specifying the paths manually. This determines the IFS objects to save from the R/3 table SPTH. Therefore, you should update the table in R/3 with the information from Table 23 on page 233 before you proceed.

To perform the backup operation, follow these steps:

1. Sign on as <SID>OFR.
2. Prompt the SAVR3SYS command and enter the SID and path to the tape drive. Specify *END for R/3 status during save as shown in Figure 168. This brings down the R/3 system automatically and restarts it when the save operation has finished.

Save R/3 System (SAVR3SYS)

Type choices, press Enter.

System ID	> R01	Character value
Device	> '/qsys.lib/tap01.devd'	
Prompt save commands	*NO	*YES, *NO
R/3 status during save	> *END	*DSCDB, *RUN, *END
IFS files to save:		
Path name	*SPTH	
Include or omit		*INCLUDE, *OMIT
+ for more values		
Save R/3 database	*YES	*YES, *NO
Save R/3 configuration	*NO	*YES, *NO
Save R/3 SQL packages	*NO	*YES, *NO
Save reference date	*ALL	Date, *ALL, *LASTSAVE
Save reference time	*ALL	Time, *ALL, *LASTSAVE
Expiration date	*PERM	Date, *PERM
Save active wait time	1200	Number, *NOMAX

More...

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

Figure 168. SAVR3SYS R3STS(*END)

11.5.3.6 Saving SAP R/3 relevant data using BRMS

You can use Backup Recovery Media Services (BRMS) to save the entire SAP R/3 environment. This section explains how to configure Links and Backup Control Groups only. For more information about how to configure and start the backup, refer to *Backup Recovery and Media Services*, SC41-5345.

Section 11.5.6, “Saving journal receivers” on page 260, tells you how to save journal receivers.

1. Use the WRKLB RM command to add a link for saving R/3 IFS objects as shown in Figure 169.

Display Link List (DSPLNKLBRM)

Type choices, press Enter.

List > R3IFS Character value

Objects:

Name > '/usr/sap/R01/DVEBMGS01'

Include or omit *INCLUDE *INCLUDE, *OMIT

Name > '/sapmnt/R01'

Include or omit *INCLUDE *INCLUDE, *OMIT

Name > '/sapmnt/trans'

Include or omit *INCLUDE *INCLUDE, *OMIT

Directory subtree *ALL *ALL, *DIR, *NONE, *OBJ

Text > 'Saving SAP R/3 IFS objects'

Bottom

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

F24=More keys

Figure 169. Link definition: BRMS offline backup

2. Use the Work with Backup Control Groups (WRKCTLGBRM) command to create a new control group with option 1. You can use the example from Figure 170 to add similar entries to your backup control group.

Edit Backup Control Group Entries

Group : SAVR3OFF

Default activity *BKUPCY

Text Offline backup of SAP R/3 without saving *JRNRCV

Type information, press Enter.

Seq	Backup Items	List Type	Weekly Activity SMTWTFS	Retain Object Detail	Save While Active	SWA Message Queue
10	R3IFS	*LNK	*DFTACT	*YES	*NO	
20	R3R01DATA		*DFTACT	*NO	*NO	
30	R346DOPT		*DFTACT	*NO	*NO	
40	R3R01400		*DFTACT	*NO	*NO	
50	R3400		*DFTACT	*NO	*NO	
60	R346DRFC		*DFTACT	*NO	*NO	
70	R346DCPIC		*DFTACT	*NO	*NO	

Bottom

F3=Exit F5=Refresh F10=Change item

F11=Display exits F12=Cancel F24=More keys

Figure 170. WRKCTLGBRM: BRMS offline backup

We recommend you use the <SID>OFR user profile to run the scheduled job.

11.5.4 Online backup with disconnect from the database

This method disconnects each of the SAP work processes from the database and effectively suspends them. This allows them to reach a commitment boundary within 5 minutes, does the save, and allows the work processes to continue once a save-while-active checkpoint has been reached. If they do not reach the commit within 5 minutes, they are rolled back.

You have basically the two options to save your R/3 system online with disconnect from the database: SAVR3SYS command and BRMS.

11.5.4.1 Saving the R/3 database and IFS objects

This section shows you how to save the R/3 database and IFS objects with the SAVR3SYS R3STS(*DSCDB) command. Refer to 11.3, “Save strategies” on page 232, for what needs to be saved. You can also refer to 11.5.6, “Saving journal receivers” on page 260, which tells you how to save journal receivers.

The SAVR3SYS command saves all the information required for an R/3 system. The intent of this command is to run on the central R/3 system (the system that has the R/3 database and the critical R/3 stream files, such as /sapmnt/trans).

You should use *SPTH for IFS files to save rather than specifying the paths manually. This determines the IFS objects to save from the R/3 table SPTH. Therefore, you should update the table in R/3 with the information from Table 23 on page 233 before you proceed.

To perform the backup operation, follow these steps:

1. Sign on as <SID>OFR.
2. Prompt the SAVR3SYS command. Enter the SID and path to the tape drive. Specify *DSCDB for R/3 status during the save as shown in Figure 171. You do not have to end the R/3 system.

Save R/3 System (SAVR3SYS)

Type choices, press Enter.

System ID	> R01	Character value
Device	> '/qsys.lib/tap01.devd'	
Prompt save commands	*NO	*YES, *NO
R/3 status during save	> *DSCDB	*DSCDB, *RUN, *END
IFS files to save:		
Path name	*SPTH	
Include or omit		*INCLUDE, *OMIT
+ for more values		
Save R/3 database	*YES	*YES, *NO
Save R/3 configuration	*NO	*YES, *NO
Save R/3 SQL packages	*NO	*YES, *NO
Save reference date	*ALL	Date, *ALL, *LASTSAVE
Save reference time	*ALL	Time, *ALL, *LASTSAVE
Expiration date	*PERM	Date, *PERM
Save active wait time	> 1200	Number, *NOMAX
More...		
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display		
F24=More keys		

Figure 171. SAVR3SYS R3STS(*DSCDB)

The SAVR3SYS command has the parameters shown in Table 28 (the default values are indicated by ">").

Table 28. SAVR3SYS parameters

Parameter	Description
SID	Specify the SAP system ID of the R/3 instance to be saved. This is a required parameter.
DEV	Specify the name of the tape device on which the R/3 system should be saved. This is specified in a "longname" format (that is, '/QSYS.LIB/TAP01.DEVD'). This is a required parameter.
PMTCMD	Specifies whether each of the save commands should be prompted by the system (allowing you to change additional options on the save). > *NO : Do not prompt the save commands. *YES : Prompt each of the save commands.
R3STS	Specify what should be done with the R/3 system during the save. > *DSCDB : Disconnect each of the SAP work processes from the database, effectively suspending them. This allows them to reach a commitment boundary within 5 minutes, does the save, and allows the work processes to continue once a save-while-active checkpoint has been reached. If they do not reach the commit within 5 minutes they will be rolled back. Using this option, you do not lose the contents of the R/3 buffers. *END : End the SAP system (and re-start it after the database and IFS saves are complete). *RUN : Allow the SAP system to keep running during the save.
IFS	Specify which IFS files to save (similar to the SAV command). > *SPTH : Determine the IFS files to save from the R/3 table SPTH. This table can be maintained in R/3. *OMIT : Use when specifying a file that should be omitted from the save. *INCLUDE : Use when specifying a file that should be included with the save.
SAVDB	Specifies whether the R/3 database (R3<SID>DATA library) should be included in the save. > *YES : Save the database library. *NO : Do not save the database library.
SAVCFG	Specifies whether the R/3 configuration libraries (R3400 and R3<SID>400) should be included in the save. > *NO : Do not save the configuration libraries. *YES : Save the configuration libraries.
SAVPKG	Specifies whether the R/3 SQL packages (containing prepared statements) should be included in the save. > *NO : Do not save the SQL package libraries. *YES : Save the SQL package libraries.
REFDATE	Specifies whether a full save or whether a save of only the changed objects (since the last save or a specified time) should be done. > *ALL : Save all specified objects, regardless of whether they have changed (REFTIME(*ALL) must also be specified). *LASTSAVE : Save all the objects that have changed since the last time the SAVR3SYS command was run (REFTIME(*LASTSAVE) must also be specified). Date : All changed objects from the specified date (and time on the REFTIME parameter) will be saved.

Parameter	Description
REFTIME	Specifies whether a full save or whether a save of only the changed objects (since the last save or a specified time) should be done. >*ALL: Save all specified objects, regardless of whether they have changed or not (REFDATE(*ALL) must also be specified). *LASTSAVE: Save all the objects that have changed since the last time the SAVR3SYS command was run (REFDATE(*LASTSAVE) must also be specified). Date: All changed objects from the specified time (and date on the REFDATE parameter) will be saved.
EXPDATE	Specifies the expiration date for the files written to the tape. All files written to the tape will be given the same expiration date. Files cannot be written over until the expiration date. >*PERM: The files will be protected permanently. Date: The files will not be allowed to be written over until the specified date.
SAVACTWAIT	Specifies the amount of time (in seconds) that the save operation should wait for the save synchronization point. This only affects the database save operations. >1200: The system will wait for the database to reach its synchronization point within 20 minutes. Number of seconds: The system will wait for the database to reach its synchronization point up to the specified number of seconds.
VOL	Specifies up to 75 volumes that will be used for the save operations. If volume names are specified, the volumes must be mounted in the order specified. >*MOUNTED: Use the mounted volumes regardless of the volume name. Volume: Use the specified volume name only.

The SAVR3SYS process consists of the following operations:

1. End R/3 or disconnect R/3 from the database (depending on the R3STS parameter).
2. Save the IFS, as specified by the IFS, REFDATE, and REFTIME parameters.
3. Save the work management objects, as specified by the SAVCFG, REFDATE, and REFTIME parameters.
4. Save the SQL packages, as specified by the SAVPKG, REFDATE, and REFTIME parameters.
5. Save the database, as specified by the SAVDB, REFDATE, and REFTIME parameters. A save while active is done.
6. Once the checkpoint is reached, the R/3 is restarted or reconnected to the database, depending on the R3STS parameter.

This command only saves SAP R/3 related objects. Any non-R/3 related or iSeries server specific objects (such as Q-libraries, user profiles, and so on) must be saved using standard iSeries server save commands or menu options.

11.5.4.2 BRMS using DSCR3SYS and RCNR3SYS

The new commands Disconnect R/3 System (DSCR3SYS) and Reconnect R/3 System (RCNR3SYS) were developed for online save operations using BRMS.

This section explained how to configure Links and Backup Control Groups only. For more information about how to configure and start the backup, refer to *Backup Recovery and Media Services*, SC41-5345. You can also refer to 11.5.6, “Saving journal receivers” on page 260, which tells you how to save journal receivers.

The concept of disconnecting the database from the work processes is exactly the same as used by the SAVR3SYS R3STS(*DSCDB) command. But direct use of the SAVR3SYS command in BRMS does not make sense.

Refer to Figure 169 on page 249 for the link list definition. We recommend that you use similar logic for the Backup Control Group as shown in Figure 172 and Figure 173.

Edit Backup Control Group Entries

Group : SAVR3DSC
 Default activity *BKUPCY
 Text Online backup of SAP R/3 w DSC, w/o *JRNRCV

Type information, press Enter.

Seq	Backup Items	List Type	Weekly Activity SMTWTFSS	Retain Object Detail	Save While Active	SWA Message Queue
10	R3IFS	*LNK	*DFTACT	*YES	*NO	
20	*EXIT		*DFTACT			
30	*EXIT		*DFTACT			
40	R3R01DATA		*DFTACT	*NO	*SYNCLIB	*LIB
50	R3R01400		*DFTACT	*NO	*SYNCLIB	*LIB
60	R346DOPT		*DFTACT	*NO	*NO	
70	R3400		*DFTACT	*NO	*NO	
80	R346DRFC		*DFTACT	*NO	*NO	

More...

F3=Exit F5=Refresh F10=Change item
 F11=Display exits F12=Cancel F24=More keys

Figure 172. WRKCTLGBRM: BRMS online with disconnect from database (Part 1 of 2)

Edit Backup Control Group Entries						
Group : SAVR3DSC						
Default activity *BKUPCY						
Text Online backup of SAP R/3 w/o DSC, w/o *JRNRCV						
Type information, press Enter.						
Seq	Backup Items	List Type	Weekly Activity SMIWTFS	Retain Object Detail	Save While Active	SWA Message Queue
90	R346DCPIC		*DFTACT	*NO	*NO	
						Bottom
F3=Exit		F5=Refresh		F10=Change item		
F11=Display exits		F12=Cancel		F24=More keys		

Figure 173. WRKCTLGBRM: BRMS online with disconnect from database (Part 2 of 2)

Remember it does not make sense to specify *YES for Save While Active for the IFS backup.

You should specify *SYNCLIB for:

- R3<SID>DATA database library
- R3<SID>400 library (because the memory-based database monitor is using files in that library)

All the other R/3 libraries can be saved successfully without using save-while-active. The benefit is that the overall backup procedure is faster compared to the one that uses checkpoint processing for all the libraries. That means the tape drive can be used for another save operation, for example if the tape drive is shared between two systems.

Add the following command to the exit points:

```
20: DSCR3SYS SID(R01)
30: MONSWABRM LIB(R3R01DATA) CMD(RCNR3SYS SID(R01)) WAITMSG(600)
```

Here's how it works:

1. The first exit disconnects the database from R/3 work processes. It waits for 5 minutes so that each work process reaches a commit boundary during that delay time. By the end of 5 minutes, if the work process does not reach the commit boundary, that job is rolled back.
2. The second exit in the control group starts the MONSWABRM job that will wait for the checkpoint processing on the library R3<SID>DATA. As soon as the check point processing is complete, this job reconnects the database to R/3 work processes.

We recommend you use the <SID>OFR user profile to run the scheduled job.

11.5.5 Online backup without disconnect

This section shows how to save the R/3 database and IFS objects. Refer to 11.3, “Save strategies” on page 232, for what needs to be saved. You should also refer to 11.5.6, “Saving journal receivers” on page 260, which explains how to save journal receivers.

This function allows iSeries server objects to be backed up while they are in use. In contrast, saving objects without this facility requires that the objects are not in use, and even the online backup with disconnect from database does not allow the work process to continue until the checkpoint has been reached. Due to the additional processing involved, save-while-active takes longer to complete than without it. However, to minimize the duration of the backup, it is a good idea to perform a save-while-active *during periods of low activity*.

The save-while-active checkpoint processing waits until all committable resources in the save request are at a commitment boundary (with respect to all jobs making committable changes to the objects being saved) prior to actually saving the objects to the save media. This is done so that no partial transactions are saved to the save media by a save-while-active operation.

Note

The save operation may not be successful (checkpoint cannot be reached) if jobs do not commit their work within the specified amount of time. See SAP Note 202593 for information on how to solve problems with backup.

When commitment control is active for any job on the system (which is the case for jobs running the SAP R/3 subsystem), the system performs the following tasks during the save-while-active checkpoint processing:

- Identifies all jobs that have one or more commitment definitions with pending committable changes related to the objects being saved by the save-while-active operation.
- For identified jobs, allows additional committable changes to be made for any commitment definitions already started or to be started. Additional committable changes are allowed for the objects so that all pending changes for the objects saved by the save-while-active operation can be committed or rolled back.
- Delays any job that attempts to make a committable change to an object being saved by the save-while-active operation. All commitment definitions for the job do not have any pending changes for any objects being saved by the save-while-active operation. The job is delayed only until the checkpoint processing is completed for the save-while active operation.

The Save active wait (SAVACTWAIT) parameter value on the save commands can be used to control the amount of time allowed for jobs to reach, and be delayed at, a commitment boundary.

When the save operation starts, the system establishes a *checkpoint image*. While the specified save is in progress and a request is received for an object to be changed, the system takes a copy of the *pages to be changed*, and the changes proceed on the original object. The copies of the pages *before the changes* were made allow the system to perform a complete backup of the object.

The *save time* of the object is the time at which the request started, which is the time the checkpoint image was established. But this process consumes a lot of resources.

11.5.5.1 Save-while-active using native commands

The SAP R/3 IFS structure can currently not be saved using the save-while-active approach. Therefore, it does not make sense to specify *YES for the SAVACT parameter. Refer to 11.5.3.4, “Saving SAP R/3 relevant data manually” on page 244, to learn how to save the R/3 IFS objects using the SAV command.

To perform the backup operation, follow these steps:

1. Sign on as QSECOFR or <SID>OFR.
2. Prompt the SAVLIB command, and specify the value as shown in Figure 174.

Save Library (SAVLIB)

Type choices, press Enter.

Library	> R3R01DATA	Name, generic*, *NONSYS...
	+ for more values	
Device	> TAP01	Name, *SAVF, *MEDDEFN
	+ for more values	
Volume identifier	*MOUNTED	
	+ for more values	
Sequence number	*END	1-16777215, *END
Label	*LIB	
File expiration date	*PERM	Date, *PERM
End of media option	*REWIND	*REWIND, *LEAVE, *UNLOAD
Use optimum block	*YES	*YES, *NO

Additional Parameters

Target release	*CURRENT	*CURRENT, *PRV, V3R2M0...
Update history	*YES	*YES, *NO

More...

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

Figure 174. SAVLIB SAVACT(*LIB): Online backup without disconnect (Part 1 of 3)

3. Specify *YES for Object pre-check, *LIB for Save active, and at least 3600 (or *NOMAX) for Save active wait time as shown in Figure 175.

Save Library (SAVLIB)

Type choices, press Enter.

Clear	*NONE	*NONE, *ALL, *AFTER, *REPLACE
Object pre-check >	*YES	*NO, *YES
Save active >	*LIB	*NO, *LIB, *SYNCLIB, *SYSDFN
Save active wait time >	3600	0-99999, *NOMAX
Save active message queue . . .	*NONE	Name, *NONE, *WRKSTN
Library	*LIBL	Name, *LIBL, *CURLIB
Save access paths >	*YES	*NO, *YES
Save file data	*YES	*YES, *NO
Storage	*KEEP	*KEEP, *FREE
Data compression	*DEV	*DEV, *NO, *YES
Data compaction	*DEV	*DEV, *NO
Libraries to omit	*NONE	Name, generic*, *NONE
+ for more values		

More...

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

Figure 175. SAVLIB SAVACT(*LIB): Online backup without disconnect (Part 2 of 3)

4. Specify *PRINT for Output as shown in Figure 176.

Save Library (SAVLIB)

Type choices, press Enter.

Objects to omit:

Object	*NONE	Name, generic*, *NONE, *ALL
Library	*ALL	Name, generic*, *ALL
Object type	*ALL	*ALL, *ALRTBL, *BNDDIR...
+ for more values		

Output > *PRINT *NONE, *PRINT, *OUTFILE

File to receive output

Library	*LIBL	Name, *LIBL, *CURLIB

Output member options:

Member to receive output . . .	*FIRST	Name, *FIRST
Replace or add records	*REPLACE	*REPLACE, *ADD
Type of output information . . .	*OBJ	*OBJ, *LIB, *MBR, *ERR

Bottom

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

Figure 176. SAVLIB SAVACT(*LIB): Online backup without disconnect (Part 3 of 3)

5. Check the job log and the save listing to be absolutely that all objects could be saved correctly even if the object pre-check was used.

11.5.5.2 SAVR3SYS R3STS(*RUN)

This save option automates the two separate save operation from the previous section (SAV and SAVLIB). Other than that, it basically does the same thing.

You should use *SPTH for IFS files to save, rather than specifying the paths manually. This determines the IFS objects to save from the R/3 table SPTH. Therefore, you should update the table in R/3 with the information from Table 23 on page 233 before you proceed.

To perform the backup operation, follow these steps:

1. Sign on as <SID>OFR.
2. Prompt the SAVR3SYS command and specify the value as shown in Figure 177.

Save R/3 System (SAVR3SYS)

Type choices, press Enter.

System ID	> R01	Character value
Device	> '/qsys.lib/tap01.devd'	
Prompt save commands	*NO	*YES, *NO
R/3 status during save	> *RUN	*DSCDB, *RUN, *END
IFS files to save:		
Path name	*SPTH	
Include or omit		*INCLUDE, *OMIT
+ for more values		
Save R/3 database	*YES	*YES, *NO
Save R/3 configuration	*NO	*YES, *NO
Save R/3 SQL packages	*NO	*YES, *NO
Save reference date	*ALL	Date, *ALL, *LASTSAVE
Save reference time	*ALL	Time, *ALL, *LASTSAVE
Expiration date	*PERM	Date, *PERM
Save active wait time	3600	Number, *NOMAX

More...

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

Figure 177. SAVR3SYS R3STS(*RUN)

11.5.5.3 BRMS using save-while-active

Section 11.5.6, “Saving journal receivers” on page 260, explains how to save journal receivers.

Refer to Figure 169 on page 249 for the link list definition. We recommend that you use similar logic for the Backup Control Group as shown in Figure 178.

Edit Backup Control Group Entries

```

Group . . . . . : SAVR3ON
Default activity . . . . . *BKUPCY
Text . . . . . Online backup of SAP R/3 with DSC, w/o *JRNRCV

```

Type information, press Enter.

Seq	Backup Items	List Type	Weekly Activity SMIWTFS	Retain Object Detail	Save While Active	SWA Message Queue
10	R3IFS	*LNK	*DFTACT	*YES	*NO	
20	R3R01DATA		*DFTACT	*NO	*SYNCLIB	*LIB
30	R3R01400		*DFTACT	*NO	*SYNCLIB	*LIB
40	R346DOPT		*DFTACT	*NO	*NO	
50	R3400		*DFTACT	*NO	*NO	
60	R346DRFC		*DFTACT	*NO	*NO	
70	R346DCPIC		*DFTACT	*NO	*NO	

Bottom

F3=Exit
F11=Display exits
F5=Refresh
F12=Cancel
F10=Change item
F24=More keys

Figure 178. WRKCTLGBRM: BRMS online backup without disconnecting

Remember it does not make sense to specify *YES for Save While Active for the IFS backup.

You should specify *SYNCLIB for:

- R3<SID>DATA database library
- R3<SID>400 library (because the memory-based database monitor is using files in that library)

All the other R/3 libraries can be saved successfully without using save-while-active. The benefit is that the overall backup procedure is faster compared to the one that uses checkpoint processing for all the libraries. That means the tape drive can be used for another save operation, for example if the tape drive is shared between two systems.

Note

BRMS uses the native OS/400 save commands for saving objects. Since the command default for the SAVACTWAIT parameter of the SAVLIB command is set to 120 seconds, you have to change the command default parameter value to at least 3600 seconds. Otherwise, your backup may always fail. There's no way to do this in BRMS itself.

You can use the MONSWABRM command in BRMS to monitor the save-while-active operation. The logic is similar to the one used by the SAVDLTRCV command. For more information, refer to “SAVDLTRCV program” on page 262.

We recommend you use the <SID>OFR user profile to run the scheduled job.

11.5.6 Saving journal receivers

The SAP R/3 application uses journaling of database tables. When a database table is journaled, the system uses a journal receiver to log a record of the changes that occur to each record in the table. Saving the journal receivers provides a method of recovering from a system failure. However, the total size of journal receivers can be quite large. If a single database record is changed many times, the journal receiver has multiple records representing each of the changes.

Here are some considerations when saving journal receivers:

- **Separate tapes for journal receivers:** Use separate backup media for the R/3 database and journal receivers to provide a better protection.
- **Threshold:** The default threshold for R/3 journal receiver is 200,000 KB. This means that if the size of the attached journal receiver exceeds the threshold value and the Manage receivers (MNGRCV) parameter is set to *SYSTEM, the system automatically detaches the receiver, and creates and attaches a new receiver.
- **User ASP overflow:** Use care in ensuring that any user ASP does not exceed the storage allocated. That is because this causes it to overflow into the system ASP and loses the protection and performance benefits of configuring the ASP on separate disks.
- **Do not change MNGRCV(*SYSTEM) and DLTRCV(*NO):** We recommend you do not change these default values with the CHGJRN command. Let the system manage the changing of journal receivers, and do not let the system delete the journal receivers.
- **Keep the backup tapes:** You have to keep the tapes of the journal receiver backup to minimize the amount of data loss in case of emergency. The journal receiver chain must not be broken. You can reuse the tapes after the R/3 database has been saved in a consistent state.
- **Do not save attached journal receivers:** We recommend you do not save attached journal receivers (status ATTACHED) because they are not fully saved. When you have to restore an incompletely saved journal receiver, the status will show PARTIAL. That means you have lost some journal entries and the journal receiver chain is broken. If you want to save the currently attached journal receiver, you can use the command:

```
CHGJRN JRN (R3<SID>DATA/QSQJRN) JRNRCV (*GEN)
```

This command generates a new journal receiver, detaches the current one, and attaches the new one.

11.5.6.1 Manual save

This section explains how to save journal receivers individually. Follow these steps:

1. Enter the command:

```
WRKJRNA JRN (R3<SID>DATA/QSQJRN)
```

Press F15 to display the receiver directory for the journal. The receiver directory tells which journal receivers have not yet been saved as shown in Figure 179.

Work with Receiver Directory

Journal : QSQJRN

Library : R3R01DATA

Total size of receivers (in kilobytes) : 750720

Type options, press Enter.

4=Delete 8=Display attributes

Opt	Receiver	Library	Number	Attach Date	Status	Save Date
	QSQJRN0008	R3R01JRN	00001	10/20/00	SAVED	10/28/00
	QSQJRN0009	R3R01JRN	00002	10/20/00	ONLINE	00/00/00
	QSQJRN0010	R3R01JRN	00003	10/23/00	ONLINE	00/00/00
	QSQJRN0011	R3R01JRN	00004	10/26/00	ATTACHED	00/00/00

Bottom

Parameters or command

====>

F3=Exit F4=Prompt F5=Refresh F9=Retrieve F11=Display size

F12=Cancel

Figure 179. WRKJRNA and F15: Work with Receiver Directory

2. Use the `SAVOBJ` command to save the receivers with the status `ONLINE`. As soon as the save process has completed, the status turns to `SAVED`.

The advantage to using this technique is that each journal receiver is saved only once. You will not have problems with duplicate names and partial receivers if you need to restore. The disadvantage to this technique is that it requires manual effort to determine the names of the journal receivers to be saved.

11.5.6.2 Automatic save using a CL program

You can use a CL program to automate the backup of journal receivers. Use the following program logic:

1. Monitor the journal message queue (by default `QSYSOPR`) for the message (`CPF7020`) that indicates that the system has successfully detached the journal receiver.
2. Your CL program can then save the receiver that was detached to tape using the `SAVOBJ` command and delete it with the `DLTJRNRCV` command.

An alternate method of automatically saving journal receivers is to use the following program logic:

1. Use the Retrieve Journal Information (`QjoRetrieveJournalInformation`) API to determine the journal receiver directory and which receivers are not saved.
2. The program could then save the journal receivers that are not marked as saved.
3. This program could be set up to run on a regular basis or as part of normal processing.

11.5.6.3 Automatic save using the `SAVDLTRCV` command

We recommend you use the backup method that is described in this section. For information about how to install the tools Save and Delete Journal Receivers

(SAVDLTRCV) and Stop Save and Delete Journal Receivers (SAVDLTRCVE), refer to SAP Note 82079.

This backup method may not work together with high availability solutions from business partners because they usually use their own message queue for the journal. In this case, we recommend you use the backup method that was presented in the previous section.

SAVDLTRCV program

This program waits until a receiver becomes detached, resends the message to the *SYSOPR message queue, saves the journal receiver, deletes it afterwards, and removes the message from the message queue.

1. Before using this program, sign on as <SID>OFR and create a message queue:

```
CRTMSGQ MSGQ (R3<SID>400/SAVDLTRCV)
```

2. Change the journal to use the message queue:

```
CHGJRN JRN (R3<SID>DATA/QSQJRN) MSGQ (R3<SID>400/SAVDLTRCV)
```

3. Submit a batch job to run this program:

```
SBMJOB CMD (SAVDLTRCV MSGQ (R3<SID>400/SAVDLTRCV) DEV (TAP01) ) JOB (SAVDLTRCV)
```

Or add these entries to your start profile:

```
_SAVDLTCMD = SAVDLTRCV MSGQ (R3<SID>400/SAVDLTRCV) DEV (TAP01)
Execute_04 = local SBJOB CMD ($(_SAVDLTCMD)) JOB (SAVDLTRCV)
Stop_Program_04 = local SAVDLTRCVE MSGQ (R3<SID>400/SAVDLTRCV)
```

You should save the journal receivers to tape rather than to a save file. This provides better protection in case of a disk failure.

However, if you want to save the journal receivers to a save file, you have to create a R3<SID>SAVF library and use the SAVDLTRCV parameters DEV (*SAVF) SAVF (R3<SID>SAVF/*JRNRCV). The library can then be cleared after the database library is saved successfully.

The SAVDLTRCV command has the parameters as shown in Table 29.

Table 29. SAVDLTRCV parameters

Parameter	Description
FULLMSGQ	Qualified message queue name as passed by command. Recommended value: R3<SID>400/SAVDLTRCV
DEV	Device like TAP01. Special value *SAVF is allowed. Recommended value: Tape device or *SAVF
FULLSAVF	Qualified save file name as passed by command. Important: Use the special value *JRNRCV to make sure that save file contents are not overwritten! Recommended value: R3<SID>SAVF/*JRNRCV
WAITITV	Wait time when receiving messages. If the time expires, the job end status is checked and either the processing is ended or the wait loop is entered again. Recommended value: Any high number of seconds such as 600 (10 minutes)

Parameter	Description
DLTRTYTIM	If the journal receiver cannot be deleted (CPF7024), the job is delayed some time before re trying to delete the journal receiver. Recommended value: Any high number of seconds such as 600 (10 minutes)

SAVDLTRCVE program

This program sends a user message to stop the SAVDLTRCV job. End the backup of the journal receivers that was started with the SAVDLTRCV command by using the command:

```
SAVDLTRCVE MSGQ(R3<SID>400/SAVDLTRCV)
```

11.6 Recovery

Once the system fails or needs to be restored, in case of a software failure or human error, use the information from the backup media to restore the system up to the point of failure to minimize the amount of data loss. The restore options may be selected from the standard Restore menu, which can be accessed by entering the `GO RESTORE` command from an iSeries server command entry line.

In this section, the restore functions produce a database fully synchronized across the entire system. However, the database is only current up to the last full database backup, for example, to the previous night.

Note

For more information on new restore options in OS/400, see Appendix B, "Apply Journaled Changes Extended" on page 545.

11.6.1 User ASP overflow recovery

When the disk units allocated to a user ASP become full, the user ASP is in overflowed status. The system sends message CPI0953 to the QSYSOPR message queue warning you that an ASP is approaching its storage threshold. The system sends message CPI0954 when the storage threshold has been exceeded and the ASP is in overflowed status.

You should reset a user ASP in overflowed status as soon as possible. An overflowed ASP affects system performance. It also makes recovery more difficult and may increase the amount of data lost if a failure occurs.

For information about how to reset an overflowed user ASP, refer to *Backup and Recovery*, SC41-5304.

11.6.2 Restoring the entire system

This section does not explain the entire system restore scenario. You can find a more complete explanation in the manual *Backup and Recovery*, SC41-5304. This section lists the major steps and commands that are needed to accomplish this:

1. Restore or install System Licensed Internal Code from the Alternate IPL device (CD-ROM or tape).

2. Install the operating system from the Alternate IPL device.
3. Restore the user profiles (RSTUSRPRF).
4. Restore the configuration (RSTCFG).
5. Restore all libraries (RSTLIB).
6. Restore the document library objects (RSTDLO).
7. Restore the integrated file system (RST).
8. Apply the journal changes (APYJRNCHG). Refer to 11.6.4.3, "Forward recovery" on page 269.
9. Restore public and private authorization (RSTAUT).

11.6.3 Restoring the SAP R/3 environment

If a save changed object (SAVCHGOBJ) command was used to save information from a library, refer to the *Backup and Recovery Guide*, SC41-5304, to learn how to restore the objects.

Attention

The sequence of restoring the SAP R/3 environment is very important. Make sure the journal receiver library R3<SID>JRN is restored before the R3<SID>DATA database library. If not, a new journal receiver chain will be created in the database library.

Do not restore individual files to the R/3 database. Such an action can result in an inconsistent database. Do this only under the direction of support personnel.

1. Stop SAP R/3 if active.
2. Sign on as QSECOFR.
3. Change your job so the message queue does not wrap when it is full. Use the command:

```
CHGJOB JOBMSGQFL(*PRTWRAP)
```
4. Restore the R/3 kernel library and configuration libraries if necessary.
5. Restore the R/3 IFS structure if necessary.
6. Restore the journal receivers for forward recovery.
7. Delete the R/3 database library using the command:

```
DLTLIB LIB(R3<SID>DATA)
```
8. Restore the R/3 database library with the command:

```
RSTLIB SAVLIB(R3<SID>DATA) DEV(TAP01) MBROPT(*ALL) ALWOBJDIF(*ALL)
```

Check the job log to make sure that all objects of the database library restored successfully.
9. Sign off and sign on as <SID>OFR.
10. Delete R/3 packages with the command:

```
DLTR3PKG SID(<SID>) PKGTYPE(*ALL)
```

11.If the library R3<REL>OPT has been restored, fix the R/3 program owners using the following command:

```
FIXR3OWNS LIB(R3<REL>OPT) OBJ(*ALL)
```

Use this only if the library R3<REL>OPT has been restored. For more information on this command, see SAP Note 173579.

12.Forward recovery using journal receivers. Refer to 11.6.4.3, “Forward recovery” on page 269, for more information.

13.Grant object authority for user <SAP><Inst> using the command:

```
GRTOBJAUT OBJ(R3<SID>DATA) OBJTYPE(*LIB) USER(<SAP><Inst>)
```

This step is only necessary for R/3 releases prior to 4.5A.

14.Start SAP R/3.

11.6.4 Recovering the R/3 database

The reasons to recover the R/3 database are:

- System failure or power failure that caused damage to the database objects
- Disk failure that caused damage to the database objects
- Program failure or human error that caused damage to the content of the database

The journal receivers that contain the changes made to the database can be used to recover the database.

11.6.4.1 Before you begin

Before you begin to recover anything, read this chapter in its entirety. If you are using OS/400 release V5R1 or later, refer to Appendix B, “Apply Journaled Changes Extended” on page 545.

Database consistency

Consider these points in regard to database consistency:

- You must ensure that commitment transaction boundaries are honored on the apply or remove journaled changes operations by using `CMTBDY(*YES)` for any `APYJRNCHG` or `RMVJRNCHG` command. If you don’t ensure this, the R/3 database will end up to be inconsistent!
- You must always recover all the database files in the R/3 library suing the parameter `FILE(R3<SID>DATA/*ALL)`. Otherwise, the R/3 database will be inconsistent!
- If the recovery process fails, it does not terminate at a commitment boundary. The R/3 database is inconsistent in this case.

Recovery restrictions

Some types of entries in the journal receiver cause the apply or remove process to stop. These entries represent events that the system cannot reconstruct.

For example, if journaling was ended for a particular object (journal code F, entry type EJ), the system cannot continue applying journaled changes simply because there is no record of the subsequent changes to that object. For more information, refer to the “Actions by journal code and entry type” table in *Backup and Recovery Guide*, SC41-5304.

When one of these events is encountered, the process ends and a message is sent that indicates the sequence number of the last journal entry that was successfully applied or removed and the reason the process ends. Certain illogical conditions, such as a duplicate key in a database file defined as unique, also cause processing to end.

DB2 UDB for iSeries offers two ways to recover a database to a specific point in time.

Table 30 tells you when to use what recovery procedure.

Table 30. When to use what recovery procedure

Type of outage	Recommended recovery procedure
System failure or power failure	Forward recovery
Disk failure	Forward recovery
Program failure or human error	Backout recovery

In case you lose the database library, backout recovery is not an option. If a program failure or human error forces you to recover the database, consider the following aspects:

- Find out at what point in time the database was last in a consistent state.
- When did you back up the data?
- It may be faster to perform a backout recovery, since you do not have to restore the database.
- It may be better to use forward recovery in case the failure happened recently after the last backup, or in case you simply cannot tell when the failure occurred and you have go back to a certain database level.

11.6.4.2 Backout recovery

Depending on the type of damage to the physical file (damage to the content, not to the object) and the amount of activity since the file was last saved, removing changes from the file can be easier than applying changes to the file. You do not have to restore the R/3 database. Therefore, it is usually the fastest way to reset the R/3 database.

Use the Remove Journalled Changes (`RMVJRNCHG`) command directly or the Work with Journal (`WRKJRN`) command, and follow the prompts to remove (or back out) changes from a file member. The changes are removed in reverse chronological order from the order in which they were originally made to the file.

You can control the changes that are removed from the file. For example, assume that an application updated entries incorrectly for a period of time. In this case, you could remove the changes from the member until that application first opened the member.

To perform the backout recovery, follow these steps:

1. Stop R/3 if active
2. Sign on as QSECOFR.
3. Make sure that no job locks the R/3 database. Use the command:

```
WRKOBJLCK OBJ(R3<SID>DATA) TYPE(*LIB)
```


4. Determine the point in time when the R/3 database was last in a consistent state. For example, let's assume the database was last OK at about 10/31/00, 18:00:00.

5. Restore the necessary journal receivers.

6. Search for F code journal entries in the corresponding receiver chain using the command:

```
DSPJRN JRN(R3R01DATA/QSQJRN) RCVRNG(*CURCHAIN) FROMTIME('10/31/00'
'18:00:00') TOTIME(<current date and time>) JRNCDE((F))
```

If F code journal entries exist in the desired range, look up the entry type in the *Backup and Recovery Guide*, SC41-5304, to find out if the process would end. Depending on the number and types of entries found, it may be more advantageous to use forward recovery.

If there aren't any entries, proceed with the next step.

7. For backout recovery, you have to find out the journal sequence number for the ending point for removing changes that were journaled. Use the following command to list the commitment boundaries in the 30 minutes prior to the failure:

```
DSPJRN JRN(R3R01DATA/QSQJRN) RCVRNG(*CURCHAIN) FROMTIME('10/28/00'
'17:30:00') TOTIME('10/28/00' '18:00:00') JRNCDE((C))
```

This command may show you a screen similar to the example in Figure 180.

```

Display Journal Entries

Journal . . . . . : QSQJRN          Library . . . . . : R3R01DATA

Type options, press Enter.
  5=Display entire entry


```

Opt	Sequence	Code	Type	Object	Library	Job	Time
	5647714	C	SC			WP00	17:57:47
	5647719	C	CM			WP00	17:57:47
	5647720	C	SC			WP01	17:58:02
	5647723	C	CM			WP01	17:58:02
	5647724	C	SC			WP01	17:59:02
	5647727	C	CM			WP01	17:59:02

```

F3=Exit  F12=Cancel

```

Figure 180. DSPJRN: Finding the journal sequence number

8. Note the journal sequence number from the last C code and CM type entry. In this case, the sequence number is 5647727.
9. Prompt the RMVCHGJRN command, and enter the values as shown in the example in Figure 181.

Remove Journalled Changes (RMVJRNCHG)

Type choices, press Enter.

Journal > QSQJRN

Library > R3R01DATA

Journalled file identification:

Journalled physical file . . . > *ALL

Library > R3R01DATA

Member > *FIRST

+ for more values

Range of journal receivers:

Starting journal receiver . . > *CURCHAIN

Library

Ending journal receiver . . .

Library

Starting sequence number *LAST

Ending sequence number > 5647727

More...

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Figure 181. RMVJRNCHG (Part 1 of 2)

10.Specify *YES for Commitment boundary as shown in Figure 182.

Remove Journalled Changes (RMVJRNCHG)

Type choices, press Enter.

Fully qualified job name

User

Number

Commitment boundary > *YES

Name

Name

000000-999999

*NO, *YES

Bottom

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Figure 182. RMVJRNCHG (Part 2 of 2)

11.Sign off and sign on as <SID>OFR.

12.Delete the SQL packages using the command:

DLTR3PKG SID(R01) PKGTYPE(*ALL)

13.Start SAP R/3.

11.6.4.3 Forward recovery

If a file member becomes damaged or is not usable, you can recover the file using the Apply Journalized Changes (APYJRNCHG) command directly or by using the Work Journal (WRKJRN) command and following the prompts. You must first reestablish the physical file member to a condition that you know is undamaged.

The journal receivers may have been deleted or saved with their storage freed since the file was last saved (or from some other point). In this case, you must restore the required journal receivers. Journal receivers do not need to be restored in a particular sequence.

The system applies the changes to the file in the same order as they were originally made. When you use the APYJRNCHG command, the file cannot be in use by anyone else.

Note

You should avoid deleting the QSQJRN journal in the R3<SID>DATA library whenever possible. Otherwise, the restore forces a chain break in the journal receiver. you cannot use the FROMENT(*LASTSAVE) or TOENT(*LASTRST) parameters on the APYJRNCHG command.

To perform the forward recovery, follow these steps:

1. Stop R/3 if active.
2. Sign on as QSECOFR.
3. Make sure that no job locks the R/3 database. Use the command:
4. Determine the point in time when the R/3 database was last in a consistent state. For example, let's assume the database was last OK at about 10/31/00, 18:00:00.
5. Restore the necessary journal receivers.
6. In case you haven't lost the journal QSQJRN in database library, perform the following steps:

- a. Lock the journal object from another job using the command:

```
WRKOBJLCK OBJ(R3<SID>DATA) TYPE(*LIB)
```

- b. Clear the database library using the command:

```
SEMJOB CMD(CLRLIB LIB(R3<SID>DATA)) JOB(CLRR3LIB)
```

The journal object may be deleted with this command if it is not locked.

- c. Display the library to make sure that only the journal remains. Use the command:

```
DSPLIB LIB(R3<SID>DATA)
```

7. Restore the R/3 database from tape using the command:

```
SEMJOB CMD(RSTLIB SAVLIB(R3<SID>DATA) DEV(TAP01) OPTION(*NEW) MBROPT(*ALL) ALWOBJDIF(*ALL)) JOB(RSTR3LIB) JOBSMSGQFL(*PRTWRAP)
```

Check the job log to make sure that all objects (except for the journal) of the database library have been restored successfully.

8. If the journal object has been deleted, use option 9 from the Work with Journals (WRKJRN) command to associate the restored receivers with the journal as shown in Figure 183.

Work with Journals

Type options, press Enter.

2=Forward recovery	3=Backout recovery	5=Display journal status
6=Recover damaged journal	7=Recover damaged journal receivers	
9=Associate receivers with journal		

Opt	Journal	Library	Text
9	QSQJRN	R3R01DATA	

Bottom

Command

====>

F3=Exit F4=Prompt F9=Retrieve F12=Cancel

Figure 183. WRKJRN: Associate receivers with journal

9. Use the WRKJRNA command, and press F15 to check whether all receivers have been added to the journal.
10. Search for F code journal entries in the corresponding receiver chain using the command:

```
DSPJRN JRN(R3R01DATA/QSQJRN) RCVRNG(*CURCHAIN) TOTIME('10/31/00'  
'18:00:00') JRNCD((F))
```

If there are F code entries, look them up in the *Backup and Recovery Guide*, SC41-5304, to check what action the system will take for these entries. If the action is “Ends”, you will need to restart the apply afterwards.

If there aren't any entries, proceed with the next step. You cannot use the RCVRNG(*CURCHAIN) command if the journal has been restored. In this case, you have to specify the starting and ending journal receiver within the same chain.

11. Type the APYJRNCHG command, and enter the values as shown in Figure 184.

Apply Journalled Changes (APYJRNCHG)		
Type choices, press Enter.		
Journal	> QSQJRN	Name
Library	> R3TSTDATA	Name, *LIBL, *CURLIB
Journalled physical file:		
Journalled physical file . . .	> *ALL	Name, *ALL
Library	> R3TSTDATA	Name, *LIBL, *CURLIB
Member	*FIRST	Name, *FIRST, *ALL
+ for more values		
Range of journal receivers:		
Starting journal receiver . .	*LASTSAVE	Name, *LASTSAVE, *CURRENT
Library		Name, *LIBL, *CURLIB
Ending journal receiver . . .		Name, *CURRENT
Library		Name, *LIBL, *CURLIB
Starting sequence number . . .	*LASTSAVE	
Ending sequence number	*LASTRST	
		More...
F3=Exit	F4=Prompt	F5=Refresh
F24=More keys	F12=Cancel	F13=How to use this display

Figure 184. APYJRNCHG (Part 1 of 2)

12.Specify the Ending date and time and *YES for Commitment boundary as shown in Figure 185.

Apply Journalled Changes (APYJRNCHG)		
Type choices, press Enter.		
Ending date and time:		
Ending date	> '10/31/00'	Date
Ending time	> '18:00:00'	Time
Fully qualified job name		Name
User		Name
Number		000000-999999
Fully qualified job name		Name
User		Name
Number		000000-999999
Commitment boundary	> *YES	*NO, *YES
		Bottom
F3=Exit	F4=Prompt	F5=Refresh
F24=More keys	F12=Cancel	F13=How to use this display

Figure 185. APYJRNCHG (Part 2 of 2)

13.In case you lost the journal object, you have to find out the starting and ending journal sequence number for the APYJRNCHG command. Use the last save entries in the journal for the journal starting sequence number. Even with specific ranges, you can specify the *LASTSAVE value.

14.Sign off and sign on as <SID>OFR.

15.Delete the SQL packages using the command:

```
DLTR3PKG SID(R01) PKGTYPE(*ALL)
```

16.Start SAP R/3.

11.7 High availability

For those customers that commit to the technology of R/3, high availability becomes an important and complex subject. It is important because the availability of R/3 is essential to the day-to-day operations of the enterprise. It is complex because the client/server environment of R/3 does not lend itself to the traditional solutions found in the centralized environment of the past.

In the R/3 environment, the focus of a highly available solution must be at the network level (that is, a focus that supports minimum disruption to any active clients while a failing element of the R/3 server environment is replaced with a backup). While the traditional replication of critical data from the primary machine to a backup remains necessary, this by itself is not sufficient.

The switch procedure, which replaces the failing elements with a backup, now comes to the forefront. This procedure determines if the entire R/3 network must come down while the failing component is replaced or whether the clients can remain up. It is also this procedure that determines whether the backup environment is prepared properly for R/3 to run or a manual intervention is required to do so.

11.7.1 Solution discussion

There are three critical components of R/3 that can affect the availability of an R/3 environment. These are:

- The R/3 database
- The central instance
- The critical IFS information

The failure of one of these components can impact the entire R/3 environment.

This is in contrast, for example, to the failure of an instance that is not a central instance. In this case, clients attached to that instance are impacted, but clients attached to other instances are unaffected.

Because of the critical nature of these three components, we advise that you keep them on the same iSeries server. The database and critical IFS information are automatically together. There is, however, no restriction that limits the location of the central instance. It may be located on a machine separate from the one that holds the database. If this is done, the R/3 solution for high availability becomes more complex.

There are now two points of failure within the R/3 network: the machine that holds the database and the machine that is running the central instance. Provisions must be made to recover from failure in either of these two nodes if R/3 high availability is to be maintained.

11.7.2 Business partner solutions

This section contains information about high availability solutions for the iSeries server SAP customers from business partners. In particular, we focus on Lakeview Technology, Vision Solutions, and DataMirror Corporation.

High availability middleware is the name given to the group of applications that provide the replication and management between iSeries servers. More recently they also provide the Cluster management middleware.

The High Availability Business Partners (HABPs) provide data resiliency tools. With OS/400 V4R4, they are heading into application resiliency. You can read more about their solutions in the following manuals and at their Web sites:

- MIMIX:

MIMIX for SAP R/3 by Lakeview Technology, Oakbrook, IL 60521, USA
<http://www.lakeviewtechnology.com>

- Vision

Vision Suite by Vision Solutions, Inc., Irvine, CA 92612, USA
<http://www.vision-solutions.com>

- DataMirror

DataMirror High Availability Suite by DataMirror Corporation, Markham, Ontario, Canada
<http://www.datamirror.com>

11.7.3 Concepts

This section discusses some high availability concepts, including LPAR, backup systems, remote journaling, and clustering.

11.7.3.1 Logical partitioning (LPAR)

Logical partitioning is the ability to make a single multiprocessing iSeries server run as if it were two or more independent systems (introduced in V4R4M0). Each logically partitioned system has one primary partition and one or more secondary partitions. Each logical partition represents a division of resources in your iSeries server. Each partition is logical because the division of resources is virtual, not physical.

The primary resources in your system are its processors, memory (main storage), I/O buses, and IOPs. Each logical partition operates as an independent logical system. However, each partition shares a few physical system attributes such as the system serial number, system model, and processor feature code. All other system attributes may vary among partitions. For example, each partition has dedicated hardware such as processors, main storage, and I/O devices. Therefore, logical partitions have some hardware fault tolerance if configured properly.

For more information, refer to Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43.

11.7.3.2 Backup system

The R/3 production system can be replicated to another machine for high availability reasons. Since the “replicated” machine does not need the same

resources (memory, disks, CPUs, and so on) as the production machine all the time, it can be configured into two LPARs, with another LPAR used for testing, for example. Performance improvements may be realized if certain batch functions (queries, reports, backups) are performed on the replicated system instead of the production system.

Figure 186 shows an example of R/3 and LPARs with replication of the production system.




Figure 186. SAP system PRD replicated to LPAR0 on a different iSeries server

11.7.3.3 Remote journaling

The high availability solutions, developed prior to V4R2M0, used local journaling and the Receive Journal Entry (RCVJRNE) command as shown in Figure 187. In a traditional environment, a user's applications that run on a source (production) system generate database changes. In turn, these changes create journal entries written to a local journal receiver. Still on the source side, the entries are received from the journal and buffered in a communications staging area.

The data is transmitted asynchronously to the target system using existing communications gear. A high availability application running on the target system receives the journal entries into a temporary storage location, usually a user space. Another job, or many jobs, replays the changes into a copy of the source database. At this point, you have an exact copy of the production database on the target machine.




Figure 187. Without remote journaling

The remote journal function provides a much more efficient transport of journal entries than the traditional approach. In the scenario shown in Figure 188, when a user application makes changes to a database file, there is no need to buffer the resulting journal entries to a staging area on the production (source) system. Efficient low-level system code is used instead to capture and transmit journal entries directly from the source system to associated journals and journal receivers on a target system. Much of the processing is done below the machine interface (MI). Therefore, more CPU cycles are available on the production machine for other important tasks.

Because the remote journal function, if activated in synchronous mode, replicates journal entries to the backup machine's main storage before updating the production's system database, the data latency is driven to zero. The high availability solutions available on the iSeries server can fully take advantage of this more efficient transport mechanism. In fact, a high availability solution is still necessary in most customer environments to apply the application-dependent data to a replica database for hot-backup scenarios. It is also necessary for providing the required management facilities for these hot-backup environments.




Figure 188. With remote journaling

When the remote journal function is activated on the source machine, the system replicates existing journal entries first as quickly as possible. This is referred to as *catch-up mode*. Once the specified journal receivers are transmitted to the target machine, the source starts continuously sending new entries either synchronously or asynchronously. The mode of operation depends on what was specified when the remote journal function was activated. The different delivery modes are discussed in the following sections.

You can consider the journal receivers on the target machine as a replica of the production machine's journal receivers. It is as if you saved the production machine's journal receivers and restored them on the target machine. The time stamps, system name, and qualified journal receiver names in the associated remote journal's journal entries are exactly the same as in the local journal's journal entries on the source system. In addition, the attach and detach times of the journal receivers are the same. However, while using remote journal support, you may see a minimal discrepancy in size between the local receiver and the associated remote receiver. Note that you are using two different machines, which pre-allocate space for the journal receivers in different operating system environments. This may result in slightly different sizes, but the data in the associated receivers is always the same.

- Synchronous mode:** In synchronous mode, journal entries are replicated to the main memory on the remote machine first. After an arrival confirmation is returned to the source machine, the journal entry is deposited to the local receiver. Next, the actual database update is made, if appropriate. The target system is updated in real-time with all of the journal entries as they are generated by a user application on the source system. Synchronous journaling allows for recovery that loses no journal entries on the target system if an outage is experienced on the source system. Sending journal

entries synchronously to a target system modestly impacts the journaling throughput on the source system.

- **Asynchronous mode:** Sending journal entries asynchronously means that the journal entry is sent to the target system at some time after control is returned to the end user application that deposited the journal entry on the source system. From a recovery standpoint, asynchronous mode is less desirable. The reason is that the source system may have journal entries ahead of those journal entries that are known to the target system. Using this method allows for recovery that may lose a number of journal entries given a failure on the source system. It should have minimal impact to the local system when compared to the synchronous delivery mode.

The business partner solutions for high availability are in the process of taking advantage of remote journaling rather than using the traditional approach.

For more information, refer to the redbook *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189.

11.7.3.4 Clustering

This offers a continuous availability solution. This solution, called *OS/400 Cluster Resource Services* (introduced in V4R4M0), provides failover and switchover capabilities for your systems that are used as database servers or application servers in a client/server environment. If a system outage or a site loss occurs, the functions that are provided on a clustered database server system can be switched over to one or more designated backup systems that contain a current copy (replica) of your critical application data. The failover can be automatic if a system failure should happen, or you can control how and when the transfer will take place by manually initiating a switch over.

The systems in a cluster can be physically connected via a LAN or a high-speed OptiConnect bus, or they can reside in different locations and communicate over telephone lines. IBM and iSeries server Cluster Middleware Business Partners have teamed together to provide state-of-the-art OS/400 cluster resource services functions along with a GUI for cluster management.

Figure 189 shows a basic cluster. There are four nodes systems A through D. The nodes are connected through a network. Systems A, B, and C are local to each other, and system D is at a remote location. The cluster management tool controls the cluster from anywhere in the network. End users work on servers in the cluster without knowing or caring where their applications are running.




Figure 189. Basic cluster

In the event of a failure, Cluster Resource Services (CRS), which is running on all systems, provides a switch over. This switch will cause minimal impact to the end user or applications that are running on a server system. Data requests are automatically rerouted to the new primary system. You can easily maintain multiple data replicates of the same data. Clusters contain more than two nodes. You can group a system's resilient data (replicated data) together to allow different systems to act as the backups for each group's resilient data. Multiple backup systems are supported. If a system fails, cluster resource services provides the means to reintroduce (rejoin) systems to the cluster and restore their operational capabilities.

Hardware and software requirements for clusters

Any iSeries server model that can run OS/400 Version 4 Release 4 or later is compatible for implementing clustering. You must have OS/400 V4R4 or later installed and TCP/IP configured on the iSeries server before you can implement clustering. In addition, you can purchase a cluster management package from a high availability business partner (HABP) that will provide the required replication functions and cluster management capabilities.

ClusterProven

ClusterProven is a new IBM brand initiative. It gives application providers the ability to apply a logo their product to say it conforms to a certain level of resilience on a particular product.

This branding enables customers to select application software appropriate to their availability needs. We provide the ClusterProven process so that customers are aware of the level of resilience their ISVs have attained and how this level was achieved.

Note

Until now, SAP R/3 had not been ClusterProven.

For more information, refer to the redbook *AS/400 Clusters: A Guide to Achieving Higher Availability*, SG24-5194.

11.7.4 Minimizing backup and recovery windows

You can minimize the backup and recovery window by using logical partitioning. For more information, refer to 3.4.6.2, “Minimizing backup and recovery windows” on page 51.

Part 3. Advanced topics

This part covers topics that will interest those of you who want to enhance your SAP R/3 installation by improving performance and adding additional functionality. In this part, you'll find these chapters:

- Chapter 12, "Coexistence with non-SAP R/3 applications" on page 283, describes techniques to enable SAP R/3 running on the iSeries to interact with iSeries non-SAP R/3 applications.
- Chapter 13, "MQSeries link for R/3" on page 323, presents a brief overview of MQSeries, the MQSeries link for SAP R/3, and Application Link Enabling (ALE). This chapter is for SAP consultants who are installing R/3 systems for customers. It is also intended for sales staff who are proposing R/3, where connection to external systems (that is, non-SAP systems) is important.
- Chapter 14, "Migration to another platform" on page 343, provides a brief overview of the concepts and processes involved in the migration of an existing SAP R/3 system to a different hardware platform, operating system, and database.
- Chapter 15, "Change and transport system" on page 349, describes the Change and Transport System, which allows you to organize your customizing and development project when implementing SAP and transport changes from one SAP R/3 system to another.
- Chapter 16, "Performance management" on page 365, introduces you to a performance management methodology for optimally running the SAP R/3 application on the iSeries. It shows you how to review the performance of an iSeries server running SAP R/3 applications.
- Chapter 17, "Access Builder for SAP R/3" on page 479, describes the Access Builder for SAP R/3, which is a toolkit that creates Java applications, applets, and JavaBeans capable of accessing an SAP system.
- Chapter 18, "mySAP.com" on page 483, covers mySAP.com, SAP's Internet strategy, which can be defined as "an open application environment for e-commerce". It consists of portals, industry solutions, and Internet applications and services, with the aim towards integrated business collaboration.
- Chapter 19, "SAP R/3 and Domino connection" on page 491, covers the integration technologies that can be used to integrate Domino and SAP applications.
- Chapter 20, "Using an IBM Network Station as an SAP front end" on page 497, shows you how to integrate a network station in your R/3 environment to start an SAP GUI.
- Chapter 21, "Problem determination" on page 503, intends to help you diagnose and manage problems that you may encounter in running an SAP R/3 system on the iSeries server.

Chapter 12. Coexistence with non-SAP R/3 applications

This chapter describes techniques to enable SAP R/3 running on the iSeries server to interact with iSeries non-SAP R/3 applications (for example RPG programs). The objective is to:

- Exchange data
- Call or trigger programs and system functions

The following examples are provided based on simple scenarios, using programs written in ABAP and RPG/400. These examples demonstrate:

- Accessing SAP R/3 data using an RPG/400 program
- Using an ABAP program to access non-SAP R/3 data
- ABAP applications calling OS/400 commands
- RPG/400 applications triggering ABAP programs
- Common Programming Interface-Communication (CPI-C) communication interfaces between an ABAP program and RPG/400

Important notice

All program examples were created and tested with OS/400 V4R5 and SAP R/3 Release 4.6C, with SAP R/3 kernel release 4.6D.

All program source code (ABAP, RPG/400, and CL) contained in this chapter are for demonstration purposes only. They are not intended for use in any production system. We recommend that a suitably qualified programmer use this program code as a basis for creating the required programs according to your specific requirements. Program names and library names should be adjusted to fit your particular environment.

12.1 Considerations

SAP R/3 on the iSeries server operates in its own environment. All SAP R/3 objects are stored in specific libraries on the iSeries, and SAP R/3 jobs run in specific subsystems, using their own work management objects (for example, subsystems and job descriptions). SAP R/3 can run concurrently with other applications on the same iSeries server or on different servers connected together.

Figure 190 shows an integrated solution. The SAP R/3 system ("R01", instance "00") and non-SAP R/3 application run concurrently on the same iSeries server. All SAP R/3 work processes run in the R3_00 subsystem. The non-SAP R/3 application uses the functions that are provided by the service-program LIBRFC for communication to the SAP system. It runs in the default iSeries subsystem environment (QINTER and QBATCH). Both the SAP R/3 system and the non-SAP R/3 application use the same EBCDIC code page.

The user can access both applications by using parallel sessions from the front-end PC:

- SAP GUI connects to SAP R/3 instance "00".
- 5250 emulation (for example, Client Access or Telnet) connects to the subsystem QINTER.




Figure 190. SAP R/3 and other applications

When planning to interface non-SAP R/3 applications with SAP R/3 or exchange data between these systems, consider the following points:

- SAP R/3 applications are developed in ABAP. Programs written in this language can only run in an SAP R/3 environment. SAP R/3 applications do not run directly in an OS/400 work management environment. They are dispatched to an instance work process (batch or dialog). You cannot call an ABAP program directly by using the OS/400 call level interface. ABAP programs do not support this interface either. SAP R/3 provides its own interfaces and tools for communication outside of the SAP R/3 environment. These include:

- Operating system command call interface
- Event handler
- CPI-C interface
- RFC interface
- SAP R/3 defines and maintains its data structures using the ABAP Dictionary interface. For example, tables and views are defined and activated in this environment. Physically, they are maintained as DB2 files within the assigned iSeries database library, using the ANSI SQL industry standard database interface of DB2 UDB for iSeries. If you are familiar with the SAP R/3 table structures, you can access ABAP Dictionary transparent tables from outside SAP R/3. However, this is not recommended since table structures may change in different SAP R/3 releases. Also, the ABAP language provides functions to access tables and stream files outside of SAP R/3. These functions include:
 - The EXEC-SQL interface to execute SQL host commands.
 - Commands for reading from and writing to files and stream files.

12.2 Example programs

The following examples show programs written in RPG/400 and ABAP. These programs read an externally described DB2 UDB for iSeries file and an ABAP Dictionary table and display the contents of these files.

12.2.1 RPG/400 example

The RPG program T8189RP1 shown in Figure 191 reads the externally-described physical file T8189DB and prints the file contents. The data definition (DDS) for file T8189DB is shown in Figure 192. Both objects are stored in the RFC library.

```

FMT FX  ....FFilenameIPEAF.....L..I.....Device+.....KExit++Entry+A.
0001.00      FT8189DB IF  E                      DISK
0002.00      FQSYSPRT O   F      132             PRINTER
0003.00      C              *IN20      DOUEQ'1'
0004.00      C              READ Z8189DB                      20
0005.00      C              *IN20      IFEQ '0'
0006.00      C              EXCPTDETAIL
0007.00      C              ENDIF
0008.00      C              ENDDO
0009.00      C              MOVE '1'          *INLR
0010.00      OQSYSPRT E 11          DETAIL
0011.00      O              ITEMNM
0012.00      O              ITEMMD

```

Figure 191. RPG program T8189RP1

```

5716SS1 V4R5M0 000526      Display File Field Description
Input parameters
  File . . . . . : T8189DB
  Library . . . . . : RFC
File Information
  File . . . . . : T8189DB
  Library . . . . . : RFC
  File location . . . . . : *LCL
  Externally described . . . . . : Yes
  Number of record formats . . . . . : 1
  Type of file . . . . . : Physical
  File creation date . . . . . : 11/16/00
Record Format Information
  Record format . . . . . : Z8189DB
  Format level identifier . . . . . : 272B1EE291453
  Number of fields . . . . . : 2
  Record length . . . . . : 30
Field Level Information
  Data      Field  Buffer  Buffer      Field  Column
Field      Type      Length Length Position      Usage  Heading
ITEMNM     CHAR        5      5         1      Both  ITEMNM
  Coded Character Set Identifier . . . . . : 37
ITEMD      CHAR       25     25         6      Both  ITEMMD
  Coded Character Set Identifier . . . . . : 37

```

Figure 192. File T8189DB DDS

12.2.2 ABAP example

In this example, the ABAP program Z8189AB1 (Figure 193) reads the SAP R/3 table Z8189DB sequentially and prints the file contents.

```

REPORT Z8189AB1.
TABLES: Z8189DB.
SELECT * FROM Z8189DB.
WRITE: / Z8189DB-ZITEMNM, Z8189DB-ZITEMD.
ENDSELECT.

```

Figure 193. ABAP program Z8189AB1

Table Z8189DB is defined as a transparent table in the ABAP Dictionary. A transparent table is stored in the assigned iSeries library as a physical file, using the same file name, record name, field names, and field attributes as defined in the ABAP Dictionary. Transparent tables can be accessed directly from the iSeries library without using an ABAP program. Figure 194 shows the ABAP Dictionary table Z8189DB.




Figure 194. ABAP Dictionary table Z8189DB

ABAP program Z8189AB1 reads the dictionary table Z8189DB sequentially, using the SAP-SQL SELECT statement. SAP-SQL is a set of commands similar to ANSI SQL (SELECT, INSERT, UPDATE, MODIFY, DELETE, COMMIT WORK, ROLLBACK WORK), which is integrated into the ABAP command set. These commands are started directly by ABAP and can only be used to access ABAP Dictionary-defined tables and views. The resulting output of ABAP program Z8189AB1 is shown in Figure 195.




Figure 195. Output of ABAP program Z8189AB1

12.3 Accessing SAP R/3 data using an RPG/400 program

The ABAP Dictionary table Z8189DB is defined as a transparent table and stored as a physical file in the iSeries library R3R01DATA. This is the assigned database library for the SAP R/3 installation (R01) used in this example.

Where Figure 194 shows the ABAP Dictionary table Z8189DB, Figure 196 shows the file layout of the associated database file Z8189DB.

```

5716SS1 V4R5M0 000526      Display File Field Description
Input parameters
  File . . . . . : Z8189DB
  Library . . . . . : R3R01DATA
File Information
  File . . . . . : Z8189DB
  Library . . . . . : R3R01DATA
  File location . . . . . : *LCL
  Externally described . . . . . : Yes
  Number of record formats . . . . . : 1
  Type of file . . . . . : Physical
  SQL file type . . . . . : TABLE
  File creation date . . . . . : 11/17/00
Record Format Information
  Record format . . . . . : Z8189DB
  Format level identifier . . . . . : 29B6A5934104D
  Number of fields . . . . . : 2
  Record length . . . . . : 30
Field Level Information
      Data      Field Buffer      Buffer      Field      Column
Field   Type    Length Length  Position    Usage    Heading
ZITEMNM  CHAR         5      5         1      Both    ZITEMNM
  Default value . . . . . :
  ' '
  Coded Character Set Identifier . . . . . : 500
ZITEMD   CHAR        25     25         6      Both    ZITEMD
  Default value . . . . . :
  ' '
  Coded Character Set Identifier . . . . . : 500

```

Figure 196. DDS for file Z8189DB

File T8189DB (Figure 192 on page 286) and file Z8189DB (SAP R/3 described) are completely compatible. Therefore, it is possible to print the contents of file Z8189DB using the RPG program defined in Figure 191 on page 285. To do this, use the OS/400 `OVRDBF` command before running the RPG program, for example:

```

OVRDBF FILE(T8189DB) TOFILE(R3R01DATA/Z8189DB) LVLCHK(*NO)
CALL RFC/T8189RP1

```

12.4 Accessing non-SAP R/3 data using ABAP programs

This section describes how to use ABAP programs to access non-SAP R/3 data. This can be achieved by using:

- SAP-SQL in the ABAP program to access a logical file in the SAP R/3 database library: The logical file is based on a non-SAP R/3 physical file
- The EXEC-SQL interface
- Sequential files (stream files and program-described physical files)

SAP-SQL

SAP-SQL consists of a set of SQL commands similar to ANSI SQL. These commands are run by ABAP. SAP-SQL can only access ABAP Dictionary objects.

The ABAP program Z8189AB1 (Figure 193 on page 286) shows how to select data using SAP-SQL.

EXEC-SQL

When you use the EXEC-SQL interface, all SQL commands are passed directly to the DB2 UDB for iSeries database manager to be run. You may embed all functions supported by SQL/400 and work with all DB2 UDB for iSeries database objects to which you have access.

Figure 197 shows ABAP program Z8189SQL selecting data from outside the SAP R/3 database (physical file T8189DB in library RFC) using the EXEC-SQL statement. The sub-routine OUTPUT-ITEM is performed for each record retrieved from the T8189DB file.

```
REPORT Z8189SQL.
DATA: BEGIN OF REC,
      RITEMNM(5) TYPE C,
      RITEMD(25) TYPE C,
      END OF REC.
EXEC SQL PERFORMING OUTPUT-ITEM.
      SELECT ITEMNM, ITEMND
            INTO :REC
            FROM RFC/T8189DB
ENDEXEC.
FORM OUTPUT-ITEM.
      WRITE: / REC-RITEMNM, REC-RITEMD.
ENDFORM.
```

Figure 197. ABAP program Z8189SQL

12.4.1 Accessing non-SAP R/3 data through the ABAP Dictionary

To access non-SAP R/3 data using the SAP R/3 data dictionary, you need to replace the ABAP Dictionary table by a DB2 UDB for iSeries logical view. The following example explains how to do this:

1. Create the ABAP Dictionary table Z8189TAB as shown in Figure 198.




Figure 198. ABAP Dictionary table Z8189TAB

2. Ensure that the table structure is identical to physical file RFC/T8189DATA (Figure 199).

```

5716SS1 V4R5M0 000526      Display File Field Description
Input parameters
  File . . . . . : T8189DATA
  Library . . . . . : RFC
File Information
  File . . . . . : T8189DATA
  Library . . . . . : RFC
  File location . . . . . : *LCL
  Externally described . . . . . : Yes
  Number of record formats . . . . . : 1
  Type of file . . . . . : Physical
  File creation date . . . . . : 11/17/00
Record Format Information
  Record format . . . . . : Z8189TAB
  Format level identifier . . . . . : 27233EE291472
  Number of fields . . . . . : 2
  Record length . . . . . : 30
Field Level Information
  Data      Field Buffer      Buffer      Field      Column
  Field     Type      Length Length Position Usage      Heading
ITEMNM     CHAR         5         5         1      Both      ITEMNM
  Coded Character Set Identifier . . . . . : 37
ITEMD     CHAR        25        25         6      Both      ITEMD
  Coded Character Set Identifier . . . . . : 37

```

Figure 199. DDS for file T8189DATA

3. SAP R/3 creates the transparent table (physical file) Z8189TAB in library R3R01DATA.

4. Delete this transparent table using the following OS/400 command:

```
DLTF R3R01DATA/Z8189TAB
```

5. Define the logical file Z8189TAB as shown in Figure 200, and create this logical file in the SAP R/3 database library, using the following OS/400 command:

```
CRTLFILE FILE(R3R01DATA/Z8189TAB) SRCFILE(RFC/QSOURCE) SRCMBR(Z8189TAB)
```

```
Columns . . . :   1  71          Edit          QGPL/QTXTSRC
SEU==>                                Z8189TAB
FMT **   ...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7
          ***** Beginning of data *****
0001.00          R 78189TAB          PFILE(RFC/T8189DATA)
0002.00          K ITEMNM
          ***** End of data *****
```

Figure 200. Logical file Z8189TAB

12.4.2 Accessing iSeries stream files using ABAP programs

ABAP supports reading and writing to iSeries stream files (object type *STMF). iSeries stream files can be accessed in either text mode or binary mode:

- **Text mode:** Data is passed using record delimiters.
- **Binary mode:** Data is passed as a byte stream without any structure.

12.4.2.1 ABAP program writing to iSeries stream file

The ABAP program shown in Figure 201 reads data from the ABAP Dictionary table Z8189DB (using an EXEC-SQL interface). It also writes data to an existing stream file, t8189seq, located in the iSeries IFS directory /t8189. The stream file t8189seq is opened in text mode.

```
REPORT Z8189AB2.
DATA: FILE(20) VALUE '/t8189/t8189seq'.
DATA: BEGIN OF REC,
      RITEMNM(5) TYPE C, RITEMD(25) TYPE C,
      END OF REC.
OPEN DATASET FILE FOR OUTPUT IN TEXT MODE.
EXEC SQL PERFORMING OUTPUT-ITEM.
      SELECT ITEMNM, ITEMMD
      INTO :REC
      FROM RFC/T8189DB
ENDEXEC.
FORM OUTPUT-ITEM.
      TRANSFER REC TO FILE.
ENDFORM.
```

Figure 201. ABAP program Z8189AB2

12.4.2.2 ABAP program reading iSeries stream file

The ABAP program shown in Figure 202 reads data sequentially from the existing stream file, t8189seq, located in the iSeries IFS directory /t8189 and prints the file contents. The stream file T8189SEQ is opened in text mode.

```
REPORT Z8189AB3.  
DATA: FILE(20) VALUE '/t8189/t8189seq'.  
DATA: BEGIN OF REC,  
       RITEMNM(5) TYPE C, RITEMD(25) TYPE C,  
       END OF REC.  
OPEN DATASET FILE FOR INPUT IN TEXT MODE.  
DO.  
  READ DATASET FILE INTO REC.  
  IF SY-SUBRC NE 0. EXIT. ENDIF.  
  WRITE: / REC-RITEMNM, REC-RITEMD.  
ENDDO.  
CLOSE DATASET FILE.
```

Figure 202. ABAP program Z8189AB3

12.4.2.3 ABAP program reading iSeries physical files

The ABAP program shown in Figure 203 reads data from an iSeries physical file and prints the file contents.

```
REPORT Z8189AB4.  
DATA: FILE(60) VALUE '/qsys.lib/rfc.lib/t8189db.file/t8189db.mbr'.  
DATA: BEGIN OF REC,  
       RITEMNM(5) TYPE C, RITEMD(25) TYPE C,  
       END OF REC.  
OPEN DATASET FILE FOR INPUT.  
DO.  
  READ DATASET FILE INTO REC.  
  IF SY-SUBRC NE 0. EXIT. ENDIF.  
  WRITE: / REC-RITEMNM, REC-RITEMD.  
ENDDO.  
CLOSE DATASET FILE.
```

Figure 203. ABAP program Z8189AB4

Program Z8189AB4 is similar to program Z8189AB3 (Figure 202). The only difference is that the file used in this example is an iSeries physical file. The physical file T8189DB is accessed using the IFS directory structure.

12.5 SAP R/3 processing OS/400 commands

This section shows how to process OS/400 commands from inside SAP R/3, using the ABAP system command interface. OS/400 commands can also be run directly by using pre-defined SAP R/3 external commands (SAP R/3 transaction SM49). Or it can be integrated directly into SAP R/3 background jobs (SAP R/3 transaction SM36). Additional OS/400 commands can be added using the SAP

R/3 transaction SM69. Refer to 12.5.3, “Working with OS/400 commands” on page 297, for more information on processing OS/400 commands.

12.5.1 ABAP system command interface

The following example shows how to invoke an OS/400 command from an ABAP program, using the ABAP system command interface.

The ABAP program Z8189AB5, shown in Figure 204, runs an OS/400 command that is passed to it as an input parameter.

```
REPORT Z8189AB5.
PARAMETERS:
  CMD(250) DEFAULT 'CALL PGM(RFC/T8189RP1)'.
DATA: BEGIN OF LISTE OCCURS 100,
      ZEILE(132),
      END OF LISTE.
CALL 'SYSTEM' ID 'COMMAND'
      FIELD CMD
      ID 'TAB' FIELD LISTE-*SYS*.
LOOP AT LISTE.
WRITE: / LISTE-ZEILE.
ENDLOOP.
```

Figure 204. ABAP program Z8189AB5

The input parameter (CMD) has a pre-defined default value (DEFAULT). When you run the program, you may accept the default input string or override it with another valid OS/400 command string. In the above example, the default parameter will run the RPG/400 program T8189RP1. See Figure 205.

The specified command is executed by calling the SYSTEM function, which in turn, invokes the OS/400 system program QCMDXEC. QCMDXEC executes the OS/400 command received from the ABAP program.

After the OS/400 command is executed, control is passed back to the ABAP program and processing continues based on the command output specifications. In the above example, the result of the OS/400 command is written to an internal table (LISTE) and displayed on the user screen.




Figure 205. ABAP program Z8189AB5: Parameter selection

Figure 206 shows the output of the ABAP program Z8189AB5.




Figure 206. ABAP program Z8189AB5 output

12.5.2 ABAP program processing CL streams

This section briefly explains how to use an ABAP program to process OS/400 commands stored in a source physical file member.

Figure 207 shows the entries created in member Z8189CL of the source physical file RFC/QCLSRC.

```

0001.00 DSPLIBL  *PRINT
0002.00 CRTLIB   Z8189LIB      *TEST TEXT('Z8189 Test Library')
0003.00 CRTSRCPF Z8189LIB/QCLSRC TEXT('Z8189 Test SRCPF')

```

Figure 207. Source physical file member Z8189CL

The ABAP program Z8189AB6, shown in Figure 208, reads this source file member (Z8189CL) sequentially and executes each OS/400 command contained in this file.

```

REPORT Z8189AB6.
DATA: FILE(60) VALUE '/qsys.lib/rfc.lib/qclsrc.file/z8189cl.mbr'.
DATA: BEGIN OF REC,
      CL(80) TYPE C,
      END OF REC.
DATA: BEGIN OF LISTE OCCURS 100,
      ZEILE(132),
      END OF LISTE.
OPEN DATASET FILE FOR INPUT.
DO.
  READ DATASET FILE INTO REC.
  IF SY-SUBRC NE 0. EXIT. ENDIF.
  WRITE: / REC-CL.
  CALL 'SYSTEM' ID 'COMMAND'
      FIELD REC-CL
      ID 'TAB' FIELD LISTE-*SYS*.
ENDDO.
LOOP AT LISTE.
WRITE: / LISTE-ZEILE.
ENDLOOP.
CLOSE DATASET FILE.

```

Figure 208. ABAP program Z8189AB6

The above ABAP program reads the source physical file member (Z8189CL) as input. Then it creates the iSeries library T8189LIB and the source physical file T8189LIB/QCLSRC. The output result of ABAP program Z8189AB6 is shown in Figure 209.

The screenshot shows the SAP R/3 Redbook: Z8189AB6 output. The window title is "IBM: SAP R/3 Redbook: Z8189AB6". The output text is as follows:

```


IBM: SAP R/3 Redbook: Z8189AB6 1
DSPLIBL *PRINT
CRTLIB Z8189LIB *TEST TEXT('Z8189 Test Library')
CRTSRC PF Z8189LIB/QCLSRC TEXT('Z8189 Test SRC PF')
5769SS1 V4R5M0 000526 Library List
Library Type Text Description##1
R3SYS SYS R/3 library for QSYS objects
TSYS SYS TOOLS product command library
QSYS SYS System Library##m
QSYS2 SYS System Library for CPI's##a
QHLPSYS SYS
QUSRSYS SYS
TUSRSYS SYS TOOLS product user object library
R3R01DATA CUR COLLECTION - created by SQL##r
R346D0PT USR
QTEMP USR
Q6PL USR
Q6Y USR
***** END OF LISTING *****

```

Figure 209. ABAP program Z8189AB6 output

12.5.3 Working with OS/400 commands

OS/400 commands can be executed directly by using pre-defined SAP R/3 external commands (SAP R/3 transaction SM49), or they can be integrated directly into SAP R/3 background jobs (SAP R/3 transaction SM36). Section 12.5.4, “External command interface for SAP R/3 background jobs” on page 299, provides more detail on using the background job command interface. Additional OS/400 commands can be created using the SAP R/3 transaction SM69, as shown in Figure 210.



The screenshot shows the SAP SM69 transaction window titled "External Operating System Commands". The window has a menu bar (Command, Goto, System, Help) and a toolbar. The main area contains the following fields:


Command name	Operating system	Type
ZDSPLIBL	OS/400	Customer

Below the table, there is a text field for "Operating system command" containing "dsplbl" and a text field for "Parameters for operating system command" which is empty.

At the bottom, there are two sets of fields for "Created by" and "Last changed by", both showing "MONTI_A" and "17.11.2000 18:36:52". There are also checkboxes for "Additional parameters allowed" (checked) and "Trace" (unchecked), and a "Check module" button.

Figure 210. SM69: Creating the OS/400 command

Use the SAP R/3 transaction SM49 to execute a predefined OS/400 command. Figure 211 shows the input parameters for this transaction. Note that transaction SM49 allows you to execute the OS/400 command on another connected host system.



The screenshot shows the SAP SM49 transaction window titled "Execute External Operating System Command". The window has a menu bar (Command, Goto, System, Help) and a toolbar. The main area contains the following fields:

Command name	Operating system	Type
ZDSPLIBL	OS/400	Customer

Below the table, there is a text field for "Operating system command" containing "dsplbl" and a text field for "Parameters for operating system command" which is empty.

At the bottom, there is a section titled "Execute" with a sub-section "Additional parameters" containing an empty text field. Below that is a "Destination host" field containing "AS23".

Figure 211. SM49: Executing the OS/400 command

The command results are shown in Figure 212.




Figure 212. SM49: Output results

12.5.3.1 SAP R/3 function modules for external commands

SAP also provides function modules (based on predefined OS/400 commands) for checking and executing external commands using a CPI-C interface. These include:

- SXPG_CALL_SYSTEM
- SXPG_COMMAND_EXECUTE
- SXPG_COMMAND_CHECK
- SXPG_COMMAND_LIST_GET
- SXPG_COMMAND_DEFINITION_GET
- SXPG_DUMMY_COMMAND_CHECK

These function modules can be used in ABAP programs to execute pre-defined OS/400 commands. Refer to the SAP online documentation for detailed information on the above function modules.

12.5.4 External command interface for SAP R/3 background jobs

The SAP R/3 transaction SM36 is used to define background jobs in the SAP R/3 system. The following example shows how iSeries commands can be executed from an SAP R/3 background job.

Figure 213 shows how to define the job name and assign the job priority. Note that you may select to run the background job on a specific server of the SAP R/3 system. Next, you define the job steps to be executed by the background job.




Figure 213. SM36: Defining an SAP R/3 background job

Figure 214 shows how to specify a pre-defined OS/400 external command as a job step. External commands are created using the SAP R/3 transaction SM69. Refer to 12.5.3, “Working with OS/400 commands” on page 297, for more information on external commands.

Create Step 0

User: MONTI_A

Program values

ABAP program External command External program

ABAP program

Name:

Variant:

Language:

External command (command pre-defined by system administrator)

Name: ZDSPLIBL

Parameters: output(*print)

Operating sys.: OS/400

Target server: AS23

External program (direct command input by system administrator)

Name:

Parameter:

Target host:

Check Control flags

Figure 214. SM36: Defining external command job steps

Figure 215 shows how to specify an external iSeries program as a job step. This can be an iSeries server program or a user program.

Figure 215 shows the 'Edit Step 1' dialog box in SAP SM36. The 'Program values' section has three tabs: 'ABAP program', 'External command', and 'External program'. The 'External program' tab is selected, displaying the following fields:

External program (direct command input by system administrator)	
Name	runqry
Parameter	qry(rfc/z8189qry) outtype(*printer)
Target host	as23

The 'ABAP program' tab shows fields for Name, Variant, and Language. The 'External command' tab shows fields for Name, Parameters, Operating sys., and Target server. At the bottom, there are buttons for 'Check', 'Control flags', and a close button.

Figure 215. SM36: Defining external program job steps

The job steps will be run in the sequence specified. Schedule the background job to run at the required date and time.

12.6 iSeries jobs starting SAP R/3 applications

The following methods can be used to start SAP R/3 jobs from the iSeries server:

- Send an event to the SAP R/3 background job scheduler, using the SAP Event (SAPEVT) CL command.
- Start an ABAP program, using the STRREPORT CL command.

Both the SAPEVT and STRREPORT commands are SAP-supplied CL commands and reside in the SAP R/3 kernel library.

The following sections briefly explain how to use these two methods.

12.6.1 The SAPEVT command

The SAPEVT command enables iSeries jobs to trigger SAP R/3 background jobs, using the SAP R/3 event interface. When you define an SAP R/3 background job (using SAP R/3 transaction SM36), you can specify the job start criteria. These include:

- Immediately
- At a specific date and time
- After a certain job has finished
- At the start of a specific operation mode
- Depending on a specific event being raised

Briefly, an event is a signal that can be sent to the SAP R/3 background job scheduler to trigger a process. This signal can be sent from either inside the SAP R/3 system or from an external iSeries job.

Use the SAP R/3 transaction SM62 to maintain events.

The following list briefly explains how to trigger an SAP R/3 job, using SAPEVT:

1. Define the SAP R/3 background job using transaction SM36.
2. Specify the job steps to be executed.
3. Select **After event** from the start time selection options.
4. Select the **Periodic job** option to automatically re-schedule the job upon completion (if required). This option will automatically schedule and release a new job (with the same name) upon completion, waiting for the event to be raised again.
5. Save and release the background job. The job cannot be triggered if it is not in released state.

This background job will start running once the event is raised. Figure 216 shows how to specify the event settings in the background job definition.

The screenshot shows the 'Start Time' dialog box in SAP SM36. It has five tabs: 'Immediate', 'Date/Time', 'After job', 'After event', and 'At operation mode'. The 'After event' tab is active. Below the tabs, there are three main sections: 'Date/Time' (empty), 'After job' (empty), and 'After event' (containing 'Event: Z8189_EVENT' and an empty 'Parameter' field). At the bottom, there is a 'Periodic job' checkbox which is checked. The very bottom of the dialog has three buttons: 'Check' (with a green checkmark), 'Save' (with a floppy disk icon), and 'Cancel' (with a red X icon).

Figure 216. SM36: Job starting after an event

The above event can then be triggered from the iSeries by issuing the following CL command:

```
R346DOPT/SAPEVT EVTID(Z8189_EVENT)
PROFILE('/usr/sap/R01/sys/profile/R01_DVEBMG S01_AS23')
```

This command may also be incorporated into a CL program that can be used to trigger the event.

12.6.2 The STRREPORT command

The STRREPORT command starts an ABAP program, using the iSeries CL command interface. This command can be called interactively or from an iSeries job stream. The following iSeries CL command executes the ABAP program, Z8189AB3, in the specified SAP R/3 system:

```
R346DOPT/STRREPORT REPORT(Z8189AB3) JOB(Z8189_STRREPORT) SID(R01)
      INSTANCE(01) CLIENT(300) USER(MONTI_A) LANGUAGE(E)
```

12.7 Interactive program communication

ABAP programs can communicate directly with iSeries non-SAP R/3 programs written in ILE C/400, ILE RPG/400, or ILE COBOL/400. Other languages may be used, but certain functions may not be available. CL programs only provide limited functionality, because the CL language does not support pointers. You may choose to invoke the interactive send/receive of the data from either the ABAP or the non-SAP R/3 program, by using either a CPI-C or RFC interface. CPI-C is a low-level interface on which the CPI-C functions are built. Both these interfaces are delivered by SAP. In this section, both techniques are shown based on ABAP and ILE RPG/400 programs. Figure 217 shows the ABAP program invoking the ILE RPG/400 to retrieve the item description from a database file.




Figure 217. CPI-C example: ABAP - RPG/400 ILE

The ABAP program reads the input parameter and receives the field ZITEMNM (item number). To retrieve the item description, the ABAP program invokes the ILE RPG/400 program, passing the item number to it. This program performs a random read on the database file T8189DB, using the item number as the key. The item-description is passed back to the ABAP program.

12.7.1 Using the CPI-C connection

The Common Programming Interface-Communication (CPI-C) delivered by SAP is a low-level interface for program-to-program communication. It can be used for these types of communication:

- Establish and control a conversation
- Exchange data
- Convert data (ASCII to EBCDIC and vice versa)
- Deallocate a conversation

SAP R/3 provides the CPI-C Software Developers Kit (SDK) with CPI-C interfaces both for the ABAP and the OS/400 ILE compiler. The communication can be established based both on the SNA LU 6.2 and TCP/IP protocols. CPI-C functions for the OS/400 ILE compiler are delivered within service programs (OS/400 object type *SRVPGM). To access the functions within your OS/400 ILE program, you have to bind the following service programs to your program:

- CPICSLIB if using the LU 6.2 protocol
- CPICTLIB if using the TCP/IP protocol

Two groups of CPI-C functions are available:

- CPI-C start, which includes the following functions:
 - Establish a connection
 - Exchange data
 - Deallocate a connection
- CPI-C advanced function calls, which include the following functions:
 - Convert data
 - Manage and change a conversation
 - Security functions

The CPI-C interface can establish a program-to-program connection between ABAP programs and other jobs running outside of SAP R/3 on the same system or another system connected to it:

- An ABAP program invokes a program outside of SAP R/3.
- A program outside of SAP R/3 invokes an ABAP program.

The example in the following section demonstrates an ABAP program evoking an RPG/400 program running on the same system. The TCP/IP protocol and only functions of the CPI-C start set are used. Refer to the relevant SAP R/3 documentation for more examples based on the C programming language.

12.7.1.1 CPI-C connection ABAP: RPG/400

Figure 218 shows the connection between ABAP program Z8189AB7 and ILE RPG/400 program T8189ILE. The ABAP program Z8189AB7 is running in a work process under the SAP R/3 instance subsystem R3_00 and invokes ILE RPG/400 program T8189ILE, within the same subsystem. The communication between the two programs is provided by the CPI-C handler of the SAP R/3 gateway process, which must be activated for this instance. The SAP R/3 gateway job spawns a new job in subsystem R3_00, which in turn, calls the ILE RPG/400 program T8189ILE. The SAP R/3 table TXCOM contains information needed by the ABAP program to establish the communication.




Figure 218. CPI-C Connection ABAP: RPG/400 ILE

12.7.1.2 Sideinfo table TXCOM

To initialize program-to-program communication, the SAP R/3 CPI-C interface needs information about the remote environment. This information is stored in *sideinfo tables*. The ABAP program accesses the ABAP Dictionary table TXCOM for an appropriate destination entry when initializing communications.

Use the SAP R/3 transaction SM54 to maintain table TXCOM (Figure 219).

The screenshot shows the 'Change XCOM Entry' dialog box in SAP. It has a table-like structure with the following fields and values:

Dest.	AS23
LU	as23
TP	/qsys.lib/rfc.lib/t8189ile.pgm
Log	E
Gateway host	
Gateway service	
Trace	OFF

At the bottom of the dialog are three icons: a green checkmark, a floppy disk, and a red 'X'.

Figure 219. SM54: Maintaining sideinfo table TXCOM

Create one record for each destination system you want to access, specify:

- **Symbolic Destination (Dest.):** A symbolic name for the communication definition. The ABAP communication interface refers to this name when initializing the communication.

- **Logical Unit (LU):** The name of the logical unit where the remote program is located.
- **Transaction Program (TP):** The path and name of the remote program you want to invoke.
- **Protocol Type (Log):** Describes the communications type, for example:
 - C: Communication with SAP R/2 ABAP program or LU 6.2 connection
 - I: Communication with SAP R/3 ABAP program
 - E: Communication with ILE program based on TCP/IP protocol
- **Gateway host:** This is only required when using a gateway service running on a different system.
- **Gateway service:** This is only required when using a gateway service running in a different SAP R/3 instance.

Note: The Gateway host and Gateway service parameters are required for R/2 systems.

12.7.1.3 CPI-C interaction between ABAP and RPG/400

For CPI-C-based interaction, SAP provides CPI-C interfaces both for ABAP and programs outside of SAP R/3:

- **ABAP:** Uses the COMMUNICATION command
- **ILE RPG/400:** Calls bound procedures (CALLB command) from the service program bound to your program (CPICSLIB for LU 6.2 or CPICTLIB for TCP/IP)

The functions in Table 31 are available to perform a CPI-C based interaction between ABAP and ILE RPG/400 programs.

Table 31. SAP CPI-C communication functions

CPI-C call from an ABAP program	CPI-C call from an ILE RPG/400 program (LU 6.2)	CPI-C call from an ILE RPG/400 program (TPC/IP)
COMMUNICATION INIT	CMINIT	STINIT
COMMUNICATION ALLOCATE	SAP_CMALLC, CMALLC	SAP_CMALLC, STALLC
COMMUNICATION ACCEPT	CMACCP	STACCP
COMMUNICATION SEND	CMSEND	STSEND
COMMUNICATION RECEIVE	CMRCV	STRCV
COMMUNICATION DEALLOCATE	CMDEAL	STDEAL

The example in Figure 220 shows the interaction between an ABAP program and ILE RPG/400 program, based on a TCP/IP connection. The ABAP program Z8189AB7 establishes the connection to the ILE RPG/400 program T8189ILE using the COMMUNICATION INIT and ALLOCATE functions. The ABAP program sends the data. Then, the ILE RPG/400 program accepts the conversation and waits for the first input from the ABAP program.




Figure 220. CPI-C interaction between ABAP and ILE RPG/400

The numbers in Figure 220 correspond to the numbers in the following explanation:

1. **Communication Init:** Communication, using the symbolic destination name in the SAP R/3 table TXCOM, is initialized and the Conversation-ID is assigned.
2. **Communication Allocate:** A session is established. and the ILE RPG/400 program T8189ILE is activated.
 - Connects to the gateway process (SAP_CMACCP).
 - Sends acceptance of the conversation back to the ABAP program (STACCP) and receives back the Conversation-ID.
 - Waits for data to be sent from the ABAP program (STRCV).
3. **Communication Send:** Data is sent from the ABAP program to the ILE RPG/400 program.
4. **Communication Receive:** Data is sent back from the ILE RPG/400 program to the ABAP program.
5. **Communication Deallocate:** The ABAP program detaches the session.

12.7.1.4 ABAP program Z8189AB7

Create and execute the ABAP program Z8189AB7 shown in Figure 221.

```
REPORT Z8189AB7.
PARAMETERS: ZITEMNM(5) DEFAULT '  1'.
INCLUDE RSCPICDF.
DATA: CONV_ID(8) TYPE C,
      DEST(8) TYPE C VALUE 'AS23',
      RC LIKE SY-SUBRC,
      ZITEMD(25) TYPE C,
      DI(4) TYPE X,
      SI(4) TYPE X,
      LEN TYPE P VALUE 5.
COMMUNICATION INIT ID CONV_ID DESTINATION DEST RETURNCODE RC.
IF RC <> CM_OK. WRITE: / RC, 'INIT'. EXIT. ENDIF.
COMMUNICATION ALLOCATE ID CONV_ID RETURNCODE RC.
IF RC <> CM_OK. WRITE: / RC, 'ALLOCATE'. EXIT. ENDIF.
COMMUNICATION SEND ID CONV_ID BUFFER ZITEMNM LENGTH LEN.
IF RC <> CM_OK. WRITE: / RC, 'ALLOCATE'. EXIT. ENDIF.
COMMUNICATION RECEIVE ID CONV_ID BUFFER ZITEMD DATAINFO DI
                     STATUSINFO SI RETURNCODE RC.
IF RC <> CM_OK. WRITE: / RC, 'ALLOCATE'. EXIT. ENDIF.
WRITE: / ZITEMNM, ZITEMD.
COMMUNICATION DEALLOCATE ID CONV_ID RETURNCODE RC.
IF RC <> CM_OK. WRITE: / RC, 'DEALLOCATE'. EXIT. ENDIF.
EXIT.
```

Figure 221. ABAP program Z8189AB7

Here is a brief overview of the ABAP program Z8189AB7:

- Defines the input-parameter field ZITEMNM (default value 1).
- Includes RSCPICDF for definitions of symbolic values (for example, return code value CM_OK).
- Initializes and allocates communications.
- Sends the item number to the remote program.
- Receives the item description from the remote program.
- Deallocates the communication.

12.7.1.5 RPG/400 program T8189ILE

This section lists ILE RPG/400 program T8189ILE that is invoked from ABAP program Z8189AB7 (by using the symbolic destination entry defined in table TXCOM, as explained in 12.7.1.2, “Sideinfo table TXCOM” on page 306). For a description on how to communicate with the ABAP program, refer to the remarks inside the program.

```
*****
FT8189DB  IF  E          K DISK
***** Definition of
basing pointers, pointing to input
* parameters, being passed back to CPI-C handler
***** DPT
DS
D PT1          *
D PT2          *
```

```

D PT3                      *
D PT4                      *
D PT5                      *
D PT6                      *
D PTNULL                   *   INZ(*NULL)
**
DITEMKEY                   5A
*****
* Definition of parameters being passed to CPI-C
* function calls
*****
D CMPARM                   DS
D CONVID                   1      8
D RTNCOD                   9     12B 0
D DATRCV                  13     16B 0
D DLCTYP                   17     20B 0
D RCVLEN                   21     24B 0
D REQLEN                   25     28B 0
D REQTSR                   29     32B 0
D SNDLEN                   33     36B 0
D SNDTYP                   37     40B 0
D STSRCV                   41     44B 0
D PRERCV                   45     48B 0
*****
* All parameters necessary for establishing the connection
* will be passed into the program. SAP_CMCCP will pass the
* field pointers back to CPI-C handler (data structure PT).
* - SAPGW (SAP Gateway), for example 'sapgw00'
* - CONVID1 (Conversation ID ABAP program)
* - SAPHOST (TCP/IP host name of iSeries),
*   for example: 'GWHOST=sysnm002.sysnmaa.ibm.com'
* - SAPRE1-3: Reserved fields, contain
*   gateway name (SAPRE1)
*   convid ABAP (SAPRE2)
*   IFS-path and name of instance profile (SAPRE3)
*   (for example: '/usr/sap/AIM/SYS/profile/AIM_DVEBMGS00')
*****
C   *ENTRY                 PLIST
C                           PARM                      SAPGW                      64
C                           PARM                      CONVID1                    64
C                           PARM                      SAPHOST                    64
C                           PARM                      SAPRE1                      64
C                           PARM                      SAPRE2                      64
C                           PARM                      SAPRE3                      64
*****
* Setting values of basing pointers, pointing to variables
* being passed to CPI-C function SAP_CMCCP
*****
C                           EVAL      PT1 = %addr(SAPGW)
C                           EVAL      PT2 = %addr(CONVID1)
C                           EVAL      PT3 = %addr(SAPHOST)
C                           EVAL      PT4 = %addr(SAPRE1)
C                           EVAL      PT5 = %addr(SAPRE2)
C                           EVAL      PT6 = %addr(SAPRE3)
C* *****
* CPI-C module SAP_CMCCP:
*   - connection to gateway process
*   - PT: data structure, points to *entry parameter fields
*****
C*
C                           CALLB      'SAP_CMCCP'
C                           PARM                      PT
C*
*****
* CPI-C module STACCP:
*   - accepts conversation
*   - getting back CPI-C conversation ID
*****
C*
C                           CALLB      'STACCP'
C                           PARM                      CONVID
C                           PARM                      RTNCOD
C*
*****
* CPI-C module STRCV:
*   - receives data (requested length=5) into field ITEMKEY
*   - REQLEN: requested length
*   - DATRCV: data received (logical value)

```

```

*      - REVLEN: received length
*      - STSRCV: status received (logical value)
*      - REQTSR: request to send
*      - RTNCOD: returncode
*****
C*
C          Z-ADD      5          REQLEN
C          CALLB      'STRCV'
C          PARM
C          PARM
C          PARM          REQLEN
C          PARM          DATRCV
C          PARM          RCVLEN
C          PARM          STSRCV
C          PARM          REQTSR
C          PARM          RTNCOD
C*
C*****
C*      DOWEQ: stops, when ABAP will send COMMUNICATION DEALLOCATE
C*****
C      RTNCOD      DOWEQ      0
C      MOVEL      *BLANKS      ITEM
C      CHAIN      T8189DB
C      99
C*
C*****
* CPI-C module STSEND:
*      - sends data (requested length=25, field ITEM)
*      - REQLEN: requested length
*      - REQTSR: request to send
*      - RTNCOD: returncode
*****
C*
C          Z-ADD      25          REQLEN
C          CALLB      'STSEND'
C          PARM
C          PARM          CONVID
C          PARM          ITEM
C          PARM          REQLEN
C          PARM          REQTSR
C          PARM          RTNCOD
C*
C      RTNCOD      IFEQ      0
C          Z-ADD      5          REQLEN
C*
C*****
* CPI-C module STSRCV:
*      - receives data (requested length=5) into field ITEMKEY
*      - REQLEN: requested length
*      - DATRCV: data received (logical value)
*      - REVLEN: received length
*      - STSRCV: status received (logical value)
*      - REQTSR: request to send
*      - RTNCOD: returncode
*****
C*
C          Z-ADD      5          REQLEN
C          CALLB      'STRCV'
C          PARM
C          PARM          CONVID
C          PARM          ITEMKEY
C          PARM          REQLEN
C          PARM          DATRCV
C          PARM          RCVLEN
C          PARM          STSRCV
C          PARM          REQTSR
C          PARM          RTNCOD
C          ENDIF
C*
C          ENDDO
C*****
C          SETON          LR
C*****
C*****

```

12.7.1.6 Creating RPG/400 program T8189ILE

To obtain an executable program, perform the following steps:

1. Create an RPG module by running the following OS/400 command:

```
CRTRPGMOD MODULE(RFC/T8189ILE) SRCFILE(RFC/QRPGLESRC)
```

2. Create the ILE RPG/400 program, using the above RPG module, and bind the SAP service program CPICTLIB to it:

```
CRTPGM PGM(RFC/T8189ILE) BNDSRVPGM(R346DOPT/CPICTLIB)
```

12.7.2 Using RFC connection

The Remote Function Call (RFC) is a consistent set of functions for program-to-program communication in a heterogeneous environment. It can be used for the following types of communications:

- ABAP <—> ABAP
- ABAP <—> external non-ABAP programs

External programs must be written in ILE compiler languages (ILE RPG/400, ILE COBOL/400, or ILE C/400 are required). Similar to the CPI-C connection, which is described in 12.7.1, “Using the CPI-C connection” on page 305, RFC connections can:

- Establish and control communications
- Exchange data
- Convert data (ASCII <—> EBCDIC)
- Deallocate communications

All of the above functions are automatically done by the CALL FUNCTION statement in the ABAP program. See ABAP program Z8189ABB (Figure 228 on page 318) for an example of how to use this statement. OS/400 ILE programs, however, must specify this information in detail, by using the RFC procedures for OS/400 ILE compilers that are delivered within the service program LIBRFC. To access these functions, service program LIBRFC must be bound with the OS/400 ILE program.

12.7.2.1 RFC interaction between ABAP and RPG/400

For RFC-based interaction, SAP provides RFC interfaces for ABAP programs outside of SAP R/3:

- **ABAP:** Use statement CALL FUNCTION.
- **ILE RPG/400:** Bind the service program LIBRFC to the ILE RPG/400 program.

Table 32 lists the procedures that must be called from within the ILE RPG/400 program to perform an RFC-based interaction.

Table 32. SAP RFC communication call functions

Procedure call	Description
RfcEnvironment	Points to an error control program
RfcAccept	Establishes a conversation
RfcInstall	Points to an RFC-definition, server program
RfcDispatch	Invokes a server program
RfcGetData	Receives data from ABAP program

Procedure call	Description
RfcSendData	Sends data from ABAP program
RfcClose	Closes conversation

ILE RPG/400

For more information on ILE RPG/400, refer to:

- *ILE RPG/400 Programmer's Guide*, SC09-2507
- *ILE RPG/400 Reference*, SC09-2508

Figure 222 shows the interaction between the programs.




Figure 222. RFC interaction between ABAP and RPG/400 ILE

The process shown in Figure 222 is explained here:

1. ABAP program Z8189ABB establishes a session to RPG/400 program T8189RFC by calling Remote Function Module Z_8189.
2. The function module is configured to send data (Exporting) and to receive data immediately afterwards (Importing).
3. RFC destination Z8189 invokes RPG/400 program T8189RFC.
4. The main RPG/400 procedure, T8189RFC, is started:

- a. **RfcEnvironment:** A connection to the error recovery program is initialized.
- b. **RfcAccept:** A session is established and the RFC-handle code is passed back to the RPG/400 program.
- c. **RfcInstall:** A connection to the Remote Function Module and server program is established.
- d. **RfcDispatch:** The server program, T8189SRV (an RPG subprocedure), is invoked and performs the following actions:
 - i. Reads data from ABAP program (**RfcGetData**).
 - ii. Sends data back to ABAP program (**RfcSendData**).
- e. **RfcClose:** Closes the conversation.

12.7.2.2 RFC and RFC destination definition

Use SAP R/3 transaction SE37 to define a Remote Function Group (if required), and the Remote Function Module Z_8189, which will be used by the ABAP program Z8189ABB. The following examples explain how to define the Remote Function Module.

Ensure that the Function Module Processing type is set to **Remote enabled module** (Figure 223).

The screenshot shows the SAP Function Builder: Change Z_8189 dialog box. The 'Function module' is Z_8189, and its status is 'Inactive(revised)'. The 'Attributes' tab is active, showing the 'Classification' section with 'Function group' Z001 and 'Short text' IBM: SAP R/3 Redbook. The 'Processing type' section has 'Remote-enabled module' selected, with sub-options 'Start immedi.', 'Immediate start, no restart', 'Start delayed', and 'Call run'. The 'RFC target host' button is visible. The 'General data' section includes fields for 'Person responsible' (MONTI_A), 'Last changed by' (MONTI_A), 'Changed on' (23.11.2000), 'Development class' (ZBASIS), 'Program name' (SAPLZ001), 'INCLUDE name' (LZ001U01), 'Original language' (EN), and 'Not released' (unchecked). There are also checkboxes for 'Edit lock' and 'Global'.

Figure 223. Defining the function module attributes (Part 1 of 3)

Figure 224 and Figure 225 show the Import and Export settings required to communicate with the RPG/400 program.

Create the RFC destinations, using the SAP R/3 transaction SM59, as shown in Figure 226 and Figure 227. The remote program and the target server name is specified in the RFC destination settings.

The screenshot shows the SAP SM59 'RFC Destination' configuration window. The title bar includes 'Destination', 'System information', 'Test', 'System', and 'Help'. The main window has a 'Test connection' button at the top left. The 'RFC destination' field is set to 'T8189'. Below this, the 'Technical settings' section shows 'Connection type' as 'T' and 'TCP/IP connection' selected. The 'Description' section contains 'T8189'. The 'Logon' section has fields for 'Language', 'Client', 'User', and 'Password' (masked with asterisks). The 'Attributes' section has fields for 'Created by' and 'Last changed by'.

Figure 226. Creating an RFC destination (Part 1 of 2)




Figure 227. Creating an RFC destination (Part 2 of 2)

12.7.2.3 ABAP program Z8189ABB

This section explains the ABAP program Z8189ABB. It performs the following steps:

- Defines the input-parameter field ZITEMNM (default value 1).
- Calls the Remote Function Module, which:
 - Establishes a connection to the RFC destination T8189 (refer to Figure 227).
 - Sends and receives data (ITEMNM, ITEMMD).
 - Passes back the return code.
 - Closes the connection.
 - Processes the results.
 - Performs error recovery.

The ABAP program Z8189ABB is shown in Figure 228.

```

REPORT Z8189ABB.
PARAMETERS: ITEMNM(5) DEFAULT '    1'.
INCLUDE RSCPICDF.
DATA: ITEMND(25) TYPE C.
CALL FUNCTION 'Z_8189'
  DESTINATION 'T8189'
  EXPORTING
    ZITEMNM = ITEMNM
  IMPORTING
    ZITEMD = ITEMND
  EXCEPTIONS
    OTHERS = 1.
IF SY-SUBRC = 0.
  WRITE: / ITEMNM, ITEMND.
  EXIT.
ENDIF.
CASE SY-SUBRC.
  WHEN OTHERS. WRITE: / SY-SUBRC, '/FEHLER'.
ENDCASE.

```

Figure 228. SE38: ABAP program Z8189ABB

12.7.2.4 RPG/400 program T8189RFC

This section lists ILE RPG/400 program T8189RFC, as well as its embedded ILE RPG/400 modules T8189SRV and T8189ERR. These modules are invoked from ABAP program Z8189ABB. Refer to the remarks inside the RPG/400 program source code for information on how to communicate with the ABAP program.

ILE RPG program T8189RFC

```

FMT H HKeywords+++++
0001.00 D*****
0002.00 D* Pointer Structure for RfcAccept (argv = Dummy-Pointer)
0003.00 D*****
0004.00 DPT DS
0005.00 D argv *
0006.00 D PT1 *
0007.00 D PT2 *
0008.00 D PT3 *
0009.00 D PTNULL * INZ(*NULL)
0010.00 D*****
0011.00 D* RfcEnvironmet Pointer Definition
0012.00 D*****
0013.00 D*
0014.00 D*
0015.00 Denvds DS
0016.00 Dptrerr * PROCPTR
0017.00 Dptrmem * PROCPTR
0018.00 D*****
0019.00 D* RfcInstall Pointer Definitions
0020.00 D*****
0021.00 Dins1 S *
0022.00 Dins2 S * PROCPTR
0023.00 Dins3 S *
0024.00 D*****
0025.00 D* RFC - Handle Definition
0026.00 D*****
0027.00 DhRfc S 8B 0
0028.00 D*****
0029.00 D* Returncode Definitions
0030.00 D*****
0031.00 DRfcRc S 8B 0
0032.00 DRcen S 8B 0
0033.00 D*****

```

```

0034.00 D* Defining Procedure Prototypes
0035.00 D* Symbolic Procedure Names are Pointing to them
0036.00 D*****
0037.00 DRfcEnvironment    PR                8B 0 ExtProc('RfcEnvironment')
0038.00 D*
0039.00 D ptrerr            *      PROCPTR
0040.00 D ptrmem            *      PROCPTR
0041.00 DRfcAccept          PR                8B 0 ExtProc('RfcAccept')
0042.00 D argv              *
0043.00 DRfcDispatch        PR                8B 0 ExtProc('RfcDispatch')
0044.00 D hRfc              8B 0 VALUE
0045.00 DRfcInst            PR                8B 0 ExtProc('RfcInstallFunction')
0046.00 D ins1              *      VALUE
0047.00 D ins2              *      PROCPTR VALUE
0048.00 D ins3              *      VALUE
0049.00 DRfcClose           PR                ExtProc('RfcClose')
0050.00 D hRfc              8B 0 VALUE
FMT C   CLON01Factor1+++++Opcode&ExtFactor2+++++Result+++++Len++D+HiLoE
0051.00 C*****
0052.00 C* RFC Communication Environment Parameters
0053.00 C* RfcAccept (Parameter argv) is pointing to them
0054.00 C*****
0055.00 C      *ENTRY      PLIST
0056.00 C              PARM              F1              64
0057.00 C              PARM              F2              64
0058.00 C              PARM              F3              64
0059.00 C*****
0060.00 C* Initializing Fields and Pointers
0061.00 C*****
0062.00 C*
0063.00 C              MOVEL      'Z_8189'      RFCC              6
0064.00 C              MOVEL      'Z8189'      RFCC1             5
0065.00 C              EVAL      argv = %ADDR(F1)
0066.00 C              EVAL      PT1 = %ADDR(F1)
0067.00 C              EVAL      PT2 = %ADDR(F2)
0068.00 C              EVAL      PT3 = %ADDR(F3)
0069.00 C              EVAL      ins1 = %ADDR(RFCC)
0070.00 C              EVAL      ins2 = %PADDR('T8189SRV')
0071.00 C              EVAL      ins3 = %ADDR(RFCC1)
0072.00 C              EVAL      ptrerr = %PADDR('T8189ERR')
0073.00 C              EVAL      ptrmem = *NULL
0074.00 C*
0075.00 C*****
0076.00 C* Evoking RFC - Procedures (Service-Program LIBRFC
0077.00 C* - RfcEnvironment : Pointing to Procedure T8189ERR
0078.00 C* - RfcAccept      : Getting RFC-Handle
0079.00 C* - RfcInstallFunction : Pointing to Procedure T8189SRV
0080.00 C* - RfcDispatch    : Evoking Server Program T8189SRV
0081.00 C* - RfcClose      : Close Conversation
0082.00 C*****
0083.00 C* env - Points to Error Routine T8189ERR
0084.00 C*
0085.00 C*
0086.00 C*
0087.00 C              CALLP      RfcEnvironment(ptrerr : ptrmem)
0088.00 C* argv - Points to Entry Parameter Fields
0089.00 C* hRfc - RFC - Handle
0090.00 C              EVAL      hRfc = RfcAccept(argv)
0091.00 C* ins1 - Points to Remote Function Call Name Z_8189
0092.00 C* ins2 - Points to RPG-Module T8189SRV (Get and Send Data)
0093.00 C* ins1 - Points to RFC Destination Name Z8189
0094.00 C* RfcRc - Returncode
0095.00 C              EVAL      RfcRc = RfcInst(ins1 : ins2 : ins3)
0096.00 C              EVAL      RfcRC = RfcDispatch(hRfc)
0097.00 C              CALLP      RfcClose(hRfc)
0098.00 C*****
0099.00 C              MOVE      '1'      *INLR
0100.00 C              RETURN
0101.00 C*****

```

RPG/400 subprocedure T8189SRV

```

FMT H   HKeywords+++++
0001.00 H*****
0002.00 H* Defines T8189SRV as a Subprocedure
0003.00 H*****
0004.00 HNOMAIN
FMT F   FFilename++IPEASFRlen+LKlen+AIDevice+.Keywords+++++

```

```

0005.00 FT8189DB IF E K DISK
FMT D DName+++++ETDsFrom+++To/L+++IDc.Keywords+++++
0006.00 D*****
0007.00 D* Prototype Definition for Entry Parmstructure
0008.00 D* Returncode : 8 Bytes Binary
0009.00 D* Parameter : 8 Bytes passed by Value
0010.00 D*****
0011.00 DT8189SRV PR 8B 0
0012.00 D hRfc 8B 0 VALUE
0013.00 D*****
0014.00 D* Procedure Begin Definition
0015.00 D*****
0016.00 PT8189SRV B EXPORT
0017.00 D*****
0018.00 D* Procedure Interface Definition
0019.00 D*****
0020.00 DT8189SRV PI 8B 0
0021.00 D hRfc 8B 0 VALUE
0022.00 D*****
0023.00 D* Field Structure RfcGetData
0024.00 D* Point : Points to Field Name (RFC-Definition)
0025.00 D* NLEN : Field Name Length
0026.00 D* DTYPE : Data Type (0=Char)
0027.00 D* DLEN : Received Data Length
0028.00 D* PTNULL : NULL - Pointer (Table Reference)
0029.00 D*****
0030.00 DPTGET DS
0031.00 D POINT *
0032.00 D NLEN 8B 0
0033.00 D DTYPE 8B 0
0034.00 D DLEN 8B 0
0035.00 D PTKEY *
0036.00 D PTNULL * INZ(*NULL)
0037.00 D*****
0038.00 D* Field Structure RfcSendData
0039.00 D* (Refer to Structure RfcGetData)
0040.00 D*****
0041.00 DPTSEND DS
0042.00 D POINTS *
0043.00 D NLENS 8B 0
0044.00 D DTYPES 8B 0
0045.00 D DLENS 8B 0
0046.00 D PTKEYS *
0047.00 D PTNULLS * INZ(*NULL)
0048.00 D*****
0049.00 DRcge S 8B 0
0050.00 D*****
0051.00 D* Prototype Definitions
0052.00 D*****
0053.00 DRfcGet PR 8B 0 ExtProc('RfcGetData')
0054.00 D hRfc 8B 0 VALUE
0055.00 D POINT *
0056.00 D PTNULL *
0057.00 DRfcSend PR 8B 0 ExtProc('RfcSendData')
0058.00 D hRfc 8B 0 VALUE
0059.00 D POINTS *
0060.00 D PTNULLS *
FMT C CL0N01Factor1+++++Opcode&ExtFactor2+++++Result+++++Len++D+HiLoE
0061.00 C*****
0062.00 C* Initializing Fields
0063.00 C*****
0064.00 C MOVE 'ZITEMNM' FNAME 7
0065.00 C Z-ADD 7 NLEN
0066.00 C Z-ADD 5 DLEN
0067.00 C Z-ADD 0 DTYPE
0068.00 D*****
0069.00 C MOVE 'ZITEMD' FNAME1 7
0070.00 C Z-ADD 6 NLENS
0071.00 C Z-ADD 25 DLENS
0072.00 C Z-ADD 0 DTYPES
0073.00 C Z-ADD 0 Rcge
0074.00 C MOVE ' ' ITEMKEY 5
0075.00 C EVAL POINT = %ADDR(FNAME)
0076.00 C EVAL POINTS = %ADDR(FNAME1)
0077.00 C EVAL PTKEY = %ADDR(ITEMKEY)
0078.00 C EVAL PTKEYS = %ADDR(ITEMD)
0079.00 C*****
0080.00 C* Evoking RFC - Procedure (Service-Program LIBRFC)

```

```

0081.00 C* - RfcGetData : Reads Data from ABAP/4 Z8189ABB
0082.00 D*****
0083.00 C EVAL Rcge = RfcGet(hRfc :POINT : PTNULL)
0084.00 C*****
0085.00 C MOVE *BLANKS ITEM
0086.00 C ITEMKEY CHAIN T8189DB 99
0087.00 C*****
0088.00 C* Evoking RFC - Procedure (Service-Program LIBRFC)
0089.00 C* - RfcSendData : Sends Data back to ABAP/4 Z8189ABB
0090.00 C*****
0091.00 C EVAL Rcge = RfcSend(hRfc :POINTS : PTNULL)
0092.00 D*****
0093.00 C CLOSE T8189DB
0094.00 C MOVE '1' *INLR
0095.00 C RETURN Rcge
FMT P O.....N01N02N03Field++++++YB.End++PConstant/editword/DTfor
0096.00 PT8189SRV E
RPG/400 sub-procedure T8189ERR
FMT H HKeywords+++++++
0001.00 HNOMAIN
FMT D DName++++++ETDsFrom+++To/L+++IDc.Keywords+++++++
0002.00 DT8189ERR PR
0003.00 D errmsg 64
0004.00 PT8189ERR B EXPORT
0005.00 DT8189ERR PI
0006.00 D errmsg 64
FMT C CL0N01Factor1++++++Opcode&ExtFactor2++++++Result++++++Len++D+HiLoE
0007.00 C*****
0008.00 C MOVE '1' *INLR
0009.00 C RETURN
FMT P O.....N01N02N03Field++++++YB.End++PConstant/editword/DTfor
0010.00 PT8189ERR E

```

12.7.2.5 Creating ILE RPG/400 program T8189RFC

To create an executable program, perform the following steps:

1. Use the OS/400 `CRTTRPGMOD` command to create each of the RPG modules T8189RFC, T8189ERR, and T8189SRV, for example:

```

CRTTRPGMOD MODULE(RFC/T8189RFC) SRCFILE(RFC/QRPGLESRC)
CRTTRPGMOD MODULE(RFC/T8189ERR) SRCFILE(RFC/QRPGLESRC)
CRTTRPGMOD MODULE(RFC/T8189SRV) SRCFILE(RFC/QRPGLESRC)

```

2. Use the OS/400 `CRTPGM` command to create the program. Specify all the RPG modules created in the first step. Remember to bind the SAP service program LIBRFC to this program as well. For example, run this command:

```

CRTPGM PGM(RFC/T8189RFC) MODULE(T8189RFC T8189ERR T8189SRV)
        BNDSRVPGM(R346DOPT/LIBRFC)

```

Chapter 13. MQSeries link for R/3

This chapter presents a brief overview of MQSeries, the MQSeries link for R/3, and using MQSeries link with Application Link Enabling (ALE). It explains the value of using the MQSeries link for R/3 where a connection is required between R/3 and external systems.

The MQSeries link for R/3 works in conjunction with the enhanced Intermediate Document (IDoc) interface. They provide customers with access from R/3 or ABAP applications using ALE to the assured, time-independent, message delivery service provided by IBM MQSeries.

This chapter is beneficial for SAP consultants who are installing R/3 systems for customers. It also helpful for sales representatives who are proposing R/3, where connection to external systems (that is, non-SAP systems) is important.

13.1 Introduction and scope

The success of SAP R/3 has meant that it is frequently installed where there are existing systems, running outside the SAP environment. Many of the R/3 applications need to exchange data with these existing systems, which means an interface or bridge is needed between the two environments. Even if the plan is to migrate all the systems to R/3 (which is unlikely to be practicable in every case), it is not feasible to convert them all at once. Staging is required.

A few typical examples drawn from real customer experiences include:

- An SAP Sales and Distribution system, running on OS/400, needs to interface with a plant order management system running on MVS. See Figure 229.
- SAP R/3 systems, running on OS/400, need to interface to existing financial systems running on AIX.
- An SAP Financial system, running on AIX, needs to interface to a manufacturing control system running on OS/400.




Figure 229. Example of application integration

The interface requirements are frequently far more complex than these simple examples suggest. One major oil company identified the need for several hundred interfaces between their various systems.

The data flow may be one-way, either from R/3 to the external system or from the external system to R/3, or it may be two-way. The data being exchanged is of high value. It is business information that can genuinely be described as mission critical. The originating application needs to be sure that its data will be delivered to the target application, and that it will be informed if delivery cannot be achieved. Both applications need to transfer and receive data under transactional control, so that their views of the data are consistent once the transfer has occurred.

The exchange of information does not always need to be immediate. The phrase “near real time” is often used to describe the requirement. The systems involved may run in different time zones, perhaps on different continents. They almost always run on different platforms, under different operating systems, and sometimes, on different network protocols. This means that the synchronous, blocking approach typified by Remote Function Call (RFC) or CPI-C is not desirable.

13.2 Customer requirements and situations

This section examines a few customer situations where the MQSeries link for R/3 for iSeries can be used.

Support connection between SAP R/3 and external systems

The major requirement is the ability to link R/3 and external systems that may be running on a variety of different hardware and operating systems. The link must support bidirectional communication – into and out of R/3 – and should be simple to use. It should provide standard facilities and calling mechanisms across all environments. It should also support communication between R/3 and R/2 systems, and between R/3 systems themselves.

Communicate in near real-time

It is not normally necessary for communication between R/3 and external systems to take place immediately, but updates do need to take place promptly in many cases. This can be characterized as “near real time”. A batch update method is not acceptable in these circumstances, although near real-time and batch may be used in parallel where appropriate.

Communicate in a standard way between R/3 systems and external systems

The requirement here is to communicate between R/3 and external systems, and between R/3 systems themselves, by a standard method.

Fast response to changing business circumstances demands great flexibility in organizations and in organizational structures. Information processing must be able to keep pace with changes in the organization and support the move from a centralized to a decentralized model, for example. Competitive pressures mean that companies need to exploit all of the resources and information available to them. This, in turn, puts requirements on information processing to support rapid access to data across organizational boundaries. For example, the sales department needs timely information about the stock of finished goods in their warehouses and at the plant.

Customers should not have to create the infrastructure for such information exchange themselves, but should be able to use a standard predefined method, that is capable of adapting to changing circumstances.

Use a single carrier system for all asynchronous communications

Asynchronous communications may be used in all of the situations cited earlier:

- R/3 to and from external systems or for either system
- R/3 to and from R/2 or for either system
- R/3 to and from R/3 or for either system

Many installations have a combination of two or more of these. For cost and efficiency reasons, it is important to keep the management of the network as simple as possible. Therefore, a single carrier system is required to support all cases.

Have a reliable messaging system with advanced communication features

The importance of the information being exchanged makes it imperative that the system is highly reliable. That is so applications can be sure that messages entrusted to the messaging system are safely delivered. Performance must also be adequate to support possible high volumes of data transfer.

In addition to these fundamental requirements, the system should have some more advanced features, for example to:

- Start an application only when there is work for it to do
- Support different levels of priority within message queues
- Notify an application that a message has been, or cannot be, delivered
- Mix persistent and non-persistent messages (or recoverable and non-recoverable messages) on a queue

Be transparent to R/3 and ABAP programs

R/3 and ABAP applications may need to use the messaging service in combination with other techniques, such as Remote Function Call (RFC) and asynchronous RFC, especially during the transition phase and while existing programs are being transferred to R/3. This means that the application program itself should not have to change when the communication carrier changes. The same calls and parameters should be used as far as possible.

SAP and IBM believe that a combination of ALE and MQSeries can satisfy these requirements for most customers. The following sections present an overview of MQSeries and ALE. They illustrate how these products meet the requirements described here.

13.3 MQSeries overview

MQSeries products incorporate the principles of messaging and queuing, a powerful technique for program-to-program communication. Along with the conversational style represented by CPI-C and remote procedure call, analogous to SAPs RFC, messaging and queuing are one of the three main communication techniques used in program-to-program communication today. Its main differentiator from conversational style and RPC is that it is asynchronous, which makes it very suitable for use in an ALE environment.

13.3.1 How messaging and queuing works

Messaging and queuing enables programs to communicate across a network, without having a private, dedicated, logical connection to link them. It does this in a way that is simple, elegant, and proven. Programs communicate by putting messages on message queues and by taking messages from message queues.

Figure 230 demonstrates this concept. In the figure, program A sends a message to program B via Queue 1. If program B needs to communicate with program A, it puts a message on Queue 2. Programs A and B could run on the same processor in a single environment, on the same processor in different environments, or on different processors in different environments.




Figure 230. How messaging and queuing work

Messages can be one-way (from program A to program B, for example), or they can be reciprocated (a message from program A can cause program B to issue a reply message, for example). However, two-way communication isn't compulsory. If no response is required, none is sent.

Three key facts about messaging and queuing differentiate it from other communication techniques:

- Communicating programs can run at different times (because MQSeries is asynchronous).
- The application structure is not constrained.
- Programs are insulated from network complexities.

13.3.2 An example of a distributed application using MQSeries

To illustrate how the MQSeries products work, consider an example of two programs: one to process customer orders and the other to bill the customer for the goods (Figure 231). For organizational reasons, these programs aren't running on the same computer. The Customer Orders program (CO) is running on a Hewlett Packard processor in a branch office in Chicago. The Customer Billing program (CB) is running on an iSeries server at company headquarters in Dusseldorf. At some point, information has to be exchanged between CO and CB, so that a bill can be produced when goods are ordered. Using the interface supplied by the MQSeries products, the programs CO and CB can communicate at any time.




Figure 231. Programs CO and CB invoke the services of the queue managers

Program CO tells program CB about a customer order by putting a message on CB's message queue (Queue 1). When CB takes the message from Queue 1, it

uses the information in the message to produce a bill for the goods. CB may also want to send a response to CO (simply an acknowledgement, perhaps, or confirmation that the goods have been charged for). CB would do this by putting a message on CO's message queue (Queue 2).

The heavy-duty work – moving messages from CO to Queue 1, and from CB to Queue 2 – is managed by the queue managers.

A program talks directly to its local queue manager (the one on the same processor as the program itself) using the Message Queue Interface (MQI). The MQI is a set of calls that programs use to ask for the services of a queue manager. There are only two basic operations: put a message on a queue (using the MQPUT call) and take a message from a queue (using the MQGET call). The MQI and the queue managers are part of MQSeries.

One of these programs, CO, could be part of an R/3 sales and distribution system, while the other could be part of an existing mainframe-based financial system. The principles remain the same.

As Figure 231 shows, there is a queue manager on each of the two processors on which communicating programs are running. In the example, the queue manager on the Hewlett-Packard processor could be provided by the MQSeries for HP-UX product. And, the queue manager on the iSeries processor could be provided by the MQSeries link for R/3 for iSeries product. When communication across a network is required, the queue managers at the different nodes in the network communicate with each other. When communication between programs at a single node is required, it can be handled by a single queue manager. The programs themselves operate in ignorance of the network and keep no part of the networking burden.

13.3.3 Data integrity and resource protection

MQSeries applications can transfer data with an extremely high degree of confidence. Message delivery can be implemented using a syncpoint mechanism – and the message-queue manager logs or journals – for the recovery of important data in the event of a system failure.

All resources, such as queues, can be protected using the security facilities available on the operating platform.

13.3.4 Message attributes

In MQSeries, messages have attributes, for example, character set and format of data. Two important attributes that are defined in the message descriptor are persistence and priority.

A message is persistent if it survives when MQSeries restarts. This implies that the message must be logged, or saved, and can be reinstated as part of the recovery procedure.

Each MQSeries message has a priority assigned to it by the sending application. The priority, which is a number in the range 0 through 9, can affect the order in which a message is retrieved from a queue and the way that trigger events are generated.

13.3.5 Triggering

Triggering allows an application to be started automatically when predefined conditions on a queue are met. These conditions include, for example, receipt of any message, messages over a particular priority, or the number of messages on a queue.

13.3.6 Message sizes

The maximum message size supported by MQSeries products is 100 MB. However, not all MQSeries products can support messages of this size. In practice, the size of message that an application program can put on a queue is limited by:

- The maximum message length defined for the receiving queue
- The maximum message length defined for the queue manager
- The maximum message length supported by the platform

If an application program is to place on a queue data that exceeds this size limit, the program must split the data into a number of pieces and put each piece on the queue as a separate message.

13.3.7 MQSeries products

MQSeries link for R/3 makes it possible for existing and new R/3 systems to connect to external non-SAP systems with the help of MQSeries products, which are the answer to your customer's immediate or short-term needs. It would, however, be shortsighted to ignore the longer term, and this is where MQSeries really demonstrates its strengths. No matter what plans the customer has for the future, the one thing that is inevitable is change. No matter how careful the planning has been done, the unexpected will alter those plans.

The MQSeries products are the base on which to build tomorrow's client/server networks. At the present time, there are more than 20 platforms.

Any network based on MQSeries can cope with future change with:

- IBM's commitment to the continued enhancement of MQSeries
- The ease with which programs written to the MQSeries Interface (MQI) can be ported from platform to platform
- The fact that MQSeries assures a message will be delivered and ends the application programmer's concern about delivery and differing network protocols

MQSeries should be considered as the basis on which any network is built, whether in stages or all at once. It has many building blocks that you can include as needs dictate.

For additional information, please refer to:

<http://www.ibm.com/software/ts/mqseries/>

13.4 ALE overview

Since R/3 Release 3.0, SAP has included an Application Link Enabling (ALE) facility to allow users to connect distributed applications. R/3 systems can use ALE to exchange data with other R/3 systems, R/2 5.0G systems, or, using a suitable interface, non-R/3 systems.

13.4.1 Concepts

ALE allows applications to be treated as independent, but integrated, parts of an R/3 implementation. Each R/3 system runs with its own data. There is no need for a central database; instead, data is distributed throughout the R/3 implementation as and when required. When data needs to be distributed between applications, the necessary communications functions are performed by ALE. Different R/3 implementations communicate with each other using synchronous or asynchronous RFC calls.

The ALE facility contains three layers to handle:

- Application interfacing
- Distribution processing
- Communication control

When using ALE, you'll find:

- The application is freed from communication processing.
- The structure of a message is not affected by its content.
- It's easy to change the configuration of ALE.

Facilities are provided to ensure that data is re-sent if the connection is interrupted and to monitor arrival of data.

13.4.2 Data distribution

The ALE interface is based on Electronic Data Interchange (EDI). While EDI communication uses an EDI subsystem situated between the communicating systems, ALE communicates directly using an RFC to the target system.

As in EDI, an Intermediate Document (IDoc) is used to contain distributed data. An IDoc is a structure of segments, such as headers and collections of data fields. SAP has predefined types of IDoc for use with ALE, but you can modify an existing IDoc or define your own.

13.4.3 Outbound processing

An application set up to send data to another system sends a generic IDoc to ALE. The IDoc contains the data to be sent together with all the information ALE needs to send the data. ALE performs the following processing:

- ALE decides where to send the data. This could be one system, several systems, or none at all. For each recipient, selected data from the generic IDoc is placed in a new IDoc called the *communication IDoc*.
- If necessary some segments are filtered out or some fields are converted.
- If the receiving system is a different version level (for example, 3.0D rather than 3.0E), the communication IDoc is generated in the correct version style.

- The method and timing of the data transmission is determined.
- The number of IDocs to be sent together is determined.

13.4.4 Communications

If the data must be sent immediately, generally an asynchronous RFC is used. Sometimes when remote systems need to communicate directly in read mode, a synchronous RFC or a CPI-C program must be used.

If data needs to be sent at specified times, the data transmission can be scheduled as a job.

13.4.5 Inbound processing

When the data is received at the target system, ALE processes the incoming IDoc. Inbound processing includes regeneration of any IDoc that is not in the correct version style, as well as any necessary segment filtering or field conversion. ALE then posts the IDoc to the application depending on the ALE settings for that IDoc. These settings identify the target application and the triggering of workflow.

The settings also differentiate between single IDoc processing and bulk IDoc processing.

13.5 How the MQSeries link for R/3 works

The MQSeries link for R/3 is an interface that provides a simple way of integrating the customer's existing R/3 applications with other applications running in environments that are not supported by R/3.

The MQSeries link for R/3 works with the ALE layer of R/3 to transmit IDocs into and out of the R/3 system. It uses MQSeries messages and queues to carry the information. It extends the scope of the business by linking R/3 applications to any other application that can be accessed through MQSeries, even when those applications require different data formats. See Figure 232.




Figure 232. MQSeries link for R/3 and ALE

There are two logical aspects to the MQSeries link for R/3: outbound and inbound. Outbound is concerned with sending information from R/3, and inbound

is concerned with receiving information into R/3. The two types flows are shown in Figure 233.




Figure 233. Inbound and outbound servers

Outbound

Outbound is when an R/3 application sends information to another application. The R/3 application passes a single transaction to the MQSeries link for R/3 as a set of one or more IDocs. The outbound server builds messages to an MQSeries queue or queues.

If the data being sent is in a format that needs to be converted, an exit handler gives access to a user exit before the message is put on MQSeries.

Inbound

Inbound is when an R/3 application receives information from another application. The inbound server retrieves messages from MQSeries and converts them into IDocs to pass to the R/3 application. If the data is not in IDoc format, an exit handler in the inbound server gives access to a user exit to convert the data received in the message from MQSeries.

The application information that the MQSeries link for R/3 passes to the R/3 system is always in the SAP IDoc format.

13.5.1 Components of the MQSeries link for R/3

The MQSeries link for R/3 components include:

- **Outbound server:** This component accepts IDocs from the local R/3 system and sends them to other applications connected by MQSeries.
- **Inbound server:** The MQSeries link for R/3 component that receives information from applications connected by MQSeries and passes the information to the local R/3 system.

- **C source header file:** A header file containing structure definitions required to process MQSeries link for R/3 message data.
- **Sample exits:** C source code for sample exits that convert IDocs to and from simple application specific formats.
- **Sample MQSC and ini files:** An MQSeries command file and an initialization file each for the inbound and outbound server.
- **Administration panels:** The panels in R/3 that allow an administrator to configure the MQSeries link for R/3 software.
- **Standard R/3 application:** An R/3 application in SAP R/3 used to test the basic configuration of the MQSeries link for R/3.

13.5.1.1 Outbound sequence of events

The following diagram and list show the sequence of events when R/3 executes a transaction consisting of one IDoc or multiple IDocs, outbound from an R/3 application. The numbers in Figure 234 correspond to the numbers in the text that follows.




Figure 234. Outbound sequence of events

1. An R/3 application program creates one or more IDocs representing a single transaction.
2. The RFC component handles the connections to the MQSeries link for R/3 and provides control information required to handle the transaction and send it to the specified destination.

The RFC component uses the RFC table (RFCDES) to identify the connection type and physical destination of the transaction. The table contains entries for all MQSeries and non-MQSeries destinations and is maintained using the standard R/3 transaction sm59.

If the RFC destination program ID matches that of a running outbound server, the MQSeries link for R/3 is used.

3. The MQSeries link for R/3 accesses the messaging queuing (MQ) ABAP function to retrieve additional information needed for an MQSeries destination.

The MQ ABAP function uses the RFC destination name as a key to access the MQ table (MQDES) containing the additional information. This includes the target queue name, target queue manager, logged-on user ID, and whether an outbound user exit should be called. The MQ table is maintained through the MQSeries link administration panels, accessed through R/3 transaction sm59.

4. The MQSeries link for R/3 outbound server uses MQSeries destination information (from the control information) and the transaction data to construct the MQSeries message. Each message contains one or more IDocs for a specific destination.

If the control information indicates that data formatting or conversion is required, the outbound exit handler calls a user exit to perform the translation.

The outbound server sends the messages to the specified MQSeries queue. All the messages are put under syncpoint. They are treated as a single logical unit of work.

13.5.1.2 Inbound sequence of events

Figure 235 and the following list explain the sequence of events when the MQSeries link for R/3 receives an inbound transaction from MQSeries for an R/3 application.




Figure 235. Inbound sequence of events

The numbers in Figure 235 correspond to the numbers in the text that follows.

1. Once the inbound server has been started, it polls the MQSeries queue for new messages. An application sending information to an R/3 application places a message on an MQSeries queue. The MQSeries link inbound server retrieves the message from the MQSeries queue, under syncpoint. The MQSeries link for R/3 creates IDocs from the incoming transaction data. If necessary, MQSeries link for R/3 calls a user exit to convert the data to IDoc format.
2. The inbound server requests a transaction ID from the R/3 system.

The MQSeries link inbound server keeps track of inbound transactions and manages any messages by making an entry in an internal transaction store queue.

3. The MQSeries link for R/3 header in the incoming message data contains information about the R/3 application that is to receive the transaction.

If the information about the remote R/3 system was not specified within the R/3 transaction SM59, the link header outbound message does not contain all the required destination information. To connect to the remote system, MQSeries link for R/3 uses the default values specified in the initialization (.ini) file for the inbound server.

The RFC program waits for an inbound message from the MQSeries link for R/3. When a message arrives, the RFC program executes the R/3 function module INBOUND_IDOC_PROCESS.

4. After the transaction is transferred successfully to the R/3 system, the MQSeries link for R/3 executes an MQSeries commit to delete the inbound message. The MQSeries link for R/3 also deletes its entry from the internal transaction store queue.

13.5.1.3 Message format

The outbound server constructs an MQSeries message to send R/3 application information. The R/3 information is always in the form of IDocs.




Figure 236. MQSeries message format and IDocs

The numbers in Figure 236 correspond to the numbers in the explanation that follows.

1. This is the generic MQSeries message format. The header contains MQSeries control information for processing the message and is followed by message data.
2. The message data begins with a link header that is used by the MQSeries link for R/3 software. The link header contains information to identify the destination of the message.

The link header is followed by application data. For inbound transactions, the application data field is in IDoc format.

For outbound transactions, the application data field should be converted to the format required by the target application. However, if the target application is another R/3 system, this field takes an IDoc format.

3. In an MQSeries message, the application data consists of one or more IDocs stored contiguously with no separators between. The maximum size of the application data is specified in the ini file, which is subject to a maximum of 100 MB for MQSeries messages.
4. Each IDoc has a control table and one or more data tables of fixed length.

Each IDoc is identified by a unique IDoc number that is stored in a field in both the control table and data tables. A new IDoc within the same message is identified by a new IDoc number contained in the control table. Figure 236 shows a message with two IDocs. The first has two data parts and the second has one data part. Sample C source code for handling messages in this format is provided in the sample exits.

13.5.1.4 Exit handlers

Outside the MQSeries link for R/3, MQSeries data conversion provides data conversion between different computing platforms and code pages.

The MQSeries link for R/3 includes exit handlers for access to user exits that allow conversion of the structure and content of the user message data. As with all MQSeries exits, the MQSeries link exit handlers execute the user-supplied software in line, without additional message flows.

The MQSeries link software provides two sample exits: one to show the typical structure of a user exit and one to provide an example of simple data reformatting. If required, the outbound exit is called before a message is placed into an MQSeries destination queue, and the inbound exit is called after a message has been retrieved from an inbound MQSeries queue. These exits allow the format of the user data in the message to be converted.

For example, an SAP R/3 application sends invoice data to a non-SAP application. The data from the SAP application is always in IDoc format. The receiving application may expect the invoice data in a different format. An exit can be written to extract the relevant data from the IDoc in the message and build a new message with the data in the format expected by the receiving application.

Figure 237 shows the processing of exits by the inbound server.




Figure 237. Processing of exits by the inbound server

Figure 238 shows the processing flow of exits by the outbound server.




Figure 238. Processing of exits by the outbound server

13.6 Installing MQSeries link for R/3

This section tells you how to install MQSeries link for R/3 on the iSeries server. The steps in the following processes assume that MQSeries is already installed and configured on your system.

13.6.1 Installing the MQSeries link for R/3 on the iSeries server

To install the MQSeries link for R/3 on the iSeries server, follow these steps:

1. Insert the product CD into the CD-ROM drive of your iSeries server.
2. Issue the following command:

```
RSTLICPGM LICPGM(5733A22) DEV(OPT01)
```

Substitute *OPT01* with the actual name of your CD-ROM drive.

3. Verify that the MQSeries link for R/3 is installed correctly on the iSeries server. Perform these steps:

- a. Issue the Display Software Resources (DSPSFWRSC) command to see a list of the software products installed on your system.
- b. Page down until you see **5733A22**. There should be two entries.
- c. If the product is not installed properly, you see an ERROR message.
- d. In case of error, check the job log for any errors that occurred during the installation and perform any necessary actions before repeating the installation procedures.

The installation directories are:

- /QIBM/ProdData/smq
- /QIBM/ProdData/smq/MRI<nnnn>
- /QIBM/ProdData/smq/samp
- /QIBM/ProdData/smq/java

The product library is QMQLINK.

13.6.2 Three stages of configuration

The three stages of configuring the MQSeries link for R/3 include:

1. Linking ALE to the outbound server
2. Connecting the outbound server to the MQSeries
3. Linking the inbound server from MQSeries to the SAP R/3 system

These stage are shown in Figure 239.




Figure 239. Three stages of installation and configuration

13.6.2.1 Stage 1: ALE to the outbound server

The ALE layer of the R/3 system must be configured to connect to the MQSeries network through the MQSeries link for R/3. Perform the following tasks:

1. Create a logical R/3 system. For example, tell system A that there is a distributed system B.
2. Define a TCP/IP-type RFC destination. RFC is the communication mechanism between ALE and the MQSeries link for R/3.
3. Create a transactional RFC communication port that is linked to the RFC destination.
4. Maintain partner profiles for the logical R/3 system that links the logical system to the port.

The R/3 configurations are done using the SAP transactions SM59 and SPRO. SM59 defines the RFC destinations, which are, in this case, specific for the MQSeries link for R/3.

For details on configuring ALE, see the *SAP Online Documentation CD*.

13.6.2.2 Stage 2: Outbound server to MQSeries

To connect the outbound server to the MQSeries, complete the following steps:

1. Define the MQSeries queues used by the outbound server to transfer messages around the MQSeries network.

There is a sample configuration file provided with the MQSeries link for R/3. To create the queues from the sample configuration file, run the following command:

```
STRMQMQSC SRCMBR(SMQSCDEF) SRCFILE(QMQLINK/QMQSC) OPTION(*RUN)
```

2. When starting the outbound server, specify the queue manager name and queue definitions either as command line options or as initialization file entries.

3. Use an SAP gateway to connect the outbound server to the R/3 system. Specify the gateway hostname and service on the command or in the outbound initialization file.

Messages flow from the outbound server to the outbound MQSeries queue that can be defined as local or remote. This allows for transportation throughout the MQSeries network as business needs dictate.

4. You must identify the outbound queue for each destination in the destination smqDestConf configuration file.

13.6.2.3 Stage 3: Inbound server from MQSeries to R/3

To link the inbound server from MQSeries to R/3, complete the following tasks:

1. Define the queues to enable the inbound server to connect to the MQSeries network. To create the queues from the sample configuration file, run the command:

```
STRMQMQSC SRCMBR(SMQSCDEF) SRCFILE(QMQLINK/QMQSC) OPTION(*RUN)
```

2. On starting of the inbound server, specify the queue manager name and queue definitions either as command line options or as initialization file entries.

The inbound server makes a direct connection to the R/3 system. The inbound server receives messages from the inbound queue and sends the IDocs into R/3. The message header normally contains information about that R/3 system to which the IDoc should be sent.

If the message header does not contain R/3 connection information, the IDoc is sent to the default R/3 system that is specified in the initialization file for the inbound server.

13.7 Ordering and installing MQSeries link for R/3

The MQSeries link for R/3 is obtainable from IBM. For more information, contact your IBM Representative or the branch office that serves your geographic region. Or go to the Web site at: <http://www.ibm.com/software/ts/mqseries/>

The appropriate MQSeries software must also be ordered if the customer doesn't already have it installed:

- The MQSeries product for the platform on which the R/3 system runs. For example, the MQSeries for iSeries product is required for R/3 systems on the iSeries platform.
- The MQSeries product for each of the systems with which the R/3 system needs to communicate. This can be any of the more than 20 platforms that MQSeries supports.
- If the customer has more than one R/3 system, defined by more than one database server, and wants to use MQSeries to communicate between the R/3 systems, an MQSeries link for R/3 is needed for each of the R/3 systems to be connected. These can be on the same or different platforms.

Additional prerequisite software may also be needed, for example, the appropriate level of TCP/IP. Details of the prerequisite software required for any MQSeries product can be found on the MQSeries home page or obtained from IBM.

Note: To run the MQSeries link for R/3, you must have at least one R/3 system installed.

13.8 Benefits of using MQSeries link for R/3

Using the MQSeries link for R/3 allows customers to connect R/3 systems to external systems in a simple manner. Customers gain the benefits that come from using MQSeries:

- Assured, once-only delivery
- Transaction coordination
- Time-independent processing

In addition, customers are freed from a considerable development and maintenance burden, compared to building their own interfaces. The MQSeries link for R/3, working with the IDoc interface, provides a standard approach where the customer does not need to write any interface code. Only the destinations and the receiving systems need to be defined. These definitions can be changed as necessary to cope with changing business circumstances.

Using the MQSeries link for R/3 with ALE also allows for a staged migration from existing systems to R/3, with minimal disruption to existing operations.

You may decide to add more R/3 systems to the network or to more non-SAP systems, in the future, or the customer may split off part of the company or takes over other companies. Whatever the situation, MQSeries provides the simplicity and flexibility to re-shape the network or networks to support the business with minimum disruption.

Benefits to SAP consultants

Customers expect SAP consultants to help them find a solution to the problem of interfacing R/3 with existing systems. Whether the non-SAP R/3 applications run on mid-range or mainframe systems, the problem of interfacing disparate applications in mixed hardware and software environments is a complex and challenging one. It can slow down the implementation of the R/3 installation, unless a solution is found. There have been examples where the sale of an R/3 system was delayed until the customer could be satisfied that a solution existed.

Because the systems involved are almost always central to the customer's business, any proposed solution must be demonstrably reliable and secure, and ensure that data is consistent throughout the enterprise.

IBM's MQSeries is the established standard for commercial messaging systems and, since its introduction in 1994, has been proved to provide secure, reliable, messaging in hundreds of customer sites. The combination of ALE, the enhanced IDoc interface, and the MQSeries link for R/3 provides a ready made, flexible solution to the application interfacing problem. This, in turn, speeds up the sale and installation cycle for new and existing R/3 business, with obvious benefits to the SAP consultant.

Where the customer plans to replace all existing systems with R/3 applications over time, using the MQSeries link for R/3 with ALE supports a staged migration where the SAP system is seen as having the integrating role.

The simplicity of the MQSeries link for R/3, and the fact that it uses standard ALE calls, and a standard SAP administration procedure, means the SAP consultant can concentrate on helping the customer to define the Distribution Model unconstrained by interface or connectivity considerations. The customer's key technical staff, too, can concentrate on defining the links required, rather than on the method to be used. The result should be a faster, smoother, implementation cycle with a model that truly meets the customer's business needs.

With the MQSeries link for R/3, the solution can be proposed with confidence, in the knowledge that IBM will support the link and the MQSeries products that may be required.

Chapter 14. Migration to another platform

This chapter provides a brief overview of the concepts and processes involved in the migration of an existing SAP R/3 system to a different hardware platform, operating system, and database. This process should be well planned, and the SAP guidelines for system migration should be closely followed.

14.1 System copy

There are generally two types of system copies:

- **Homogenous system copy:** A homogenous system copy is performed when an existing SAP R/3 system is to be replicated onto a hardware platform that runs the same operating system and database system as the source system. This type of system copy does not require the SAP Operating System/Database (OS/DB) Migration Service.
- **Heterogeneous system copy:** A heterogeneous system copy entails copying an existing SAP R/3 system onto another hardware platform that runs a different operating system, different database system, or both. Migrating an SAP R/3 system can, therefore, be classified as a heterogeneous system copy.

14.2 Migration tools

To ensure a smooth transition to the new platform and optimum functioning of the new system, only the SAP migration tools, together with the SAP OS/DB Migration Service, should be used for system migrations. SAP provides specially designed software for migrating to and from various combinations of hardware platforms, operating systems, and database systems. For example, the software for migrating an existing SAP R/3 system running on a UNIX operating system with an Oracle database to an iSeries server is different than the software that allows migration to an iSeries server from a Microsoft NT operating system with a Microsoft SQL Server database.

SAP R/3 release restriction

SAP only supports the migration for supported database releases. Supported releases at the moment are: 3.1H, 3.1I, 4.0B, 4.5B, 4.6B, 4.6C, and 4.6D. Other systems must be upgraded to a later supported release of SAP R/3 first. SAP recommends at least six weeks of productive work on the upgraded system prior to the migration of the system.

14.3 Sap Migration Service

SAP requires that you register your migration project with the local SAP offices, who will provide details of this service, together with the relevant costs involved. The purpose of the SAP Migration service is to assist in the following areas:

- Planning the migration project.
- Ensuring optimal utilization and performance of system resources after the migration.
- Minimizing system downtime.
- Providing specialist support and practical expertise.
- Providing reports with action lists and improvement recommendations. Refer to SAP Note 0198466 for detailed information about SAP Remote Services reports.

14.4 Requirements

To minimize the risk involved in migrating an SAP R/3 system, SAP recommends that you adhere to the following guidelines:

- Only suitably certified SAP Basis consultants with knowledge of the operating system, database system, and migration tools should perform an OS/DB Migration.
- The SAP OS/DB Migration Service must be used.

The SAP OS/DB Migration Service can be ordered through SAPNet. This service consists of several remote service sessions, migration tools, documentation, and project planning guides. During these sessions, SAP provides general input and advice, assists in sizing the target hardware platform, and does performance analysis exercises on the target system upon successful migration. SAP produces a report with findings and recommendations at the end of the migration project.

The SAP OS/DB Migration Service supports the migration of an SAP R/3 production system, together with its related systems (development, test, quality assurance, and so forth). SAP recommends that you increase the project time line when migrating to a new operating system and database system at the same time. This allows additional time to adjust and tune the target system.

- Only the SAP-approved migration tools should be used to migrate an SAP R/3 system.
- The target operating system and database system combination must be certified by SAP as a valid SAP R/3 environment. SAP supported operating system and database combinations can be found at:
<http://service.sap.com/dbosplatforms>
- The target system must contain the required operating system and database system patches for the specific SAP R/3 release. The iSeries APAR list is available at: <http://www.as400.ibm.com/service/bms/support.htm>
- The required hardware for the target system must be adequately sized and procured at the beginning of the project. The migrated SAP R/3 system is tested on the target system, while the source system continues normal productive work.
- A separate project plan should be drawn up for each SAP R/3 system that will be migrated. Refer to the *SAP OS/DB Migration Service Planning Guide* for a sample project plan.

Code page setting restrictions

If the following restrictions are not observed, data inconsistencies may occur:

- The code page setting of the source and target systems must be identical.
- If either the source or the target system runs on an iSeries server, only the Latin 1 code page (ISO 8859-1) is supported by default.
- Contact SAP if you want to perform a heterogeneous system copy for a system that uses non-Latin-1 code pages.

14.5 Preparation

Prior to commencing the migration process, you should verify that:

- The source system contains enough free disk space to accommodate a copy of the source database. For the source database dump, compression rates between 5 and 10 have been achieved.
- The target system has enough disc capacity for the operating system, as well as the source database dump and the new database.
- Remote communications between the source and target systems are set up and functioning correctly.
- Performance statistics are collected on the source system and saved for a later comparison with the target system.
- Backup and recovery strategies, as well as high availability solutions, are in place for the target system.

14.6 System migration testing

To determine the feasibility of migrating an SAP R/3 production system, a simulated migration of the source system has to be completed.

Testing is typically done using a copy of the source database. OS/DB migration involves a series of iterative test runs before the final migration of the SAP R/3 production system. This is done to ensure an efficient migration process.

During each test run, the source database is exported from the source system and imported into the target system. Afterwards, data consistency between the two systems is checked and verified. System performance is also checked by executing critical SAP R/3 transactions. All interfaces to and from the SAP R/3 system have to be checked and adapted to the new hardware platform if necessary.

Note

The profile and parameter settings of the source system should not be copied onto the target system. Some of these settings depend on the operating system and the hardware configuration.

Printing and front-end operations should also be tested.

These steps are repeated until data consistency between the source and target systems can be verified and the target system functions suitably.

Errors, inconsistencies and concerns should be clearly documented during each test run. Also, measures should be in place to resolve these errors in subsequent test runs.

The migration process can be adjusted manually to achieve higher data transfer rates. However, this requires highly experienced basis and migration consultants.

Figure 240 shows an overview of the components and processes involved in migrating an SAP R/3 system.




Figure 240. Migration process overview

14.7 Final system migration

Allow at least three to six months between the first test run and the final migration of the production database. During this period, ensure proper testing of all attached devices for the SAP R/3 system, as well as backup and recovery procedures. Be sure to also conduct user testing. Once you are satisfied with the result of the test runs, the production SAP R/3 system can be migrated.

The source system is stopped, and the source database is exported and imported into the target system. The procedures followed and documented during the test runs are used to set up the target system.

For more information on SAP OS/DB Migration and services, refer to:

<http://service.sap.com/osdbmigration>

14.8 SAP R/3 license

After the SAP R/3 system is migrated, you must request a new SAP R/3 license key from SAP. The new system will not function with the license key of the source system, due to the fact that the database and operating system combination had changed. General information on SAP R/3 licensing is described in the SAP documentation *Installing R/3 on IBM AS/400*. You can find further information on license keys at: <http://service.sap.com/Licensekey>

Chapter 15. Change and transport system

The change and transport system allows you to organize your customizing and development project when implementing SAP and transport changes from one SAP R/3 system to another. *Transporting* is the act of moving change requests (which may contain program changes or customizing or configuration changes) from one SAP R/3 system to another. This is accomplished by using the Transport Organizer, the Transport Management System (TMS), and the TP and R3TRANS commands.

This chapter begins with a basic explanation of a R/3 three-system landscape, followed by an example of setting up the Transport Management System, using the Transport Organizer, the customizing process, and the TP command. Then, the chapter examines a detailed step-by-step example of creating, releasing, and transporting a change request. Finally, it discusses connections to non-iSeries environments.

For the purposes of this discussion, we assume that the naming conventions of the R/3 systems are:

- DEV (development system)
- TST (test system)
- PRD (production system)

Note: OS/400 supports the network file system (NFS) as part of the standard operating system. This enables iSeries servers to share stream file systems with non-OS/400 operating systems such as AIX. This makes it easier to transport changes between the iSeries server and other systems.

15.1 SAP R/3 in a three-system landscape

The number of SAP R/3 systems in a landscape is determined by the complexity of the implementation and the requirements of the organization. A three-system landscape provides you with a development environment, sufficient testing or quality assurances, and a separate productive system. To establish this landscape, the appropriate systems must be installed and clients created. Once established, all changes made in the development system need to be “bundled” using a change request. Then, they must be propagated over the SAP R/3 landscape using the transport system.

In a three-system landscape, a test system (TST) exists between the development system (DEV) and the production system (PRD). These systems are also referred to as the integration (DEV), consolidation (TST), and delivery (PRD) systems. A two-system landscape containing only a development (DEV) and production (PRD) system is recommended for smaller implementations where little development work will take place.

It is also possible to implement SAP R/3 in a one-system landscape, consisting only of a production system (PRD). If you are considering this option, you should carefully consider the advantages and disadvantages that are discussed in Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43.

There are many ways to distribute a three SAP R/3 system landscape, such as TST, DEV, and PRD over iSeries machines. One way is to have all three on a

single iSeries server containing two logical partitions, one for DEV and TST, and a second partition for PRD. Another way is to have each R/3 system run on a separate iSeries server. The third possibility is to run two SAP R/3 systems on one iSeries server and run the third SAP R/3 system on another iSeries server, but it is the same at the conceptual level. These options are explored in depth in Chapter 3, “SAP R/3 system landscapes on iSeries” on page 43.

We recommend you isolate your SAP production system as far as possible from your test and development systems. There are several important disadvantages of running your production system on the same iSeries server as your development or test systems:

- Development activity may adversely impact the response time for the production users
- Testing that requires major amounts of system resources, such as a trial conversion load of master data, will impact production users
- Training classes run in a development or testing client may degrade performance in the production R/3 system
- OS/400 PTFs once applied will impact all environments: development, test, and production. Ideally all changes to the production system should be first tested in a non-production environment

Each SAP R/3 system in the landscape is installed separately using R3SETUP as described in Chapter 5, “Installation and configuration” on page 67. A key /sapmnt/trans directory used by the Transport Management System is located on only one of the iSeries servers in the landscape. This is referred to as the *global transport directory*. It consists of all the transport files that contain the changes that are transported between R/3 systems.

Typically, the DEV R/3 system is installed first so the link is created from /usr/sap/trans directory on this iSeries server to the /sapmnt/trans directory. When additional R/3 systems, such as TST and PRD, are installed at a later date on different iSeries400 servers, the Change R/3 Shared Location (CHGR3SHLOC) command must be used to specify the server that contains the global transport transport directory. This then creates links using QFileSvr.400 file system to point back to the server with the global transport directory.

For example, DEV is installed on host DEVSYS, which contains the global transport directory, TST is installed on host TSTSYS, and PRD is installed on host PRDSYS. After the CHGR3SHLOC command is run for TST and PRD, with DEVSYS specified as the hostname for the global transport directory, the links shown in Table 33 are created.

Table 33. Links for the global transport directory

R/3 system	Directory	Link
DEV	/usr/sap/trans	/sapmnt/trans
TST	/usr/sap/trans	/QFileSvr.400/DEVSYS/sapmnt/trans
PRD	/usr/sap/trans	/QFileSvr.400/DEVSYS/sapmnt/trans

The creation of development objects, such as ABAP programs, and customizing changes is completed in the development client in the DEV system and moved to the TST system. This allows you to perform testing, training, quality assurance,

and user acceptance prior to releasing the changes to the production environment. This three-system landscape also allows development teams to work on further development phases without impacting the TST system.

Note

Normally we recommend you use the /usr/sap/trans directory on the production system. Since PRD system may not exist at the beginning, the /usr/sap/trans directory can be moved to the PRD system later. The advantages are the higher performance for imports into the PRD system, which brings reduced downtime for this (for example, support packs), and the higher availability of the trans-dir on the PRD system as opposed to the DEV system.

As shown in Figure 241, the standard SAP R/3 client 001 is copied to a new client 200 where the development activity takes place. A client 210 is created as a sandbox or playpen that is regularly refreshed from client 200. Upon completion of a development phase, the change requests generated in client 200 are exported and then imported using TMS to a clean client 300 in the TST system that is used for continuous testing. Typically, client 300 is copied to a new client 310 in the TST system to perform training of end users. When the testing is verified, the changes are imported to client 410 in the PRD system for pre-production integration testing. Once live, further changes are imported into client 400, the production client, and then client 410 can be deleted.




Figure 241. Client structures in a three-system landscape

The appropriate client change settings must be specified for each client and also for the system as a whole in the system change options. These settings define the types of objects that can be modified and whether changes made in a particular client will be saved in a change request to enable them to be transported.

When an SAP R/3 system is installed, it comes standard with three clients: 000, 001, and 066. These reference clients should not be deleted under any circumstances and should not be used for customizing or development activities, which should be done in a copy of client 001.

The Transport Management System provides a central administration point for transports, called the *transport domain controller*. It defines the R/3 systems and the transport paths between them. The terminology *transport group* is used to refer to those R/3 systems that share a common transport directory. *Transport domain* refers to one or more transport groups. In our example, one transport domain DOMAIN_R01 consists of three systems R01 (the transport domain controller), TST, and PRD as shown in Figure 242.




Figure 242. Transport domain for a three-system landscape

15.2 Global transport directory

The Transport Management System uses the subdirectories under the /usr/sap/trans directory for storing objects that have to be transported from one SAP R/3 system to another SAP R/3 system. The /usr/sap/trans directory is actually a symbolic link to one physical directory called /sapmnt/trans. The /usr/sap/trans directory also contains a config subdirectory that contains all of the SAP R/3 system definitions.

If you have all three SAP R/3 systems on a single iSeries server, each SAP R/3 system's /usr/sap/trans directory has a symbolic link to the /sapmnt/trans directory on that system. If you have multiple iSeries servers for three SAP R/3 systems, you have the /sapmnt/trans global directory on one of the iSeries servers, for example, the iSeries server for DEV. Other iSeries servers access this shared directory, which is referred to as the *global transport directory*, through the QFileSvr.400 file system.

You must create the QFileSvr.400 subdirectory for the host names of each iSeries server for which access is required by using the command:

```
MKDIR DIR('/QFileSvr.400/<hostname>')
```

Important

After an IPL, these directories no longer exist. Therefore, the MKDIR commands should be added to the OS/400 startup program of the iSeries server.

The R/3 installation tool, R3SETUP, has a menu option for changing the location of the /sapmnt/trans directory that executes the CHGR3SHLOC command. The TCP/IP hostname of the iSeries server with the global transport directory must be entered.

Figure 243 shows an example of using the /sapmnt/trans shared directory when having multiple iSeries servers in a three-system landscape.




Figure 243. Example of the /sapmnt/trans shared directory

In this example, the /sapmnt/trans directory exists on only one iSeries server, and a symbolic link /usr/sap/trans points to the /sapmnt/trans directory on the DEV system. As we mentioned earlier, we recommended you move the /sapmnt/trans directory to the PRD system once the PRD system is up and running. To understand fully how to work with the QFileSvr.400 system, refer to *OS/400 Integrated File System Introduction*, SC41-5711.

Important

The OS/400 user profiles <SID>OFR and <SID>nn (where <SID> is the SAP System Identifier and *nn* is the instance number) must exist on the servers that share the common transport directory. They should be created using the CRTSAPUSR *OFR and CRTSAPUSR *SIDINST commands.

The passwords of QSECOFR, <SID>OFR, and <SID>nn must be the same on all servers sharing a common transport directory. This is a requirement for access via the QFileSvr.400 file system.

Refer to SAPNet Note 67213 for requirements for transports between R/3 systems on different iSeries servers.

For three-tier system landscapes, where you have separate database and application iSeries servers, we recommend you locate the global transport directory on the database server.

15.2.1 Setting up the Transport Management System

The Transport Management System must be set up after the installation of an R/3 system. From R/3 Release 4.5A onwards, this is performed using the Transaction STMS. The `/usr/sap/trans/bin/TP_DOMAIN_DEV.PFL` file, where *DOMAIN_DEV* is the name of the transport domain, is automatically updated by STMS with the required parameters for each system as shown in Figure 244.

```
#TMS:0048:DOMAIN_DEV
#
TESTIMPORT  = 0
TRANSDIR    = /usr/sap/trans/
#
DEV/DBHOST  = DEVSYS
DEV/DBNAME  = dev
DEV/DBTYPE  = db4
#
PRD/DBHOST  = PRDSYS
PRD/DBNAME  = prd
PRD/DBTYPE  = db4
#
TST/DBHOST  = TSTSYS
TST/DBNAME  = tst
TST/DBTYPE  = db4
```

Figure 244. Example of the `TP_DOMAIN_DEV.PFL` file

Prior to R/3 Release 4.5A, the TP command used with the change and transport system requires parameters within the *global transport parameter* (TPPARAM) file. Therefore, prior to R/3 Release 4.5A, you need to check that the TPPARAM file in the `/usr/sap/trans/bin` directory includes an entry for each SAP system in the transport domain with the following information, as shown in Figure 245.

To edit the TPPARAM file, issue the command:

```
EDTF '/usr/sap/trans/bin/TPPARAM'
```

```
#####
# Global parameters                                     #
#####
transdir    = /usr/sap/trans/
testimport  = 0
#####
# Specific parameters                                   #
#####
DEV/dbhost  = DEVSYS
DEV/dbtype  = db4
TST/dbhost  = TSTSYS
TST/dbtype  = db4
PRD/dbhost  = PRDSYS
PRD/dbtype  = db4
```

Figure 245. Example of the `TPPARAM` file

DEV, TST, and PRD in Figure 245 represent <SID>. DEVSYS, TSTSYS, and PRDSYS are the TCP/IP host names of the database server on which the <SID> runs.

As of release 4.5B, the TPPARAM should no longer become maintained. A link to TP_DOMAIN_DEV.PFL should be set up to ensure that the data doesn't run out of sync. This is useful because every TP from the OS-level, where "PF=.." is not explicitly used, uses TPPARAM. Therefore, as of release 4.5B, perform the following actions:

```
DLITF /usr/sap/trans/bin/TPPARAM
ADDLNK OBJ('/usr/sap/trans/bin/tp_domain_dev.pfl')
NEWLNK('/usr/sap/trans/bin/TPPARAM')
```

The SAP R/3 WRKLNKSAP command provides options for changing, copying, removing, displaying, renaming, moving, printing, and working with stream files (Figure 246).

```

Work with Object Links (WRKLNKSAP)

Directory . . . . : /usr/sap/trans/bin

Subset: *

Type options, press Enter.
  2=Change  3=Copy  4=Remove  5=Display  6=Print  7=Rename
  9=Edit authority 11=Move  12=Work with

Opt  Object link      Type      Owner      Size  Date
   .          *DIR      QSECOFR    53248  Oct 18 14:34
  ..         *DIR      QSECOFR    77824  Oct 13 18:04
  DOMAIN.CFG *STMF      DEV01       705    Oct 14 20:20
  TP_DOMAIN_DEV.BAK *STMF      DEV01       879    Oct 14 20:20
  TP_DOMAIN_DEV.PFL *STMF      DEV01       903    Oct 14 20:20
  2  TPPARAM    *STMF      QSECOFR     401    Oct 18 14:34

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F11=Alt.view  F12=Previous level  F13=Repeat
F14=Sort by column  F16=User options  F17=Top  F18=Bottom  F21=System command

```

Figure 246. Work with Object Links (WRKLNKSAP)

15.2.2 Defining the transport domain and transport routes

The Transport Management System (TMS) is used to define the R/3 systems that are part of the transport domain. To create the transport domain, you log in to client 000 as user DDIC and choose **Tools-> Administration-> Transports-> Transport Management System**, or transaction code **STMS**.

The basic settings and standard configuration are generated. The transport domain controller is set as the current system. You can then proceed to define the transport routes between systems in the transport domain. For systems that are not yet installed, they can be defined as "virtual" systems initially and are changed later to real systems.

Figure 247 shows the setup of TMS after the installation of the development system for a three-system landscape where the test and production systems were defined as virtual systems. Icons indicate the type of system.




Figure 247. TMS setup of three-system landscape showing transport routes

Once the test and production systems are installed, they can apply to be accepted into the transport domain via RFC connections to the transport domain controller. The changed configuration is then re-distributed to all systems in the transport domain again via RFC from the domain controller.

The Transport Management System setup before R/3 Release 4.6A created the following table entries:

- **Known SAP systems in the landscape** (table TSYST): This defines all of the systems in the system landscape. SAP systems that are yet to be installed can be created as “virtual” systems in TMS and be fully defined later.
- **Delivery routes or paths** (table TASYs): This defines the systems for which transports are delivered automatically after successfully importing them into the consolidation system.
- **Consolidation routes or paths** (table TWSYS): This allows for transporting SAP standard repository objects into the consolidation system.
- **Transport layers** (table DEVL): This defines the transport path from the development system to the consolidation system.

To accept the changes and imports from SAP directly such as via SAP Support Packages (prior to release 4.5, referred to as “hot packages” and “legal change patches (LCPs)”), TMS also includes an entry for a delivery system with a <SID> of SAP.

From R/3 Release 4.6A onwards, the four tables are obsolete and replaced by the tables:

- System List (table TCESYST)
- Deliveries (table TCEDELI)
- Transport Layers (table TCETRAL)

The above tables contain a field for the Version number which allows different TMS configurations to be stored and retrieved as needed. The TMS version information is stored in table TCEVERS.

15.3 Customizing and development process in an R/3 system

Customizing is an SAP term that means making appropriate changes to the SAP-supplied system that parallels your company's needs. The objects affected by these changes to the development system are moved to the test and production systems through the Transport Management System. The term *objects* may mean a single ABAP program, a newly defined table, individual table entries, or even an entire client. The term *modification* is used to refer to changes made to SAP objects such as reports or tables in the R/3 repository.

All objects and data within SAP are classified as either *client dependent* or *client independent*. Client dependent objects are valid only for an individual client and have the client number (MANDT field) as part of the key for the table. Client independent objects cross all clients (a change to this table affects all clients in the system). The client dependent tables have the first field called MANDT. This is a key field that designates an entry to a specific client.

15.3.1 Client attributes

Client attributes allow development efforts in each client to be controlled and are assigned for each client in an R/3 System. To check client attributes, select transaction **SCC4**. Or select **Tools-> Administration-> Administration-> Client Administration-> Client Maintenance**.

The client attributes determine how customizing or client-dependent changes are recorded in each client and whether client-independent changes can be made when logged into the client.

The client-dependent attributes include:

- **Changes without automatic recording:** Deactivates automatic saves to change requests and allows customizing table entries to be changed (default). Changes to customizing and table entries can be included in a task manually or through table view selection.
- **Automatic recording of changes:** Allows customizing table contents to be changed and automatically included in a task.
- **No changes allowed:** Prevents users from making changes to customizing entries (only display mode).
- **No transports allowed:** Customizing table entries can be changed but cannot be included in a change request (this option can be used to isolate a demonstration or training client).

Customizing transactions are used by configurators to maintain customizing using the Implementation Guide (IMG). System parameters, table views, and so on are configured in the customizing client and saved in assigned tasks. Multiple tasks can be included in change requests. A change request can be created manually by the project leader or automatically by the system. This change request is transported to the test and then production systems.

Change requests are numbered in the format <SID>K9xxxxx, where <SID> is a three-character system identifier and XXXXX is a SAP supplied number. An example is DEVK900065. The tasks under this change request typically follow sequentially (that is, DEVK900066 would be the change request number and DEVK900067 would be the task). More than one task may be created under a change request.

At the start of a development project, the project leader has the option of creating a change request manually and assigning the project team members to it. Another option is to let SAP R/3 automatically create the change request. The Transport Organizer also creates a task for each project member. During the customizing process, the team members assign their changes to a task number.

The task contains all of the objects that a team member is currently processing in the development project. As the team members finish their work, they release their tasks. The task objects are passed to the change request. When all team members release their tasks, the change request can be released and exported.

15.3.2 Tasks and change requests

The Transport Organizer provides a hierarchical view of all the change requests for a specific user. To call the Transport Organizer, select **Tools-> Administration-> Transports-> Transport Organizer**. Or, use transaction codes **SE09** or **SE10**.

Tasks are the lowest level of the transport system. All saved customizing and configuration changes are copied to tasks. Each task is owned by an individual user master record. The task contains the actual transport objects that may consist of table entries or objects such as SAPscripts. The owner of the task must *release* the task to the change request. Prior to releasing a change request, all tasks under that change request must be released. Once released, a task can no longer be modified. It is now ready to be imported into the next system in the transport landscape.

In the definition window or transport requests, you can view the results of a transport by choosing the menu option **Goto-> Transport Logs**.

15.3.3 Transporting change requests

Once a change request is released and exported, it needs to be transported into the test system. Transporting is the act of moving change requests into another system. This is accomplished via the transaction STMS, which calls the SAP R/3 command, TP, in the kernel library. You can also run the TP command manually on the iSeries command line, although this is not recommended. STMS ensures that the steps in exporting and importing objects are performed in the correct, sequential order.

After change request is released, the objects it contains are extracted from the database or the source system and stored in files at the operating system level under the /usr/sap/trans global transport directory. The files listed in Table 34 are created.

Table 34. Transport files created on release and export of change request

Directory	File	Contents
/usr/sap/trans/data	Rnnnnnn.<SID>	Data file
/usr/sap/trans/data	Dnnnnnn.<SID>	Dictionary file (only for changes to dictionary objects)
/usr/sap/trans/cofiles	Knnnnnn.<SID>	Control file

At the same time, the change request number is added to the transport buffer of the test system (TST). In the import phase, the requests in the transport buffer are applied to the TST system. At the same time, the requests are added to the buffer of the production system (PRD). Once the change is tested in the test system, it can then be imported into the production system.

15.4 Example of creating a repository object change request

The following example shows how to create a repository object (for example, an ABAP program), save it in a system generated change request, release the assigned task and change request, and import the ABAP program into the test system. Again, we assume that the system names in the landscape are DEV, TST, and PRD.

To perform this example, complete the following steps:

1. Click **Tools-> ABAP Workbench-> Development-> ABAP Editor**. Or, call transaction **SE38**.

On the ABAP Editor window, give your program a name (for example, ZTMSPGM). Click the **Create** button. The ABAP Program Attributes window appears.

Note: You must first register your R/3 user ID with SAPNet/OSS as a Developer and enter the Developer key assigned by SAP when prompted. Thereafter, you are not prompted to enter the Developer key each time you make a change or create a new object.

2. Fill in a short description. Under attributes, fill in Type with *Executable program*, and Application with *Basis*. Click the **Save** icon. At this point, the Create object catalog entry appears.
3. Select the appropriate development class (at least one development class must be created using transaction SE80 before change requests can be created) and click **Save**. The Change Request Query window is shown.
4. Click the **Create request** button. The Create Request window appears. Enter a short description, and click **Save**. You return to the Change Request Query window, and the request number is automatically created for you. Press Enter.
5. You return to the ABAP Program Window. Click the **Source code** button.
6. The ABAP editor is shown. Enter a sample program as shown in the following example and click **Save**. The program is now in the Change Request mode.

```
report ztmshpgm .  
write: / 'This is a test of SAP TMS'.
```

15.4.1 Transport Organizer

The Transport Organizer can be reached through transaction **SE09** or **SE10**:

1. Type transaction SE10. The Transport Organizer window is shown.
The Modifiable and Released options are both selected by default. To select only requests that have not yet been released, deselect the **Released** box and click the **Display** button. All requests and tasks that are still modifiable and created under your user ID are now shown. You can only release tasks or change requests that are owned by your ID.
2. Click the plus icon next to the change request number that was created in the previous steps. The task number now appears.
3. Click the task and choose **Release**. You are prompted to enter any pertinent documentation for the task. Click **Save** and then the back arrow. The task is now released to the change request and is highlighted in blue.
4. Release and export the change request to the TST system buffer in /usr/sap/trans. Again, click the change request and choose **Release**.
5. A message indicating that your objects are now being exported appears. Press Enter.
6. You return to the Transport Organizer where your change request is reflected under the "Released" category. To view the export logs to check the return codes of the export, click the change request. Then choose **Goto-> Transport Logs**.

Double-click the **Export** and **Test** import categories to view logs.

15.4.2 Strategies for importing requests

Depending on your system landscape and the requirements of your project team, you can choose from three methods for controlling the import of change requests into the consolidation and delivery systems. Refer to the SAP documentation online help under **Change and Transport System - Overview-> Transport strategy in the CTS** for a more detailed explanation of the different options available.

The three methods are:

- Import all requests in the import queue (mass transports)
- Import all requests in a project
- Import individual requests in the import queue (single transports)

Refer to the SAP documentation online help under **Transport Management System-> Performing Transports** for more details on how to proceed in each case.

15.4.3 Importing single requests using STMS

After successfully releasing and exporting the change request, it is now time to import it into the TST system. Follow these steps:

1. Sign on to the SAP R/3 target system.
2. Call Transaction STMS.

3. Select the **truck** symbol.
4. The import overview appears.
5. Position the cursor on the SAP system for which you want to import requests. Choose the **glasses** symbol.
The import queue of the selected SAP System appears.
6. Position the cursor on the request that you want to import.
7. Choose the **truck** symbol.
The dialog box *Import Transport Request* appears. The box displays the request you chose and the import target.
8. After you have made your choices for settings, choose **Continue**.
9. The Start Import dialog box appears. Check and confirm the information on this dialog box. The box shows you which system and which target client you have chosen for the import.
10. After the import is finished, you should check the import history for the return codes of each of the import steps. Select **Goto-> History-> Import History**. From the import history, you can go to the transport logs to check the reason for any errors.

15.4.4 Importing single requests using TP

To import a single request into the TST system using TP, follow these steps:

1. Sign on to the iSeries server of the target system (TST) as <SID>OFR (TSTOFR in this example). Make sure the kernel library for SAP is in your library list (normally R3xxxOPT where xxx is the R/3 release). This will always be the case if you log in under the <SID>OFR user profile.
2. Change your current directory (using the CHGCURDIR command or CD) to the location of the TP parameter file TPPARAM by typing:

`cd /usr/sap/trans/bin`
3. Alternatively, select the **Change** and **Transport System** options from the R/3 Main Menu. Then change current directory to /usr/sap/trans/bin.
4. Type the command:

`tp showbuffer TST`

The TST import buffer appears. The transport number is shown in Main Import. Your change request is listed in the list of imports in the buffer. Press Enter to return to command prompt.
5. Type the following command, where xxxxx is your change request number:

`tp import DEVK9xxxxxx TST`

You receive an indication that the TP import completed successfully.
6. Go back to SAP transaction SE10 and make sure the box labeled **Released** is selected. Press Enter.
7. Your change request is listed in the Released category. Click the change request and select **Goto-> Transport Logs** from the pull-down menu view the log.

Congratulations! You just completed your first release and import of a change request.

15.5 Transporting objects between SAP systems on different host systems

In the following example, a program was exported and released for transport (using SAP transaction SE09) on the machine with host name SYSNM001 and SAP R/3 source system R01. The transport request R01K900240 contains cofiles member K900240.R01 and data member R900240.R01. The files must be transported to the machine with host name SYSNM002. After the export on the source machine, you can start the import on the target machine.

15.5.1 iSeries-only environments

In a homogeneous environment (iSeries source and iSeries target), the /sapmnt/trans directory is on one machine (SYSNM001). The /usr/sap/trans symbolic link and QFileSvr.400 make the contents of the directory available for the other machine (/QFileSvr.400/SYSNM001/sapmnt/trans).

15.5.2 Heterogeneous environments (NFS-connected machines)

When using Windows NT, you don't have to purchase NFS. You can use iSeries NetServer or iSeries QNTC file system.

Attention

Windows NT user names may be longer than 10 bytes. Therefore, you may need a guest user profile to obtain the authority from Windows NT.

In a heterogeneous environment where NFS is available (OS/400 source and AIX target), one machine (in our example, source machine SYSNM001) has the /usr/sap/trans directory and exports this directory to the other machine:

1. Start NFS server daemons.
2. Export the directory tree under the '/usr/sap/trans' path name to SYSNM002:

```
STRNFSSRV *ALL
CHGNFSEXP OPTIONS('-I') DIR('/usr/sap/trans') HOSTOPT((SYSNM002))
```

3. The target machine SYSNM002 mounts this directory:

```
ADDMFS TYPE(*NFS) MFS('/usr/sap/trans') MNTOVRDIR('/sapmnt/trans')
```

In this example, both the export and import function are shown based on OS/400 commands.

In the case of a heterogeneous environment, replace the non-OS/400 side with the import command that is supported by that system. If the global transport directory is located on the OS/400 system, proceed as follows:

1. Export all subdirectories of /sapmnt/trans/ with the "Data file code page *ASCII" option, except for the directory /data when you use the EXPORTFS command:

```
HOSTOPT(SYSNM002 *ASCII *ASCII)
```

2. Export the /data subdirectory as BINARY when you use the EXPORTFS command:

```
HOSTOPT(SYSNM002 *BINARY *ASCII)
```


Note: NFS can also be used in homogeneous iSeries environments. However, we recommend that you always use QFileSvr.400 because it performs better.

For more information about OS/400 Network File System support, see *OS/400 Network File System Support*, SC41-5714.

15.5.3 Heterogeneous environments (no NFS available)

In heterogeneous environments where the NFS function is not available to connect both file systems (for example, Windows NT and OS/400), the Windows NT machine and the iSeries server both have a /usr/sap/trans directory. The files in the /usr/sap/trans/cofiles and /usr/sap/trans/data directories must be transported to the OS/400 integrated file system using FTP.

We want to sign on to the iSeries server and move the exported data from the Windows NT system to the OS/400 integrated file system. Use the following steps:

1. Sign on to iSeries server SYSNM002 and start the FTP command to the remote host SYSNM001.
2. Log on to SYSNM001 (user ID and password).
3. Set the name format 1 (IFS name format).
4. Change the local and remote current directory.
5. FTP the cofiles member (subcommand get).
6. FTP the data member (subcommand get, binary transfer type).

```
ftp SYSNM001
userid: yyyyyy
password: xxxxxx
namefmt 1
cd /usr/sap/trans/cofiles
lcd /usr/sap/trans/cofiles
get K900240.R01
cd ../data
lcd ../data
binary
get R900240.R01
quit
```

Note: The cofiles member is transported using ASCII - EBCDIC translation, while the data member is transported in binary format.

For the cofiles file (text), the CPYTOSTMF command must use the ENDLINFMT(*LF) parameter.

For the data file, the CPYTOSTMF command must use the ENDLINFMT(*FIXED) parameter.

7. Fix authority after FTP:

```
CHGAUT OBJ('/usr/sap/trans/cofiles/K900240.CUS') USER(R3GROUP)
DTAAUT(*RWX) OBJAUT(*OBJMGT *OBJREF)
```

For more information about the FTP command, see the *TCP/IP Configuration and Reference*, SC41-5420.

15.6 Testing the Transport Management System

To test the Transport Management System, export one of your own developed programs on the R/3 development system (DEV). You receive a transport request number (DEVK9xxxxx). Sign on to the iSeries containing the R/3 test system (TST), and issue the following commands using user ID <SID>OFR, where <SID> is the target R/3 system.

```
cd '/usr/sap/trans/bin'
tp 'ADDTOBUFFER DEVK9xxxxx <SID>'
tp 'import DEVK9xxxxx <SID>'
```

DEVK9xxxxx is your transport request number (for example, DEVK900240), and <SID> is the system identifier of the target system (for example, TST).

Note: The user ID performing the TP command must be authorized to the source database R3<SID>DATA library, where <SID> is the source system. The user ID must also have the supplemental group of <SID>GROUP of the source system in their OS/400 user profile.

To check that communication between the two systems is working correctly, use the following command:

```
tp 'connect <SID>'
```

If the connection is successfully established, the following message is displayed:

```
Connection to database of <SID> was successful
```

Chapter 16. Performance management

This chapter introduces you to a performance management methodology for optimally running the SAP R/3 application on the iSeries server. It shows you how to review the performance of an iSeries server running SAP R/3 applications. This includes:

- Introducing the more important aspects that influence performance in the iSeries environment.
- Monitoring and tuning iSeries system performance with OS/400 commands.
- The concepts of monitoring and optimizing SAP R/3 performance with SAP R/3 tools.

SAP R/3 is a client/server ERP application that extensively uses SQL on the iSeries. Therefore, this chapter also emphasizes the DB2 UDB for iSeries database performance monitor. The recommended approach is based on the ongoing measurement and analysis of system performance to enable you to perform capacity planning and specific performance problem analysis.

16.1 Introduction to performance

This section presents an overview of some key factors that affect system, application, and network performance. A brief overview of iSeries work management is also included.

16.1.1 Performance concepts

This section introduces you to the subject of system performance. It enables you to understand the performance and tuning concepts that are discussed later in this redbook.

16.1.1.1 Queuing concepts

The work of a single job, or the transactions within that job, is comprised of several tasks. The invitation to perform the work required by the task is called a *request*. The requested work is performed by the *server*. The time taken to complete the requested work of the task is called *service time*.

Queuing is a concept that applies equally to requests waiting for computer resources and to people waiting in line at the supermarket or bank. In general, the length of time it takes for a work request to be completed, whether it's a request to complete the purchase at the supermarket counter, complete a transaction at the bank, perform a disk I/O operation, or use the CPU, depends on three primary parameters:

- The number of “waiters” in the line ahead of a new request
- The number of servers responding to requests
- The service time to complete the request given to the server, which is a function of the speed of the server and the amount of work to do

Consider a service point where certain tasks are performed for requestors of service. Generally, requests for service are responded to in the order in which they arrive. Therefore, those arriving first are responded to first and leave the service point first.

If the rate of arrival of requests is greater than the rate at which they leave after being serviced, a queue is built at the server. The total response time to have a request serviced is the sum of:

- The time spent in the queue waiting to be serviced
- The time it takes the server to perform the requested work

When the queue grows longer, more time is spent waiting in the queue, and the total time taken for a request to be serviced becomes longer.

The following basic principles govern queuing:

- A single server can service only one request at a time.
- Multiple concurrent requests are queued for service.
- The higher the server utilization is, the longer the queue is, and the longer the queue waiting time and total response time are.

In the iSeries environment, examples of service requestors are:

- Applications
- System tasks

Examples of service providers are:

- CPU
- I/O processors
- Disk arms

The equivalent functions of requestors and servers are also present within the client system and within the communications network.

In an SAP R/3 environment, an additional server at which a queue may arise is at the SAP R/3 dispatcher (refer to 2.4, “SAP R/3 work processes, dispatcher, sessions, and modes” on page 30), which assigns work to SAP R/3 work processes. This may happen if you do not define sufficient SAP R/3 work processes for dialog (interactive) and background (batch) work.

When requests arrive randomly at a single server with a single queue, there is an equation that predicts the expected service times as a function of server utilization with a good degree of accuracy over time. This equation expresses the expected service times as a multiple of the time it takes to process a request once it has finished waiting in the queue and is actually processed by the server. The equation is:

$$QM = 1 / (1 - U)$$

Here, U stands for utilization, and QM represents the Queuing multiplier.

For example, if server utilization is at 50%, $1 / (1 - .5) = 2$, this indicates that a new request arriving in the queue is expected to take twice as long to be completed as it would if the server was not being utilized at all.

Response time is directly related to queue length, and the queuing multiplier expresses the effect of queuing on response times. A graph of the queuing multiplier against utilization of the server is shown in Figure 248.




Figure 248. Queuing multiplier versus utilization

As the utilization of a server increases (more work for the server), queuing can account for a much longer elapsed time for work (or request) completion.

The queuing multiplier is an important factor when projecting the impact of adding work or additional hardware on current system performance. Systems with performance problems often show resources with high queuing multiplier factors.

The simplified queuing theory discussed here assumes a single queue of requestors for a single server. In the iSeries model range, multiprocessor (N-way) servers have more than one central processor (CPU) executing instructions, even though there is only a single queue of requestors (Task Dispatching Queue). In this situation, the increased number of servers reduces the queuing multiplier and the average queue length somewhat, but the effect of queuing on response times is still a significant factor.

16.1.1.2 Response time curve

Response time is the elapsed time between the request for a service and the completion of that service. In an interactive iSeries environment, it is the time between the user pressing the Enter key (or a function key) on a 5250 display session and the keyboard unlocking with the information on the display. For an SAP R/3 user, the interactive response time is the time between pressing the Enter button (or clicking the process icon with the mouse) and the information appearing on the SAP GUI screen.

The queuing multiplier is a measure of the effect of queue length on response time, and U is the utilization of the resource providing the service.

Another important concept is highlighted in Figure 248. The curve shows the utilization at various rates and the significance of the knee of the curve. The knee of the curve is the point where a change in utilization produces a corresponding greater change in the queuing multiplier. That is, the change along the Y-axis (queuing multiplier) is significantly greater than the change along the X-axis (utilization). The knee of this curve is the maximum utilization point to which a

certain resource should be driven. After this knee, service time becomes less stable and may increase dramatically for small utilization increases.

Not all resources react the same. There are different recommended maximum values for the different resources, such as CPU, disk, memory, controller, remote line, IOPs, and so on.

You can find more information in *OS/400 Performance Tools/400*, SC41-5340. The graph shows a simplified queuing formula and a curve derived from it that highlights the effect of increasing utilization on the queuing multiplier for a single server.

16.1.1.3 Response time components

In this section, the differences between the components of response time in a client/server iSeries environment such as SAP R/3 and for a traditional interactive (such as 5250 terminal emulation) iSeries user are examined.

Client/server response time

In a client/server environment, the response time perceived by the user is the total response time of the following service providers:

- **Client system:** When a user at a client system, such as a PC, requests information, that request is first processed by the PC and translated to a request to the server system.
- **Communication network (to the server system):** The request is sent through the line to the server (such as a database or application or file server).
- **Server system:** The server system accepts the request and performs the requested functions.
- **Communication network (from the server system):** The server response is sent back to the client.
- **Client system:** The client receives the information, performs further processing as necessary, and presents the final response to the user's request.

Therefore, the total response time experienced by a client/server application user is the total of the service times of the:

- Client
- Server
- Network

Typically, a server system functions in an environment with multiple requestors. The response time experienced by a requestor is affected not only by the function of the particular task, but also by the workload introduced by other concurrent requestors and the relative servicing priority assigned to them.

Client PCs, on the other hand, are single-user systems where the contention for resources is minimal. However, with the introduction of multitasking operating systems and more concurrent activity on the PCs, resource contention is becoming a significant contributor to overall client/server performance.

The number of times information has to move between the client and server (communications flows) and the amount of data moved before a response is completed also increase the response time.

Components of iSeries interactive response time

Within a system, there are many functions contributing to the response time experienced by the interactive user, including CPU time and disk I/O time. There are wait times associated with these servers, including waiting for the CPU and waiting for disk I/O. Each transaction is affected by these.

The interactive response time experienced by an iSeries user is the total of many components as shown in Figure 249.




Figure 249. Components of iSeries response time

During the interactive response time, the following actions occur:

- There is a transmission time delay for the transaction to reach the CPU. This is significant in situations such as a remote workstation.
- Once the transaction reaches the system, the system's response time measurement begins.
 - The job may have to wait for an activity level at the system.
 - Once the activity level is entered, resource utilization begins, which includes:
 - CPU processing time (including queuing)
 - Disk I/O time (including queuing)
 - There may also be periods of inactivity when the transaction is waiting in a variety of states that are reported in the performance tools report. They are also discussed in *Performance Management/400 V4R4*, SC41-5347, and *Performance Tools for AS/400*, SC41-5340. The states may include:
 - Excess Time in Activity Level (Excs ACTM)
 - Short waits
 - Short waits – extended
 - Object or record seize or lock conflict
- There is a transmission delay in the response reaching the user.
- Finally, the time is taken by the user workstation to process the information for presentation.

The components of the response time diagram show that the CPU is only one of the resources (servers) involved in the response time. Disk service time must also be included in response time expectations. Additional wait times (such as exceptional wait times including waiting for record or object locks, waiting for communication line data transmission, and so on) can play an important part in actual response times experienced by a user. They must be considered when analyzing performance problems.

16.1.1.4 Impact of database growth on iSeries performance

The iSeries uses a *binary radix tree* structure to implement indexes. This organization allows a search for a specific index key to be completed using a binary search approach that minimizes the number of searches in finding the required index. Also, when a page segment of a binary radix tree becomes full, the tree is partitioned at a higher level in the tree to minimize the number of page segments that have to be searched. This structure is referred to as a *partitioned binary radix tree*.

On average, a specific key can be located in approximately 20 tests in an index of one million entries! Therefore, the size of a database file has a relatively low impact on accessing a particular row by index.

Naturally, if the entire file has to be processed (as in creating a full customer list, for example), the total processing time is directly related to the file size.

16.1.2 iSeries work management concepts

Work management on the iSeries is the method used to manage iSeries resources to achieve optimum throughput. The iSeries has many objects that interact with each other and the applications to process information efficiently. An understanding of the concepts of work management is an important prerequisite to maximizing system performance.

This section discusses iSeries work management concepts that affect performance. However, it is not the intent of this section to cover all aspects of the subject. Education courses, such as OS/400 Structure, Tailoring, and Basic Tuning (Course Code S6023), are available for a greater understanding of the subject. You can also find more information on all aspects of work management in *OS/400 Work Management*, SC41-5306.

16.1.2.1 Subsystems

A subsystem in OS/400 is an operating environment used to allocate main storage and provide a degree of isolation for jobs with similar processing characteristics (such as batch or interactive) to run in. This minimizes contention for system resources and increases efficiency.

You can look at the components of one of the standard OS/400 subsystems, QBATCH, using the DSPSBSD QBATCH command. The SAP R/3 subsystem descriptions can be found in the library R3<SID>400, where <SID> is the three-character SAP R/3 system identifier. The SAP R/3 subsystem names are R3_<nn>, where <nn> is the instance number assigned during installation.

Some of the important components of a subsystem are:

- **Storage/memory pools:** Storage pool definitions provide information on:
 - Pool identification (within subsystem)
 - Pool size
 - Maximum activity level
- **Autostart job:** This performs a one-time initialization job or a repetitive activity associated with a subsystem. The QSERVER subsystem uses autostart job QPWFSEVER to initiate file and database servers. The SAP R/3 job QXDAEDRSQL is an example of an autostart job in the QSYSWRK subsystem.

- **Routing entries:** In a subsystem, these provide a means of selecting the environment and program that is to be run. The environment includes:

- Entry selection criteria
- Program to run
- Memory pool number (within the subsystem)
- Execution class

- **Class:** A class identifies aspects of the execution environment such as:

- Run priority
- Time slice
- Purge option

The SAP R/3 execution classes have names R3_ <nn>, where <nn> is the instance number. They are located in the library R3<SID>400, where <SID> is the SAP R/3 system identifier.

- **Communications job:** A communications job, from an OS/400 work management standpoint, is a batch job started by a program start request from a remote system. In the case of servers, the start request is initiated by a client or PC application. Communications work entries identify the sources from which the subsystem accepts start requests. A communications entry includes:

- Device
- Mode
- Job description
- Default user
- Maximum active jobs

For example, a program start request using a *mode* entry of QSERVER is routed to the QSERVER subsystem. Users of other modes, including QCASERV and QPCSUPP, are routed to the QCMN subsystem.

- **Prestart job:** This is a batch job that is started in anticipation of a program start request. The objective of the prestart job is to have as much “startup” activity as possible before the remote request is received.

16.2 SAP memory management concepts

This section introduces the memory management techniques used with SAP R/3. To simplify the discussion, we briefly explain the memory management concepts. The parameter settings shown here refer to memory management for SAP R/3 Releases 4.6A and higher only and may change as enhancements to the R/3 kernel and OS/400 are developed. They also depend on the type of R/3 instance (production or development), the total amount of physical memory and storage available, the number of R/3 work processes, and the R/3 modules and transactions being used. For the latest information, you should refer to the SAP notes:

- Note 139326 Memory management in releases as of 4.6A, AS/400
- Note 103747 Performance 4.0/4.5/4.6: Parameter recommendations
- Note 44695 AS/400: Memory management as of release 3.0C (for R/3 releases before 4.6A; valid for releases 3.0C to 4.5B)

Please refer to the appropriate SAP online help documentation if you need more detailed information.

The SAP architecture has *user contexts* containing information on each active user. The user context is mapped into the *roll area* of the work process during the processing of a dialog step. The work process always processes using the current roll area. The user context as shown in Figure 250 includes:

- User-specific data that is shared across all sessions
- Up to six session contexts including:
 - Session-specific data
 - Up to nine internal modes (internal sessions) that are created by ABAP programs (each of these is considered a user context and this internal mode is what is copied to or from the roll area)

User Data	Session-1 data	Session-2 data	Session-3 data		Session-6 data
	Mode - 1 Mode - 2 Mode - 3 Mode - 4 : Mode - 9				

Figure 250. User contexts

16.2.1 Early implementations

This section explains memory management in the implementations of SAP R/3 before version 3.0x and SAP R/3 on the iSeries to help you better understand the subject:

- The SAP work process uses an address space that includes the following areas:
 - **Roll area** is a part of the address space of each work process where the segments of a user context are copied into at the beginning of a dialog step.
 - **Private memory** is allocated when the work process requires more memory than is available in the roll area and paging area. Private memory is allocated as required until the limits specified in the *abap/heap_...* parameters for the instance are reached.

When a user task requires the assignment of private memory, that user can run tasks in that work process only. Also, a work process with private memory allocated does not participate in the work process dispatch. Therefore, the allocation of private memory effectively reduces the number of work processes available to other users.

The allocated private memory is released only when the task using the work process has completed and the work process is restarted. The number of processes that are allowed to run in “private mode” can be limited by *rdisp/wppriv_max* and *rdisp/max_priv_time*, but this mechanism applies only to dialog work processes. See Figure 251.




Figure 251. R/3 memory paging: Early implementation

- In addition to the work process address space, the following areas of memory are shared by the work processes:
 - **SAP paging memory** is an extension to the roll area and is used by ABAP programs to hold large amounts of data such as internal tables. This is also a part of the work process address space. The paging memory can also be accessed by all work processes.
 - **Shared roll memory/file** is a common resource accessible by all work processes. It contains user context information for users not currently processing a dialog step.

Figure 252 shows that required segments of user context are copied from it at the beginning of a dialog step (*roll-in*) and are copied to it at the end of the dialog step (*roll-out*) from the corresponding work process address space.




Figure 252. R/3 memory paging and shared roll memory

- **Shared memory pool** is a pool of memory defined to be used by approximately 55 logical shared memory segments. These logical segments contain various table and program buffers (refer to SAP R/3 transaction ST02). The allocation of the memory pool to the logical segments is managed by an SAP administrator using a bank of approximately nine registers.

This was introduced because early systems only had access to a limited number of shared memory segments. Therefore, multiple logical memory

segments were mapped to a single real memory segment. Each logical segment contains information for administering SAP R/3 or to hold buffered data.

16.2.2 Later enhancements

In more recent enhancements of the architecture, the following components were included in the memory management design:

- **Extended memory** is an enhancement to the concept of using roll, paging, and private memory. It is based on a sharing of memory but does not involve the copying of the contents of the shared resource. The required user context is mapped into the virtual address space of the work process before a dialog step is run.

The maximum total extended memory and the maximum extended memory per dialog process are defined using the *em/initial_size_MB* and *ztta/roll_extension* parameters in the instance profile. Non-dialog work processes use extended memory since release 4.5A.

- **Implementation of the roll area** was enhanced to allow a part (defined as *ztta/roll_first*) to be used initially. The balance (equal to *ztta/roll_area* minus *ztta/roll_first*) is allocated after using the specified extended memory.
- **SAP paging memory** is no longer included in the work process address space. It is still used for some functions such as export to memory.

With the inclusion of the extended memory and roll extension memory segments, and the change in implementation of paging memory, a work process address space may be represented as shown in Figure 253.




Figure 253. R/3 memory paging: Current implementation

Some of the key SAP R/3 values that affect memory and disk utilization, and are specified in the SAP R/3 instance profile (for example, *<SID>_DVEBMGS<nn>_<hostname>*, where *<SID>* is the SAP R/3 system identifier, *<nn>* is the instance number, and *<hostname>* is the TCP/IP hostname) are listed here:

- **ipc/shm_psize:** Shared memory pool descriptor sizes that are used to define the shared memory segments. In some platform implementations, the number of shared memory segments per work process is limited. Shared memory pools are not used in the iSeries implementation.

- **abap/use_paging**: Specifies that SAP's new memory management facilities should be used.
- **ztta/roll_area**: Specifies the total roll memory available to each work process. This value is set to 16773120 on the iSeries server.
- **ztta/as4/roll_extent_size**: Determines the unit of allocation of roll_area in the iSeries implementation.
- **ztta/roll_first**: The new SAP R/3 memory management technique subdivides the roll area into two parts. The roll_first memory is allocated initially. The advantage of this is to minimize the size of the user context to be rolled in and out. The remaining part of roll memory is used after extended memory allocation.
- **ztta/roll_extension**: Limits the maximum amount of extended memory that can be allocated by any dialog user context and is specified in bytes. A large value of 2 GB is recommended to prevent dialog work processes entering private mode.
- **em/initial_size_MB**: Determines the total amount of extended memory available for use by all dialog user contexts and is specified in megabytes. After roll_first memory is allocated, user contexts use extended memory up to the limit specified in the roll_extension parameter until the extended memory is exhausted.
- **rdisp/ROLL_SHM**: Specifies the maximum amount of shared roll memory in blocks. This is not used in the iSeries implementation.
- **rdisp/ROLL_MAXFS**: Denotes the size of the file for shared roll memory in blocks. This is not used in the iSeries implementation.
- **rdisp/PG_SHM**: Specifies the maximum amount of paging buffers in blocks. In the iSeries server, all paging is done in shared memory. This value should be set equal to rdisp/PG_MAXFS.
- **rdisp/PG_MAXFS**: Denotes the size of the file for paging buffers in blocks. In the iSeries server, all paging is done in shared memory. This value should be set equal to rdisp/PG_SHM.
- **abap/heap_area_dia**: Identifies the maximum amount of process-specific memory (bytes) that may be used by all dialog work processes.
- **abap/heap_area_nondia**: Identifies the maximum amount of process-specific memory (bytes) that may be used by the non-dialog work processes. Background work processes do not use extended memory. This value is larger than the corresponding value for dialog work processes.
- **abap/heap_area_total**: Identifies the total amount of process-specific memory (bytes) available to all work processes.
- **abap/heaplimit**: Specifies the size above which the work process is restarted after processing of the current user context. The restart releases any private memory allocated to the work process.
- **rdisp/wppriv_max_no**: Specifies the maximum number of dialog work processes that can be in private mode.
- **rdisp/max_priv_time**: Once the limit of work processes in private mode is reached, the system waits for the number of seconds specified before restarting the oldest work process in private mode.

16.2.3 SAP memory management (iSeries)

The iSeries architecture uses the concept of *single-level storage (SLS)* where the main memory and disk are managed by the system as a single address space. Therefore, every addressable byte has the same address regardless of the process accessing it. In contrast, “process local storage” has different processes having the same address, but accessing different bytes in memory. When an object is created on the iSeries server, space is allocated from a *large* total address space using 64-bit addressing. Mapping this virtual address space to real memory and disk is managed by the iSeries server.

Therefore, because of dynamic memory management by the iSeries, even if a large address space is assigned for an SAP R/3 memory unit, the unused pages of memory are pageable and available for use by other requesters. In addition, copying data to and from the user context from and to a roll is not necessary.

With OS/400 V4R4, *teraspace* memory management is also available, which was introduced to support very large linear address spaces and to avoid the 16 MB memory segment size restriction in SLS. Internally, teraspace uses SLS to make available large linear address spaces.

From SAP R/3 Release 4.6, extended memory, heap memory, paging memory and the R/3 internal buffers are located in teraspace. Only heap memory is implemented directly in the SLS.

OS/400 memory pools are allocated to subsystems. The memory requirement of an SAP R/3 instance is managed within a subsystem. The total memory available to an SAP R/3 instance is the memory that is allocated to the corresponding R3_<nn> subsystem, where <nn> is the instance number. Naturally, SAP R/3 jobs compete for resources with other jobs assigned to the same memory pools. The corresponding OS/400 *subsystem description* defines the amount of memory assigned to each SAP R/3 instance.

The results from the iSeries architecture are:

- The roll_first space is recommended to be set to 1 byte. The roll_area is specified as a high value (16773120).
- The need for management of user context sizes through manipulation of parameters, such as ztta/roll_area and ztta/roll_extension, is eliminated.
- The maximum extended memory that a session can request is specified in ztta/roll_extension as a large value (2 GB) to prevent dialog work processes from entering private mode. Even though 2 GB is specified, only the required amount is allocated. Extended memory is allocated in blocks of the size specified in em/block_size_KB from the extended memory pool (em/initial_size_MB).
- The parameter for em/initial_size_MB defines the amount of extended memory that is preallocated at the startup of the R/3 system. For releases before R/3 4.6, it is only allocated when it is really used.
- The need to allocate private memory (and the resulting exclusion of the work process from availability) is minimized if the extended memory limit (ztta/roll_extension) is defined large enough.

- The iSeries uses all the memory assigned to each SAP R/3 instance. Memory management on the iSeries is designed to minimize memory faulting through dynamic paging algorithms.
- The need to allocate specific “swap space” on disk does not arise in the iSeries because of the single-level storage architecture. Some amount of “spare” disk space is required to allocate temporary work files and to provide for database growth.
- The need for complex management of “table spaces” and file system extents in the installation and maintenance of the database does not arise in an iSeries environment.
- The automatic balancing of disk space when objects are created or extended by the iSeries provides a more or less even distribution of the workload over all available disks.
- Allocation of roll memory buffers and file size is not required.
- No shared memory pools (*ipc/shm_psize*) need be defined or managed (that is, mapping the logical shared memory segment to the 10 shared memory pools is not required). The *ipc/shm_** parameters are ignored.

Table 35 contains a selection of SAP R/3 parameters that should be set initially and reviewed after implementation to optimize performance. The values shown in the recommended value column apply for R/3 Release 4.6 or higher and are typical of a production SAP instance.

Table 35. SAP parameter settings (iSeries)

Parameter name	Description	Recommended value
ztta/roll_area	Roll area (bytes)	16773120
ztta/roll_first	First roll area (bytes)	1
ztta/roll_extension	Extended memory/context (bytes)	2000000000
em/blocksize_KB	Extended memory allocation block size (KB)	4096 (cannot be changed)
em/initial_size_MB	Extended memory initial allocation size	8096
rdisp/ROLL_SHM	Roll memory	Not applicable
rdisp/ROLL_MAXFS	Roll file	Not applicable
rdisp/PG_SHM	Paging memory	32768
rdisp/PG_MAXFS	Paging file	32768
abap/heaplimit	Work process restart threshold (bytes)	40000000
abap/heap_area_dia	Dialog work process memory (bytes)	2000000000
abap/heap_area_nondia	Non-dialog work process memory (bytes)	2000000000
abap/heap_area_total	Total work process memory (bytes)	2000000000

Parameters indicated as “not applicable” are not used in the iSeries implementation and are ignored.

Important

We recommend that SAP R/3 memory parameters only be changed on the advice of SAP or IBM. The SAP Early Watch service and Going Live functional upgrade check examine the settings of these parameters and make recommendations that are optimal for your workload and system resources. It is vital that these sessions are scheduled before you “go live” with your production SAP environment.

Use the SAP R/3 transaction ST02 and Current Parameters push button, or the SAP R/3 transaction RZ11 to determine the current parameter values set for the instance profile. The values can be changed by R/3 transaction RZ10 for the instance profile (for example, `<SID>_DVEBMGS<nn>_<hostname>`, where `<SID>` is the system identifier, `<nn>` is the instance number, and `<hostname>` is the TCP/IP hostname). A complete list of the current values for all R/3 profile parameters, including ones not explicitly defined in the instance or default profiles, can be displayed via transaction RZ10 by clicking from the menu bar **Goto-> Profile values-> Of a server**, and then double-clicking the server name.

Note: The effect of these values in an iSeries SAP R/3 implementation is different from other platforms. Large values in the parameters result mainly in increased disk utilization, with “runaway programs” using excessive amounts of system resources cause problems. These jobs can significantly impact other users of the system by running longer and using up more system resources before they reach the specified limits and are terminated.

You should also consult SAPNet Notes 121625, 139326, or 44695 for more for buffer size parameter settings specific to the iSeries. Please refer to 16.6.1.9, “SAP R/3 internal buffers” on page 450, for information on buffer sizes and additional parameters that need to be reviewed for good system performance.

16.3 A approach to performance analysis

This section outlines an approach to addressing performance issues with SAP R/3 on the iSeries server:

1. Monitor the OS/400 system (refer to 16.4, “iSeries system monitoring” on page 380):

Perform a “snapshot” analysis using the WRKSYSSTS, WRKDSKSTS, WRKACTJOB, WRKSBSJOB, and WRKSYSACT commands. For a more complete review, collect OS/400 performance data with the OS/400 Performance Monitor, and use OS/400 Performance Tools to analyze the data. Check the:

- CPU usage, memory faulting, and job queuing
- Disk arm activity and space usage
- Job activity (high CPU/disk usage)
- Job status (waiting on a message, in “snapshot” mode only)

2. Tune the OS/400 system as necessary (one change at a time) (refer to 16.5, “iSeries system tuning” on page 412):
 - Pools and activity levels
 - Job priority
 - Subsystem definition
3. Repeat step 1 and step 2 until you achieve an acceptable level of OS/400 resource tuning.
4. Monitor the SAP R/3 system’s performance (refer to 16.6, “SAP R/3 application performance” on page 437):
 - a. Review the SAP R/3 workload and response times (ST07):
 - In-house developed functions: Trace SQL (ST05)
 - b. Perform a workload analysis (ST03):
 - High CPU usage?
 - Dialog processes busy/high wait time?
 - High program load time?
 - DB request time?
 - c. Check the SAP buffer analysis (ST02; also run SAP R/3 report RSDBBUFF):
 - Hit ratio high around 96%?
 - Buffer space free?
 - Free directory entries?
 - d. Analyze the database activity:
 - Analyze SQL activity (ST04).
 - Analyze database activity (ST04/DB02).
 - e. Analyze the high CPU users (ST06).
 - f. Check for printing problems:
 - Printing to other servers?
 - Printing priority?
5. Tune the SAP R/3 system (refer to 16.7, “SAP R/3 tuning” on page 472):
 - a. Adjust the SAP buffers.
 - b. Adjust the SAP run-time parameters.
 - c. Adjust the number of SAP work processes.
 - d. Adjust the priority of the SAP R/3 processes.
 - e. Spread the load (move users to other application servers, and distribute users by application module).
6. You may have to repeat the entire procedure if necessary.

Note: Please review the latest performance information in the following “AS/400 Latest News” SAPNet notes:

- SAPNet Note 103123 for SAP R/3 3.1I
- SAPNet Note 77864 for SAP R/3 4.0A
- SAPNet Note 99814 for SAP R/3 4.0B
- SAPNet Note 103805 for SAP R/3 4.5A
- SAPNet Note 139942 for SAP R/3 4.5B
- SAPNet Note 146836 for SAP R/3 4.6A
- SAPNet Note 179665 for SAP R/3 4.6B

- SAPNet Note 300105 for SAP R/3 4.6C
- SAPNet Note 319471 for SAP R/3 4.6D

You should also refer to these Notes:

- SAPNet Note 101113 AS/400: Analysis of performance problems
- SAPNet Note 103747 Performance 4.0/4.5/4.6: Parameter recommendations

16.4 iSeries system monitoring

This section introduces the concept of iSeries performance monitoring and tuning. These facilities can be used to assist in optimizing iSeries resource utilization and to minimize possible contention between applications and tasks running on the system.

Prior to making any performance tuning changes, it is necessary to review key iSeries performance measurement parameters. You can make performance tuning adjustments to your iSeries server in one of two ways:

- Set up the system to make performance adjustments dynamically by using the OS/400 system value QPFRADJ, which:
 - Monitor the memory pool faulting.
 - Adjust the memory pool sizes and activity levels.
- Make the performance adjustments manually after setting the OS/400 system value QPFRADJ=0.

When you make adjustments, change only one parameter at a time. Then review the impact of the change by monitoring the system before making more changes.

16.4.1 iSeries performance indicator guidelines

This section looks at some guidelines for key iSeries performance indicators. For further information, refer to the *OS/400 Work Management*, SC41-5306.

The key OS/400 resource utilization indicators that affect performance are:

- **Memory faulting rates** indicating the frequency of faulting in a memory pool. In general, attempt to keep these values as low as possible.

The important consideration for memory faulting is the *number of memory faults per second that each active thread can experience without impacting response times*. Therefore, the total number of faults that occur in a pool is not the only important consideration. A pool used by many threads can show a much higher average number of faults per second compared to a pool of the same size with a lower number of active threads (each SAP R/3 work process is represented by a single thread).

For example, if there are 10 active threads in a pool showing 100 faults per second, the number of faults per thread per second is 10. Assuming a disk service time of 0.01 seconds, the delay resulting from memory faulting is 0.1 seconds.

On the other hand, if there are only two active threads in the same pool, there is an average of 50 faults per thread, at a “cost” of 0.5 seconds per thread. The following calculation may help you evaluate the impact of memory pool faulting on overall response time:

- **Machine pool:** This is always System Pool 1, and the faulting rates should be maintained below *10 faults per second*.
- **Other pools:** The other pools should have faulting rates preferably under 200 faults per second.

Each R/3 system runs in its own OS/400 subsystem. Each subsystem runs in a memory pool. By default, SAP R/3 runs in the *BASE pool, which is System Pool 2, and the amount of memory allocated to this pool should be maximized. However, if you run more than one R/3 system on one iSeries, you have to consider potential performance impacts of one system on the other. By default, the jobs in each subsystem will run at the same priority, which, therefore, contend equally for the same resources.

With multiple R/3 systems, they may be separated into different memory pools, and the system value QPFRADJ is set to 0 for no automatic adjustment of the memory and activity level of a pool. In this case, specific amounts of memory are allocated to each pool via the WRKSYSSTS command. This memory is not available to other subsystems running outside that memory pool.

If the system value QPFRADJ is set to 2, OS/400 performs automated dynamic performance tuning. This checks the paging rates periodically and will adjust the memory pool sizes to reduce paging. In this case, you use the OS/400 WRKSHRPOOL command to set the priority of each pool.

In this scenario, the WRKSHRPOOL command is also used to set the minimum percentage (for example, 25%) of system memory for a pool, to protect the R/3 system from losing so much memory that it can no longer function. The absolute minimum for an R/3 system to function is approximately 250 MB. Also, when the STOPSAP command is used to end the SAP system running in a separate memory pool, the OS/400 subsystem for that R/3 system is not ended. The pool is still active and will consume some memory. If the R/3 system is the only thing in the pool, then the subsystem should also be ended so the memory can become available to other pools.

If you are running other non-SAP applications on the same iSeries server as SAP R/3, define a separate pool for the SAP R/3 system:

Faults in the memory pool used by SAP R/3 system	= 200 per second
Number of active (not total) work processes	= 10
Average number of faults per work process	= 20 per second
Faulting overhead (assuming disk response time=10ms)	= 200ms per second

Therefore, the page faulting overhead is 200ms per second or 20% for each work process. For more information, see SAP Notes 44695 and 139326 *AS/400: Memory Management*.

- **Disk space usage** is expressed as a percentage of total space available in the system auxiliary storage pool (ASP1). When the SAP R/3 system and any other applications are active, the usage indicated by the displays includes all permanent objects as well as any temporary objects created by OS/400 to manage system activity. Review your disk usage during peak activity. Excessive disk space usage can impact performance, but is not directly related to the percentage used, only the available disk space. Ensure that you allow adequate disk space for growth of the database and the usage does not exceed 90%.

- **Disk arm utilization** represents an important component of overall system performance. It is displayed on the WRKDSKSTS screen as the %Busy column. High disk arm utilization can result in queuing for disk requests to be serviced. While higher utilization can be supported by the latest iSeries disk drives, we recommend that you maintain arm utilization under 40%.
- **Number of activity levels** is the number of process threads that can be active in the memory pool. The maximum that can be active for a pool is displayed on the WRKSYSSTS and WRKSHRPOOL screens. A task must allocate an activity level prior to being serviced by the CPU. When jobs queue for an activity level, you notice transitions from *Wait-to-Ineligible*.

The number of activity levels for R/3 systems does not need to be more than the total number of R/3 threads and work processes in the instance if only the R/3 system is running in the pool. If you are running R/3 in *BASE pool, which is normally the case, you need to make some allowances for the OS/400 subsystem monitors that run in this pool by default. The value should be high enough to keep the transitions *Wait-to-Ineligible* and *Active-to-Ineligible* at zero. However, in situations where the system is memory constrained, you may consider reducing the number of activity levels to a lesser value. If the memory pool is shared by other applications, a suitable value needs to be set to provide adequate activity levels to utilize the CPU effectively.

16.4.2 OS/400 commands versus SAP R/3 transactions

For monitoring performance of the overall iSeries server, use either of the following methods, which provide identical information:

- OS/400 commands
- R/3 transactions

Note

The SAP R/3 transactions do not allow changes to OS/400 system parameters. Use native OS/400 commands to make any adjustments.

The SAP transactions provide historical data to compare periods during the day or to monitor trends. The OS/400 interactive monitoring commands show only current information. However, you can use the OS/400 STRPFRMON command, the Performance Collector in Operations Navigator, or other third-party software tools, such as Snapshot/400 or Insight (discussed in 16.4.13, “Insight SAP performance reports” on page 412), to collect and store performance data for later analysis and comparison.

16.4.2.1 System monitoring using OS/400 commands

The OS/400 WRKSYSVAL command enables you to work (display or change) on some parameters that provide many default values for the entire system. There are other OS/400 commands that enable you to work with statistical information gathered during an identified time interval (shown as *elapsed time* at top of the display). These commands include:

- WRKSYSSTS (OS/400 System Status)
- WRKDSKSTS (OS/400 Disk Status)
- WRKACTJOB (OS/400 Active Jobs)
- WRKSBSJOB (OS/400 Subsystem Jobs)

The information is:

- Updated by extending the measurement time (F5).
- Reset and a new time interval started (F10).

We recommend the time interval for observation to be *not less than five minutes*. For a more detailed analysis, use the data collected by the OS/400 Performance Monitor (STRPFRMON) command, which collects data over specified periods into system-created database files in the specified library. You can use Performance Tool/400, 5769-PT1, to analyze the collected data (see 16.4.12, “iSeries Performance Tools” on page 409) if you purchased this licensed program. The performance data collection capability is a part of the standard operating system.

16.4.2.2 System monitoring using SAP R/3 transactions

The SAP R/3 transaction ST06 provides you with data gathered by the OS Data Collector, SAPOSCOL. The OS Collector is started automatically when the SAP instance is started through the inclusion of the following statement in the SAP R/3 Start Profile for the instance (but only one collector is active per iSeries server):

```
Start_Program_05 = local $(DIR_EXECUTABLE)/SAPOSCOL
```

The start profile file is in the directory /usr/sap/<SID>/SYS/profile and has the name START_DVEBMGSnn_<hostname>.

It can be stopped and restarted using the OS Collector push button on ST06, or it can be stopped via the R/3 Main Menu as <SIDOFR>. Note that if you have more than one R/3 system on an iSeries, the SAPOSCOL job will run in only one of the active R/3 subsystems, usually that of the first R/3 system to be started.

The initial display from transaction ST06 presents summary information on:

- CPU utilization
- Memory
- Pool transitions
- Disk utilization

An example of the displays from transaction ST06 appears in Figure 254 and Figure 255.




Figure 254. ST06 OS monitor: Summary (Part 1 of 2)




Figure 255. ST06 OS monitor: Summary (Part 2 of 2)

Note: Some values are not applicable (marked N/A) in the OS/400 environment, and others are treated differently from other operating system platforms.

CPU Utilization shows average CPU utilization over the last one minute, five minutes, and 15 minutes. This is sometimes referred to as the *load average*.

Pages ins/outs appear as equal on the iSeries server due to the paging algorithm where all available memory is used. If information is paged in, an equivalent amount must be paged out.

Clicking the Detail analysis menu push button presents a menu that allows you to review performance data through:

- A “Snapshot Analysis” of the last 59 seconds (Table 36)

Table 36. SAP R/3 transaction ST06: Snapshot analysis

Snapshot analysis		
Button	System information	Reference
CPU	CPU utilization	
Memory	Pages per second	
Disk	Disk drive statistics for each disk unit	
LAN	LAN usage	
Filesys (File system)	Free disk space in each ASP	
Pool	Paging/queuing statistics	
Top CPU	High CPU processes	

- Data collected in “Previous Hours” when the Collector was active (Table 37)

Table 37. SAP R/3 transaction ST06: Previous hours

Previous hours		
Button	System information	Reference
CPU	CPU utilization	
Memory	Pages per second	
Disk	Disk drive usage	
LAN	LAN usage	
Filesys (File system)	Free disk space in each ASP	
Pool	Paging/queuing statistics	
OS log	OS/400 QSYSOPR messages	
HW Info (Hardware information)	WRKSYSSTS, WRKDSKSTS, and Hardware configuration	

- A performance database comparing previous periods (Table 38)

Table 38. SAP R/3 transaction ST06: Performance database

Performance database		
Button	System information	Reference
Compare recent days	CPU, memory, disk hourly statistics	
Compare all servers	CPU, memory, disk hourly statistics	

- Additional functions showing OS/400 system values, parameter changes, and the server network (Table 39)

Table 39. SAP R/3 transaction ST06: Additional functions

Additional functions		
Button	System information	Reference
System configuration	OS/400 system values	
Parameter changes	Changes to relevant system values	
LAN check by ping	Test LAN communications	
PTF check	Compare PTFs installed with IBM Info APAR for SAP R/3	

These options provide information on the iSeries server through the SAP GUI interface for users who may not have access to an IBM 5250-type display session.

The display you see when you click Detail analysis menu on transaction ST06 is shown in Figure 256.




Figure 256. ST06 OS Monitor: Detail analysis menu

16.4.3 OS/400 system values

The following section describes a few values that you should review to optimize overall iSeries server performance.

Use the OS/400 WRKSYSVAL command to view and change system values, or the SAP R/3 ST06 transaction to view the OS/400 system values.

16.4.3.1 WRKSYSVAL display

Refer to *OS/400 Work Management*, SC41-5306, for more information on setting OS/400 system values. You should consider the values in Figure 257.

```

Work with System Values
System: DEVSYS
Position to . . . . . Starting characters of system value
Subset by Type . . . . . *ALL F4 for list

Type options, press Enter.
2=Change 5=Display

Option Value Type Description
QACTJOB *ALC Initial number of active jobs
QADLACTJ *ALC Additional number of active jobs
...
QADLTOTJ *ALC Additional number of total jobs
...
QBASACTLVL *STG Base storage pool activity level
QBASPOOL *STG Base storage pool minimum size
...
QDYNPTYADJ *SYSCTL Dynamic priority adjustment
QDYNPTYSCD *SYSCTL Dynamic priority scheduler
...
QJOBMSGQFL *ALC Job message queue full action
QJOBMSGQMX *ALC Maximum size of job message queue
...
QPFRADJ *SYSCTL Performance adjustment
...
QMCHPOOL *STG Machine storage pool size
...
QTOTJOB *ALC Initial total number of jobs

```

Figure 257. Work with System Values (WRKSYSVAL)

Recommendations for these system values for iSeries servers running SAP R/3 are published in the SAP document *SAP system installation: IBM AS/400* for each SAP R/3 release.

The QPFRADJ system value controls automatic tuning of the iSeries server. It adjusts memory pool sizes and activity levels based on the extent of system activity and memory faulting:

- 0 = Switch off automatic tuning for all values.
- 1 = Set up the system only at IPL time.
- 2 = Set up the system at IPL time and tune dynamically.
- 3 = Tune dynamically; do not set up the system at IPL time.

Note

We recommend that SAP R/3 systems are operated with QPFRADJ=0, as R/3 runs in the BASE pool and usage in other pools INTERACT and SPOOL should be minimal. The amount of memory in the BASE pool should be maximized. It is important that the MACHINE pool, where OS/400 runs, be assigned sufficient resources to minimize any page faults in this pool. Refer to 16.5.1.2, “Automatic tuning of pools and activity levels” on page 413, for information on using QPFRADJ values of 2 or 3 for automatic memory pool and activity level tuning.

16.4.3.2 SAP R/3 system on iSeries: System values display

From transaction ST06 Detail analysis menu, click the **System configuration** push button for the SAP R/3 system values display (Figure 258 only shows a subset of iSeries system values).

Local (ASM23) / Configuration Parameter

Refresh display		
Tue Nov 28 15:55:51 2000# interval 59 sec.		
Parameter	Value defined	Value used
QACTJOB		50
QADLACTJ		20
QADLTOTJ		20
QBASACTLVL		137
QCCSID		65,535
QJOBMSGQMX		32
QMAXACTLVL		32,767
QMCHPOOL		247,912
QTOTJOB		500
Parameter	Value defined	Value used
QCHRID		00000006970000000037
QCTLSBSD		QCTL QSYS
QDYNPTYSCD		1
QIPLDATTIM		*NONE
QJOBMSGQFL		*WRAP
QMODEL		170
QPFRADJ		0
QPWDEXPITV		*NOMAX
QPWRRSTIPL		0
QQRVDEGREE		*NONE
QQRVITLMT		*NOMAX
QSECURITY		40
QSRTSEQ		*HEX
QUPSDLYTIM		*CALC 0000000200
QUTCFFSET		+0000

Figure 258. ST06 Detail Analysis menu: System configuration (values)

16.4.4 OS/400 system status and disk status

This function provides information on OS/400 system resource utilization for a particular time interval, including:

- Disk capacity and usage of the system auxiliary storage pool (ASP1)
- Faulting rate per second in each memory pool
- Queuing of jobs running in each memory pool

16.4.4.1 WRKSYSSTS display

Select an elapsed time of at least five minutes for an interactive analysis of OS/400 resource utilization (F10 restarts elapsed time and F5 extends the

elapsed time). Choose the Advanced display mode by pressing F21 to select the Assistance level. Run this command during typical levels of activity, and check the items highlighted in Figure 259.

Work with System Status										DEVSYS		
										11/28/00	10:56:19	
% CPU used		:	90.3	System ASP		:	488.2 G					
% DB capability		:	24.7	% system ASP used		:	68.4641					
Elapsed time		:	00:14:12	Total aux stg		:	496.6 G					
Jobs in system		:	4384	Current unprotect used		:	38287 M					
% perm addresses		:	.027	Maximum unprotect		:	40407 M					
% temp addresses		:	.500									
Sys	Pool	Reserved	Max	----DB----		--Non-DB---	Act-	Wait-	Act-			
Pool	Size M	Size M	Act	Fault	Pages	Fault	Pages	Wait	Inel	Inel		
1	489.87	325.80	+++++	.6	1.7	16.7	19.4	21.4	.0	.0		
2	5318.12	.08	81	45.2	4193	496.2	631.7	1546	.0	.0		
3	32.00	.00	12	.0	.7	.5	1.2	67.3	.0	.0		
4	256.00	.00	16	.0	.2	.1	.1	2.3	.0	.0		

Figure 259. Work with System Status display

Note the following explanations of some of the parameters and columns in the screen in Figure 259:

- **% CPU used** indicates the amount of CPU used during the elapsed time. If any batch processes is running at the time, this value can be high. If CPU usage is sustained at a high level, you may have an overloaded processor.
- **% DB capability** indicates the amount of CPU used for database processing during the elapsed time.
- **% System ASP used** indicates the amount of disk space used of the total available in the system ASP. This value includes the amount of disk space occupied by the temporary objects created during normal OS/400 operations.
- **Total aux stg** refers to the total disk space in the entire system and excludes any disk space that is used up by disk protection such as RAID-5 or mirroring.
- **Current unprotect used** indicates the amount of disk currently used by temporary objects. Because SAP uses temporary objects, you will notice this value rise after an IPL when SAP is started. When SAP is stopped, most of this space is returned to the system ASP. The amount of temporary storage used is a function of the number of active R/3 work processes, the instance profile parameters for the SAP system, and the R/3 transactions executed. An estimate of the initial size of temporary storage that will be used when SAP is started can be made using the program SAPPFPAR in the kernel library. Call this command with the parameter pf=instance profile parameter name as follows:

```
CALL PGM(SAPPFPAR)
PARM('pf=/usr/sap/<SID>/SYS/profile/<SID>_<instancename>_<hostname>')
```

You may then enter the memory profile parameter names and see their values.

- **Maximum unprotect used** indicates the maximum amount of disk space used by temporary objects since the last IPL.
- **DB and Non DB Faulting** is measured in faults per second and is the only measure to determine how the iSeries server is using the available memory in

each pool. Make sure you review the total number of faults per second in each pool and not the number of pages.

- **Wait->Inel** measures the transition of jobs for each memory pool from the wait state to ineligible state where a job is ready to run but cannot secure an activity level. The measurement is in number of threads per minute. If this is greater than 0, it indicates that the value set for *maximum active* jobs in that memory pool is too low.
- **Act->Inel** measures the transition of jobs for each memory pool from the active state to ineligible state. If this is greater than 0, it may also indicate that the value set for *maximum active* jobs in that memory pool is too low. Threads that are in the ineligible state have work to do, but the system is unable to accept more work at that time.

Note: By default, all SAP R/3 jobs run in pool 2 (*BASE).

16.4.4.2 WRKDSKSTS display

The Work with Disk Status command shows the iSeries server's disk activity and helps to determine any potential impact on performance. Use the OS/400 WRKDSKSTS command to display the information as shown in Figure 260.

Work with Disk Status										DEVSYS
Elapsed time: 00:01:05										11/28/00 10:58:11
Unit	Type	Size (M)	% Used	I/O Rqs	Request Size (K)	Read Rqs	Write Rqs	Read (K)	Write (K)	% Busy
1	6717	6442	68.4	21.1	7.0	16.0	5.1	6.9	7.6	4
2	6607	3670	68.4	9.2	12.9	8.0	1.1	13.1	11.3	5
3	6607	3670	68.4	8.9	12.1	8.1	.7	12.4	8.8	4
4	6607	3670	68.4	11.7	11.6	10.4	1.2	11.3	14.6	10
5	6607	3670	68.4	10.8	6.3	8.4	2.3	5.8	8.0	5
8	6607	3670	68.4	5.7	8.2	3.7	2.0	9.5	5.9	3
9	6607	3670	68.4	7.1	11.6	5.9	1.2	12.4	7.5	4
10	6607	3670	68.6	6.0	15.9	5.3	.6	16.8	8.8	5
11	6607	3670	68.4	6.2	11.4	4.7	1.5	12.8	7.0	7
12	6607	3670	68.4	8.8	8.2	8.1	.6	8.2	8.3	1
13	6607	3670	68.4	6.6	9.7	5.3	1.3	9.8	9.1	7
14	6607	3670	68.5	8.8	7.4	6.5	2.2	7.4	7.5	4
15	6717	6442	68.4	17.9	8.2	14.7	3.2	8.5	7.0	4
										More...
Command										
====>										
F3=Exit F5=Refresh F12=Cancel F24=More keys										

Figure 260. Work with Disk Status display

Note the following explanations:

- **% Used** indicates the amount of space used on the disk unit listed. It should not exceed 90% on any disk. Ensure you have *adequate spare disk space to allow for growth in the database*.
- **% Busy** is a key factor in disk performance that measures disk utilization and should remain under 40%.

16.4.4.3 SAP R/3 system on iSeries: System status information

The SAP R/3 transaction ST06 provides OS/400 system status information similar to the OS/400 WRKSYSSTS command and the WRKDSKSTS command.

From the SAP R/3 Detailed analysis menu (Figure 256), click the **HW Info** (hardware information) push button within *Previous hours* to display the values:

- System status
- Disk status
- Installed hardware

The System Status Information screen is shown in Figure 261.

Local (ASM23) / Operating System Monitor: OS/400

System Status Information

Page 1

5769SS1 V4R5M0 000526AS2311/28/00 16:22:20

% CPU used1.0System ASP60.13 G

% DB capability.0% system ASP used81.7840

Elapsed time00:21:24Total aux stg68.71 G

Jobs in system196Current unprotect used6280 M

% perm addresses.009Maximum unprotect7334 M

% temp addresses.010

SysPoolSize MReservedMaxDBNon-DBAct-Wait-Act-PoolSubsystemLibraryPaging

Pool1242.10163.07++++.0.0.0.11.8.0.0*MACHINEOption

22703.17.37137.0.0.53.9537.7.0.0*BASE*CALC

330.71.006.0.0.0.02.1.0.0*SPOOL*FIXED

496.00.008.0.0.11.02.6.0.0*INTERACT*FIXED

***** END OF LISTING *****

Figure 261. ST06 System status information

The Work with Disk Status screen is shown in Figure 262.

Local (ASM23) / Operating System Monitor: OS/400



5769SS1 V4R5M0 000526Work with Disk Status11/28/00 16:22:21

Elapsed time00:21:24System nameAS23

Unit	Type	Size (M)	Used	I/O	Request Size (K)	Read Rqs	Write Rqs	Read (K)	Write (K)	%	ASP	Type	Status	Protection	Compression
1	6713	6442	81.7	.5	19.9	.1	.3	20.9	19.4	0	1	DPY	ACTIVE		
2	6713	6442	81.7	.2	9.7	.0	.2	29.8	8.8	0	1	DPY	ACTIVE		
3	6713	8589	81.8	.1	7.0	.0	.1	10.5	6.3	0	1	DPY	ACTIVE		
5	6713	6442	81.7	.1	6.8	.0	.1	9.7	6.5	0	1	DPY	ACTIVE		
6	6713	6442	81.8	.1	7.4	.0	.1	22.0	6.2	0	1	DPY	ACTIVE		
7	6713	6442	81.7	.3	8.2	.0	.3	5.8	8.6	0	1	DPY	ACTIVE		
8	6713	6442	81.7	.4	9.2	.0	.3	24.4	5.9	0	1	DPY	ACTIVE		
9	6713	6442	81.7	.5	25.6	.4	.0	30.2	4.0	0	1	DPY	ACTIVE		
10	6713	6442	81.7	.1	5.1	.0	.1	7.5	4.4	0	1	DPY	ACTIVE		
4	6713	8589	46.7	.0	5.3	.0	.0	5.3	0	2	DPY	ACTIVE			

***** END OF LISTING *****

Figure 262. ST06 Disk status information

Figure 263 shows an example of the Display Hardware Resources screen.

Local (ASM23) / Operating System Monitor: OS/400

5769SS1 V4R5M0 000526

Display Hardware Resources

AS23

Processor Resources List

Processor feature code . . . : 2388
 Processor feature : 2388
 Interactive feature :

Resource	Type-Model	Serial Number	Part Number	Location			Text
				Frame ID	EIA	Card Pos	
CEC01	9406-170	10-3BHLM					Main Card Enclosure
PN01	2468-001	00-0000000	0000021F5772	1			System Control Panel
MP01	2388-000	10-8106027	0000097H7855	1		M01	System Processor Card
MP02	2388-000	10-8106027	0000097H7855	1		M01	System Processor Card
SP01	6757-003	10-8307010		1		C01	Service Processor Card
BCC01		00-0000000					Bus Extender
MS05	3004-000	10-8194003	0000090H9359	1		M02M	256MB Main Storage Card
MS06	3004-000	10-8194006	0000090H9359	1		M02B	256MB Main Storage Card
MS07	3004-000	10-8194012	0000090H9359	1		M02K	256MB Main Storage Card
MS08	3004-000	10-8194013	0000090H9359	1		M02G	256MB Main Storage Card
MS09	3004-000	10-8194021	0000090H9359	1		M02C	256MB Main Storage Card
MS10	3004-000	10-8194034	0000090H9359	1		M02H	256MB Main Storage Card
MS11	3004-000	10-8194037	0000090H9359	1		M02L	256MB Main Storage Card
MS12	3004-000	10-8194039	0000090H9359	1		M02A	256MB Main Storage Card
MS13	3004-000	10-8194043	0000090H9359	1		M02D	256MB Main Storage Card
MS14	3004-000	10-8194047	0000090H9359	1		M02J	256MB Main Storage Card
MS15	3004-000	10-8260004	0000090H9359	1		M02F	256MB Main Storage Card
MS16	3004-000	10-8260020	0000090H9359	1		M02E	256MB Main Storage Card

Figure 263. ST06 installed hardware

16.4.5 SAP R/3 system on iSeries: Snapshot information

In addition to the preceding information, ST06 provides additional OS/400 resource utilization information measured over *one minute* before selecting the transaction. However, these measurements alone do not present adequate information to support changing tuning parameters, because the duration of the measurement is too short.

16.4.5.1 CPU snapshot

From the SAP R/3 Detailed analysis menu (Figure 256), click the **CPU** push button within *Snapshot analysis* to see the information in Figure 264.

Local (dfsapbak) / CPU Snapshot

Refresh display

Tue Jan 2 15:26:11 2001# interval 59 sec.

CPU	User in %	System in %	Idle in %
0	41	0	59
1	1-	1-	1-
2	1-	1-	1-
3	1-	1-	1-
4	1-	1-	1-
5	1-	1-	1-
6	1-	1-	1-
7	1-	1-	1-

Figure 264. ST06 CPU Snapshot

A CPU snapshot shows the average CPU load for all main processors in an iSeries server. This is similar to the OS/400 WRKSYSSTS command, which also shows a single value for CPU utilization. The automatic load balancing on the OS/400 system ensures that all processors are evenly utilized.

16.4.5.2 Memory snapshot

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Memory** push button within *Snapshot analysis* to view the information in Figure 265.

Local (dfsapbak) / Memory Snapshot

Refresh display

Tue Jan 2 15:29:15 2001# interval 59 sec.

Physical Memory		Pages in	Pages out	paged in	paged out
free [Kb]	avail [Kb]	[/s]	[/s]	[Kbyte/s]	[Kbyte/s]
N/A	10485760	556	556	2,224	2,224

Figure 265. ST06 Memory Snapshot

16.4.5.3 Disk snapshot

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Disk** push button within *Snapshot analysis* to view the information in Figure 266.

Local (dfsapbak) / Local Disks Snapshot

Refresh display

Tue Jan 2 15:29:15 2001# interval 59 sec.

Disk	Resp. [ms]	Util. [%]	Queue Len.	Wait [ms]	Serv [ms]	Kbyte [/s]	Oper. [/s]
ASP1DSK16	22	3	0	0	22	0	1
ASP1DSK09	19	4	0	0	19	1	2
ASP1DSK20	16	10	0	0	16	3	6
ASP1DSK23	15	6	0	0	15	2	4
ASP1DSK32	15	6	0	0	15	2	4
ASP1DSK51	15	7	0	0	15	2	4
ASP1DSK19	13	4	0	0	13	1	3
ASP1DSK39	13	5	0	0	13	2	4
ASP1DSK40	13	7	0	0	13	2	5
ASP1DSK11	12	2	0	0	12	0	1
ASP1DSK26	12	7	0	0	12	3	6
ASP1DSK48	12	5	0	0	12	1	3
ASP1DSK58	12	6	0	0	12	2	5
ASP1DSK30	11	7	0	0	11	3	6
ASP1DSK77	11	1	0	0	11	0	1
ASP1DSK25	10	6	0	0	10	3	6
ASP1DSK66	10	7	0	0	10	3	7
ASP1DSK68	10	9	0	0	10	4	8
ASP1DSK18	9	3	0	0	9	1	3

Figure 266. ST06 Disk Snapshot

16.4.5.4 LAN snapshot

The LAN Interface display shows the data passing the iSeries LAN adapter. The activity of the LAN interface is not displayed via the native OS/400 interactive commands.

From the SAP R/3 Detailed analysis menu (Figure 256), click the **LAN** push button within *Snapshot analysis* to see the information in the display in Figure 267.

Local (dfsapbak) / LAN Interfaces Snapshot

Refresh display

Tue Jan 2 15:29:15 2001# interval 59 sec.

LAN	Packets/s		Errors/s		Collisions [/s]
	in	out	in	out	
ETHLINE	0	0	0	0	0
	11	0	0	0	0

Figure 267. ST06 LAN Interface Snapshot

16.4.5.5 File system snapshot

From the SAP R/3 Detailed analysis menu (Figure 256), click the **File system** push button within *Snapshot analysis* to see the information in the display shown in Figure 268.

Local (dfsapbak) / Filesystem Snapshot

File system monitor

Tue Jan 2 15:31:19 2001# interval 59 sec.

Filesystem	Capacity [MB]	Free space [MB]	Size in % free
ASP1	488,216	266,481	55
ASP2	8,388	5,293	63

Figure 268. ST06 Filesystem Snapshot

The file system snapshot shows the disk capacity utilization of all auxiliary storage pools (ASPs) defined on the iSeries such as:

- ASP1 = System ASP (SAP R/3 database and executables)
- ASP2 = User ASP (SAP R/3 journal receivers)

16.4.5.6 Memory pool snapshot

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Pool** push button within *Snapshot analysis* to see the information in the display in Figure 269.

Local (dfsapbak) / Pool Snapshot

Refresh display

Tue Jan 2 15:32:21 2001# interval 59 sec.

Pool	max. Act	Size	Mach.resv.	DB-Flt	non-DB-F	DB_Pg	non-DB-Pg	Ac->W	Ac->Ine	Wt->Ine
1	0	585,864	397,280	0.03	4.66	0.03	4.66	11	0	0
2	81	9,604,984	96	42.86	161.43	798.38	236.81	2,475	0	0
3	12	32,768	0	0.06	2.35	0.25	5.95	76	0	0
4	16	262,144	0	0.00	0.00	0.00	0.00	0	0	0

Figure 269. ST06 Pool Snapshot

The pool snapshot shows the memory pools defined on the iSeries, together with the database and non-database faults per second.

Use the scroll bar to see the job transition information.

16.4.6 SAP R/3 system on iSeries: Statistics from previous hours

The following options enable you to see OS/400 performance data over the last 24-hour period if the **OS Collector** is active (job SAPOSCOL, which runs in the R3_<nn> subsystem).

16.4.6.1 CPU: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **CPU** push button within *Previous hours* to see the information in the display in Figure 270.

Local (dfsapbak) / CPU Last 24 Hours

Graphics by column								
Tue Jan 2 15:34:24 2001#								
Hour	CPU utilization in %				Interrupts	System	Context	
	User	System	Idle	IO Wait	/h	calls /h	switches /h	
15	21	0	0	80	N/A	N/A	N/A	
14	21	0	0	79	N/A	N/A	N/A	
13	19	0	0	93	N/A	N/A	N/A	
12	31	1	0	91	N/A	N/A	N/A	
11	27	1	0	83	N/A	N/A	N/A	
10	23	0	0	87	N/A	N/A	N/A	
9	20	0	0	67	N/A	N/A	N/A	
8	10	1	15	79	N/A	N/A	N/A	
7	5	0	0	84	N/A	N/A	N/A	
6	3	0	0	97	N/A	N/A	N/A	
5	4	0	0	98	N/A	N/A	N/A	
4	13	0	0	99	N/A	N/A	N/A	
3	17	0	0	73	N/A	N/A	N/A	
2	13	0	0	87	N/A	N/A	N/A	
1	15	0	0	86	N/A	N/A	N/A	
0	20	1	0	85	N/A	N/A	N/A	
23	2	0	0	97	N/A	N/A	N/A	
22	18	0	0	99	N/A	N/A	N/A	
21	24	1	0	75	N/A	N/A	N/A	
20	3	0	0	100	N/A	N/A	N/A	
19	1	0	0	100	N/A	N/A	N/A	
18	3	0	0	99	N/A	N/A	N/A	
17	3	0	0	100	N/A	N/A	N/A	
16	0	0	0	100	N/A	N/A	N/A	

Figure 270. ST06 CPU Last 24 Hours

16.4.6.2 Memory: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Memory** push button within *Previous hours* to see the information in the display in Figure 271.

Local (dfsapbak) / Memory Last 24 Hours

Graphics by column

Tue Jan 2 15:35:25 2001#							
Hour	Pages in /h	Pages out /h	Paged in [Kb/h]	Paged out [Kb/h]	Free memory/h in Kbyte		
					minimum	maximum	average
15	3256,451	3256451	0	0	N/A	N/A	N/A
14	2732,189	2732189	0	0	N/A	N/A	N/A
13	3707,805	3707805	0	0	N/A	N/A	N/A
12	6867,662	6867662	0	0	N/A	N/A	N/A
11	4055,405	4055405	0	0	N/A	N/A	N/A
10	2867,878	2867878	0	0	N/A	N/A	N/A
9	2529,268	2529268	0	0	N/A	N/A	N/A
8	2000,767	2000767	0	0	N/A	N/A	N/A
7	659,512	659,512	0	0	N/A	N/A	N/A
6	621,941	621,941	0	0	N/A	N/A	N/A
5	694,049	694,049	0	0	N/A	N/A	N/A
4	253,225	253,225	0	0	N/A	N/A	N/A
3	279,410	279,410	0	0	N/A	N/A	N/A
2	915,722	915,722	0	0	N/A	N/A	N/A
1	417,185	417,185	0	0	N/A	N/A	N/A
0	809,866	809,866	0	0	N/A	N/A	N/A
23	148,295	148,295	0	0	N/A	N/A	N/A
22	787,859	787,859	0	0	N/A	N/A	N/A
21	1050,893	1050893	0	0	N/A	N/A	N/A
20	2590,123	2590123	0	0	N/A	N/A	N/A
19	56,253	56,253	0	0	N/A	N/A	N/A
18	200,551	200,551	0	0	N/A	N/A	N/A
17	263,685	263,685	0	0	N/A	N/A	N/A
16	54,210	54,210	0	0	N/A	N/A	N/A

Figure 271. ST06 Memory Last 24 Hours

16.4.6.3 Disk: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Disk** push button within *Previous hours* to view the information shown in Figure 272.

Local (dfsapbak) / Disk ASP1DSK01 Last 24 Hours

Disk-- Disk++ Graphics by column								
Tue Jan 2 15:36:27 2001#								
Hour	Utiliz. [%]	Queue length	Wait [ms]	Srv. Time [ms]	Oper. [Mbyte/h]	Oper. [Kbyte/s]	Oper. [Kbyte/s]	Oper. [Kbyte/s]
15	1	0	0	3	9	19,916	3	6
14	1	0	0	3	9	19,956	3	6
13	2	0	0	3	12	24,363	3	7
12	3	0	0	3	19	39,327	5	11
11	2	0	0	2	13	28,012	4	8
10	1	0	0	3	11	22,736	3	6
9	2	0	0	3	12	24,830	3	7
8	1	0	0	3	7	15,308	2	4
7	0	0	0	2	4	8,813	1	2
6	1	0	0	3	6	12,636	2	4
5	1	0	0	4	4	9,824	1	3
4	0	0	0	2	2	5,931	1	2
3	0	0	0	4	2	4,603	1	1
2	0	0	0	2	3	6,617	1	2
1	0	0	0	3	3	6,665	1	2
0	0	0	0	3	4	8,365	1	2
23	0	0	0	3	1	4,008	0	1
22	0	0	0	4	2	5,711	1	2
21	1	0	0	5	3	7,252	1	2
20	1	0	0	3	7	15,003	2	4
19	0	0	0	2	1	3,567	0	1
18	0	0	0	3	2	5,132	1	1
17	0	0	0	2	2	6,013	1	2
16	0	0	0	1	1	2,793	0	1

Figure 272. ST06 Disk Last 24 Hours

16.4.6.4 LAN: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **LAN** push button within *Previous hours* to view the information shown in Figure 273.

Local (dfsapbak) / LAN Interfaces Snapshot

Refresh display

Tue Jan 2 15:29:15 2001# interval 59 sec.

LAN	Packets/s		Errors/s		Collisions
	in	out	in	out	[/s]
ETHLINE	0	0	0	0	0
	11	0	0	0	0

Figure 273. ST06 LAN Interfaces Snapshot last 24 hours

16.4.6.5 File system: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **File system** push button within *Previous hours* to view the information shown in Figure 274.

Local (dfsapbak) / File ASP1 Last 24 Hours

Previous file system | Next file system | File system monitor

Tue Jan 2 15:40:33 2001#

Hour	Filesystem	Capacity [MB]	Free space [MB]	Size in % free
15	ASP1	488,216	266,482	55
14	ASP1	488,216	266,646	55
13	ASP1	488,216	266,647	55
12	ASP1	488,216	264,993	54
11	ASP1	488,216	266,738	55
10	ASP1	488,216	267,725	55
9	ASP1	488,216	268,196	55
8	ASP1	488,216	268,904	55
7	ASP1	488,216	269,264	55
6	ASP1	488,216	269,307	55
5	ASP1	488,216	267,774	55
4	ASP1	488,216	267,859	55
3	ASP1	488,216	266,977	55
2	ASP1	488,216	267,310	55
1	ASP1	488,216	267,230	55
0	ASP1	488,216	267,318	55
23	ASP1	488,216	267,395	55
22	ASP1	488,216	267,389	55
21	ASP1	488,216	267,094	55
20	ASP1	488,216	269,201	55
19	ASP1	488,216	269,204	55
18	ASP1	488,216	269,250	55
17	ASP1	488,216	269,364	55
16	ASP1	488,216	269,506	55

Figure 274. ST06 File system Last 24 Hours

The file system shows the disk capacity usage of all auxiliary storage pools (ASPs) defined on the iSeries server such as:

- ASP1 = System ASP (SAP R/3 database and executables)
- ASP2 = User ASP (SAP R/3 journal receivers)

16.4.6.6 Memory pool: Previous hours

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Pool** push button within *Previous hours* to view the information shown in Figure 275.

Local (dfsapbak) / Pool No: 1 Last 24 Hours									
POOL-- POOL++									
Tue Jan 2 15:41:35 2001#									
Hour	max.Act	Size	Mach.resv.	DB-Flt non-DB-F		DB_Pg	non-DB-Pg	Ac->W	Ac->Ine
15	81	9,604,984	1,632	28.18	209.04	1082.29	256.85	1,167	0
14	81	9,604,984	96	22.69	232.76	934.83	279.09	1,274	0
13	81	9,604,984	96	27.15	354.24	1170.95	418.62	1,135	0
12	81	9,604,984	96	36.70	805.71	1995.91	1026.59	1,061	0
11	81	9,604,984	96	29.55	287.04	1426.18	363.09	1,159	0
10	81	9,604,984	1,632	26.54	217.05	996.76	271.45	1,252	0
9	81	9,604,984	96	22.81	172.56	894.75	211.05	1,226	0
8	81	9,604,984	1,632	15.79	143.49	677.09	186.78	858	0
7	81	9,604,984	96	11.25	29.97	224.25	52.30	527	0
6	81	9,604,984	96	31.62	22.39	213.27	37.99	400	0
5	81	9,604,984	96	27.43	24.49	224.36	41.28	394	0
4	81	9,604,984	96	5.04	13.81	65.49	39.73	324	0
3	81	9,604,984	96	5.74	11.90	86.01	38.59	319	0
2	81	9,604,984	96	7.82	14.24	378.21	39.38	273	0
1	81	9,604,984	1,632	4.73	47.12	97.42	75.58	312	0
0	81	9,604,984	96	20.39	40.43	272.93	75.39	279	0

Figure 275. ST06 Pool Last 24 Hours

The pool snapshot shows the memory pools defined on the iSeries server together with the database and non-database faults per second. Use the scroll bar to see the job transition information.

16.4.7 SAP R/3 system on: Performance database

SAP R/3 also provides additional information from a performance database on the overall performance of the iSeries servers.

16.4.7.1 Server performance: Recent days

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Compare recent days** push button within *Performance database* to view the information shown in Figure 276.

Local (dfsapbak) / Compare Recent Days For Server dfsapbak

Previous server Next server Details Sort Graphics by column											
Compare recent days for server dfsapbak Date/time 02.01.2001 15:43:40 Average values for peak workload 10-13 and 14-15											
Date	Usr	Sys	Id1	Pages in/h		Pages out/h		Swap space		Lan packets/h	
	avg	avg	avg	max	avg	max	avg	max size	min free	max	avg
20010101	1	0	100	271,462	101,204	271,462	101,204	1-	1-	288,484	288,484
20001231	5	0	95	420,501	223,584	420,501	223,584	1-	1-	288,484	288,484
20001230	5	0	95	403,278	265,081	403,278	265,081	1-	1-	288,562	288,562
20001229	21	1	79	3,201,196	2,652,128	3,201,196	2,652,128	1-	1-	283,841	283,841
20001228	13	1	87	2,214,483	1,815,594	2,214,483	1,815,594	1-	1-	0	0
20001227	18	0	82	1,999,049	1,604,788	1,999,049	1,604,788	1-	1-	0	0
20001226	6	0	94	354,453	232,994	354,453	232,994	1-	1-	0	0
20001225	5	0	95	256,057	149,598	256,057	149,598	1-	1-	288,562	288,562
20001224	6	0	94	200,791	180,551	200,791	180,551	1-	1-	288,562	288,562
20001223	11	0	90	265,466	206,734	265,466	206,734	1-	1-	293,281	293,281
20001222	28	1	72	3,143,800	1,997,463	3,143,800	1,997,463	1-	1-	0	0
20001221	20	1	80	2,468,479	2,150,877	2,468,479	2,150,877	1-	1-	0	0
20001220	26	1	73	4,111,823	3,314,229	4,111,823	3,314,229	1-	1-	293,363	290,963
20001219	29	1	71	3,210,558	2,524,195	3,210,558	2,524,195	1-	1-	0	0
20001218	26	1	74	4,636,459	3,665,554	4,636,459	3,665,554	1-	1-	288,562	288,536
20001217	3	0	97	371,121	176,964	371,121	176,964	1-	1-	288,562	288,562
20001216	3	0	97	307,544	269,839	307,544	269,839	1-	1-	0	0
20001210	26	0	74	656,478	529,527	656,478	529,527	1-	1-	0	0
20001209	24	7	74	283,834	188,305	283,834	188,305	1-	1-	0	0
20001208	37	1	62	5,479,513	2,845,665	5,479,513	2,845,665	1-	1-	0	0
20001130	36	1	63	5,137,080	4,386,729	5,137,080	4,386,729	1-	1-	0	0
20001129	72	3	25	8,001,269	5,663,021	8,001,269	5,663,021	1-	1-	0	0
20001128	88	6	6	19447,989	13527,901	19447,989	13527,901	1-	1-	0	0
20001127	83	4	13	14967,204	12877,267	14967,204	12877,267	1-	1-	0	0
20001126	9	0	91	582,297	410,945	582,297	410,945	1-	1-	0	0
20001125	6	0	94	329,039	258,758	329,039	258,758	1-	1-	0	0
20001124	69	4	28	11304,705	8,144,298	11304,705	8,144,298	1-	1-	0	0
20001123	66	4	31	14635,309	9,413,114	14635,309	9,413,114	1-	1-	0	0
20001122	48	3	49	7,134,043	5,707,341	7,134,043	5,707,341	1-	1-	0	0
20001121	80	2	23	12486,989	7,461,390	12486,989	7,461,390	1-	1-	0	0

Figure 276. ST06 Server statistics: Daily summary

16.4.7.2 Server comparison: Previous day

From the SAP R/3 Detailed analysis menu (Figure 256), click the **Compare all servers** push button within *Performance database* to view the information in Figure 277.

Local (dfsapbak) / Compare All Servers For Date 20001229

Previous day Next day Details Sort Graphics by column											
Compare servers for date 20001229 Date/time 02.01.2001 16:32:43 Average values for peak workload 10-13 and 14-15											
Server	Usr	Sys	Id1	Pages in/h		Pages out/h		Swap space		Lan packets Lan errors/h	
	avg	avg	avg	max	avg	max	avg	max size	min free	max	avg
dfsapbak	21	1	79	3,201,196	2,652,128	3,201,196	2,652,128	1-	1-	283,841	283,841
										0	0

Figure 277. ST06 server statistics: All servers

16.4.8 Active jobs on the iSeries server

Use the OS/400 WRKACTJOB command, the WRKSBSJOB SBS (R3_<nn>) command, or the SAP R/3 transaction SM50 to review the active SAP R/3 jobs.

16.4.8.1 OS/400 WRKACTJOB command

The WRKACTJOB command (Figure 278) shows all of the jobs running on the iSeries server within each of the active subsystems. It includes the following information:

- Subsystem
- User identifier
- Job type
- Memory pool
- Job priority
- CPU usage (%)
- CPU usage (seconds)
- Disk I/O count

Work with Active Jobs						PRDSYS
						12/02/00 11:14:05
CPU %:	18.7	Elapsed time:	00:00:16	Active jobs:	421	
Opt	Subsystem/Job	User	Type	CPU %	Function	Status
	QBATCH	QSYS	SBS	.0		DEQW
	QCMN	QSYS	SBS	.0		DEQW
	QCTL	QSYS	SBS	.0		DEQW
	QSYSSCD	QPGMR	BCH	.0	PGM-QEZSCNEP	EVTW
	QINTER	QSYS	SBS	.0		DEQW
	LID02526	QSYSOPR	INT	.0	CMD-DSPMSG	DSPW
	QPADEV0005	PRDOFR	INT	.0	MNU-R3MAIN	DSPW
	QPADEV0006	QSYSOPR	INT	.0	CMD-WRKSBSJOB	DSPW
	QPADEV0008	PRDOFR	INT	.0	CMD-WRKACTJOB	RUN
	QSERVER	QSYS	SBS	.0		DEQW
	QPWFSERVSD	QUSER	BCH	.0		SELW
	QSERVER	QPGMR	ASJ	.0		EVTW
	QZDASRVSD	QUSER	BCH	.0		SELW
	QSOC	QSYS	SBS	.0		DEQW
	SOCMGR	QSOC	ASJ	.0	PGM-QYYCMGR	DEQW
	QSPL	QSYS	SBS	.0		DEQW
	ENINETPJL	QSPLJOB	WTR	.0		SIGW

Figure 278. Work with active jobs

Use F11 to view additional information (Figure 279) on the jobs running on the iSeries server.

Work with Active Jobs										PRDSYS
										12/02/00 11:14:05
CPU %:	18.7	Elapsed time:		00:00:16	Active jobs:		421			
Opt	Subsystem/Job	Type	Pool	Pty	CPU	Int	Rsp	AuxIO	CPU %	
	QBATCH	SBS	2	0	.5			0	.0	
	QCMN	SBS	2	0	.3			0	.0	
	QCTL	SBS	2	0	.2			0	.0	
	QSYSSCD	BCH	2	10	.2			0	.0	
	QINTER	SBS	2	0	1.2			0	.0	
	ENT01824	INT	4	20	.4	0	.0	0	.0	
	LID02526	INT	4	20	.0	0	.0	0	.0	
	QPADEV0005	INT	4	20	.0	0	.0	0	.0	
	QPADEV0006	INT	4	20	.1	0	.0	0	.0	
	QPADEV0008	INT	4	20	.1	2	.0	17	.0	
	QPADEV0009	INT	4	20	69.3	0	.0	0	.0	
	QPGMR	SBS	2	0	.0			0	.0	
	QSERVER	SBS	2	0	.3			0	.0	
	QPWFSEVRSD	BCH	2	20	.0			0	.0	
	QSERVER	ASJ	2	20	.0			0	.0	
	QZDASRVSD	BCH	2	20	.0			0	.0	
	QSOC	SBS	2	0	1.2			0	.0	
	SOCMGR	ASJ	2	15	15.2			0	.0	
	QSPL	SBS	2	0	.9			0	.0	
	ENINETPJL	WTR	3	15	108.1			0	.0	

Figure 279. WRKACTJOB F11 (Display elapsed data)

By default, all SAP R/3 jobs run in memory pool 2 (*BASE). With the exception of the Dispatcher, Gateway, Message Server, and Enqueue jobs, which run at priority 12, the R/3 work process jobs run at a priority of 20. The memory pool that *all* work processes run in can be changed by modifying the appropriate subsystem description for pool definition and the routing entry.

The run priority for *all* the SAP R/3 work processes in an SAP R/3 instance can be modified by changing the associated values in the *class* used by the OS/400 subsystem description. It is also possible to assign priorities by work process class.

An SAP R/3 instance has a unique number within each iSeries server and is assigned during creation of the instance. The OS/400 subsystem under which an SAP R/3 instance runs is denoted by R3_<nn>, where <nn> is the instance number.

You can limit the jobs that appear when using the WRKACTJOB command to a specific SAP R/3 instance by specifying the parameter SBS(R3_<nn>), where <nn> is the instance number.

If you want to see the same information presented in order (by CPU usage, for example), position the cursor on the CPU % column and press F16.

16.4.8.2 SAP R/3 system on iSeries: Top CPU user's display

An SAP R/3 display gives you similar information about all jobs running on the iSeries system sorted by CPU consumption. This function can be selected by clicking the Top CPU processes button on the SAP R/3 Detailed Analysis menu (Figure 256) from within *Snapshot analysis*. The screen shown in Figure 280 appears.

Local (dfsaphak) / Top CPU Users

Refresh display						
Tue Jan 2 15:41:35 2001# interval 59 sec.						
Job no.	Username	Command	CPU Util [%]	CPU Time [s]	Resident size [kB]	Prior.
883478	PRD02	DISP	6.95	1131:36	N/A	23
891733	PRD02	DISP	6.62	25:56	N/A	20
883493	QUSER	QXDARECVR	2.75	666:12	N/A	23
890636	PRD02	DISP	1.46	148:29	N/A	20
890637	QUSER	QXDARECVR	1.22	77:42	N/A	20
891835	PRD02	DISP	1.19	10:47	N/A	23
883772	ODSOWNER	ODS_SEND01	0.89	3128:57	N/A	20
883490	PRD02	DISP	0.66	63:16	N/A	20
890737	PRD02	DISP	0.55	89:23	N/A	20
883511	PRD02	DISP	0.49	319:36	N/A	20
891734	QUSER	QXDARECVR	0.39	3:22	N/A	20
883544	QUSER	QXDARECVR	0.28	564:23	N/A	20
891836	QUSER	QXDARECVR	0.25	4:34	N/A	20
890739	QUSER	QXDARECVR	0.20	48:15	N/A	20
883766	ODSOWNER	PRDLNK_RDR	0.19	102:44	N/A	20
8	*SYS	CFINT01	0.18	154:58	N/A	0
883484	PRD02	DISP	0.17	287:13	N/A	20
883481	PRD02	DISP	0.16	597:24	N/A	20
883531	PRD02	DISP	0.15	117:40	N/A	20
883494	PRD02	DISP	0.13	101:19	N/A	20
883508	PRD02	DISP	0.13	29:19	N/A	20
883473	PRD02	AS4RDIREXE	0.12	64:26	N/A	12
883486	PRD02	DISP	0.10	166:35	N/A	20
883773	ODSOWNER	ODS_EXCEPT	0.09	1051:28	N/A	20
883558	QUSER	QXDARECVR	0.09	49:22	N/A	20
7	*SYS	CFINT02	0.09	67:59	N/A	0

Figure 280. ST06 Top CPU processes

The values under “Command” are actually OS/400 job names.

16.4.8.3 OS/400 WRKSBSJOB command

The OS/400 Work with Subsystem Job (WRKSBSJOB) command with the parameter SBS(R3_<nn>) shows the SAP R/3 subsystem for instance <nn> together with the active SAP R/3 jobs.

To identify the SAP R/3 work process type and process ID (PID) of each OS/400 server job, from R/3 Release 4.6, you can use R/3 transaction SM50. The number *nn* in the OS/400 job name WPnn equates to the first column of the SM50 Process overview display.

16.4.8.4 OS/400 job number and SAP R/3 PIDs

Using SAP R/3 transaction SM50, you can see the SAP R/3 work processes and their process identifiers (PID). This is in contrast to the WRKACTJOB command, where the jobs are listed by their OS/400 job names and job numbers. If you want to work with the OS/400 job attributes of an R/3 work process, you need to determine the corresponding PID using SM50 (Figure 281), which shows the active SAP R/3 work processes.

No	Ty.	PID	Status	Reasn	Start	Err	Sem	CPU	Time	Report	Cl.	User	Action
0	DIA	2427	waiting		Yes								
1	DIA	2428	Running		Yes					SAPLTHFB	300	MICHAEL	
2	DIA	2429	waiting		Yes								
3	DIA	2430	waiting		Yes								
4	DIA	2431	waiting		Yes								
5	DIA	2432	waiting		Yes								
6	DIA	2433	waiting		Yes								
7	DIA	2434	waiting		Yes								
8	UPD	2435	waiting		Yes								
9	UPD	2436	waiting		Yes								
10	ENQ	2437	waiting		Yes								
11	BGD	2438	waiting		Yes								
12	BGD	2439	waiting		Yes								
13	BGD	2440	waiting		Yes								
14	SPO	2441	waiting		Yes								
15	UP2	2442	waiting		Yes								

Figure 281. SM50 Process overview

From SAP R/3 Release 4.6, the OS/400 job names of SAP R/3 work processes are in the format "WP<nn>", where <nn> is the work process number. For example, the OS/400 job name for work process 1 would be WP01. The job name DISP is used for the R/3 Dispatcher job.

For releases before SAP R/3 Release 4.6, use the OS/400 Work with a Job by PID (WRKPID) command (delivered with SAP R/3 kernel) to obtain OS/400 job information for the required PID. This OS/400 command shown in Figure 282 provides you with the job details about active SAP R/3 work processes. Note that all R/3 work process jobs are called DISP in releases before SAP R/3 Release 4.6.

Work with Job by PID (WRKPID)

Type choices, press Enter.

Process ID	2427	Character value
Prompt command	*YES	*YES, *NO, Y, N
xxxJOB command	WRK	Character value

Bottom

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

Figure 282. Work with Job by PID (WRKPID)

You can change the value in the prompt for the xxxJOB command from *WRK* to *DSP* or *CHG*.

Press the Enter key, and you see the prompts for the corresponding command for the specified PID as shown in Figure 283.

Work with Job (WRKJOB)

Type choices, press Enter.

Job name	> WP00	Name, *
User	> R0101	Name
Number	> 018160	000000-999999
Output	*	*, *PRINT
Option	*SELECT	*SELECT, *STSA, *DFNA...

Bottom

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys

Figure 283. Work with Job (WRKJOB)

16.4.9 OS/400 system log

The iSeries maintains a log that provides you with information about the activity on the system including starting and ending jobs and special messages.

16.4.9.1 OS/400 DSPLOG command

This command shown in Figure 284 allows you to specify the beginning and ending dates and times for the period in which you want to review the system log.

Display Log (DSPLOG)

Type choices, press Enter.

Log	QHST	QHST
Time period for log output:		
Start time and date:		
Beginning time	*AVAIL	Time, *AVAIL
Beginning date	*CURRENT	Date, *CURRENT, *BEGIN
End time and date:		
Ending time	*AVAIL	Time, *AVAIL
Ending date	*CURRENT	Date, *CURRENT, *END
Output	*	*, *PRINT, *PRIWRAP

Figure 284. Display Log (DSPLOG)

The screen in Figure 285 shows the messages displayed in the history log when an R/3 system is ended.

```

Display History Log Contents

Job 018170/R0101/WP10 ended on 11/29/00 at 10:48:02; 1 seconds used; end code
Job 018175/R0101/WP15 ended on 11/29/00 at 10:48:02; 1 seconds used; end code
Job 018167/R0101/WP07 ended on 11/29/00 at 10:48:02; 1 seconds used; end code
Job 018166/R0101/WP06 ended on 11/29/00 at 10:48:02; 1 seconds used; end code
Job 018164/R0101/WP04 ended on 11/29/00 at 10:48:03; 2 seconds used; end code
Job 018174/R0101/WP14 ended on 11/29/00 at 10:48:04; 133 seconds used; end co
Job 018163/R0101/WP03 ended on 11/29/00 at 10:48:06; 6 seconds used; end code
Job 018169/R0101/WP09 ended on 11/29/00 at 10:48:06; 1 seconds used; end code
Job 018168/R0101/WP08 ended on 11/29/00 at 10:48:07; 1 seconds used; end code
Job 018162/R0101/WP02 ended on 11/29/00 at 10:48:08; 34 seconds used; end cod
Job 018159/R0101/GWRD ended on 11/29/00 at 10:48:09; 12 seconds used; end cod
Job 018173/R0101/WP13 ended on 11/29/00 at 10:48:19; 9 seconds used; end code
Job 018176/QUSER/QXDARECVR ended on 11/29/00 at 10:48:36; 1 seconds used; end
Job 018165/R0101/WP05 ended on 11/29/00 at 10:48:37; 56 seconds used; end cod
Job 018172/R0101/WP12 ended on 11/29/00 at 10:48:39; 130 seconds used; end co
Job 019460/QUSER/QXDARECVR ended on 11/29/00 at 10:48:58; 1 seconds used; end
Job 018161/R0101/WP01 ended on 11/29/00 at 10:48:59; 583 seconds used; end co
More...

Press Enter to continue.

F3=Exit   F10=Display all   F12=Cancel

```

Figure 285. Display History Log Contents

More detailed information for any message is available through the F1 help function key by positioning the cursor on the required line.

16.4.9.2 OS/400 DSPMSG command

This command allows you to display messages for any OS/400 user profile. When you use the MSGQ(QSYSOPR) parameter, you can view the messages sent from the iSeries to the system operator as shown in Figure 286.

```

Work with Messages
System:  DEVSYS

Messages in:  QSYSOPR

Type options below, then press Enter.
5=Display details and reply

Opt  Message
Messages needing a reply
(No messages available)

Messages not needing a reply
Storage limit exceeded for ASP 2.
Load form type 'X_65_80' device HP4050TN writer HP4050TN.  (G B I H R C)
Reply . . :  G
Cleanup has completed.
Cleanup of job logs and other system output successfully completed.
Cleanup of system journals and system logs successfully completed.
30 days used for problem log cleanup value.
Cleanup of system journals and system logs started.
More...

F1=Help   F3=Exit   F5=Refresh   F12=Cancel   F17=Top   F18=Bottom
F21=Select assistance level   F22=Display list details

```

Figure 286. Display system operator messages

The screen in Figure 286 shows the display when using Basic Assistance level. If using Intermediate assistance (change level using the F21 key), more detailed information is available by using the F1 Help function key. Full message details are available including the message identifier and the program sending the message.

16.4.9.3 SAP R/3 system on iSeries: Operating system log

From the ST06 SAP R/3 Detailed Analysis menu (Figure 256), click the **OS log** push button within *Previous hours* to see the information shown in the display in Figure 287.




Figure 287. ST06 Operating System Log

This information is the same as the information shown by the DMSMSG MSGQ(QSYSOPR) command. However, it is not possible to drill down further on an individual message or respond to messages needing a reply.

16.4.10 SAP R/3 system on iSeries: Additional functions

The SAP R/3 transaction ST06 has some additional options that are described in the following section.

16.4.10.1 SAP R/3 system on iSeries: Parameter changes

From the ST06 SAP R/3 Detailed analysis menu (Figure 256), click the **Parameter Changes** push button within *Additional functions* to see the changes made to the operating system parameters relevant for R/3 (Figure 288).




Figure 288. ST06 Parameter Changes

16.4.10.2 SAP R/3 system on iSeries: LAN check

From the ST06 SAP R/3 Detailed analysis menu (Figure 256), click the **LAN Check by Ping** push button within *Additional functions* to see the presentation servers (PCs or front-end stations), application servers, and database servers in the R/3 system. You can then see the results in milliseconds of a 1 or 10 try ping of the server (Figure 289).




Figure 289. ST06 LAN Check

16.4.10.3 SAP R/3 system on iSeries: PTF check

From the ST06 SAP R/3 Detailed analysis menu (Figure 256), click the **PTF check** push button within *Additional functions* to check the installed operating system PTFs against the required PTFs for SAP R/3 that are listed in the IBM

Informational APAR for the operating system release. It is checked against the contents of the /usr/sap/trans/config/infoapar.450 file if your operating system release is V4R5. The latest version of the Informational APAR for V4R5 can be downloaded to your system providing you have Electronic Customer Support, using the following OS/400 commands:

```
SNDPTFORD PTFID((II12399 INFOAS4 V4R5M0))
```

```
CPYTOSTMF FROMMBR ('/QSYS.LIB/QGPL.LIB/QAPZCOVER.FILE/QII12399.MBR')
TOSTMF ('/usr/sap/trans/config/infoapar.450')
```

Refer to SAP Note 144149 *AS/400: Compare PTF status with Informational APAR* for the correct command for earlier OS/400 releases. See Figure 290.




Figure 290. ST06 PTF check

16.4.11 Collecting iSeries performance data

Performance data can be collected on the iSeries server with the Performance Collector. This facility is available as a standard feature of the operating system and is accessed via Operations Navigator. We recommend you use this feature of Operations Navigator to collect performance data. Refer to Chapter 10, "iSeries system administration using GUI" on page 215, for more information on the Performance Collector.

In *OS/400 release V4R5 and earlier*, it is also possible to use the STRPFRMON command. In OS/400 V5R1, the STPFRMON command has been replaced with Collection Services.

In this case, we recommend you start the OS/400 Performance Monitor *after you start the SAP R/3 systems* and allow the R/3 application to run for a while to allow the SAP internal buffers to stabilize. End the OS/400 performance monitor before you stop SAP R/3 system. This gives you data that covers normal user activity and excludes any overhead involved in starting and ending the SAP R/3 system.

To start the performance monitor from the OS/400 Main menu, type the `STRPFRMON` command and press F4. Figure 291 shows the parameters.

Start Performance Monitor (STRPFRMON)

Type choices, press Enter.

Member	> PRD1	Name, *GEN
Library	QPFRDATA	Name
Text 'description'	> 'SAP Production system peak time'	

Time interval (in minutes)	> 5	5, 10, 15, 20, 25, 30, 35...
Stops data collection	*ELAPSED	*ELAPSED, *TIME, *NOMAX
Days from current day	0	0-9
Hour	> 1	0-999
Minutes	0	0-99
Data type	*ALL	*ALL, *SYS
Select jobs	*ALL	*ALL, *ACTIVE
Trace type	*NONE	*NONE, *ALL
Dump the trace	*YES	*YES, *NO
Job trace interval	.5	.5 - 9.9 seconds
Job types	*DFT	*NONE, *DFT, *ASJ, *BCH...

+ for more values

More...

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys

Figure 291. Start Performance Monitor (STRPFRMON)

The arrows indicate the fields that you may want to change on this display, such as the duration of the collection and the time interval between each collection of performance monitoring information. We advise that you obtain the instructions from someone with OS/400 performance analysis expertise to provide you with guidance on collecting and analyzing the information.

16.4.12 iSeries Performance Tools

The IBM licensed program Performance Tools/400, 5769-PT1, is a set of commands and programs that provides tools for analysis, printing detailed reports, and graphical views of iSeries resource usage. This section provides an overview of analyzing the performance of an iSeries server running SAP R/3. Please note the following points:

- Any transaction count and average response for interactive jobs that you may see in any of the reports or displays *do not* refer to the SAP R/3 user transactions.
- All SAP R/3 jobs on the iSeries server are server jobs and appear as non-interactive jobs in all the iSeries Performance Tools reports.
- The iSeries Performance Tools reports give you an indication of resource utilization across the entire system and SAP R/3 as a whole.
- Use SAP R/3 transactions, such as ST03 and SM04 shown in 16.6.1.8, “ST03 Workload analysis” on page 443, and 16.6.1.2, “SM04 User list” on page 439, to obtain information about SAP R/3 user transactions (dialog steps) and response times.
- You cannot be sure that the sequence of the OS/400 server jobs corresponds to the sequence in which SAP R/3 creates them when R/3 is started. That is

because some work processes may be terminated and recreated during SAP R/3 operation. Use the `WRKPID` command to identify the OS/400 job corresponding to an SAP R/3 work process. From R/3 Release 4.6, the work processes jobs are renamed WPnn, where *nn* is the work process number shown in SM50.

Apart from the `WRKSYSACT` command, all options use data gathered by the Performance Monitor.

16.4.12.1 **WRKSYSACT**

The Work with System Activity (`WRKSYSACT`) command, which is installed with the Performance Tools licensed program, allows you to interactively work with the jobs and tasks currently running in the system. It gives a snapshot of user jobs and system tasks using CPU and disk resources being refreshed automatically. See Figure 292.




Figure 292. Work with System Activity

16.4.12.2 **Performance Tools review options**

This section introduces a few of the important features of the IBM licensed program Performance Tools, 5769-PT1. For more information about using these tools, please refer to *OS/400 Performance Tools/400*, SC41-5340.

If you installed Performance Tools on your system, the menu shown in Figure 293 appears when you enter the `GO PERFORM` command.

PERFORM	IBM Performance Tools for AS/400	System: DEVSYS

Select one of the following:

1. Select type of status
2. Collect performance data
3. Print performance report
4. Capacity planning/modeling
5. Performance utilities
6. Configure and manage tools
7. Display performance data
8. System activity
9. Performance graphics
10. Advisor

70. Related commands

Selection or command
====>

F3=Exit F4=Prompt F9=Retrieve F12=Cancel F13=Information Assistant
F16=System main menu
(C) COPYRIGHT IBM CORP. 1981, 2000.

Figure 293. GO PERFORM: Performance data

Some of the main options are explained here:

- **Print performance report** enables you to print useful information such as:
 - **System report** that shows the major system resources and their utilizations, including processor, memory, and disk.
 - **Component report** that presents a more detailed analysis of the system resources and their utilization. This report also shows details of each job (or work process) running on the iSeries and the amount of resources used.
 - **Pool report**, which gives useful information on processor activity by memory pool for each measured interval, which can help identify possible resource contention.
- **Display performance data** allows you to review performance information interactively. However, we recommend that you do this on a heavily loaded system, since you may impact user's response times.
- **Performance graphics** enables you to view and print a selection of graphs to assist in identifying potential iSeries resource bottlenecks.
- **Advisor** analyzes performance data that you collect with the performance monitor. It can also produce conclusions and recommendations to help improve performance. The advisor can help you improve system performance, but does not identify or fix all performance problems.

For more detailed information on facilities available with the Performance Tools, please refer to *OS/400 Performance Tools/400*, SC41-5340.

16.4.13 Insight SAP performance reports

The Insight program is available from IBM to collect performance data from the iSeries server running the R/3 production system and download this to a dedicated PC. The data is usually collected for a period of between 4 and 10 days and is then compressed and sent to IBM for analysis. Reports showing periods of peak activity and breakdown of workload by SAP R/3 module can be produced.

The software program can be downloaded from the Web site

<http://www.ibm.com/erp/sap/insight> where you can find more information, or you can obtain more information by sending e-mail to: ibmerp@us.ibm.com

16.5 iSeries system tuning

This section describes a few techniques you may use to improve your iSeries system performance. Naturally, before making any tuning adjustments, review the OS/400 performance monitor data both interactively using WRKSYSSTS and WRKDSKSTS (refer to 16.4.4, “OS/400 system status and disk status” on page 388), and using reports from the OS/400 Performance Tools (refer to 16.4.12, “iSeries Performance Tools” on page 409).

16.5.1 Initial iSeries tuning

Using the guidelines for the key performance indicators on the iSeries system (refer to 16.4.1, “iSeries performance indicator guidelines” on page 380), make changes to OS/400 parameters one at a time. Monitor the effect after each change to ensure that your changes are improving performance.

16.5.1.1 Manual tuning of pools and activity levels

If you tune the memory pools manually, ensure that QPFRADJ=0 (automatic tuning off) so your changes are not overridden by the operating system. Use the OS/400 `WRKSYSSTS` command and make changes to:

- **Pool Sizes:** Memory allocated to each memory pool.
- **Max Active:** Number of threads that can be active in a storage pool.

Make sure that the number of memory faults is minimized and within guidelines. This screen is shown in Figure 294.

```

Work with System Status
DEVSYS
11/28/00 10:52:59
% CPU used . . . . . :      88.7  Auxiliary storage:
% DB capability . . . . :      24.4  System ASP . . . . . :    488.2 G
Elapsed time . . . . . : 00:10:52  % system ASP used . . :    68.4450
Jobs in system . . . . :    4384  Total . . . . . :    496.6 G
% perm addresses . . . . :    .027  Current unprotect used :    38232 M
% temp addresses . . . . :    .500  Maximum unprotect . . :    40407 M

Type changes (if allowed), press Enter.

System   Pool   Reserved   Max   -----DB-----  ---Non-DB---
Pool   Size (M)  Size (M)  Active Fault Pages  Fault  Pages
  1     489.87   325.53   +++++   .6   2.0   16.6   19.0
  2    5318.12    .08     81  45.1  4729  475.9  599.0
  3     32.00    .00     12   .1   .9    .6    1.5
  4     256.00    .00     16   .0   .1    .1    .1

Command
===>
F3=Exit  F4=Prompt      F5=Refresh  F9=Retrieve  F10=Restart
F11=Display transition data  F12=Cancel  F24=More keys

```

Figure 294. Work with System Status

16.5.1.2 Automatic tuning of pools and activity levels

Alternatively you can use OS/400 automatic performance adjustment facility by setting the system value QPFRADJ to a value other than 0 using the CHGSYSVAL or WRKSYSVAL command. A change to this system value takes effect immediately.

Values of 1 or 2 for QPFRADJ result in a second memory pool of the QSPL and QBASE subsystems to be changed to shared pools *SPOOL and *INTERACT respectively. QPFRADJ values of 1, 2, or 3 change the machine, interactive, and spool memory pool sizes. Values of 2 or 3 in QPFRADJ include the shared memory pools in the memory management process.

Table 40 gives an indication of the effect of the various values of QPFRADJ.

Note: For more information, refer to the OS/400 online help text.

Table 40. QPFRADJ values

Change to settings	QPFRADJ value			
	0 (Automatic)	1 (IPL Only)	2 (IPL and automatic)	3 (automatic)
Subsystems QSPL, QBASE second memory pools changed to *SPOOL, *INTERACT		X	X	
QMCHPOOL		X	X	X
QBASACTLVL		X	X	X
*SPOOL size and activity level		X	X	X
*INTERACT size and activity level		X	X	X

Change to settings	QPFRADJ value			
	0 (Automatic)	1 (IPL Only)	2 (IPL and automatic)	3 (automatic)
*SHRPOOLs size and activity levels			X	X

When you use the automatic tuning feature (QPFRADJ=1,2 or 3), also specify the following parameters using the OS/400 `WRKSHRPOOL` command followed by pressing F11. You then see the screen in Figure 295.

Work with Shared Pools							System: DEVSYS
Main storage size (M) . . :							6096.00
Type changes (if allowed), press Enter.							
Pool	Priority	-----Size %-----		-----Faults/Second-----			
		Minimum	Maximum	Minimum	Thread	Maximum	
*MACHINE	1	7.50	100	10.00	.00	10.00	
*BASE	1	4.99	100	5.00	.50	200	
*INTERACT	2	5.00	100	10.00	2.00	100	
*SPOOL	2	1.00	100	5.00	1.00	100	
*SHRPOOL1	2	1.00	100	10.00	2.00	100	
*SHRPOOL2	2	1.00	100	10.00	2.00	100	
*SHRPOOL3	2	1.00	100	10.00	2.00	100	
*SHRPOOL4	2	1.00	100	10.00	2.00	100	
*SHRPOOL5	2	1.00	100	10.00	2.00	100	
*SHRPOOL6	2	1.00	100	10.00	2.00	100	
							More...

Figure 295. Work with Shared Pools

The columns on this screen are explained here:

- **Size Limits:** Establishes the range of sizes that a particular pool may have. Specifying a suitable lower limit prevents critical memory pools from being drained during periods of low activity. This impacts the system's ability to react rapidly to a sudden surge in activity.
- **Pool Priority:** Identifying a pool as high priority (with a lower numeric value) increases its tendency to be given more memory over a pool with a lower priority. Therefore, sequence your pools in the order of importance with the more important memory pools having relatively lower values. More than one pool can have the same priority if they are equally important.

For example, in an SAP R/3 environment where there is a production system and a development system on one iSeries, consider setting the pool supporting the production system (*BASE pool for example) at a priority of 2. Setting the pool running the development system at priority 3.

- **Faults per second:** Indicates the number of acceptable faults that can be experienced by the pool:
 - **Minimum faults:** Specifying a low value for the minimum faults per second for a pool reduces the likelihood of the pool surrendering memory to another pool. Until the pool hits this lower threshold, it is acceptable for this pool to further reduce the faulting rate.

- **Maximum faults:** Specifying a low value for the maximum faults per second for a pool increases the likelihood of the pool being given memory from another pool to reduce its faulting rate to below the specified value.
- **Thread:** Determines an acceptable faulting rate for each thread. The sum of the minimum faults and the total faults of all the active threads must be lower than the maximum specified. Faults per thread per second is a computed average. It is not possible to measure this.

For example, consider an SAP memory pool with interactive workload. If the average response time is one second, and you require the memory faulting overhead to be under 10%, the memory faulting time should be under 0.1 second. If the disk response time is 10 milliseconds, the allowable memory faults per interactive dialog will be 10 (that is, $0.1/0.01$). If there are 7,200 dialogs per hour (2 per second), the total allowable faults per second for the pool is 20 (faults per request x requests per second = $10 * 2$). If it is estimated that there are 40 active threads in the pool, the value to be entered here is 0.5 (total faults per number of active threads).

16.5.1.3 Disk space and disk arm utilization

The values at the upper right of the WRKSYSSTS display give you an indication of the disk space used. Ensure that this provides adequate capacity for growth in the database. You should look at the system when there is high activity because this can increase the amount of temporary disk space used by the system that is indicated by the Current Unprotect Used amount. If you look at this display soon after an IPL of the machine, the Current Unprotect Used amount will be low.

If disk usage is high, for example over 90%, you should take steps to reduce disk space usage by deleting objects that are not required. You should make sure that the automatic cleanup jobs are run each day. Use the OS/400 `GO CLEANUP` command and select option 1 to display or change the cleanup options as shown in Figure 296.

Change Cleanup Options		DEVSYS
		11/27/00 21:19:15
Type choices below, then press Enter.		
Allow automatic cleanup	Y	Y=Yes, N=No
Time cleanup starts each day	22:00:00	00:00:00- 23:59:59, *SCDPWROFF, *NONE
Number of days to keep:		
User messages	7	1-366, *KEEP
System and workstation messages	4	1-366, *KEEP
Job logs and other system output	7	1-366, *KEEP
System journals and system logs	30	1-366, *KEEP
OfficeVision for AS/400 calendar items	30	1-366, *KEEP

Figure 296. Cleanup options

As a regular exercise, you should also review the contents of user-defined output queues with the WRKOUTQ command and remove unnecessary spooled files.

The OS/400 commands RCLSTG and RCLSPLSTG should be run as part of your regular maintenance schedule for the iSeries server. The Reclaim Storage (RCLSTG) command is used to perform a general cleanup of auxiliary storage. Note that RCLSTG command has to be run with the system in a restricted state (all subsystems ended) and may take considerable time to run. This depends on the number and type of objects in the system, the amount of damaged objects, the amount of auxiliary storage configured, and the percentage of that storage in use. The Reclaim Spool Storage (RCLSPLSTG) command reclaims unused storage for spooled files that have not been used for more than the number of days specified.

You can set up thresholds for each disk auxiliary storage pool (ASP) which sends a message to the QSYSOPR message queue whenever the threshold value is exceeded. The threshold values are set in the System Service Tools. To set the threshold values, enter the STRSST command. The screen shown in Figure 297 is presented.

System Service Tools (SST)

Select one of the following:

1. Start a service tool
2. Work with active service tools
3. Work with disk units
4. Work with diskette data recovery
5. Work with system partitions

Selection

F3=Exit F10=Command entry F12=Cancel

Figure 297. System Service Tools (SST)

Choose option 3 (Work with disk units), next select option 2 (Work with disk configuration), and then choose option 3 (Work with ASP threshold). The next screen that appears (Figure 298) allows you to set the threshold values for each ASP.

Select ASP to Change Threshold							
Type option, press Enter.							
1=Select							
Option	ASP	Threshold	Overflow	----Protected---		---Unprotected---	
				Size	%Used	Size	%Used
	1	90%	No	60129	17.58%	0	0.00%
	2	90%	No	8589	0.03%	0	0.00%
F3=Exit F12=Cancel							

Figure 298. ASP thresholds

You can also attempt to recover any space occupied by deleted records. Even though SAP R/3 database files on the iSeries server are defined to reuse deleted records (REUSEDLT=*YES), depending on the activity of your database, there can be an accumulation of deleted records that may increase your disk usage beyond what is necessary.

Examine the database for deleted records using SAP R/3 transaction ST04 and selecting **Detail analysis menu-> State on disk** or transaction DB02 by clicking the **Deleted row analysis** push button. You can use the OS/400 RGZPFM command to recover the space occupied by deleted records. A method of selecting only files with more than a specific number of deleted records is described in SAP Note 84081: *Reorganization of an R/3 system on AS/400*.

Useful data about the amount of disk space used by each file system and library on the iSeries can be obtained by using the OS/400 RTVDSKINF command. This command should be scheduled to run on a regular basis, for example weekly, and when major changes are made to your system. It should be scheduled during a period of low activity on the server. Depending on the amount of data on the system, it may take several hours to complete. Each time the command is run, it will replace the existing data previously collected.

Once the RTVDSKINF command has collected the information about disk space, a report can be generated via the OS/400 PRTDSKINF RPTTYPE(*LIB) command. A sample of the output is shown in Figure 299.

```

Disk Space Report
5769SS1 V4R5M0 000526
System Information                      DEVSYS 11/28/00 11:00:24
Information collected . . . . . : 11/21/00 19:00:03

Total disk space on system in 1,000,000
bytes . . . . . : 496606
Main storage size in megabytes . . . . . : 6096
Machine type-model . . . . . : 9406-830
System serial number . . . . . : 65-7ABUN

Description                                % of      Size in
                                           Disk      1,000,000 bytes
User libraries                           57.29      284497.60
User directories                         .85        4196.58
Folders and documents                    .00         24.65
QSYS                                     .43       2154.18
Other IBM libraries                      1.10       5442.63
Licensed Internal Code                   .26       1312.28
Temporary space                          7.56      37523.01
Unused space                            28.48     141445.52
Internal objects                         .89        4397.89
Objects not in a library                 .00          .00
TOTAL                                   96.86     480994.34

```

Figure 299. Disk space report

The Work with Disk Status (WRKDSKSTS) command shows the current performance information for every disk unit in auxiliary storage (Figure 300).

```

Work with Disk Status                      DEVSYS
                                           11/28/00 10:58:11
Elapsed time: 00:01:05

Unit  Type  Size  %   I/O  Request  Read  Write  Read  Write  %
      (M)  Used Rqs  Size (K) Rqs  Rqs   (K)  (K)   (K)   Busy
1  6717  6442  68.4  21.1   7.0   16.0   5.1   6.9   7.6    4
2  6607  3670  68.4   9.2   12.9   8.0   1.1  13.1  11.3    5
3  6607  3670  68.4   8.9   12.1   8.1   .7  12.4   8.8    4
4  6607  3670  68.4  11.7   11.6  10.4   1.2  11.3  14.6   10
5  6607  3670  68.4  10.8    6.3   8.4   2.3   5.8   8.0    5
8  6607  3670  68.4   5.7    8.2   3.7   2.0   9.5   5.9    3
9  6607  3670  68.4   7.1   11.6   5.9   1.2  12.4   7.5    4
10 6607  3670  68.6   6.0   15.9   5.3   .6  16.8   8.8    5
11 6607  3670  68.4   6.2   11.4   4.7   1.5  12.8   7.0    7
12 6607  3670  68.4   8.8    8.2   8.1   .6   8.2   8.3    1
13 6607  3670  68.4   6.6    9.7   5.3   1.3   9.8   9.1    7
14 6607  3670  68.5   8.8    7.4   6.5   2.2   7.4   7.5    4
15 6717  6442  68.4  17.9    8.2  14.7   3.2   8.5   7.0    4

```

Figure 300. Work with Disk Status

If your disk arm utilization is approaching the recommended threshold values and there are no memory constraints resulting in excessive memory faulting and the associated disk arm activity, you might consider installing additional or faster disk drives to reduce the disk arm utilization. As higher capacity disk (DASD) devices (8 GB and 18 GB) for the iSeries become available, fewer arms are needed to satisfy the capacity requirements. This can lead to configuring too few disk arms (actuators) to meet the workload placed on them. A lack of disk arms can bottleneck the processor's performance. To avoid such a bottleneck, a minimum

number of disk arms is needed for optimum performance. This information is published in *AS/400 Disk Arm Requirements Based on Processor Model Performance* by Harold Rosenberg (IBM Rochester Lab, August 1999).

If you recently added more disk drives to an ASP, the new drives, at first, will have very little data stored on them. Over time, OS/400 will gradually even this out, but you will achieve better read and write performance if the data is spread evenly over all the available disk units. You can use the OS/400 STRASPBAL command to evenly distribute the data in an ASP across all the disk units by using the *CAPACITY option. Another alternative is to save the data in the ASP, delete it, and then restore the data from tape. Other options can be considered with disk balancing. For example, one option is the distribution of data by usage balancing to move data from busy disks to less used disks. Another option is by Hierarchical Storage Management (HSM) balancing to move data from low performance drives to high performance drives if you have different speed disk drives in an ASP. These two options require the collection of trace data using the TRCASPBAL command.

During OS/400 system operation, System Storage Management is responsible for placing data on available disk units. This is done automatically and does not require any actions from a user. The user has limited control of placing new data on disks by specifying the target library in a user ASP.

When an OS/400 object is created, the system storage management algorithms locate the disk unit with the largest percentage of free space within an ASP. Then the object is created on that disk unit and the data for that object is written in the sectors on the disk. Once written, the object stays in the same disk location until the object is deleted. The data is spread as evenly as possible on each disk unit within ASP. Generally this method delivers the best system performance and makes the best use of disk capacity for all objects within the ASP.

However, there are situations when these algorithms do not provide the best performance, for example, when a new disk unit is added to an existing ASP. At the beginning, this disk is empty and contains no data. Therefore, all newly created objects or additions of new data to existing objects are primarily done on this disk. This causes the new disk to be used at a higher rate than other disks within ASP, until the new disk reaches the same percentage of used space as other disks. Most likely, the new disk will also be used more heavily later because it contains newer data, which is typically accessed more frequently. A solution is to use the new ASP balancing options that enable the authorized user to balance the distribution of data in an ASP.

16.5.1.4 ASP balancing options

OS/400 V4R4 contains three ASP balancing options:

- Capacity Balance
- Usage Balance
- Hierarchical Storage Management Balance

These balancing options redistribute the existing data throughout the disks within the ASP. They do not change the storage management algorithms, which still work the same way.

Capacity Balance

Capacity Balance offers the user a mechanism to redistribute the data *evenly* across all disk units in the selected ASP. The goal of this option is to move data so that each disk unit has approximately equal percentage of used and unused space. However, since some objects (such as journals, journal receivers and temporary objects) are not moved, it is possible that this percentage may not be quite even.

Usage Balance

This option balances the data on the disk units in an ASP according to the *frequency of usage* of this data. As opposed to the Capacity Balance, the percentage of used space will no longer be even after running the Usage Balance. This is OK since the final goal of using the Usage Balance is to have less data on highly utilized disks and more data on less utilized disks. Use Usage Balance when your system has a wide range of different disk utilizations within an ASP. This difference can be due to a mixture of different capacity disks or when some data is accessed more frequently. There are several ways to determine disk utilization:

- Performance Manager/400 (PM/400) report for customers who sign up for the IBM PM/400 service offering.
- OS/400 Work with Disk Status (WRKDSKSTS) command. This command shows performance and status information about the disk units on the system. It displays the number of units currently on the system, the type of each disk unit, the size of disk space, the percentage of disk space used, the I/O requests per second, the average size of the I/O requests, the average number of read and write requests, the average amount of data read and written, and the percentage of time that the disk is being used.
- Component Report of the Performance Tools licensed program. Performance Tools provide a means for a detailed performance analysis of all system components. The performance data used for analysis can be collected by:
 - OS/400 feature Performance Monitor that you run with the Start Performance Monitor (STRPFRMON) command
 - Operations Navigator feature Performance Collection Services

Usage Balance moves cold data to highly utilized disks, and therefore, tries to influence future allocations of data away from those units. Usage Balance uses the results of previous ASP traces to determine the disk unit usage. Therefore, you must perform an ASP trace before you can run the Usage Balance.

Hierarchical Storage Management Balance

With HSM Balance, frequently used data is moved to high-performance disk units, while less frequently used data is moved to compressed (lower performance) disk units. As with Usage Balance, the percentage of used space will no longer be even after running HSM Balance. HSM Balance uses the results of previous ASP traces to determine the disk unit usage. Therefore, you must perform an ASP trace before you can run the HSM Balance. The ASP selected for an HSM Balance *must* contain compressed disk units *and* non-compressed disk units.

Disk unit compression was introduced with V4R3. Compressed disks have a larger capacity (because they contain compressed data), but are somewhat

slower than non-compressed disks. This is due to the overhead of compression and decompression, and the variations in the length of the data that is written to disk. The compression must be explicitly defined for a specific disk unit within an ASP (only a user ASP can have a compressed disk), and requires compression-capable disk controllers, for example, #6533 (SPD) or #2741 (PCI). If the number of compressed disks in an ASP is much lower than the number of non-compressed disks, then the compressed disks can become a performance bottleneck since more high-use data ends up there. For this reason, we advise that you have a reasonable number of high performance units, depending on your intended use of ASP.

16.5.1.5 When to perform ASP balancing

Use balancing options to improve system performance in situations where standard, single-level storage management algorithms do not bring optimum results. Perform ASP balancing when:

- New disk units are added to an existing ASP.
- Several records or tables are accessed more frequently, which causes higher utilization of certain disk arms.
- A single ASP has a mixture of large capacity disk units (such as 17 GB) and smaller capacity disk units (such as 4 GB or 2 GB).
- There is a mixture of compressed and non-compressed disk units within the same user ASP.

16.5.1.6 How to perform ASP balancing

The balancing options can be managed through three different interfaces:

- The OS/400 Control Language (CL) commands Start ASP Balancing (STRASPBAL), End ASP Balancing (ENDASPBAL), and Trace ASP Balancing (TRCASPBAL).
- System Service Tools (SST). You start SST with the Start System Service Tools (STRSST) command.
- Dedicated Service Tools (DST). You access the DST after a manual IPL of the iSeries server.

These interfaces are described in more detail in the following sections.

16.5.1.7 ASP balancing with SST and DST

In V4R4, the SST and DST contain a new option, option 8 (Add units to ASPs and balance data), to perform Capacity Balance automatically after a disk unit is added to the ASP. Figure 301 shows the SST (or DST) Work with Disk Configuration screen with the new option 8.

Work with Disk Configuration

Select one of the following:

1. Display disk configuration
2. Add units to ASPs
3. Work with ASP threshold
4. Include unit in device parity protection
5. Enable remote load source mirroring
6. Disable remote load source mirroring
7. Start compression on non-configured units
- 8. Add units to ASPs and balance data**
9. Start device parity protection

Selection
8

F3=Exit F12=Cancel

Figure 301. SST with Work with Disk Configuration function

This screen appears after you select the *Work with disk units* option on the SST (or DST) main menu and select *Work with Disk Configuration*.

If you do not use the *Add units to ASPs and balance data* option, we recommend that you run capacity balancing as soon as possible after the addition of a new disk unit to ASP.

16.5.1.8 STRASPBAL

The Start ASP Balance (STRASPBAL) command allows the user to run the ASP balancing function for one or more ASPs. In OS/400 releases prior to V4R1, the only straightforward way to redistribute the data on disks is through this process:

1. Save the object.
2. Rename or destroy the object.
3. Create dummy files to fill the space that the object occupied.
4. Restore the object, which is then evenly spread across all disk units in an ASP.

V4R3 contains the Disk Balance (DSKBAL) command that schedules the balancing to occur at the next IPL. The DSKBAL command is also available as a PTF for V4R1 and V4R2. In V4R4, the STRASPBAL command replaces the DSKBAL command.

The STRASPBAL command can run for 1 to 9,999 minutes or with unlimited time. The ASP balancing ends when:

- The balancing is completed.
- The specified time is expired.
- The ENDASPBAL command is run.

If the balance function stops before the balancing is completed, it will continue from where it left off when the balance function restarts. This allows the balancing to be run during off hours over a several day period.

Recommendation

Since ASP Balance generates additional disk activity, run it when the system is not used much.

Figure 302 shows an example of the STRASPBAL command used to start Capacity Balance.

```
Start ASP Balance (STRASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID . . . > *ALL          1-16, *ALL
      + for more values
Balance type . . . . . > *CAPACITY      *CAPACITY, *USAGE, *HSM
Time limit . . . . . > *NOMAX          1-9999 minutes, *NOMAX


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

Figure 302. Start ASP Balance command with Capacity Balance

Consider using this option after a new disk unit is added to the existing ASP and the SST (or DST) option 8 was not performed.

Note

DST option 8 is more effective in balancing the units than the STRASPBAL command with the *CAPACITY option, since it runs in the dedicated environment. Other users do not use objects that are being redistributed.

Figure 303 shows an example of the STRASPBAL command used to start Usage Balance.

```

                                Start ASP Balance (STRASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID . . . > *ALL          1-16, *ALL
      + for more values
Balance type . . . . . > *USAGE          *CAPACITY, *USAGE, *HSM
Time limit . . . . . > *NOMAX          1-9999 minutes, *NOMAX


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

```

Figure 303. Start ASP Balance command with Usage Balance

Usage Balance uses the results of previous ASP traces to determine the disk unit usage. You must perform an ASP trace before you run a Usage Balance. ASP trace is described in 16.5.1.10, “TRCASPBAL” on page 425.

Figure 304 shows an example of the STRASPBAL command used to start HSM balancing.

```

                                Start ASP Balance (STRASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID . . . > *ALL          1-16, *ALL
      + for more values
Balance type . . . . . > *HSM          *CAPACITY, *USAGE, *HSM
Time limit . . . . . > *NOMAX          1-9999 minutes, *NOMAX


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

```

Figure 304. Start ASP Balance command with HSM Balance

16.5.1.9 ENDASPBAL

If you want to end balancing before the time limit requested is reached, use the End ASP Balance (`ENDASPBAL`) command. You can end ASP balancing any time and restart it later.

Figure 305 shows the `ENDASPBAL` command screen.

End ASP Balance (`ENDASPBAL`)

Type choices, press Enter.

Auxiliary storage pool ID . . . > *ALL 1-16, *ALL
+ for more values

Bottom

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

Figure 305. End ASP Balance (`ENDASPBAL`)

16.5.1.10 TRCASPBAL

Before you use Usage Balance or HSM Balance, you must run the Trace ASP Balance (`TRCASPBAL`) command. This command monitors the frequency that data is accessed on the disk units in the ASP. Each I/O to the disk units is monitored, and the results are recorded for use by the `STRASPBAL` command. The statistics that are collected are cumulative. Therefore, you can collect trace information from several `TRCASPBAL` command runs. This allows you to accumulate statistics for peak hours from several days.

For example, you may want to run one trace on Monday for 35 minutes. On Tuesday, you run another trace on the same ASP for 15 minutes. The second group of statistics is added to the first collection, and the cumulative result is then used by the `STRASPBAL` command to balance the ASP. Once a trace is used by the `STRASPBAL` command, it cannot accept additional trace accumulations. It must be cleared before a new trace cumulation can be gathered.

You start the trace by running the `TRCASPBAL` command with the Trace option setting parameter set to `*ON` as shown in Figure 306.

Trace ASP Balance (TRCASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID . . .	> *ALL	1-16, *ALL
+ for more values		
Trace option setting	> *ON	*ON, *OFF, *CLEAR
Time limit	> *NOMAX	1-9999 minutes, *NOMAX

Bottom

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

Figure 306. Trace ASP Balance command set to *ON

The TRCASPBAL command can run for 1 to 9,999 minutes or without limit. The trace can be stopped before the time limit is expired by running the TRCASPBAL command with the Trace option setting parameter set to *OFF. Trace data is collected cumulatively until it is cleared. Trace data can be cleared in two ways:

- After the HSM balancing has completed, the system automatically clears the trace information.
- By running the TRCASPBAL command with the Trace option setting parameter set to *CLEAR as shown in Figure 307.

Trace ASP Balance (TRCASPBAL)

Type choices, press Enter.

Auxiliary storage pool ID . . . > *ALL

1-16, *ALL

+ for more values

Trace option setting > *CLEAR

*ON, *OFF, *CLEAR

Bottom

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

F24=More keys

Figure 307. Trace ASP Balance command set to *CLEAR

Clear the old trace data if you do not want this data to be used in determining the locations of the high-use and low-use data.

How long should you run the trace? Here are the guidelines you should follow:

- In case of Usage Balance, run the trace during the periods of high disk utilization that you determined by using the WRKDSKSTS command or PM/400 report or Performance Tools. Usage Balance will move infrequently used (cold) data to highly utilized disk units and try to influence future allocations away from those units. The target units percent full will be higher than other units.
- For HSM Balance, the traces should be longer. HSM Balance moves cold data from the non-compressed disk units to the compressed units. It also moves active data off the compressed disk units to non-compressed units.

16.5.1.11 Disk compression

For systems running OS/400 V4R3 or higher with appropriate DASD controllers, instead of adding more disk drives, the total disk capacity available can be increased by utilizing the iSeries's integrated hardware disk compression. If you are using V4R3 on your system, you can start or stop disk compression only on nonconfigured disk units. If you are using V4R4 or later, you can start or stop disk compression on configured and nonconfigured disk units. This is only the case for User ASPs, however, because compression cannot be used on the system ASP. Data is dynamically compressed or uncompressed by the DASD controller as data is written to or read from the disk. It has no effect on the main CPU utilization since compression is performed by the DASD controller IOP. Compression rates of approximately two times, which will double the disk capacity of a drive, have been achieved for SAP R/3 data. The read times from disks with compression are approximately the same as uncompressed drives, although write times may be a little slower.

The performance of compressed disk drives is slower than the performance of non-compressed disk drives. This is due to the overhead of compression and decompression, and the variations in the length of the data that is written to disk. Disk compression is started and ended using the Dedicated Service Tools (DST) menu. For more detailed information refer, to Section 8.5 of *OS/400 Backup and Recovery V4R5*, SC41-5304.

16.5.1.12 Disk reorganization with STRDSKRGZ and ENDDSKRGZ

These two commands are new in OS/400 V4R2. The STRDSKRGZ command is used to start disk reorganization. The ENDDSKRGZ command ends disk reorganization. Disk reorganization is a system function used to collect together all unused space on each disk unit in the selected ASP (or ASPs). It allows future allocations of large space on the disk unit to be done more efficiently.

The screen in Figure 308 shows the parameters for the STRDSKRGZ command.

Start Disk Reorganization (STRDSKRGZ)

Type choices, press Enter.

Time limit	> *NOMAX	1-9999 minutes, *NOMAX
Auxiliary storage pool ID . . .	*ALL	1-16, *ALL
+ for more values		

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

F24=More keys

Bottom

Figure 308. Start Disk Reorganization (STRDSKRGZ) parameters

The user specifies a time limit that the function must run for each ASP being reorganized. When the time limit is reached, the function ends. The time limit specified is for each ASP being reorganized. For example, if ASP(*ALL) is specified and the machine has four ASPs configured and TIMLMT(60) is specified, four reorganization functions are started, and each can run 60 minutes.

If the reorganization of any ASP has not completed after 60 minutes, it will be forced to end. This allows you to do disk reorganization incrementally.

The screen in Figure 309 shows the parameters for the ENDDSKRGZ command.

End Disk Reorganization (ENDDSKRGZ)

Type choices, press Enter.

Auxiliary storage pool ID . . .

*ALL

1-16, *ALL

+ for more values

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Bottom

Figure 309. End Disk Reorganization (ENDDSKRGZ) parameters

The user can select to end disk reorganization for all ASPs or for one or more specific ASPs. A message is sent to the system history log when the reorganization function is ended for each ASP.

Recommendation

Run the STRDSKRGZ command if you receive the CPI0999 system message (Storage directory threshold reached).

16.5.2 Expert cache

When you turn on expert cache, the system automatically adjusts storage pool paging characteristics and determines the best approach for handling data in the pool.

- If objects are referred to sequentially, the system brings larger blocks of data into the memory and delays writing changes of data to the disk. This reduces the number of I/O operations and reduces the disk arm contention.
- If objects are referred to randomly, the system does not bring in large blocks of data because that does not reduce the number of I/O operations.

Always use expert cache for the SAP R/3 pool, which by default is *BASE. Turn off expert cache only in case of main storage constraints.

Use the OS/400 WRKSYSSTS or WRKSHRPOOL command to turn on expert cache (*CALC) for the *BASE pool as shown in Figure 310.

```

Work with Shared Pools
System:  DEVSYS

Main storage size (M) . :      6096.00

Type changes (if allowed), press Enter.

Pool          Defined   Max   Allocated   Pool   -Paging Option--
              Size (M) Active Size (M)   ID     Defined  Current
*MACHINE      489.87   +++++   489.87     1     *FIXED  *FIXED
*BASE         5318.12     81   5318.12     2     *CALC   *CALC
*INTERACT     256.00     16   256.00     4     *FIXED  *FIXED
*SPOOL        32.00     12    32.00     3     *FIXED  *FIXED
*SHRPOOL1      .00      0
*SHRPOOL2      .00      0
*SHRPOOL3      .00      0
*SHRPOOL4      .00      0
*SHRPOOL5      .00      0
*SHRPOOL6      .00      0
More...

Command
====>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F11=Display tuning data
F12=Cancel

```

Figure 310. Expert cache

16.5.3 Maximum active threads per memory pool

Use the OS/400 WRKSYSSTS or WRKSHRPOOL command to see the values for the maximum number of threads that can use the CPU concurrently for each memory pool. The values for the base (*BASE), interactive (*INTERACT), and spool (*SPOOL) memory pools can be adjusted. If these are set too low, threads may transition to the ineligible state, and if set too high, excessive page faulting in that pool may occur. These values will be automatically adjusted when the performance adjustment system value is turned on. However, when running SAP on the iSeries, the automatic performance adjustment should be switched off after initial tuning is performed. These values may need to be changed when for example, more printers are added. As a starting point, allow one thread in *SPOOL per active printer, and one thread for each interactive (not SAP R/3) user who will sign on concurrently to the iSeries via a 5250 session in *INTERACT.

The value of Max Act for the *BASE pool can be initially set using the following calculation:

- **Central server** (database and application): Number of R/3 work processes x 1.2
- **Database server only:** (Number of R/3 work processes on DB server + number of work processes on each remote application server) x 1.2
- **Application server only:** Number of R/3 work processes on application server x 1.2

This number may need to be adjusted in line with the recommendations for the transitions and page faulting described above. For example, for a central server with a total of 68 work processes, the calculation for the *BASE pool would be:

$$68 \times 1.2 = 81.6$$

Therefore, the value of Max Act can be set to 81 as shown in Figure 311.

Work with System Status									
								DEVSYs	
								11/28/00	10:56:19
% CPU used	:	90.3		System ASP	:	488.2	G		
% DB capability	:	24.7		% system ASP used	:	68.4641			
Elapsed time	:	00:14:12		Total aux stg	:	496.6	G		
Jobs in system	:	4384		Current unprotect used	:	38287	M		
% perm addresses	:	.027		Maximum unprotect	:	40407	M		
% temp addresses	:	.500							
Sys	Pool	Reserved	Max	---DB----	--Non-DB---	Act-	Wait-	Act-	
Pool	Size M	Size M	Act	Fault	Pages	Fault	Pages	Wait	Inel
1	489.87	325.80	+++++	.6	1.7	16.7	19.4	21.4	.0
2	5318.12	.08	81	45.2	4193	496.2	631.7	1546	.0
3	32.00	.00	12	.0	.7	.5	1.2	67.3	.0
4	256.00	.00	16	.0	.2	.1	.1	2.3	.0

Figure 311. Maximum active threads

16.5.4 Run priority

All SAP R/3 server jobs in an iSeries subsystem run at the same priority, which is determined by the corresponding class. Therefore, an SAP R/3 dialog (interactive) work processes competes for the CPU at the same priority as the SAP R/3 batch work processes within an SAP R/3 instance.

16.5.4.1 SAP R/3 system priority

The SAP R/3 execution classes have the names R3_<nn>, where <nn> is the instance number. They are in the library R3<SID>400, where <SID> is the SAP R/3 system identifier. The OS/400 DSPCLS command displays the values that apply to the class as shown in Figure 312.

Display Class Information		System: DEVSYs
Class	R3_01	
Library	R3DEV400	
Run priority	20	
Time slice in milliseconds	2000	
Eligible for purge	*YES	
Default wait time in seconds	120	
Maximum CPU time in milliseconds	*NOMAX	
Maximum temporary storage in megabytes	*NOMAX	
Maximum threads	*NOMAX	
Text		

Figure 312. OS/400 class for R/3 instance 01

By changing the Run Priority in a particular class, it is possible to change the run priority of all the work processes within an SAP R/3 instance.

For example, you may have installed a separate R/3 system using the homogeneous copy method to be used for training users in R/3. In this case, it is possible to run all the work processes of the training SAP R/3 system at a higher priority (for example, priority=20) than the Development system (for example,

priority=40), if both systems are installed on the same iSeries system. Each instance has its own subsystem description and execution class.

16.5.4.2 Work process priority

It is also possible to alter the run priorities of individual work processes or types of work processes. For example, you can lower the priority of *batch jobs* in an SAP R/3 instance by using the OS/400 WRKPID command (Figure 313). This requires the SAP R/3 work process identifier as input (refer to 16.4.8.4, “OS/400 job number and SAP R/3 PIDs” on page 402).

Work with Job by PID (WRKPID)

Type choices, press Enter.

Process ID > 2429

Character value

Prompt command *YES

*YES, *NO, Y, N

xxxJOB command CHG

Character value

F3=Exit

F4=Prompt

F5=Refresh

F12=Cancel

F13=How to use this display

F24=More keys

Bottom

Figure 313. Work with Job by PID (WRKPID)

Make sure you do not downgrade the priorities of other SAP R/3 work processes because this can severely impact throughput of SAP R/3 transactions in the subsystem.

From SAP R/3 Release 3.1I, there are four levels of prioritization allowed for each type of work processes in each SAP R/3 instance on the iSeries server:

- HH: Class priority minus 8
- H: Class priority minus 4
- M: Class priority
- L: Class priority plus 4

By default, the class priority is 20, and all jobs in the R/3 subsystem will run at this priority with the exception of the following four jobs that run at priority 12.

The following work processes are assigned to run at the priority **HH** level and cannot be changed using instance parameters:

- Dispatcher (DISP)
- Message Server (MSG_SERVER)
- Gateway (GWRD)
- Enqueue (WPnn)

The following work processors are assigned to run at the priority **M** level and cannot be changed using instance parameters:

- Dialog
- Update

Three R/3 instance parameters allow user prioritization of update, batch and spool work processes. They are set to M by default:

1. rdisp/prio/upd = M (or H)
2. rdisp/prio/btc = M (or L)
3. rdisp/prio/spo = M (or L)

To change the priority of the update, batch or spool R/3 work processes from the default of 20, enter the appropriate parameter in the instance profile using transaction RZ10. You then need to restart the SAP instance for the change to take effect.

You must also have set the OS/400 dynamic priority scheduler system value, QDYNPTYSCD, to 1. If you need to make this change, you must IPL the iSeries server for the change to take effect.

16.5.5 Separate memory pool for an SAP R/3 instance

As an installation default, all SAP R/3 instance jobs are routed into System Pool 2 (*BASE pool) and possibly may compete for system resources (activity levels, memory, and CPU) with other applications running in this pool.

If you are running a non-SAP application on the same iSeries as an SAP R/3 system, separate memory pools for each application may help minimize contention for memory. This section shows how to route SAP R/3 instance jobs into a different pool.

Stop the SAP R/3 instance and change the pool definition for the instance subsystem (Figure 314). Assign either a private pool or a shared pool (this example). If you assign a separate private memory pool, the performance adjuster (QPFRADJ) does not manage the pool even if the function is turned on.

Change Subsystem Description (CHGSBSD)

Type choices, press Enter.

Subsystem description	> R3_01	Name
Library	> R3DEV400	Name, *LIBL, *CURLIB
Storage pools:		
Pool identifier	> 1	1-10, *SAME
Storage size	> *SHRPOOL1	Number, *BASE, *NOSTG...
Activity level		Number
	+ for more values	
Maximum jobs	*SAME	0-1000, *SAME, *NOMAX
Text 'description'	*SAME	

Bottom

F3=Exit
F4=Prompt
F5=Refresh
F10=Additional parameters
F12=Cancel

F13=How to use this display
F24=More keys

Figure 314. CHGSBSD: Change pool definition for R/3 subsystem

Use the OS/400 WRKSHRPOOL command to set the size of the shared memory pool.

16.5.6 Network communication

The network design, setup, and tuning is key to achieving good performance. The network frequently can be the slowest link in the system. This needs to be considered in the original design and implementation to achieve the performance goals. Typically, this is achieved by reducing the number of turns across the WAN network to a minimum at the expense of potentially increasing network traffic on the LAN.

Consider implementing some of the following changes to reduce the network communications traffic and improve performance:

- Change the TCP/IP configuration using the CFGTCP command option 12 to have the system search the local host table before the remote name server. Set the value Host Name Search Priority to *LOCAL.
- Change the SAP global directory to bypass the remote file system (QFileSvr.400) when the global directory is local as described in SAP OSS Note 68487.

For TCP/IP communications, the key performance parameters that can be changed on the iSeries are the maximum transmission unit size and the size of the send and receive buffers.

Note the following parameters:

- **Maximum Transmission Unit (MTU) Size:** Specifies the maximum size (in bytes) of IP datagrams that can be transmitted through this route.
Use the CHGTCPRIE command (CFGTCP option 2) to set the MTU value of *DFTRROUTE to *IFC, which means that the maximum transmission unit is the MTU of the interface that is associated with this route. The MTU value for the TCP/IP interface should likewise be set to *LIND.
- **Buffer Size:** Use the CHGTCPA command (CFGTCP option 3) to set the TCP/IP send and receive buffer sizes to 64K (65536 bytes) for OS/400 V4R3 and earlier releases. From V4R4 onwards, this value should be increased to a higher figure, typically 1MB. Values up to a maximum of 8 MB (8388608 bytes) are possible. This is shown in Figure 315.

Change TCP/IP Attributes (CHGTCPA)

Type choices, press Enter.

TCP keep alive	120	1-40320, *SAME, *DFT
TCP urgent pointer	*BSD	*SAME, *BSD, *RFC
TCP receive buffer size	1048576	512-8388608, *SAME, *DFT
TCP send buffer size	1048576	512-8388608, *SAME, *DFT
UDP checksum	*YES	*SAME, *YES, *NO
Path MTU discovery:		
Enablement	*YES	*SAME, *DFT, *NO, *YES
Interval	10	5-40320, *ONCE
IP datagram forwarding	*NO	*SAME, *YES, *NO
IP source routing	*YES	*SAME, *YES, *NO
IP reassembly time-out	10	5-120, *SAME, *DFT
IP time to live	64	1-255, *SAME, *DFT
ARP cache timeout	15	1-1440, *SAME, *DFT
Log protocol errors	*NO	*SAME, *YES, *NO

Bottom

Figure 315. Change TCP/IP Attributes

You should examine the performance using NETSTAT option 3, F10 (Display connection totals) for retransmitted segments before and after changing this setting. Setting the TCP/IP buffer size too high can cause storage allocation problems on your server since it will allocate the same size receiving buffer for each connection. This may happen if you have some remote connections on low speed ISDN lines and others with high speed local connections. The exact buffer size that provides the best throughput will be dependent on several factors, including the types of network switches, ACK timing, error rate and the network topology. This is discussed in *Performance Capabilities Reference Version 4, Release 5*, which is available on the Web at:

<http://publib.boulder.ibm.com/pubs/pdfs/as400/V4R5PDF/AS4PPCP3.PDF>

- Line Descriptions:** To take advantage of the Ethernet performance improvements since V4R4, the Ethernet line description parameter TCPONLY should be set to *YES if no other protocols, such as APPC, will be active on the line. If the Ethernet adapter has a dedicated IOP (feature code 2842 or 2843 on the 2xx and 8xx iSeries servers) the IOP and device driver will run the TCP/IP optimized version of its microcode. Ensure also that the duplex mode for the Ethernet line description is set correctly. For example, if a switch is set to full duplex, the TCP/IP interface parameter for the line should also be set to full duplex.

For token-ring networks, the Maximum frame size parameter for the token-ring line description can be increased from the default value of 1994 to its maximum of 16393 to allow for larger transmissions.

- AnyNet:** If AnyNet support (such as for SNA communications) is not required, the network attribute should be turned off. Use the CHGNETA ALWANYNET(*NO) command to do this.

16.5.7 Installing additional SAP R/3 instances

You may also consider having multiple SAP R/3 instances on the same iSeries server with the objective of isolating the interactive (dialog) work processes from

the batch work, for example. This gives you the additional flexibility of setting up the class of the subsystem where the R/3 batch jobs run so they have a lower priority than the subsystem running the R/3 dialog jobs.

You need to be aware that running additional instances on a system can result in an increased demand for memory for instance buffers when all instances are started.

16.5.8 Temporary storage

SAP R/3 uses large amounts of temporary storage for R/3 user contexts and DB2 UDB for iSeries database operations such as ODPs, temporary indexes, joins, and sorting. You can monitor this using the DSPTMPSTG command and the WRKSYSSTS command in the Current unprotect used field.

If you have more than one R/3 instance active, the DSPTMPSTG command shows the amount of temporary storage used by the OS/400 subsystem for a particular instance by specifying the subsystem name, R3_<nn>, where <nn> is the instance number.

The DETAIL parameter allows you specify the *FULL value to get a detailed listing from the DSPJOBLOG F10 Detail screen as shown in Figure 316.

```
DSPTMPSTG SBS(R3_01) DETAIL(*FULL)
Ownership of object GETSBSSTMPs in QTEMP type *USRSPC changed.
Subsystem:      R3_01  Job:   R3_R01_01  Temporary storage:    14336 KB.
Subsystem:      R3_01  Job:   MSG_SERVER  Temporary storage:    23552 KB.
Subsystem:      R3_01  Job:      R3_01    Temporary storage:     2048 KB.
Subsystem:      R3_01  Job:      DISP     Temporary storage:   1110016 KB.
Subsystem:      R3_01  Job:      GWRD     Temporary storage:    48128 KB.
Subsystem:      R3_01  Job:      WP01     Temporary storage:   532480 KB.
Subsystem:      R3_01  Job:      WP02     Temporary storage:   154624 KB.
Subsystem:      R3_01  Job:      WP03     Temporary storage:    47104 KB.
Subsystem:      R3_01  Job:      WP04     Temporary storage:    37888 KB.
Subsystem:      R3_01  Job:      WP05     Temporary storage:    40960 KB.
Subsystem:      R3_01  Job:      WP06     Temporary storage:    43008 KB.
Subsystem:      R3_01  Job:      WP07     Temporary storage:   103424 KB.
Subsystem:      R3_01  Job:      WP08     Temporary storage:    39936 KB.
Subsystem:      R3_01  Job:      WP09     Temporary storage:   674816 KB.
Subsystem:      R3_01  Job:      WP10     Temporary storage:    35840 KB.
Subsystem:      R3_01  Job:      WP11     Temporary storage:   110592 KB.
Subsystem:      R3_01  Job:      WP12     Temporary storage:    44032 KB.
Subsystem:      R3_01  Job:      WP13     Temporary storage:    36864 KB.
Subsystem:      R3_01  Job:      WP14     Temporary storage:   117760 KB.
Subsystem:      R3_01  Job:      WP15     Temporary storage:    36864 KB.
Subsystem:      R3_01  Job:      SAPOSCOL  Temporary storage:    13312 KB.
Subsystem:      R3_01  Job:      WP00     Temporary storage:   296960 KB.
-----
Subsystem:      R3_01  Job:      *ALL     Temporary storage:  3564544 KB.
```

Figure 316. Display temporary storage job list

The longer an R/3 instance is active, the more temporary storage will be used. The temporary storage will be released when R/3 is ended or when the server is IPLed.

The SAPPFPAR program in the kernel library can be used to obtain an estimate of the amount of temporary storage that will initially be used by an SAP R/3

system. This is a function of the number of R/3 users, the R/3 transactions used, and some R/3 instance memory parameters, most importantly `em/initial_size_MB` and `abap/buffersize`.

For example, the following command results in the display shown in Figure 317:

```
CALL PGM(SAPPPFAR)
PARM('pf=/usr/sap/<SID>/SYS/profile/<SID>_<instancename>_<hostname>')
```

```
Parameter-Name....:
> ztta/roll_extension
Result.....: 2000000000
Parameter-Name....:

===>

F3=Exit F4=End of File F6=Print F9=Retrieve F17=Top
F18=Bottom F19=Left F20=Right F21=User Window
```

Figure 317. SAPPPFAR display

It can also be estimated by adding together the following components:

- The number of sessions at peak time (SM04) x 1 MB
- The Max. use value of the extended memory from ST02
- The sum of the buffer values from ST02
- The number of work processes (SM50) x 120 MB

16.5.9 Logon load balancing

Logon load balancing is a feature of SAP R/3 that allows for the distribution of users across multiple instances of application servers at the time they log on to SAP R/3 system. Load balancing is achieved through the definition of logon workgroups in transaction SMLG. When the users log on to the system, they will automatically be logged on to the server that currently has the best performance statistics or the lowest number of logged on users.

For more information on this subject, refer to the *Logon Load Balancing* section in the SAP R/3 online documentation.

Ensure that the iSeries system value QDYNPTYSCD (Dynamic Priority Scheduler) is set to 1 using the command:

```
DSPPSYVAL SYSVAL(QDYNPTYSCD)
```

16.6 SAP R/3 application performance

This section is primarily for iSeries performance specialists who are monitoring the performance of an SAP R/3 system running on an iSeries server. It outlines tools and facilities provided by SAP to monitor the activity of an SAP R/3 system. This is only an introductory discussion and is not intended to provide you with all the information you need to analyze and tune an SAP R/3 system.

For more information, you should attend the SAP course *Workload Analysis*. You can also find additional information in the SAP R/3 Online Documentation CD under the path **Computing Center Management System-> CCMS Monitoring-> Workload Monitor**, which includes guidelines for response times.

16.6.1 SAP R/3 monitoring

The following SAP R/3 transactions are some of the transactions that were used to analyze performance at the application level:

- **SM66** Global (System-Wide) Work Process Overview (see 16.6.1.1, “SM66 Global system-wide work processes” on page 438)
- **SM04** Logged on User List (see 16.6.1.2, “SM04 User list” on page 439)
- **ST07** Application Monitor: User Distribution (see 16.6.1.3, “ST07 Application monitor” on page 440)
- **ST03** Workload Analysis
- **ST03N** Workload Monitor
- **ST02** Tune Summary - buffers
- **ST04** DB2/400 Performance Monitor
- **DB02** Database Performance: State on Disk
- **ST05** SQL Trace - ABAP level

The SAP R/3 Operating System Monitor transaction ST06, which deals with monitoring the total resources of the iSeries, and the Work Process Overview transaction SM50 are discussed in 16.4.7, “SAP R/3 system on: Performance database” on page 398.

Note: Native OS/400 measurements do not recognize SAP R/3 users, dialogs steps, or response times. You must use tools provided by SAP through the *Computing Center Management System* (CCMS) to determine the number of dialog steps generated by SAP R/3 users and the response time on the iSeries server.

The features of CCMS in SAP R/3 provide information to assist you in determining causes that can impact performance. These considerations apply even if the iSeries server is not constrained by computing resources such as CPU, memory, and disk. You also need to keep in mind that network delays can also contribute significantly to response times perceived by the end users.

16.6.1.1 SM66 Global system-wide work processes

This transaction allows you to view all of the work processes in all of the instances within an SAP R/3 system (Figure 318).

Systemwide Work Process Overview

CPU

Debugging

Long <-> short names

Select process

Settings

Sort: Server

Server	No	Typ	PID	Status	Reason	Se	Start	Err	CPU	Time	CLI	User	Report
dfsapbak_PRD_02	8	DIA	24	running			Yes				100	HEGGARTP	RMMMBEST
dfsapbak_PRD_02	2	DIA	6411	running			Yes		4866	100	CONNN	RM07MMAT	
dfsapbak_PRD_02	34	BTC	7778	running			Yes		2967	100	BUSHL	SAPLKKRV	

Figure 318. SM66 Global work process overview

You can filter the work processes you want to view by using the **Select Process** push button. This allows you to focus on those work processes you think need to investigate (Figure 319).

Systemwide Work Process Overview

CPU Debugging Long <-> short names Select process Settings														
Select Process														
Work process selection														
Type ☒ Dialog ☒ Update ☒ Background ☒ Enqueue ☒ Spool					Status ☐ wait ☒ running ☒ stops ☒ finished Restriction ☐ only waiting on semaphore									
Run-time selection														
CPU time >= sec					Run-time >= sec									
Application selection														
Dialog Transaction Function code Screen Report (CUA)														
Reporting Report Job														
User selection Client User name					Action selection Action on table									
☒ ☐														

CPU	Time	CLI	User	Report
			100 HEGGARTP	RMMMBEST
	4866	100	CONNN	RM07MMAT
	2967	100	BUSHL	SAPLKKRV

Figure 319. SM66 Process selections

16.6.1.2 SM04 User list

This transaction allows you to view all of the users logged onto a SAP R/3 instance (Figure 320).

User list

Sessions   User info						
C1.	User	Terminal	Tcode	Time	Sessions	Int.ssess. Trace
000		ent0569		12.44.42	1	2 OFF
330	ARANHAV	AC7813BU120514	SM30	12.18.18	1	2 OFF
330	AVVARIP	ENT01642.dairyfarme		12.39.50	4	5 OFF
330	DALEC	ENT02183.dairyfarme	SE38	12.02.18	2	4 OFF
330	DELROSARIOI	ent01640.dairyfarme	SE10	12.20.59	1	2 OFF
330	DEWETW	ent02184.dairyfarme		12.37.50	4	6 OFF
330	DOTHANHM	ent01952.dairyfarme	SE12	12.14.43	4	11 OFF
330	GANL	ent02065	SE38	11.12.32	1	2 OFF
330	KARLEKARV	ent03148.dairyfarme		12.14.38	1	2 OFF
330	LINESA	1id01643.dairyfarme	SE38	12.03.14	3	9 OFF
330	PARISID	AC003327		12.21.00	2	3 OFF
330	PETERSL	3J04DM74H0BF	SPR0	09.28.00	1	2 OFF
330	POWERM	1idmap	SM04	12.44.57	1	2 OFF
330	SINGHAR	AC3J08DTQ6S96K	SE10	12.43.33	1	3 OFF
330	SUKKARC	1id01951.dairyfarme		08.47.58	1	2 OFF
330	ZHAQZ	ent05003.dairyfarme	SE10	12.18.29	2	4 OFF
*** 16 users logged on with 30 modes ***						

Figure 320. SM04 User list

Use this function prior to shutting down the R/3 system so users can be warned in advance. A system message can be sent to all users currently logged on via transaction SM02.

If you select **Goto-> Memory**, you can see if any of the users allocated *private memory*, which impacts performance by preventing a work process from being used by other users.

16.6.1.3 ST07 Application monitor

This transaction enables you to view the overall activity of your SAP R/3 system (Figure 321). The initial display shows:

- Total number of users
- Number of servers
- Number of clients
- Summary of users by application module

Application Monitor: User Distribution					
Choose Sort SAP buffer DB accesses DB memory Response time Quantity structure History					
Database	Name	PRD	R/3 Release	45B	
	Server	dfsapprd	Time	12:10:31	
	System	DB400	Date	03.01.2001	
User	1,095		all clients		
Number of servers	1		Work processes	45	
Application	Number of users			Sess. per	Appl.
	loggedOn	active	in WP	User	Server
Basis Components	17	5	1	2.00	1
Controlling	1	0	0	5.00	1
Financial Accounting	22	4	0	1.86	1
Logistics - General	11	0	0	1.91	1
Logistics Execution	12	5	2	1.83	1
Materials Management	23	5	0	1.87	1
Production Planning and Control	13	4	0	2.15	1
Quality Management	6	3	1	1.50	1
Sales and Distribution	13	4	1	2.85	1
Treasury	3	0	0	2.00	1
Other	43	2	3	1.26	1
total	164	32	8	1.83	1

Figure 321. ST07 Application monitor: Summary by application

16.6.1.4 ST07 Users of an application

You can select an SAP R/3 application module and use the **Choose** push button from the initial ST07 display, shown in Figure 321, to see a detailed analysis of the users within the selected application module (Figure 322).

Application Monitor: User Distribution

Choose	Sort	Basic list	SAP buffer	DB accesses	DB memory	Response time	Quantity structure

Database	Name	PRD	R/3 Release	45B
	Server	dfsapprd	Time	12:10:31
	System	DB400	Date	03.01.2001

User	1,095	all clients	
Number of servers	1	Work processes	45

Navigation
1 Materials Management

Application	Number of users			Sess. per	Appl.
	loggedOn	active	in WP	User	Server
Inventory Management	13	4	0	2.46	1
Invoice Verification	4	1	0	1.25	1
Purchasing	6	0	0	1.00	1
total	23	5	0	1.87	0

Figure 322. ST07 Application monitor: Module details

16.6.1.5 ST07 R/3 buffers

You can select an SAP R/3 application module and use the **R/3 Buffers** push button from the initial ST07 display (in Figure 321) to see a detailed analysis of the SAP R/3 internal buffer usage by the application module (Figure 323).

Application monitor: SAP buffer

Choose Sort Absolute <-> %					
Database	Name	PRD	R/3 Release	45B	
	Server	dfsapprd	Time	12:13:56	
	System	DB400	Date	03.01.2001	
Appli. Server APPN.DFSAPBAK					
Application	Program	Generic	Single	CUA	Total
	(kByte)	(kByte)	(kByte)	(kByte)	(kByte)
Accounting - General	292	4	0	1	297
Basis Components	102,860	25,031	13,764	2,133	143,787
Controlling	71,700	1,518	92	1,418	74,728
Cross-Application Components	46,608	4,690	356	865	52,519
Customer Service	1,304	0	0	0	1,304
Enterprise Controlling	19,224	1,136	0	232	20,592
Environment Management	180	0	0	0	180
Financial Accounting	50,556	1,076	1,208	989	53,829
Investment Management	764	0	0	0	764
Logistics - General	58,520	739	28	1,031	60,318
Logistics Execution	47,064	785	12	943	48,804
Materials Management	44,156	963	84	1,418	46,621
Payroll Accounting	112	0	0	0	112
Plant Maintenance	2,252	1	4	6	2,262
Production Planning and Control	95,140	1,757	380	1,478	98,754
Project System	7,492	1	0	9	7,503
Quality Management	19,880	123	4	687	20,694
Real Estate Management	1,088	0	0	0	1,088
Sales and Distribution	53,472	48,695	380	896	103,443
Treasury	1,508	4	12	45	1,569
Other	127,412	86	4	1,670	129,172
Sum	751,584	86,609	16,328	13,821	868,340

Figure 323. ST07 Application monitor: R/3 buffers

A more detailed analysis is possible by using the **Choose** push button after selecting an application module (Figure 324).

Application monitor: SAP buffer

Choose Sort Absolute <-> %					
Database	Name	PRD	R/3 Release	45B	
	Server	dfsapprd	Time	12:17:02	
	System	DB400	Date	03.01.2001	
Appli. Server APPN.DFSAPBAK					
Navigation					
1 Materials Management					
Application	Program	Generic	Single	CUA	Total
	(kByte)	(kByte)	(kByte)	(kByte)	(kByte)
Consumption-Based Planning	336	0	0	0	336
External Services Management	2,936	5	0	43	2,985
Information System	344	3	0	13	359
Inventory Management	13,280	725	4	626	14,635
Invoice Verification	7,548	28	0	85	7,660
Purchasing	19,712	202	80	652	20,646
Sum	44,156	963	84	1,419	46,621

Figure 324. ST07 Application monitor: R/3 buffers (by function)

16.6.1.6 ST07 Database access

The Database accesses push button on the display in Figure 321 shows you an analysis of the total database calls and database activity in the application module.

A more detailed analysis is possible for a specific application module by using the Choose push button.

16.6.1.7 ST07 Response time

The Response time push button on the display in Figure 321 shows you an analysis of the SAP R/3 response times averaged by the application module. In addition, the display shows the number of dialog steps and the resource usage (CPU and disk) per dialog.

Application monitor: Response times

Choose Sort Graphic Total<->avg Task type						
DB name PRD SAP Release 45B Server dfsapbak Analysis dat 03.01.2001 We DB syst DB400 Date/Time 03.01.2001 12:20:32 Filter No filter						
Application	Dialog steps	Resp. time avg(ms)	CPU time avg(ms)	Wait time avg(ms)	DB time avg(ms)	Requested kB avg
Basis Components	18,502	819	264	7	338	369
Controlling	479	33,129	18,594	8	5,495	4,755
Cross-Application Components	638	6,214	3,883	2	1,945	1,179
Enterprise Controlling	3	1,358	39	0	1,142	200
Financial Accounting	6,050	171	40	10	97	25
Logistics - General	2,194	276	57	7	173	19
Logistics Execution	8,877	1,087	302	4	672	150
Materials Management	13,437	290	60	11	189	54
Production Planning and Control	8,341	485	111	8	260	71
Quality Management	5,187	261	80	7	142	75
Sales and Distribution	3,709	447	147	5	203	51
Treasury	195	461	58	49	288	12
Other	17,963	453	103	5	259	28
Total	85,575	765	279	7	327	157

Figure 325. ST07 Application monitor: Response times

16.6.1.8 ST03 Workload analysis

This transaction assists you in reviewing the workload on an SAP R/3 system (Figure 326).

Workload: Analysis of SAP System PRD

Choose for analysis Refresh Survey graphics Detail analysis menu

Database system DB400 Currently active servers 1
Data base server dfsapbak
This application server dfsapbak Performance database

Application servers with workload data

Central system

nonactive servers
dfsapprd

Figure 326. ST03 Performance workload analysis

The Choose for analysis push button allows you to select a specific server for review in a multi-server environment (Figure 327). This selection is followed by a time period selection window.

Workload: Analysis of SAP System PRD

Choose for analysis Refresh Survey graphics Detail analysis menu

Database system DB400 Currently active servers 1
Data base server dfsapbak
This application server dfsapbak Performance database

Application servers with workload data

Central system

nonactive servers
dfsapprd

Choose Server Name

Select application server
SAP system PRD

Server	Last record
DFSAPBAK	TODAY 10:51
TOTAL	TODAY 10:51

Choose Time Period

Choose a period for:
Reporting on dfsapbak

☒ Today
☐ Previous days
☐ This week
☐ Previous weeks
☐ This month
☐ Previous months

Figure 327. ST03 Performance workload (selections)

ST03 Workload overview

This window displays a summary of resource usage information and response time for the time period and server that you selected (Figure 328). The initial display is for all dialog tasks during the selected interval.

Performance: Workload Overview for Server dfsapbak




Figure 328. ST03 Performance workload overview

Use the scroll bar to access more information in the lower portion of the window. If you need to look at data for the Dialog or Background tasks only, select the appropriate push button at the bottom of the display.

ST03 Top 40 by response time

A display showing the dialog steps with the highest response times for a period can be obtained starting from the Workload Overview display by selecting the path **Goto-> Hit List-> Top 40 response**.

ST03 Top 40 by database requests

A display showing the dialog steps with the highest database requests for a period can be obtained starting from the Workload Overview display by selecting the path **Goto-> Hit List-> Top 40 DB requests**.

ST03 Response time analysis

A display showing the dialog steps for a period analyzed into groups by response time can be obtained starting from the Workload Overview display by selecting the path **Goto-> Summary reports-> Workload overview**. This information can be printed if required. See Figure 329.

Workload Summary Report




Figure 329. ST03 Workload: Response time analysis

ST03 Time profile

Use the Time Profile push button to see detailed hourly information (Figure 330) of the particular dialog tasks you selected including:

- Dialog steps
- Average response times
- Average CPU seconds used
- Time spent waiting for a work process
- Database response time

Workload: Time Profile Report




Figure 330. ST03 Workload: Time profile

ST03 Memory profile

The Memory Profile push button shows detailed information of memory utilization (Figure 331) including:

- Extended memory
- Private memory

Workload: Memory Profile

Long/short names

User names

Instance

SAP SystemPRDFirst record00:00:13Date03.01.2001

ServerdfsapbakLast record10:51:33

Instance no. 02Elapsed time10:51:20Task typeTotal

Sort: Tcode/Report

Entries: 392

Tcode or Report	Background job name	Dialog steps	Mem. total avg [Bytes]	Extended memory [Bytes]			Workproc. reservations	Priv. mem. avg [Bytes]
				avg used	max/session	max/t.actn.		
TOTAL		85575	3,288,271	2,236,224	190,227,814	190,227,814	70	0
<adm message		1,163	0	0	0	0	0	0
<no buffer>		1,887	0	0	0	0	0	0
(B) OTHER		2,334	0	0	0	0	0	0
AutoABAP		131	1,906,429	676,526	3,303,332	3,303,332	0	0
ALE		1	3,908,421	2,678,521	3,043,816	2,683,929	0	0
AL11		34	1,945,501	715,598	2,751,849	890,925	0	0
Buf_Sync		652	0	0	0	0	0	0
BD07		7	2,764,036	1,709,812	21,667,765	19,043,851	0	0
CA10		1	1,753,043	523,138	6,204,732	536,450	0	0
CA51		148	5,779,975	4,550,072	9,800,474	9,800,474	0	0
CK13		6	2,752,573	1,522,666	6,062,959	3,340,175	0	0
CM25		113	5,548,959	4,329,940	7,979,005	7,979,005	0	0
CO6I		113	2,290,404	1,289,066	8,648,404	8,609,108	0	0
CO10		1,987	4,698,420	3,476,560	43,063,435	43,063,435	0	0
COPI		83	2,415,400	1,630,040	51,248,088	51,248,088	0	0
CORM		1	1,979,757	749,852	762,533	762,533	0	0
CORS		36	3,777,757	2,616,182	9,695,753	9,452,825	0	0
CORT		32	3,210,001	2,172,270	7,662,575	3,726,084	0	0
COR1		34	2,870,777	1,795,569	12,621,538	6,776,053	0	0
COR2		1,132	4,808,421	3,779,517	44,657,233	25,292,116	0	0
COR3		56	3,226,261	1,996,378	3,516,498	3,516,498	0	0
COR5		107	3,490,062	2,639,474	8,434,268	7,160,615	0	0
COR6		1,555	4,219,993	3,197,313	21,344,362	14,820,315	0	0

Figure 331. ST03 Workload: Memory profile

ST03 Detailed analysis

The Detailed Analysis menu push button on the initial window of transaction ST03 (Figure 326) shows a menu in four parts (Figure 332). This allows you to make a detailed analysis of the workload through a series of push buttons:

- **Today:**

- *Workload* shows you the performance workload similar to that for transaction ST03.
- *Statistical records* presents the option to choose the records you want to investigate.
- *Last minutes workload* allows you to specify the last number of minutes you want to review.
- *Monitor response times* provides a graphical representation of the performance information.

- **Performance history:**

- *Recent period* allows the selection of the period and shows a workload analysis window.
- *Compare recent periods* enables you to compare the workload by:
 - Days of the week
 - Week
 - Month

- **Performance history - other server** allows you to compare another server on your SAP R/3 system.

- Recent period
- Compare recent periods
- **Global** compares all servers in the system:
 - Recent period
 - Compare all servers
 - Compare recent periods
 - Servers and users

Workload: Analysis of SAP System PRD




Figure 332. ST03 Detailed analysis

Use the scroll bar to see more of the menu options.

Note: The Workload push button shows a display similar to that for Workload Analysis in Figure 328.

ST03 Today's statistics

The Statistic records push button from the menu selection in Figure 332 shows a window to make a selection from the day's statistic records.

When you make your selections, the window in Figure 333 shows details by transaction.

Workload: Display Statistical Records

Long/short names Records Records 1 hour 1 hour Expansion

Server : dfsapbak
Statistic file: /usr/sap/PRD/DVEBMGS02/data/stat
Analyzed time : 03.01.2001/11:00:00 - 03.01.2001/11:01:33

End time	Tcod	Program	T	Scr.	Wp	User	Response time(ms)	Memory used(kB)	Wait time(ms)	CPU time(ms)	DB req. time(ms)	Load/Gen kBytes transfer	Phys. db changes
11:00:00	*	*	*	*	*	*							
11:00:00	Z_WI	Z_SD_WIS	D	1000	3	DESMONDE	294	2,204	288	5	0	1	0.0
11:00:00	VA01	SAPMV45A	D	4001	3	VECCHIETD	719	6,244	1	218	504	1	6.9
11:00:00	SART	RCOPCA02	D	0120	0	LEVENT	6	4,273	0	6	0	1	0.0
11:00:00	SART	RCOPCA02	D	0120	0	LEVENT	8	4,273	0	7	0	1	0.0
11:00:00	SART	RCOPCA02	D	0120	0	LEVENT	7	4,273	0	8	0	1	0.0
11:00:00	F-26	SAPDF05X	D	1104	0	NEWMANJ	10	2,897	0	10	0	0	0.0
11:00:00	SART	Z_FI_BOM	D	0120	9	LEONGS	8	5,492	0	6	0	0	0.0
11:00:01	SART	Z_SD_L0A	D	0120	1	SHAWD	165	4,005	0	37	0	2	0.0
11:00:01	MB51	RM07MMAT	D	1000	3	KIRANS	9	2,713	0	8	0	2	0.0
11:00:01	VT02	SAPMV56A	D	0020	4	HUDSONJ	162	4,776	1	82	90	1	16.9
11:00:01	SART	RCOPCA02	D	0120	0	LEVENT	7	4,273	0	7	0	1	0.0
11:00:01	VFB3	SAPMV60A	D	0104	2	SKELTONMI	816	2,957	1	70	47	10	21.5
11:00:02	F-26	SAPDF05X	D	1104	0	NEWMANJ	12	2,897	0	12	0	0	0.0
11:00:02	MB51	RM07MMAT	D	1000	3	KIRANS	19	2,713	0	13	5	1	1.2
11:00:02	VT02	SAPMV56A	D	0020	4	HUDSONJ	133	4,822	0	87	58	1	22.7
11:00:02	ME22	SAPMM06E	D	0210	3	CHERINR	15	4,735	0	11	5	0	0.0
11:00:02	ME22	SAPMM06E	D	0105	3	SUPPORT_PW	22	2,082	0	21	0	4	0.0
11:00:03	VA03	SAPMV75F	D	0120	2	SKELTONMI	12	2,301	0	11	0	0	0.0
11:00:04		RSM13000	2	3000	41	DFBATCH	121	0	0	81	40	4	29.9
11:00:04	ME22	SAPMM06E	D	0210	1	CHERINR	17	4,735	1	12	5	0	0.0
11:00:04	ME22	SAPMM06E	D	0101	3	SUPPORT_PW	1,337	2,719	1,207	30	102	6	3.4
11:00:04		MainMenu	D	0040	0	WARDP	715	1,630	0	28	6	193	0.1
11:00:04		(B)OTHER B			1	SAPSYS	23	0	0	12	13	2	0.2
11:00:04	VA03	SAPMV45A	D	4001	2	SKELTONMI	80	5,929	0	79	0	1	0.0
11:00:04	SART	Z_SD_L0A	D	0120	1	SHAWD	14	3,885	0	6	0	0	0.0
11:00:04		RSM13000	U	3000	25	DFBATCH	123	0	0	54	70	3	30.1
11:00:04	Z_WI	Z_SD_WIS	D	0120	3	DESMONDE	1,494	2,263	0	19	1,418	15	1.1
11:00:05	VA03	SAPMV45A	D	4001	2	SKELTONMI	79	5,866	0	75	5	2	0.0
11:00:05		MainMenu	D	0040	1	NAYLORS	14	1,630	1	13	0	2	0.0

Figure 333. ST03 Today's statistics

From this display, you can select a particular transaction and use the Expansion push button for a detailed view of the transaction (Figure 334).

Workload: Display Statistical Records

Long/shortnames Records Records 1 hour 1 hour Expansion												
Server : dfsapbak Statistic file: /usr/sap/PRD/DVEBMGS02/data/stat Analyzed time : 03.01.2001/11:00:00 - 03.01.2001/11:01:33												
End time	Tcod	Program	T	Scr.	Wp	User	Response time(ms)	Memory used(kB)	Wait time(ms)	CPU time(ms)	DB req. time(ms)	Load/Gen kBytes transfer
11:00:00	*	*	*	*	*	*						
11:00:00	Z_WI	Z_SD_WIS	D	1000	3	DESMONDE	294	2,204	288	5	0	1 0.0
							CPU time	5 ms	Number	Roll ins	1	
							RFC+CPIC time	0 ms		Roll outs	1	
							Response time	294 ms		Enqueues	0	
							Wait for work process	288 ms	Load time	Program	0 ms	
							Processing time	4 ms		Screen	0 ms	
							Load time	1 ms		CUA interf.	1 ms	
							Generating time	0 ms	Roll time	Out	0 ms	
							Roll (in+wait) time	1 ms		In	1 ms	
							Database request time	0 ms		Wait	0 ms	
							Enqueue time	0 ms				
11:00:00	VA01	SAPMV45A	D	4001	3	VECCHIETD	719	6,244	1	218	504	1 6.9
							CPU time	218 ms	Number	Roll ins	1	
							RFC+CPIC time	0 ms		Roll outs	1	
							Response time	719 ms		Enqueues	0	
							Wait for work process	1 ms	Load time	Program	0 ms	
							Processing time	211 ms		Screen	1 ms	
							Load time	1 ms		CUA interf.	0 ms	
							Generating time	0 ms	Roll time	Out	2 ms	
							Roll (in+wait) time	2 ms		In	2 ms	
							Database request time	504 ms		Wait	0 ms	
							Enqueue time	0 ms				

Figure 334. ST03 Today's statistical records (detail)

ST03 Compare recent periods

The Compare recent periods push button from the menu selection in Figure 332 shows you a performance comparison by day of the week. You can select the appropriate push button for a comparison by week or by month. See Figure 335.

Workload: Compare all Periods for Selected Server

Choose period

<<

>>

Graphics

Sort

By days

Compare by weeks

By months

Performance of SAP System PRD

List creation 03.01.2001 / 11:36:03

Analyzed server dfsapbak

Sorted by Dialog steps

Sum/average of all task types

Summary of average workload

Day	Dialog steps	Av. resp. time(ms)	Avg CPU time(ms)	Avg wait time(ms)	Avg load time(ms)	Av.DB rq time(ms)	Avg DB requests	Av. phys DB calls	Avg Kbyte requested
We 20.12.2000	273179	900	312	6	16	399	695.9	252.4	159.0
Tu 19.12.2000	264478	803	293	5	13	378	543.8	186.2	143.4
Th 21.12.2000	255763	712	252	5	17	297	283.1	70.4	127.7
Tu 02.01.2001	240429	789	306	10	15	335	333.1	87.1	165.6
Th 28.12.2000	217624	637	246	3	11	275	301.4	79.5	122.3
Mo 18.12.2000	212920	688	240	7	12	365	383.5	116.6	171.7
We 27.12.2000	203185	647	259	20	11	277	301.6	74.9	162.0
Fr 29.12.2000	181078	904	303	7	17	365	336.7	88.0	154.0
Fr 22.12.2000	175541	1,177	321	5	16	466	357.1	100.9	136.3
We 03.01.2001	86,860	759	276	9	13	323	256.6	66.5	169.8
Tu 26.12.2000	55,421	979	434	1	9	315	445.6	126.2	176.0
Su 31.12.2000	53,639	1,095	504	3	12	356	572.8	158.3	174.2
Sa 23.12.2000	47,157	2,323	619	3	16	743	691.2	188.9	236.3
Sa 30.12.2000	46,492	1,231	578	4	14	395	589.6	166.5	216.0
Su 17.12.2000	41,760	1,186	544	2	12	384	582.3	162.6	196.6
Su 24.12.2000	40,252	1,254	573	2	12	395	562.1	159.9	228.7
Mo 01.01.2001	31,342	1,394	639	2	15	419	608.5	168.9	220.4
Mo 25.12.2000	19,374	1,911	947	1	11	520	738.1	262.5	374.4
Sa 16.12.2000	17,024	1,686	769	5	14	683	844.6	249.3	316.8

Figure 335. ST03 Comparing recent periods (for a server)

16.6.1.9 SAP R/3 internal buffers

SAP R/3 implementation is supported by internal buffers that are used to store frequently used data, which typically remains unchanged such as:

- ABAP programs
- Windows
- ABAP Dictionary data
- Company specific table data

These buffers are maintained at the application servers. Statistical information on the usage of these buffers is available through SAP R/3 transaction ST02.

In a three-tier (separate database and application servers) implementation of SAP R/3, ensure that changes to buffered data are propagated to other application servers to maintain data consistency. This propagation does not apply to a two-tiered implementation. There are two parameters that control this:

- rdisp/bufrefmode
 - For central system: sendoff,exeauto
 - For three-tier system (database and application servers): sendon,exeauto
- rdisp/bufreftime: Specify time between buffer synchronization (seconds)

There are seven main groups of buffers:

- **Repository buffers** (also referred to as *dictionary* or *nametab buffers*) contain the active table and field definitions. These buffer sizes are defined in the instance profile by:

- rsdb/ntab/entrycount
- rsdb/ntab/ftabsize
- rsdb/ntab/irbdsz
- rsdb/ntab/sntabsz

- **Table buffers** are the single, generic, and resident table buffers that store records (rows) from a table. This classification determines if only a single row, group of rows, or an entire table is buffered. The attribute of a particular table is specified through SE13. The buffer sizes are specified by:

- Generic Key and 100% Buffered Tables (TABL)

- zcsa/db_max_buhtab
- zcsa/table_buffer_area

- Single Record Partial Buffered Tables (TABLP)

- rtbb/max_tables
- rtbb/buffer_length

- **Program buffers** store the executable ABAP programs. The programs that are preloaded are listed in stream files pxastat and pxanew in /usr/sap/<SID>/DVEBMGS<nn>/work, where <SID> and <nn> are the SAP R/3 system identifier and instance number respectively. The size of the buffer is defined in the abap/buffersize instance profile parameter. The abap/buffersize is set at a default of 600 MB (600000).

- **GUI buffers** store the screen presentation and menu buffers, and the sizes are defined by:

- sap/buflr_entries
- zcsa/presentation_buffer_area
- rsdb/cua/buffersize

- **Roll and paging buffers** store the user contexts, large lists, and long internal tables respectively. These values are specified in:

- rdisp/PG_SHM
- rdisp/PG_MAXFS

The rdisp/ROLL_SHM and rdisp/ROLL_MAXFS parameters can be ignored on the iSeries implementation.

- **Calendar buffers** store the factory and holiday calendars, and the size is specified in zcsa/calendar_area.
- **Cursor cache** is used to improve system performance by reducing the amount of SQL parsing. On the iSeries server, these values do not apply, and the display shows a hit ratio of 100%. On the iSeries, SQL statements are stored in SQL packages on disk. They are accessed using a package mapping algorithm and work process cursor cache is not used.

The following tables give an indication for setting these values, but you need to reset them after analysis of the buffer quality using R/3 transaction ST02. After the SAP R/3 application runs for a few hours, the buffers should be fully loaded. Once a stable operating environment is reached, the target should be able to achieve a hit ratio at the buffer of approximately 96%. The following settings were made for a production system with 8 GB of main memory and approximately 100 active SAP R/3 users with R/3 Release 4.6. The settings for a development system would normally be quite different. The latest information about the size

limitations of various R/3 buffers is contained in the SAP Note 121625 *AS/400: Buffer sizes*.

Important

The correct setting of the R/3 buffer sizes is crucial to the performance of your SAP system. We recommend that this task be conducted by an experienced Basis consultant with knowledge of the iSeries SAP environment. In the worst case, incorrect settings may prevent your SAP system from functioning.

Table 41. SAP R/3 buffer size indications

Parameter name	Value	Unit
rsdb/ntab/entrycount	30000	records
rsdb/ntab/ftabsize	30000	Kbytes
rsdb/ntab/irbdsiz	4000	Kbytes
rsdb/ntab/sntabsize	2500	Kbytes
zcsa/db_max_bufstab	10000	records
zcsa/table_buffer_area	64000000	bytes
rtbb/max_tables	1000	tables
rtbb/buffer_length	25000	Kbytes
abap/buffersize	600000	Kbytes
sap/bufdir_entries	10000	entries
zcsa/presentation_buffer_area	16773120	bytes
rsdb/cua/buffersize	16383	Kbytes
rdisp/PG_SHM	32768	Kbytes
rdisp/PG_MAXFS	32768	Kbytes
zcsa/calendar_area	500000	bytes

Note: These values are valid from R/3 Release 4.6 onwards.

Additional SAP R/3 parameters

The following parameters (also shown in Table 42) are some additional SAP R/3 parameters that need to be considered for good overall performance and functioning of the system:

- **rdisp/btctime:** The time interval (in seconds) at which the batch scheduler via the dispatcher attempts to start up batch jobs (default value 60).
- **rdisp/autoabaptime:** The time interval (secs) after which certain ABAP programs defined for the R/3 system are run automatically (default value 300).
- **rdisp/keepalive:** The time interval (secs) after which connections are checked using pinging when devices are not active (default value 1200).

- **rdisp/gui_auto_logout:** The time interval (secs) after which if the front end has not sent any data to the server the front-end connection is closed (default value 0).
- **rdisp/max_wprun_time:** The maximum run time (secs) for a process step within a dialog work process (default value 600). When this is exceeded, the dialog transaction is terminated and a short dump generated. *Prior to R/3 Release 4.6x*, for a long running SQL statement, the clock is reset once, so this time is doubled.
- **login/system_client:** Default logon client.
- **rdisp/tm_max_no:** Maximum number of open connections per instance. When this is exceeded, a user trying to logon will receive the message *Maximum number of connected terminals reached* (the default value is 200).

Table 42. Additional SAP R/3 parameters

Parameter name	Value	Unit
rdisp/btctime	60	seconds
rdisp/autoabaptime	300	seconds
rdisp/keepalive	1200	seconds
rdisp/gui_auto_logout	7200	seconds
rdisp/max_wprun_time	900	seconds
login/system_client	Default client	
rdisp/tm_max_no	Depends on number of users	

16.6.1.10 ST02 tune summary

The SAP R/3 transaction ST02 provides information on the activity of the SAP buffers discussed in 16.6.1.9, “SAP R/3 internal buffers” on page 450. An example of the information displayed by this transaction is shown in Figure 336.




Figure 336. ST02 Tune summary

Use the scroll bar to move across and down to see more information on the various buffer utilizations, including:

- Hit ratio
- Allocated size (KB)
- Free space (KB and%)
- Number of directory entries
- Free directory entries (number and %)
- Swaps
- Number of accesses

ST02 Current parameters

Information on the current SAP R/3 buffer parameters together with a brief description can be viewed by using the Current parameters push button shown in Figure 337.

Tune: Profile parameters for SAP buffers (dfsapbak_PRD_02)

System : dfsapbak_PRD_02 Profile parameters for SAP buffers
Date & time : 02.01.2001 14:55:20

Profile parameter	Value	Unit	Comment
Program buffer			
abap/buffersize	752000	kB	Size of program buffer
abap/pxa	shared		Program buffer mode
CUA buffer			
rsdb/cua/buffersize	16383	kB	Size of CUA buffer
The number of max. buffered CUA objects is always: size / (2 kB)			
Screen buffer			
zcsa/presentation_buffer_area	16773120	Byte	Size of screen buffer
sap/bufdir_entries	10000		Max. number of buffered screens
Generic key table buffer			
zcsa/table_buffer_area	96000000	Byte	Size of generic key table buffer
zcsa/db_max_buftab	15000		Max. number of buffered objects
Single record table buffer			
rtbb/buffer_length	16383	kB	Size of single record table buffer
rtbb/max_tables	500		Max. number of buffered tables
Export/import buffer			
rsdb/obj/buffersize	20000	kB	Size of export/import buffer
rsdb/obj/max_objects	20000		Max. number of objects in the buffer
rsdb/obj/large_object_size	8192	Byte	Estimation for the size of the largest object
Table definition buffer			
rsdb/ntab/entrycount	60000		Max. number of table definitions buffered
The size of the ITAB is nearly 100 bytes * rsdb/ntab/entrycount			
Field description buffer			

Figure 337. ST02 Tune summary: Profile parameters for SAP buffers

Use the scroll bar to view all of the current values of the SAP buffer parameters.

ST02 Detail analysis menu

Selecting the Detail analysis menu push button on the display shown in Figure 338 enables you to make a detailed study of the various SAP internal buffers. Use the scroll bar to see more push buttons on the Detail Analysis menu.

Tune Summary (dfsapbak_PRD_02)									
Current parameters Detail analysis menu									
System : dfsapbak_PRD_02 Tune summary									
Date & time of snapshot: 02.01.2001 14:53:36 Startup: 16.12.2000 10:48:24									
Buffer	Hitratio [%]	Allocated [kB]	Free space [kB]	Dir. size [%]	Free directory Entries	Swaps	Database accesses		
Nametab (NTAB)	99.71	10,082	6,968	84.94	60,000	50,967	84.95	0	668,113
Table definition	99.99	64,692	27,622	46.04	120,001	110,428	92.02	0	10,892
Short NTAB	100.00	7,192	1,892	75.68	120,001	116,235	96.86	0	3,772
Initial records	100.00	8,692	2,094	52.35	120,001	113,140	94.28	0	6,870
Program	99.97	765,582	316	0.04	94,208	85,164	90.40	30,020	117,192
CUA	99.99	16,383	1,058	6.86	8,191	4,247	51.85	0	4,111
Screen	99.26	16,380	5,096	32.51	10,000	9,135	91.35	0	864
Calendar	99.99	488	355	74.27	200	28	14.00	0	688
Tables	99.72	93,750	3,525	3.85	15,000	11,370	75.80	73	1,743,092
Generic key	99.78	16,383	0	0.00	500	131	26.20	10,350	384
Single record									
Export/import	98.65	20,000	1,834	11.09	20,000	15,227	76.14	60,119	0
SAP memory	Current use [%]	Max. use [kB]	In memory [kB]	On disk [kB]	SAP cursor cache	Hitratio [%]	Entries	Free [%]	
Paging area	19.20	100,640	366,224	524,288	IDs	100.00	67,500	100.00	
Extended Memory	24.26	2,034,688	7,557,120	8,387,584	Statements	100.00	22,500	100.00	
Heap Memory		0	1,953,121						
Call statistics	Hitratio [%]	ABAP/4 Requests	Processor Fails	Total calls	Database AvgTime[ms]	Rows affected			
2428 tables buffered									
Select single	91.14	481928,854	90382,156	573,130	0.000	391,546,698			
Select	78.49	150520,282	0	5,387,263	0.000	3,106,852			

Figure 338. ST02 Tune summary: Detail analysis menu

The window in Figure 339 presents an example of the detailed information that is displayed by selecting the Table definition buffer push button.

Tune: Detail Analysis (dfsapbak_PRD_02)			
Current parameters Buffered objects			
System : dfsapbak_PRD_02 Table definition buffer			
Date & time of snapshot: 02.01.2001 14:53:36 Startup: 16.12.2000 10:48:24			
Efficiency	Hitratio	%	99.71
	Hits		227,794,258
	Requests		228,466,927
	DB access quality %		99.71
	DB accesses		668,113
	DB accesses saved		227,794,407
	Reorgs		0
Size	Allocated	kB	10,082
	Available	kB	8,203
	Used	kB	1,235
	Free	kB	6,968
Directory entries	Available		60,000
	Used		9,033
	Free		50,967
Swaps	Objects swapped		0
	Frames swapped		0
Resets	Total		0

Figure 339. ST02 Tune summary: Table definition buffer

Figure 340 shows the Call Statistics window presented by using the Call Statistics push button.

Table Call Statistics

Tables
☐ all tables
☒ not buffered
☐ single key buffered
☐ generic key buffered

Time frame
☒ since startup
☐ today ☐ previous day
☐ this week ☐ previous week
☐ this month ☐ previous month

Server
☐ all servers ☒ this server
☐

☒ show statistics

Figure 340. ST02 Tune summary: Call statistics

Click **Show statistics** to view the performance analysis for each table.

You can see more detailed information by using the Overview<->Detail toggle push button (Figure 341).

Performance analysis: Table call statistics

Choose | All tables | Generic buffer | Single record buffer | Sort | Analyze table | Reset | Refresh

System : dfsapbak_PRD_02 Not buffered tables
 Date & time of snapshot: 02.01.2001 15:13:38 System Startup: 16.12.2000 10:48:24

Table	Changes/ Total (%)	ABAP/IV Total	Processor requests Direct reads	Seq. reads	Changes	DB activity Calls	Rows affected
Total		2721583347	2532139356	146132922	43311,069	464687337	381398141
AFRU	0.54	680,382	8	676,701	3,673	681,030	4,330
AFVC	1.48	789,875	112	778,099	11,664	804,023	36,088
AFVU	1.58	738,184	0	726,520	11,664	749,677	27,927
AFVY	1.58	738,235	0	726,571	11,664	749,762	27,927
AFWI	9.26	11,335	0	10,285	1,050	12,211	2,248
AGDB	0.00	610	0	610	0	98	462
AGKO	99.26	12,496	0	93	12,403	12,496	12,408
AKKB	12.90	31	0	27	4	31	4
AKKD	28.57	7	0	5	2	7	6
AKKP	3.81	105	101	0	4	206	105
ALMONISETS	0.00	44	16	28	0	69	44
ALTSTLOD	50.00	9,894	0	4,947	4,947	14,841	514,488
ANEX	0.00	297	0	297	0	433	0
ANIA	0.00	12	0	12	0	12	0
APQD	15.43	486,978	25,974	385,862	75,142	913,717	542,971
APQI	37.71	200,998	13,649	111,555	75,794	321,092	193,410
APSRV	0.00	1	1	0	0	1	0
AQGBB	0.00	10,899	0	10,899	0	2,576	276
AQGTB	0.00	658	0	658	0	658	0
AQLDB	0.00	3,370	0	3,370	0	1,436	501
AQXINT	0.00	49	0	49	0	49	0
AQXINTA	0.00	49	0	49	0	49	0
AQXINTT	0.00	49	0	49	0	49	0
ARCH_DEF	0.00	18	0	18	0	20	0
ARCH_IDX	0.00	7	7	0	0	14	7
ATAB_ARCH_NUM	50.00	18	9	0	9	32	18

Figure 341. ST02 Tune summary: Call statistics - Detail

Use the Storage push button to see information on SAP R/3 memory usage (Figure 342).

Storage Usage and Requirements			
Shared memory detail		Shared memory technical	
System	dfsapbak_PRD_02	Date	02.01.2001
Virtual storage usage and memory allocations		Time	15:08:06
Storage shared between work processes			
Memory key	used	free	allocated
Not in pool	824,682 K		824,682 K
Total	824,682 K	0 K	824,682 K
User storage for each of 41 work processes			
Roll area	16,380 K	Cursor Cache	2,000 K
Paging buffer	0 K	d+w user bss	NO DATA
Short storage	1,367 K	d+w user data	NO DATA
Total	19,747 K		
User storage for all work processes			809,627 K
Size of SAP executables (Disp+Work)		Shared text	NO DATA
Size of Extended Memory			0 K
Virtual memory allocated		*TOTAL*	1,634,309 K
ATTENTION : The size for the SAP executables (disp+work) IS NOT INCLUDED !			

Figure 342. ST02 Tune summary: Storage usage

Select the Parameter push button to view the installed instances for the SAP R/3 system (Figure 343).

Parameter Changes in SAPSYSTEM			
Select Period	Active parameters	History of file	Server names
02.01.2001 15:19:00 PRD dfsapbak			
Dates of most recent SAPSYSTEM parameter changes			
HostName	System I	Mod. date	
dfsapbak	02	08.12.2000	
dfsapprd	02	01.12.2000	

Figure 343. ST02 Tune summary: Parameters - Select instance

Selecting an instance and clicking the appropriate push button shows you:

- All active parameters
- History of parameter changes

See the display in Figure 344.




Figure 344. ST02 Tune summary: Instance parameters

16.6.2 Database monitor

The DB2 UDB for iSeries database monitor gathers information on iSeries database performance and is shown using SAP R/3 transactions such as ST04 and DB02.

Apart from providing you with information on the activity of the DB2 UDB for iSeries database, running the database monitor is important for gathering information for SAP's EarlyWatch support. In the past, this facility tended to be heavy on CPU and disk resources. However, since SAP R/3 Release 4.5, a more efficient version has been provided that enables data to be collected more frequently and retained over longer periods of time.

The database monitor provides valuable information on functions that:

- Consume a lot of system resources.
- Take an extremely long time to run.
- Create a temporary keyed access path during execution.
- Use the query sort during execution.
- May perform better by creating a permanent index.

The "old" version of the database monitor is retained for its value as a debugging tool.

16.6.2.1 Starting and stopping the 'old' DB2 UDB for iSeries monitor

The "old" DB2 UDB for iSeries monitor can be started and stopped using the native OS/400 STRDBMON and ENDDDBMON commands. The data is collected in a database file (QAQQPRF) in the R3<SID>DATA library. Note that this is the only file that is not journaled in the R3<SID>DATA library.

Important

The “old” database monitor should never be run continuously for lengthy periods of time because it will consume a large amount of disk storage. We recommend that you run for no more than 30 minutes.

- To start the DBMON program, type:

```
STRDBMON OUTFILE (R3<SID>DATA/QAQQPRF) JOB (*ALL)
```

- To stop the DBMON program, type:

```
ENDDDBMON JOB (*ALL)
```

16.6.2.2 The ‘new’ DB2 UDB for iSeries database monitor

The start and stop functions of the memory-based SQL database monitor are enabled through OS/400 APIs.

More information on the database monitor is contained in the SAP Online Documentation CD, under the section **Computing Center Management System-> CCMS Monitoring-> Database Monitor-> SAP/DB2/400 Database Monitor**.

Update database monitor statistics

For the memory-based SAP releases starting with release 4.6, the RSDB4UPD and RSDB4UPH programs are run to collect the information for ST04. For SAP releases before 4.6, the RSDB4T6M and RSDB4090 programs have to run before statistical information (about tables, indexes, etc.) for ST04 and DB02 transactions is available.

The memory-based database monitor gathers statistics about database usage and retains this information in memory until the ABAP RSDB4DMP program is executed (by the autoABAP program SAPMSSY6, which runs every NNN seconds as specified by the profile parameter rdisp/autoabaptime). RSDB4DMP writes the statistics to output files in the R3<SID>400 library. Alternatively, this can be forced to run from the transaction ST04 using the option **Edit-> Dump new outfiles**.

The data in the output files is then merged with the R/3 database performance tables by the RSDB4UPD program, which is scheduled by default to run every 5 hours (300 minutes). This value can be changed in the transaction ST04 by the menu path Edit-> DB monitor configuration-> Change, which displays the screen shown in Figure 345.




Figure 345. Transaction ST04: Database Monitor Configuration

The database history dump (RSDB4DBH) runs, by default, every 60 minutes to update the overview of database workload statistics that is displayed in transaction ST04.

For further information about the DB2 UDB for iSeries monitor, see *DB2 for OS/400 Database Programming*, SC41-4701.

16.6.3 ST04 Database performance

SAP R/3 transaction ST04 enables you to review the performance of the R/3 database on the iSeries server. Prior to presenting the initial window, the system gathers the required information. Therefore, you may notice a delay before the window in Figure 346 is shown.




Figure 346. ST04 Database performance monitor


This window shows the overview of the DB2 UDB for iSeries database performance monitor statistics. The date and time the database was last started (IPL) is shown on the first line, and then the date and time of the last merge of database monitor statistics are shown.

The database monitor will actively monitor the SAP R/3 work processes as long as the as4/dbmon/enable instance profile parameter is set to the value 1. Set the as4/dbmon/memory_size parameter to 1000 in order to reserve up to 1000 MB memory for the data in the memory.

To check the R/3 work processes and equivalent OS/400 jobs that are being monitored, choose the path **Goto-> Display database monitor activity**. The screen shown in Figure 347 shows the server name, OS/400 job number, user and name, and the DBMon Type field, which is the type of DB2 UDB for iSeries database monitor that is active for the job. This field can have the values:

- 0 Job is not currently monitored
- 1 Job is monitored by the old file-based DB2 UDB for iSeries database monitor that uses the STRDBMON and ENDDBMON OS/400 commands
- 2 Job is monitored by the new memory-based SQL database monitor

- Job is monitored by both the old file-based DB2 UDB for iSeries database monitor and the new memory-based SQL database monitor



DB2/400 Performance Monitor Overview (new DBMON)

17.11.2000 15:13:43 - Since Database Start

Database start	2000-11-16-15.40.58
Last merge of database monitor tables	2000-11-17-12.01.18
Last reset	- - -
Total database processing time (s)	202

Server name	Job number	Job user	Job name	DBMon.Type	DBMHANDLE
RCHASM23_R01_01	015839	R0101	WP03	2	R01
RCHASM23_R01_01	015840	R0101	WP04	2	R01
RCHASM23_R01_01	015836	R0101	WP00	2	R01
RCHASM23_R01_01	015844	R0101	WP08	2	R01
RCHASM23_R01_01	015846	R0101	WP10	2	R01
RCHASM23_R01_01	015892	R0101	WP01	2	R01
RCHASM23_R01_01	015841	R0101	WP05	2	R01
RCHASM23_R01_01	015838	R0101	WP02	2	R01
RCHASM23_R01_01	015843	R0101	WP07	2	R01
RCHASM23_R01_01	015842	R0101	WP06	2	R01
RCHASM23_R01_01	015849	R0101	WP13	2	R01
RCHASM23_R01_01	015852	R0101	WP15	2	R01
RCHASM23_R01_01	015847	R0101	WP11	2	R01
RCHASM23_R01_01	015848	R0101	WP12	2	R01
RCHASM23_R01_01	015845	R0101	WP09	2	R01
RCHASM23_R01_01	015850	R0101	WP14	2	R01

Figure 347. ST04 Display database monitor activity

Summary

To ensure that the new database monitor is set up and working correctly, you need to check that the instance parameter profile as4/dbmon/enable is set to 1, the SAP work processes are being monitored by the new database monitor in ST04 **Goto-> Display database monitor activity**, and that the RSDB4UPD and RSDB4DBH background jobs have been run successfully in SM37.

The initial ST04 overview window also shows the summary information collected by the database monitor on:

- User calls
- File activity
- Wait statistics
- Table scans
- Sorts
- Index creates
- Temporary files
- Indexes advised

16.6.3.1 ST04 Detail analysis menu

Select the **Detail Analysis** menu push button on the initial ST04 transaction as shown in Figure 348 for a more detailed review of database access performance.




Figure 348. ST04 Database monitor detail analysis menu

The detail analysis menu has the following options that you can select through push buttons:

- Overall activity:
 - Database lock monitor
 - Wait situations
 - File activity
- Resource consumption analysis:
 - DB monitor history
 - SQL requests
 - 50 slowest queries
- Query analysis:
 - Table scans
 - Sorts
 - Temporary files
 - Table locks
 - Index advised
 - Index usage
 - Index creates
- Additional functions:
 - State on disk
 - System catalog views
 - Performance tables
 - SQL packages

16.6.3.2 ST04 Database monitor: Wait statistics

The Wait situations push button shows a summary of the wait statistics for all of the instances on the system. The number of waits, total and average wait times in milliseconds, are displayed (Figure 349). These can be caused by:

- Internal machine activity
- Database locks
- Non-database locks

Database performance: Wait statistics by instance

Instance detail   Graphics						
03.01.2001 15:22:40 Wait Statistics by instance (times in milliseconds)						
Instance	DB lock waits	DB lock wait time	Average DB wait time	Non-DB lock waits	Non-DB lock wait time	Average non-DB wait time
dfsapbak_PRD_02	45,713	20,443,614	447.2166	421	40,178	95.4347
Total:	45,713	20,443,614	447.2166	421	40,178	95.4347

Figure 349. ST04 Database monitor: Wait statistics

You can select an instance for more detailed information by using the Instance detail push button. The statistics are then displayed for each SAP work process (the PID number and OS/400 job number are both displayed).

16.6.3.3 ST04 Database monitor: File activity

This display provides an analysis of activity for all physical files (tables) in the database. You can sort the information by selecting the column followed by the Sort push button, or filter the columns presented via the Set filter push button (Figure 350).

Database Performance: File Activity

03.01.2001 07:22:00 Statistics of physical accesses on DB2/400 database							
Filename	Log. Reads	Phys. Reads	Opens	Closes	Inserts	Updates	Deletes
KONP	2147,483,647	66,024,099	2,645	2,422	141,306	27,986	3,244
EKBE	815,741,344	319,223,268	1,456	1,308	76,319	5,599	432
BSIS	302,096,050	3,003,459	788	738	1,058,547	48	52,813
LTAK	250,454,112	4,047,527	1,338	1,272	3,732	2,667	0
D010INC	221,829,966	135,986	686	663	3,783	0	1,784
YBRP	184,769,928	21,861,009	1,934	1,786	245,733	0	0
CBPR	170,411,746	2,685,501	466	434	0	1	0
MARA	152,334,983	1,914,504	7,515	7,132	447	6,198	0
AUFK	128,630,839	432,290	9,463	9,203	2,621	12,506	0
TFDIR	93,936,466	36,518	1,096	1,014	48	14	17
CE4DF01	93,335,134	5,411,570	1,861	1,786	7,741	1,885	0
VBEP	93,178,995	14,537,975	1,362	1,252	17,284	4,406	278
JEST	92,849,403	1,284,260	2,167	2,009	214,125	43,365	4,557
COSP	79,756,281	19,270,457	5,028	4,698	108,427	4,339,282	65
AUSP	77,823,979	5,029,867	5,373	5,092	16,407	3	3,831
CE3DF01	74,037,356	6,489,900	3,033	2,911	104,101	200,666	0
APQI	66,345,000	307,165	4,210	4,074	13,259	55,105	13,281
VBAP	60,021,898	2,715,603	3,609	3,318	50,350	19,156	496
PLPO	56,799,134	8,093	1,220	1,138	36	43	0
MSEG	55,965,860	4,782,771	3,528	3,396	591,779	1	0
CBPT	55,096,849	1,243	899	851	0	1	0
VBPA	52,827,019	3,281,998	3,360	3,038	167,267	76,465	798
COSB	46,426,747	17,803,987	2,303	2,259	109,244	256,487	0
EDIDC	41,278,112	10,574,830	3,434	3,316	93,308	354,463	0
COSS	32,852,402	5,549,803	3,491	3,393	47,753	1,584,075	1
RESB	32,811,007	11,706,989	4,136	3,838	323,222	97,733	315,772
MARC	31,670,085	374,967	4,624	4,345	5,415	5,190	0
DBSTATHDB4	24,769,890	36,103	111	110	927,740	0	674,626

Figure 350. ST04 Database monitor: File activity

16.6.3.4 ST04 Database monitor: SQL requests

The SQL request push button shows details of SQL activity based on the selection information you entered. The information includes:

- SQL statement
- Count of executions
- Maximum duration
- Location of SQL package
- Module name

You can access additional data by using the More information push button.

Selecting an SQL request and using the Explain push button shows detailed information about the statement, including the access path and access method used in the execution.

Use the More information push button for detailed information.

16.6.3.5 ST04 Database monitor: 50 slowest queries

Information similar to that for SQL requests can be obtained for the 50 slowest queries during the measured period. Use the More information push button and the Explain push button for details.

16.6.3.6 ST04 Database monitor: Table scan

This option provides information on database accesses through the table scan access method. It shows detailed information on the number of executions, total time taken, and the number of rows read and selected.

Queries that use a table scan to select relatively few rows from a large table need to be analyzed further. Creation of an index over the columns used in the “where”

clause maybe considered to improve overall performance of the query. However, the frequency of use of the index created must also be considered in the evaluation.

16.6.3.7 ST04 Database monitor: Sorts

The Sort push button provides information on the occurrence of table sorting by the database. Information on the duration of the sort and the alternative options considered by the SQL optimizer are also indicated.

If you frequently encounter queries that sort the selected rows, consider creating an index over the columns used in the “order by” clause.

16.6.3.8 ST04 Database monitor: Index creates

Any temporary indexes created by the system during the time the database monitor was active are identified by this option. You are provided with information on the number of indexes created for each table and the average rows in the index.

Using the selection facilities and the Explain push button, you can identify the criteria used in the creation of the index and the time taken. Consider creating a permanent index that corresponds to the system-created indexes if they are used frequently and the average creation times are high.

16.6.3.9 ST04 Database monitor: Temporary files

This function provides an overview of temporary files created on the database. Temporary files are used for intermediate results when the optimizer determines that it needs to implement a query in two steps. These temporary files are discarded immediately after the job completes. The OS/400 Query Optimizer indicates why it needed to create the temporary file. This gives you hints on how to rewrite your query to avoid this.

16.6.3.10 ST04 Database monitor: Index usage

This function lists the indexes used by queries. It indicates the frequency of usage of the index and the number of rows read using this index.

Examine this option before you decide to delete any indexes that you have created. Never delete indexes that were created by the SAP R/3 system.

16.6.3.11 ST04 Database monitor: Index advised

This function lists the tables used by queries where the OS/400 SQL optimizer expects performance improvements by creating a new permanent index. It also provides the names of the fields to use as key fields in the “create index” statement. You can now select a particular SQL request after which an Explain button is visible. Click the **Explain** button, click the **More information** button, and scroll to the bottom. The index advised support must be used with caution. Adding indexes adds overhead and may not justify the cost savings for one particular query.

16.6.3.12 ST04 Database monitor: State on disk menu

This enables you to analyze the database and include all tables and indexes. Similar data is presented through the SAP R/3 transaction DB02. SAP R/3 tables and indexes of a particular SAP R/3 system are located in a single library (R3<SID>DATA) and are usually stored in the system auxiliary storage pool

(ASP1). Details of the system ASP are shown at the bottom of the display (Figure 351).

Database performance: Tables and Indexes

Refresh Screen

Data base system

DB system

DB400

File Statistic Analysis Date

03.01.2001 07:22:00

SAP system

PRD

Refresh File Statistics

Consistency Checks

Tables and indexes

	Tables	Primary Keys	Indexes	
Total number	16,868	16,867	2,908	Detailed obj.analysis
Total size (KB)	115,288	17,033	8,834	Deleted row analysis
				List damaged files
				Missing indexes
				Space statistics

	Tables	Primary Keys and Indexes	Views
Missing in database	0	0	0
R/3 DDIC	5	0	3

Current disk status

System ASP Size (GB)	488,217	% System ASP used	46.46420
Total ASP Size (GB)	496,606		

Figure 351. ST04 Database monitor: State on disk menu

Use the **Consistency check** push button to validate your database and ensure that all required database objects are included:

- *Find database tables without primary (indexes)* lists all database tables without any indexes. Define at least the primary key.
- *Database <—> ABAP Dictionary consistency* lists all objects defined in the ABAP Dictionary but not available in the DB2 UDB for iSeries database library.
- *R/3 kernel <—> ABAP Dictionary consistency* checks release consistency between the SAP kernel (library R3<rel>OPT) and the ABAP Dictionary (library R3<SID>DATA).

As shown in Figure 352, select one of the check boxes to see the names of the missing objects of that type. For the checks against the ABAP Dictionary, you are prompted to either display the results of the last check or to run a new check. Follow the information presented via the Information button on the results screen or contact SAP in case of inconsistency.

Note: This consistency check can only be used to detect missing objects. A more detailed check can be performed through SE12.

Database performance: Tables and Indexes

Refresh Screen

Data base system			
DB system	DB400	File Statistic Analysis Date	03.01.2001 07:22:00
SAP system	PRD		
		Refresh File Statistics	Consistency Checks

Tables and indexes	
Tables	
Total number	16,868
Total size (kB)	115,288
Missing in	
database	0
R/3 DDIC	5

Database checks

Installation checks:

☐ Find database tables w/o primary keys
☒ Check: Database <-> ABAP/4 dictionary
☐ Check: SAP kernel <-> ABAP/4 dictionary

☒ ☐

Current disk status			
System ASP Size (GB)	488,217	% System ASP used	46.46420
Total ASP Size (GB)	496,606		

Figure 352. ST04 Database monitor: Consistency check

A summary window of inconsistencies is presented (Figure 353). You can proceed to make a detailed review of the relevant tables.

Database performance: Tables and Indexes

Consistency check of 03.01.2001 15:42:16

- Objects missing in the database
 - Primary indexes 0
 - Secondary indexes 0
 - Tables 0
 - Views 0
- Unknown objects in ABAP Dictionary
 - DB tables 5
 - DB indexes 0
 - DB views 3
 - DB tables without unique index 1
- Optional indexes
 - Too many indexes created 0
 - Indexes not created 87

Figure 353. ST04 Database monitor: Consistency check (continued)

Use the **Deleted rows analysis** push button to see a list of database tables with deleted rows that can potentially occupy a lot of disk space (Figure 354).

The disk space occupied by deleted records may be retrieved using the OS/400 RGZPFM command.

Tables with deleted records

03.01.2001 07:22:00 Tables by deleted space					
Name	Rows	Deleted rows	%deleted	Row length	Deleted bytes
S00S	24,890	1,286,627	98	364	468,332,228
S571	1,250,376	747,176	37	240	179,322,240
ZDGL06	170,906	572,076	77	1,132	647,590,032
S0FM	25,215	550,269	96	141	77,587,929
S0C3	15,851	441,151	97	2,929	1,292,131,279
S00D	10,122	406,998	98	493	200,650,014
S970	760,202	195,914	20	172	33,697,208
UST12	150,941	177,455	54	116	20,584,780
ATAB	1,343,581	175,435	12	516	90,524,460
AGR_1251	76,627	150,318	66	165	24,802,470
TSPEVJOB	105,879	120,247	53	758	91,147,226
TST03	151,638	116,789	44	10,029	1,171,276,881
DBSTATDB4	607,240	84,340	12	187	15,771,580
APQD	402,650	74,303	16	7,938	589,817,214
ZMLEG00003	1,416	55,515	98	7,938	440,678,070
ZMSAP00002	1,361	50,036	97	7,938	397,185,768
T512T	190,974	47,705	20	43	2,051,315
SERPT	180,189	46,693	21	80	3,735,440
SNAP	7,000	46,651	87	1,642	76,600,942
TBTCF	19,166	45,407	70	1,161	52,717,527
UGR13	51,816	41,851	45	87	3,641,037
AGR_SELECT	22,114	40,753	65	169	6,887,257
AGR_TC0DES	15,775	37,171	70	91	3,382,561
RTXTL	147,033	37,125	20	312	11,583,000
T52C5	142,077	35,548	20	90	3,199,320
RF048	165,132	35,452	18	72	2,552,544
T52D7	140,313	35,183	20	13	457,379
UST10S	34,357	34,632	50	38	1,316,016
TC50C	39,911	33,221	45	114	3,787,194
DSYS_00007	129,816	32,511	20	114	3,706,254
TST01	38,159	31,969	46	213	6,809,397

Figure 354. ST04 Database monitor: Tables with deleted records

16.6.3.13 ST04 Database monitor: Detailed object analysis

With this function, details about one or more tables can be retrieved. After selecting a particular table, you can click the Analyze button and select:

- Table and its indexes
- History
- Structure

16.6.3.14 ST04 Database monitor: List damaged files

Damaged files are rarely encountered in a system. In rare situations (for example, after an abnormal termination of your system), a table may be damaged. When something unusual happens to your machine, we advise that you perform this function. The report of a backup also lists any damaged files. If an SAP table is damaged, consult SAP to determine what action to take.

16.6.3.15 ST04 Database monitor: Missing indexes

This function shows that indexes defined in SAP are not present in the database, or indexes present in the database are not defined in the ABAP Dictionary. If the indexes are missing immediately after an installation, this normally means that there is something wrong with your database, especially when it is a primary index (they end with a 0).

Note: All objects in the library R3<SID>DATA are examined, so non-SAP tables with no index or non-SAP indexes not defined in the ABAP Dictionary are also reported. If an SAP index is missing, consult SAP to determine what action to take.

16.6.3.16 ST04 Database monitor: System catalog tables

The tables shown in Figure 355 provide all of the information about the DB2 UDB for iSeries objects such as table and view definitions, indexes, interrelationships between objects, and so on.

System Catalog Views

The screenshot shows a dialog box titled "DB2/400 System Catalog Views". At the top left is a "Choose" button. Below the title bar, there is a list of radio buttons, each followed by a table name and a description. The first option, "SYSTABLES - Information about tables and views", is selected and highlighted with a dashed border. The other options are "SYSCOLUMNS - Information about column attributes", "SYSXCST - Information about constraints", "SYSKEYCST - Information about constraints keys", "SYSINDEXES - Information about indexes", "SYSKEYS - Information about index keys", "SYSVIEWS - Information about view definitions", and "SYSVIEWDEP - Information about view dependencies on tables".

Figure 355. ST04 Database monitor: System catalog tables

Note: These tables are large, take a while to return a display, and use a lot of CPU.

16.6.3.17 ST04 Database monitor: Performance tables

Tables used by the Performance Monitor functions can be used for your own analysis queries (Figure 356).

Database Performance Tables

The screenshot shows a dialog box titled "DB2/400 Performance Tables". At the top left is a "Choose" button. Below the title bar, there is a list of radio buttons, each followed by a table name and a description. The first option, "SQL Query Information (SQ1000 DB4)", is selected and highlighted with a dashed border. The other options are "SQL Arrival Sequence (SQ3000 DB4)", "SQL Index Used (SQ3001 DB4)", "SQL Index Created (SQ3002 DB4)", "SQL Sorts (SQ3003 DB4)", "SQL Temporary Files Created (SQ3004 DB4)", "SQL OptimizerTimed Out (SQ3007 DB4)", "SQL Subqueries (SQ3008 DB4)", "SQL Statement Text (SQTEXT DB4)", and "Hostvariable Values for most expensive Stmt. (SQHVAR DB4)".

Figure 356. ST04 Database monitor: Performance tables

16.6.4 ST05 SQL trace (ABAP level)

Use this function to trace the execution of SQL statements in an ABAP program (Figure 357). It is particularly useful where a user-written application is identified via ST03 Workload analysis as being heavy on system resources.




Figure 357. ST05 Trace SQL database requests

The recommended use of this function is to:

1. Identify a single transaction to trace.
2. Select **Trace on**.
3. Run the selected transaction.
4. Select **Trace off**.
5. Select **List trace**.

The list assists you in identifying the possible causes of delay in the transaction response time.

16.7 SAP R/3 tuning

This section introduces a few options to enhance the performance of SAP R/3. Attend applicable courses run by SAP to develop skills in managing your SAP R/3 environment.

Ensure that your iSeries is correctly tuned and there are no overall resource constraints on the OS/400 system. Refer to 16.5, “iSeries system tuning” on page 412, for more information.

16.7.1 SAP R/3 buffers

When the hit ratio for an SAP internal buffer is poor (less than 96%) and there is no free space left, you can increase the size of the buffer. If the hit ratio is good and there is still a large amount of free space in the buffer, you can reduce the allocation of space to this buffer.

Notes

- After the STARTSAP command is run, the system needs about one hour of intensive usage before the buffers reach a good hit ratio.
- You can be generous with the allocation of buffer space on the iSeries server because this memory is also pageable and, therefore, under control of the system's memory management algorithms.

Prior to SAP R/3 Release 4.6, there are restrictions on the maximum size to which certain memory buffers can be set. These restrictions are documented in SAP Note 121625: AS/400 Buffer sizes. Refer to 16.6.1.9, "SAP R/3 internal buffers" on page 450.

16.7.2 SAP work processes

It is important to minimize the queuing time of an SAP dialog waiting on a work process. Review the number of available work processes during periods of peak activity and ensure that there are adequate numbers of each type (dialog, batch, and update). Transaction SM50 shows all active SAP R/3 work processes, and it indicates the type and status of each process.

Transaction ST07 also assists you in forming an opinion regarding the adequacy of work processing by looking at the response time display. A high wait time can indicate the need for more dialog work processes.

The component report from OS/400 Performance Tools shows you the SAP R/3 jobs together with the corresponding resource utilization. Jobs with low resource utilization in each group can indicate if there are adequate work processes in that group. On the other hand, if all jobs of a group are more or less equal, some queuing for work processes may occur. However, you need to know which OS/400 jobs belong to each type.

16.7.3 File reorganizing

Reorganization of the table is normally not necessary in a DB2 UDB for iSeries environment. All files in the SAP R/3 Database library R3<SID>DATA have the attribute Reuse Deleted Records (REUSEDLTRCD) set to *YES, so as rows are deleted. They will be reused later with new rows. In case of performance problems with large tables containing a lot of deleted rows, for example, following a client delete, we advise that you reorganize the tables with a high amount of deleted rows. Tables with many deleted rows can degrade performance because of the paging in of blocks into memory when performing SQL SELECTs become inefficient if a large number of deleted records are returned in the block.

You can only reorganize the database library files when SAP R/3 is down. The files are not available during the RGZPFM command. A backup of the database library should be taken immediately after the reorganization is complete for

database recovery reasons. The APYJRNCHG and RMVJRNCHG journaling commands used in recovery to apply and remove journal changes do not work across tables that have been reorganized.

16.7.4 Build required indexes

When the system creates temporary indexes during execution, it impacts the response time of the transactions that require the index. Also, it can add a significant CPU overhead if it occurs frequently. This increased CPU overhead can impact other jobs. However, system built temporary indexes on the iSeries do not lock the table involved, unlike certain other database implementations.

Review information produced by the DB2 UDB for iSeries Database monitor through transaction ST04. If you notice certain temporary indexes are being built and used frequently, you may consider creating them as permanent indexes to improve performance.

However, avoid creating too many indexes because they add to the system overhead of index maintenance. Evaluate the cost to the system of maintaining additional indexes against the benefit of improved performance and response time.

16.7.5 Table locks

Investigate any extended table locks that can result in degraded response time. This can be caused by running conflicting applications at the same time.

The Table Lock push button available through transaction ST04 shows information on the occurrence of these conflicts.

16.7.6 Long running job

Discourage users from running long jobs interactively in a dialog work process. The `rdisp/max_wprun_time` instance parameter is used to terminate dialog programs that run for longer than the time specified in seconds. The default value is 600 seconds or 10 minutes. This type of job is best run in the background. This allows the user to work more productively without the workstation being locked up by one job.

16.7.7 Resource-intensive functions

Review resource-intensive functions, particularly if they are user-developed. Use transactions ST03 and ST04 to determine the highly resource-intensive transactions, programs and reports to examine. The SQL trace transaction ST05 assists you in tracing the SQL statements in an ABAP program. Refer to 16.6.4, “ST05 SQL trace (ABAP level)” on page 472, for more information on SQL trace.

16.7.8 Deleting SQL packages

Certain changes that occur on any iSeries server running SAP R/3 will effect the way the DB2 UDB for iSeries database optimizer works. SQL packages that contain the information used by the optimizer may, therefore, become out-of-date. It is possible, and sometimes recommended, to delete all R/3 SQL packages, or at other times, to delete specific packages that may be causing a problem with an individual SQL statement.

In the following cases, we advise you to delete SQL packages using the command:

```
DLTR3PKG SID(<SID>) PKGTYPE(*ALL)
```

The major system changes include:

- Processor upgrade, memory upgrade
- R/3 kernel upgrade
- OS/400 version upgrade
- OS/400 CUM package install
- OS/400 DB2 UDB for iSeries fix pack install
- Individual PTF installation if the PTF cover letter advises to delete SQL packages or recompile programs that contain embedded SQL
- Before and after an R/3 version upgrade
- After a R/3 installation
- After client copy, client delete R/3 operations, or a language import

Delete specific packages in cases such as the creation of new index over a table.

16.7.9 Short dumps

Examine the system daily for program terminations using the SAP R/3 transaction ST22 ABAP dump analysis. A large number of program terminations may degrade overall system performance. The display also indicates long running programs that have terminated due to a time-out.

Drill down to see the cause of the termination and the recommended action to follow. In many cases, you need to use the SAP Notes information database to search for the keywords suggested in the dump report.

16.7.10 Reporting performance problems

When reporting performance problems to SAP or IBM, you should include the following information:

- iSeries server model: Memory, CPUs, and disk
- SQL statement, SQL package and statement name, or SQL trace data (SAP R/3 transaction ST05)
- Table statistics for the tables used such as size, number of rows, and deleted rows
- Any known changes to the environment, for example, operating system fixes and kernel patches
- Quantification of the prior performance versus the current performance
- The steps you have tried already, such as index creates

16.8 LPAR performance

One of the most important performance factors in any multi-computer installation is the speed of the communications links. Prior to V4R4 and the LPAR configurations, the fastest inter-system communication link for the iSeries server was OptiConnect at 1063 Mbps. To achieve the 1063 Mbps speed, OptiConnect uses external optical bus hardware and software.

Logical partitions communicate using *Virtual OptiConnect*, which can achieve similar or better throughput rates than OptiConnect. One of the advantages of implementing LPAR in the R/3 environment is a performance improvement of inter-system functions due to the fast communication link. You may expect the following functions to particularly benefit from the high throughput rates of Virtual OptiConnect:

- FTP
- Correction and transports
- Remote client copy and client transport
- Save Restore (OptiConnect) processes between systems
- Objects, such as libraries, programs, data files, configurations, and directories and stream files in the IFS, can be saved from one partition and restored onto another partition over this communications link
- Data replication

16.8.1 Virtual OptiConnect and DMA

Virtual OptiConnect provides a hardware path to communications software (TCP/IP or SNA) that connects a logical partition with another partition. This path does not involve any input/output processors. This is shown in Figure 358.




Figure 358. Interpartition communication with Virtual OptiConnect

Virtual OptiConnect emulates external OptiConnect hardware by providing a virtual bus between logical partitions. This is done without any additional hardware requirements. To use Virtual OptiConnect, you only need to purchase either OptiConnect (a chargeable feature of OS/400) or OptiMover (a PRPQ).

The OptiConnect software will choose the virtual path over an external path if both paths are available.

When you create a logical partition, you must select whether you want Virtual OptiConnect (Interpartition OptiConnect) enabled on that logical partition. You may enable Virtual OptiConnect for a partition at any time. But you must install either OptiConnect or OptiMover before you use the OptiConnect function. When

you enable or disable Virtual OptiConnect, you must restart the system for the change to take effect.

Data is transferred directly from a memory in one partition to a memory in another partition. This is done using the Direct Memory Access (DMA) function, which provides high data throughput between partitions. DMA is not directly available to user programs, so the user programs have to use TCP/IP or SNA functions.

Figure 359 shows how data is transferred from partition to partition using DMA.




Figure 359. Virtual OptiConnect using DMA

The data resides on the disk in partition 0 and has to be written to the disk on partition 1. The following series of actions occurs (note the corresponding numbers to those in Figure 359):

1. A read request arrives at a disk IOP on partition 0. The IOP reads the data, and transfers it (using the DMA function) into memory pages in partition 0.
2. Pages are forwarded using DMA into an intermediate buffer on partition 0.
3. To move the data to the memory of partition 1, once again, a DMA is used. Pages are transferred from the intermediate buffer on partition 0 into memory on partition 1. No communication link is used for this memory copy. Virtual OptiConnect emulating the real OptiConnect hardware moves the data, from the buffer on partition 0 to the memory of partition 1, using a Virtual OptiConnect IOP.
4. Finally, DMA transfers the pages from the memory of partition 1 to the disk IOP who writes them to disk. This is done again by a DMA function.

16.8.2 Performance tests

For a performance comparison of Virtual OptiConnect versus LAN connection in an R/3-LPAR environment, two tests were performed in two different lab environments:

- Remote client copy
- Save and restore a library using the Save Restore Library (SAVRSTLIB) command

Note: These tests should serve as *examples* of possible performance gains of transferring data using Virtual OptiConnect instead of a token-ring connection. In the sample test results listed here, the numbers are actual timing results between logical partitions. Use these results only as a reference. Actual results in your particular environment will depend on the amounts of data to be copied.

16.8.2.1 Test 1: Remote client copy

This example demonstrates possible performance gains of doing a remote client copy over the Virtual OptiConnect as opposed to the token-ring connection. The test environment included:

- A two-processor iSeries server (one processor for each partition)
- 4 GB main memory (2 GB assigned to each partition)
- 200 GB total disk space (100 GB for each partition)
- 16 Mbps Token Ring adapter for each partition
- SAP R/3 Release 4.0B
- R/3 production system in the primary partition
- R/3 test system in the secondary partition

Two remote client copy tests were then performed between the two partitions. In test 1, the partitions were connected with TCP/IP over the token-ring LAN. In test 2, the partitions were connected with TCP/IP over the Virtual OptiConnect. Both tests used TCP/IP over the Token Ring for the SAP GUI connection between the PC and R/3 central instance. The test produced these results:

- **Test 1:** TCP/IP running exclusively through the token-ring.
Remote client copy took approximately six hours to complete.
- **Test 2:** TCP/IP running over Virtual OptiConnect
The remote client copy took approximately 3.5 hours.

16.8.2.2 Test 2: Save and restore a library

The second test used the ObjectConnect commands to save a 4 GB library in one partition and restore it to another partition. Again, two tests were run: one over the token-ring and one over Virtual OptiConnect. The test environment included:

- A four processor iSeries server (two processors for each partition)
- 8 GB main memory (4 GB assigned to each partition)
- 200 GB total disk space (100 GB for each partition)
- 16 Mbps Token Ring adapter for each partition
- SAP R/3 Release 4.5B
- R/3 production system in the primary partition
- R/3 replicated production system in the secondary partition

Two tests with the SAVRSTLIB command were performed between the two partitions using a library containing a 4 GB save file. Test 1 used TCP/IP over the token-ring LAN to connect the partitions. Test 2 used TCP/IP over the virtual OptiConnect between partitions. The tests produced these results:

- **Test 1:** TCP/IP exclusively over token-ring.
SAVRSTLIB took approximately one hour.
- **Test 2:** TCP/IP over Virtual OptiConnect.
SAVRSTLIB took approximately eight minutes to complete.

Chapter 17. Access Builder for SAP R/3

Access Builder for SAP R/3 is a toolkit that creates Java applications, applets, and JavaBeans capable of accessing an SAP system. Using the SAP Business Objects technology, which includes a Business Application Programming Interface (BAPI), applications can be quickly created in Java to access business data. This Access Builder will be of interest to you if you are currently using the SAP system to implement and track your business transactions.

This chapter provides an introduction on how Access Builder works conceptually and also an overview of the major components of Access Builder for SAP R/3.

17.1 Overview

Access Builder for SAP R/3 works in a similar manner to the Remote Method Invocation (RMI) Access Builder. You begin by creating proxy beans in VisualAge for Java that are derived from SAP business objects. Then you or your customers use the beans to access the business data in SAP. Figure 360 shows the process that starts with creating the proxy bean and ends with using the bean for business purposes.




Figure 360. Access Builder for SAP R/3 providing front-end access to SAP R/3

As shown at the top of the diagram, the SAP system usually resides on its own server. It includes a Business Object Repository (BOR) plus the actual business data from a corporation. In the Business Object Repository are objects with associated methods typically used in business. Interfaces are already defined for them. Collectively, these interfaces are referred to as the Business Application

Programming Interface (BAPI). In the diagram, the SalesOrder business object is shown along with its CreateFromData, GetList, and GetStatus methods.

Using the Access Builder, as shown in the center of the diagram, a Hockey Equipment Sales Form is created. This form is a proxy bean; specifically, it is a business object proxy bean. In the Hockey Equipment Sales Form are listed the items sold by the corporation: hockey sticks, helmets, and goalie pads. The sales form is now ready to be used by the customers.

At the bottom of the diagram, the Hockey Equipment Sales Form is accessed by customers using the World Wide Web. The sales form passes their requests back to the SAP system, updating the business data controlled by SAP accordingly.

17.2 Using SAP business objects and BAPIs

You use the Access Builder tool to browse meta information about SAP business objects and BAPIs of the SAP R/3 system. The Access Builder allows you to:

- Retrieve complete meta information from the Business Object Repository (BOR) within R/3
- Keep meta information of multiple R/3 systems
- Access meta information locally without R/3 connection

Access Builder for SAP R/3 generates proxy beans for SAP business objects, their BAPIs, and RFC modules. You can use these beans to design your application visually or to code an application or applet.

Access Builder is used to:

- Generate HTML documentation for specific SAP business objects and RFC modules
- Generate proxy beans for SAP business objects
- Generate proxy beans for RFC modules

17.3 Accessing the SAP system

The Access Builder for SAP R/3 provides an easy way to access SAP systems and provides the following facilities:

- Switch to different middleware technologies for development, intranet, and Internet environments with a common Java interface
- Manage R/3 access with a basic access library
- Implement an SAP logon panel by integrating a prebuilt Logon bean
- Dynamically access meta information with an additional run-time library

The SAP Access Builder interface provides access to the SAP system and its business objects.

The Common R/3 Access Interface defines a middleware-independent layer. You can leverage different ways to communicate with R/3 in your application without recoding. All generated JavaBeans are based on this interface and provide you with the same flexibility. Access Builder for SAP R/3 includes an actual layer via

Java Native Interface (JNI). This communication layer is best suited for the development environment, but can also be used for server-side Java or Java applications that are installed on fat clients.

You can also use another communication layer – CORBA/IIOP. This communication layer is often used for Internet/intranet environments with client-side Java.

R/3 Access Classes build the basic run-time access to R/3 Business Objects and BAPIs. They are used by all other components of the Access Builder package and by the generated JavaBeans. You can use them to dynamically access SAP Business Objects and BAPIs.

17.4 Logging on to an SAP system

To implement the SAP logon, you can integrate the Logon bean into your application, by:

- Including the user interface of an SAP Logon in your application
- Visually customizing the logon panel with your default values using the Visual Composition Editor of VisualAge for Java
- Setting the national language of the logon panel to your local needs

The online help shows you specifically how to log on through the VisualAge for Java interface.

17.5 Business Object Repository (BOR)

The BOR Access Classes retrieve information from the Business Object Repository (BOR) within the R/3 application server. The BOR contains all information about the SAP business objects, such as attributes, methods, and events. With the BOR Access Classes, you retrieve information on SAP business objects and BAPIs at run time. The Access Builder tool uses these classes to retrieve the information required to generate the Java proxies as beans. The generated beans do not depend on the BOR Access Classes, but only on the R/3 Access Classes.

Chapter 18. mySAP.com

mySAP.com, SAP's Internet strategy, can be defined as "an open application environment for e-commerce". It consists of portals, industry solutions, and Internet applications and services, with the aim toward integrated business collaboration. The mySAP.com product covers the entire range of SAP solutions. Existing SAP products form part of this integrated solution in the form of components.

mySAP.com contains integrated software components and industry-specific add-ons, including:

- SAP Workplace
- SAP R/3
- SAP APO
- SAP BW
- SAP BBP
- SAP CFM
- SAP CRM
- SAP EH&S
- SAP KW
- SAP SEM

18.1 SAP Workplace

The SAP Workplace is the front end to all of mySAP.com. It is an enterprise portal (installed on desktop computers, hand-held devices, and laptop computers) that acts as the users' gateway to all IT systems required to perform a specific task. SAP Workplace consists of Roles (logical collections of activities), a Launch pad (for navigation to activities pertaining the user role), and Mini Apps (information for specific roles). It is possible to connect to any URL, internal SAP R/3 systems, and all other SAP components from the Workplace, thereby offering integration across disparate systems. Figure 361 shows an overview of the SAP Workplace configuration and its components.




Figure 361. SAP Workplace configuration

SAP Workplace architecture consists of the Workplace server and middleware, with a Web browser installed on the front-end PC. The Workplace server manages information to and from the systems (both SAP, as well as non-SAP systems) that are accessed from the SAP Workplace. The middleware processes all Internet requests from the Web browser, SAP ITS, SAP GUI for HTML, and so on. It acts as a bridge between the SAP landscapes and the Web-based user front-end. The middleware currently requires the Microsoft Windows NT operating system. In a system environment with a Workplace server and more than one attached SAP components (for example, SAP R/3, BW and APO), several middleware servers must be installed. With the concepts of virtual instances, these multiple middleware servers can be installed on a single machine. You can determine the number of instances required by adding together:

- One Workplace instance
- One administration instance
- One instance for each attached SAP

Other SAP Workplace features include:

- **Single sign-on:** This feature eliminates the need for the user to sign on to each system separately in order to perform tasks. A user can perform all activities, across all systems, assigned to the user role after a single sign-on. SAP provides two methods for a single sign-on:
 - User ID and password
 - X.509 client certificates

The use of digital certificates and asymmetric encryption methods are also available through the mySAP Trust Center Service.

- **Mini-Apps:** This gives the user access to information and required functions immediately after sign-on. The Workplace includes a number of standard SAP Mini-Apps and additional ones can also be defined to suit user-specific roles.
- **Personalization:** This allows the user to structure the Workplace contents according to the specific user role. Internet addresses and “favorites” can be added as required.
- **Drag and Relate:** The user can select an object and drag it to another object in the launchpad, which will trigger a specific activity. For example, if the user selects and drags a customer number to the “Display customer details” object on the launchpad, the customer details will be displayed automatically. You can also link objects with other objects, as well as external Web pages.

The interface and exchange of data between the Workplace server and the Workplace components (for example, SAP R/3, SAP APO, and so on) is facilitated by an SAP Workplace Plug-In. SAP supplies a Workplace Plug-in with each release of the SAP Workplace.

The iSeries server fully supports the SAP Workplace by providing the database and application server environments for the SAP Workplace.

18.2 SAP R/3

This is the standard SAP R/3 system. SAP does not require any conversion from the existing release of SAP R/3 to mySAP.com

Both the database and application server environments of SAP R/3 run native on the iSeries server.

18.3 SAP Advanced Planner and Optimizer (APO)

SAP APO is part of SAP's comprehensive answer to the needs of supply chain management. It can be used as a planning system to replace MRP in SAP R/3 or as an Available to Promise (ATP) system. This can either replace or work with MRP in SAP R/3, depending on the level of ATP being used. APO contains a set of tools that facilitate planning and decision support throughout the supply chain. These tools include Strategic planning, Demand planning, Supply Network planning, Transportation planning, Production planning and Detailed scheduling, and Global Available-to-Promise, taking into account all components of the supply chain.

APO Demand planning employs the same technology as SAP BW and uses the SAP BW InfoCubes. Data exchange between SAP APO and SAP BW is, therefore, relatively simple.

APO also features integrated availability checks with SAP CRM and supports the roles and scenarios associated with the SAP Workplace. Some of the components of APO are:

- A core SAP R/3 system

SAP APO integrates fully into SAP R/3 by means of an SAP R/3 Plug-In. This Plug-In allows for data exchange (master data and transaction data) between

SAP R/3 and APO. This can be an SAP R/3 system running native on the iSeries. The SAP R/3 Plug-In for APO is also provided for the iSeries.

- APO Database

It is possible to run the APO Database on the iSeries as another SAP R/3 system.

- APO application server

APO currently requires a Microsoft Windows 2000 application server. The database can run on the iSeries server.

- APO Optimization server

This works closely with the front-end planning tools that an APO planner would use and allows planning changes to be propagated through the APO system. It can be installed on the planner's workstation or on separate application servers attached to the APO Database. Or, it can share an application server with the APO Database in the Central Instance. The APO Optimization server, therefore, does not run on the iSeries.

- APO Live Cache system

This allows for Rapid MRP calculations. It contains the same data as the SAP R/3 system, but omits many reliability components (for example, rollback) as the master data is contained in the SAP R/3 system. All the data is copied into memory. The aim of this component is to enhance performance. It allows multiple planning scenarios and the calculation of production schedules by taking into account changing constraints and other complex calculations that would not be feasible in SAP R/3. APO Live Cache currently requires the Microsoft Windows 2000 operating system.

18.4 SAP Business Information Warehouse (BW)

SAP BW is SAP's data warehousing application for SAP R/3 and other data sources. SAP BW integrates with SAP APO, SAP SEM, and SAP R/3, and serves as a reporting tool, complimenting SAP R/3 reporting.

BW consists of:

- BW front end (including the SAP GUI)
- BW server and administrator workbench
- BW extractors (installed in the SAP R/3 system by means of a Plug-In)

The iSeries server supports every BW release that is supported by SAP since release 1.2B.

18.5 SAP Business-to-business Procurement (BBP)

SAP BBP enables procurement across enterprise boundaries. BBP functions include managing purchase requisitions, purchase orders and reservations, approval or rejections, as well as invoicing and reporting. Tender and bidding processes through Internet Marketplaces are also possible, thereby, providing purchasing functions for e-commerce.

It possible to implement BBP in the following modes:

- A stand-alone system
- Linked to an existing SAP R/3 system
- Linked to a non-SAP R/3 system
- A combination of the above modes

BBP has a direct interfaces to SAP BW. You can access to BBP functions directly from the SAP Workplace.

18.6 SAP Corporate Finance Management (CFM)

SAP CFM manages financial resources and financial business processes. It consists of the following components:

- Transaction manager
- Market Risk analyzer
- Credit Risk analyzer
- In-house Cash management
- Portfolio analyzer
- Liquidity planner

CFM integrates into SAP R/3 by means of an SAP R/3 Add-on. It is also possible to run CFM as a stand-alone system in a distributed system environment. Specific roles exist in the SAP Workplace for SAP CFM.

18.7 SAP Customer Relationship Management (CRM)

SAP Customer Relationship Management is a set of applications designed to manage and extend customer and channel relationships and integrate these relationships with the back-end systems of any company. It provides a consolidated view of customer relationship data throughout the enterprise by using a central customer information database. In addition, it includes an Internet Pricing configuration for the following Internet Application components:

- Internet Sales - BBP (business to business)
- Internet Sales - B2C (business to consumer)
- Internet Sales - B2R (business to reseller)
- Internet Customer Sales Service

There is also a Mobile Sales and Mobile Service component where data is kept up to date by regular data exchange between the central CRM system and the local database on the mobile clients.

The iSeries fully supports CRM. The Message Hub, which communicates via RFC with the SAP R/3 OLTP system, can be run on the iSeries. An SAP R/3 Plug-In that supports data exchange with other mySAP.com components, such as SAP BW or SAP APO, is installed on the OLTP SAP R/3 system that runs on the iSeries server.

ITS is one of the architecture components of CRM, Internet Sales, and SAP R/3 Online Store. The CRM client is connected to the Communication Station, which converts DCOM calls to RFC calls. The RFC calls communicate with the CRM middleware component, or message hub. Network design is an important issue in the design between the different CRM components.

The CRM mobile client requires the Microsoft Windows 98 or Microsoft Windows NT operating system and runs Microsoft SQL. The front end to CRM (excluding the mobile client) has the same requirements as the SAP GUI.

The CRM Communication Station is an external middleware component required for communication between the CRM mobile clients and the CRM database and application servers. The CRM Communication Station requires the Microsoft Windows NT operating system. An additional Microsoft Windows NT server will, therefore, be required for components such as the DCOM connector and MTS. The CRM Administrator Workbench can run on this same server.

18.8 SAP Environment, Health, and Safety (EH&S)

SAP EH&S manages the environmental protection and industrial hygiene and safety requirements of an organization. The tools and applications for this component are integrated into mySAP.com, with specific links to the materials management, distribution, manufacturing, plant maintenance, HR, and cost accounting components of SAP R/3.

18.9 SAP Knowledge Management (KM)

SAP Knowledge Management is the follow-on to the SAP Advanced Training System (ATS). KM integrates with SAP HR components and SAP Workplace to allow personalized training. It also contains SAP knowledge for the implementation team, operational team, and end users. KM consists of a single repository and tools for authoring, production, translation, distribution, delivery, and retrieval. The KM content (SAP documentation, SAP QM manual, as well as SAP training courses and instructor guides) is available in all languages.

The SAP KM functions are contained in the SAP Workplace, allowing users to access KM using their own SAP Workplace roles. The iSeries server supports the Knowledge Warehouse tables or "Info Objects".

18.10 SAP Strategic Enterprise Management (SEM)

SAP SEM is a set of software functions and processes that enable executives and senior managers to implement and operate strategic management processes across the organization. SEM helps executives to simulate, analyze, monitor, optimize, and communicate the strategic aspects of the enterprise. It endorses value-based management methodologies to improve corporate value.

SAP SEM consists of the following components:

- Business consolidation
- Corporate performance monitor
- Business information collection
- Stakeholder relationship management

SAP SEM uses accounting information from the SAP R/3 system and requires a Plug-In to be installed in the SAP R/3 system. SEM also integrates with SAP BW and SAP APO. The SAP Workplace offers specific roles for SAP SEM.

18.11 Internet Business Framework Architecture

The design of mySAP.com is based on Internet Business Framework (IBF). This is SAP's concept of how to integrate all SAP business applications.

IBF is based on standard interfaces and allows separate SAP solutions to integrate with each other. New SAP products (for example APO) can be integrated into existing SAP environments, such as SAP R/3, without changing the existing infrastructure. Internet Business Framework, therefore, allows SAP applications running on iSeries servers to integrate with other SAP applications that are not currently available on the iSeries (such as SAP APO).

18.12 SAP Internet Transaction Server (ITS)

Internet Transaction Server is the SAP connection to the Internet. It consists of two main parts: the A-gate and the W-gate. The W-gate Web server currently supports the following products:

- Microsoft Internet Information Server 3.0
- Microsoft Internet Information Server 4.0
- Netscape Enterprise Server 3.51

ITS requires the Microsoft Internet Explorer 5 Web browser. Currently, ITS requires the Microsoft Internet Information Server (IIS) product, which only runs on the Microsoft NT 4.0 operating system.

The ITS components form the front end of mySAP.com Workplace. You need to install a separate ITS for each SAP system in the Workplace landscape.

18.13 SAP Business Connector Interface (BC)

The Business Connector is an open interface that allows for the integration of SAP systems and non-SAP systems through the Internet. It is an XML-based interface, allowing data exchange (for example, transmitting orders and invoices between buyers and sellers) through the mySAP.com Workplace portal, using FTP or SMTP as its communications protocol. It also includes an RFC client and server.

The SAP proprietary RFC format is converted to XML (or HTML). Therefore, this eliminates the need for SAP software on the other end of the communication line.

SAP BC has built-in support for SAP's specification of IDoc-XML and RFC-XML. The corresponding compiler (JDK, C-compiler, VB) is, however, required if you want to develop your own program modules. SAP BC is written in Java and runs on any machine for which a Java Virtual Machine of Release 1.1.6 or higher is available. A shared version of SAP's C-library *librfc* is needed for the connection to SAP R/3.

The SAP Business Connector is available on the iSeries and requires no additional software to be installed. For installation details, refer SAP Note 0350575 (SAP BC 3.1 on AS/400).

For more information on mySAP.com, go to the Web site:

<http://service.sap.com/mysapcom>

Chapter 19. SAP R/3 and Domino connection

Lotus provides a variety of integration technologies that can be used to integrate Domino and SAP applications. They address customer requirements to extend SAP R/3 managed data to internal staff, customers, and business partners accessing Lotus Domino applications. The resulting integration provides seamless integration framework for optional utilization of Domino and SAP R/3.

This chapter introduces different techniques to connect your SAP R/3 with Domino server, using one of the following solutions:

- Lotus Domino Connector for SAP R/3
- Lotus Script Extensions for SAP R/3
- Lotus Domino Access to SAP R/3 business workflow
- Domino Mail Transfer Agent (MTA) for SAP R/3

19.1 Lotus Domino Connector for SAP R/3

The Lotus Domino Connector for SAP R/3 is a system file developed by Lotus Development to manage data authentication and translation between SAP R/3 server data and other Lotus Domino Connectors. This Domino Connector for SAP R/3 may be used with the following Lotus Enterprise Integration products:

- **Domino Enterprise Connector Services (DECS)**

This is a technology supplied with the Domino server release 4.6.3 and higher server that enhances Domino applications with real-time data access or update capabilities to external source systems. It includes Enterprise Resource Planning systems such as SAP R/3 without programming. Domino Server 4.6.5a or 5.01 or higher server release is required to use DECS with the Lotus Connectors for Enterprise Resource Planning Applications, which includes the Connector for SAP 1.5.

- **Lotus Enterprise Integration (LEI)**

A server-based data transfer product facilitating scheduled high volume transfer and synchronization of data across Lotus Domino Connector sources. LEI includes data transfer template forms for sophisticated scheduled data transfer without programming. It also includes support for Lotus Script and Java programmatic transfers. Domino server release 4.6.3 or higher and LEI server release 3.0 or higher required to use LEI with Lotus Domino Connector for SAP.

- **Lotus Connector Lotus Script Extensions (LC LSX)**

The LC LSX enables programmatic access and manipulation of Lotus Domino Connector source data, allowing full programmatic control over data transfer. All supported Domino Connectors may utilize the same Lotus Connector API object model, exposed in Lotus Script classes, to syntactically access a wide variety of enterprise data sources.

- **Lotus Connector Java Class Library**

The LC Java Class library enables programmatic access and manipulation of Lotus Domino Connector source data for the development of server-based Java programs to control data transfer operations across Lotus Domino Connector sources. The LC Java Classes are available from the Lotus Domino Web site at: <http://www.lotus.com/domino>

For detailed information on installation and configuration, refer to *New Enterprise Integration Functions for Lotus Domino for AS/400*, SG24-6203.

The Lotus Domino Connector for SAP R/3 controls authentication and data transfer from Domino to SAP R/3 application data. The Connector was developed using SAP R/3's Remote Function Call Software Development Kit (RFCSDK). It enables the execution of any SAP R/3 RFC, Business Application Programming Interface (BAPI), and updates to SAP R/3 applications using Batch Data Input. Using the Domino Connector for SAP R/3 technology ensures that data transfers and queries are processed through the SAP R/3 application layer. This preserves the business logic and data validations contained in SAP R/3 RFC and transaction interfaces that comprise SAP R/3 server processes. Therefore, reading and writing SAP R/3 data is always performed through the application layer, not by directly accessing back-end database tables. All the business rules provided by RFCs and SAP R/3 transactions are maintained.

19.2 Integration of OS/400, Lotus Domino, and SAP R/3

Domino for iSeries is a full-function Domino server that combines the industry's leading messaging and collaboration solution with the iSeries server's inherent value of integration. As a native OS/400 application, Domino for iSeries is the first Domino server that uses the IBM 64-bit PowerPC RISC technology.

The combination of Domino for iSeries with the help of the Lotus LSX and SAP R/3 provides a highly integrated solution. The LSX for SAP's R/3 system enables a bi-directional data flow between Lotus products, such as Lotus Notes or any SmartSuite application, and the R/3 system.

19.3 Architecture

SAP supports a whole set of Remote Function Calls (RFCs) and Business Application Programming Interfaces (BAPIs). These interfaces enable external systems to run transactions and access data on the R/3 server. SAP provides these interfaces with the RFC SDK Toolkit.

LotusScript is an embedded, object-oriented, BASIC-compatible structured programming environment similar to Microsoft's Visual Basic. The LSX for R/3 is a set of plug-ins to LotusScript that enable application developers to access RFC and BAPI objects from the LotusScript Integrated Development Environment.

19.4 Benefits of Lotus Connector Lotus Script Extensions (LC LSX)

Lotus LSX in an R/3 environment offers the following benefits:

- **Improved analysis of R/3 data:** Updates, manages, and distributes data across the enterprise via already familiar tools. Eliminates geographic boundaries related to enterprise wide collaborative efforts.
- **Improved reporting of R/3 data:** R/3 output can be created, maintained, and distributed through existing Lotus Notes infrastructure.
- **Access to R/3 information via Web browsers:** Makes R/3 data available via inter and intra networks, using widely accepted and existing user interfaces.

The existing Domino workflow allows users to update tasks, have them approved, and then automatically entered into the R/3 environment.

- **Support for remote and mobile users:** R/3 reports and extracted data can be stored in Lotus Notes databases allowing remote and mobile users access to R/3 information locally while not connected to the network.

19.5 Sample scenarios for Lotus Connector Lotus Script Extensions

The Lotus Corporation provides the LSX for R/3 as a download from their Web site (see Section C16, for details). Within this download kit, there is a sample database. After you install the kit, you can set up the following scenarios.

19.5.1 Scenario 1

The LSX for R/3 and the Notes DB are installed on a client workstation. In this configuration, a Notes application runs on a client workstation. The application accesses the R/3 System using the LSX installed locally. See Figure 362.




Figure 362. LSX extensions locally on the workstation

19.5.2 Scenario 2

The LSX for R/3 and the Notes DB are installed on the Domino server. In this configuration, the Domino server accesses the SAP R/3 system. Lotus Domino and SAP R/3 can be on separate iSeries servers (Figure 363).

The Notes clients are connected to the Domino server. In this case, retrieving or updating data happens between the Domino server and R/3 System. Since Domino is also a Web server, it is possible for a user to connect to Domino with a Web client.

This function opens up your R/3 data to users on the Internet or intranet. Simply create a Notes application (following the instructions about how to Web-enable a Notes DB). Then, you have an Internet-ready interface to your R/3 system.




Figure 363. Domino server and R/3 system on separate iSeries servers

Note: If you plan to run Lotus Domino and SAP R/3 on the same system, you must take performance considerations into account.

The Domino server and the SAP R/3 system can be on one iSeries server (Figure 364).




Figure 364. LSX extensions on the Domino server and SAP R/3 on one iSeries server

You can also set up a server agent. You can view an agent as a kind of batch job that asynchronously exchanges data between the Domino server and the SAP R/3 system. If you are maintaining data on the Domino server database, the synchronization between the Domino and the SAP R/3 can happen at different times.

19.6 Domino Access to SAP R/3 business workflow

When a Domino user thinks of workflow, normally that person thinks of e-mail circulating through the various steps necessary for the workflow to occur. When an R/3 user thinks of workflow, often e-mail is not even involved, except for completion notification. When a SAP GUI user looks into the integrated inbox, there may be work items there, but these are really pointers to the workflow application. They are not e-mail. How can these two different views of workflow be integrated?

We have created a modified Domino mail template as one part of this answer. By using Lotus Script and the LSX for R/3, the periodic server agents in this template download the work items found in a given R/3 user's work item inbox and create what appear to be e-mail documents in the Domino user's mail database. Depending on what release of R/3 you have, there will be different amounts of functionality available to the Domino user for these work items. The user may then perform some processing of these work items from within Domino, or the user may invoke the SAP GUI and be taken directly to the screen that would have been presented if the user was working completely from within the SAP GUI. The end result is that the user only has to look in one location for all that they need, and that location is the Domino mail database.

For installation information, refer to *New Enterprise Integration Functions for Lotus Domino for AS/400*, SG24-6203.

19.7 Domino Mail Transfer Agent (MTA) for SAP R/3 R1.5

For many years, Lotus Domino and SAP R/3 have both supported SMTP for e-mail connectivity between the systems. Why is SMTP not good enough? When you define a business process with e-mail components, you have to be sure that the e-mails reach their destinations, and they are actually read and processed at that point.

With SMTP, you don't get this kind of information. It is a black hole of uncertainty. SAP realized this to be true and created SAP Connect, a mail interface where you can be sure that mail gets where it is intended to go. Lotus and SAP jointly developed this new MTA using the SAP Connect technology. It provides robust logging and tracking and good Rich Text support for your Windows and OS/2 based desktop applications.

For installation information, refer to *New Enterprise Integration Functions for Lotus Domino for AS/400*, SG24-6203.

19.8 Further information

For more information about Lotus LSX, refer to *Enterprise Integration with Domino.Connect*, SG24-2181. The Lotus LSX for R/3 is available as a free download from: <http://www.edge.lotus.com>

This site contains information and documentation, sample databases, and installation instructions.

Chapter 20. Using an IBM Network Station as an SAP front end

This chapter shows you how to integrate a network station in your R/3 environment to start a SAP GUI. It also provides the following information:

- Short presentation on network station
- An overview of three alternatives using a network station as an SAP front end
- Step-by-step configuration

20.1 IBM Network Station: Basic description

The IBM Network Station is a desktop network computer with:

- Local processor and memory
- Terminal emulation capabilities
- No disk (neither hard drive, nor diskette)
- Some form of network connection
 - Token-Ring 16/4 Mbps
 - Ethernet 100/10 Mbps
 - Twinaxial
- IP-based protocols, such as:
 - TCP
 - FTP
 - Telnet
 - NFS
- A Web browser
- A Java Virtual Machine (JVM)

In short, a network computer can be thought of as a *smart terminal*.

The most important characteristic is that all of the software required by the Network Station can reside on one or multiple servers in the network. Network Station Management Version 3.0 allows you to configure the Network Station to access boot information from different servers. The servers can be:

- Kernel server
- Configuration server
- Logon server

The software is downloaded on demand when the network station is powered on or when the user activates new functions. To boot the network station, the user must first load the operating system, configuration, emulation, and application programs from a host that acts as the boot and authentication server.

Anyone of the following protocols can be used to boot an IBM Network Station from the iSeries server:

- Boot protocol (Bootp)
- Dynamic Host Configuration Protocol (DHCP)

The chosen protocol determines how much information is stored on the Network Station's Non-Volatile Random Access Memory (NVRAM). Bootp requires more information than DHCP from the network station.

Trivial File Transfer Protocol (TFTP) is for other data communication purposes, such as the read or write access to files on a remote server. The Network File System (NFS), if available, can be used for mapping remote disk drives so that they appear to be local. The NFS server stores common configuration files and makes them available to an NFS client (the IBM Network Station).

An IBM Network Station offers the benefits of a non-programmable terminal, plus:

- Low cost of ownership
- Central management of software and data
- Simplicity in installation and administration
- Data access and security
- Graphical interfaces with browser-based administration features

20.1.1 Introducing various scenarios

To connect your network station in an R/3 environment, such as an SAP front end, there are three different proven configurations that are described in detail in the following sections of this chapter. We give only a short description of them for an overview.

20.1.1.1 Configuration one

The first configuration includes an iSeries server, which is your R/3 server and booting server for your network station, and a PC server for your SAP GUI application for the network station. The network station kernel, the network station configuration file, and some applications, such as a 5250 emulation program, are stored in the iSeries server's IFS. They are downloaded during the boot process.

To access your R/3 system, you need SAP GUI, which is stored on the PC server and runs on top of the Windows-based Terminal Server (WTS) from Microsoft. WTS is installed in place of a Windows NT operating system and allows multiple concurrent users and desktop export to the network station using Metaframe Independent Computing Architecture (ICA) protocol from Citrix Systems, Inc. WinCenter is the program, that, when called, transfers the WTS desktop window to your network station. The WTS display appears with the SAP GUI icon. By selecting the icon, you can start SAP GUI on the PC server from your network station.

20.1.1.2 Configuration two

The second configuration is the same as in the previous configuration, except it uses an Integrated xSeries Server instead of a PC running Windows-based Terminal Server (WTS) for the network station and using the Metaframe (ICA) protocol. In this situation, you do not need to use a separate PC server. To learn more about the advantages of using an Integrated xSeries Server, visit the IBM Web site: <http://www.ibm.com/servers/eserver/series/windowsintegration/>

In this configuration, the SAP GUI is stored on the Integrated xSeries Server. As in the previous configuration, the SAP GUI icon is on the WTS desktop window that is displayed on the network station.

20.1.1.3 Configuration three

The third configuration uses the Java version of the SAP GUI. The iSeries configuration of the SAP R/3 database and application server are the same as in the previous configurations.

Normally, a SAP GUI data stream is used to communicate between an application server and the SAP GUI application on the PC. The Java SAP GUI adds a conversion between the normal SAP GUI data stream and Java and makes the Java available via HTTP protocol. Since the data now comes from HTTP in a Java format, the traditional SAP GUI application on the PC is not needed. Instead, you can use any Java-capable browser to connect to the HTTP server that is connected to the application server. On the network station, the Java Virtual Machine is available for this purpose. User input is submitted back to the application server using browser and is converted from the Java data stream back into the SAP GUI data stream. The early version of Java SAP GUI does not have all of the functions of SAP GUI. It is also slower due to the conversion and Java interpretation.

20.1.2 WTS with separate PC Server

The network station boots the necessary software from the iSeries server that is stored in the iSeries server's IFS in a directory called /networkstation. You should already have the network station up and running.

For further information on IBM Network Station, refer to:

- *AS/400 - IBM Network Station - Getting Started*, SG24-2153
- *IBM Network Station Manager for AS/400 Users*, SC41-0632
- *IBM Network Station Guide for Windows NT*, SG24-2127

To access the network station manager's administration HTML file located on your iSeries booting server, configure your network station browser session default URL to:

`http://hostname/networkstation/admin`

You can use any Web browser for the administration of the network station.

We assume that the SAP GUI is already installed on the Windows-based Terminal Server. Complete the following steps:

1. Start the browser session by clicking **IBM browser**.
2. Sign on to the Network Station Manager.
3. Click **Startup** and then **Menus**.
4. On the Menus window, choose to make your settings available for the entire system or for a specific user only. After making your selection, click **Next**.
5. On the next window, scroll down to the **Remote program menu items** section:
 - a. Add a meaningful name to the **Menu item label** field.
 - b. Add the IP address of your PC server running WTS to the **Remote host** field.
 - c. Add `wincenter` to the **Program to run** field.

- d. Add the following string to the **Optional parameters** field:

```
-display ${IP}:0 -username :${user}  
-password ${password}
```

Using these general parameters, you can use the defined menu item from any network station where you sign on. Otherwise, you must have the authority to call the WinCenter program. For security reasons, we recommend that you sign on manually instead of specifying a user name and password in the menu definition.

6. When you complete the preceding task, click **Add a Remote Program**, and then click **Finish**.
7. Reboot your network station. After the network station is up, click the menu item you previously defined. If your settings are correct, the WinCenter window appears. If not, go back to the preceding steps and review your configuration to fix it.
8. Find your SAP GUI icon, start it, and sign on to your R/3 system.

20.1.3 Windows-Based Terminal Server on the Integrated xSeries Server

Instead of a PC server, an Integrated xSeries Server can run WTS and SAP GUI. With such configuration, you can use the iSeries server's integrated environment to manage your Integrated xSeries Server and SAP GUI application.

We assume that you configured your Integrated xSeries Server, the Windows-based Terminal Server is running on it, and you installed SAP GUI. With these prerequisites, follow the steps described in 20.1.2, "WTS with separate PC Server" on page 499.

Two configurations have been tested. One configuration has the Integrated xSeries Server installed on the same system as the R/3 application server. In this configuration, you can take advantage of the internal LAN between the iSeries and the Integrated xSeries Server. By using the internal LAN, you reduce the performance loss you would experience by contending with the network traffic on the external LAN.

In the other tested configuration, the Integrated xSeries Server and the R/3 application server were stored on two different iSeries servers. In this configuration, the external LAN is used to communicate between the SAP GUI and the application server.

20.1.4 Java SAP GUI

The Java SAP GUI emulates the standard SAP GUI interface and provides the ability to run R/3 transactions over the Internet or intranet using a Java-enabled browser. Java is a machine-independent programming language and forms the basis of SAP's new R/3 ultra-thin client. This makes it possible to access the R/3 system from any network station.

The Java SAP GUI can be downloaded from SAPNet at:

<http://sapnet.sap.com/javagui>

The software download is free for all SAP customers and business partners. A SAPNet user ID and password are required to access SAPNet. If you have access to SAPNet, you can download the Java SAP GUI from any of the available

SAPNet servers. Refer to SAP Note 77429. In the future, the Java SAP GUI will be provided on the SAP R/3 Presentation Server CD. To download SAP GUI, perform these steps:

1. Download the Java SAP GUI into a temporary directory.
2. Before installing the Java SAP GUI, read the readme.txt file for any last minute installation instructions.
3. Run the program that unpacks the Java SAP GUI files.
4. Run the **setup.exe** file and follow the installation instructions.

For more information, refer to 'BcJavGui.hlp' in the SAP GUI directory. The SAP GUI is installed under the root directory of your selected Web server. For example, the \InetPub\wwwroot\SAP GUI path is used with Internet Information Server.

To display the Java SAP GUI on the IBM network station, a JVM is downloaded from the server when booted. Running the SAP GUI for Java from a PC client requires a Java-enabled Web browser to be installed. To set up or verify that the following Web browsers are Java-enabled, perform the following steps:

- For **Microsoft Internet Explorer 3.0**, click **View-> Options-> Security-> Enable Java Programs**
- For **Netscape 4.0**, click **Edit-> Preferences-> Advanced-> Enable Java and Java Scripts**

Before running the Java SAP GUI, edit the SAP GUI HTML file in the <drive>\InetPub\wwwroot\SAP GUI directory and change the value of the sap_connect parameter. For example, if your hostname is *SAPtest* and the instance number is *00*, the sap_connect parameter is specified as sap_connect value=/H/SAPtest/S/3200.

Specify the hostname exactly as it appears in the HOSTS table and observe uppercase and lowercase characters.

Ensure that the ORBIX server daemon is running on the Windows NT server. If not, run **startORB.BAT** in the <drive>\InetPub\wwwroot\SAP GUI directory. As an alternative, click **Start-> Programs-> SAP GUI in Java-> Start the ORB Daemon**.

To call the Java SAP GUI from a PC client, start the Web browser and specify the URL (for example, <http://hostname/directory/htmlname.html>). To run Java SAP GUI from an IBM network station, we manually edited the user's Startup.nsm file on the boot server.

Note: Since the downloaded Java SAP GUI is a beta release, changes can be expected that may effect the setup of the IBM Network Station or PC client.

20.2 Further information

For more information, see:

- *AS/400 - IBM Network Station - Getting Started*, SG24-2153
- *IBM Network Station Manager for AS/400 Users*, SC41-0632
- *IBM Network Station Guide for Windows NT*, SG24-2127

Chapter 21. Problem determination

This chapter is intended to help you diagnose and manage problems that you may encounter in running an SAP R/3 system on the iSeries server.

Refer to Chapter 16, “Performance management” on page 365, for information regarding tuning for performance. Refer to 9.12, “Resolution tips for printing problems” on page 212, for the problem determination of printer problems.

21.1 Implementation of SAP R/3 on the iSeries server

This section describes how SAP R/3 is implemented on the iSeries server. That means which jobs on operating system level are used to run R/3.

The application is started with the STARTSAP command and ended with the STOPSAP command. The STARTSAP command has the DLTSPLF parameter with a default of *YES, which is causing the deletion of all spooled files created during the previous run. If R/3 is stopped and restarted for debugging, use the command:

```
STARTSAP DLTSPLF(*NO)
```

21.1.1 Job structure

When you run the WRKACTJOB SBS(R3_<nn>) command, where <nn> is the instance number, the display shows you all of the R/3 jobs for that instance running on the iSeries server in one particular subsystem. These jobs correspond to the R/3 processes:

- Startup job
- Message server
- Dispatcher
- Gateway
- Work processes

The number of each type of work processes started in an instance is controlled by the instance profile. If it is a central instance, a message server and an enqueue process are active. The first instance started on each iSeries server also runs the SAP OS collector process.

The STARTSAP command first starts the R/3 subsystem (instance) and submits a background job to this subsystem that runs the R/3 start program SAPSTART. The SAPSTART program starts the R/3 services and processes associated with each instance as shown in Figure 365.

```

Work with Active Jobs
AS23
11/10/00 18:15:27
CPU %: 8.9 Elapsed time: 00:47:57 Active jobs: 191

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
R3_01 QSYS SBS .0 DEQW
DISP R0101 BCI .0 SELW
GWRD R0101 BCI .0 SELW
MSG_SERVER R0101 BCI .0 SELW
R3_R01_01 R0101 BCH .0 PGM-SAPSTART EVTW
SAPOSCOL R0101 BCI .0 SIGW
WP00 R0101 BCI 1.6 SEMW
WP01 R0101 BCI .6 SEMW
WP02 R0101 BCI 1.2 SEMW

Parameters or command
====>
F3=Exit F5=Refresh F7=Find F10=Restart statistics
F11=Display elapsed data F12=Cancel F23=More options F24=More keys
More...

```

Figure 365. R/3 jobs in the subsystem

The STOPSAP command does not end the subsystem job and the SAP performance collector job per default. If you want to end these jobs, specify the ENDSBS (*YES) parameter.

21.1.1.1 Job status

Table 43 explains all jobs in the subsystem for the R/3 central instance. The Initial status field shows the status of the jobs after the instance has been started.

Table 43. Jobs, job type, and status

Subsystem/job	Description	Job type	Initial status
R3_<instance>	Subsystem job	SBS	DEQW
DISP	Dispatcher	BCI	SELW
GWRD	Gateway (reader)	BCI	SELW
MSG_SERVER	Message server	BCI	SELW
R3_<SID>_<instance>	SAPSTART program	BCH	EVTW
SAPOSCOL	SAP performance collector	BCI	SIGW
WRK<nn>	Work processes	BCI	SEMW

Also you should see in the QSYSWRK subsystem (starting with V4R4M0) the QXDAEDRSQ job. Here, listener (Port 4402) for remote database requests from application servers, which are started by the STRTCPSVR SERVER (*EDRSQ) or the STARTSAP SID (*DB) command.

If the system is running in a three-tier environment, with the OptiConnect or OptiMover product, the database server has jobs named APIAnnnnnn in

subsystem QSOC that correspond to the work process jobs on the application server.

21.2 Working with job logs

This chapter explains how you can find the corresponding job and its log on the iSeries server based on the R/3 work process. Every R/3 work process has its corresponding job. And every job has its associated job log. The job log for a job records those messages that were sent between the R/3 job and the operating system.

The job log is initialized when the job is started and remains in existence until the job ends. When the job ends, the job log may be written to a spooled file where it can be viewed or printed.

21.2.1 Changing the job attributes

It depends on the job attributes and the end code of the job whether a spooled file will be generated when a job ends. Because the R/3 kernel is monitoring for many messages, you receive an end code 0 even if an error has occurred. To receive a job log in any case, you can change the job descriptions used by the R/3 work processes using the command:

```
CHGJOB JOB(R3400/R3_<nn>) LOG(4 00 *SECLVL)
```

This change will be in effect after the next restart of the R/3 system.

During the installation or upgrade of an R/3 system, the install jobs inherit the attributes from the job that started the installation. To make sure that a spooled file will be created, before (re)starting the installation, change your interactive job using the command:

```
CHGJOB JOB(*) LOG(4 00 *SECLVL)
```

In case of an upgrade, the R3UP program will set the job attributes automatically. To force a job log here, press F21 (User Window) from the R3UP screen to get a command line and enter the command:

```
CHGJOB LOG(4 00 *SECLVL)
```

Then leave the window with F12 and continue with the upgrade.

21.2.2 Work process overview

Use the R/3 transaction SM50 (Work process overview) to display the R/3 work processes as shown in Figure 366.

No	Ty.	PID	Status	Reason	Start	Err	Sem	CPU	Time	Report	Cl.	User	Action
0	DIA	7349	Running		Yes					SAPLTHFB	300	TEST	
1	DIA	7350	waiting		Yes								
2	DIA	7351	waiting		Yes								
3	DIA	7352	waiting		Yes								
4	DIA	7353	waiting		Yes								
5	DIA	7354	waiting		Yes								
6	DIA	7355	waiting		Yes								
7	DIA	7356	waiting		Yes								
8	UPD	7357	waiting		Yes								
9	UPD	7358	waiting		Yes								
10	ENQ	7360	waiting		Yes								
11	BGD	7361	waiting		Yes								
12	BGD	7362	waiting		Yes								
13	BGD	7363	waiting		Yes								
14	SP0	7364	waiting		Yes								
15	UP2	7365	waiting		Yes								

Figure 366. Work Process Overview (SM50)

Use the R/3 dispatcher monitor (DPMON) if access to R/3 through SAP GUI is not available to process transaction SM50. You can see the same information from a 5250 screen by running the following command:

```
CALL PGM(DPMON) PARM('pf=/usr/sap/<SID>/SYS/profile/<instance profile>')
```

Figure 367 shows the output of the command.

Workprocess Table											
=====											
No	Ty.	Pid	Status	Cause	Start	Err	Sem	CPU	Time	Program	Cl User

0	DIA	20	Wait		yes						
1	DIA	21	Wait		yes						
2	DIA	22	Wait		yes						
3	DIA	23	Wait		yes						
4	DIA	24	Wait		yes						
5	DIA	25	Wait		yes						
6	DIA	26	Wait		yes						
7	DIA	27	Wait		yes						
8	UPD	28	Wait		yes						
9	ENQ	29	Wait		yes						
10	BTC	30	Wait		yes						
11	BTC	31	Wait		yes						
12	BTC	32	Wait		yes						
13	SPO	33	Wait		yes						
14	UP2	34	Wait		yes						
q - quit											
m - menu											
-->											
====>											
F3=Exit F4=End of File F6=Print F9=Retrieve F17=Top											
F18=Bottom F19=Left F20=Right F21=User Window											

Figure 367. Work process table (DPMON)

There are three different ways to find the job log of a dialog work process on the iSeries server. In this example, we want to display the job log of work process number 0.

21.2.3 The WRKPID command

Note the process ID (PID) of that job from transaction SM50, which in this example, is 7349. Make sure that the R/3 kernal library is in the library list of your job. Use the Work with Job by PID (**WRKPID**) command on the operating system level to map the R/3 work process to the job. In our example, we use the **WRKPID PID(7349)** command. This command runs the **WRKJOB** command for the corresponding job as shown in Figure 368.

Work with Job (WRKJOB)

Type choices, press Enter.

Job name	> WP00	Name, *
User	> R0101	Name
Number	> 013088	000000-999999
Output	*	*, *PRINT
Option	*SELECT	*SELECT, *STSA, *DFNA...

Bottom

F3=Exit	F4=Prompt	F5=Refresh	F10=Additional parameters	F12=Cancel
F13=How to use this display	F24=More keys			

Figure 368. WRKPID command

Press Enter and type option 10 to display the job log of the job.

21.2.4 The database lock monitor (DB01)

Note the work process number from transaction SM50, in this case, is 0. Follow these steps:

1. Call transaction DB01 (Database lock monitor).
2. Select the instance, and click the **Execute** button. The window in Figure 369 appears.

Instance	Jobname	Jobuser	Jobno.	DB PID	Status
RCHASM23_R01_01	QXDARECVR	QUSER	013107	7369	TIMW
RCHASM23_R01_01	WP07	R0101	013095	7356	SEMW
RCHASM23_R01_01	WP01	R0101	013089	7350	SEMW
RCHASM23_R01_01	WP04	R0101	013092	7353	SEMW
RCHASM23_R01_01	WP03	R0101	013091	7352	SEMW
RCHASM23_R01_01	WP10	R0101	013099	7360	SEMW
RCHASM23_R01_01	WP12	R0101	013101	7362	SEMW
RCHASM23_R01_01	WP05	R0101	013093	7354	SEMW
RCHASM23_R01_01	WP02	R0101	013090	7351	RUN
RCHASM23_R01_01	WP11	R0101	013100	7361	SEMW
RCHASM23_R01_01	WP08	R0101	013096	7357	SEMW
RCHASM23_R01_01	WP13	R0101	013102	7363	SEMW
RCHASM23_R01_01	WP00	R0101	013088	7349	RUN
RCHASM23_R01_01	WP06	R0101	013094	7355	SEMW
RCHASM23_R01_01	WP09	R0101	013097	7358	SEMW
RCHASM23_R01_01	WP14	R0101	013103	7364	SEMW
RCHASM23_R01_01	WP15	R0101	013104	7365	SEMW
RCHASM23_R01_01	QXDARECVR	QUSER	013105	7367	TIMW

Figure 369. Database Lock Monitor (DB01)

3. Select WP00 and click the **Display job log** button.

Note

You can save the job log from that window by selecting **System-> List-> Save-> Local file** or by entering %pc. You should always use the unconverted format. This may be very helpful in case you have to forward the job log to support personnel.

21.2.5 The WRKACTJOB command

Since SAP R/3 kernel release 4.6, you can map the R/3 work process to the job by simply using the WRKACTJOB command. Follow this process:

Note: The work process number from transaction SM50, in this case, is 0.

1. Enter the command:

```
WRKACTJOB SBS (R3_<nn>)
```

Refer to Figure 365 for the output.

2. Search for job WP00. Use option 5 (Work with) and option 10 on the next screen to display the job log.

21.2.6 Printing and locating the job log

If the job found by these steps is still active, you can use the command to obtain a spool file of the job log:

```
DSRJOBLOG JOB(qualified job name) OUTPUT(*PRINT)
```

Use the `WRKSPFL SELECT(*CURRENT)` command or the `WRKJOB JOB(*)` command and option 4 to locate the generated spooled file. If the job is no longer active, use option 4 from the `WRKJOB` command and look into the spooled file QPJOBLOG.

If you cannot identify a certain job by the above steps (this mainly happens when problems occur during installation or upgrade of R/3), you can also scan through all spooled files generated by the user profiles `<SID>OFR` (for startup, installation, or upgrade problems), `SAP<nn>` (for SAP releases up to 4.0B), and `<SID><nn>` (for SAP releases 4.5A and later). As usual, replace `<SID>` with the R/3 system ID and `<nn>` with the instance number.

When you look at the list of spooled files, you may notice that some files are named QPRINT. Some R/3 functions send output to the standard output device (STDOUT) or the standard error device (STDERR). Since all of the jobs under the R/3 subsystem are batch type jobs, output sent to STDOUT or STDERR appear as spooled files. When debugging a problem, look at these type of files, as well as the job logs.

21.3 R/3 profiles

To locate the R/3 profiles, use the iSeries server command:

```
WRKLNK OBJ('/usr/sap/<SID>/SYS/profile') DETAIL(*EXTENDED) DSPOPT(*ALL)
```

There are three profiles used to control the settings during the startup of R/3:

- **Default profile**

```
DEFAULT.PFL
```

The default profile indicates basic information about the SAP R/3 system such as the system name, database host name, and so on. It also contains defaults for all instances in an SAP R/3 system. Anything specified in the instance profile overrides information in the default profile.

- **Start profile**

```
START_<instance>_<host>
```

The programs identified in the start profile are started first. This includes the message server, the dispatcher, and the performance collector. Other jobs are started from the dispatcher.

- **Instance profile**

```
<SID>_<instance>_<host>
```

The instance profile indicates the R/3 parameter overrides for the instance including the number of work processes to be started.

Usually these profiles should be maintained through R/3 transaction RZ10 (Profile Maintenance). But in emergency cases (if R/3 doesn't start), it is also possible to change them with the `EDTF` command. When adding or changing a profile, use caution to enter the data in the correct format.

21.4 Trace and log files

This section describes SAP R/3 and IBM traces and log files that can be used for problem analysis.

21.4.1 System log (SM21)

The main source of error information in R/3 is the system log that can be displayed with transaction SM21. In case of errors, it shows the error message, some information about where in the application the error occurred, and often a reference to a so-called short dump and the qualified job name of the job that received the error.

The system log entries for an instance are in the `'/usr/sap/<SID>/<instance>/log/SLOG<nn>'` file. The system log entries are encoded. Looking at them using a 5250 session may give some insight. In general, the entries are unreadable, unless you are familiar with the encoding scheme.

21.4.2 Short dumps (ST22)

When an R/3 job behaves unexpectedly, the application usually generates a short dump, which allows you to analyze ABAP program dumps. These short dumps can be reviewed through the transaction ST22.

The short dump has very detailed information about the failing source statement (including the contents of some variables). However, the SQL statement shown in the short dump is written in ABAP syntax, which is different from the statement that is actually by the OS/400 database code.

21.4.3 Developer traces (ST11)

The various R/3 trace files can be viewed using transaction ST11. However, if the SAP GUI presentation layer is unavailable, the logs can also be viewed from a 5250 session. The data in the trace files is in plain text and is not encoded in any form. All trace files are located under the `'/usr/sap/<SID>/<instance>/work'` path.

You see a number of files in the work directory. The file name `dev_disp` is the trace file for the dispatcher task. A name of the form `dev_w<nn>` is the trace file for work process `<nn>`. The file `dev_ms` is for the message server, while the file `dev_rd` is for the gateway reader. In addition, there are trace files for RFC, startup, and shutdown.

21.4.4 SQL trace (ST05)

The SQL trace can be switched on to protocol the SQL statements being passed to the database management system (DB2 UDB for iSeries).

21.4.5 Startup log

If R/3 does not start correctly or not at all, look at the startup log that can be found in `'/usr/sap/<SID>/<instance>/work/sapstart.log'` file.

21.4.6 Upgrade logs

The upgrade to a higher R/3 database release is done in many phases. Each phase writes its own log into the `'/usr/sap/put/log'` directory.

An appendix of the SAP upgrade manual contains a list of the upgrade phases and the logs created during each phase.

21.4.7 Transport logs

The '/usr/sap/trans/log' directory contains logs from operations such as:

- Client import/export/copy
- Import/export of SAP R/3 corrections

21.4.8 XDA trace

With PTF SF64765 in OS/400 V4R5M0, a new trace facility has been implemented to help debugging problems in the interface between the SAP R/3 application and database or operating system functions. The trace can be activated in three different levels for one specific job, for all jobs in an R/3 instance, or system wide. The primary intention of the trace facility is to make it easier to gather information about failing function calls and to reduce the need for specific trap code in case of a problem. It is not assumed that end-users or system administrators use the data. To gain a better understanding about the traced data, it is helpful to become familiar with the OS/400 File APIs QxdaProcessExtDynEDRS, QSQPRCED, QxdaProcessImmediateEDRS, QxdaCallProgramEDRS, QxdaConnectEDRS, QxdaDisconnectEDRS, QxdaCommitEDRS, and QxdaRollbackEDRS.

21.4.8.1 Controlling the trace level

Data can be traced at the levels NONE, ERROR, INFO, and VERBOSE. If no trace level is set, NONE is assumed as the default. The trace level is defined with the help of an environment variable named QIBM_COMPONENT_TRACE_LEVEL. Environment variables can be set in a job or system-wide with the ADDENVVAR, CHGENVVAR, or WRKENVVAR commands. The value of this variable defines the trace level to be used. For the four available trace levels, the values shown in Table 44 need to be set.

Table 44. Values for trace levels

Trace level	Value of environment variable
NONE (0)	'XDA,NONE'
ERROR (1)	'XDA,ERROR'
INFO (2)	'XDA,INFO'
VERBOSE (3)	'XDA,VERBOSE'

In the SAP R/3 environment, there are several methods to control the trace level by environment variable QIBM_COMPONENT_TRACE_LEVEL:

- Set the environment variable QIBM_COMPONENT_TRACE_LEVEL at the system level. The setting will be used in each job that is started afterwards. This option should be used with care, and it should be ensured that the value is reset (or the environment variable is deleted at system level) after the analysis is completed. Otherwise, the system will continue to create user spaces named QP0Z<nnnnnn> in library QUSRSYS and fill up the system.
- Set the environment variable QIBM_COMPONENT_TRACE_LEVEL in the job that issues the STARTSAP command. The setting will be used for all SAP work processes. In a three-tier environment, the setting will also be used for the shadow processes on the database server.
- For a job that is already running, change the trace level by setting the environment variable QIBM_COMPONENT_TRACE_LEVEL in an interactive

job to the required level and then using the CHGUSRTRC JOB(<qualified job name>) command to activate the change in the requested job. The CHGUSRTRC command can also be used to set the trace buffer size and the wrap option.

In all three cases, the change of the trace level is documented by the message CPF9898 in the job log with the text "XDA TRACE LEVEL CHANGED FROM <x> TO <y>". If the trace level of the shadow process on the database server is changed during STARTSAP because of the setting on the application server (see method number 2 from above), it will be indicated by a message CPF9898 in the job log with the text "XDA TRACE LEVEL CHANGED FROM <x> TO <y> PER CLIENT REQUEST". The trace can be used for the SAP work processes, the database shadow processes (named QXDARECVR in a TCP/IP environment or APIA<nnnnnn> with OptiConnect), and for the QXDAEDRS SQL server job.

21.4.8.2 Selecting the right trace level

The right trace level depends on the kind of problem to be analyzed. A higher trace level adds more overhead to the application, and the trace buffers may wrap too soon to catch the interesting data if the level is too high.

The following overview is intended to help you make the right decision based on the problem description. A higher trace level includes all the data of the lower levels. Each trace entry has a time stamp associated with it. The following list of traced information at each level may help to select the right trace level based on the problem to catch.

Trace level 0 (NONE)

No information will be traced.

Trace level 1 (ERROR)

At this trace level, information is only traced if an error is assumed.

- Whenever an XDA API returns an error in the error code structure to the caller, this message is also sent to the job log.
- If the API QxdaConnectEDRS is failing, the input parameters are traced.
- If an SQL statement execution via QxdaProcessImmediateEDRS is failing, some basic information is traced.
- If an SQL operation via QxdaProcessExtDynEDRS is failing, information about the call to QSQPRCED is traced. However, because the SAP R/3 application is tolerating a certain set of SQL codes (they usually do not indicate a real problem), there will be no information in case of those SQL codes. The SQL codes to be ignored are 100 (record not found); -204 (object not found); -514, -516, and -518 (prepared statement not found); -601 (object already exists); and -803 (duplicate key).
- If QxdaCallProgramEDRS receives an error, it traces the program name and library and the program parameters as they are after the failing call.
- If any internal heap allocation fails, information about the heap amount and the job's activation groups are traced.
- If the internal Alarm Handler was called and forced a rollback, a trace entry is written.

- If the QXDAEDRS SQL job receives an error code other than 3420 (address already in use), the error code and some information about the originating function will be traced.

Trace level 2 (INFO)

At this level, most XDA APIs write trace data, even if no error occurred during execution. This may be used to track down some history that may have lead to an error.

- API QxdaCallProgramEDRS is tracing the program name and library, as well as the number of parameters in any case (not only in case of an error).
- API QxdaConnectEDRS traces the input format, server name, connection type, commitment control information, and the SQLDA cache size.
- API QxdaCommitEDRS traces the commit option it was called with (0 = COMMIT WORK, 1 = COMMIT HOLD).
- API QxdaDisconnectEDRS traces whether the disconnect was requested with the commit or the rollback option.
- API QxdaFindEDRSJob is tracing how many jobs were found and how much information was returned to the caller, based on the provided space.
- API QxdaGetDBTime traces the DB time it was received from the database server.
- API QxdaProcessCommandEDRS traces up to 30 characters of the command string if no error happened during the command execution. If an error happened, the full command string and the error message are shown.
- API QxdaProcessExtDynEDRS traces the following information for each executed statement:
 - SQL package
 - Statement name
 - Cursor name
 - Function (for the QSQPRCED API)
 - Basic information about the SQLDA

If an error happened (SQL code not 0 or 100), the format for the QSQPRCED API, pointer information, the host variables (sqlvar structure in the SQLDA), the SQL code, and the additional error information are shown. If a “real” error happened (not one of the codes listed above under trace level 1), the complete SQLP0300 or SQLP0400 format is dumped, and information about the cached SQLDA will be given.

- API QxdaProcessImmediateEDRS traces up to 30 characters of the statement text if no error happened (SQL code is 0 or 100). If an error happened, the full statement text is traced.
- API QxdaRollbackEDRS traces the rollback option it was called with (0 = ROLLBACK WORK, 1 = ROLLBACK HOLD).
- Signal handler calls are traced with information whether the job was at a commit boundary.
- The QXDAEDRS SQL job traces information about incoming Connect requests. If the job to initiate the connection is running with the same or a higher level of the interface, the qualified job name is written into the trace file. Otherwise, only the user ID of the caller is traced. It is also traced whether the request was a TCP/IP (T) or UNIX Domain (U) socket.

Trace level 3 (VERBOSE)

This is the most detailed level of information, so the trace buffer may fill (and wrap) quickly. While the other levels may help finding problems that are not in the XDA code itself but in the caller's code or in the underlying DB code, this level is primarily designed to catch problems that happen within the XDA component.

- API QxdaCallProgramEDRS additionally traces the program parameters before the actual program call.
- API QxdaConnectEDRS additionally traces the job associated user data and the job suspension user data.
- API QxdaFindEDRSJob additionally traces the job user data value.
- API QxdaProcessCommandEDRS always traces the full command.
- API QxdaProcessExtDynEDRS always traces the format for the QSQPRCED API, pointer information, the host variables (sqlvar structure in the SQLDA), and the SQL code. In case of any error (SQL code not 0 or 100), the SQLP0300 or SQLP0400 format is dumped.
- API QxdaProcessImmediateEDRS always traces the full statement text.
- All heap operations (allocate, reallocate and free) are traced.

21.4.8.3 Dumping the trace data

The trace data is collected in user space objects (type *USRSPC) in library QUSRSYS with the name QP0Z<nnnnnn>, with <nnnnnn> as the job number of the job that generated the trace data. The text description of these user spaces shows the full qualified job name. The data in these spaces can be formatted and dumped with the DMPUSRTRC JOB(<qualified job name>) command. This command creates a member named QP0Z<nnnnnn> (with <nnnnnn> as the job number of the job that was traced) in file QTEMP/QAP0ZDMP of the job that issues the DMPUSRTRC command. This member or file should be copied to some other library to make sure it is not lost after signing off.

21.4.8.4 Cleaning up trace data

The user space objects QUSRSYS/QP0Z<nnnnnn> are not deleted automatically. For each new job to be traced, a new space is created (the initial size is 300 KB). We recommend you turn off the trace feature when it is no longer needed. You should also delete the old user space objects with the DLTUSRSPC command. Such a user space should not be deleted if the job it was used for is still active.

21.5 Problem analysis

This section describes what needs to be considered when a problem occurs.

21.5.1 Where to look first

Perform the following checks in the order given:

1. **PTFs:** IBM is maintaining Informational APARs to keep track of the PTFs that were found to be useful for running R/3 on iSeries server. Please make sure

that you have applied all the PTFs listed in the corresponding Informational APAR. Table 45 lists the Informational APARs based on the OS/400 release.

Table 45. Informational APARs

OS/400 release	APAR number
V4R3M0	II11296
V4R4M0	II11832
V4R5M0	II12399
V5R1M0	II12833

2. **Kernel patch level:** Make sure you are at the most recent patch level for the R/3 kernel. To find out the installed patch level, use transaction SM51. Position the cursor on the line with your server, and click the **Release Information** button. Information about available kernel patches and installation can be found in OSS Note 49365.
3. **Short dumps:** Start to analyze the problem by displaying the short dump. If you do not have a short dump go to the next step. The short dump contains usually some information about the work process and the job on the iSeries server like the example shown in Figure 370.




Figure 370. Short dump

4. **System log:** The system log may contain additional information about the error. Use the time stamp from the short dump to specify the From data/time for the transaction SM21. The short dump shows the work process in any case and the job name for some types of errors.

5. **Developer trace:** Use the work process name from the short dump or the system log entry to check the corresponding developer trace using transaction ST11.
6. **Job log:** Use the job log to check the error messages that have been sent from work process to the job.

You should be able to identify the type of the problem by looking at the R/3 trace and log files, the job logs, and the system configuration. The following sections describe the data that can be reviewed or should be collected when you are going to report the problem.

21.5.2 Database error

In many cases, the end user receives an R/3 short dump for a database error condition. The short dump often gives an SQL error code in conjunction with the qualified job name of the work process in which the error happened, for example:

```
SQL error "-913" occurred when accessing table "SFLIGHT "  
Database error text: "Row or object SFLIGHT in R3R01DATA type  
*FILE in use. Job=016470/R0101/WP01"  
Internal call code .....: "RSQL/OPEN/SFLIGHT"
```

When debugging, it is important to find out if messages in the job log are related to the problem seen by the end user. If a short dump with an SQL error code is generated, you can search for the associated message ID. For example, if the short dump shows "SQL error "-913" occurred...", you can search for the term "SQL0913". First, you need to verify if the time stamp in the job log matches the time stamp of the short dump (a few seconds of tolerance are allowed) to make sure this is the same event. In any case, but especially in case of an SQL0901, check which messages were sent prior to this message. If no short dump is written or the short dump does not contain an SQL error code, you can look for escape-type messages around the time of the error.

Note

You should delete the SQL packages using the `DLTR3PKG` command to avoid database error conditions after:

- Kernel patch installation
- Cumulative PTF package installation
- Database fix package installation

There are some types of messages that generally can be found in the job logs and *do not indicate* an error condition (unless they appear in a short dump). Among these are:

- SQL0204: ... in ... type *SQLPKG not found. – for type *SQLPKG
- SQL0514: Prepared statement ... not found.
- CPC2206: Ownership of object ... in ... type *SQLPKG changed.
- SQL0904: Resource limit exceeded. – for resource type 7
- CPF5009: Duplicate record key in member ...
- CPF5034: Duplicate key on access path for based-on member of ...
- SQL0803: Duplicate key value specified.

The first four messages are caused and handled by R/3, while checking for SQL packages that exist and creating them if they do not. The last three messages are typically caused by the ABAP statement MODIFY. It first tries to insert a record into a file and, in a second attempt, performs an update if a record with the specified primary key already exists.

For database related problems, you should collect the following information:

- R/3 system log
- R/3 short dump
- Developer trace of the R/3 work process
- Job log of the corresponding job (Refer to 21.2, “Working with job logs” on page 505)
- R/3 SQL trace (ST05) to identify the failing SQL statement if the problem is reproducible

21.5.3 Performance problems

For all kinds of performance problems, you should ask yourself the question: “What has changed since the performance is poor?” Think about recent changes in the system configuration, network, increased number of users, kernel patch, support packages, PTFs, modifications, SAP release, or OS/400 release.

If you cannot solve the performance problem by means of this analysis, contact SAP and describe the problem as precisely as possible.

21.5.3.1 System-wide performance

The general performance of R/3 may be poor. The following list gives clues of where to look for bottlenecks in the system-wide performance.

For OS/400, check:

- **Sizing:** Make sure the sizing of the iSeries server is still correct.
- **Disk capacity:** Check what percent of the system ASP is used.
- **CPU:** Verify whether the CPU is constantly busy.
- **Work management:** Use the `WRKSYSSTS` command to check the size of the storage pools, the activity levels, and the page faults.
- **User ASP:** An overflowed journal user ASP causes performance problems.
- **Disk balancing:** Use the command `WRKDSKSTS` to make sure that all disk are equably busy. If one disk sticks out, use the `STRDSKRGZ` command (introduced in V4R4M0; use the `TRCASPBAL` command on previous releases) to balance the usage and the capacity of the disks in an ASP.
- **Interactive load:** The interactive load can be too high on a server model. Use the `WRKSYSACT` command from the Performance Tools (5769-PT1) licensed program, if installed, to check if the `CFINT<nn>` tasks consumes most of CPU time.

For R/3, check:

- **Trace level:** The trace level for the developer traces should be set to 1.

- **DBMON:** Use transaction ST04 to check whether the database monitor is running. Note that even the new memory-based database monitor needs some system resources.
- **Buffering:** Check the buffering quality with transaction ST02.
- **Workload:** Transaction ST03 shows some information about today's workload.

21.5.3.2 Transaction-based

The performance of a certain transaction may be poor. If the problem is reproducible, use the following steps to collect all data that is necessary for a detailed analysis:

1. From the R/3 session where you want to run poor performing transaction, complete these steps:
 - a. Go to transaction SE38, enter `RSTRC000` for the report, and click the **Execute** button.
 - b. Place an "X" in the Keep work process box, and click the **Change** button. You now see a message that says the work process is locked.
2. From another R/3 session, using transaction SM50, you should see the work process with a status of "Stopped" and a reason of "Lock" for your first session. You need the PID for the session for the next step.

3. From OS/400, run the command:

```
WRKPID PID(<nnn>)
```

Here `<nnn>` is the PID for the stopped work process. This gives you the associated job for the work process that needs to be debugged.

4. Change the job attributes of the work process job with the command:

```
CHGJOB JOB(<jobname>) LOG(4 00 *SECLVL) LOGCLPGM(*YES)
```

Here, `<jobname>` is the qualified job name you received from the WRKPID command.

5. Start the remote service operation with the command:

```
STRSRVJOB JOB(<jobname>)
```

6. Put the job into debug mode using the command:

```
STRDBG UPDPDPROD(*YES)
```

7. From the second R/3 session, start the SQL trace for the user using transaction ST05.

8. Run the poor performing transaction. It may run even slower than before because you have two levels of tracing going, so please be patient.

9. Once the transaction completes, end the SQL trace with transaction ST05, end the debug mode with the `ENDDBG` command, and end the remote service operation with the `ENDSRVJOB` command.

10. Do not forget to unlock the work process using report RSTRC000.

11. Copy the job log to a spool file with the command:

```
DSPJOBLOG JOB(<jobname>) OUTPUT(*PRINT)
```

Here `<jobname>` is the qualified job name from step 3.

21.5.4 IFS problems

SAP R/3 uses the IFS and especially the /QFileSvr.400 file system. If problems occur, it should be checked if a directory for each system in the R/3 landscape exists under /QFileSvr.400. In addition, verify the following command has been run:

```
STRHOSTSVR SERVER(*FILE)
```

Make sure port 8473 is in state "Listen". It should also be verified that the instance user profile is enabled (SAP<nn> up to SAP Release 4.0B, <SID><nn> from SAP Release 4.5B on), does not have an expired password, and has the same password on all systems.

21.5.5 Printing problems

Refer to 9.12, "Resolution tips for printing problems" on page 212, in the R/3 system printing section for more information.

21.5.6 OptiMover/400 or OptiConnect problems

It is required that the QSOC subsystem runs on the application server machines and the database server machine when using OptiMover/400. If it is not running on all of the machines in a three-tier network, connections between the multiple machines cannot be established. To check if the subsystem is up, run the command:

```
WRKACTJOB SBS(QSOC)
```

If the subsystem is not up, run the command:

```
STRSBS SBSD(QSOC/QSOC)
```

The same job description and user profile must be used by both the application server and the database server to which it is connected. The user profile is named SAP<nn> or <SID><nn> and the job description is R3_<nn>, where <nn> is the instance number. Moreover, the job description must be in library R3<SID>400. The user profile must have the same password on both systems.

To verify that the user profile is present, use the command:

```
WRKUSRPRF
```

To verify that the job description exists, use the command:

```
WRKJOB D R3<SID>400/R3_<nn>
```

In addition, you may want to perform these steps:

1. For three-tier systems, you need to change the default profile '/usr/sap/<SID>/SYS/profile/DEFAULT.PFL' using the `EDTF` command or transaction RZ10. Change the `rdisp/bufrefmode` parameter from `sendoff`, `exeauto` to `sendon`, `exeauto`.
2. For three-tier systems to use OptiConnect, set the `db4/db4/opticonnect=1` profile parameter, in the '/usr/sap/<SID>/SYS/profile/<SID>_<instance><host>' instance profile using the `EDTF` command or transaction RZ10.
3. Change the passwords for the user profiles.

21.5.7 Unable to start R/3

These problems are mainly due to errors made in the setup and startup of R/3, not errors in R/3.

21.5.7.1 TCP/IP and server support

Ensure that TCP/IP support is started before the QSERVER subsystem. To make sure this is always the case, you can add an autostart job to the QSERVER subsystem. Check if TCP/IP support is active using the command:

```
WRKACTJOB SBS(QSYSWRK)
```

Then, examine the jobs running under that subsystem. The job name for the required TCP/IP jobs is QTCPIP. Also, you can ping the host using the IP address. If this job is not running, you must start TCP/IP using the `STRTCP` command.

You could also use the iSeries server `NETSTAT` command to review the status of TCP/IP network.

Before attempting to run the `STARTSAP` command, run the following command to ensure that QSERVER and QPWFSERVD are running:

```
WRKACTJOB SBS(QSERVER)
```

If not, end the QSERVER subsystem and TCP/IP support. Restart TCP/IP followed by the QSERVER subsystem using the `STRHOSTSVR` command.

After `STARTSAP` is started correctly, the jobs QSERVER, QPWFSERVD, and QPWFSERVS0 should be running in the QSERVER subsystem.

21.5.7.2 TCP/IP host name table

If the host name table does not contain the correct name for the database host and message server host, the `STARTSAP` command fails. The host name that R/3 expects is in the default profile file `'/usr/sap/<SID>/SYS/profile/DEFAULT.PFL'`. View this file to see what the names need to be. The database host is specified by the `sapdbhost` profile parameter. The message server host is specified by the `rdisp/mshost` parameter.

To ensure that the host name table is correct, ping using the host names specified in the R/3 default profile.

To check or correct the host name table entry, enter the `CFGTCP` command and choose option 10 to work with the host table entries. Also use option 13 to ensure that the Searched first parameter specifies `*LOCAL`. This avoids any errors in host table specifications in the remote name server and improves the performance.

The system name in the default profile `DEFAULT.PFL` is case sensitive and must match the host table entry.

21.5.7.3 Remote file system authority

In a three-tier environment, the application server uses the remote file system to read or write IFS stream files to the database server. It is required that the iSeries server user profile be replicated on the database server. The copy on the database server must be identical to the original, including the password. The remote file system determines the user profile and password for the job under

which the read or write operation is being performed. It attempts to run that operation on the database server using that same user profile and password. If the profile does not exist on the database server or the password is different, the operation fails.

Therefore, the iSeries server user profile for an application instance must exist on the database server. The user profile for an instance has the name SAP<nn>, where <nn> is the instance number. If, for example, instance 00 exists on the database server and instance 01 exists on an application server, the application server must have user profile SAP01. The database server must have user profile SAP00 and user profile SAP01. If you are using the standard installation procedure, this is done automatically. If you are not using the standard installation procedure, you need to manually create the iSeries server user profiles from the application server to the database server. You can use the CRTSAPUSR command to do this.

Note

Since R/3 Release 4.5A, SAPnn is replaced with SIDnn.

21.5.7.4 MTXW

If you encounter a situation where many work processes show a status of MTXW after issuing the STARTSAP command, and they remain for a long period of time, use the WRKACTJOB command, option 5, and then option 19 to find out what mutex is causing the wait situation. This is valuable information for SAP to find out the root cause of the problem. You could try terminating the subsystem and restarting R/3 once to see if this situation still stays the same.

21.6 Reporting the problem

We want you to provide the following information if you want to report the problem to either SAP or IBM. By doing so, you help us to solve your problem faster. Table 46 shows the information that needs to be collected in any case, regardless of the type of problem you've encountered.

Table 46. General information that needs to be collected

Type of information		How to find it	Example
R/3	R/3 database release	System -> Status	4.6C
	R/3 kernel release	SM51 -> Release info	4.6D
	R/3 kernel patch level	SM51 -> Release info	280
	SID, instance and client	-	P01, 00, 300
	R/3 landscape	DSPPR3SYS <SID>	-

Type of information		How to find it	Example
OS/400	System name	DSPNETA	AS23
	System model number	DSPSYSVAL QMODEL	840
	System serial number	DSPSYSVAL QSRLNBR	100ABCD
	Cumulative PTF package level	DSPPTF	C0294450
	Database fix package level	DSPDTAARA SF9910<r>	SF99105-03

21.6.1 R/3 environment

You must always provide the following information about the R/3 system where the error occurred:

- R/3 database release
- R/3 kernel release
- R/3 kernel patch level
- SID, instance, and client
- R/3 landscape

To determine the kernel patch level, follow these instructions:

1. Call transaction SM51, and select the server name.
2. Click the **Release info** button, and note the Support package level number under section SAP R/3 Kernel information (for example, 280).

If SAP GUI is not available, use the command:

```
DSPDTAARA DTAARA (R3<REL>OPT/KERNEL)
```

21.6.2 OS/400 environment

You must always provide the following information about your OS/400 environment:

- **System name:** The system name of the iSeries server identifies the right system.
- **System model number:** Every system has its system model number.
- **System serial number:** IBM needs the system serial number in order to open a PMR.
- **Cumulative PTF package level:** This PTF package contain a collection of PTFs for various licensed program products.

Use the `DSPPTF` command to find the cumulative PTF package level of your system. Note the first PTF ID, which is at least temporarily applied, like TL00294. You can ignore the first three characters. The fourth character indicates the year (like 0 for 2000 or 9 for 1999). The next three characters stands for the day in the year. This is called Julian date format. We recommend you use the format Cydddvrn (C=cum PTF package, y=year, ddd=day in year, v=version, r=release, m=modification). In our example, this would be C0294450.

- **Database fix package level:** The database fix package contains PTF that fixes known database problem.

Use the command:

```
DSPDTAARA DTAARA (SF9910<r>)
```

Here, *r* is the OS/400 release. Note the level of the database fix package, such as SF99105-03.

21.6.3 Saving spooled files

To save all problem-related spooled files (such as job logs, etc.) to a physical file and to save the physical file to a save file, enter the following commands in the order presented:

1. CRTLIB LIB(<library>) TYPE(*TEST)
2. CRTPF FILE(<library>/<file>) RCDLEN(132) MBR(*NONE) MAXMBRS(*NOMAX) SIZE(*NOMAX)
3. CPYSPLF FILE(<spooled file name>) TOFILE(<library>/<file>) JOB(<qualified job name>) SPLNBR(<spooled file number>) TOMBR(<spooled file name>)
4. CRTSAVF FILE(<library>/<save file>)
5. SAVOBJ OBJ(<file>) LIB(<library>) DEV(*SAVF) OBJTYPE(*FILE) SAVF(<library>/<save file>) DTACPR(*YES)

21.6.4 Reporting the problem to SAP

Report the problem to SAP using SAPNet - R/3 Frontend. You can send the information you've already collected to sapservX (/incoming directory).

21.7 Additional information

You can find additional information in:

- **IBM Informational APARs**

- II12632: Terminology – A short overview about terms used with R/3 Implementation – Job and object names used in R/3
- II12633: General debug information – How to find job logs etc
- II12634: Typical errors – Examples how to handle typical cases
- II12635: Documentation collection

- **IBM Web sites**

- General SAP R/3 on the iSeries server:
<http://www.as400.ibm.com/service/bms/support.htm>
- AS/400 SAP Forum: http://as400service.ibm.com/s_dir/SAPDiscuss.nsf
- Support Line Knowledge Base:
<http://as400service.ibm.com/supporthome.nsf/document/10000051>

- **SAP Notes**

- 205699: General recommendations for problems
- 101113: Analysis of performance problems
- 142023: High temporary storage utilization
- 36578: Problems accessing stream files
- 37987: Importing transports
- 63993: Apply R/3 Fix (APYR3FIX) tips
- 107116: Info during termination of work process
- 121625: Buffer sizes

Part 4. Appendices

This part contains appendixes and complementary information to the chapters:

- Appendix A, “OS/400 user tools” on page 527, covers some very useful tools that allow you to do such things as editing and displaying stream files, maintain optimal disk performance, and use stream files to store compressed savefiles of objects.
- Appendix B, “Apply Journalized Changes Extended” on page 545, explains the Apply Journalized Changes Extended PRPQ (product number 5799-AJC), which provides an extended version of the Apply Journalized Changes (APYJRNCHG) command for OS/400. It provides the ability to replay object-level changes (for example, create a file or change a file) based on a restored library and journal entries that were deposited during the original operation.
- Appendix C, “Support for SAP R/3” on page 553, outlines the support available for SAP R/3 on the iSeries server.

Appendix A. OS/400 user tools

IBM and SAP provide some very useful tools. The method of obtaining these tools varies. Starting with V4R4, most of the tools are available through OS/400. Others are available in the R/3 kernel library, via SAPSERVx, or via PTFs. Some of the commands are listed in the following tables along with their location, owner, and release in which they were incorporated.

Table 47. New OS/400 commands

Command name	Location	Owner	SW release
Edit File (EDFT)	Library QSYS	IBM	OS/400 V4R4
Display File (DSPF)	Library QSYS	IBM	OS/400 V4R4
Start ASP Balance (STRASPBAL)	Library QSYS	IBM	OS/400 V4R4
End ASP Balance (ENDASPBAL)	Library QSYS	IBM	OS/400 V4R4
Trace ASP Balance (TRCASPBAL)	Library QSYS	IBM	OS/400 V4R4
Start Disk Reorganization (STRDSKRGZ)	Library QSYS	IBM	OS/400 V4R4
End Disk Reorganization (ENDDSKRGZ)	Library QSYS	IBM	OS/400 V4R4
Save and delete journal receivers (SAVDLTRCV)	FTP server SAPSERVx	SAP	
Stop save and delete journal receiver (SAVDLTRCVE)	FTP server SAPSERVx	SAP	
Create SAP User (CRTSAPUSR)	Library R3<SID>OPT ¹	SAP	R/3 3.0F
¹ <SID> will be often used in this document, it stands for a particular SAP system ID			

Table 48. Modified OS/400 commands

Command name	Location	Owner	SW release
Copy to Stream File (CPYTOSTMF)	Library QSYS	IBM	OS/400 V4R4
Copy from Stream File (CPYFRMSTMF)	Library QSYS	IBM	OS/400 V4R4
Start SAP system (STARTSAP)	Library R3<SID>OPT	SAP	R/3 4.5B
Stop SAP system (STOPSAP)	Library R/3<SID>OPT	SAP	R/3 4.5B
Save the R/3 System	Library R/3<SID>OPT	SAP	R/3 4.5B

Table 49. Unchanged OS/400 commands

Command	Location	Owner	SW release
Work with Object Links(WRKLNKSAP)	R3<SID>OPT	SAP	
Scan Directory (SCANDIR)	R3<SID>OPT	SAP	
Remove Directory Recursively (RRM)	R3<SID>OPT	SAP	
Copy Stream File (CPYSTMF)	R3<SID>OPT	SAP	
Display Temporary Storage Size (DSPTMGSTG)	R3<SID>OPT	SAP	
Show SQLPKG Information (LSTPKGINF)	R3<SID>OPT	SAP	
Work with Job by PID (WRKPID)	R3<SID>OPT	SAP	
Start Report in R/3 Batch (STRREPORT)	R3<SID>OPT	SAP	
Run R/3 Command (RUNR3CMD)	R3<SID>OPT	SAP	
Run Distributed R/3 Command (RUNDR3CMD)	R3<SID>OPT	SAP	
SQL Utility (SQLUTIL)	QGPTOOLS PTF ¹	IBM	
Notes: 1. Information regarding how to download and apply the QGPTOOLS PTF is listed in the Informational APAR for your corresponding OS/400 release. They are V4R3 - II11296, V4R4 - II11832, V4R5 - II12399, and V5R1 - II12844.			

For detailed information on these commands, please refer to:

- *AS/400 CL Reference V4R4 - Part 4*, SC41-5722
- *OS/400 Backup and Recovery V4R4*, SC41-5304
- SAP Internet site at: <http://service.sap.com>

A.1 Edit File (EDTF)

The EDTF command allows you to edit an integrated file system (IFS) stream file or a database physical file member. This command was previously available as a PTF in a user tool library QGPTOOLS.

Figure 371 shows the EDTF command with parameters for editing a stream file.

```

                                Edit File (EDTF)

Type choices, press Enter.

Stream file, or . . . . . > '/usr/sap/r01/dvebmgs01/log/alalerts'

Data base file . . . . . Name
Library . . . . . *LIBL Name, *LIBL, *CURLIB

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

```

Figure 371. EDTF: Edit stream file parameters

After you press Enter on this screen, the Edit File screen appears from which you can edit your file (Figure 372).

```

Edit File: /usr/sap/r01/dvebmgs01/log/alalerts
Record . : 1 of 9441 by 10 Column: 1 of 66 by 126
Control :

CMD
.....1.....2.....3.....4.....5.....6.....7.....8..
..+...0.....1.....2.....
*****Beginning of data*****
#MONITORING_SEGMENT
# ALSYSID="R01"
# SEGMENT_NAME="SAP_CCMS_ASM23_R01_01"
# STARTED_AT="Wed Oct 18 11:01:13 2000"
# SEGMENT_TYPE=APPLSERVER
# OWNER="SAP_CCMS"
# LIBRARY_VERSION=20000320
# SEGMENT_VERSION=20000320
# UPLOAD_STARTED_AT="Wed Oct 18 11:01:42 2000" by AlPid 3
# DOWNLOAD_STARTED_AT="Wed Oct 18 11:11:42 2000" by AlPid 0
# NEXT_DOWNLOAD_TREE_AT="Wed Oct 18 11:41:42 2000"
#NEXT_DOWNLOAD_ALERTS_AT="Wed Oct 18 11:41:42 2000"
# SAP_DEFAULT_VERSION="Fri Mar 24 10:15:03 2000"
#

OLD_ALERT
NAME="\###\MoniInfra_ASM23_R01_01\Sapmssy8\Sapmssy8_Runtime"

F2=Save F3=Save/Exit F12=Exit F15=Services F16=Repeat find
F17=Repeat change F19=Left F20=Right
(C) COPYRIGHT IBM CORP. 1980, 2000.

```

Figure 372. EDTF: Edit stream file screen

Beginning in V4R5, the EDTF option is also available from the WRKLNK screen. You can produce the same results shown in Figure 372 by selecting option 5 next to the desired file (Figure 373).

```

Work with Object Links

Directory . . . . : /usr/sap/R01/DVEBMGS01/log

Type options, press Enter.
  2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type      Attribute  Text
-----
2    ALALERTS          STMF
    ALMITREE           STMF
    ALPERFHI           STMF
    ENQBCK             STMF
    SLOG01             STMF

Parameters or command                                     Bottom
====>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field      F23=More options

```

Figure 373. WRKLNK: EDTF option

A.2 Display File (DSPF)

The DSPF command allows you to display the contents of a stream file or a database file member. This command functionally combines the Display Stream File (DSPSTMF) command, formerly available in QGPTOOLS, and the Display Physical File Member (DSPPFM) command.

Figure 374 shows the DSPF command with the parameters for displaying a stream file.


```

Display File (DSPF)

Type choices, press Enter.

Stream file, or . . . . . > '/usr/sap/r01/dvebmgs01/log/alalerts'

Data base file . . . . .
Library . . . . . *LIBL      Name
                                Name, *LIBL, *CURLIB

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

```

Figure 374. DSPF: Display stream file parameters

After you press Enter on this screen, the Display File screen appears from which you can display your file (Figure 375).

```

Browse : /usr/sap/r01/dvebmgs01/log/alalerts
Record . : 1 of 9737 by 18      Column: 1 of 66 by 131
Control :

....+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8..
..+....0....+....1....+....2....+....3..
*****Beginning of data*****
#MONITORING_SEGMENT
#   ALSYSID="R01"
#   SEGMENT_NAME="SAP_CCMS_ASM23_R01_01"
#   STARTED_AT="Wed Oct 18 11:01:13 2000"
#   SEGMENT_TYPE=APPLSERVER
#   OWNER="SAP_CCMS"
#   LIBRARY_VERSION=20000320
#   SEGMENT_VERSION=20000320
#   UPLOAD_STARTED_AT="Wed Oct 18 11:01:42 2000" by AlPid 3
#   DOWNLOAD_STARTED_AT="Wed Oct 18 12:21:42 2000" by AlPid 0
#   NEXT_DOWNLOAD_TREE_AT="Wed Oct 18 12:51:42 2000"
#NEXT_DOWNLOAD_ALERTS_AT="Wed Oct 18 12:51:42 2000"
#   SAP_DEFAULT_VERSION="Fri Mar 24 10:15:03 2000"
#

OLD_ALERT
NAME="\###\MoniInfra_ASM23_R01_01\Sapmssy8\Sapmssy8_Runtime"
VALUE=YELLOW SEVERITY=20 UID=5161 TIME=971867212

F3=Exit F10=DisplayHex F12=Cancel F15=Services F16=Repeat find F19=Left F20=Right
(C) COPYRIGHT IBM CORP. 1980, 2000.

```

Figure 375. DSPF: Display stream file screen

Beginning in V4R5, the DSPF command is also available from the WRKLNK screen. You can produce the same results that are shown in Figure 375 by selecting option 5 next to the desired file (Figure 376).

```

Work with Object Links

Directory . . . . : /usr/sap/R01/DVEBMGS01/log

Type options, press Enter.
  2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type      Attribute  Text
-----
5    ALALERTS          STMF
    ALMITREE           STMF
    ALPERFHI           STMF
    ENQBCK             STMF
    SLOG01             STMF

Parameters or command
====>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field  F23=More options
  
```

Figure 376. WRKLNK: DSPF option

A.3 STRASPBAL, ENDASPBAL, and TRCASPBAL

Although not directly related to the SAP R/3 environment, these three commands prove very useful to any SAP system administrator for maintaining an optimal disk performance. The STRASPBAL command starts the redistribution of data on disks within an auxiliary storage pool (ASP). The ENDASPBAL command ends the redistribution of data. The TRCASPBAL command runs the trace that is sometimes a prerequisite to run the STRASPBAL command.

A.4 SAVDLTRCV and SAVDLTRCVE

The new Save and Delete Journal Receivers (SAVDLTRCV) and Stop Save and Delete Journal Receivers (SAVDLTRCVE) journal management commands are available through a download from the SAPSERVx FTP server. The SAVDLTRCV command is used to save detached journal receivers on the tape. You run the SAVDLTRCV command in a batch job. The SAVDLTRCVE command is used to end the SAVDLTRCV job.

The SAVDLTRCV command continuously monitors the status of the R/3 journal receivers and backs them up after they are detached from the journal. The iSeries server detaches the journal receiver when the receiver reaches a certain size. This is defined at the creation of the journal QSQJRN in the library R3<SID>DATA, with the *Manage receivers* and *Receiver size options* attributes. When the old receiver is detached, the system automatically creates a new receiver, attaches it to the journal, and sends a message to the message queue, which you need to create for this purpose. After the SAVDLTRCV command

receives a message that the journal receiver was detached, it saves the old receiver.

Follow these steps to use the SAVDLTRCV and SAVDLTRCVE commands:

1. Download the SAVDLTRCV save file from the SAPSERVx FTP server using the same procedure as for downloading SAP patches. The save file is located in the /general/R3server/patches/COMMON/OS400 directory.
2. Create the SAVDLTRCV and SAVDLTRCVE commands, which are stored in the save file. Please refer to SAP Note 82079 for information on how to create these two commands.
3. Use the Work with Journal Attributes (*WRKJRNA*) command to make sure that the R/3 journal receivers are managed by the iSeries server (see Figure 377).

Work with Journal Attributes (WRKJRNA)

Type choices, press Enter.

Journal	> QSQJRN	Name, *INTSYSJRN
Library	> R3R01DATA	Name, *LIBL, *CURLIB
Output	*	*, *PRINT

Bottom

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

Figure 377. *WRKJRNA*: Work with R/3 journal attributes parameters

In this example, the System ID is R01. Press Enter and check that the *Manage receivers* attribute is set to *SYSTEM (Figure 378).

Work with Journal Attributes			
Journal	: QSQJRN	Library	: R3R01DATA
Auxiliary storage		Journal type	
pool	: 1	Journal state	: *LOCAL
Message queue	: QSYSOPR	Receiver size options:	: *INACTIVE
Library	: *LIBL		
Manage receivers	: *SYSTEM		
Delete receivers	: *YES		
Text	: *BLANK		
Type options, press Enter.			
8=Display attributes			
Attached			
Option	Receiver	Library	
	QSQJRN0013	R3R01JRN	
			Bottom
F3=Exit	F5=Refresh	F12=Cancel	F13=Display journaled files
F14=Display journaled access paths	F24=More keys		

Figure 378. Work with R/3 Journal Attributes screen

If the attribute *Manage receivers* is *not* set to *SYSTEM, then change the journal using the Change Journal (CHGJRN) command as shown in Figure 379.

Change Journal (CHGJRN)			
Type choices, press Enter.			
Journal	> QSQJRN	Name	
Library	> R3R01DATA	Name, *LIBL, *CURLIB	
Journal receiver:			
Journal receiver	*SAME	Name, *SAME, *GEN	
Library		Name, *LIBL, *CURLIB	
Journal receiver		Name, *GEN	
Library		Name, *LIBL, *CURLIB	
Sequence option	*CONT	*RESET, *CONT	
Journal message queue	QSYSOPR	Name, *SAME	
Library	*LIBL	Name, *LIBL, *CURLIB	
Manage receivers	*SYSTEM	*SAME, *USER, *SYSTEM	
Delete receivers	*YES	*SAME, *NO, *YES	
Receiver size options	*NONE	*SAME, *NONE, *RMVINTENT...	
Journal state	*SAME	*SAME, *ACTIVE, *INACTIVE	
			More...
F3=Exit	F4=Prompt	F5=Refresh	F12=Cancel F13=How to use this display
F24=More keys			

Figure 379. CHGJRN: Changing the R/3 journal

4. Create the message queue for the SAVDLTRCV command (Figure 380).

Create Message Queue (CRTMSGQ)

Type choices, press Enter.

Message queue	> SAVDLTRCV	Name
Library	> R3R01400	Name, *CURLIB
Text 'description'	> 'MSGQ for SAVDLTRCV command'	

Bottom

F3=Exit	F4=Prompt	F5=Refresh	F10=Additional parameters	F12=Cancel
F13=How to use this display	F24=More keys			

Figure 380. Create Message Queue (CRTMSGQ) for the SAVDLTRCV command

5. Use the Grant Object Authority (GRTOBJAUT) command to grant user R3OWNER the authority to this message queue as shown in Figure 381.

Grant Object Authority (GRTOBJAUT)

Type choices, press Enter.

Object	> SAVDLTRCV	Name, generic*, *ALL
Library	> R3R01400	Name, *LIBL, *CURLIB, *ALL...
Object type	> *MSGQ	*ALL, *ALRTBL, *BNDDIR...
Users	> R3OWNER	Name, *PUBLIC
+ for more values		
Authority	> *ALL	*CHANGE, *ALL, *USE...
+ for more values		
Authorization list		Name, *NONE
Reference object		Name
Library	*LIBL	Name, *LIBL, *CURLIB
Reference object type	*OBJTYPE	*OBJTYPE, *ALRTBL, *BNDDIR...
Replace authority	*NO	*NO, *YES

Bottom

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

Figure 381. GRTOBJAUT: Granting authority for the SAVDLTRCV message queue

6. Use the Change Journal (CHGJRN) command to associate the new message queue with the R/3 journal as shown in Figure 382.

Change Journal (CHGJRN)

Type choices, press Enter.

Journal	> QSQJRN	Name
Library	> R3R0DATA	Name, *LIBL, *CURLIB
Journal receiver:		
Journal receiver	*SAME	Name, *SAME, *GEN
Library		Name, *LIBL, *CURLIB
Journal receiver		Name, *GEN
Library		Name, *LIBL, *CURLIB
Sequence option	*CONT	*RESET, *CONT
Journal message queue	> SAVDLTRCV	Name, *SAME
Library	> R3R01400	Name, *LIBL, *CURLIB
Manage receivers	*SYSTEM	*SAME, *USER, *SYSTEM
Delete receivers	*YES	*SAME, *NO, *YES
Receiver size options	*NONE	*SAME, *NONE, *RMVINTENT...
Journal state	*SAME	*SAME, *ACTIVE, *INACTIVE

More...

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

Figure 382. CHGJRN: Associating the message queue with the R/3 journal

- Use the Submit Job (SBMJOB) command to submit a batch job that will start the SAVDLTRCV command as shown in Figure 383. Make sure that you signed on with user ID <SID>OFR or <SID>OPR.

Submit Job (SBMJOB)

Type choices, press Enter.

Command to run > SAVDLTRCV MSGQ(R3R01400/SAVDLTRCV) DEV(TAP 01)

...

Job name	SAVDLTRCV	Name, *JOB
Job description	*USRPRF	Name, *USRPRF
Library		Name, *LIBL, *CURLIB
Job queue	*JOB	Name, *JOB
Library		Name, *LIBL, *CURLIB
Job priority (on JOBQ)	*JOB	1-9, *JOB
Output priority (on OUTQ)	*JOB	1-9, *JOB
Print device	*CURRENT	Name, *CURRENT, *USRPRF...

More...

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys

Figure 383. SBJJOB: Submitting the SAVDLTRCV command

After the job is submitted, the SAVDLTRCV job is active in the Message Waiting status as shown in Figure 384.

```

Work with Active Jobs
TSCSAP02
07/21/99 14:17:10
CPU %: 1.1 Elapsed time: 00:00:02 Active jobs: 164

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
QBATCH QSYS SBS .0 DEQW
SAVDLTRCV R01OFR BCH .0 CMD-SAVDLTRCV MSGW
QCMN QSYS SBS .0 DEQW
QCTL QSYS SBS .0 DEQW
QSYSSCD QPGMR BCH .0 PGM-QEZSCNEP EVIW
QINTER QSYS SBS .0 DEQW
QPADEV0026 H41OFR INT .9 CMD-WRKACTJOB RUN
QSERVER QSYS SBS .0 DEQW
QPWFSESVSD QUSER BCH .0 SELW

More...

Parameters or command
===>
F3=Exit F5=Refresh F7=Find F10=Restart statistics
F11=Display elapsed data F12=Cancel F23=More options F24=More keys

```

Figure 384. WRKACTJOB: Status of the SAVDLTRCV job

The SAVDLTRCV job should not be active when the SAVR3SYS command is run. You can stop the job by running the SAVDLTRCVE command before you run the SAVR3SYS command (Figure 385).

```

Stop Save and delete receivers (SAVDLTRCVE)

Type choices, press Enter.

Message queue . . . . . > SAVDLTRCV Name
Library . . . . . > R3R01400 Name, *LIBL, *CURLIB

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

```

Figure 385. Stop Save and delete receivers (SAVDLTRCVE)

After the SAVR3SYS command has completed, start the SAVDLTRCV job again.

A.5 CRTSAPUSR

The Create SAP User (CRTSAPUSR) SAP command has become available in R/3 Release 3.0F. It is of great importance especially for the Transport Management System (TMS).

Important

Always use the CRTSAPUSR command to manually create OS/400 user profiles <SID>OFR, <SID>OPR, and <SID><INST>. These user profiles will be used by the shared transport directory across iSeries servers or logical partitions. Creating these user profiles in any other way (by copying an existing profile, for example) may result in TMS failures.

Figure 386 shows the parameters of the CRTSAPUSR command that are set to create user profile R01OFR.

Create an SAP user profile (CRTSAPUSR)

Type choices, press Enter.

User to create	> *OFR	*OWNER, *OPR, *OFR...
SAP system id	R01	Character value
Instance number	10	00-97

Bottom

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

Figure 386. CRTSAPUSR command

A.6 CPYTOSTMF and CPYFRMSTMF

The Copy to Stream File (CPYTOSTMF) command was previously used to copy data from a physical file member into an IFS stream file. Now it is enhanced to copy data from a save file as well. Figure 387 shows the parameters of the enhanced CPYTOSTMF command when copying from a save file into a stream file.


```

Copy To Stream File (CPYTOSTMF)

Type choices, press Enter.

From file member or save file . > '/QSYS.LIB/QGPL.LIB/SAVSRC.FILE'

To stream file . . . . . > '/tmp/savsrc'

Stream file option . . . . . > *NONE          *NONE, *ADD, *REPLACE
Data conversion options . . . . *AUTO          *AUTO, *TBL, *NONE
Database file CCSID . . . . . *FILE          1-65533, *FILE
Stream file code page . . . . . *STMF         1-32767, *STMF, *PCASCII...

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

```

Figure 387. CPYTOSTMF: Copying from a save file into a stream file

The Copy from Stream File (CPYFRMSTMF) command was previously used to copy data from an IFS stream file into a physical file member. Now it is enhanced to copy data into a save file as well. Figure 388 shows the parameters of the enhanced CPYFRMSTMF command when copying into a save file.

```

Copy From Stream File (CPYFRMSTMF)

Type choices, press Enter.

From stream file . . . . . > '/tmp/savsrc1'

To file member or save file . . > '/QSYS.LIB/QGPL.LIB/SAVSRC1.FILE'

Member option . . . . . *NONE          *NONE, *ADD, *REPLACE
Data conversion options . . . . *AUTO          *AUTO, *TBL, *NONE
Stream file code page . . . . . *STMF         1-32767, *STMF, *PCASCII
Database file CCSID . . . . . *FILE          1-65533, *FILE
End of line characters . . . . . *ALL          *ALL, *CRLF, *LF, *CR...
Tab character expansion . . . . *YES          *YES, *NO

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

```

Figure 388. CPYFRMSTMF: Copying from a stream file into a save file

A.7 SAVR3SYS

The Save R/3 System (SAVR3SYS) command is available via sapservX, as described in SAP Note 86305. The SAVR3SYS command can save all the information required for an SAP system except for the R/3 kernel library. Figure 389 shows an example of the SAVR3SYS command that saves the R/3 database, IFS files, configuration, and SQL packages.

Save R/3 System (SAVR3SYS)

Type choices, press Enter.

System ID	> R01	Character value
Device	/QSYS.LIB/TAP01.DEVD	
Prompt save commands	*NO	*YES, *NO
R/3 status during save	*RUN	*DSCDB, *RUN, *END
IFS files to save:		
Path name	*SPTH	
Include or omit		*INCLUDE, *OMIT
+ for more values		
Save R/3 database	*YES	*YES, *NO
Save R/3 configuration	> *YES	*YES, *NO
Save R/3 SQL packages	> *YES	*YES, *NO
Save reference date	*ALL	Date, *ALL, *LASTSAVE
Save reference time	*ALL	Time, *ALL, *LASTSAVE
Expiration date	*PERM	Date, *PERM
Save active wait time	3600	Number, *NOMAX

More...

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

Figure 389. Save R/3 System (SAVR3SYS)

A.8 STARTSAP and STOPSAP

The STARTSAP command is enhanced to use the environment variables. The STOPSAP command is enhanced to use the environment variables and to end the R/3 subsystem.

The *SAP System ID* and *R/3 instance* parameters in the STARTSAP and STOPSAP commands now include the value *ENV. When the *ENV value is used, the STARTSAP (and STOPSAP) command will determine, from the environment variables, which SID and instance will be started (or stopped). Figure 390 shows an example of the STARTSAP command using the environment variables.

Start R/3 System (STARTSAP)

Type choices, press Enter.

SAP System ID	*ENV	Character value, *ENV, *DB
R/3 instance	*ENV	Character value, *ENV, *ALL
R/3 instance hostname	*LOCAL	Character value, *LOCAL, *ALL
Delete existing spool files . .	*YES	*YES, *NO, Y, N
Start profile name	*DFT	

Bottom

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

Figure 390. STARTSAP: Start R/3 System

During the STARTSAP and STOPSAP commands, all necessary environment variables have to be set properly. When you are using different users than SIDOPFR and SIDOPR, issue the following command first. To set up the environment variables, add the kernel-library and the database library to your library list automatically:

```
CALL R3<SID>400/R3INLPGM
```

The new *End subsystem* parameter has been added to the STOPSAP command to permit the ending of the OS/400 subsystem that runs the SAP R/3. Ending the subsystem releases the memory allocated to that subsystem. To end the OS/400 subsystem, set the *End subsystem* and *Wait for instance to end* parameters to *YES.

Note

Do not end the subsystem with this option when the SAP Operating System Collector (SAPOSCOL) is running in that subsystem. Always end the SAPOSCOL first by calling the program SAPOSCOL with parameter "-k".

Figure 391 shows an example of the STOPSAP command, with the *End subsystem* and the *Wait for instance to end* parameters set to *YES.

Stop R/3 System (STOPSAP)

Type choices, press Enter.

SAP system ID	*ENV	Character value, *ENV
R/3 instance	*ENV	Character value, *ENV, *ALL
R/3 instance hostname	*LOCAL	Character value, *LOCAL, *ALL
Wait for instance to end	*YES	*NO, *YES
Maximum wait time (seconds)	120	10-999999
End subsystem	*YES	*NO, *YES

Bottom

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

Figure 391. STOPSAP: Stop R/3 System with the End subsystem option

A.9 SQL Utility (SQLUTIL)

The SQLUTIL command is used to issue SQL commands interactively. This tool is useful if you do not have the SQL product installed on your iSeries server.

Parameter	Description
OUTPUT	Specify whether you want the output from the command displayed at the workstation, printed with the job's spooled output, or placed in a database file. The following special values may be specified: • * – Output is displayed. • *PRINT – Output is printed with the job's spooled output. • *OUTFILE – Output is directed to the database file specified on the OUTFILE parameter.
COMMIT	Specify the commitment isolation level for the level of work. The following special values may be specified: • *NONE – No commitment control • *CHG – Uncommitted read • *CS – Cursor stability • *ALL – Read stability • *RR – Repeatable read
NAMING	Specify the naming convention you want to use. The following special values may be specified: • *SYS – Use system naming convention. • *SQL – Use SQL naming convention.
OUTFILE	Specify the physical database file where the SQL command results are directed. If the output file already exists, the

command attempts to use it. Existing data in the file member is replaced depending on the OUTMBR parameter. If the output file does not exist, the command creates a physical database file (with the name specified on the OUTFILE parameter) in the designated library. A member is created for the file with the name specified in the OUTMBR parameter.

OUTMBR

Specifies the name of the database file member where the output of the command is directed. A second value specifies whether the new data replaces the existing data or is added to the end of the data already in the file member.

Appendix B. Apply Journalized Changes Extended

The Apply Journalized Changes Extended PRPQ (product number 5799-AJC) provides an extended version of the Apply Journalized Changes (APYJRNCHG) command for OS/400. It provides the ability to replay object-level changes (for example, CREATE FILE or CHANGE FILE), based on a restored library and journal entries that were deposited during the original operation.

This product's key benefit is to improve recoverability of an application running on the iSeries server. Prior to Version 5 Release 1 (V5R1), many object-level operations were not journaled. Now they are journaled, but the OS/400 APYJRNCHG command is not capable of applying object-level journal entries. The command supplied with this PRPQ, APYJRNCHGX, does apply the object-level operations.

This PRPQ is especially applicable in environments where object-level changes occur between database backups. If an application creates or alters tables (or otherwise makes object-level changes) during productive operations, then this PRPQ provides the ability to more fully recover the database in the event of a disaster.

If you are recovering the whole R/3 system (and not just an individual file) and are using the forward recovery technique, we recommend that you use the PRPQ. It is available free of charge from the standard IBM software order process.

Once you have the PRPQ, be sure to view the “readme” file and the online help.

B.1 Example of recovering a database with APYJRNCHGX PRPQ

Let's look at an example of how to recover a database with the Apply Journalized Changes Extended PRPQ. Here is the test scenario:

1. Perform a full offline save of the database (original save).
2. Start an SAP version upgrade using the AON switch (productive activity allowed during the initial upgrade phases).
3. Simulate productive workload during the upgrade.
4. Stop the upgrade at the MODPROF_TRANS phase.
5. Save the database (second save).
6. Restore the original save.
7. Apply journaled changes up to the time of the second save.
8. Compare the databases

Note: We ran the test on a 4-way S20 with 50 disk arms (193.5 GB total), 2G memory.

B.1.1 Original save

We saved the database to a save file. Of course, in a productive environment, we would save to tape. But the technique is pretty much the same. Note the time the database library was saved by using Display Save File (DSPSAVF) command. We can use this date/time to find the starting journal receiver.

DSPSAVF DDL45B

The result of this command is shown in Figure 392.

Display Saved Objects - Save File

Library saved . . . : R3DDLDATA

ASP : 1

Save file : DDL45B

Library : TOMDDL SAV

Records : 19390469

Save command . . . : SAVLIB

Save active : *NO

Save date/time . . . : 04/18/01 10:45:02

Release level . . . : V5R1M0

Data compressed . . : Yes

Objects displayed . : 6

Objects saved . . . : 23599

Access paths : 0

Type options, press Enter.

5=Display saved data base file members

Opt	Object	Type	Attribute	Owner	Size (K)	Data
	R3DDLDATA	*LIB	PROD	DDLOWNER	15168	YES
	QSQJRN	*JRN		DDLOWNER	2312	YES
	"/SAP0001"	*FILE	PF	DDLOWNER	160	YES
	"/SAP0002"	*FILE	PF	DDLOWNER	168	YES
	"/SAP0003"	*FILE	PF	DDLOWNER	160	YES
	"/SAP0004"	*FILE	PF	DDLOWNER	160	YES +

F3=Exit

F12=Cancel

Figure 392. Display Saved Objects - Save File screen on the original save

B.1.2 Second save

In a disaster scenario, we would not have this save. Rather, we would have the original save and whatever journal receivers were available. In our test scenario, we made a second save of the database after the SAP upgrade reached the MODPROF_TRANS phase.

Since we have a second save, we cannot use the *LASTSAVE feature of the APYJRNCHG command. So we have to determine the range of entries ourselves. The date/time of this “second save” is really the point in time to which we need to recover (Figure 393).

Display Saved Objects - Save File

Library saved . . . : R3DDLDATA
ASP : 1
Save file : DDLUPG
Library : TOMDDL SAV
Records : 27961073
Save command . . . : SAVLIB
Save active : *NO
Save date/time . . . : 04/20/01 17:39:18

Release level . . . : V5R1M0
Data compressed . . : Yes
Objects displayed . : 6
Objects saved . . . : 30085
Access paths : 0

Type options, press Enter.

5=Display saved data base file members

Opt	Object	Type	Attribute	Owner	Size (K)	Data
	R3DDLDATA	*LIB	PROD	DDLOWNER	19808	YES
	QSQJRN	*JRN		DDLOWNER	3084	YES
	"/SAP0001"	*FILE	PF	DDLOWNER	160	YES
	"/SAP0002"	*FILE	PF	DDLOWNER	168	YES
	"/SAP0003"	*FILE	PF	DDLOWNER	160	YES
	"/SAP0004"	*FILE	PF	DDLOWNER	160	YES +

F3=Exit
F12=Cancel

Figure 393. Display Saved Objects - Save File screen on the second save

B.2 Planning the recovery

It is important to understand the events that occurred on the system prior to the failure. We describe one scenario, but your particular circumstances may differ.

In our case, we restore the database to the ORIGINAL SAVE version and then apply journal entries up to the time of the SECOND SAVE.

B.2.1 Finding the starting receiver

If you have an uninterrupted chain of receivers, and you have not lost the QSQJRN *JRN object in the database library, you can let the system determine the starting receiver, by specifying *LASTSAVE. In some cases, this is not possible. Therefore, you have to determine the starting receiver by comparing the Save date/time to the Attach/Detach date and time.

Another consideration for explicitly specifying the starting receiver and sequence number is performance. If you have a large set of receivers containing many millions of journal entries, it may be faster to find the starting point yourself. The system scans backwards through the entire receiver range looking for the last saved entry. By examining attach times, you can generally get there faster. Perform the following command:

```
WRKJRNA R3DDLDATA/QSQJRN
```

Select option 8 to display the attributes (Figure 394).

```

Work with Receiver Directory

Journal . . . . . : QSQJRN      Library . . . . . : R3DDLDATA

Total size of receivers (in kilobytes) . . . . . : 14089184

Type options, press Enter.
  4=Delete   8=Display attributes


Opt Receiver  Library  Number Attach      Status      Save
   _8 QSQJRN0109 R3DDLJRN  00001 04/17/01 SAVED      04/24/01
      QSQJRN0110 R3DDLJRN  00002 04/19/01 SAVED      04/24/01
      QSQJRN0111 R3DDLJRN  00003 04/19/01 SAVED      04/24/01
      QSQJRN0112 R3DDLJRN  00004 04/19/01 SAVED      04/24/01
      QSQJRN0113 R3DDLJRN  00005 04/19/01 SAVED      04/24/01
      QSQJRN0114 R3DDLJRN  00006 04/19/01 SAVED      04/24/01
                                           More...

Parameters or command
====>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F11=Display size
F12=Cancel

```

Figure 394. Work with Receiver Directory

The first receiver we have in our chain is QSQJRN0109. If we display the attributes using the command DSPJRNA, we can see the attach and detach date and time (Figure 395).

```

Display Journal Receiver Attributes

Receiver . . . . . : QSQJRN0109  Library . . . . . : R3DDLJRN

Journal . . . . . : QSQJRN      Library . . . . . : R3DDLDATA
Threshold (K) . . . . : 200000   Size (K) . . . . . : 205128
Attach date . . . . . : 04/17/01  Attach time . . . . : 15:15:28
Detach date . . . . . : 04/19/01  Detach time . . . . : 10:54:48
Save date . . . . . : 04/24/01   Save time . . . . . : 17:03:52
Text . . . . . : *BLANK

Auxiliary storage pool . . . . . : 1
Status . . . . . : SAVED
Number of entries . . . . . : 240663
Minimized fixed length . . . . . : NO
Receiver maximums option . . . . . : 0
Maximum entry specific data length . . . . . : 32689
Maximum null value indicators . . . . . : 12
First sequence number . . . . . : 1586624
Last sequence number . . . . . : 1827286
                                           More...

F3=Exit   F5=Refresh   F6=Display associated receivers
F10=Work with journal attributes   F12=Cancel

```

Figure 395. Display Journal Receiver Attributes

Since the save took place on 04/18/01 at 10:45:02, we see that QSQJRN0109 is the correct starting point, because it was attached before the save and detached after the save.

B.2.2 Finding the ending receiver

Now look at the last receiver we have in our range using the WRKJRNA command (Figure 396).

Work with Journal Attributes			
Journal	: QSQRN	Library	: R3DDLDATA
Auxiliary storage pool	: 1	Journal type	: *LOCAL
Message queue	: QSYSOPR	Journal state	: *ACTIVE
Library	: *LIBL	Receiver size options:	: *NONE
Manage receivers	: *SYSTEM	Minimize entry data :	: *NONE
Delete receivers	: *NO		
Text	: *BLANK		
Type options, press Enter.			
8=Display attributes			
Attached			
Option	Receiver	Library	
8	QSQRN0177	R3DDLJRN	
Bottom			
F3=Exit F5=Refresh F12=Cancel F13=Display journaled files			
F14=Display journaled access paths F24=More keys			

Figure 396. Work with Journal Attributes

Note the attach time is before the *second* save. This will be our ending journal receiver for the APYJRNCHGX operation. See Figure 397.

Display Journal Receiver Attributes			
Receiver	: QSQRN0177	Library	: R3DDLJRN
Journal	: QSQRN	Library	: R3DDLDATA
Threshold (K)	: 200000	Size (K)	: 151616
Attach date	: 04/20/01	Attach time	: 17:00:02
Detach date	: 00/00/00	Detach time	: 00:00:00
Save date	: 00/00/00	Save time	: 00:00:00
Text	: *BLANK		
Auxiliary storage pool : 1			
Status : ATTACHED			
Number of entries : 299116			
Minimized fixed length : NO			
Receiver maximums option : 0			
Maximum entry specific data length : 3039			
Maximum null value indicators : 12			
First sequence number : 36059640			
Last sequence number : 36358755			
More...			
F3=Exit F5=Refresh F6=Display associated receivers			
F10=Work with journal attributes F12=Cancel			

Figure 397. Display Journal Receiver Attributes

If the “Attach time” was later than the *second save*, we could look at the previous journal receivers until we found the correct receiver.

B.2.3 Finding the starting journal entry

We know that the starting journal receiver is QSQJRN0109. Even though we have to specify the explicit range of journal receivers, we could still let the system figure out the starting sequence number by using the FROMENT(*LASTSAVE) option. But let’s find the information ourselves and use an explicit sequence number. We can easily find this information by using the DSPJRN command as follows:

```
DSPJRN JRN (R3DDLDATA/QSQJRN) RCVRNG (R3DDLJRN/QSQJRN0109 R3DDLJRN/QSQJRN0109)
ENTTYP (MS)
```

The results of this command are shown here:

1749391	C	EC			DISP	10:18:11
1749392	C	EC			DISP	10:18:59
1749393	C	EC			DISP	10:20:23
1749394	F	MS	"/SAP0001"	R3DDLDATA	SAVDDL1	11:13:57
1749395	F	MS	"/SAP0002"	R3DDLDATA	SAVDDL1	11:15:01

From this information, we see that the entry with journal sequence number 1749394 is the first “Member Saved” (MS) entry.

Note: This example shows the “MS” entries related to SAVLIB. If you employ a save-while-active strategy, you need to search for the “SS” entries instead. Again, you could let the system find the starting point by using FROMENT(*LASTSAVE).

B.2.4 Finding the ending journal entry

We know that our ending journal receiver is QSQJRN0177. In our example, the ending point is the start of the “second save”. In a more typical recovery scenario, the ending point would simply be *LAST. It could also be the last restore or the point just before the database library was deleted prior to the restore. But the same technique may be used.

Note: If you explicitly deleted your database library prior to restoring a backup version, it is absolutely essential that you do not use TOENT(*LASTRST). This would have the unfortunate result of replaying all of the journaled changes, including the delete file entries!

We perform the following command:

```
DSPJRN JRN (R3DDLDATA/QSQJRN) RCVRNG (R3DDLJRN/QSQJRN0177 R3DDLJRN/QSQJRN0177)
ENTTYP (MS)
```

The result is shown here:

Sequence	Code	Type	Object	Library	Job	Time
36336424	C	EC			DISP	17:29:00
36336425	C	EC			DISP	17:32:21
36336426	F	MS	"/SAP0001"	R3DDLDATA	SAVDDLUPG	18:01:25
36336427	F	MS	"/SAP0002"	R3DDLDATA	SAVDDLUPG	18:01:29

From this information, you see that the entry with journal sequence number 36336426 is the first “Member Saved” (MS) entry from the “second save”.

B.3 Looking inside the journal

In a disaster recovery scenario, you would start the apply at this point. But let's dig a bit deeper.

We are going to scan the journal to see if there are any entries that will cause the apply to be interrupted. Since we will look in the journal multiple times for specific information, we decided that it would be faster to use the DSPJRN command and dump the output to an output file. This would allow us to perform queries against the output file.

We perform the following command:

```
DSPJRN JRN(R3DDLDATA/QSQJRN) RCVRNG(R3DDLJRN/QSQJRN0109 R3DDLJRN/QSQJRN0177)
DEPENT(*NONE) OUTPUT(*OUTFILE) OUTFILEMT(*TYPE4) OUTFILE(R3DDLJRN/JRNOUTF)
```

Now we create an index to speed up searches on the entry type field:

```
CREATE INDEX R3DDLJRN/joenttidx ON R3DDLJRN/JRNOUTF (JOENTT ASC)
```

Start SQL and run the following command:

```
SELECT * FROM r3ddljrn/jrnoutf
WHERE (JOENTT = 'AY') OR (JOENTT = 'EJ') OR (JOENTT = 'IV')
      OR (JOENTT = 'MF') OR (JOENTT = 'MR') OR (JOENTT = 'RC')
      OR (JOENTT = 'SA') OR (JOENTT = 'SR')
```

With V5R1 and the Apply Journalized Changes Extended PRPQ, any of the above journal entries would cause the apply processing to end.

Note: 'EJ' during ALTER TABLE does not cause processing to end. The Ignore on Apply field in the journal entry is set to 1. 'EJ', which is a result of the ENDJRNPF command, causes the apply to end.

If you want to check the journal to see if there are any object-level entries (and therefore require the use of the PRPQ), look for the following entry types:

```
SELECT * FROM r3ddljrn/jrnoutf
WHERE (JOENTT = 'AC') OR (JOENTT = 'CG')
      OR (JOENTT = 'CT') OR (JOENTT = 'DC') OR (JOENTT = 'DM')
      OR (JOENTT = 'DT') OR (JOENTT = 'FM') OR (JOENTT = 'FN')
      OR (JOENTT = 'GC') OR (JOENTT = 'GO') OR (JOENTT = 'GT')
      OR (JOENTT = 'MC') OR (JOENTT = 'MD') OR (JOENTT = 'MM')
      OR (JOENTT = 'MN') OR (JOENTT = 'RM') OR (JOENTT = 'RV')
      OR (JOENTT = 'TC') OR (JOENTT = 'TD') OR (JOENTT = 'TG')
```

B.3.1 Counting the record level changes

This is not really necessary, but if you want to see how many record level changes will be applied, start SQL and perform the following query:

```
SELECT JOENTT, COUNT(*) FROM r3ddljrn/jrnoutf
WHERE (JOENTT = 'BR') OR (JOENTT = 'DL')
      OR (JOENTT = 'DR') OR (JOENTT = 'IL') OR (JOENTT = 'PT')
      OR (JOENTT = 'PX') OR (JOENTT = 'UB') OR (JOENTT = 'UP')
      OR (JOENTT = 'UR')
group by JOENTT
```

TYPE	COUNT (*)
BR	55
DL	1,689,901
DR	2,057
PT	30,301,736

PX	566,017
UB	1,040,174
UP	1,040,174
UR	55

B.4 The APYJRNCHGX command

Now we put it all together and start applying the journaled changes using the APYJRNCHGX command:

```
>> APYJRNCHGX JRN (R3DDLDATA/QSQJRN)
      FILE ( (R3DDLDATA/*ALL) )
      RCVRNG (R3DDLJRN/QSQJRN0109 R3DDLJRN/QSQJRN0177)
      FROMENT (1749394)
      TOENT (36336426)
      CMTBDY (*YES)
```

Note that in our example, we understood the events that occurred and thoroughly checked the range of journal entries to apply. We really could have specified CMTBDY(*NO), which would eliminate one pass through the journal receivers (and save time). That's because the R/3 system was down and there were no updates happening to the database. If you are unsure about this, or if you are recovering to a point in time when the R/3 system was active, you should always use CMTBDY(*YES).

B.5 Performance hints

There are several ways to speed up the applying of journaled changes. First, watch out for interactive performance. If you are running in an interactive job, make sure you are running in a pool with a lot of memory. If you have a server model, you should run the APYJRNCHGX command in batch.

You should also be aware that there are potentially three passes through the journal. We discussed earlier how to skip the *LASTSAVE pass through the journal by figuring out the receiver range and starting sequence number. If you are planning to apply a large number of journal entries, this technique can definitely save time.

The second pass through the journal is for commit boundary processing. In most cases, this pass cannot be avoided. However, if you are sure the ending point is at a commit boundary (for example, if the R/3 system was ended at the point to which you will recover), then it is possible to save another big chunk of time by specifying CMTBDY(*NO). However, you have to be sure about this. It is absolutely essential to recover the database to a consistent point. If you are unsure, always select CMTBDY(*YES). The third pass actually applies the journal entries to the database.

Appendix C. Support for SAP R/3

This appendix outlines the support available for SAP R/3 on the iSeries server.

C.1 Marketing and technical support

When you need help with marketing situations and opportunities or if you need more technical assistance with SAP R/3 on iSeries, you may find this appendix useful. Contact your National SAP IBM Competence Center as your first-level support on SAP R/3 issues. There are “champions” to assist you in locations where there are no competence centers. In addition to this, the iSeries Division has established the iSeries Technology Center (iTC) based in Rochester, Minnesota, to assist you.

C.1.1 Defect support

All problems and defect reports should be channelled through the SAP support organization using SAPNet - R/3 Frontend (formerly known as OSS) or your local IBM Support Center. The normal process of defect management and escalation is followed by IBM and SAP. The iSeries account team may contact the iTC in special situations to seek assistance and information on problems reported by their customers.

C.2 Competence centers

There are competence centers in several countries worldwide. The purpose of these centers is to provide the sales and technical support needed for successful SAP R/3 implementations. Many of these centers are jointly staffed by IBM and SAP to achieve the required synergy.

The *National SAP IBM Competence Centers* provide the first-level support for the joint IBM-SAP sales force. If you need help in an SAP-related customer situation, you can contact these centers in your country.

C.2.1 European Competence Center

ERP Solutions Sales EMEA
c/o Holger Rasig
Altrottstrasse 31
D-69190 Walldorf,
Germany
Phone 49-6227-73-1045

C.3 North American Competence Center contact

IBM/ERP North American Competency Center
Corporate Center
1400 N. Providence Rd.
Building 1, Suite 400N
Media, PA 19063
Phone: +1-610-892-3009
Phone in USA: +1-800-IBM-0222
Fax: +1-610-892-3035

C.4 Latin America contacts

Wakiyama, Eduardo Hideki
ERP Solutions Sales Manager
Phone: +55-11-30503172
Notes ID: Eduardo H Wakiyama/Brazil/IBM@IBMBR
Internet ID: waki@br.ibm.com

Tosatto, Alessandro Gino Secioso
LA Midrange ERP Segment Manager
Telephone: +55-21-849-3053
Internet ID: atosatto@br.ibm.com

C.5 Asia Pacific Competence Center contacts

Japan
IBM Japan
c/o Toshi Sakayori
Fulfillment Competence Center Mgr
19-21 Nihonbashi Hakozaiki-cho
Chuo-ku, Tokyo
Japan
Phone: +81-3-3808-8982
Fax: +81-3-3664-4878

Korea
IBM Korea
Hanil Bldg
c/o Her, Jin Uk
Team leader, Production Ind Mktg Sol & SVC
25-11 Yoido, Yeongdeungpo-gu
Seoul Korea
Phone: +82-2-781-6306
Fax: +82-2-782-9146

Singapore
IBM Business Solutions Center
c/o Jalean Han
IBM World Trade Asia Corporation
80 Anson Road, IBM Towers
Singapore (0207)
Republic of Singapore
Phone: +65-320-1208
Fax: +65-227-2969

Australia
ISSC Australia
c/o Mark Engel
GPO Box 1798Q
Melbourne 3001
Australia
Phone: +61-3-626-6546
Fax: +61-3-626-6010

IBM Australia Limited
c/o Sally Broughton
GPO Box 3318
Sydney 2001
Australia
Phone: +61-2-353-3060
Fax: +61-2-353-3433

IBM Australia Limited
c/o Noel McClean
GPO Box 3318
Sydney 2001
Australia
Phone: +61-2-353-3132
Fax: +61-2-353-3433

C.6 Africa and Middle East Competence Center contacts

Saudi Arabia
Saudi Business Machines Ltd
c/o Dieter R. Moeller
General Marketing and Services Rep
for IBM SEMEA (GMSR)
Juffali HQ Bldg on Medina Road
P.O. Box 5648
Jeddah 21432
Saudi Arabia
Phone: +966-2-6600007
Fax: +966-2-6651163

South Africa
SAP SA (Pty) Ltd
c/o IBM SAP Competence Center
Dunkeld Crescent East,
Cnr Albury & Jan Smuts Ave.
DUNKELD WEST
Johannesburg
South Africa
2196
Phone: +27-11-880-6775
Fax: +27-11-880-6535

IBM South Africa (Pty) Ltd
c/o Vincent Oliver
Building IE11
IBM PARK
Private Bag X9907
SANDTON
South Africa
2146
Phone: +27-11-320-8368
Fax: +27-11-320-8578

C.7 IBM SAP International Competence Center (ISICC) support

The ISICC is the second-level support for the national competence centers. Among its many tasks are:

- To provide marketing support material (presentations, brochures, customer references)
- To provide technical support (benchmarks, sizing recommendations)
- To provide second-level regional support (Infoservice)
- To be the second-level customer briefing center
- To communicate news to the field
- To initiate/lobby for the development of solution packages

IBM SAP International
Competence Center (ISICC)
Altrottstraße 31
D-69190 Walldorf
Germany Phone: 49-6227-73-1001
Fax: 49-6227-73-1055

C.8 Regional support

There are two organizations providing international support: IBM and SAP. The decision of where to submit a request has to be based on the following guidelines:

- The ISICC Infoservice or the Philadelphia Support should be used for questions related to:
 - IBM hardware platforms in the SAP environment (PC server, pSeries server, and iSeries server)

- IBM software and add-ons in the SAP environment
- IBM pre-sales support (sizing, references, availability, and so on)
- The SAP regional support should be used by SAP customers (having an SAP customer number) when they have questions related to all problems and defects in an SAP installation. SAP wants to handle all requests centrally. If it proves to be a problem related to a specific system platform, SAP routes it to the respective partner company.

C.8.1 ISICC InfoService (Walldorf)

Since the International SAP IBM Competence Center (ISICC) was established in 1993 in Walldorf, the *Infoservice* has been one of the center's major support offerings for the SAP/IBM world.

The ISICC InfoService serves as a point of entry for all IBM SAP-related ERP pre-sales questions directed to the ISICC. Their main focus is on pre-sales issues of IBM products within the SAP world. They are also working to improve the flow of information between both companies. As a managed question and answer service the ISICC InfoService ensures that questions will be assigned to the most appropriate experts and will be answered as soon as possible despite restrictions in peoples' individual ability.

The ISICC InfoService offers support for:

- Second-level technical support
- Second-level marketing support
- Second-level sizing support/entry point for first level sizings
- Second-level SAP Industry Solution Sizings
- Briefings
- Access to "SAP Service Marketplace"

The ISICC InfoService can be contacted via by:

- E-mail: ISICC@DE.IBM.COM or Infoserve@de.ibm.com
- Phone: +49/6227-73-1099 (from 9 AM to 5 PM CET Monday through Friday)
- Fax: +49/6227-73-1052

It is understood that time is often critical. However, the request should be sent as early as possible so that the Infoservice can provide a quality answer. Requests received from customers are forwarded to the appropriate country organizations.

C.8.2 Philadelphia IBM SAP Competency Center outline

Within the U.S., a toll-free "800" number is managed in the competency center located in Philadelphia, PA (US) at 1-800-IBM-0222.

C.8.3 SAP regional support

The SAP regional support is organized as follows:

- Priority 1 messages: +49 (0) 180/534 34 36
- Non-technical inquiries: +49 (0) 180/534 34 34
- Problem messages: +49 (0) 180/534 34 31
- R/3 Fax: +49 (0) 180/534 34 30
- R/3 EarlyWatch*: +49 (0) 180/534 34 35
- R/3 Weekend stand-by: +49 (0) 180/534 34 36

- Remote consulting*: +49 (0) 180/534 34 37
- Network service (remote connection): +49 (0) 180/534 34 38

Note: The * (asterisk) indicates a chargeable service.

C.9 Information access

This section highlights the options available to you for having your questions answered.

C.9.1 SAP Service Marketplace

The SAP Service Marketplace is set up to provide an immediate source of information that is relevant to everyone working in the SAP environment. It covers news about SAP, news about competitors, and allows questions of a marketing nature to be asked and answered by participants.

To access SAP information through the Internet, go to: <http://service.sap.com>

Access to the SAP Service Marketplace is restricted only to SAP's customers and business partners, with valid SAPNet user ID and password. If your company does not have SAPNet access, you can apply for this through SAPNet Support at +49 (0) 180 534-3433.

The current requirement to receive an SAPNet user ID and password is a licensed R/3 installation at your site. The NET provides an easy to use front end for making comments and responses.

C.9.2 SAPNet - R/3 Frontend

SAP provides an online service called SAPNet - R/3 Frontend that each customer is required to access to report and monitor problems. SAPNet - R/3 Frontend also provides the developers with keys to enable them to define or change objects in the SAP system. This requires an established connection to SAP through X.25, ISDN, or frame relay. SAP must be notified about the relevant network data.

Please read SAP's *Remote Connection to the R/3 Online Services* document for details. This can be found in the R/3 system online help.

C.9.3 iSeries Informational APARs

An iSeries Informational Authorized Program Analysis Report (APAR) has been created for each SAP supported release that lists the current recommended PTFs and cumulative package level that SAP R/3 customers should have on their system. These Informational APARs are:

- V4R4 II11832
- V4R5 II12399
- V5R1 II12833

C.9.4 iSeries Technology Solutions Center

The iSeriesTechnology Solutions Center's ERP Team provides the following services for SAP R/3 customers on the iSeries platform.

C.9.4.1 Service offerings

The Service offerings include:

- **iSeries SAP R/3 Performance Evaluation Offering**

Validate the iSeries performance parameters prior to or after placing an iSeries SAP R/3 installation into production.

- **iSeries SAP R/3 Capacity Planning Offering**

IBM can assist in determining the expected iSeries hardware requirements. This will help you meet the information processing objectives based on current utilization and workloads on the system. This also serves as assessment of its additional requirement in the future resulting from upgrading the SAP R/3 software from the current version of increasing the size of a companies overall operations.

Additional information on these offerings can be viewed on the iTC's informational Web page at: <http://iws.as400.ibm.com/Service/bms/support.htm>
This page also contains valuable information on tools, performance tips, and fixes.

C.9.5 SAP Support Services

SAP provides the following services free of charge within the maintenance:

- **GoingLive Check**

The GoingLive Check is done when the system will shortly go productive. It checks if all prerequisites are met for going live with the hardware and parameter settings. Otherwise, recommendations are given to improve the performance.

- **EarlyWatch**

EarlyWatch service supports R/3 system operation by remote diagnosis for R/3 installations worldwide. It carries out detailed analysis of SAP applications and configurations in addition to database, operating system, and network components. Specific customer problems are analyzed and appropriate solutions are developed.

- **GoingLive Functional Upgrade Check**

This check is recommended when the customer decides to upgrade the productive SAP system. It consists of two parts. In the first part, the old release is checked (this is done before the upgrade). In the second part, a check is done after the upgrade to improve and adjust the settings of the system.

- **EarlyWatch Alert**

EarlyWatch Alert gathers information on the customers system and transfers the data to SAP. SAP analyzes this data automatically, and the generated report is sent to the customer. The customer can collect and send the data as often as they want. A lot of optimization can be done before a manual check has to be performed. On the other hand, this improves SAP's ability to give advice without looking at the customer's system, because a lot of the data is already available at SAP.

- **Migration Service**

The Migration Service is needed when a customer wants to migrate the SAP system to another database. The customer then receives the necessary software to export the old database and then import on the new database afterwards. Before the migration on the old productive system and after the productive import, checks are done to optimize the performance.

To arrange TCC services, contact SAP at:

Tel: +49 62277 43766

Fax: +49 62277 44214

Or, refer to CSS Note 35360.

Appendix D. Special notices

This publication is intended to help customers, SAP R/3 basis consultants, Business Partners, and IBM specialists who are implementing SAP R/3 for iSeries. The information in this publication is not intended as the specification of any programming interfaces that are provided by SAP R/3 and OS/400. See the PUBLICATION 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:

e (logo)® 	Redbooks
IBM ®	Redbooks Logo 
Advanced Function Printing	OfficeVision
AFCCU	OfficeVision/400
AFP	Operating System/400
AIX	OS/2
AnyNet	OS/400
Application System/400	PartnerWorld
APPN	Print Services Facility
AS/400	PROFS
AS/400e	RPG/400
AT	SAA
C/400	Service Director
CICS	SNAP/SHOT
ClusterProven	SP
COBOL/400	SP1
CT	SP2
Current	SQL/400
DB2	VisualAge
DB2 Universal Database	XT
Distributed Relational Database Architecture	400
DRDA	Lotus
IBM.COM	1-2-3
Infoprint	Approach
Information Warehouse	Lotus Notes
Intelligent Printer Data Stream	SmartSuite
IPDS	Domino
LPDA	Notes
MQSeries	Tivoli
Network Station	TME

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.

MIMIX, MIMIX/400, MIMIX/Switch, and MIMIX/Object are trademarks or registered trademarks of Lakeview Technology.

OMS/400, ODS/400, and SAM/400 are trademarks or registered trademarks of Vision Solutions, Inc.

SwitchOver System, Transformation Server, DataMirror, High Availability

Suite, and ObjectMirror are trademarks or registered trademarks of DataMirror Corporation.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

Notes, Lotus Notes, LotusScript, and NotesPump are trademarks or registered trademarks of Lotus Development Corporation in the U.S. and other countries.

SAP, R/2, R/3, ABAP/4, EarlyWatch, and SAPscript are trademarks or registered trademarks of SAP SG 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 E. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

E.1 International Technical Support Organization publications

For information on ordering these ITSO publications, see “How to get IBM Redbooks” on page 571.

- *DB2 UDB for AS/400 Visual Explain: Illustrating the Secrets*, REDP0505
- *Using AS/400 Database Monitor and Visual Explain To Identify and Tune SQL Queries*, REDP0502
- *DB2 UDB for AS/400: Database Administration with Operations Navigator V4R5*, SG24-5993 (Redpiece)
- *iSeries Handbook Version 5 Release 1*, GA19-5486
- *AS/400 Printing IV*, GG24-4389
- *IBM Network Station Guide for Windows NT*, SG24-2127
- *AS/400 - IBM Network Station - Getting Started*, SG24-2153
- *iSeries and AS/400e System Builder Version 5 Release 1*, SG24-2155
- *AS/400 Server Capacity Planning*, SG24-2159
- *AS/400 Printing V*, SG24-2160
- *Enterprise Integration with Domino.Connect*, SG24-2181
- *AS/400 AnyNet Scenarios*, SG24-2531
- *UNIX C Applications Porting to AS/400*, SG24-4438
- *Complementing AS/400 Storage Management Using Hierarchical Storage Management APIs*, SG24-4450
- *AS/400 Client/Server Performance using the Windows Client*, SG24-4526
- *AS/400 Performance Management V3R6/V3R7*, SG24-4735
- *Backup Recovery and Media Services for OS/400: A Practical Approach*, SG24-4840
- *UNIX C Applications Porting to AS/400 Companion Guide*, SG24-4938
- *AS/400 TCP/IP Autoconfiguration: DNS and DHCP Support*, SG24-5147
- *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189
- *AS/400 Client Access Express for Windows: Implementing V4R4M0*, SG24-5191
- *AS/400 Clusters: A Guide to Achieving Higher Availability*, SG24-5194
- *Slicing the AS/400 with Logical Partitioning: A How to Guide*, SG24-5439
- *New Enterprise Integration Functions for Lotus Domino for AS/400*, SG24-6203

E.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web Site for information about all the CD-ROMs offered, updates and formats.

E.3 Other publications

These publications are also relevant as further information sources:

- *MQSeries link for R/3 User's Guide*, GC33-1934
- *MQSeries for OS/400 Administration Guide*, GC33-1956
- *AS/400 System Concepts*, GC41-9802
- *AS/400 Guide to Advanced Function Presentation and Print Service Facility*, S544-5319
- *AFP Utilities for AS/400*, S544-5349
- *SAP R/3 AFP: Printing on the AS/400*, S544-5412
- *ILE RPG/400 Programmer's Guide*, SC09-2074
- *ILE RPG/400 Reference*, SC09-2077
- *ILE RPG/400 Programmer's Guide*, SC09-2507
- *ILE RPG/400 Reference*, SC09-2508
- *ILE COBOL/400 Reference*, SC09-2539
- *ILE COBOL/400 Programmer's Guide*, SC09-2540
- *OptiMover for AS/400*, SC41-0626
- *IBM Network Station Manager for AS/400 Users*, SC41-0632
- *OS/400 Security - Basic*, SC41-3301
- *AS/400 SNA Distribution Services*, SC41-3410
- *AS/400 TCP/IP Configuration and Reference*, SC41-3420
- *OS/400 TCP/IP Fastpath Setup*, SC41-3430
- *Client Access for Windows 95/NT - Setup*, SC41-3512
- *Client Access/400 for Windows 3.1 TCP/IP Setup*, SC41-3580
- *Data Description Specification*, SC41-3712
- *Printer Device Programming*, SC41-3713
- *OS/400 Server Concepts and Administration*, SC41-3740
- *AS/400 Publications Reference*, SC41-4003
- *AS/400 System Operation*, SC41-4203
- *Getting Started with AS/400*, SC41-4204
- *AS/400 System Startup and Problem Handling*, SC41-4206
- *OS/400 Security - Reference*, SC41-4302
- *OS/400 Backup and Recovery - Basic*, SC41-4304
- *OS/400 Backup and Recovery - Advanced*, SC41-4305
- *OS/400 Work Management*, SC41-4306
- *Performance Tools/400*, SC41-4340
- *Performance Tools/400-Getting Started*, SC41-4343
- *AS/400 Backup Recovery and Media Services for OS/400*, SC41-4345
- *OptiConnect for AS/400*, SC41-4414

- *DB2 for OS/400 SQL Programming*, SC41-4611
- *DB2 for OS/400 Database Programming*, SC41-4701
- *OS/400 Data Management*, SC41-4710
- *OS/400 Integrated File System Introduction*, SC41-4711
- *OS/400 NFS Support*, SC41-4714
- *AS/400 System API Reference*, SC41-4801
- *AS/400 National Language Support*, SC41-5101
- *Basic System Operation, Administration, and Problem Handling*, SC41-5206
- *Security - Basic*, SC41-5301
- *Security - Reference*, SC41-5302
- *Backup and Recovery Guide*, SC41-5304
- *AS/400 Work Management Guide*, SC41-5306
- *OS/400 Work Management*, SC41-5306
- *Distributed Data Management*, SC41-5307
- *Automated Tape Library Planning and Management*, SC41-5309
- *IBM Content Manager for OnDemand for AS/400 User's Guide*, SC41-5325
- *OS/400 Performance Tools/400*, SC41-5340
- *Performance Management/400 V4R4*, SC41-5347
- *Communications Configuration*, SC41-5401
- *X.25 Network Support*, SC41-5405
- *Communications Management*, SC41-5406
- *SNA Distribution Services*, SC41-5410
- *TCP/IP Configuration and Reference*, SC41-5420
- *TCP/IP Fastpath Setup*, SC41-5430
- *OS/400 Integration of Lotus Notes*, SC41-5431
- *AS/400 Toolbox for Java Setup Guide*, SC41-5438
- *OS/400-AS/400 Integration with Windows NT Server*, SC41-5439
- *DB2 for OS/400 SQL Programming*, SC41-5611
- *DB2 for OS/400 Database Programming*, SC41-5701
- *Data Management*, SC41-5710
- *Integrated File System Introduction*, SC41-5711
- *DDS Reference*, SC41-5712
- *Printer Device Programming*, SC41-5713
- *OS/400 Network File System Support*, SC41-5714
- *CL Programming*, SC41-5721
- *CL Reference*, SC41-5722
- *Publication Reference Book*, SC41-5003
- *System API Programming*, SC41-5800

- *System API Reference*, SC41-5801
- *AS/400 Disk Arm Requirements Based on Processor Model Performance* by Harold Rosenberg (IBM Rochester Lab, August 1999).

E.4 SAP documentation

These publications from SAP may also provide further information:

- *SAP R/3 Online Documentation CD*
- *SAP Software in PC Networking*
- *SAP-Interfaces for Program to Program Communication*
- *SAP-Supported Network Products*
- *SAP System Installation: IBM AS/400*
- *SAP Printing Guide*
- *SAP OS/DB Migration Service Planning Guide*
- *BC Computing Center Management System*
- *BC Users, Authorization, and System Security*
- *BC-SAP Printing Guide*
- *BC SAProuter*
- *BC R/3 Database Guide: DB2/400*
- *R/3 Administration*
- *R/3 Language Transport*
- *R/3 Language Combination*
- *Connection to the Online Services*
- *Installing R/3 on IBM AS/400*
- *Installing SAP Frontend Software for PCs*
- *Installing the SAP Library*
- *Online Service System*

E.5 Web sites

- IBM SAP International Competence Center ISICC home page:
<http://w3.isicc.de.ibm.com/isicc/intranet.nsf/ContentPageAliases/Welcome>
- Solutions packaging program for iSeries:
<http://www.as400.ibm.com/developer/packaging/sapr3.html>
- IBM Insight for SAP R/3 download site: <http://www.ibm.com/erp/sap/insight>
- IBM ITSO Redbooks home page: <http://www.redbooks.ibm.com>
- Legacy System Migration Workbench sign-on page: <http://service.sap.com/LSMW>
- iSeries Porting Team in Partnerworld for Developers:
<http://www.iseries.ibm.com/developer/porting/>
- iSeries home page: <http://www.ibm.com/servers/eserver/series/>
- iSeries Online Publications: <http://as400bks.rochester.ibm.com/>

- **iSeries Information Center:**
<http://publib.boulder.ibm.com/pubs/html/as400/v4r5/ic2924/index.htm>
- **iSeries Technical Studio:** <http://www.iseries.ibm.com/tstudio/>
- **iSeries Windows Integration:**
<http://www-1.ibm.com/servers/eserver/series/windowsintegration/>
- **Business solutions by IBM and SAP:** <http://www.developer.ibm.com/erp/sap/index.html>
- **iSeries Solution Team:** <http://www.iseries.ibm.com/btobpartner/>
- **SAP home page:** <http://www.sap-ag.de/>
- **SAPNet:** <https://www.sap-ag.de/notes>
- **SAP Online Help:** <http://cofsv01.yasu.ibm.com/SapHelp/HELPDATA/JA/ca/4d513456061303e10000009b38f83b/frameset.htm>

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** ibm.com/redbooks

Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

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.

- **E-mail Orders**

Send orders by e-mail including information from the IBM Redbooks fax order form to:

In United States or Canada	e-mail address pubscan@us.ibm.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.link.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at <http://w3.ibm.com/> for redbook, residency, and workshop announcements.

IBM Redbooks fax 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

ABAP	Advanced Business Application Programming Language	DYNPRO	Dynamic program concept
AFCCU	Advanced function common control unit	EBCDIC	Extended Binary Coded Decimal Interchange Code
AFP	Advanced Function Presentation	ECS	Electronic Customer Support
AFPDs	Advanced Function Printing Data Stream	EDI	Electronic Data Interchange
ALE	Application Link Enabling	EIS	Executive information system
APA	All points addressable	FDDI	Fiber distributed data interface
APAR	Authorized program analysis report	HTTP	Hypertext Transfer Protocol
API	Application program interface	IAC	Internet Application Components
ASCII	American National Standard Code for Information Interchange	IBM	International Business Machines Corporation
ASP	Auxiliary storage pool	ICA	Independent Computing Architecture
AWT	Abstract Windowing Toolkit	IDOC	Intermediate document
AS/400	Application System/400	IFS	Integrated file system
BAPI	Business application programming interfaces	ILE	Integrated language environment
BCOCA	Bar Code Object Content Architecture	I-listed PRPQ	IBM-listed PRPQ
CCMS	Computing Center Management System	IMG	Implementation guide
CGI	Common Gateway Interface	IMM	Invoke Media Map
CL	Control language	IPCS	Integrated PC Server
CLP	Control Language Program	IPDS	Intelligent printer data stream
CPI-C	Common Programming Interface-Communication	ISDN	Integrated Services Digital Network
CTS	Correction and transport system	ISICC	International SAP IBM Competence Center
DCE	Distributed computing environment	ISV	Independent Software Vendor
DDA	Domain define attributes	ITS	Internet Transaction Server
DDL	Database Definition Language	ITSO	International Technical Support Organization
DDM	Distributed data management	JVM	Java virtual machine
DDS	Data description specification	LAN	Local area network
DHCP	Dynamic Host Configuration Protocol	LAPD	Link Access Procedure for D-channel of ISDN
DMS	Document management system	LF	Logical file
DST	Dedicated Service Tools	LIC	Licensed Internal Code
		LPD	Line printer daemon
		LPO	Licensed program offering
		LPQ	Line printer queue

LPR	Line printer requester	SLS	Single-level storage
LSX	LotusScript Extension	SMAPP	System Managed Access Path Protection
MI	Machine interface	SQL	Structured Query Language
MM	Materials Management (R/3 Module)	SST	System Service Tools
MQI	Message queue interface	TCP/IP	Transmission Control Protocol/Internet Protocol
NFS	Network file system	TFTP	Trivial File Transfer Protocol
NVRAM	Non-volatile random access memory	TIMI	Technology Independent Machine Interface
OCR	Optical character recognition	TME	Tivoli Management Environment
OLE	Object linking and embedding	UPS	Uninterruptible power supply
OPM	Original Program Model	URL	Universal resource locator
OS/400	Operating System/400	WTS	Windows-based Terminal Server
OSS	Online Service System		
OTF	Output text format		
OV/400	OfficeVision/400		
PCL	Printer Control Language		
PF	Physical file		
PID	Process identifier		
POP	Post Office Protocol		
PPFA	Page Printer Formatting Aid		
PROFS	Professional Office System		
PRPQ	Programming request for price quotation		
PSF	Print Services Facility		
PTF	Program temporary fix		
RAID	Redundant array of independent disks		
RFC	Remote Function Call		
RFS	Remote file system		
SAP	Systems, Applications, Products in Data Processing		
SAP GUI	SAP graphical user interface		
SCS	SNA character string		
SD	Sales and Distribution (R/3 Module)		
SD	Solution Developers		
SDK	Software Developers Kit		
SDLC	Synchronous Data Link Control		
SID	SAP system identifier		

Index

Symbols

*JRN 226
*JRNRCV 226
/sapmnt/trans directory 75
/usr 75
/usr/sap// 75
/usr/sap//SYS/exe/run 75
/usr/sap//sys/profile 75
<SID><instance number> 79
<SID>GROUP 79
<SID>OFR 79, 364
<SID>OPR 79
<SID>OPRGRP 79
<SID>OWNER 79

Numerics

5799-AAS 117
64-bit 3, 9

A

ABAP 137, 293
abbreviations 573
abnormal end 222
Access Builder for SAP R/3 479
access method 188
 for AFP printers 188
access method L 188
access permission table 105
acronyms 573
active job 399, 400
 working with 382, 401
active user 56
active->wait 390
activity level 380, 387
add TCP/IP route information 101, 110
adding printers to R/3
 device type 157
 output device 157
adding TCP/IP remote system 103
address 11
Advanced Function Presentation (AFP) 183
Advanced Function Printing (AFP) 17
Advanced Planner and Optimizer (APO) 485
AFP (Advanced Function Presentation) 183
AFP (Advanced Function Printing) 17
AFP driver for R/3 186
AFP resources 186
A-gate process 45
Alternate IPL device 264
APARs 81
APO 485
Application Benchmark Performance Standard (SAPS) 62

Application Gate 45
application performance 437
application server 31
application software 236
Apply Journaled Changes (APYJRNCHG) command 264, 265, 546
Apply Journaled Changes Extended (APYJRNCHGX) command 545, 552
ASCII application support 117
ASCII printers 207
ASCII runtime 117
ASCII/Unicode installation 120
ASP (auxiliary storage pool) 12
ASP overflow 263
ASP1 394, 397
ASP2 394, 397
authority management 15
authorization 15
automatic tuning 387
autostart job 370
auxiliary storage pool (ASP) 12, 468
availability solution 223

B

backout recovery 266
backup and recovery 221
backup considerations 236
backup media options 236
backup performance 237
backup requirements 236
backup scheduling application 17
barcode 201
barcode printing definitions 173
barcode.tab 187
Base device type field 162
basic tuning 370
batch 370
batch immediate jobs 33
batch input 135
batch input program 137, 138
batch input session 137
battery backup unit (BBU) 4
BBP (Business-to-business Procurement) 486
BBU (battery backup unit) 4
BC (Business Connector) 489
BCI (batch immediate job) 33
BDC_CLOSE_GROUP 137
BDC_INSERT 137
BDC_OPEN_GROUP 137
BDCDATA 138
BDCDATA structure 139
benchmark user 56
benchmarks 57
bibliography 565

- binary radix tree 370
- BM Insight for SAP R/3 60
- BM Network Station 497
- BOR (Business Object Repository) 480
- buffer size 129, 434
- buffers 473
- build required indexes 474
- Business Connector (BC) 489
- Business Connector Interface 489
- Business Information Warehouse 486
- Business Object Repository (BOR) 480, 481
- Business-to-business Procurement (BBP) 486
- buttons
 - current parameters (ST02) 454
 - detail analysis (ST02) 455
 - expansion (ST03) 449
 - explain (ST04 - Index Advised) 467
 - statistic records (ST03) 448
- BW (Business Information Warehouse) 486
- byte-string 11

C

- calendar buffers 451
- CALL DIALOG statement 137
- CALL TRANSACTION USING statement 137
- capacity planning 55
- catch-up mode 276
- CCSID (coded character set identification) 113
- central instance 29
- centralized implementation 31, 43
- CFM 487
- Change Journal (CHGJRN) command 534
- changed objects 243
- character set 42, 157, 165
- characters 165
- class 36, 370
- client system 368
- Cluster Resource Services 277
- clustering 18
- coded character set identification (CCSID) 113
- command, CL
 - Apply Journaled Changes (APYJRNCHG) 265
 - APYJRNCHG (Apply Journaled Changes) 265
 - Change Authority (CHGAUT) 363
 - CHGAUT (Change Authority) 363
 - Convert Printer Data (CVTPRTDTA) 188, 212
 - Copy from Stream File (CPYFRMSTMF) 539
 - Copy to Stream File (CPYTOSTMF) 538
 - CPYFRMSTMF (Copy from Stream File) 539
 - CPYTOSTMF (Copy to Stream File) 538
 - Create Controller Description (Network) (CRTCTL-NET) 96
 - Create Output Queue (CRTOUTQ) 183, 210
 - Create Printer Device Description (CRTDEVPRT) 182
 - Create SAP User (CRTSAPUSR) 538
 - Create X.25 Line Description (CRTLINX25) 94
 - CRTCTLNET (Create Controller Description (Network)) 96
 - CRTDEVPRT (Create Printer Device Description) 182
 - CRTLINX25 (Create X.25 Line Description) 94
 - CRTOUTQ (Create Output Queue) 183, 210
 - CRTSAPUSR (Create SAP User) 538
 - CVTPRTDTA (Convert Printer Data) 188, 212
 - Delete R/3 SQL Package (DLTR3PKG) 475
 - Display File (DSPF) command 530
 - Display Network Attributes (DSPNETA) 73
 - Display Stream File (DSPSTMF) command 530
 - DLTR3PKG (Delete R/3 SQL Package) 475
 - DSPF (Display File) command 530
 - DSPNETA (Display Network Attributes) 73
 - DSPSTMF (Display Stream File) command 530
 - Edit File (EDTF) 528
 - EDTF (Edit File) 528
 - End ASP Balance (ENDASPBAL) 532
 - End Disk Reorganization (ENDDSKRGZ) 428
 - ENDASPBAL (End ASP Balance) 532
 - ENDDSKRGZ (End Disk Reorganization) 428
 - file transfer protocol (FTP) 363
 - FTP (file transfer protocol) 363
 - GO BACKUP 231
 - GO CMDRST 231
 - GO CMDSAV 231
 - GO RESTORE 231
 - GO SAVE 231
 - SAV (Save Integrated File System Objects) 231
 - SAVCFG (Save Configuration) 231
 - SAVCHGOBJ (Save Changed Objects) 231
 - SAVDLO (Save Document Library Objects) 230, 231
 - Save Changed Objects (SAVCHGOBJ) 231
 - Save Configuration (SAVCFG) 231
 - Save Document Library Objects (SAVDLO) 230, 231
 - Save Integrated File System (SAV) 231
 - Save Library (SAVLIB) 230, 231
 - Save Objects (SAVOBJ) 230, 232
 - Save R/3 System (SAVR3SYS) 540
 - Save Security Data (SAVSECDTA) 231, 232
 - Save Storage (SAVSTG) 231
 - Save System (SAVSYS) 231
 - SAVLIB (Save Library) 230, 231
 - SAVOBJ (Save Objects) 230, 232
 - SAVR3SYS (Save R/3 System) 540
 - SAVSECDTA (Save Security Data) 231, 232
 - SAVSTG (Save Storage) 231
 - SAVSYS (Save System) 231
 - Start ASP Balance (STRASPBAL) 532
 - Start Disk Reorganization (STRDSKRGZ) 428
 - Start Performance Monitor (STRPFRMON) 382, 408
 - Start Print Writer (STRPRTWTR) 183
 - Start Remote Writer (STRRMTWRT) 183
 - Start TCP (STRTCP) 74
 - Stop SAP (STOPSAP) 39
 - STOPSAP (Stop SAP) 39
 - STRASPBAL (Start ASP Balance) 532
 - STRDSKRGZ (Start Disk Reorganization) 428
 - STRPFRMON (Start Performance Monitor) 382, 408
 - STRPRTWTR (Start Print Writer) 183
 - STRRMTWRT (Start Remote Writer) 183
 - STRTCP (Start TCP) 74

- Trace ASP Balance (TRCASPBAL) 425, 532
- TRCASPBAL (Trace ASP Balance) 425, 532
- Work with Active Jobs (WRKACTJOB) 382, 400, 401, 509
- Work with Disk Status (WRKDSKSTS) 382, 390
- Work with Job (WRKJOB) 402, 404
- Work with Job by PID (WRKPID) 432, 507
- Work with Link (WRKLNK) 76
- Work with Object Links (WRKLNKSAP) 76, 530
- Work with Shared Pools (WRKSHRPOOL) 429
- Work with Subsystem Jobs (WRKSBSJOB) 382, 402
- Work with System Activity (WRKSYSACT) 410
- Work with System Status (WRKSYSSTS) 382, 388, 429
- Work with System Values (WRKSYSVAL) 382, 386
- Work with Writers (WRKWTR) 183
- WRKACTJOB (Work with Active Jobs) 382, 400, 401, 509
- WRKDSKSTS (Work with Disk Status) 382, 390
- WRKJOB (Work with Job) 402, 404
- WRKLNK (Work with Link) 76
- WRKLNKSAP (Work with Object Links) 76, 530
- WRKPID (Work with Job by PID) 432, 507
- WRKSBSJOB (Work with Subsystem Jobs) 382, 402
- WRKSHRPOOL (Work with Shared Pools) 429
- WRKSYSACT (Work with System Activity) 410
- WRKSYSSTS (Work with System Status) 382, 388, 429
- WRKSYSVAL (Work with System Values) 382, 386
- WRKWTR (Work with Writers) 183
- communication deallocate 309
- communication IDoc 330
- communication job 371
- communication line 368
- Competence Centers 553
- component 369
- concept, performance 365
- configuration 67
- configuration files 76
- configuring a new device 166
- configuring OSS 106
- configuring TCP/IP 98
- continuously powered mainstorage (CPM) 4
- Corporate Finance Management (CFM) 487
- CPI-C connection 305
- CPM (continuously powered mainstorage) 4
- create PSF configuration 205
- create remote output queue 210
- CRM (Customer Relationship Management) 487
- Cryptographic Access Provider 134
- cursor cache 451
- Customer Relationship Management (CRM) 487
- customizing 357
- cyclic fields 148

D

- DASD (direct access storage device) 3
- data export 148
- data porting 135
- data porting services 139

- data porting tools 139
- data transfer program 135, 136
- database 4, 236
- database consistency 265
- database error 517
- database lock monitor 508
- database monitor 459
- database server 30
- date 87
- DB and non DB faulting 390
- DB01 508
- DB02 459
- DB2 UDB for iSeries 14
- DBCS (double-byte character set) 116
- DECS (Domino Enterprise Connector Services) 491
- dedicated system 231
- default profile 76
- defcp.tab 187
- defect support 553
- defining field relations 147
- deleting SQL packages 474
- detail analysis menu 383
- developer traces 511
- Device class parameter 158
- device type 159
 - format 164
 - SAPGOF_E 188
- Device type parameter 158
- dialog instance 29
- dialog step 56
- direct access storage device (DASD) 3
- Disconnect R/3 System (DSCR3SYS) command 252
- disk 385
- disk activity 390
- disk fragmentation 12
- disk reorganization 428
- disk space 12
- disk status 390
 - working with 382, 390
- dispatcher 30
- Domino access to SAP R/3 business workflow 495
- Domino and SAP R/3 491
- Domino Enterprise Connector Services (DECS) 491
- Domino Mail Transfer Agent (MTA) for SAP R/3 R1.5 495
- double-byte character set (DBCS) 116
- DSCR3SYS (Disconnect R/3 System) command 252
- DSPCLS 431
- Dynamic Priority Scheduler 437
- dynamic tuning 380
- dynpro 138

E

- EarlyWatch 559
- EarlyWatch Alert 559
- ECS (Electronic Customer Support) 18
- EH&S (Environment, Health, and Safety) 488
- Electronic Customer Support (ECS) 18, 81
- em/block_size_KB 376
- em/initial_size_MB 376
- enqueue server 29

- enqueue work process 30
- Enterprise Resource Planning (ERP) 9
- Environment, Health, and Safety (EH&S) 488
- ERP (Enterprise Resource Planning) 9
- exceptional wait time 369
- EXEC-SQL 290
- expert cache 20, 429, 430
- extended memory 374
- extended memory buffer 377

F

- File Activity (ST04) panel 466
- file reorganizing 473
- file system 74, 385
- font 186
- fonts.tab 188
- form definition 186
- format types 162
- formdefs 191
- forward recovery 269
- FTP 363

G

- gateway 30
- Generic Output Format 188
- global transport directory 350, 352
- global transport parameter file 354
- GoingLive Check 559
- GoingLive Functional Upgrade Check 559
- Grant Object Authority (GRTOBJAUT) command 535
- GRTOBJAUT (Grant Object Authority) command 535
- GUI buffers 451

H

- hardware information 390
- hardware requirements 67
- hierarchical directory structure 16
- Hierarchical Storage Management (HSM) 20
- high availability 272
 - solution 273
- horizontal scaling 26
- host table 70
- hub machine 124
- Hypervisor 47

I

- IBM Performance and Testing Support/Analysis Services 61
- IBM SAP Capacity Planning Service Offering 60
- IBM SAP International Competence Center (ISICC) 556
- IBM server positioning 6
- IBM sizing support 60
- IFS (integrated file system) 15, 74
- IFS problems 520
- IMM (Invoke Media Map) 191
- index 474
- ineligible activity time 369
- information user 56

- Informational APAR 524, 558
- initializing the tape 240
- installation 67
- installing a DBCS system 118
- installing OptiMover 127
- instance 29, 376
- instance priority 431
- instance profile 29, 76
- integrated file system (IFS) 15, 74
- integration 3
- interactive 370
- internal buffers 450
- Internet Business Framework Architecture 489
- Internet Transaction Server (ITS) 489
- Invoke Media Map (IMM) commands 191
- IP address 92
- IPDS printers 203
- iSeries
 - 64-bit RISC technology 9
 - availability 4, 18
 - client/server capability 5
 - communications protocols 5
 - database 14
 - database management 4
 - DB2 UDB for iSeries 14
 - integrated file system 15
 - integration 3
 - interoperability 5
 - job management 17
 - object-based architecture 11
 - Performance Tools 409
 - printer management 17
 - security 15
 - SQL standards 14
 - system management 4, 17
 - user management 18
 - user profiles 15
 - work management 23
- iSeries Informational APAR 558
- iSeries integration 8
- ISICC (International SAP IBM Competence Center) support 556
- ISICC infoservice (Walldorf) 557
- iStar 10
- ITS (Internet Transaction Server) 489

J

- Japanese 118
- Java 23, 567
- Java SAP GUI 500
- Java Virtual Machine (JVM) 501
- job 17
 - attributes 505
 - management 17
 - status 504
 - structure 503
 - working with 402, 404
- job logs 505
- journal 226
- journal entries 225

journal receiver 226, 260, 394, 397
journaled changes 265
JVM (Java Virtual Machine) 501

K

kernel 78
KM 488
Knowledge Management (KM) 488

L

LAN 385
LAN check 407
 by ping 385
LAN performance 393
LAN-attached ASCII printers using PJI drivers 208
LAN-attached IPDS printers 203
LAN-attached printers 202
Latin-1 86, 114
layered architecture 3
layout set 195
LC LSX (Lotus Connector Lotus Script Extensions) 491, 492
Legacy System Migration Workbench (LSMW) 139
LEI (Lotus Enterprise Integration) 491
licensed programs 3
line printer daemon (LPD) 208
line printer requestor 203
line printer requestor (LPR) 208
load average 384
locating R/3 developer trace files 511
lock printer 158
log files 510
logged-on user 56
logical partitioning (LPAR) 18
logical partitions (LPAR) 234
logon load balancing 437
long running job 474
loopback 69
Lotus Connector Java Class Library 491
Lotus Connector Lotus Script Extensions (LC LSX) 491, 492
Lotus Domino Connector for SAP R/3 491
Lotus Enterprise Integration (LEI) 491
LPAR (logical partitioning) 18
LPAR (logical partitions) 234
LPAR performance 475
LPD (line printer daemon) 208
LPR (line printer requestor) 208
LSMW (Legacy System Migration Workbench) 139
LSX for the SAP R/3 system 492

M

main storage 370
Management Central 19
manual tuning 412
max. active 390
maximum active job 390
maximum active threads 430

maximum activity level 370
Maximum Transmission Unit (MTU) Size 434
MDMP (Multi-Display Multi_Process) 122
measured customer data analysis 59
Media Definition Object 237
media map selection 195
media maps 191
media style 191
memory 376, 385
memory management 371
memory pool snapshot 394
message server 29, 30
Microsoft Windows 5
MIDAS400 147
middleware servers 484
Migration Service 560
mode 30, 372
mode entry 371
modification 357
monitor 383
MTA 495
MTU Size 434
MTXW 522
Multi-Display Multi_Process (MDMP) 122
multiple language support 122
multiple overlays 191
multi-tier landscape 32, 44
mySAP.com 483

N

named user 56
national language support (NLS) 18, 86, 113
native command 382
NLS (national language support) 18, 86, 113
NLS for SAP R/3 34

O

object 15
object persistence 13
ObjectConnect 234
objects 11, 357
 created during the R/3 installation 78
Online Service System (OSS) 91, 558
operating system 8
Operations Navigator 215
OptiConnect 123, 130
 problems 520
OptiMover 127
OptiMover API 124
OptiMover for OS/400 (5799-FWQ) 123
OptiMover/400 problems 520
orientation 163
OS Collector 395
OS Data Collector 383
OS/400 Cluster Resource Services 277
OS/400 integration 8
OSS (Online Service System) 91, 558
OTF 191
OTF user barcode 201

- OTF user font 199
- output device 157, 189
- output queue 210
- output requests 181
- overlay 186

P

- page definition 186
- page format 157, 163
- page segment 186
- page-by-page media map 195
- pagedef.tab 188, 191
- paging 380, 387
- paging file sizes 377
- paging memory 373
- paper type 162, 163
- parallel backup and restore 237
- parameter changes history 406
- partitioned binary radix tree 370
- password 18, 79
- performance 5, 380
 - concept 365
 - data 408
 - data from previous hours 395
 - database 385, 398
 - review 410
 - tips 129
 - tools 409, 412
- performance monitor 382, 408
 - starting 382, 408
- Performance Tools/400 409
- personalization 485
- ping 407
- PJL driver 208
- Plaut 150
- pool 385
- pool identification 370
- pool size 370
- pool snapshot 395, 398
- post-production phase 58
- preload option 80
- pre-production performance evaluation 59
- pre-production phase 58
- pre-sales phase 58
- presentation services 30
- prestart job 371
- previous hours 395
- print 17
- print control 157
- print driver 157
- printer 17
- printer commands 182
- printer definition 156
- printing problem resolution tips 212
- priority 431
- private memory 372, 376
- problem analysis 515
- problem determination 503
- problem management 18
- process overview 402

- program buffers 451
- program temporary fix (PTF) 18
- PRPQ 545
- PSF configuration object 205
- PSF/400 185
- PTF (program temporary fix) 18
- Pulsar 10

Q

- QADRT 119
- QBATCH 370
- QDATE 87
- QDYNPTYSCD 437
- QFileSvr.400 file system 75
- QPFRADJ 387, 388
- QPWFSEVER 370
- QSOC 520
- QSTRUP program 74, 82
- QSYS.LIB file system 74
- QTIME 87
- quadword 11
- quantity-structure-based sizings 61
- questionnaire based input analysis 58
- queue 366
- queuing 369
- queuing concept 365
- Quick Sizer 61
- QUTCOFFSET 87

R

- R/3
 - printing 151
 - profiles 510
 - spool system 152
- R/3 system
 - customizing 357
- R/3 user
 - active user 56
 - benchmark user 56
 - information user 56
 - logged-on user 56
 - named user 56
- R/3-KOMPAKT 150
- R3<SID> 79
- R3<SID>400 78
- R3<SID>DATA 79
- R3<SID>JRN 79
- R3_ 79
- R3_<nn> 79
- R3400 78
- R3DATA 79
- R3GROUP 79
- R3JRN 79
- R3OPT 78
- R3OWNER 79
- race 512
- RCNR3SYS (Reconnect R/3 System) command 252
- recent periods 398
- Reconnect R/3 System (RCNR3SYS) command 252

- recovery 221, 263
- recovery restrictions 265
- regional support 556
- relational database 14
- Remote Connection Data Sheet 104
- Remote Function Call (RFC) 330
- remote journaling 274
- remote output queue 210
- reporting performance problems 475
- repository buffers 450
- request 365, 366
- required indexes 474
- resource 376
- resource-intensive functions 474
- response time 369
- restore 231
- restoring
 - the entire system 263
 - the SAP R/3 environment 264
- restricted state 231
- RFC (Remote Function Call) 330
- RFC connection 312
- RLSOUTQ 212
- roll 377
- roll and paging buffers 451
- roll area 372, 376
- roll_first 374
- route permission table 105
- routing entry 370
- RSCPINST 122
- RST 231, 264
- RSTAUT 231, 264
- RSTCFG 231, 264
- RSTDLO 231, 264
- RSTLIB 231, 264
- RSTOBJ 231
- RSTUSRPRF 231, 264
- run priority 431
- Russian 114
- RZ10 378
- RZ11 378

S

SAP

- Advanced Planner and Optimizer (APO) 485
- Application Benchmark Performance Standard (SAPS) 57, 62
- Business Connector Interface (BC) 489
- Business Information Warehouse (BW) 486
- Business-to-business Procurement (BBP) 486
- characters and character sets 165
- code page number 113
- Corporate Finance Management (CFM) 487
- Customer Relationship Management (CRM) 487
- Environment, Health and Safety (EH&S) 488
- Generic Output Format 188
- Internet Transaction Server (ITS) 489
- Knowledge Management (KM) 488
- paging memory 373
- Quick Sizer 61

- regional support 557
- Strategic Enterprise Management (SEM) 488
- SAP GUI 30
- SAP OS/DB Migration Service Planning Guide 344
- SAP R/3
 - and Domino 491
 - and Domino connection 491
 - buffers 473
 - centralized implementation 31
 - client 29
 - client/server implementation 31
 - database 29
 - instance 29
 - landscapes 31
 - monitoring 438
 - system 28
 - transaction 382, 383
 - tuning 472
- SAP R/3 Basis 27
- SAP Workplace 483
- SAP2AFP 191
- sapconf 76
- SAPDBA tool 39
- SAPGOF (SAP Generic Output Format) 188
- SAPGOF_E device type 188
- SAPNet 90
- SAPNet - R/3 Frontend 84, 558
- SAPnn 79
- SAPOSCOL 395
- SAPROUTER 106
- SAPROUTTAB 105
- saprouttab 105
- SAPS (SAP Application Benchmark Performance Standard) 57, 62
- SAPscript data 191
- SAPscript driver 162
- SAP-SQL 289
- satellite 124
- SAV 231
- SAVDLTRCV (Save and Delete Journal Receivers) command 262, 532
- SAVDLTRCVE (Stop Save and Delete Journal Receivers) command 263, 532
- Save and Delete Journal Receivers (SAVDLTRCV) command 262, 532
- save command 230
- save journal receiver 260
- Save R/3 System (SAVR3SYS) command 540
- save strategies 232
- saving all user data 243
- saving changed objects 243
- saving journal receivers 260
- saving logical partitions (LPAR) 234
- saving the entire system 241
- SBCS (single-byte character set) 113
- scalability 5
- SE06 355
- SE12 468
- SE38 122
- secondary instance 29

- security 15
- SEM (Strategic Enterprise Management) 488
- send and receive buffer size 129
- sequential file 135
- server 365, 366
- server performance 398
 - statistics 398
- server system 368
- Service Director 19
- service time 365, 366
- session 30, 372
- shared memory pool 373, 377
- shared roll memory/file 373
- short dumps 475, 511
- short name 158
- short wait 369
- short wait extended 369
- single-byte character set (SBCS) 113
- single-level storage (SLS) 11, 376
 - benefits 11
- sizing 55
 - terminology 56
- sizing and capacity planning 55
- SLIC (System Licensed Internal Code) 8, 263
- SLS (single-level storage) 11
- SM04 User list 439
- SM21 511
- SM50 399, 402, 473, 505
- SM50 transaction 399
- SM50/SM51 212
- SM66 438
- SM66 Global system-wide work processes 438
- SMLG 437
- SMLT 122
- snapshot analysis 383, 385
- snapshot information 392
- software requirements 67
- SP01 181, 212
- SPAD 156, 177
- spool
 - administration 156
 - request 179
 - server 158
- spool work process 153
- spooled files 179
- spooled request 179, 181
- SQL standards 14
- SQL trace 511
- SQL trace transaction 474
- SQL Utility (SQLUTIL) command 542
- SQLUTIL (SQL Utility) command 542
- ST02 378
- ST02 Current parameters 454
- ST02 Detail analysis menu 455
- ST02 tune summary 453
- ST03 474
- ST03 Detailed analysis 447
- ST03 Memory profile 447
- ST03 Response time analysis 445
- ST03 Time profile 446

- ST03 Top 40 by database requests 445
- ST03 Top 40 by response time 445
- ST03 Workload analysis 443
- ST03 Workload overview 444
- ST04 459, 474
- ST04 Database monitor
 - detailed object analysis 470
 - fifty slowest queries 466
 - file activity 465
 - index advised 467
 - index creates 467
 - index usage 467
 - list damaged files 470
 - missing indexes 470
 - performance tables 471
 - sorts 467
 - SQL requests 466
 - state on disk menu 467
 - system catalog tables 471
 - table scan 466
 - temporary files 467
 - wait statistics 465
- ST04 Database performance 461
- ST04 Detail analysis menu 463
- ST05 474
- ST05 SQL trace 511
 - ABAP level 472
- ST06 383, 386, 388
- ST06 transaction 399
- ST07 473
- ST07 Application monitor 440
- ST07 Database access 442
- ST07 R/3 buffers 441
- ST07 Response time 443
- ST11 511
- ST22 475, 511
- Start Performance Monitor (STRPFMON) command 408
- start profile 76
- Start SAP (STARTSAP) command 540
- Stop SAP (STOPSAP) command 540
- Stop Save and Delete Journal Receivers (SAVDLTRCVE) command 263, 532
- storage design 11
- storage management 11, 12
- storage pool 380, 387
 - definition 370
- stream files 16
- STRPFMON (Start Performance Monitor) command 408
- subsystem 23, 370
- subsystem definition 370
- subsystem description 376
- subsystem jobs 382, 402
- system 29, 236
- system activity 410
- system configuration 385
- System Licensed Internal Code (SLIC) 8, 263
- system log 406, 511
- System Managed Access Path Protection (SMAPP) 20

- system management 17
- system monitoring 380
- system priority 431
- system status 382, 388, 390, 392
- system value 382, 386, 388

T

- table buffers 451
- table lock 474
- task 366
- TCP/IP 5, 69, 129
- TCP/IP spooled file 208
- Technology Independent Machine Interface (TIMI) 3, 8
- Technology Solutions Center (TSC) 558
- Temporary storage 436
- TemSe 153
- TemSe data 153
- teraspace 13, 376
- three-tier configuration 31
- time 87
- TIMI (Technology Independent Machine Interface) 3, 8
- TMS (Transport Management System) 355
- top CPU processes 385
- trace files 510
- transaction 57, 366, 369
 - ST04 database performance 461
 - ST05 SQL Trace (ABAP level) 472
- transaction-based sizing 59
- transport domain 352
- transport domain controller 352
- transport group 352
- Transport Management System (TMS) 355
- transport request 362
- transport system 349
- TSC (Technology Solutions Center) 558
- tuning 380, 412
- two-tier 43
- two-tier landscape 31, 43

U

- Unicode 117
- uninterruptible power supply (UPS) 4
- UNIX 3
- upgrade of OS/400 90
- upgrade of SAP R/3 database 90
- upgrade of SAP R/3 kernel 90
- UPS (uninterruptible power supply) 4
- user context 372
- user data 243
- user ID 18
- user management 18
- user profile 15
- user-based sizing 59, 61

V

- V4R4 18
- vertical scaling 26
- virtual address 11

- Virtual OptiConnect 133, 476
- virtual private network (VPN) 134
- VPN (virtual private network) 134

W

- wait time 369
- Wait->Inel 390
- Web Gate 45
- W-gate process 45
- Windows NT 363
- work management 23
 - concept 370
 - response time curve 367
- work process overview 505
- work process priority 432
- workbench organizer 349
- working with
 - active jobs 382, 401
 - disk status 382, 390
 - job 402, 404
 - subsystem jobs 382, 402
 - system activity 410
 - system status 382, 388
 - system values 382, 386
- Workplace server 484

X

- X.25 hardware requirements 91
- X.25 line 91
- X.25 network provider 92
- XDA trace 512
- xxxxyyyy.tab 187

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at ibm.com/redbooks
- Fax this form to: USA International Access Code + 1 845 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number	SG24-4672-03
Redbook Title	Implementing SAP R/3 on OS/400
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	☐ Very Good ☐ Good ☐ Average ☐ Poor
Please identify yourself as belonging to one of the following groups:	☐ Customer ☐ Business Partner ☐ Solution Developer ☐ IBM, Lotus or Tivoli Employee ☐ None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	☐ Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. ibm.com/privacy/yourprivacy/



Implementing SAP R/3 on OS/400

(1.0" spine)
0.875" x 1.498"
460 x 788 pages



Implementing SAP R/3 on OS/400

Learn how to use SAP R/3 with the IBM @server iSeries server

Explore and implement the best practices, tips, and hints

Run R/3 on the iSeries faster, more smoothly, and easily

This IBM Redbook features a collection of knowledge gained from IBM and SAP solution experts who work with customers that use SAP R/3 on the IBM @server iSeries server. It was written to assist R/3 basis consultants and other IT professionals in implementing a total business solution consisting of iSeries and AS/400 servers, OS/400 Version 4 Release 5 (V4R5), SAP R/3 Release 4.6C, DB2 UDB for iSeries database, and complementary solution products. The primary content of this Redbook is divided into three parts.

Part 1, "Understanding the solution", presents the concepts and other basic knowledge necessary to understand the structure, features, and functions of the SAP R/3 solution on the iSeries server.

Part 2, "Implementation", describes the implementation techniques necessary to install and properly set up R/3 in the iSeries environment. It contains detailed guidance and explanations of the specific tasks associated with the implementation. Professionals involved in implementing R/3 on OS/400 may, at some stage, face all the topics covered in this part.

Part 3, "Advanced topics", covers topics that will be of interest to those who want to enhance their SAP R/3 installation by improving performance and adding additional functionality.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks