Implementing SAP Applications on the IBM System i Platform with IBM i5/OS

Comprehensive yet simple SAP business applications with i5/OS

Successful business solutions in an on demand complex world

SAP applications and System i integration

Susan Powers
Matthias Bartusch
Manfred Engelbart
Dieter Huber
Gert Ruland
Sabine Schloesser
Ron Schmerbauch
Steve J. Tlusty

Redbooks

ibm.com/redbooks

International Technical Support Organization

Implementing SAP Applications on the IBM System i Platform with IBM i5/OS

August 2006

SG24-7166-00

Note: Before using this information and the product it supports, read the information in "Notices" on page xxv.

First Edition (August 2006)

This edition applies to Version 5, Release 3, Modification 5 of i5/OS (product number 5722-SS1) and mySAP ERP 2004, SAP ECC 5.0 based on SAP NetWeaver(R) ´04.

Contents

Figures xv

Tables xxiii

Notices xxv
Trademarks xxvi

Preface xxvii
The team that wrote this redbook. xxvii
Become a published author xxx
Comments welcome. xxx

Part 1. SAP applications on System i models with i5/OS 1

Chapter 1. IBM System i solution for SAP 3
1.1 The integrated approach 4
1.2 DB2 UDB for i5/OS: Integrated database for System i family, optimized for SAP 4
1.3 Storage management 5
1.4 Consolidate homogeneous workloads under i5/OS 5
1.5 Consolidate heterogeneous workloads on a System i server 6
1.5.1 Heterogeneous operating systems 6
1.5.2 Heterogeneous workloads 6
1.5.3 Heterogeneous PTF levels 6
1.6 IBM System i value proposition for SAP 6

Chapter 2. SAP Solution Portfolio 9
2.1 History 10
2.2 mySAP Business Suite 11
2.3 SAP NetWeaver architecture 12
2.3.1 People integration 13
2.3.2 Information integration 13
2.3.3 Process integration 13
2.3.4 Application integration. 14
2.4 Complexity of SAP NetWeaver implementation 14
2.5 Critical success factors 14

Chapter 3. SAP system landscapes 17
3.1 SAP system architectural dimension. 19
3.1.1 2-tier implementation 19
3.1.2 3-tier implementation 20
3.2 Operational stages of SAP systems (change management) 20
3.3 SAP system dimension 20
3.4 Naming conventions 21

Chapter 4. SAP system landscape scenarios 23
4.1 Example environment descriptions 24
4.2 Small configuration scenario 24
4.2.1 Components 24
4.2.2 Schema. 24
4.3 Medium configuration scenario 25

4.3.1 Components . . . 25
4.3.2 Schema . . . 26
4.4 Large configuration scenario . . . 27
4.4.1 Components . . . 27
4.4.2 Schema . . . 27
4.5 Availability considerations and influence on landscape design . . . 28
4.6 System i value proposition for SAP NetWeaver landscapes . . . 29

Chapter 5. System i virtualization . . . 31
5.1 Virtualization . . . 32
5.1.1 Static LPAR . . . 32
5.1.2 Dynamic LPAR . . . 32
5.1.3 Micro partitions . . . 33
5.2 Implementation options . . . 33
5.3 Summary . . . 34

Chapter 6. Encoding data used in SAP systems . . . 35
6.1 What are encoding schemes? . . . 36
6.2 SAP direction . . . 36
6.3 What encoding schemes are available for SAP applications on i5/OS? . . . 36
6.3.1 Unicode . . . 37
6.3.2 GLS . . . 37
6.4 Choosing between Unicode and GLS . . . 37
6.5 Converting an installed system to Unicode . . . 38
6.5.1 Moving from GLS to Unicode . . . 38
6.5.2 Moving from EBCDIC to Unicode or GLS . . . 38

Chapter 7. SAP sizing considerations . . . 41
7.1 Overview of the sizing process . . . 42
7.2 Size with the help of the SAP Quicksizer . . . 43
7.2.1 Input to Quicksizer . . . 43
7.2.2 Considerations for the Quicksizer output . . . 44
7.2.3 Work with Quicksizer output . . . 44
7.2.4 Considerations for sizings not based on the Quicksizer . . . 45
7.2.5 Size upgrades . . . 46
7.2.6 Summary . . . 46

Part 2. Operational topics . . . 47

Chapter 8. i5/OS from an SAP application view . . . 49
8.1 System i architecture and overview . . . 50
8.1.1 Single-level storage . . . 50
8.1.2 Layered architecture . . . 51
8.1.3 Overview of i5/OS . . . 52
8.1.4 Subsystems . . . 52
8.1.5 Auxiliary Storage Pool . . . 53
8.1.6 Independent ASP . . . 53
8.1.7 Logical partitions . . . 54
8.2 i5/OS object structure . . . 55
8.2.1 Object types . . . 55
8.2.2 SQL collection objects . . . 56
8.2.3 Database access from the SAP system . . . 60
8.2.4 The Integrated File System . . . 63
8.2.5 Language support . . . 70

8.2.6 TCP/IP 71
8.2.7 Security and authorization 72
8.3 System i software products for an SAP implementation 72
8.3.1 Software product names and options 72
8.3.2 Work with licensed programs 73
8.3.3 System i software product licensing 74
8.3.4 Required IBM software products for an SAP implementation 74
8.3.5 Useful and optional software products for SAP systems 76
8.3.6 Program temporary fixes 76
8.4 Additional tools for SAP systems 78
8.4.1 SQLUTIL 79
8.4.2 Other tools 80
8.4.3 Interactive SQL 81
8.5 i5/OS commands 82
8.6 i5/OS Portable Application Solutions Environment 84
8.6.1 i5/OS PASE on System i models 84
8.6.2 i5/OS PASE with SAP systems 85
8.7 Special operating system settings for SAP applications 86
8.7.1 Multiple SAP systems on one server 86
8.7.2 Adjusting the run priority of SAP work processes 86
8.8 Java 88
8.9 Java database connectivity 88

Chapter 9. SAP storage and database considerations 89
9.1 Hardware disk configuration 90
9.1.1 Internal and external disks 91
9.2 System i integrated storage 93
9.3 IBM TotalStorage solutions 94
9.3.1 Fibre Channels for external disks in an SAP system environment 96
9.4 Planning for disk storage in an SAP installation 97
9.5 Basic principles of SAP databases and SAP systems 98
9.5.1 SAP databases and SAP journals in standard installations 98
9.5.2 SAP databases with shared or common SAP journals 99
9.5.3 Multiple components in one database and their journals 101
9.6 SAP applications and journal management 102
9.6.1 The SAP journal and SAP journal receiver 102
9.6.2 Default SAP application *JRN and *JRNRCV settings 106
9.6.3 Manual procedure to set up journal logging 107
9.6.4 Journal switch 108
9.6.5 Reduction in journal receiver size 109
9.7 ASP considerations for SAP journal management 110
9.7.1 Using a separate ASP for the journal receiver 110
9.7.2 General recommendations 111
9.7.3 Sizing the SAP journal receiver. 112
9.7.4 ASP overflow 113

Chapter 10. Users and authorities 115
10.1 Authorization concepts 116
10.1.1 Objects, owners, and users 116
10.1.2 Solutions for some authority situations 118
10.1.3 The i5/OS concept of adopted authority 121
10.2 i5/OS system users with SAP applications 124
10.2.1 SAP system administrator 125

10.2.2 SAP system operator 126
10.2.3 ABAP and JDBC user profiles. 126
10.2.4 Characteristics of the SAP system user profiles on i5/OS 127
10.3 Environment variables. 131

Chapter 11. Installation overview 133
11.1 SAP documentation 134
11.2 IBM Redpaper 134
11.3 Set up an SAP system by system copy 134
11.4 Differences between 2-tier and 3-tier landscapes 135
11.5 Upgrade guides 136

Chapter 12. The SAP kernel on the System i server 137
12.1 Introduction to SAP kernels. 138
12.1.1 What is a kernel? 138
12.1.2 Naming the kernel library 139
12.1.3 Using separate kernels for each SAP system. 140
12.1.4 Other libraries that include runtime objects (J2EE and RFC) 141
12.2 SAP kernel, operating system and database release 144
12.2.1 Released operating systems for SAP kernel 4.6x and DB2 UDB for iSeries . . 145
12.2.2 Released operating systems for SAP kernel 6.x and DB2 UDB for iSeries . . . 145
12.2.3 Maintenance for SAP kernels 145
12.3 Managing the SAP kernel 147
12.3.1 SAP kernel when planning an operating system upgrade 147
12.3.2 Loading an SAP kernel 149
12.3.3 Applying an SAP kernel 152
12.3.4 Updating a Java kernel 154
12.3.5 SAP kernel patches 154
12.3.6 Downloading SAP kernel patches. 158
12.3.7 Applying the SAP kernel patches 164
12.4 Kernel 6.40 content 171

Chapter 13. SAP instance profile parameters 173
13.1 SAP instance profile parameters for i5/OS installations 174
13.2 SAP instance profile parameter values for i5/OS installations 176

Chapter 14. Software logistics 181
14.1 SAP system landscape, clients, and roles 182
14.2 SAP system landscapes 183
14.2.1 SAP system environments 183
14.2.2 One-system landscape 185
14.2.3 Two-system landscape 185
14.2.4 Three-system landscape. 185
14.3 Manage clients after the installation of an SAP system 187
14.4 Global transport directory 189
14.4.1 Setup the Transport Management System 192
14.4.2 Define the transport domain and transport routes. 194
14.5 Customization and development process in an SAP system 197
14.5.1 Client attributes 197
14.5.2 Transport change requests 199
14.6 Handling of a repository object change request 200
14.6.1 The SAP Software Change Registration (SSCR) 200
14.6.2 Creating a repository object change request 202
14.6.3 Transport Organizer 203

14.6.4 Strategies for importing requests . . . 204
14.6.5 Workflow strategy for imports . . . 204
14.6.6 Use STMS to import single requests. . . . 205
14.6.7 Use TP to import single requests . . . 207
14.7 Transport objects between SAP systems on different host systems. . . . 208
14.7.1 Homogeneous environment . . . 208
14.7.2 Heterogeneous environments . . . 208
14.8 Test the Transport Management System . . . 211

Chapter 15. The SAP system of printing . . . 213
15.1 Introduction to printing with SAP applications . . . 214
15.2 Spool support architecture . . . 216
15.2.1 Overview of i5/OS spool support. . . . 216
15.2.2 Overview of the SAP spool system . . . 216
15.2.3 TemSe data . . . 217
15.2.4 The spool work process . . . 217
15.3 Spool request and spooled files . . . 220
15.4 Printing with i5/OS systems . . . 221
15.4.1 i5/OS printer data streams . . . 221
15.4.2 AFP printing . . . 222
15.4.3 i5/OS printer commands . . . 223
15.4.4 i5/OS printer configuration . . . 224
15.5 Administration of the SAP spool system . . . 234
15.5.1 Components of a printer definition . . . 234
15.5.2 Output devices . . . 235
15.5.3 Configure Access Method C . . . 237
15.5.4 SAP device types . . . 239
15.5.5 Paper types . . . 241
15.5.6 Print controls . . . 242
15.5.7 SAP characters and character sets. . . . 243
15.6 Configure new SAP system device . . . 244
15.7 Barcode printing definitions . . . 246
15.8 Manage SAP spooled and output requests. . . . 247
15.8.1 Attributes of a spool request . . . 247
15.8.2 Spool request actions . . . 248
15.9 Resolving printing problems . . . 248
15.10 Printing documentation references . . . 249

Chapter 16. SAP system administration . . . 251
16.1 SAP system monitoring . . . 252
16.1.1 Monitoring session information . . . 253
16.1.2 Monitoring SAP information . . . 254
16.1.3 Monitoring an SAP server: Transaction SM51 . . . 256
16.1.4 Monitoring SAP work processes: Transactions SM50 and SM66. . . . 257
16.1.5 Monitoring an SAP user: Transactions SM04 and AL08. . . . 257
16.1.6 Monitoring asynchronous update and update processing: Transaction SM13 . 257
16.1.7 Monitoring SAP locks: Transaction SM12. . . . 258
16.1.8 Monitoring ABAP dumps: Transaction ST22. . . . 259
16.1.9 Monitoring the SAP system log: Transaction SM21 . . . 260
16.2 SAP system administration tasks . . . 262
16.2.1 Global administration tasks. . . . 262
16.2.2 Activities before stopping an SAP system. . . . 265
16.2.3 SAP profiles . . . 266

16.2.4 SAP background processing. 268
16.2.5 SAP and DB2 UDB for iSeries data base administration in CCMS. 270
16.2.6 Maintaining users and authorities 271
16.3 SAP Computing Center Management System (CCMS) 272
16.3.1 CCMS overview 272
16.3.2 Setting up CCMS 272
16.4 System Administration Assistant (SSAA transaction) 273
16.4.1 Using keyword search to find tasks. 273
16.4.2 Maintaining the System Administration Assistant 273
16.4.3 Defining a new task. 274
16.4.4 Configuring a system landscape 275
16.4.5 Accepting system landscape settings 275
16.4.6 Allowing remote access 275
16.5 SAP Solution Manager 276
16.5.1 Monitoring an SAP system 277
16.5.2 SAP EarlyWatch Alert 281
16.6 Daily administration tasks - sample lists 281
16.7 Useful SAP system administration transactions 287

Chapter 17. General topics. 293
17.1 Control network access with SAProuter 294
17.2 Using the IBM Job Scheduler in an SAP system environment 302
17.3 Automated startup. 305
17.3.1 i5/OS functions 305
17.3.2 QSTRUP in an SAP system environment. 307
17.3.3 Starting SAP applications during an IPL 309
17.3.4 Concerning Linux Application Server 310
17.3.5 Automatically start and stop multiple instances and the application server . . . 310
17.3.6 Starting Windows AppServer from System i central host on remote systems . 312
17.4 Copying CDs and DVDs to hard disk 313

Chapter 18. Life cycle 317
18.1 Growth management. 318
18.1.1 Storage capacity growth 318
18.1.2 Workload growth 320
18.2 Tools to track resource utilization 320
18.3 System i resource upgrades 321
18.3.1 Adding memory or disk resources to a partition 321
18.3.2 Adding CPU resources to a partition. 321
18.3.3 Instance profile tuning changes 322
18.4 Software release updates 322
18.4.1 i5/OS release upgrade 322
18.4.2 i5/OS database upgrade 323
18.4.3 SAP release upgrade 323
18.5 Add SAP systems 324
18.6 Adding languages to SAP applications 324
18.7 Migration and code page conversion 324
18.7.1 Migration 325
18.7.2 Code Page Conversion (CPC) 325

Chapter 19. Problem avoidance. 331
19.1 Prevent instead of react 332
19.2 Hardware. 332
19.3 IBM software 332

19.3.1 IBM Informational APARs for SAP customers 332
19.3.2 Frequency of PTF application 333
19.3.3 IBM software support 334
19.4 SAP software corrections 334
19.4.1 SAP Support Marketplace for Notes and Hot News 334
19.5 Testing 334
19.6 History 335
19.7 Avoid more problems once one is identified 335
19.8 Education 337

Chapter 20. Problem determination and management 339
20.1 Implementing SAP applications on i5/OS 340
20.1.1 Job structure 340
20.2 Working with job logs 342
20.2.1 Changing the job attributes 342
20.2.2 Work process overview 342
20.2.3 The WRKPID command 344
20.2.4 The database lock monitor (DB01) 345
20.2.5 The WRKACTJOB command 346
20.2.6 Printing and locating the job log 346
20.2.7 SAP instance profiles 347
20.3 Trace and log files 347
20.4 The CHKXDA SAP tool 349
20.5 Special considerations for SAP Java 351
20.6 Problem analysis 351
20.6.1 Where to look first 351
20.6.2 Database error 353
20.6.3 Program exceptions 354
20.6.4 Performance problems 354
20.6.5 Integrated File System problems 356
20.6.6 Printing problems 356
20.6.7 Damaged objects on the System i server 356
20.6.8 Unable to start the SAP application 357
20.6.9 Inability to stop a job or work process 359
20.6.10 SAPOSCOL 360
20.7 Reporting the problem 360
20.7.1 SAP system environment 361
20.7.2 i5/OS environment 361
20.7.3 Saving spooled files 362
20.7.4 Reporting the problem to SAP 363
20.8 Additional information 363

Chapter 21. Correctable scenarios 365

Part 3. Availability, backup and recovery 371

Chapter 22. Availability, backup and recovery concepts 373
22.1 Scheduled outages 374
22.2 Unscheduled outages 374
22.3 Availability solutions for unscheduled outages 375
22.3.1 Hardware solutions 375
22.3.2 Software solutions 377
22.3.3 Backup solutions 379
22.3.4 Availability solutions 380

22.4 Backup and recovery . . . 380
22.4.1 i5/OS save and restore commands . . . 382
22.4.2 Save strategies . . . 384
22.4.3 Backup considerations . . . 388
22.4.4 Recommendations prior to backup . . . 389

Chapter 23. Backup and recovery . . . 391
23.1 Backup . . . 392
23.1.1 Backup methods . . . 392
23.1.2 Initializing the tape . . . 393
23.1.3 Offline backup . . . 394
23.1.4 Online backup of the SAP database . . . 398
23.1.5 Saving journal receivers . . . 405
23.1.6 Saving SAP Integrated File System objects . . . 408
23.2 Recovery . . . 412
23.2.1 User ASP overflow recovery . . . 412
23.2.2 Restoring the entire system . . . 413
23.2.3 Restoring the SAP system environment . . . 413
23.2.4 Recovering the SAP database . . . 415
23.3 Using BRMS for save and restore . . . 422

Chapter 24. Disaster recovery and high availability . . . 433
24.1 Solution discussion . . . 434
24.1.1 Disaster recovery . . . 434
24.1.2 High availability . . . 434
24.2 Concepts and basic technologies . . . 435
24.2.1 Logical partitioning (LPAR) . . . 435
24.2.2 Backup system . . . 435
24.2.3 Clustering . . . 436
24.2.4 Independent Auxiliary Storage Pool . . . 437
24.2.5 Mirrored system environments . . . 438
24.2.6 Remote journaling . . . 441
24.3 SAP system implementation with a switchable IASP . . . 447
24.3.1 Switchable device environment . . . 451
24.3.2 Cross site mirroring . . . 451
24.3.3 System i copy services . . . 452
24.4 Logical replication . . . 453
24.4.1 Business Partner solutions . . . 454

Part 4. Optimizing performance . . . 455

Chapter 25. Performance concepts and monitoring . . . 457
25.1 System i performance concepts . . . 458
25.1.1 Client/server response time . . . 458
25.1.2 Components of System i interactive response time . . . 458
25.1.3 Impact of database growth on System i performance . . . 459
25.2 System i work management concepts . . . 460
25.2.1 Subsystems . . . 460
25.3 Memory management of an SAP system . . . 462
25.3.1 Functions of the SAP memory management system . . . 462
25.4 Performance monitoring . . . 470
25.4.1 System i monitors . . . 470
25.4.2 WRKSYSSTS . . . 471
25.4.3 WRKDSKSTS . . . 473

25.4.4 WRKACTJOB 473
25.4.5 WRKSYSACT 478
25.4.6 Start Data Collection Services 479
25.4.7 iSeries Navigator 483
25.4.8 SAP transaction ST06 485
25.5 Monitoring SAP systems 485
25.5.1 CCMS 485
25.5.2 Solution Manager 492
25.5.3 Alert Monitor (RZ20) 494
25.5.4 Transaction SSAA 495
25.5.5 Global (System Wide) Work Process Overview (SM66) 496
25.5.6 Process overview (SM50) 498
25.5.7 System Log - Analysis (SM21) 500
25.5.8 Logged on users (SM04) 501
25.5.9 Application Monitor(ST07) 502
25.5.10 Workload analysis (ST03) 505
25.5.11 Tune summary and buffers (ST02) 511
25.5.12 Performance trace (ST05) 514
25.5.13 Monitoring Java 515
25.6 Database monitoring 516
25.6.1 Database performance: State on Disk (DB02) 516
25.6.2 Enabling database monitoring 517
25.6.3 Database performance monitor (ST04) 519
25.6.4 Database monitor: Transaction DB4COCKPIT 522
25.7 SAP Business Intelligence performance 530

Chapter 26. Performance tuning 531
26.1 Tuning System i models 532
26.1.1 Tuning an SAP system 539
26.1.2 Database tuning 545
26.2 Analyzing and optimizing single transactions 554
26.2.1 Transaction ST05: performance analysis 555
26.2.2 SQL trace 557
26.3 Java optimization 562

Chapter 27. Performance management tools and services 565
27.1 SAP EarlyWatch 566
27.1.1 IBM Insight 568
27.2 Performance aspects in 3-tier environments 575
27.3 Performance aspects in complex system landscapes 577
27.3.1 SAP OS Collector 577
27.3.2 Multiple system considerations 578
27.3.3 Testing systems with lower priority 581
27.4 Performance management services 583
27.4.1 PM iSeries 583
27.4.2 Performance Tools 585
27.4.3 iDoctor for iSeries 590
27.4.4 Workload Estimator 593
27.4.5 SAP services 593

Chapter 28. SAP Business Information Warehouse 597
28.1 Query performance 598
28.1.1 InfoCubes 598
28.1.2 Operational Data Store 608

28.2 Troubleshooting poor performing queries 608
28.2.1 Query Monitor 609
28.2.2 Visual Explain 610
28.3 Data load performance 612
28.3.1 Initial InfoCube data loads 612
28.3.2 Number range buffering 613
28.3.3 Index maintenance 614
28.3.4 ODS activation settings 615
28.4 SAP notes for SAP BW 616
28.5 SAP BW naming convention for database tables 617
28.6 Additional DB2 information 618

Part 5. Data exchange between applications 619

Chapter 29. Data exchange scenarios and examples 621
29.1 Considerations 622
29.2 Example programs 625
29.2.1 RPG/400 example 625
29.2.2 ABAP example 626
29.3 Accessing SAP systems data using an RPG/400 program 628
29.3.1 CRTSAPLF: Accessing SAP Unicode or GLS Latin-1 code page data 632
29.3.2 CRTSAPVIEW: Accessing SAP GLS data with non Latin-1 code pages 637
29.4 Accessing non-SAP data with ABAP programs 641
29.4.1 Accessing non-SAP data with EXEC-SQL 642
29.4.2 Accessing non-SAP data with database multiconnect 643
29.4.3 Accessing non-SAP data with open dataset 648
29.5 Process i5/OS commands from SAP applications 649
29.5.1 Defining i5/OS commands 650
29.5.2 Executing i5/OS commands 650
29.5.3 Defining i5/OS job steps for SAP system background jobs 652
29.6 System i jobs to start SAP applications 655
29.6.1 The SAPEVT command 655
29.6.2 The STRREPORT command 656
29.7 Interactive program communication 657
29.7.1 Using an RFC connection 657
29.7.2 ABAP calling ILE or Java: flowchart of RFC connection 658
29.7.3 ILE RPG or Java program calling ABAP: Flowchart of RFC connection 680
29.7.4 Monitoring communication status 688

Chapter 30. Data porting 689
30.1 Concept of data porting to SAP applications 690
30.2 Programming aspects for data porting to an SAP system 691
30.2.1 Data transfer program 691
30.2.2 Batch input program 691
30.3 Data porting services 694
30.4 Data porting tools 694
30.4.1 Before working with LSMW 694
30.4.2 Legacy System Migration Workbench 695

Part 6. Appendices 709

Appendix A. Support for SAP applications 711
A.1 Marketing and technical support 712
A.2 Defect support 712

A.3 IBM SAP International Competence Center (ISICC) support . . . 712
A.3.1 ISICC InfoService . . . 713
A.3.2 SAP regional support . . . 713
A.4 Sizing support . . . 713
A.5 Information access . . . 714
A.5.1 SAP Service Marketplace . . . 714
A.5.2 SAPNet - Frontend . . . 714
A.5.3 System i Informational APARs . . . 714
A.5.4 iSeries Technology Solutions Center . . . 714
A.5.5 SAP Support Services . . . 715
A.6 Education . . . 716
A.6.1 Basis Administration for SAP on the System i platform . . . 716
A.6.2 Code Page Conversion . . . 716

Abbreviations and acronyms . . . 717

Related publications . . . 721
IBM Redbooks . . . 721
Other publications . . . 722
Online resources . . . 722
How to get IBM Redbooks . . . 723
Help from IBM . . . 723

Index . . . 725

Figures

1-1 IBM System i - An integrated system. 3
2-1 History of mySAP Business Suite 10
2-2 mySAP Business Suite 11
2-3 SAP NetWeaver Architecture. 12
2-4 SAP NetWeaver Architecture, ABAP and Java 14
3-1 The complexity of SAP system landscapes 18
3-2 SAP system architecture layers 19
3-3 Technical structure of SAP systems 21
4-1 Small landscape example 25
4-2 Medium landscape example 26
4-3 Large landscape example 28
8-1 Single-level storage 50
8-2 The System i layered architecture 51
8-3 The System i Independent ASP. 53
8-4 Logical partition concept 54
8-5 SQL collection 58
8-6 SQL long name and i5/OS names of objects and file systems 59
8-7 The Integrated File System 63
8-8 Remote file system objects and file systems. 65
8-9 SAP Integrated File System directory structure 66
8-10 Integrated File System structures for the SAP Web Application Server, ABAP stack 68
8-11 Integrated File System structures for the Web Application Server, ABAP and Java . 69
8-12 Object and file system language support. 71
8-13 Options of the Work with Licensed Programs menu 73
8-14 The i5/OS PTF menu. 77
8-15 Interactive SQL user interface, STRSQL and SQLUTIL 82
9-1 Planning for storage. 96
9-2 SAP ABAP and Java database and SAP journals in a standard SAP installation . . . 98
9-3 SAP databases with shared or common journals 99
9-4 SAP databases in a MCOD implementation 102
9-5 The Create Journal Receiver command CRTJRNRCV 103
9-6 The Create Journal CRTJRN command (1/2) 104
9-7 The Create Journal CRTJRN command (2/2) 104
9-8 The Change Journal command CHGJRN (1/2) 105
9-9 The Change Journal command CHGJRN (2/2) 105
9-10 The i5/OS command EDTRCYAP (Edit Recovery for Access Paths) 109
10-1 The SAP command FIXR3OWNS 119
10-2 The SAP command AS4FIXFILE. 120
10-3 The i5/OS CHGPGM command to change the USEADPAUT parameter 122
10-4 i5/OS authority search order 123
10-5 SAP System Administrator. 125
10-6 SAP system operator. 126
10-7 SAP ABAP and JDBC user profiles 127
10-8 Operating system user's settings. 128
10-9 SAP system environment variables within SAP system environments and i5/OS. . 131
12-1 The Load SAP kernel Library (LODR3KRN) command. 149
12-2 SAP Software Download Center ("swdc") on the SAP Service Marketplace. 151
12-3 SAP kernel component concept ("Binary Patches"). 154

12-4 Applying SAP kernel components (“Binary Patches”) 155
12-5 SAP support package stacks 157
12-6 The SAP Service Marketplace Software Distribution Center (SWDC). 158
12-7 Overview SAP support packages and patches grouped by applications. 159
12-8 The SAP 64-bit kernels 159
12-9 Selection of 6.40 GLS kernel patches from the SAP “SWDC”. 160
12-10 6.40 ASCII kernel patches (extract) from SAP Software Distribution Center. 161
12-11 6.40 GLS kernel patches in the basket of the SAP Download Manager 162
12-12 Download Manager from the SAP “SWDC”. 164
12-13 Extract from the developer trace 165
12-14 APYR3FIX command (1 of 2) 166
12-15 APYR3FIX command (2 of 2) 166
14-1 Client structures in a three-system landscape. 188
14-2 Transport domain for a complex SAP system landscape 189
14-3 The SAP transport landscape 190
14-4 /sapmnt/trans shared directory 191
14-5 TMS setup of three-system landscape showing transport routes (graphical) 195
14-6 TMS setup of three-system landscape showing transport routes (hierarchical) . . . 196
14-7 Client-dependent attributes 198
14-8 The SAP Software Change Registration (SSCR) for an SAP object 201
14-9 SAP transport organizer starting window (transaction SE10 or SE09) 203
14-10 TMS Workflow Engine 204
14-11 Import a change request: Date selection 206
14-12 Import a change request: Execution selection. 206
14-13 Import a change request: Options selection 206
15-1 Flow of SAP output data on a System i server 214
15-2 SAP printing with i5/OS: Information flow 217
15-3 Spool work process on the database server 218
15-4 Spool work process on the application server 219
15-5 SAP spool request and output request 220
15-6 EBCDIC and ASCII based data stream 221
15-7 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 1 of 8) . . 226
15-8 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 2 of 8) . . 226
15-9 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 3 of 8) . . 227
15-10 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 4 of 8) . . 227
15-11 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 5 of 8) . . 228
15-12 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 6 of 8) . . 228
15-13 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 7 of 8) . . 229
15-14 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 8 of 8) . . 229
15-15 CRTOUTQ command (part 1 of 4). 230
15-16 CRTOUTQ command (part 2 of 4). 230
15-17 CRTOUTQ command (part 3 of 4). 231
15-18 CRTOUTQ command (part 4 of 4). 231
15-19 Create a WSTCST or a device description 232
15-20 The i5/OS WRKOUTQ command 233
15-21 Components of a printer definition. 234
15-22 Configure access method “C” 237
15-23 Access method C and U 238
16-1 Monitoring the SAP system 252
16-2 Drop-down menu on the SAP GUI front-end window 253
16-3 SAP session information 253
16-4 The SAP system status (1/2) 255
16-5 The SAP system status (2/2) 255

16-6 Update processing 259
16-7 The CCMS Monitor templates (Syslog) from transaction RZ20 261
16-8 SAP NetWeaver Technical Operations Manual part 1 of 3 - ABAP 263
16-9 SAP NetWeaver Technical Operations Manual (part 2 of 3 - Java) 264
16-10 SAP NetWeaver Technical Operations Manual (part 3 of 3 - other components) . . 265
16-11 Sequence of assigning the SAP profile parameters 266
16-12 Maintaining SAP profiles 267
16-13 SAP background jobs selected for execution 269
16-14 Users in the SAP system environment 271
16-15 SAP Solution Manager - Monitoring, Reporting and Data Collection 277
16-16 SAP Solution Manager - Alert Monitoring 279
17-1 SAProuter as a firewall 294
17-2 How to find and get the SAProuter patch from SAP Service Marketplace 296
17-3 i5/OS CHGIPLA command 306
18-1 Code pages for SAP systems running on i5/OS - Source SAP 329
20-1 SAP jobs in the subsystem 341
20-2 Work process overview (SM50) 343
20-3 Work process table (DPMON) 344
20-4 WRKPID command 345
20-5 Database Lock Monitor (DB01) 346
20-6 Short dump 352
20-7 WRKPTFGRP, F11 - Display descriptions 362
22-1 Save commands and Save menu options 381
22-2 Restore commands and Restore menu options 382
22-3 SAP WebAS ABAP system - directory structure 385
22-4 SAP WebAS ABAP and Java system - directory structure 386
23-1 Initialize tape (INZTAP) 394
23-2 SAVLIB: Save the SAP database (Part 1 of 3) 396
23-3 SAVLIB: Save the SAP database (Part 2 of 3) 397
23-4 SAVLIB: Save the SAP library (Part 3 of 3) 397
23-5 SAVLIB: Save-While-Active (Part 1 of 3) 399
23-6 SAVLIB: Save-While-Active (Part 2 of 3) 400
23-7 SAVLIB: Save While Active (Part 3 of 3) 400
23-8 SAVLIB: Rapid Checkpoint Save-While-Active (Part 1 of 3) 401
23-9 SAVLIB: Rapid Checkpoint Save-While-Active (Part 2 of 3) 402
23-10 SAVLIB: Rapid Checkpoint Save-while-Active (Part 3 of 3) 403
23-11 STRSQL: analyzing associated journal receivers 404
23-12 Tables with pending transactions and associated journal receivers 404
23-13 WRKJRNA and F15: work with receiver directory 406
23-14 Save object 409
23-15 OBJ parameter 410
23-16 Save Object 411
23-17 Save object, continued 412
23-18 DSPJOBLOG: message CPI3731 414
23-19 DSPJRN: Finding the journal sequence number 417
23-20 RMVJRNCHG (Part 1 of 2) 418
23-21 RMVJRNCHG (Part 2 of 2) 418
23-22 WRKJRN: associate receivers with journal 420
23-23 APYJRNCHGX (Part 1 of 2) 421
23-24 APYJRNCHGX (Part 2 of 2) 421
23-25 Define recovery policy 423
23-26 LInk definition: BRMS Integrated File System backup 423
23-27 WRKCTLGBRM: BRMS offline backup 424

23-28 WRKCTLGBRM: BRMS Save-While-Active 425
23-29 WRKCTLGBRM: option 8, change backup control attributes 426
23-30 WRKCTLGBRM: BRMS Rapid Checkpoint Save-while-Active 427
23-31 WRKCTLGBRM: option 8 page 3 of 4 Change Backup Control group attributes . . 427
23-32 WRKMEDIBRM: select libraries. 428
23-33 WRKMEDIBRM: confirm selected libraries 429
23-34 WRKMEDIBRM: restore command defaults 429
23-35 STRRCYBRM: print recovery plan. 430
23-36 SAP recovery plan, step 003 431
24-1 SAP production system with additional server. 436
24-2 SAP applications running in a System i cluster 437
24-3 Switchable and standalone IASPs. 438
24-4 Mirrored internal disks (RAID-1) 439
24-5 Mirrored external disks (RAID-1) 440
24-6 Mirrored external disks (Metro Mirror) 441
24-7 Logical replication of SAP 442
24-8 Remote journal environment for an SAP system. 442
24-9 With remote journaling. 443
24-10 Relational database entry for remote journaling 445
24-11 ADDRMTJRN: add remote journal for an SAP system environment. 446
24-12 WRKJRNA: activate remote journal. 446
24-13 CHGRMTJRN 447
24-14 SAP system environment running in the system ASP 448
24-15 SAP implementation with an Independent Auxiliary Storage Pool. 449
24-16 SAP system running in a switchable device environment 451
24-17 SAP applications running in cross-site mirroring environment. 452
24-18 DS storage server with Metro Mirror and System i copy services 453
25-1 Components of System i response time 459
25-2 SAP memory management system 463
25-3 SAP work process and shared ressources 464
25-4 Allocation sequence for dialog work processes. 469
25-5 Allocation sequence for non-dialog work processes 469
25-6 WRKSYSSTS (view 1). 471
25-7 WRKSYSSTS (view 2) - F11 transition data 472
25-8 WRKDSKSTS 473
25-9 WRKACTJOB (view 1). 474
25-10 WRKACTJOB (view 2). 475
25-11 Work with job locks 476
25-12 WRKSBSJOB 477
25-13 WRKSYSACT 478
25-14 Methods to start Data Collection Services. 479
25-15 Manage Data Collection Services with iSeries Navigator 480
25-16 Setup parameters of Data Collection Services 481
25-17 CHGPFRCOL 482
25-18 iSeries Navigator monitors. 483
25-19 iSeries Navigator - define a monitor. 484
25-20 SAP transaction ST06 485
25-21 CCMS in the SAP easy access menu 487
25-22 CCMS Monitor Sets (RZ20) - view CCMS Monitor Templates 488
25-23 Buffers - view current system status 489
25-24 Buffers - view open alerts 490
25-25 CCMS monitor legend 490
25-26 Buffers - single record - alert display 491

25-27 Buffers - single record - alert - detailed view 491
25-28 Buffers - current view 492
25-29 Solution Manager - Landscape Components 493
25-30 SAP Solution Manager as a system-wide central CCMS 494
25-31 Transaction SSAA - alert view 495
25-32 Global Work Process Overview (SM66) 496
25-33 SM66 - Work process selection 497
25-34 SM66 - Work process detail 497
25-35 Process Overview (SM50) 498
25-36 SM50 - Detail display 499
25-37 SM50 - Choose other tasks from Utilities 499
25-38 Example - Runtime analysis of a transaction 500
25-39 Choose a system log 500
25-40 System Log - Analysis (SM21) 500
25-41 System log - detail 501
25-42 User list (ST04) 501
25-43 Memory list (ST04) 502
25-44 Application Monitor (ST07) 502
25-45 Application Monitor - Response times 503
25-46 Application Monitor - Database accesses 503
25-47 Application Monitor - Database accesses - detailed 504
25-48 Application Monitor - SAP buffer 504
25-49 Workload Monitor (ST03) 506
25-50 ST03 - choose a mode 506
25-51 ST03 - save as initial view 506
25-52 ST03 - upper navigation area 507
25-53 ST03 - lower Navigation area 507
25-54 ST03 - dialog workload analysis by time 508
25-55 ST03 - memory usage statistic of all work processes 508
25-56 ST03 - Percentage Response Time distribution for all task types such as details . 509
25-57 ST03 - top list of dialog transactions such as details for the first transaction 509
25-58 ST03 - filters to generate a last minutes workload analysis report 510
25-59 ST03 - format your report - sort in descent order 510
25-60 ST03 - export function 511
25-61 Tune summary and buffers (ST02) 511
25-62 ST02 - detail analysis 512
25-63 ST02 - Current parameters 512
25-64 ST02 - buffer history 513
25-65 ST02 - Detail Analysis menu 514
25-66 Database performance: State of Disk (ST02) 517
25-67 DB02 - Deleted rows analysis 517
25-68 Database Performance Monitor (ST04) 519
25-69 ST04 - file activity 520
25-70 ST04 - Analyze database activity 520
25-71 ST04 - System catalog view 521
25-72 ST04 - System catalog SYSINDEXES 522
25-73 ST04 - Performance tables 522
25-74 Database monitor (DB4COCKPIT) 523
25-75 DB4COCKPIT - SQL statements 524
25-76 DB4COCKPIT - SQL statement detail 525
25-77 DB4COCKPIT - Used Indexes 525
25-78 DB4COCKPIT - Advised Indexes 526
25-79 DB4COCKPIT - Advised Indexes - detail 526

25-80 DB4COCKPIT - List of advised Indexes 527
25-81 DB4COCKPIT - Database Monitor Activity - Monitored Jobs and Last Analysis . . . 528
25-82 DB4COCKPIT - Database Monitor Configuration 528
25-83 DB4COCKPIT - Select your Java instance 529
25-84 DB4COCKPIT - Database Monitor for the Java instance 529
25-85 DB4COCKPIT - Table scans 530
26-1 Tuning steps in general 532
26-2 An approach to System i tuning. 533
26-3 Steps to tune an SAP system 539
26-4 SAP system response time 540
26-5 An approach to analyze database performance 542
26-6 Database tables SPLFI and SFLIGHT. 546
26-7 Report ZFLIGHT_SQLTEST1 results 546
26-8 SQL Trace (ST05) of ZFLIGHT_SQLTEST1 547
26-9 Query: record selection 548
26-10 Query: join condition 549
26-11 Query: Group By 550
26-12 Query: Order By. 550
26-13 SQL explain results 552
26-14 Use views to optimize a query 553
26-15 Query optimization: Star join (snowflake) schema. 553
26-16 Time line: use the SQL trace for ABAP 554
26-17 Session 2: Transaction ST05. 555
26-18 Session 2: Activate Trace 555
26-19 Activate Trace with filters: selection screen. 556
26-20 Session 2: Activate trace with Filter: status 556
26-21 Session 2: Deactivate trace: status 557
26-22 Session 2: selection for SQL Trace 557
26-23 Session 2: SQL Trace: an ABAP Trace list (view 1) 557
26-24 Session 2: SQL trace: an ABAP trace list (view 2) 558
26-25 Session 2: SQL Trace: an ABAP Trace list (view 3) 558
26-26 ABAP SQL Trace: Explain 559
26-27 SQL Trace: DDIC Information 560
26-28 DDIC Information for Index DD03L+1 560
26-29 Java SQL trace 561
26-30 Java SQL trace: output 562
27-1 SAP EarlyWatch Alert: system status 567
27-2 SAP EarlyWatch: trend of response time for the top five transactions 567
27-3 SAP EarlyWatch: history of ABAP dumps 568
27-4 IBM Insight: send data mail dialog window 571
27-5 IBM Insight: select geography for the send data e-mail. 571
27-6 IBM Insight: active users 572
27-7 IBM Insight: dialog steps 573
27-8 IBM Insight - Dialog steps versus CPU time by code type. 574
27-9 IBM Insight - Dialog Steps versus CPU Time by Module. 574
27-10 IBM Insight: top transaction by memory consumption 575
27-11 Shadow processes in a 3-tier installation. 576
27-12 WRKACTJOB 578
27-13 WRKSYSSTS 579
27-14 CHGSBSD 579
27-15 Pool Report 581
27-16 i5/OS: Change Class command. 582
27-17 PM iSeries homepage 584

27-18 Performance Tools main menu 585
27-19 System Report (view 1) 587
27-20 System Report (view 2) 588
27-21 System Report (view 3) 589
27-22 Performance data sets used by iDoctor. 591
27-23 Controlling interfaces for Collection Services database. 592
28-1 Extended star schema 599
28-2 Use of surrogate keys in SAP BW 600
28-3 Encoded vector index example for column COLOR 601
28-4 Star join example query 602
28-5 Example of query rewrite process to pushdown predicates. 603
28-6 Creating dynamic bitmaps 604
28-7 Visual Explain for statement using LPG. 605
28-8 ST05 output for statement using LPG 606
28-9 Visual Explain example 611
29-1 SAP system and other applications 623
29-2 RPG program T8189RP1. 625
29-3 DSPFFD: File T8189DB. 626
29-4 ABAP dictionary table Z8189DB 627
29-5 Output of ABAP program Z8189AB1 628
29-6 Transparent SAP database table Z8189DB. 629
29-7 DSPFD for file Z8189DB 630
29-8 WRKQRY: Work with Unicode tables 631
29-9 Accessing SAP dictionary tables from outside the SAP system 632
29-10 CRTSAPLF: Create Logical File for SAP Latin-1 or Unicode data 633
29-11 DSPFD for logical file T8189DB in library R3ERPLF. 634
29-12 DSPFFD for logical file Z8189DB 635
29-13 CRTSAPLF: Data access from RPG/400 636
29-14 CRTSAPLF: Source file for Z8189DB 637
29-15 CRTSAPVIEW for table Z8189DB. 638
29-16 CRTSAPVIEW for logical file Z8189DB. 639
29-17 Z8189DB1: View for RPG with fixed field length 640
29-18 CRTSAPVIEW: Data access from RPG/400 641
29-19 ABAP data access 642
29-20 SM30: Maintain table DBCON 644
29-21 Visual administrator: JDBC connector, register connection. 645
29-22 Visual administrator: JDBC connector, define SQL connection. 646
29-23 SM69: SAP transaction creating the i5/OS command 650
29-24 SM49: SAP transaction executing the i5/OS command. 651
29-25 SAP program output results. 652
29-26 SM36: Defining an SAP background job 653
29-27 SAP program defining external command job steps 654
29-28 SAP program defining external program job steps 655
29-29 SAP job starting after an event 656
29-30 Flowchart: ABAP program triggering an ILE program 658
29-31 Defining the function module attributes (Part 1 of 3) 660
29-32 Defining the function module import parameters (Part 2 of 3) 661
29-33 Defining the function module export parameter (Part 3 of 3). 662
29-34 Creating an RFC destination 663
29-35 Registering connection to ABAP program 664
29-36 SMGW: Registered RFC destinations T8189RFC. 665
29-37 Flowchart: RPG program RPG2ABAP. 680
29-38 SM37: Monitoring status of SAP jobs 688

30-1 Concept of data porting to an SAP system 690
30-2 Batch input technique 693
30-3 Data migration in three steps 696
30-4 How LSMW works 697
30-5 Transaction LSMW (Welcome Legacy System Migration Workbench) 699
30-6 Defining the project, subproject, and object 700
30-7 Project structure 701
30-8 LSMW project steps 702
30-9 Object attributes 703
30-10 Entering the structure for the existing system in an SAP application (source fields) 704
30-11 Defining conversion rules and field mapping 705
30-12 ABAP code for the generated conversion program 706
30-13 Interface for the input file (from the PC or from the Integrated File System) 707

Tables

6-1 GLS versus Unicode 38
7-1 Modifications to SAP Quicksizer output 45
8-1 Name and options of select System i licensed programs 72
8-2 IBM software products required for an SAP system (ABAP and Java) 74
8-3 Additional IBM software products useful for SAP system environment (ABAP) 75
8-4 Useful and optional software products for SAP systems 76
8-5 For SAP relevant i5/OS PTF groups 78
8-6 The SQLUTIL tool 80
8-7 i5/OS object management commands. 83
8-8 i5/OS Integrated File System commands 83
8-9 i5/OS system control commands 83
8-10 i5/OS system control menus 84
8-11 i5/OS communication configuration commands. 84
8-12 Priority classes for the SAP work processes 87
9-1 Conversion rates between System i objects and internal objects; limits of a journal 100
12-1 Additional libraries for kernel component patches 142
12-2 SAP kernel patch level for i5/OS and OS/400 releases. 148
13-1 SAP instance profile parameters for an i5/OS implementation 174
13-2 Typical SAP instance profile parameters for i5/OS installations 176
14-1 Links for the global transport directory in the DEV system before relocation 186
14-2 Links for the global transport directory in the PRD system after relocation 186
14-3 Transport files created on release and export of change request 200
14-4 The tp individual or unconditional modes. 208
15-1 Access method C or U. 238
16-1 General database administration tasks 270
16-2 CCMS specific database administration tasks. 270
16-3 Example daily administration checklist for System i configurations 282
16-4 Daily SAP system administration checklist (example) 282
17-1 SAProuter options 297
17-2 Additional saprouter options 297
17-3 System i niping test 298
19-1 Info APARs for i5/OS and OS/400. 333
20-1 Jobs, job type, and status 341
20-2 SAP notes for SAPOSCOL 360
20-3 General information to be collected 361
22-1 Hardware availability solutions. 376
22-2 System i libraries included in the SAP product 384
22-3 When to save an object 387
22-4 What to save in the Integrated File System 387
22-5 Recommended backup methods 388
23-1 Methods of saving the database 393
23-2 SAVDLTRCV parameters 407
23-3 When to use what recovery procedure 416
25-1 SAP parameter settings on System i models. 467
25-2 System i monitors 470
26-1 Guidelines for long-time CPU utilization 534
26-2 Guidelines for short-time CPU utilization 534
26-3 Fault rate factor of response time guideline 536

26-4 Guidelines for disk space 537
26-5 Guidelines for disk arm utilization - long time 537
26-6 Guidelines for disk arm utilization - short time / peak 537
26-7 Guidelines for disk IOP utilization - long time 538
26-8 Guidelines for disk IOP utilization short time / peak. 538
26-9 Combine the columns of an SQL query 551
26-10 Possible key columns for table SPLFI and table SFLIGHT 551
28-1 Recommended QAQQINI file options for star join support 607
28-2 Instance profile settings for star join support 607
28-3 Display statistics data column descriptions 609
28-4 Primary SAP notes related to SAP BW on System i configurations 616
28-5 Other SAP notes related to SAP BW on System i configurations 616
28-6 SAP BW table naming convention. 618
29-1 SAP RFC communication call functions 658

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

eserver®
Redbooks (logo) ™
eServer™
iSeries™
i5/OS®
pSeries®
xSeries®
z/OS®
z/VM®
zSeries®
Advanced Function Presentation™
Advanced Function Printing™
AFP™
AIX 5L™
AIX®
AS/400e™
AS/400®
C/400®
COBOL/400®
CUA®
Domino®
DB2 Universal Database™
DB2®
Electronic Service Agent™
Enterprise Storage Server®
FlashCopy®
Infoprint®
Integrated Language Environment®
Intelligent Printer Data Stream™
IBM®
IPDS™
Language Environment®
Lotus®
MQSeries®
NetServer™
OS/390®
OS/400®
PartnerWorld®
PowerPC®
POWER™
POWER4™
POWER5™
POWER5+™
Redbooks™
RETAIN®
RPG/400®
System i™
System i5™
System z™
System/38™
SQL/400®
Tivoli®
TotalStorage®
WebSphere®

The following terms are trademarks of other companies:

SAP Solution Manager, SAP BW, SAP Business Information Warehouse, SAP NetWeaver Business Intelligence, mySAP Supply Chain Management, mySAP Supplier Relationship Management, mySAP Product Lifecycle Management, mySAP ERP, mySAP CRM, mySAP Customer Relationship Management, mySAP Business Suite, BAPI, ABAP, SAP ArchiveLink, SAP EarlyWatch, SAP GoingLive, SAP NetWeaver, mySAP, mySAP.com, SAP R/3, and SAP are the trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Enterprise JavaBeans, EJB, Java, Java Naming and Directory Interface, JavaBeans, JDBC, JDK, JMX, JVM, J2EE, Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows server, Windows NT, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Pentium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

IBM® i5/OS® and IBM System i™ technology is a proven platform for SAP applications. This combination has more than 10 years of success across more than 2500 SAP installations worldwide in small, medium, and large enterprises. System i models are an ideal platform for SAP customers who are looking for easy usability, high performance, reliability, and carefree operation of their SAP applications. The unique value of System i lies in its ability to reduce information technology complexity and to simplify SAP system landscapes.

This IBM Redbook features a collection of knowledge gained from IBM and SAP system experts who work with customers that use SAP applications on IBM System i models with i5/OS. It is written to assist SAP basis consultants and other IT professionals in implementing a total business solution, consisting of System i servers, i5/OS Version 5 Release 4 (V5R4), DB2® UDB for i5/OS database, and SAP products.

The ideas within this redbook apply to the whole spectrum of business solutions within the mySAP Business Suite. These solutions share the common technical application platform SAP NetWeaver that includes the SAP Web Application Server and SAP Business Information Warehouse.

This book first provides an overview of SAP applications and landscape options when running on i5/OS. After discussing SAP and i5/OS implementation and infrastructure topics, it provides information regarding operational topics, such as administration, printing, and how to avoid problems. Next comes a detailed discussion about availability and backup and recovery, which is followed by a section dedicated to performance topics. The book ends by looking at different ways to exchange data between your SAP applications and other applications.

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.

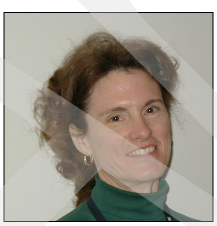

Susan Powers is a Consulting IT Specialist at the ITSO, Rochester Center. Prior to joining the ITSO in 1997, she was an AS/400® Technical Advocate in the IBM Support Center with a variety of communications, performance, and work management assignments. Her IBM career began as a Program Support Representative and Systems Engineer in Des Moines, Iowa. She holds a degree in Mathematics, with an emphasis in Education, from St. Mary's College of Notre Dame. She is a Project Manager for complex projects in the ITSO.

Matthias Bartusch is an IT Specialist for SAP on IBM eServer™ iSeries™ at IBM Germany. He holds a graduate degree in Engineering from the Technical University of Dresden. He has 15 years experience with IBM as an AS/400 System Engineer for midrange market customers, as well as consulting and implementing support for SAP R/3 Basis on the AS/400 server. Matthias currently provides technical presales support for customers, business partners and the IBM Sales Team in Germany, Austria and Switzerland.

Manfred Engelbart is an IT Specialist at the IBM eBusiness Technical Sales Support Team in the Integrated Marketing Team (IMT) of Germany providing technical presales support for the SAP on iSeries solution. He has worked at IBM for 32 years, has 18 years of experience with iSeries, and has coauthored IBM Redbooks™ and Redpapers about SAP and i2 implementation on iSeries systems.

Dieter Huber is an SAP/IBM Solution Architect and Senior IT Consultant from basycs GmbH, Germany. He started his company as CEO in 2002 with the intention to increase the SAP on System i market. His areas of expertise include 11 years of development, implementation and consultancy with SAP on iSeries, and also on UNIX® and Windows® with different databases. Dieter holds a degree of Computer Science and Economics from the University of Karlsruhe in Germany, and he combines theoretical knowledge with practical experience.

Gert Ruland is from IBM Germany. He has 28 years of IT experience, including five years with SAP products. He has been a member of the IBM SAP International Competence Center since 2000. He holds a degree in Mathematics from the University of Münster in Germany. His areas of expertise include i/OS, z/OS®, relational database management systems, and sizing. He has coauthored three Redbooks on a variety of subjects. Gert is a Certified SAP R/3 Basis Consultant.

Sabine Schloesser is working as a Solution Sales Specialist with the responsibility for SAP on System i sales, in Germany. She joined IBM in 1984, starting with System /38 and, from the very beginning, AS/400. For the last 10 years, she has worked with customers in sizing and design of SAP system landscapes.

Ron Schmerbauch is the Technical Leader of the SAP on System i Team, as part of i5/OS Development in Rochester. He has 12 years of experience with SAP on System i, beginning with the original porting effort of SAP R/3 to OS/400®. Before joining the SAP Team, Ron worked in Database development. He holds a degree in Computer Science from the University of Southern California.

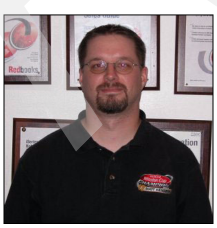

Steve J. Tlusty is a member of the SAP on System i Team, as part of i5/OS Development in Rochester. He has five years of experience with SAP Business Intelligence. Before joining the SAP Team, Steve worked in IBM Global Services developing business intelligence solutions for retail and manufacturing customers.

Thanks to the following people for their contributions to this project:

Eric Kass
Christoph Langer
Barbara Roth
Edelgard Schittko
Sabine Schloesser
Marco Wermann
Anja Winkler
IBM Germany

Michael J. Frost
Kolby Hoelzle
IBM Rochester

Joachim Kern
Thomas Obermeier
Sally Power
Axel Stein
Bernhard Wolf
SAP

Jürgen Zeh
basycs GmbH, Germany

Erica Wazewski
Linda Robinson
IBM ITSO

A special thank you to the following people for their thorough review and continual support of this ITSO residency, people, project, and redbook:

Christian Bartels
IBM Germany

Dorothea Rink
SAP

Volker Gueldenpfennig
consolut.gmbh, Germany

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Part 1

SAP applications on System i models with i5/OS

The chapters in Part 1 discuss the value which IBM System i servers and the integrated i5/OS operating system provide for running an SAP application. If you are familiar with the integrated approach of the IBM System i concepts, you also will understand the value of integration that this platform brings for the SAP system environment. If you are new to System i platforms, you are soon to discover what features make the System i models unique for SAP applications.

1

IBM System i solution for SAP

As a business owner, you need to run your business and you need your information technology (IT) system to work, every day. An i5/OS and IBM System i solution can help you meet these needs. One key factors that differentiates the System i platform from other platforms is the level of hardware and software integration. For example, in a UNIX environment, you must select software components from different vendors (operating system, database software, system management software, and so on) and integrate them to build a working environment for an SAP system.The System i architecture takes a simpler integrated approach, as shown in Figure 1-1.

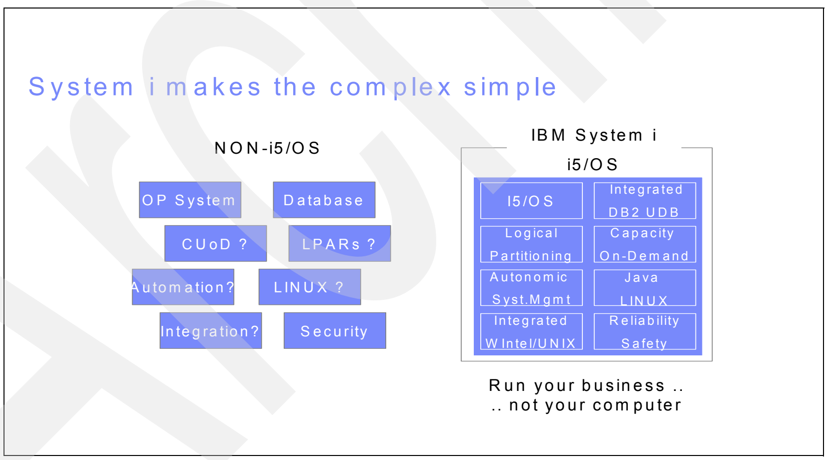

Figure 1-1 IBM System i - An integrated system

The System i family is a fully integrated IT solution that runs its own i5/OS and is based on the IBM POWER5+™ processor technology. An integrated DB2 UDB database is part of System i. With i5/OS, this DB2 UDB offers automatic functions for system management, such as fully automated disk management and workload management. System i can help you simplify your SAP system landscapes, by allowing them to operate with a small staff in a user-friendly way.

1.1 The integrated approach

IBM i5/OS is fully integrated and highly secure right out of the box. This means that relational database, communication and networking capabilities, online help, Web enablement technologies, easy enterprise management and much more are fully integrated into the operating system as an entire whole system. The user communicates with all components of i5/OS using a single command language or a simple unified graphical user interface (GUI).

i5/OS helps you simplify your IT environment with support for Linux®, Windows Server®, Java™ and UNIX applications within a single, easy-to-manage system. It combines high availability with superior application workload management and logical partitioning (LPAR). The next generation of applications can be quickly deployed and managed in a single, partitioned system alongside current business applications.

This chapter describes how these features ease the effort to implement and run SAP applications on the IBM System i server. Visit the System i Web site for further information:

http://www-03.ibm.com/systems/i/

1.2 DB2 UDB for i5/OS: Integrated database for System i family, optimized for SAP

Many System i users are not aware that they have a full relational database system integrated in their system. Consider the following:

- The System i server is shipped with the i5/OS operating system installed. The database system is integrated in the operating system.
- Storage management is done by i5/OS and provides a solid basis for running an SAP application.
- Data can be accessed with Structured Query Language (SQL for short). This is the primary access method used by SAP applications. Database access is also possible via read and write methods as used by core applications.
- DB2 UDB does not have to compete with SAP applications for i5/OS and hardware resources because it is part of i5/OS. Other databases are just additional applications on top of Unix or Windows operating systems, so they compete with SAP applications for operating system resources.

The integration of the database into the operating system is one of the key advantages and distinguishing characteristics of the System i platform compared to other platforms. Besides delivering the complete functionality of a database system, the integrated database:

- Complies fully to the latest SQL standard 2003
- Is accessible from all types of applications, from core business RPG code to new Java development. These applications can reside in the same operating system image or on other systems.
- Easily allows you to install and administer multiple SAP database schemes under one i5/OS operating system. This also easily allows you to install and run multiple SAP applications in one i5/OS partition.
- Is trusted by i5/OS and has extra opportunities for resiliency down to the hardware components. Since other databases are applications, the operating system underneath them can terminate them if the operating system believes the database application is not behaving correctly.

The i5/OS database implementation adheres to the System i philosophy of ease of use. Consider the following factors:

- The database optimizer collects the information it needs automatically
- Storage management is done by the operating system
- Database administration is minimal, as compared to what is required on other platforms.

The integrated database is a key advantage of i5/OS. See 8.2, "i5/OS object structure" on page 55.

1.3 Storage management

Managing storage can be one of the most demanding tasks to maintain good system performance. Managing storage can involve:

- The placement of the data on disk, to avoid hot spots
- The reorganization of the data to insure sequential read performance and reclamation of space
- Providing sufficient free space for data volume growth

i5/OS on the System i platform performs each of these tasks. Once a collection of disks is defined to host the data (called an Auxiliary Storage Pool (ASP)), i5/OS and the System i server do the rest. The data is evenly spread over the storage pool to optimize disk access times. This helps to eliminate hot spots on the disk devices. The architecture of the storage layer of System i servers ensures that reorganization is done continuously.

Storage is further described in Chapter 9, "SAP storage and database considerations" on page 89.

1.4 Consolidate homogeneous workloads under i5/OS

Under i5/OS you can easily run and administer multiple SAP systems in parallel. Each system requires one or more instance. Each instance is started as a subsystem containing multiple i5/OS jobs. Each job represents an SAP work process to process either Advanced Business Application Programming (ABAP) or Java programs. i5/OS knows the interdependencies and presents them to the operator. All processes belonging to a subsystem can be easily identified.

You can customize each subsystem with different attributes. For example, you can specify the CPU priority for jobs within the subsystem and you can specify a priority for allocation of main memory to this subsystem. These are ways in which you can choose to manage and control each subsystem according to your service levels. Performance data can be reported on a subsystem level.

It is easy to run and manage multiple SAP components (for example Enterprise Resource Planning (ERP), Customer Relationship Management (CRM) or Enterprise Portal (EP)) under one i5/OS image, and common to exploit this option. Refer to Chapter 3, "SAP system landscapes" on page 17 and Chapter 4, "SAP system landscape scenarios" on page 23 to read more about running multiple instances of SAP systems on the System i server. Also see the SAP notes for this topic, including 416994.

1.5 Consolidate heterogeneous workloads on a System i server

A System i server can host several operating systems in parallel by the virtue of logical partitions (LPAR). Each LPAR is independent of the other LPARs and can host i5/OS, AIX® or Linux on POWER™. You can also use the System i server as a host or file server for systems running Windows.

1.5.1 Heterogeneous operating systems

You can have a variety of applications running in your environment, for example, several SAP applications, WebSphere® applications, Domino® and core business (existing) applications. The applications can require different operating systems as hosts.

The functionality of System i servers to support different operating systems on one server provides the means to consolidate several different applications onto one system. SAP Advanced Planning and Optimizing (APO) may be another example for taking advantage of the System i multiple operating system approach. You can run the SAP APO database in one LPAR under i5/OS and run SAP application server and Live Cache in another LPAR under AIX or Linux on Power.

1.5.2 Heterogeneous workloads

Many SAP components have the workload characteristics of online transaction processing (OLTP) systems, like SAP ERP, where transactions are quick and data accesses are usually operating on only a few rows in the database. However, there are other SAP components like SAP Business Intelligence that have workload characteristics of online analytical processing (OLAP) systems, where transactions can run for long durations because of very complex queries involving millions of rows in the database.

Customers have successfully run mySAP Business Information Warehouse (BW) along with other SAP components within the same partition, especially in development and test partitions. We recommend that you run multiple LPARs on your System i configuration to isolate the OLTP and OLAP workloads from each other, especially in production systems.

1.5.3 Heterogeneous PTF levels

The virtualization options of System i models allow you to consolidate applications which may require different PTF levels under the same i5/OS operating system. In this case you can install multiple operating system environments with different PTF levels in multiple LPARs. For example, you can run your SAP production environment on one LPAR and the development and test environment on another one. Both LPARs can contain multiple SAP components with homogeneous applications, as described in 1.4, “Consolidate homogeneous workloads under i5/OS” on page 5.

Virtualization options are discussed in more detail in Chapter 5, “System i virtualization” on page 31.

1.6 IBM System i value proposition for SAP

IBM System i models support the entire mySAP Business Suite and SAP NetWeaver ´04s. It offers a variety of benefits to SAP customers.

These benefits are:

- System i is easy and simple

  Due to System i integrated operating system and database technology and its features, SAP customers require only a small staff to manage even complex landscapes.

- System i servers are powerful, scalable and reliable

  The latest IBM 64-bit POWER5™ + processor technology makes the System i server even more powerful. It offers smooth options for 2-tier and 3-tier SAP system landscapes. Single System i servers can scale from 1 to 64 processors. Its availability is nearly 100%.

- System i is flexible and versatile

  System i leading IT virtualization functions allow highly efficient operations. Less hardware is required to run any given SAP application. System i multi-platform support allows customers to run concurrent operations of i5/OS, AIX and Linux on one server, hence simplifying their IT-landscape.

- System i models are strategic and proven

  A successful more than 10-year history and excellent customer references prove the System i architecture to be solid and reliable. The entire mySAP Business Suite and SAP NetWeaver ´04s are supported.

- System i can be operated carefree

  The i5/OS operating system and its integrated DB2 UDB for i5/OS database allows SAP customers to run a highly automated SAP management system with high flexibility and availability. DB2 offers the latest database technology with superior performance proven by leading SAP BW benchmarking at no extra cost.

- System i servers are optimized for SAP applications

  Each and every release of the i5/OS operating system and its integrated DB2 UDB database since OS/400 V3R6 in 1996 have included numerous features specifically to meet the demands of SAP applications.

- System i servers are cost effective

  Exclusively for SAP customers the Solution Edition for mySAP ERP offers great performance at low cost.

The following reports can help you understand the total cost of ownership and value of System i implementations:

- A report from International Data Corporation (IDC) outlining the cost savings and enablement benefits AIX 5L™, Linux and Windows Integration in System i implementations:

  `http://www.ibm.com/servers/eserver/iseries/idcroi/`

- A consultant report which explains the value delivered by technology when balancing management and acquisition costs:

  `ftp://ftp.software.ibm.com/common/ssi/rep_wh/n/ISL02177USEN/ISL02177USEN.PDF`

For more information about IBM System i servers and i5/OS concepts refer to Chapter 8, “i5/OS from an SAP application view” on page 49.

2

SAP Solution Portfolio

SAP is one of the world's largest software vendors. mySAP Business Suite with all its components is well suited for companies of all sizes and industries. SAP started their success within large enterprises, but mid market business has become more and more important.

Although SAP applications offer a broad spectrum of functionality which address all of a company's needs, applications of different vendors can be combined.

For collaboration purposes between SAP applications or to external non-SAP applications, the Exchange Infrastructure (XI) component of mySAP Process Integration (mySAP PI) can be used. IBM WebSphere Business Integrator or other third party vendor products are also options.

2.1 History

The success story of the modern SAP company began with providing SAP R/3 as an ERP application for all kinds of companies. To provide platform independency, basic components are separated from the commercial application.

Figure 2-1 shows a history of mySAP Business Suite:

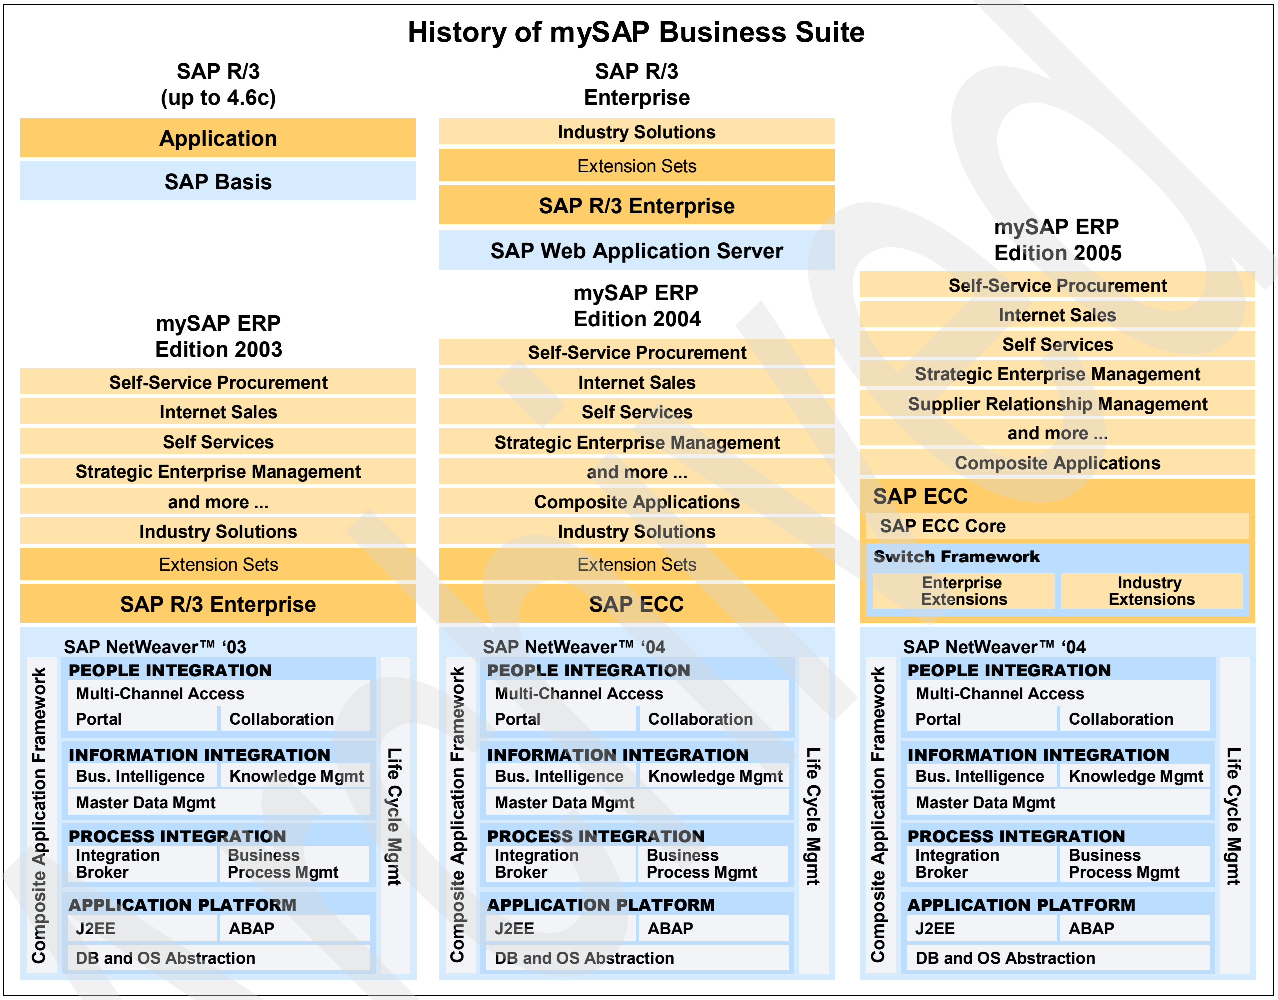

Figure 2-1 History of mySAP Business Suite

Over time, the application part of SAP R/3 was enhanced with additional functionality like Business Warehousing capabilities or Customer Relationship Management (CRM) functions, becoming an all-embracing enterprise system, which today is called mySAP Business Suite.

On the way to SAP NetWeaver the basic components have been extended with all sorts of integrating functionality, as described in this Redbook. Whenever reasonable, functionality is transferred from the application layer to SAP NetWeaver components.

2.2 mySAP Business Suite

mySAP Business Suite is a family of adaptive business applications, providing “best-of-breed” functionality, based on the SAP NetWeaver platform.

Figure 2-2 shows a view of the mySAP Business Suite:

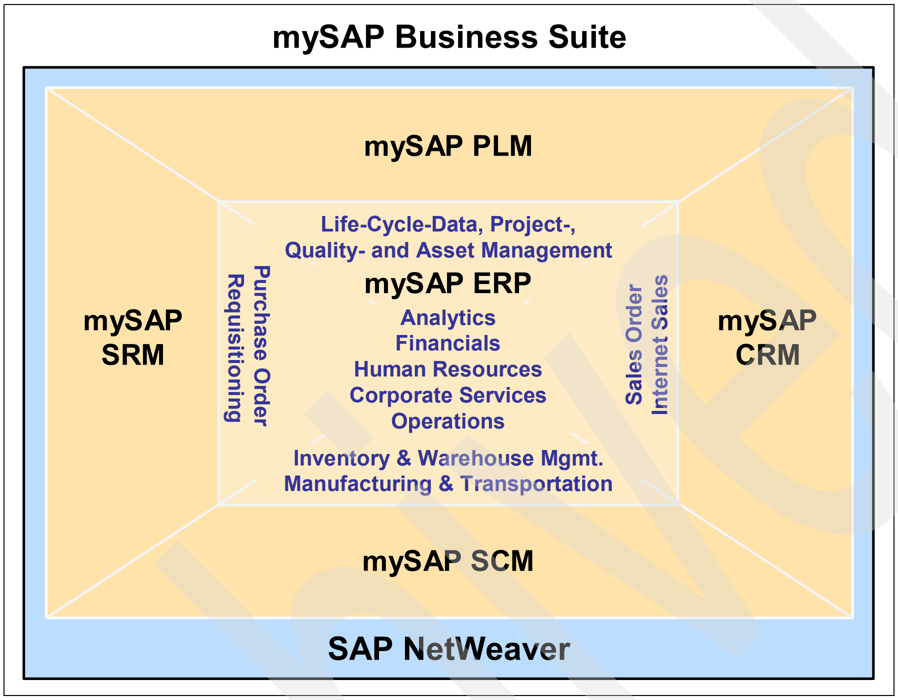

Figure 2-2 mySAP Business Suite

From an application point of view, mySAP Business Suite consists of Enterprise Resource Planning (ERP) Central Components (ECC) and Business Suite components:

- mySAP ERP
- mySAP Customer Relationship Management (CRM)
- mySAP Product Lifecycle Management (PLM)
- mySAP Supplier Relationship Management (SRM)
- mySAP Supply Chain Management (SCM)

This portfolio is complemented by industry-specific solutions which also run platform-independent. For System i businesses, IS-Retail, IS-Health and IS-Utilities are of major importance.

All-in-One solutions address specific mid-market business demands with customized solutions, for example for the automotive, pharmacy or beverage business. Nevertheless, All-in-One implementations also provide the full flexibility of SAPs application functionality to respond quickly to whatever upcoming changes. These applications are mainly provided by SAP business partners.

Additional third party software is available for specific business requirements.

SAP ECC cannot be installed individually. Nevertheless, for each subcomponent patches may be provided separately, if needed. For more information about subcomponents of SAP ECC or industry solution add-ons inside SAP ECC, refer to mySAP ERP Master Guide and SAP note 700778.

In typical scenarios, a combination of above mentioned application components are used. Those landscapes consist of more than one SAP ECC installation, each with its own database.

As Java becomes more important and is used for new developments, it can happen that both ABAP and Java components are included. In such a case, at least two databases are installed.

To keep these environments manageable, SAP recommends using SAP Solution Manager. SAP Solution Manager is defined as “a centralized, robust solution management toolset that facilitates technical support for distributed systems, with functionality that covers all key aspects of solution deployment, operation and continuous improvement”. SAP Solution Manager keeps track of your patch levels. It will become more critical as landscapes become more complex.

2.3 SAP NetWeaver architecture

One of the main objectives of SAP NetWeaver is providing a business-focused logic infrastructure that supports continuous business process evolution and change with minimal risk and cost when introducing new, innovative business processes as existing systems remain unchanged. It also represents the technical enabling layer of SAP Enterprise Services Architecture (ESA), defined by SAP as the basically service-oriented architecture (SOA) for adaptive business solutions. Neither ESA nor SOA are discussed in this redbook.

The release of SAP NetWeaver referenced in this redbook is SAP NetWeaver ´04s.

Figure 2-3 shows an overview of SAP NetWeaver.

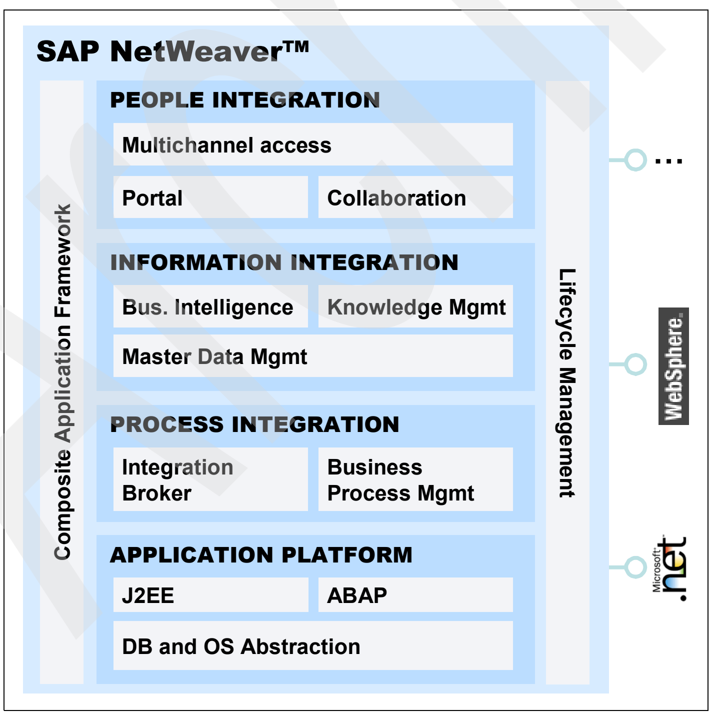

Figure 2-3 SAP NetWeaver Architecture

This chapter outlines the major parts of SAP NetWeaver. More details can be found in SAP publications.

The main parts are:

- People Integration
- Information Integration
- Process Integration
- Application Platform

There are other components framing these major parts, like composite application framework and life cycle management.

Although SAP NetWeaver provides an integrated framework, many components may be replaced or enriched with third party products. IBM WebSphere business integrator for information exchange is one example. IBM WebSphere portal or Lotus® Domino for people integration is another one.

2.3.1 People integration

The main goal of people integration is to allow them to be connected, enable collaboration, and provide information every time there is a need.

The portal infrastructure delivers unified, personalized and role based user access to the information technology environment. Collaboration promotes dynamic communication within teams or communities.

Multichannel access provides connection to enterprise systems through voice or mobile infrastructure.

2.3.2 Information integration

Integration of unstructured and structured information is key for customer success. There needs to be a way between providing data in a correct and timely manner without overloading these data warehouses.

Knowledge Management (KM) provides a single access point to SAP content management system and third party repositories.

Business Intelligence (BI) enables customers to integrate, analyze and disseminate relevant and timely information.

Master Data Management (MDM) enables customers to store and consolidate master data, while ensuring consistent distribution to all applications and systems within the IT landscape.

2.3.3 Process integration

Within process integration, business process management provides capabilities to model and drive processes in a dynamic IT environment, for example in using workflow techniques.

Integration broker capabilities enables XML/SOAP based communication between application components from various sources and vendors. This is provided by SAP exchange infrastructure (XI, now named PI), which provides a technical infrastructure for XML-based message exchange in order to connect SAP components with each other, as well as with non-SAP components.

2.3.4 Application integration

2.3.1, “People integration” on page 13 to 2.3.3, “Process integration” on page 13 reflects more on business enabling techniques, while application integration is the major basic technical layer for application development and database interfacing.

Besides ABAP as the development environment, Java is the upcoming development environment.

2.4 Complexity of SAP NetWeaver implementation

For future SAP NetWeaver implementations, Java will become an important development environment. There is a steady transformation between ABAP and Java developed components. As those environments become more and more complex, the System i architecture is well prepared to help reduce complexity.

Figure 2-4 shows a more technical picture of SAP NetWeaver implementation.

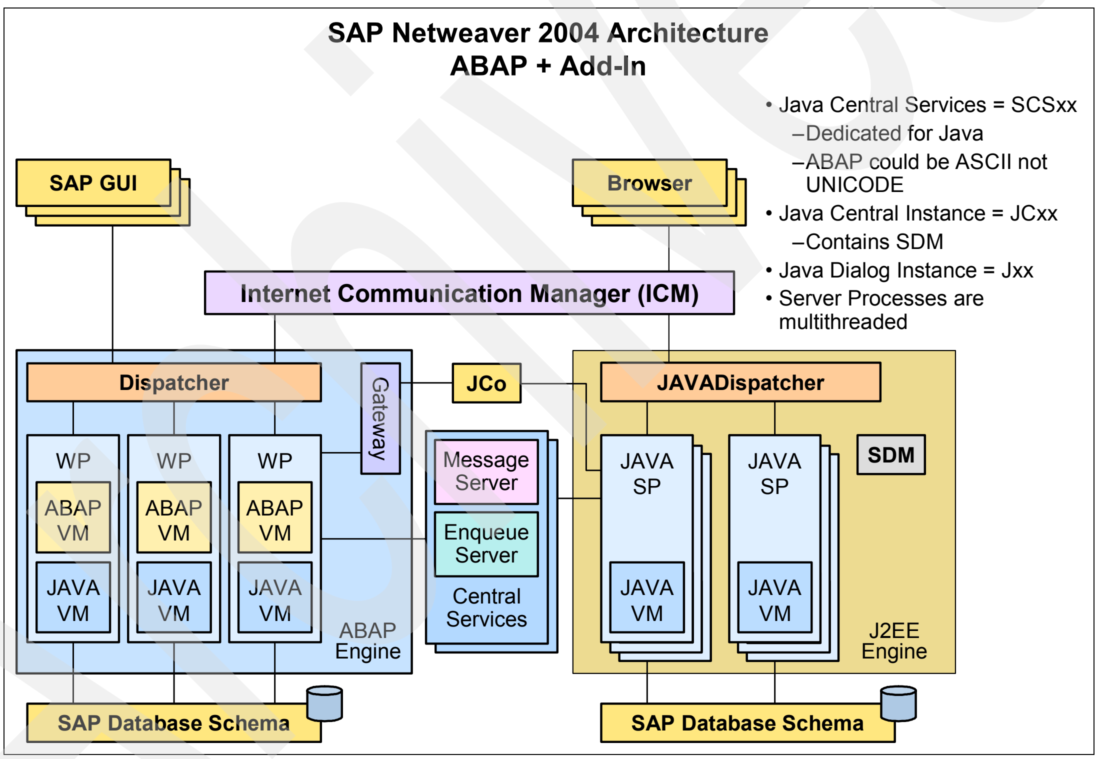

Figure 2-4 SAP NetWeaver Architecture, ABAP and Java

As shown in Figure 2-4, every development environment, ABAP or Java, has its own database schema. So in an SAP system where both ABAP and Java components are included, you find two database schemas.

2.5 Critical success factors

To run a landscape of SAP applications, you need to consider the following management challenges. Some are easy to solve, others can require more management effort. IBM System i architecture with the integrated capabilities of i5/OS provides several advantages compared to our competition. This redbook provides more information in the following chapters.

The following list can give you some guidance regarding what to consider in your design and planning work. It is not intended to be a complete list, nor may all points apply to all environments:

- Integrity and data resiliency of multiple databases
  - High Availability (HA) aspects
  - Log and data integrity
  - Backup and recovery
- Multiple SAP systems and instances
  - Prioritization of jobs and application environments
  - Performance optimization
  - Flexibility of multiple workload
- Management of heterogeneous workloads
  - ABAP versus Java performance
  - OLTP versus OLAP processing.

This redbook provides answers to these critical issues, showing the IBM System i value proposition for customer environments running SAP applications.

3

SAP system landscapes

To better understand the complexity of the SAP system landscape, consider that there are several dimensions involved. First, a single SAP system consists of one or more servers in a LAN. Multiple users access these servers with presentation clients. Next, there are likely to be development and test SAP systems so that changes can be created and implemented safely before being moved to the SAP production system. Finally, remember that SAP has multiple solution offerings. Most SAP system landscapes use more than one component part of SAP NetWeaver ´04. The combined view of these dimensions shows that an SAP system landscape can easily consist of multiple servers and multiple SAP systems that are interconnected. Figure 3-1 on page 18 shows these dimensions working together.

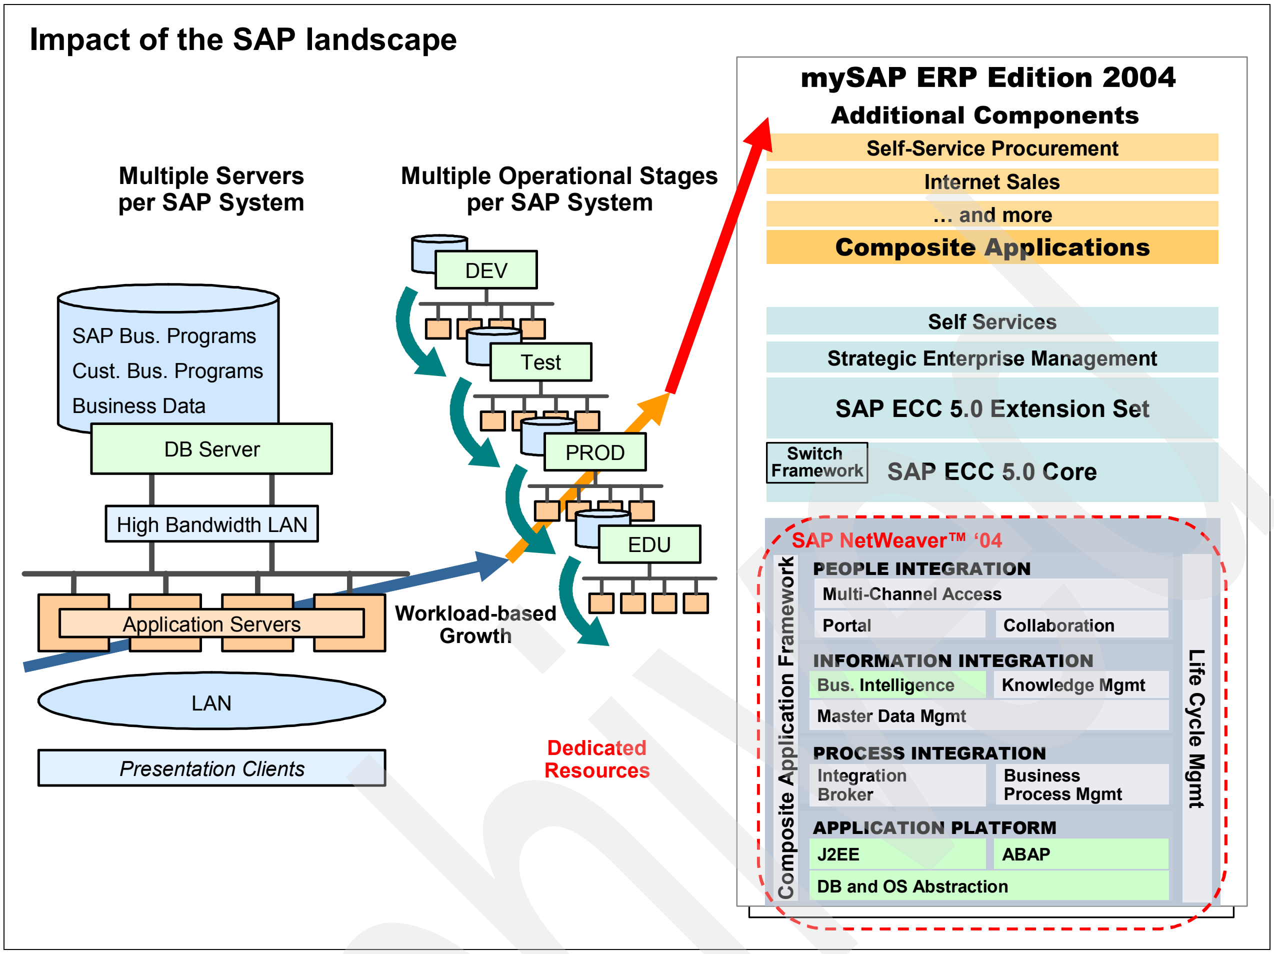

Figure 3-1 The complexity of SAP system landscapes

There are three dimensions that you need to take care of in a typical SAP system landscape:

- SAP Architectural dimension (layered structure)
- SAP Operational dimension (change management)
- SAP System dimension (horizontal view)

These attributes are discussed in the sections that follow.

3.1 SAP system architectural dimension

A SAP system consists of three main layers (Figure 3-2):

- **Presentation**: This layer is the frontend for user input and is passed to the next layer for processing. It is possible to attach all kind of frontend to an SAP system like SAPGUI, Web-Browser, PDA and other mobile devices.
- **Application**: In the application layer, the calculations and evaluations are performed. Incoming data is processed and passed to the database layer. No data is stored here.
- **Database**: All data is stored only in the database layer. Data exchange between application processes and the database is proceeded through an SQL interface.

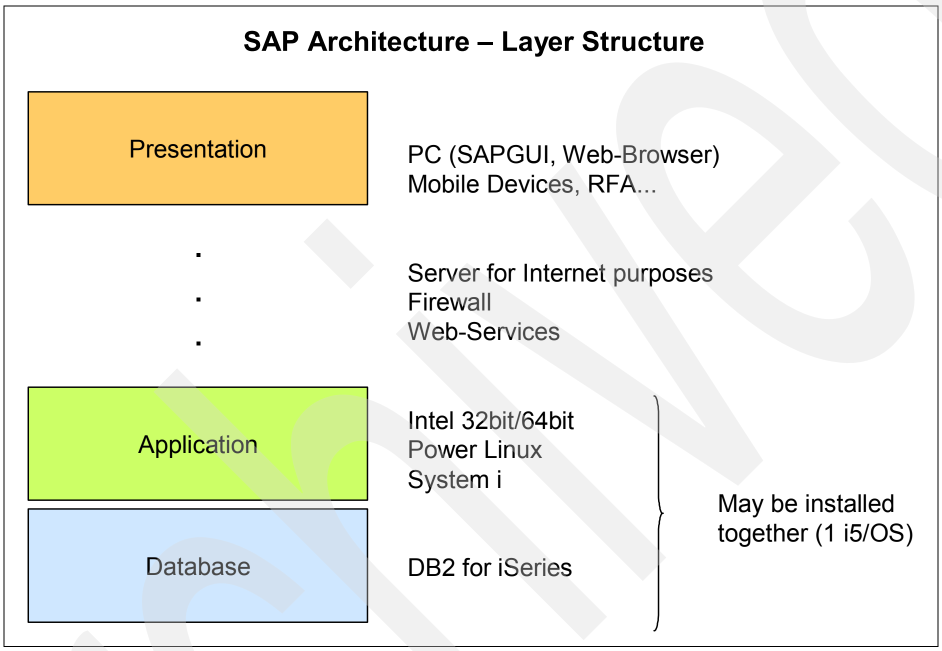

Figure 3-2 SAP system architecture layers

The Application layer and Database layer can run together in one operating system image, which then is named a 2-tier, opposed to having it separated onto different hardware or LPAR, called a 3-tier implementation. To keep it simple, we reflect on 2-tier and 3-tier implementations. All other layers (tiers) such as Web services, firewall and others are ignored, as those concepts look relatively similar for any IBM Business Partner solution.

You need to determine whether 2-tier or 3-tier is the better implementation choice for your situation. There are various reasons outlined in the next sections that describe why you may decide to choose one implementation instead of the other.

3.1.1 2-tier implementation

In many implementations, the application and database servers run together within one i5/OS partition. This is called a 2-tier or central installation because you have only two layers to use, Application and Database together as the first layer, and Presentation as the second.

For many companies, this represents the best kind of implementation, because:

- POWER5 scalability to 32 processors, including i5/OS Capacity Upgrade on Demand for high workload phases or an easy upgrade
- Best performance behavior (direct interface between application and database)

- Lowest administration costs
- Low Total Cost of Ownership (TCO)

3.1.2 3-tier implementation

3-tier implementation means to run application and database servers separately as the first two tiers, with the presentation layer being the third. This solution can be chosen because of:

- Flexible scalability in peak workloads
- Scaling requirements beyond 32 processors
- High availability solutions only for the database part can have price advantages
- Application requirements, where optional niche middleware components can require non-i5/OS implementations (for example Search and Classification Engine (TREX) in Customer Relationship Management (CRM) or LiveCache in Supply Chain Management (SCM) environments)
- Low Total Cost of Acquisition (TCA) with Windows/Linux application servers.

3.2 Operational stages of SAP systems (change management)

Besides the application architecture, an SAP system landscape consists of several SAP systems provided for different kind of usage. Each SAP system is installed separately, including separate libraries for application logic and database files.

At minimum, the following SAP systems are recommended:

- Development System (DEV): In this system, the application is customized with the business needs. All ABAP/Java development is done. In most DEV systems, only a small amount of data is used. Changes are passed to the Quality Assurance System (QAS) using the SAP transport-system
- Quality Assurance System: This system is used for validation of all changes which have been implemented in the DEV system. Changes should only be done in DEV, never in QAS. In many cases, QAS data is a copy of the production data, so even seldom used constellations can be tested. After successful validation, changes are transported to the production system. In some cases, QAS-system is also used for education purposes.
- Production System (PRD): This is the SAP system environment for running the productive application. For a smaller business, DEV and QAS may run integrated in one SAP system.

Beside this, there is no limit in using additional SAP systems, such as for:

- Sandbox
- Training
- Validation
  (often needed in pharmaceutical or health care environments to fulfill legal requirements).

3.3 SAP system dimension

Nearly all business needs of a company can be addressed by an SAP application. That does not mean that only one SAP system with its DEV and QAS need to be installed. Most landscapes consist of several independent SAP installations for different systems.

As discussed in Chapter 2, “SAP Solution Portfolio” on page 9, it is common to run a separate SAP system for each of those different applications of mySAP Business suite. For small environments, some functionality can be mixed and used in one SAP installation. The customer requirements should be discussed to determine which implementation fits best.

Keep in mind that sometimes it is hard to change the environment, so if some functionality of CRM is included and used in ERP it may be very challenging to separate this afterwards into two separate SAP systems. Also keep growing requirements in mind.

3.4 Naming conventions

For a better understanding of the contents of this book, some often-used names are introduced in this chapter. Figure 3-3 shows the major components of an SAP system from a technical point of view:

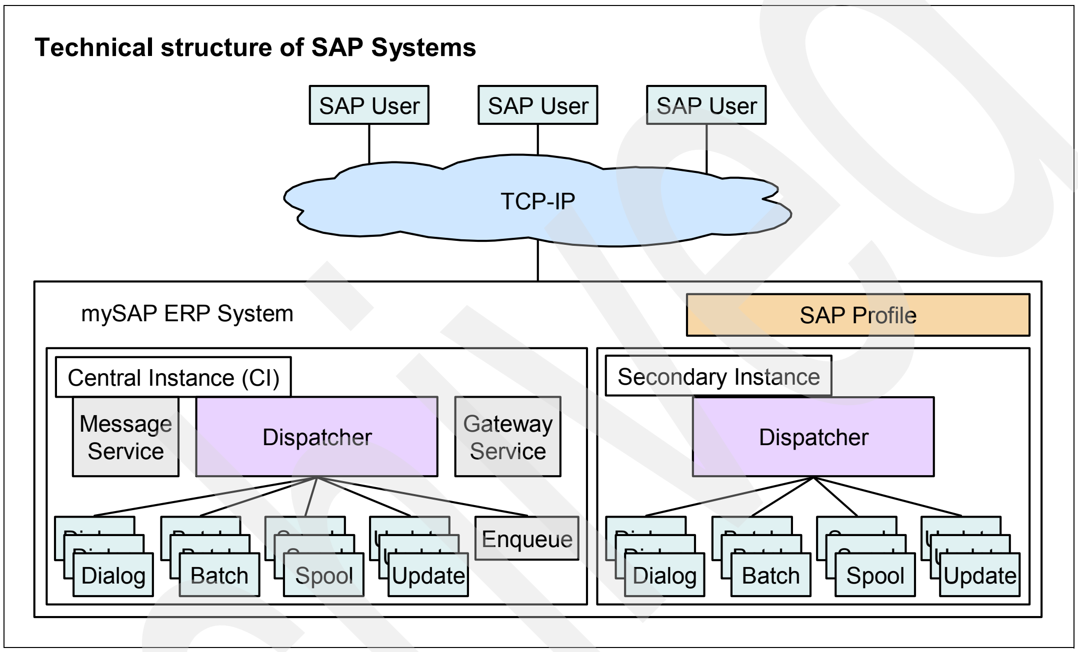

Figure 3-3 Technical structure of SAP systems

While installing and running an SAP system environment, you should be familiar with the names of objects and their usage:

- SAP system:

  An SAP system is a logical entity consistent of database and kernel. The kernel contains the executable code. The database includes the SAP database and the ABAP program modules. One SAP System is serviced by one or more instances. Multiple SAP systems can be installed on a single System i server using a unique three-character system identifier (<SID>), for example, D01, Q01 or P01.

  A system must contain at least one database. If ABAP and Java is used both in an SAP system, two databases are installed.

- SAP Instance:

  An SAP instance is a collection of processes and resources within a single SAP system. An instance provides resources to end-user-requests. An instance is connected to only one SAP database. Therefore, if using ABAP and Java together in one SAP system, at least two instances are installed. The characteristics of an instance are specified by its instance profile.

  With the System i platform, 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. It is started with the STARTSAP command with specifying the SAP system and instance number to be started.

Each SAP system has exactly one central instance with the only occurrence of a message service and an enqueue service. Only the central instance has complete information of all SAP system resources and their status.

Besides a central instance, there can be several dialog instances, which are dependent on the central instance in usage of message and enqueue services.

In a 3-tier environment, in minimum the central instance needs to be installed at the database server. Other instances can be separated onto several application servers. It is not possible to spread one instance over more than one application server, regardless if those are separate boxes or LPARs.

- SAP instance profile:

  The SAP instance profile contains the characteristics of an instance. It is stored in the integrated file system in flat file format and used whenever an SAP instance is started. For example, it contains how many work processes are started or initial buffer dimensions.

- Work process:

  For the System i work management, SAP work processes are jobs within the R3_nn subsystem that actually perform the work.

  Message and enqueue services are running only in the central instance. The enqueue work process is a special job that is responsible for handling the special locking requirements of SAP. There can be more than one enqueue work process if necessary. The message work process is specialized for the communication between instances within the same SAP system and identifies services in the SAP system to external jobs. There is only one message work process for each SAP system.

  Every instance consists of one dispatcher work process. The dispatcher controls the type of work that is performed by a work process. It assigns end-user requests and units of work to the work process of an instance for execution. Those work processes are:

  – Dialog
  – Batch
  – Update
  – Spooling

Although it is not too important for technical implementation, for planning reasons it is important to know the following terms:

- Client (also known as Mandant)

  A client is a way to organize and isolate data in an SAP system. For example, a separation by organization unit may be helpful for some companies. In small companies, development and quality assurance may be separated only by client, not by a separate SAP system.

  The SAP system consists of client-dependent and client-independent data. Client-dependent data applies only to a specific client (for example application data). Client-independent data applies to all clients (for example transaction codes).

- Session

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

- Mode

  Modes correspond to the nesting of displays within a window. Using several modes or sessions may have impact on main memory capacity planning.

4

SAP system landscape scenarios

Over time, the business requirements of most companies continues to grow. This results in growing SAP applications and perhaps adding non-SAP system landscapes. From an SAP view, most of these requirements result in a set of multiple SAP systems.

This discussion is based on our business experience and illustrated with landscape scenarios which are typical of many customer situations. Although every customer has individual requirements, we have categorized the requirements into three types of landscapes – small, medium and large, with an increasing complexity from one example to the next. In this brief discussion we introduce how an IBM System i implementation can reduce complexity and increase manageability in such an environment.

4.1 Example environment descriptions

To run an environment of SAP applications, we assume an installation of at least one SAP system for production and one for development/quality assurance. As SAP highly recommends the use of SAP Solution Manager, this SAP application is implemented in each scenario, regardless of the complexity.

The SAP applications are not described in the scenario, as they are our idea of a typical scenario and the SAP applications can be exchanged with individual customer requirements.

We intentionally keep these examples as simple as possible. In particular, we do not take into consideration any high availability requirement, which can significantly influence the landscape configuration. You can find information about how the concepts of implementing high availability in Chapter 22, “Availability, backup and recovery concepts” on page 373.

4.2 Small configuration scenario

This example may be valid for small customer scenarios using SAP applications for a low number of users, or for customers starting small with growth requirements.

4.2.1 Components

This example contains the following SAP applications:

- mySAP ERP including small BW reports

  BW requirements are of minor complexity, mainly standard predefined reports with small amount of data. Implementation is practical with SAP ECC including SAP BW functionality.

- SAP Solution Manager

Environments for each component are:

- Production
- Development and Quality Assurance in one SAP system

4.2.2 Schema

All components and environments are running under one i5/OS system, as shown in Figure 4-1 on page 25:

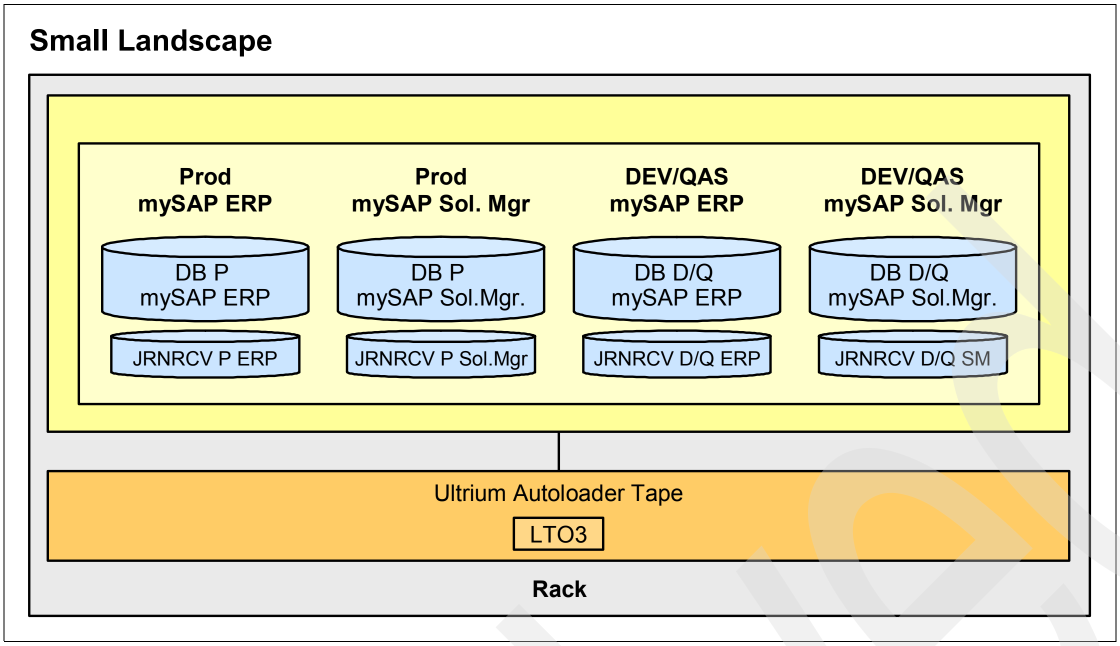

Figure 4-1 Small landscape example

We have four separate subsystems containing the four different SAP systems. Production subsystems can be privileged by adjusting the priority (the class parameter).

Data of all SAP systems is installed in separate libraries on the same disk environment. Therefore, full optimization capabilities of i5/OS are available. Each SAP system has its own journal management operating. Data backup and recovery is executed by one Linear Tape-Open (LTO) drive, managed with i5/OS native commands or tools like Backup and Recovery Media Services (BRMS).

4.3 Medium configuration scenario

Many customers run a medium configuration or are facing challenges as described in this section.

4.3.1 Components

Our example for medium sized customers contains the following SAP applications:

- mySAP ERP
- mySAP BW as a a separate SAP system with more complex queries
- mySAP CRM (small business configuration), for example Customer Interaction Center (CIC)
- mySAP Solution Manager

The environments for each solution component are:

- Production
- Dev
- QAS
- Dev/QAS integrated in one SAP system for mySAP Solution Manager

4.3.2 Schema

The production environment is separated from the Dev/QAS environments using logical partition (LPAR) technology. Also, mySAP BW is batch oriented. Therefore, the mySAP application is separated from dialog oriented applications such as mySAP ERP and mySAP CRM, by additional partitions. mySAP Solution Manager runs in the same partition as mySAP ERP. The solution example is shown in Figure 4-2.

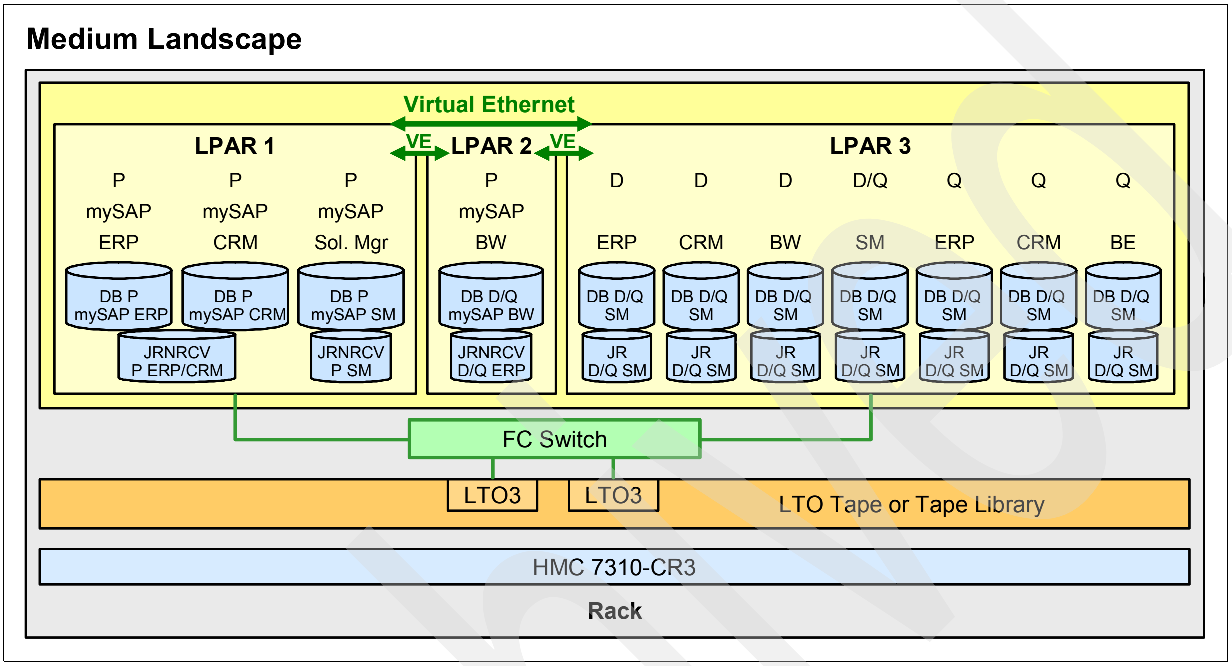

Figure 4-2 Medium landscape example

There are three partitions in this medium landscape example. The first partition is a production environment for mySAP ERP, mySAP CRM and mySAP Solution Manager, separated by subsystems. The second partition runs mySAP BW as the only application. Unless otherwise determined by customer requirements, the third partition runs all development and QAS environments.

In each LPAR, the data of the assigned SAP system is stored in separate libraries on the same disk environment. For data consistency purposes, it can be helpful to combine the journal management on dependent solutions. Therefore, journals for mySAP ERP database can run in combination with the journals of mySAP CRM. In the situation of backup, data recovery results in both databases (libraries) running on a consistent logical level.

Data backup and recovery in this example is accomplished by one Fibre channel attached IBM LTO drive. We recommend IBM Backup Recovery and Media Services (BRMS) software to backup the data. BRMS is the strategic IBM backup and recovery product for i5/OS.

Full advantage of System i virtualization can be used in this LPAR environment, including a dynamic resource assignment between LPARs. By using the capped and uncapped methodology and a shared processor pool, processor utilization can be optimized. Virtual Ethernet connections provide the fastest TCP/IP connections possible between System i partitions.

All configurations are kept in one server footprint. Other solutions are possible.

4.4 Large configuration scenario

More and more companies are growing into large complex scenarios. Such a scenario can consist of many components. One example is shown in this section.

4.4.1 Components

Our example is for a large customer environment containing the following SAP system components:

- mySAP ERP
- mySAP BW
- mySAP CRM including Search and Classification (TREX)
- Portal Solution (for example, mySAP Enterprise Portal or IBM WebSphere Portal)
- mySAP Process Integration (Exchange Infrastructure) or an adequate solution (for example, IBM WebSphere Business Integrator)
- mySAP Solution Manager

The environments for each solution component are:

- Production
- Dev
- QAS
- Dev/QAS together in one SAP system for mySAP Solution Manager

4.4.2 Schema

Although you can run such a scenario in one server footprint. In many larger customer cases, the applications are divided and installed on two System i servers. This concept allows you to implement a high availability or disaster recovery concept as well.

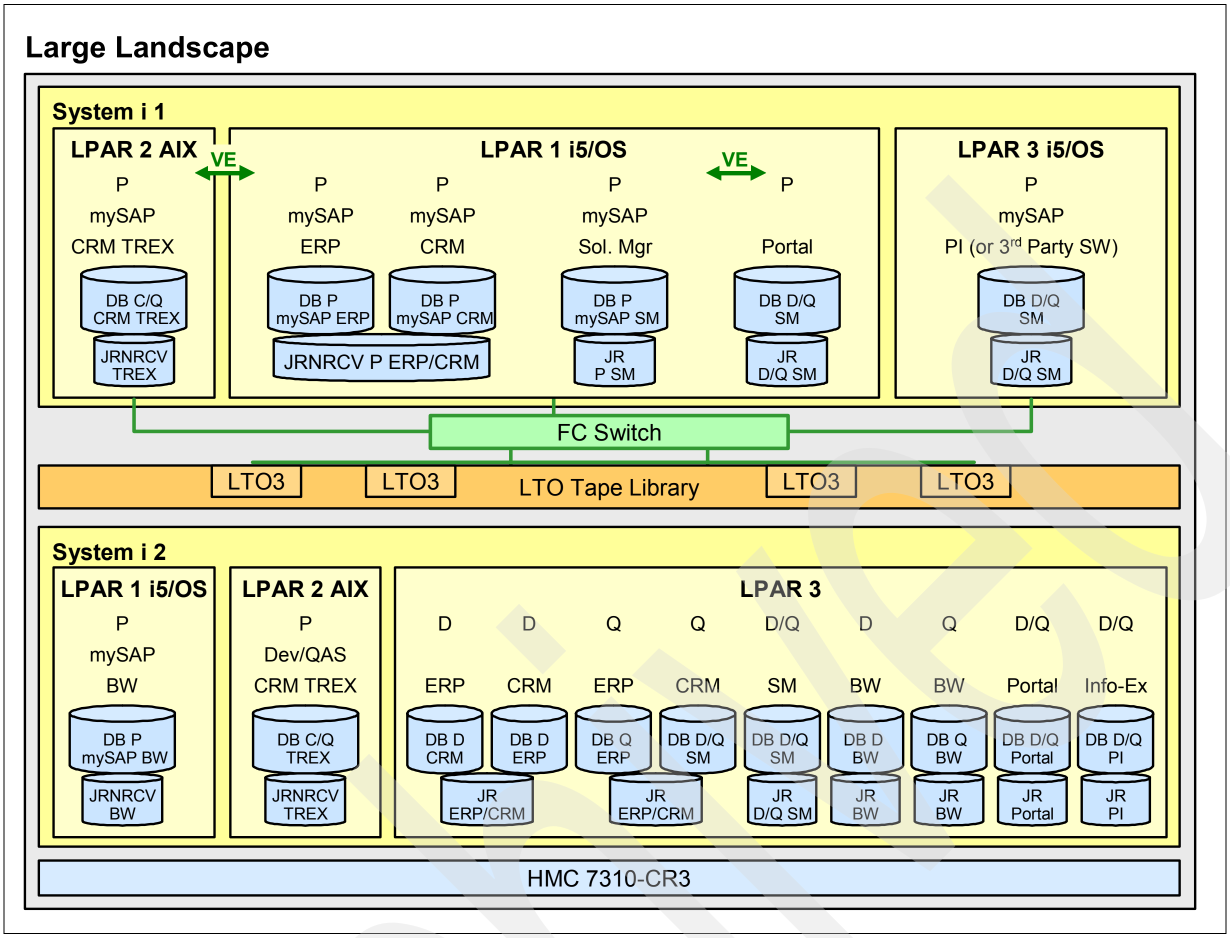

Figure 4-3 Large landscape example

For each System i server, System i LPAR capabilities are used to realize all of the advantages described in 4.3.2, “Schema” on page 26. Also included is a component which is not running natively on i5/OS in this example. A strength of System i architecture is that a requirement for an AIX or Linux partitions can be easily met since both of these operating systems are supported on System i models.

The components are distributed such both hardware systems are of the same size.

4.5 Availability considerations and influence on landscape design

IBM System i models offer numerous possibilities to increase data resiliency and application availability. Before those concepts can be designed, customers need to define (and discuss) their requirements. There needs to be a decision regarding high availability or disaster recovery requirements, which are quite different goals even if they look similar on a first view. There needs to be a decision whether a solution nearer to the hardware based on cluster technology is favored (like cross site mirroring (XSM) or storage area network (SAN) based) or a more software based solution is needed like third party vendor solutions for shadow databases.

In many cases, regardless of the solution, a second System i server is used. Our recommendation is to run development and QAS on this second server with a separate partition for availability purposes. During normal business, this backup partition can be run with a minimum amount of CPU and memory. Development and quality assurance partition benefit from these resources. In the backup case, resources are shipped from the DEV and QAS partition to the backup partition. Either the users working in development or QAS end their work, run on a minimum configuration, or operate with lesser performance, as defined in the design phase.

Refer to Chapter 22, “Availability, backup and recovery concepts” on page 373 and Chapter 24, “Disaster recovery and high availability” on page 433 for more information about high availability and disaster recovery.

4.6 System i value proposition for SAP NetWeaver landscapes

The examples in this chapter show how the integrated concept of IBM System i architecture helps to simplify SAP NetWeaver landscapes, to reduce administration costs, and to optimize the Total Cost of Ownership (TCO) of your SAP system. The major benefit of a System i implementation with SAP applications are:

- Reduced amount of physical servers and an optimal use of the installed hardware by POWER 5 virtualization tools (LPAR and Capacity on Demand).
- Reduced number of operating system environments due to the System i capability to run and administer multiple and even heterogeneous workload under one i5/OS operating system.
- Ease-of-use administration for SAP systems running in i5/OS
- Reduced administration effort to run multiple database schemes under one i5/OS image and to optimize and administer SAP database environments
- Multiple operating system integrated functions for backup and recovery, and high availability.

In this redbook you find more information about how these advantages are implemented in a System i configuration and how SAP customers benefit from this.

5

System i virtualization

The fundamental concept of virtualization is a logical partition or LPAR. The advantage of virtualization is resource sharing. The latest developments in micro partitions provide the most flexible virtualization solution for users of System i servers.

This chapter provides an overview of how virtualization features are implemented on System i servers and the options to exploit these features when SAP applications are implemented on System i configurations. For further information, refer to the iSeries Information Center:

http://www.ibm.com/eserver/iseries/infocenter

5.1 Virtualization

Virtualization as a concept in the information technology (IT) industry that serves to exploit resources to a higher degree. This contributes to a lower cost of installation. For example, if you have several servers running different operating systems with an average utilization of 20%, occasionally each of the servers run at a higher utilization. The system runs at a higher utilization when servers are combined in one hardware system. The peak utilization of the combined utilization curve of the servers determines the size of the hosting server.

To implement such a concept, the ability to run several operating system images in parallel is required. We call this the collection of resources, such as processor resources, memory, and I/O connectors, which are used to serve one of these operating system images a logical partition. A controller, called a hypervisor, manages the logical partitions, similar to how an operating system manages jobs.

> **Definition:** A logical partition (LPAR) is a part of a computer system that can run an operating system. Processor, memory and I/O connections are allocated to a partition. A logical partition can be defined without restrictions posed by hardware boundaries as, for example, processor cards or memory banks.

The flexibility of partitions is illustrated with the facilities described in this section. A simple server consolidation scenario is used to show the differences of each concept.

5.1.1 Static LPAR

The basic concept of a partition is as a static LPAR. An LPAR is configured with a fixed amount of processor, with a fixed amount of memory allocated to this LPAR, and specific adapter cards for connecting I/O devices. To change the configuration, the impacted partitions must be shut down, the available resources redefined, and partitions restarted.

Different operating systems can run in parallel, however the degree of sharing is limited with this configuration. The number of processors required is the same as when you run the operating system on different servers.

5.1.2 Dynamic LPAR

A dynamic LPAR (dLPAR) can have only a fraction of a processor allocated. Memory can be dynamically deallocated from one partition and allocated to another partition. This can be done without a need to restart the partitions. If you want LPARs to share physical processors, you need the concept of a virtual processor.

On the other side, an LPAR can have several virtual processors defined. The hypervisor must assign the available physical resources for this partition.

The sharing of real processors (server consolidation) means fewer processors are required to run the workload compared to static LPARS.

> **Note:** IBM POWER5+ technology provides the concept of simultaneous multithreading (SMT) which provides two threads per processor. Do not confuse this with virtual processors.

An operating system views a thread as a logical processor. This means one real processor is seen by an operating system as two logical processors. A virtual processor is seen by the operating system as a "real" processor on which it can run. With SMT there are two virtual logical processors.

5.1.3 Micro partitions

The latest development in virtualization is micro partitions. Micro partitions are logical partitions that can dynamically shrink and expand. Each micro partition has CPU resources defined in processing units. One processing unit represents one processor. The granularity of a processing unit is a 1/100 of a processor. A minimum of 1/10 of a processor is assigned to a micro partition.

For each 1.00 processing unit you have to define at least one virtual processor. For a virtual processor you have to define at least 0.10 processing unit. The theoretical maximum of processing units that can be allocated to a micro partition is equal to the number of virtual processors defined.

The following properties must be assigned for each micro partition:

- The minimum, desired and maximum processing units
- The minimum, desired and maximum number of virtual processors

Capped and uncapped mode

If the desired processing unit for a micro partition is equal to the maximum processing units, the partition runs in capped mode. The partition never gets more processing units allocated, even if there are processing resources available. On the other hand, if the system runs at 100% utilization, the partition has the desired processing units.

If a partition does not run in capped mode, it can be assigned more resources than desired. During initialization the partition is assigned an entitlement. The entitlement of a shared partition is the granted share that the partition gets if the system runs at 100% utilization. If the partition requires more resources than its entitlement and there are idle processing resources, a proportion of the free resources are allocated to this partition.

Granularity of dispatching

If a partition is eligible for dispatching, a time slice is allocated for its processing. This time slice is a fair share of the available resources.

Each virtual processor gets its proportion of the time slice. The partition does not have to use its time slice completely.

Effect on server consolidation

Implementing server consolidation with micro partitions can yield the highest saving in CPU capacity.

Further information can be found here:

- *Logical Partitions on the IBM PowerPC®*, SG24-8000
- iSeries Information Center at:

  http://publib.boulder.ibm.com/iseries/

5.2 Implementation options

The option to run operating system images on one system in parallel is enabled with partitions. Several SAP systems can be run under one i5/OS image.

This section discusses when to use the LPAR approach and when to implement several SAP systems in one partition.

- Different operating systems

  Different operating systems can run on the same System i server with the use of LPARs. Some SAP applications require features which are not available natively under i5/OS, for example with Supply Chain Management (SCM). SCM requires livecache, which is available under Linux on Power. In this case, the LPAR technique provides a simple implementation option.

  Consider this approach for 3-tier implementations with application servers running under Linux on Power.

- Running several i5/OS images

  SAP recommends implementing a test system and a quality assurance system implemented for each production system. If the test and quality assurance systems are in one LPAR and the production system in another LPAR, follow the same approach SAP recommends for the operating system. Install fixes for i5/OS first in the test LPAR, test, then implement the fixes in the production system. This contributes to a more stable production system.

  If the installation is medium to large, separate the OLTP systems (Enterprise Resource Planning (ERP) and the OLAP systems (BW and SEM) into different partitions. These systems have such a different resource requirement pattern, which is hard to control, if they are running in one operating system image.

- Running everything in one i5/OS image

  This is the recommended approach for a majority of implementations.

5.3 Summary

There are several features available to exploit virtualization for SAP applications on System i configurations:

- The majority of SAP application implementations on System i servers are small to medium in size. Using virtualization features can be, but is not necessarily, advantageous in these implementations. These customers might want to implement their SAP systems in one i5/OS image.
- Consider separate development and production systems for larger SAP implementations.
- Separate the production systems into separate LPARs if the SAP implementation is large and includes OLTP and OLAP systems.
- Run different operating systems on one server if parts of the application require different operating systems as hosts.

See *Logical Partitions on the IBM PowerPC*, SG24-8000, for further information about System i partitions.

6

Encoding data used in SAP systems

This chapter discusses the concepts of encoding data used in Systems, Applications and Products in Data Processing (SAP) systems. It also provides some information about the differences between two encoded versions that are available for many SAP applications: the Global Language Solution (GLS) version and the Unicode version, and when to use them.

6.1 What are encoding schemes?

One byte of data can represent up to 256 different combinations of data representations. When used to represent languages, each combination is assigned to a specific character in that language. However, when you take into consideration the many different characters used in the languages of the world, you can quickly realize that one byte is not sufficient to represent all of them, let along multiple languages. Many different encoding schemes have been created to take a single byte and assign each combination to a character.

Since there are different encoding schemes, they conflict with each other. Two schemes may use a different combination to represent the same character and the same character can be represented by two different combinations. The encoding scheme must be known when recording and reading data so that the correct characters are understood.

Most SAP systems use a single encoding scheme consistently throughout the database and application server. However, there are types of SAP systems known as multi-data multi-processing (MDMP) systems, in which a blended encoding scheme attempts to use more than one encoding scheme within the same system. The negative to this approach is that to read a piece of data correctly, you need to know which scheme was used when it was recorded. If you attempt to read a piece of data with the other scheme, misinterpretation of each character is possible, which can lead to many types of problems.

If two bytes are used to represent a character, there are enough combinations to cover all of the languages that SAP systems support. This type of universal encoding scheme is called Unicode. A detailed technical explanation of character representations can be found in “A technical description of GLS and Unicode” on page 327.

6.2 SAP direction

Now that all SAP applications are available in Unicode, SAP is clearly supports Unicode as opposed to MDMP. In a letter to customers dated April of 2006, SAP states that “New releases of SAP NetWeaver and SAP applications based on SAP NetWeaver, which are released in 2007 or later, will no longer support new installations of Non-Unicode systems. Non-Unicode upgrades and system copies of Single Codepage systems will still be possible.” Based on this statement, we recommend that all new customer installations be done using the Unicode version of the SAP application.

The letter further states that “new SAP product releases as of 2007 will be available for 64-bit server operating systems only, including Windows and Linux.” This 64-bit requirement is not an issue for System i configurations, as POWER5 is the tenth generation of 64-bit processors for System i models.

For more information regarding SAP support of Unicode, MDMP systems and 64-bit servers, see:

`http://service.sap.com/platforms`

6.3 What encoding schemes are available for SAP applications on i5/OS?

With SAP NetWeaver ´04 media from SAP, two versions are included. One is Unicode, and the other is called the GLS.

6.3.1 Unicode

Also known as “full Unicode”, Unicode is available for all SAP NetWeaver components, including SAP ECC 5.0 and all other applications built with the 6.x kernel. Some SAP applications are only available in Unicode.

In a Unicode SAP application on i5/OS systems, both the application server and database use Unicode (UCS2) encoding. UCS2 is one of several implementations of Unicode. With Unicode, additional memory is required for the application infrastructure because data within the application server is now represented with two bytes per character, as compared to single byte systems which represents data within the application server with one byte per character.

A Unicode solution offers the most flexibility for using multiple languages in a single system. It can use i5/OS, Windows or Linux on POWER as application server platforms connected to a DB2 database on i5/OS systems.

6.3.2 GLS

Prior to SAP application availability in Unicode, SAP applications were made available on i5/OS systems in a GLS version. The GLS version originated with Basis release 4.6C and is available for Enterprise Central Component (ECC) 5.0 and most other current applications that use the 6.x kernel.

The GLS version features an ASCII application server with Unicode (UCS-2) database files. It is sometimes referred to as the “ASCII version”. It offers several advantages when compared to the original SAP implementations on i5/OS systems which are based on EBCDIC encoding.

GLS is compatible with certain ASCII-based SAP components such as the SAP LiveCache. Like Unicode it can use i5/OS, Windows or Linux on POWER as application server platforms connected to a DB2 database on i5/OS systems. One other benefit of GLS is an increased capability to handle different language combinations, including Asian languages. It can also support MDMP, but again, MDMP is no longer recommended. If you are using MDMP systems, we recommend that you convert to Unicode as quickly as possible to minimize the effort that is required to resolve potential data conflicts that are caused by MDMP.

6.4 Choosing between Unicode and GLS

Before determining whether to install a new system with Unicode or GLS, consider the following factors:

- Unicode provides all of the benefits of GLS, and more.
- Certain SAP applications are only available in Unicode, including all of the Java-based SAP applications
- Running ABAP in GLS and Java in UNICODE within the same SAP system is technically possible, however not recommended, because there can be a loss of information if text data is transferred from Unicode (Java) to the GLS (ABAP) system.
- If you do not foresee requirements for multiple languages or SAP Java applications, GLS systems should be sufficient
- GLS systems offer one tangible benefit. The application server infrastructure is single-byte, so less memory is required as compared to a Unicode system. This results in a savings in hardware investment and marginally better performance.

Table 6-1 summarizes these concepts.

Table 6-1 GLS versus Unicode

Criteria	GLS	Unicode
Multiple languages	Limited	Recommended
ABAP and Java in same system	Possible, not recommended	Recommended
Resource required	Less	More

6.5 Converting an installed system to Unicode

If you already have an SAP system installed, you are still supported on your current code page. SAP's direction indicates that eventually you need to convert the application to Unicode when you upgrade to a release that is no longer available in anything but Unicode.

If you are moving to the System i platform from another platform, we recommend that you choose a Unicode solution rather than choosing GLS, then moving to Unicode eventually.

See 18.7, "Migration and code page conversion" on page 324 for a more technical discussion of conversion and migration details.

6.5.1 Moving from GLS to Unicode

Switching from a GLS system to a Unicode system can be a simple migration depending on the languages involved, because the database tables in a GLS system are already primarily in Unicode. A few fields need some internal format conversion that is handled by a tool that is available from SAP. Some validation should be performed after the conversion, but this is a straightforward process.

If you have implemented the MDMP option, a Unicode conversion can require a significant amount of manual effort, depending on the languages used, because multiple characters can be mapped to the same single byte combination. You need to determine which character was intended in each case. Data validation is critical when moving from an MDMP system.

6.5.2 Moving from EBCDIC to Unicode or GLS

SAP applications on i5/OS systems are delivered with a single byte EBCDIC implementation through SAP Basis release 4.6D. SAP applications based on the 6.x and later Web Application Server based releases are available only in the GLS version or in a Unicode version. They are not available in EBCDIC. If you presently use an EBCDIC based solution and want to upgrade to a current SAP release level, a Codepage Conversion (CPC) needs to be included in the process.

When you perform the SAP release upgrade, you upgrade to release level 4.6C first, then perform the conversion to GLS. From 4.6C on GLS, you can continue to release upgrade to the desired release, then complete with a simple conversion to Unicode.

The resulting database is approximately 50-70% larger than the original database since a large number of fields are converted from single byte (EBCDIC) type to a double byte (Unicode UCS2) type. Note that the extra capacity requirement does not cause an extra disk expense line of 50-70% because higher capacity disk units are usually only marginally more expensive than lower capacity drives, and that a single new disk unit of larger capacity might be able to replace the capacity of multiple older disk units if extra or replacement disk units are needed.

You can choose to remain on GLS if you see no future requirements for any SAP Java systems or for multiple languages, or just because you want to minimize the hardware resources required. However, we recommend that you move completely to Unicode during your upgrade project because of SAPs stated direction for the future and so that you may be in the best position should your company requirements suddenly change.

Be aware that the sorting sequence for EBCDIC is different from GLS or Unicode. This can require application changes if you have created ordering dependencies.

Since MDMP was never allowed in EBCDIC systems and is no longer allowed in GLS systems, the complications of MDMP systems should not be a concern for anyone converting from EBCDIC to GLS or Unicode.

7

SAP sizing considerations

A system configuration must be sized properly in order to perform well. An overview of the process involved in providing a sizing estimate, terminology, and considerations are described in this chapter.

The factors influencing performance of a System i and SAP system are discussed in Part 4, “Optimizing performance” on page 455.

7.1 Overview of the sizing process

Determining the correct configuration of a platform for an SAP system installation is not a trivial task. Sizing means to determine the hardware requirements of an SAP system. A system is sized in terms of processor speed, memory, disk space, and configuration design, such as the:

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

There are two dimensions involved in determining the size of a system relative to system performance:

- **Throughput**: What capacity is required to deliver a certain amount of work in a given period of time.
- **Response time**: What capacity is required to return an answer, or information, in the shortest possible amount of time.

SAP uses the throughput approach.

Analyze each planned SAP system with respect to the following sizing factors:

- Individual customer requirements
- Total number of users
- Number of users active at any one time
- Number of transactions per time period
- SAP modules to implement
- Customization required
- Production system
- How to implement a test system or environment
- How to implement a development system or environment

SAP provides information and recommendations to size a system. An SAP Quicksizer tool is available to help in sizing an SAP system. A special user ID (S-userid) is required for access to the sizing tool. Customers and Business Partners of SAP are entitled to get such a user ID. IBM employees should send a request to ISICC@de.ibm.com to get a S-userid assigned.

You can find the SAP Quicksizer on the Marketplace of SAP at:

http://service.sap.com/sizing

Guidelines and the SAP Quicksizer tool provide a capacity estimation expressed in terms of an *SAP Application Benchmark Performance Standard* (SAPS). An SAPS capacity estimate is represented as a hardware-independent value.

IBM uses Commercial Processing Workload (CPW) values to express the capacity of IBM System i processors. CPW values allow a comparison of the capacity of each System i server. This values exhibit a different behavior than the rating of System i configurations in SAPS.

IBM maintains tables to express the capacity of each System i model in SAPS values. These tables include the minimum memory recommended for each System i configuration for an SAP implementation. These tables are called *SAPS tables* in the following discussion.

The resources recommended for a System i implementation of SAP are derived as follows:

- Main memory recommendations are identified in the SAPS tables for System i servers
- Base disk requirements are derived from calculations generated with the SAP Quicksizer
- Network bandwidth can be derived from SAP guidelines

Contact your IBM Business Partner or IBM for assistance in sizing an SAP application. IBM Business partners are also supported through IBM Techline.

7.2 Size with the help of the SAP Quicksizer

The SAP Quicksizer provides two sizing categories:

- User-based
- Throughput-based

The objective of sizing is to calculate the number of transactions in a given period of time. Each transaction requires a certain CPU capacity, expressed in terms of SAPS. Multiply the number of transactions by the required capacity. The result is the required capacity in SAPS. Use this information to calculate the required minimum disk space and main memory.

The values output from the Quicksizer are:

- The number of SAPS
- Disk storage in GB, based on the amount of data generated in one year
- Required main memory, expressed in GB

Note: SAP's sizing process assumes a well tuned system. The result of the sizing process might not match your system requirements if your system is not well tuned or if your applications differ significantly from SAP standard.

7.2.1 Input to Quicksizer

This section discusses input to the SAP Quicksizer and what is done with the Quicksizer output to derive a System i configuration for SAP.

The SAP Quicksizer asks for the number of concurrent active users. Each user can be categorized as a:

- *Light user* who enters ten transactions (dialog steps) per hour
- *Medium user* who enters 120 transactions per hour
- *Heavy user* who enter 360 transactions per hour

The number of users entered can vary for each category for different SAP applications.

Note: In case you do not have direct access to the SAP Quicksizer you can ask your Business Partner or IBM to provide you with a offline questionnaire, which closely reflects the input fields of the SAP Quicksizer. The questionnaire can be obtained from the ISICC.

Throughput-based sizing uses a different approach. To determine the load in the specified amount of time, the number of specific documents is entered, or the number of processes to be handled in a given amount of time.

A thorough understanding of the meaning of the entry fields is required. This discussion is beyond the scope of this redbook.

7.2.2 Considerations for the Quicksizer output

Consider the following when interpreting the output from the SAP Quicksizer:

- The amount of SAPS
  - Sizing is an approximation relative to the amount of SAPS.
  - User-based input is less accurate than throughput-based sizing. Hence, the Quicksizer assumes a utilization of 33%. Another 32% is reserved for activities not accounted for elsewhere.
  - For throughput-based sizings the quicksizer assumes a utilization of 65%. If the Quicksizer recommends 1000 SAPS for a system which is rated at 1000 SAPS, the utilization would be 65%.
  - Spare capacity is required to ensure that response time does not deteriorate to an unacceptable level. The remaining 35% serves this purpose, and is available for workload that is not factored into the sizing model of the Quicksizer.
  - Most of the Quicksizer results are based on an Global Language Solutions (GLS) system running on the ABAP stack. If there is no GLS version of the system available or if the results refer to a Java-based system, the results are based on the Unicode system. For example XI or EP results are reported as Unicode systems.

  Note: GLS systems are sometimes incorrectly referred to as ASCII systems.

- Disk storage in GB

  The recommended disk storage is expressed in GB, and is based on the amount of data generated in one year. Factor this in if you plan to keep several years of data on the system. The amount of disk space required for the operating system and the database is not included in this number.

- Main memory in GB

  The main memory expressed in GB is not a good indicator what is required. Experience with customer installations shows that the main memory proposed by the Quicksizer is too small to run the defined workload. For these processors, adopt another approach. Consider the size the main memory as a function of the number of processors and the capacity expressed in SAPS. This estimation produces a number which scales, albeit not linearly, with SAPS as a minimum amount required for the main memory.

Important: Documentation for the Quicksizer provided by SAP can be found on the SAP Marketplace at:

`http://service.sap.com/sizing`

Use the documentation in order to understand and interpret SAP Quicksizer results.

7.2.3 Work with Quicksizer output

Consider the following selections when using the SAP Quicksizer to size a solution:

- ABAP GLS system
- ABAP Unicode system

- Java system

Processor, main memory, and the disk requirement considerations include the following:

- Processor requirements
  – Use the Quicksizer output with no modifications to size an ABAP GLS system
  – Multiply the Quicksizer results by a factor of 1.3 for an ABAP Unicode system.
  – Use the Quicksizer output without modifications for Java based systems.
- Main memory requirements
  – Follow the standard IBM approach for ABAP GLS systems. It is based on a GB per processor recommendation. This recommendation reflects the capacity per processor expressed in SAPS. You first determine how many processors you would need. Secondly you apply the multiplier “c GB/processor” for the system under consideration.

    Be aware that these are minimum requirements, if you want to run more than one SAP system under one system image or have additional demand you have to add more memory to the system. For example, the current recommendation is c = 6 GB per processor as the minimum requirement for a System i solution running on POWER5 technology. IBM Techline or the ISICC can provide you the information of the current factors.
  – ABAP Unicode: Follow the standard IBM approach then multiply the results by 1.2.
  – Java: Follow the standard IBM approach then multiply the results by 1.5.
- Disk requirements. Multiply the Quicksizer results by:
- 1.7 for ABAP GLS systems
- 1.8 for ABAP Unicode systems
- 1.8 for Java systems
- Determine the number of disk drives

  In case the user wants to retain more data, than build this assumption as a factor of the QuickSizer. Use the output of the Workload Estimator (WLE) to determine the recommended number of disk arms to configure for a balanced system.

Table 7-1 summarizes all modifications, which have to be applied to the SAP Quicksizer Output.

Table 7-1 Modifications to SAP Quicksizer output

SAP system category	SAPS modification	Main memory modification	DISKS modification
ABAP GLS based system	None	Use recommendation GB/processor	Multiply the results by 1.7
ABAP Unicode based system	Multiply the results by 1.3	Use recommendation GB/processor	Multiply the results by 1.8
Java based systems	None	Use recommendation GB/processor	Multiply the results by 1.8

7.2.4 Considerations for sizings not based on the Quicksizer

The most common SAP applications are reflected in the Quicksizer. However, there are other solutions, which are sized with other tools, for example paper sizing tools (presentations), spreadsheets. Still the output of these tools yields SAPS and the considerations using these results are the same as using the Quicksizer results. You can ask your IBM Business Partner for assistance. IBM Business Partner may ask IBM Techline or the ISICC for assistance and ultimately IBM Techline will propose an appropriate configuration.

7.2.5 Size upgrades

This section discusses how to determine the resources required to run your workload on a new version of software when upgrading a production SAP system to newer versions of the application software. The sizing process for upgrades is in some ways easier, because it is based on a system in production where you can better determine actual workload requirements. Use IBM Insight tool or SAP Early Watch Reports to help in this sizing effort.

IBM Insight tool collects statistical data provided by your SAP production system. The output is a report which provides information about the current status of your system relative to workload. This data can be sent to IBM. For details, see:

http://www.ibm.com/erp/sap/insight

The sizing process is based on the workload data of your system. SAP provides information about how many resources are required to run the same workload on a more current release.

The requirement for running the same workload on a more current release is derived with some straightforward calculations. The following guidelines prevail:

- Use your current workload information and extrapolate from it to estimate the system resource required to provide service for more users with a similar workload pattern.
- Use the traditional sizing approach and add the additional requirements to the workload to estimate the system resource required to provide service for additional functionality.

Consult with your IBM Business Partner or IBM for an appropriate configuration in case your system needs additional resource to manage the increase of requirements.

7.2.6 Summary

Use the SAP's Quicksizer or other sizing procedures from SAP and provide the output to your IBM Business Partner or IBM. If you are not familiar with the procedure, ask your IBM Business Partner or IBM for further assistance. With the help of IBM Techline your partners can propose an appropriate configuration.

Part 2

Operational topics

Part 2 describes actions that you should take in the ongoing operation of your SAP system with the System i family of products. After a technical introduction of IBM i5/OS from an SAP point of view, we provide a brief discussion of installation topics. From there, we discuss the SAP infrastructure objects including the kernel and instance structures that you just installed. Next we move on to discuss how to handle the logistics of SAP software changes across your newly installed landscape. Next, we provide information regarding how to print from your new systems.

Now that you have several systems installed and connected together, discussions of the day-to-day SAP administration is appropriate. We discuss some significant life cycle events that eventually occur over the life span of your SAP system environment. We discuss how to avoid having problems in your environment, what to do if you should encounter an issue, and finally have some example scenarios that show you what you should do to correct certain problems.

8

i5/OS from an SAP application view

In order to perform system administration, you need a general understanding of the concepts of i5/0S. You also need to know about the user interface and the most important operating system commands.

This book represents an i5/OS V5R4 view and is written for two audiences:

- Those familiar with SAP applications and new to the System i product line
- Those familiar with the System i models and new to SAP applications

This chapter serves as an introduction to the i5/OS features for both audiences.

This chapter describes i5/OS features from an SAP system perspective. However, it does not provide an exhaustive description of i5/OS and System i features. For a more detailed description of i5/OS, refer to the IBM eserver iSeries Information Center at:

http://publib.boulder.ibm.com/iseries/

8.1 System i architecture and overview

With the announcement of i5/OS V5R4 in January of 2006, *System i* is the term that describes the combination of IBM System i5™, IBM eServer i5, IBM eServer iSeries, and AS/400e™ product families. The *server*, or sometimes referred to as *the machine,* is the hardware server component in the system unit of System i models.

Note: When the System i product line's predecessor was first announced in 1988, it was named *AS/400*, which stands for Application System/400. The term AS/400 is still a common name for this platform today, especially in the user community. With the announcement of reduced instruction set computer (RISC) technology, the AS/400 became known as the *AS/400e*. Among other features and benefits, the "e" signifies e-business. When the Models 270, 820, 830, 840, SB2, and SB3 servers were announced in June of 2004, the product line became known as the iSeries. Some products still carry the name "iSeries" as the official name of the product, for example, iSeries Navigator.

There is a similar corollary with the "name" of the database. When the AS/400 family was designed, the database was not designated with a special or separate name. As part of the operating system which is packaged, ordered, priced, and installed as one entity, the database is integrated in the operating system. Users take it for granted that a reference to the operating system includes the database, and vice versa. Then the database became known as *DB2/400*. Today it is called *DB2 UDB for i5/OS*, still an integrated part of i5/OS.

8.1.1 Single-level storage

IBM i5/OS views the disk space on the server and in the server's main memory as one large storage area. This addressing architecture is known as *single-level storage*. Each object and each piece of temporary storage has its own unique 64-bit virtual address. In the understanding of i5/OS, main store or main storage is actually storage. On disk and in main storage, the data is organized in logical pages. The pages that belong to one object are automatically distributed over multiple disks. Figure 8-1 illustrates this concept.

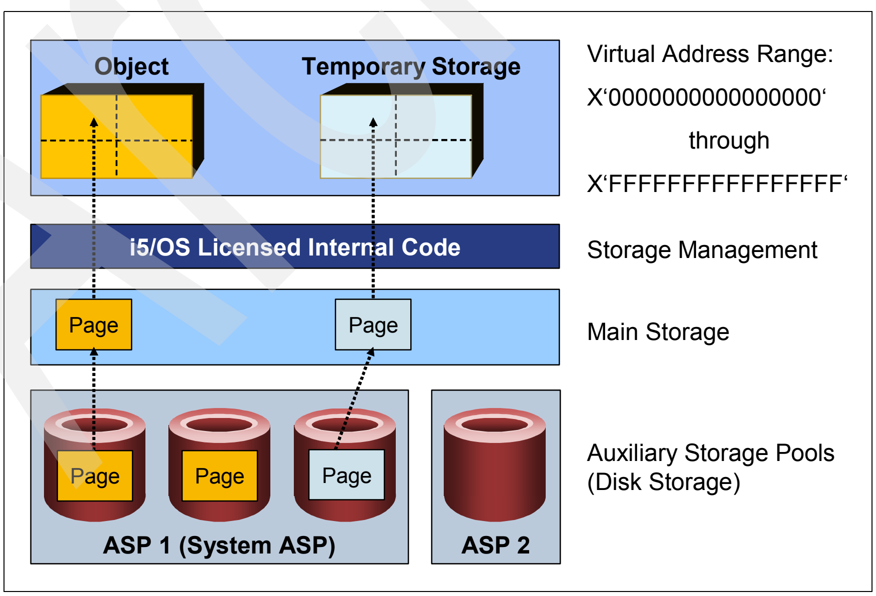

Figure 8-1 Single-level storage

Main storage is divided into memory pools. One pool, the **MACHINE pool*, is for system functions. In a typical configuration, another pool is the **BASE memory* for batch jobs and applications. Often, the **INTERACT pool* is available for interactive jobs. Memory pool assignments are made at the subsystem level. SAP-instance subsystems usually run in the *BASE pool. Depending on the configuration, there can be more pools than the ones mentioned here.

SAP note 49201 recommends that you give as much memory to the *BASE pool as possible. If only SAP applications run on this system, you can resolve the memory assignment of the *INTERACT pool by assigning all pools of other subsystems, such as QINTER, which use the related pool number of the *BASE pool. Use the i5/OS CHGSBSD command to make the pool assignment.

SAP note 49201 also recommends that you activate **Expert Cache** in the *BASE pool, if you have installed the IBM Performance Tool product (5722-PT1). Use the i/5OS WRKSHRPOOL command to set the paging option to ***CALC**.

8.3, "System i software products for an SAP implementation" on page 72 discusses the subsystems. Refer to Chapter 13, "SAP instance profile parameters" on page 173 for a discussion of SAP instances.

8.1.2 Layered architecture

System i architecture is based on layers, as illustrated in Figure 8-2.

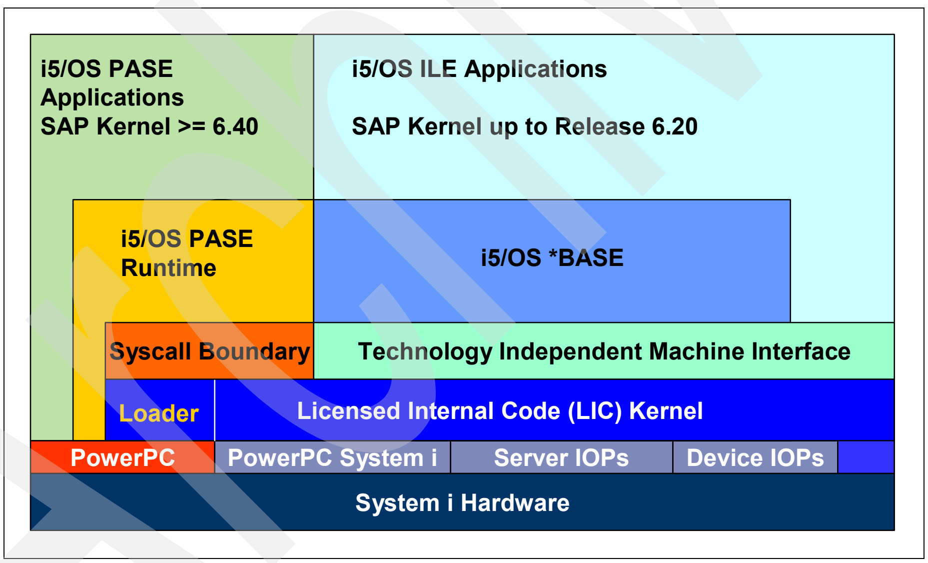

Figure 8-2 The System i layered architecture

The right side of Figure 8-2 illustrates the layered architecture of i5/OS. All applications and the base operating system are based on an architecture known as *technology independent machine interface* (TIMI) or more simply referred to as the *machine interface* (MI). If the underlying hardware changes with new models, you only need to adjust the system licensed internal code (SLIC), more simply called the *licensed internal code* (LIC). You do not need to recompile the applications. You may need to retranslate only the internal code representation. This is usually done at "first touch" and does not require user intervention.

Note: Beginning with i5/OS V5R3, devices without input/output processors (IOPs) are also supported with specific hardware configuration.

Starting with kernel release 6.40, SAP applications use the i5/OS Portable Application Solutions Environment (i5/OS PASE) as a runtime environment. i5/OS PASE provides a broad set of interfaces, in a runtime that allows PowerPC binary programs to execute directly on the PowerPC processor of IBM System i models. While many of the executables in the 6.40 kernel are PowerPC binaries, the application uses the integrated DB2 UDB for i5/OS database and keeps the look and feel of a System i application.

8.6, “i5/OS Portable Application Solutions Environment” on page 84 describes i5/OS PASE in further detail.

8.1.3 Overview of i5/OS

> **Note:** Prior to i5/OS V5R3, the operating system on System i models was named *OS/400*. Starting with V5R3, the operating system is called *i5/OS*. As this book is based on i5/OS V5R4 and V5R3, we use the term i5/OS.

i5/OS is a 64-bit object-based operating system for today’s IBM System i models. Most addressable entities in the system are considered objects. *Libraries* are a special type of object that contain other objects, such as structured query language (SQL) tables, source files, and compiled programs.

The i5/OS Command Language (CL) is very powerful and easy to learn because of the object-based nature of the operating system. In CL, you can use intuitive commands to act upon objects, such as with the Work with Objects (WRKOBJ) command.

The type of object determines which operations can be performed on it. Objects created or owned by the operating system usually start with the letter “Q”. You can prompt CL commands by pressing the **PF4** key at the i5/OS command line to show all the parameters and values, so you do not have to memorize cryptic parameter strings. You can use the Help function available in nearly all fields and commands by pressing the **PF1** key. 8.2, “i5/OS object structure” on page 55 describes the object structure further.

i5/OS has some features common to UNIX operating systems, including a directory-based file system known as the *Integrated File System*. This file system is used for many of the common SAP application files, and follows the same structure used on other platforms. Many UNIX-like commands are applicable to Integrated File System objects, yet i5/OS CL commands are also available. The Work with Links (WRKLNK) command is the most commonly used CL command to access the Integrated File System and its directory structure. 8.2.4, “The Integrated File System” on page 63 describes the Integrated File System further.

8.1.4 Subsystems

Jobs are organized in subsystems. Several subsystems are created and managed by i5/OS. They usually start with the letter “Q” and the jobs within them provide system services.

Some system jobs do not run in a subsystem. The SCPF job is the system control process. It is the first job that starts when the system comes up. QDBSRVXR and QDBSRVXR2 are the system jobs to update the database catalog files asynchronously. QDBFSTCCOL is the statistics collector for the database.

If you have SAP applications installed on a System i system, a subsystem named R3_*nn* is created by the SAP installation tools, where *nn* is the instance number, for example, R3_00, R3_01. All processes that are needed for an SAP instance run in that subsystem. Refer to

Chapter 13, “SAP instance profile parameters” on page 173 for a discussion of SAP instances.

8.1.5 Auxiliary Storage Pool

The disk units that are attached to a System i server are combined into one or multiple Auxiliary Storage Pools (ASPs). All systems have a System ASP named *ASP 1*. You can configure additional ASPs called *user ASPs*.

Each disk unit is uniquely assigned to an ASP. Pages that belong to the temporary storage reside in the system ASP. Non-system objects can reside either in the system ASP, or in a user ASP.

Whenever a program accesses a page for the first time, the page is copied from disk into main storage. It stays there, but it is paged it out if you do not continually use it. The Storage Management component of the System i architecture is responsible for allocating and freeing space on disk. New data is distributed evenly over the available disk units in the ASP.

You can display the amount of storage used in the system ASP with the i5/OS Work with System Status (WRKSYSSTS) command. Refer to 9.7, “ASP considerations for SAP journal management” on page 110.

See “Auxiliary storage pools” on page 378 for further information about ASPs.

8.1.6 Independent ASP

Another feature of the System i architecture is an Independent Auxiliary Storage Pool (IASP) where the data and system information stored in it can be handled as an entity.

An IASP contains user data and the necessary system information associated with the data. You can make it available (varied on) and unavailable (varied off) to the server without restarting the system. Figure 8-3 illustrates this concept.

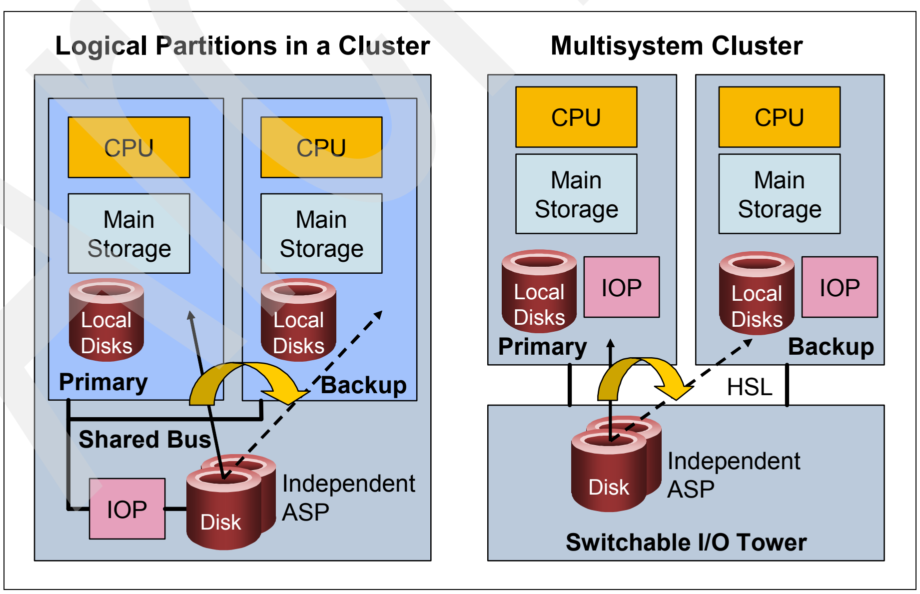

Figure 8-3 The System i Independent ASP

Switched disk support in a logical partition (LPAR) environment (where disk devices switch between partitions) is available on all systems that support LPAR, however, you need additional hardware such as a dedicated bus. For example, you can set up a test or backup environment only by switching the disk devices of a special environment to the test environment. Additionally, you save resources because you need only one physical disk system. System i high-speed link (HSL) configurations with System Area Network (SAN), available with OS/400 V5R1 onward, support switching disk devices between systems.

In order to perform an automatic switchover or failover from a primary to a backup system, you can configure a *System i cluster*. A System i cluster is a collection or group of one or more servers that work together as a single server. A *cluster node* is a server or logical partition that is a member of a cluster. In cluster resource groups (CRGs), you can define actions to be taken during a switchover or failover. Refer to 24.2.3, “Clustering” on page 436 and 24.3.2, “Cross site mirroring” on page 451 for a further discussion of clustering and cross-site mirroring.

SAP note 568820 explains how to configure the use of IASPs with SAP software.

8.1.7 Logical partitions

The System i server offers the ability to partition one server into several independent servers. Each logical partition operates as an independent logical server. However, each partition shares some physical system attributes such as the system serial number, system model, and processor feature code. All other system attributes can vary among partitions. Refer to Figure 8-4 for an illustration.

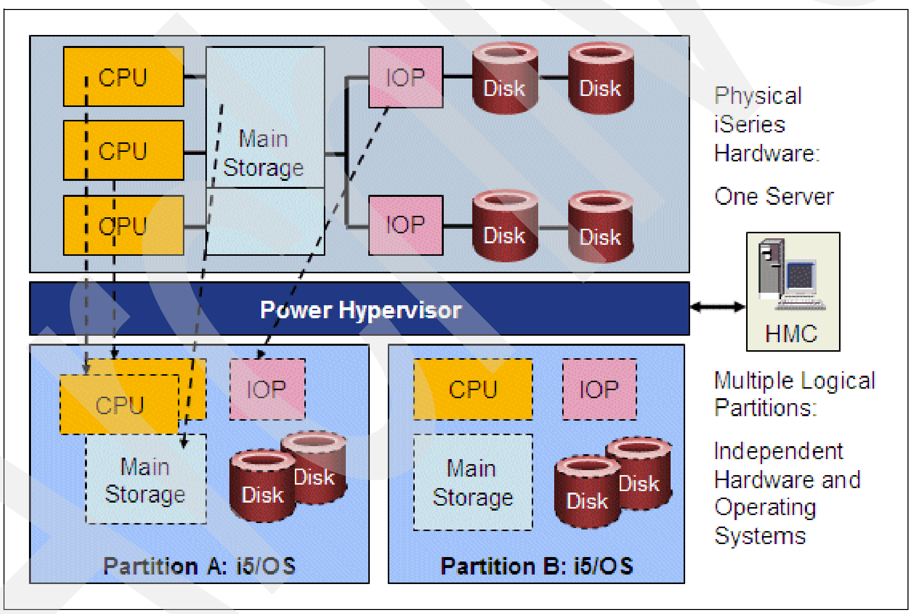

Figure 8-4 Logical partition concept

On systems prior to i5/OS V5R3 and POWER5 hardware, you have to define one primary partition that manages the other partitions. If the primary partition goes down, the secondary partitions go down as well. Beginning with i5/OS V5R3, with POWER5 or POWER5+ hardware, partition management is done by the Hardware Management Console (HMC). A primary partition is no longer necessary.

Logical partitions can use their own hardware resources. You can set up shared resources, such as CPU or main storage. You can move shared resources dynamically between partitions. This is useful if one partition contains a system that you only use occasionally, for example, for data warehousing. You can assign fractions of a CPU to partitions, down to 0.1

of a CPU. System i POWER5 and POWER5+ models support i5/OS, AIX 5L, and Linux operating systems. Each partition can run any of the supported operating systems.

There is only one IBM serial number for each System i machine (“box”), even if there are a multiple partitions installed.

You can find further information about LPAR in Chapter 5, “System i virtualization” on page 31.

8.2 i5/OS object structure

System i models have two main areas where data is stored:

- Libraries
- Integrated File System

The Integrated File System is further organized with directories and stream files. Refer to 8.2.4, “The Integrated File System” on page 63 for more information.

Libraries contain objects. i5/OS objects provide the means through which all data processing information is stored and processed by the System i processor. An i5/OS object is a named unit that exists (occupies space) in storage and on which the operating system performs operations. Object names can be 10 characters long. Usually object names start with a letter and can contain letters, the numbers zero to nine, as well as some characters like underscore, the dollar sign, or the pound sign. It is possible to create objects with other characters if you enclose the name in double quotes, which leaves only eight characters for the object name.

i5/OS objects have the following in common:

- Each object is uniquely identified by an object name, object type, and the library it resides in.
- You can rename all objects and move most objects to another library. You cannot modify the type of an object.
- Each object has a generic header template that is the same for all object types. It contains the object name and type, the owner, the library that the object resides in, the object size, backup information and more.
- Depending on the object type, the object contains a type-specific header and the data that is associated with that type.
- The object type determines the types of operations that can be performed on an object.

8.2.1 Object types

Many objects can be in any library, but some can only be in the QSYS library. Among those that can only be in the QSYS library, is the object type library (*LIB). This means that you do not have a hierarchical but a flat structure of objects in the system. Processes or jobs maintain a library list that contains only the libraries of interest for their environment.

The following is a list of some of the more important object types in an SAP system environment.

Object types in any library

The following object types can be stored in any library:

- Program (*PGM)
- SQL package (*SQLPKG)
- Service program (*SRVPGM)
- Journal receiver (*JRNRCV)
- Journal (*JRN)
- Job queue (*JOBQ)
- Output queue (*OUTQ)
- Message file (*MSGF)
- *FILE:
  - Physical file, SQL table (PF)
  - Logical file, index, view (LF)
- Message queue (*MSGQ)
- Job description (*JOBD)
- Command (*CMD)
- Subsystem description (*SBSD)

Object types only in the QSYS library

The following object types can only be stored in QSYS, the i5/OS system library:

- Library (*LIB)
- User profile (*USRPRF)
- Device description (*DEVD)
- Line description (*LIND)
- Controller description (*CTLD)
- Authorization list (*AUTL)

Special library QTEMP

There is a special library that is called *QTEMP*. Each process or job has its own QTEMP library. Objects in this library are only visible to the owning process or job.

The QTEMP library is:

- Available in each process or job
- Only visible for this process

When a process ends, the QTEMP library is deleted. Therefore, it is primarily used for temporary objects.

8.2.2 SQL collection objects

SAP applications use relational databases, and therefore, DB2 UDB for i5/OS on System i models. Together with a database, SAP applications also use commitment control mechanism for their applications.

Commitment control

Commitment control is a function that ensures data integrity. It allows you to define and process a group of changes to resources, such as database files or tables, as a transaction.

Commitment control allows you to process a group of changes to resources as a logical unit of work (LUW).

Commitment control ensures that either the entire group of individual changes occur on all systems that participate, or that none of the changes occur. IBM DB2 UDB for iSeries uses the commitment control function to commit and rollback database transactions that are running with an isolation level other than *NONE (no commit).

You can use commitment control to design an application so that the system can restart the application if a job, an activation group within a job, or the system ends abnormally. With commitment control, you can be assured that when the application starts again, no partial updates are in the database due to incomplete transactions from a prior failure.

The commitment control feature is also mentioned in "Journal management" on page 377 in conjunction with high availability topics.

SAP applications need commitment control at all events even when there is no need for high availability. The following SAP applications use commitment control:

- The program-controlled transaction mechanism
- The SAP enqueue mechanism
- The SAP application flow with the user control buttons "Back", "Exit", and "Cancel"

For the SAP system, i5/OS commitment control is always active. This means that the system processes a group of changes to database files as a single transaction. While journaling itself allows you to ensure that only changes to the last record processed are lost in case of a system failure, journal management combined with commitment control enables the system to ensure that all changes within a transaction are completed or rolled back if processing is interrupted.

For more details specific to the SAP application use of the commitment control feature, type the key words `Journaling and Commitment Control` in the SAP Help Portal on the Web at:

http://help.sap.com/saphelp_nw04/helpdata/en/start.htm

For more details about i5/OS features, refer to an example in the iSeries Information Center at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp

Alternatively, refer to the book, *Advanced Functions and Administration on DB2 Universal Database for iSeries*, SG24-4249.

SQL collection

An *SQL collection* or *schema* is a library containing a number of specific objects as illustrated in Figure 8-5 on page 58.

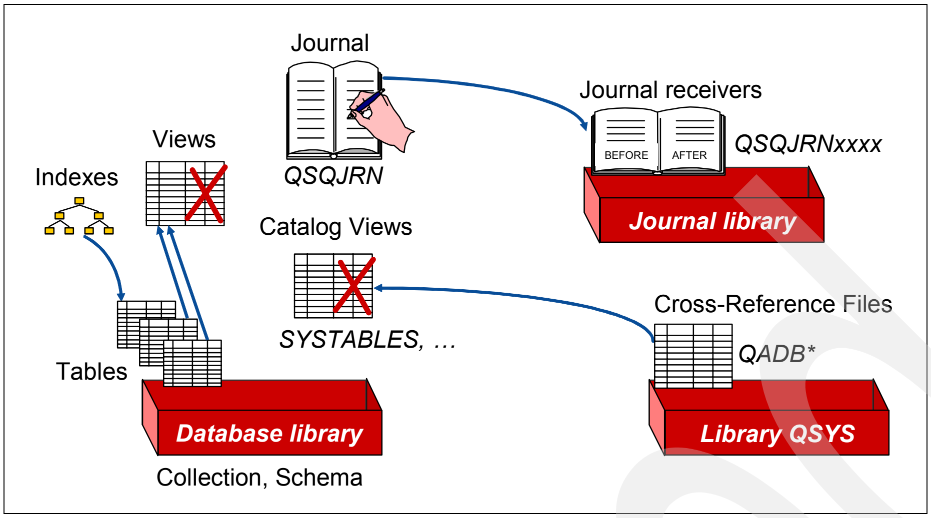

Figure 8-5 SQL collection

An SQL collection consists of the following i5/OS object types:

- Library (*LIB)

  All tables, views, and indexes, as well as the catalog views, and the journal are in one library that has the same name as the collection (or schema). Journal receivers can be in the same library or in a different one.

- Tables (*FILE, PF)

  Physical files that contain data.

- Views (*FILE, LF)

  Logical files over one or multiple tables.

- Indexes (*FILE, LF)

  Logical files over one or multiple columns in a table for faster access.

- Catalog Views (*FILE, LF)

  Logical files over system cross-reference files, for example, used to link long SQL names to short i5/OS names. You can find the SYSTABLES file in this catalog view.

- Journal (*JRN)

  Named QSQJRN, located in the collection library. All tables that are created in this collection are automatically journaled.

- Journal receiver (*JRNRCV)

  Name QSQJRN*nnnn*, where *nnnn* is an automatically generated number. The journal receiver can be located in any and in different libraries.

A given SQL collection is assigned to an ASP. i5/OS automatically spreads new SQL objects in the collection across the disk units in the ASP. A benefit of i5/OS is that the ASP is essentially the one and only tablespace for a complete SQL collection. On other platforms, you must monitor and very carefully manage multiple tablespaces.

QSQJRN is being used in the SAP system environment because SAP system data is located into an SQL collection, and this naming is used by the **Create Collection** command.

SQL long name mapping

SQL permits names of up to 128 characters, but i5/OS operating system objects are named with a maximum length of 10 characters. Therefore, different names exist in SQL and on the operating system level. The mapping between short operating system names and long SQL names is done with the help of catalog views. You can use the SYSTABLES catalog view for the mapping of table and view names. You should use the SYSINDEXES catalog view to map index names. Figure 8-6 illustrates this convention.

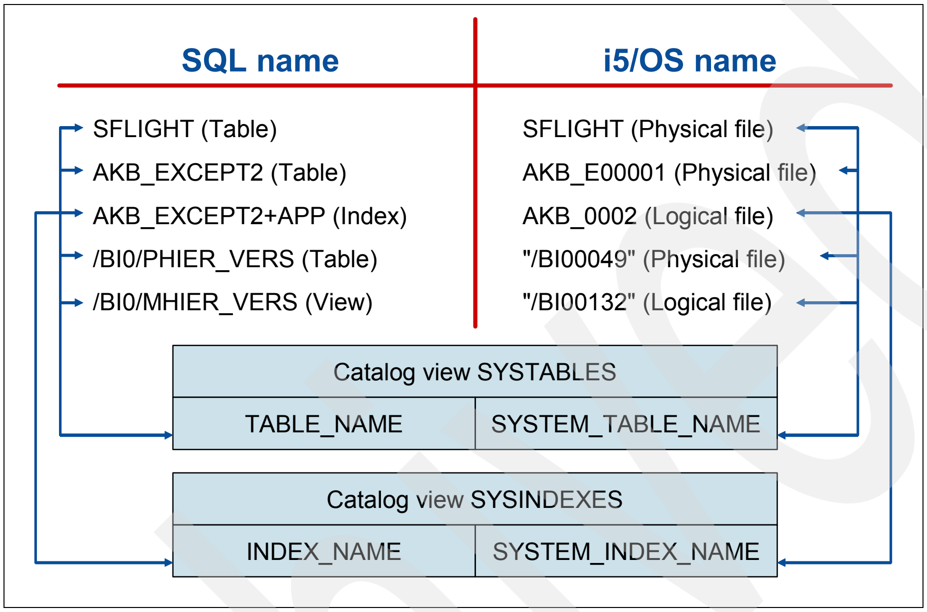

Figure 8-6 SQL long name and i5/OS names of objects and file systems

If the SQL name has a maximum of 10 characters and represents a valid i5/OS object name, the SQL name and the operating system name are the same. Refer to the example named `SFLIGHT` in Figure 8-6. If the SQL name exceeds 10 characters, but starts with a letter and uses no special characters, the operating system builds a name based on the first five characters of the SQL name and generates a five-digit number. Refer to the 11-character long `AKB_EXCEPT2` example in Figure 8-6.

If the name contains special characters such as "+" or "/", the operating system builds a name based on the first four characters and generates a four-digit number. Refer to the `AKB_EXCEPT2+APP` example in Figure 8-6. If the resulting name is not a valid i5/OS name, it is enclosed in double quotes. Refer to the last two examples, **"**`/BI00049`**"** and **"**`/BI00132`**"**. The operating system ensures that no duplicate names are generated.

To find the corresponding i5/OS name for an SQL name, query the catalog views named SYSTABLES or SYSINDEXES. If you know the short i5/OS name of an object, you can find the corresponding SQL name in the output of the Display File Description (DSPFD) command. Primary keys are not implemented as objects in a library, but are part of the SQL table.

Primary key constraint

The primary index has two functions:

1. It defines a set of columns whose values are unique for each row in the table. By specifying all columns of a primary key in the WHERE clause of an SQL statement, you can ensure that only one row is returned.

2. It provides an index to the table which you can use for faster access, ordering or grouping of rows.

The primary index is stored within the physical file object (the SQL table) in the database, not as a separate logical file. SAP's naming convention for the primary key in the Advanced Business Application Programming (ABAP) stack is to start with the name of the table and end with the number zero.

The Java dictionary has primary keys, but they do not have a name. Consequently the database creates the primary key name for itself.

Information about the primary key constraint can be found in the SYSCST and SYSKEYCST catalog views in the file description of the SQL table (the physical file). Use the i5/OS DSPFD command or the Work with Primary File Constraint (WRKPFCST) command. Use DSPFD to determine which columns build the primary key.

Journals and journal receivers

Journals and journal receivers are objects used for journaling. *Journaling* is a chronological recording of database changes. These records are required for commitment control, that is the capability of treating a set of SQL statements as one entity.

Journaling also enables:

- Forward recovery of the database from the latest backup up to a specific point in time.
- Backward recovery of the database from its current state backwards to a specific point in time.

The following two object types enable journaling in i5/OS:

- The QSQJRN journal, located in the collection library

  QSQJRN defines which objects are to be recorded in the journal. In an SAP system, all tables are journaled.
- The QSQJRNnnnn journal receivers

  These journals are either in the collection library or in another library. The last four characters in the name are a four-digit number. The receivers contain the journal entries, and can grow quickly on a busy system.

  Journal receivers are typically defined with a threshold size. When the defined threshold is reached, the system creates a new journal receiver, detaches the old journal receiver, and attaches the journal to the new receiver. This is called a *journal switch*. You can then save the old receiver and delete it from the disk.

For every update request, the database management system writes a copy of the data record to the journal receiver and then changes the database. The journal receiver then contains two images: the data record before the change (a “before image”), and the record after the change (an “after image”).

Refer to “Journal management” on page 377 for a further discussion of journals.

8.2.3 Database access from the SAP system

SAP ABAP applications (the “ABAP stack”) access the i5/OS database with the LIB_DBSL kernel component. The kernel translates the open SQL statements from ABAP to DB2/400 SQL statements.

Open SQL consists of a set of ABAP statements that perform operations on the central database in the SAP system. The results of the operations and any error messages are independent of the database system in use. Open SQL thus provides a uniform syntax and semantics for all of the database systems supported by SAP. ABAP programs that only use Open SQL statements work in any SAP system, regardless of the database system in use. Open SQL statements can only work with database tables that are created in the ABAP dictionary.

In the ABAP dictionary, you can combine columns of different database tables to a database view (or view for short). In Open SQL statements, views are handled in the same way as database tables. Any references to database tables in the following sections can equally apply to views.

SQL packages

An SQL package contains access information needed by the DB2 UDB optimizer for iSeries. The SQL package complements other access path information contained in the SQE plan cache. Optimizing an SQL statement is quite complex and time consuming, especially when the query is joining multiple tables. In order to reduce the processing overhead, significant parts of the access plan are persistently saved in the SQL package and reused when the same statement is executed again. This is even possible across IPLs. An SQL statement can be uniquely referenced knowing the library and package it is stored into and its statement ID.

The information that is stored in an SQL package can be viewed or printed by using the i5/OS command Print SQL Information (**PRTSQLINF**).

There are a number of libraries on an SAP system that are named R3<sid>nnnnn. Many of these libraries contain SQL packages for the ABAP stack of the SAP system. One or more SQL statements are in each package.

These libraries, packages and statements are created and handled automatically by the SAP system. If SQL packages have been deleted, they are automatically rebuilt during normal operation with the performance penalty of having to rebuild the access information. Consequently, packages do not need to be saved. Furthermore, obsolete packages are automatically removed by the SAP system during startup. Under normal operation it is normally not necessary to delete packages; however, you might be asked to do so by SAP or IBM via notification after having applied major changes to your system. The **DLTR3PKG** command from the kernel library gives you the ability to delete all packages belonging to a specific SAP *SID*.

Note: Never delete SQL packages while the SAP system is running.

Java is different. JDBC™ uses pure dynamic SQL, it does not use SQL Packages.

Reorganization of SQL tables

Reorganization of SQL tables is normally not necessary in a DB2 UDB for an iSeries installation. All files in the database library R3<SID>DATA have the attribute Reuse Deleted Records (REUSEDLTRCD) set to *YES, so rows that are deleted are reused later for new rows. Tables with many deleted rows can degrade performance because of the paging of blocks into memory when performing SQL SELECTs becomes inefficient if a large number of deleted records are returned in the block. Reorganize the tables to improve the performance with large tables containing a lot of deleted rows.

Prior to i5/OS V5R3, files are not available during the Reorganize Physical File Member (RGZPFM) command. This command can only run when the application is down. Beginning with i5/OS V5R3, you can use the RGZPFM command with the <> option while the SAP system is running.

SQL table statistics

DB2 UDB for i5/OS uses statistics about the columns involved in a particular query during the optimization of that query. Set the QDBFSTCCOL system value to *AUTO. Statistics are gathered automatically and are stored with the table. If the table is saved and restored, the statistics are still there. As the file changes, the statistics can be automatically recollected at the discretion of the database. They also be recollected as a result of an operating system upgrade or with the application of a Program Temporary Fix (PTF).

SQL Interfaces

There are different database access methods for the SQL statements:

- Static SQL
- Dynamic SQL
- Extended dynamic SQL

The SAP Web Application Server ABAP database interface uses *Extended Dynamic SQL*. In the following we explain this database access method.

Traditional SQL interfaces used an embedded SQL approach. SQL statements are placed directly in an application's source code, along with high-level language statements written in C, COBOL, RPG, and other programming languages. The source code is precompiled, which translates the SQL statements into code that the subsequent compile step can process. This method sometimes is referred to as *Static SQL*. One performance advantage to this approach is that SQL statements are optimized in advance, rather than at runtime while the user is waiting.

ODBC, however, is a call level interface (CLI) that uses a different approach. Using a call level interface, SQL statements are passed to the database management system (DBMS) within a parameter of a runtime application programming interface (API). Because the text of the SQL statement is not known until runtime, you must perform the optimization step each time an SQL statement is run. This approach commonly is referred to as *Dynamic SQL*.

> **Note:** Some DBMS use cache technics to intermediately store the database access information. So the difference between static and dynamic SQL is blurred. However, if the cache is cleaned or the system stops the database access information is lost. In contrast to that with the usage of SQL packages the database access informations are persistent.

So the ABAP stack of the SAP Web Application Server uses the so-called "*Extended Dynamic SQL*" as best of the "static" and the "dynamic" SQL to run the SQL statements. The SAP Database Service Layer (DBSL) uses dynamic SQL and additionally SQL packages with their prepared statements to benefit from these persistent stored database access paths as an additional static SQL part.

The use of extended dynamic SQL can improve response times, and can also dramatically improve server utilization. This is because it is better to perform this step only once as optimizing SQL queries can be costly. This works well with a unique feature of DB2 UDB for i5/OS. Unlike other DBMSs, the DB2 UDB for i5/OS database access with the usage of SQL packages ensures that statements which are stored in so-called SQL packages are kept up to date in terms of optimization, without administrator intervention. Even if a statement is prepared for the first time in weeks or months, DB2 UDB for iSeries automatically regenerates the access plan when it determines that sufficient database changes require optimizing again.

QAQQINI Query options

The IBM Query Optimizer automatically uses the settings from the QAQQINI Query Options File for SQL Query optimization. You can create and configure a local copy of this file for every SAP system, as described in *SAP note 820325* and related notes.

SAP DBCON

The SAP ABAP language normally uses the default connection to its own database. It also allows native SQL access to the SAP database or to other databases using the *SAP DBCON method*, as described in *SAP note 146642*.

Java database access method

For Java, the method is similar, except that one Java or Java 2 Platform, Enterprise Edition (J2EE™) application can open multiple connections at the same time in parallel, because it can handle parallel transactions on different connections. SQL statements and others are associated with a connection object (similar as for ODBC). Java also allows multi-connect. It is easier since every database has a Java Database Connectivity (JDBC) driver that runs on any operating system. (This is because type 4 JDBC drivers are completely written in Java.) You can define connections via the *JDBC Connector* Service in the Config Tool. *SAP note 924753* is the System i equivalent to the ABAP multi-connect note.

> **Note:** SAP does not use SQL packages when accessing the database through JDBC. Dynamic SQL is used in Java environments.

8.2.4 The Integrated File System

The Integrated File System is a part of i5/OS that supports stream input/output similar to PCs and UNIX operating systems.

Figure 8-7 illustrates the Integrated File System hierarchy.

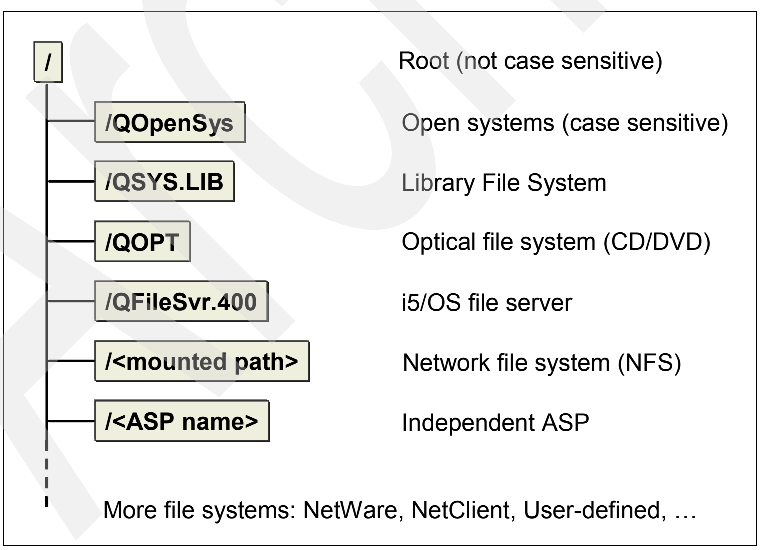

Figure 8-7 The Integrated File System

The following are the most commonly used file systems in an SAP system environment:

- Root File System

  All directories and stream files that are not part of another file system. Each component of file and path names can be up to 255 characters long and are not case sensitive.

- Open Systems File System (/QOpenSys)

  This is a UNIX-style file system. Path names and file names are case sensitive, can be up to 255 characters long, and start with /QOpenSys.

- Library File System (/QSYS.LIB)

  This is the path to libraries and objects as described earlier. For example you can access a physical file member by using the following path:

  `/QSYS.LIB/`*`MYLIB`*`.LIB/`*`MYDATA`*`.FILE/`*`MYDATA`*`.MBR.`

- Optical File System (/QOPT)

  This file system is used to access data on CDs or DVDs that are loaded into the built-in CD or DVD drive. The path name starts with /QOPT.

- i5/OS File Server (/QFileSvr.400)

  Allows you to access data remotely between System i servers.

- Network File System (NFS)

  This file system uses the standard Transmission Control Protocol over Internet Protocol (TCP/IP) protocol for NFS and can be used to access data on or from remote UNIX or System i servers.

- User-defined file system (UDFS)

  When an independent ASP is varied on, a user-defined file system (UDFS) is created automatically in that ASP. The user-defined file system is named /dev/*ASP name*/QDEFAULT.UDFS. In addition, a directory with the name of the ASP is created in the root directory, and the user-defined file system is mounted over that directory (/*ASP name*).

Symbolic links

There are two types of links:

- Hard links
- Symbolic links

Whenever a new directory or file is created, a hard link to that object is defined. When multiple hard links point to the same object and one of the links is removed, the object itself is not deleted. The object is removed when the last hard link to an object is removed. The object is deleted as well.

A soft link or symbolic link is created as an object with an object type of *SYMLNK. A symbolic link contains a path name as reference to an object. When a symbolic link is created, the target is not verified. When the symbolic link is used by an application, the content of the symbolic link is resolved. Symbolic links can cross file system boundaries, for example, a symbolic link /usr/sap/*sid*/sys/exe/run/AS4RMTCCMS (the root file system) can contain /QSYS.LIB/R3640A0056.LIB/AS4RMTCCMS.PGM (the library file system). When removing a symbolic link, the object that it points to is not deleted.

Path names can be absolute (starting with the symbol “/” in the root directory) or relative to the current directory. When creating symbolic links, you should be careful with relative path names. Each user has a home directory, which can be specified as “~” in relative path names.

Remote file systems

Figure 8-8 demonstrates two methods to access Integrated File System objects on remote hosts. In both cases it is assumed that a TCP/IP connection exists between the two hosts, and that the hosts are known by their names (HOST2 or HOST4). It is also assumed that all necessary TCP/IP services are started automatically during system startup. The Work with Links (WRKLNK) command illustrated in Figure 8-8 is only used to show how to access the directories.

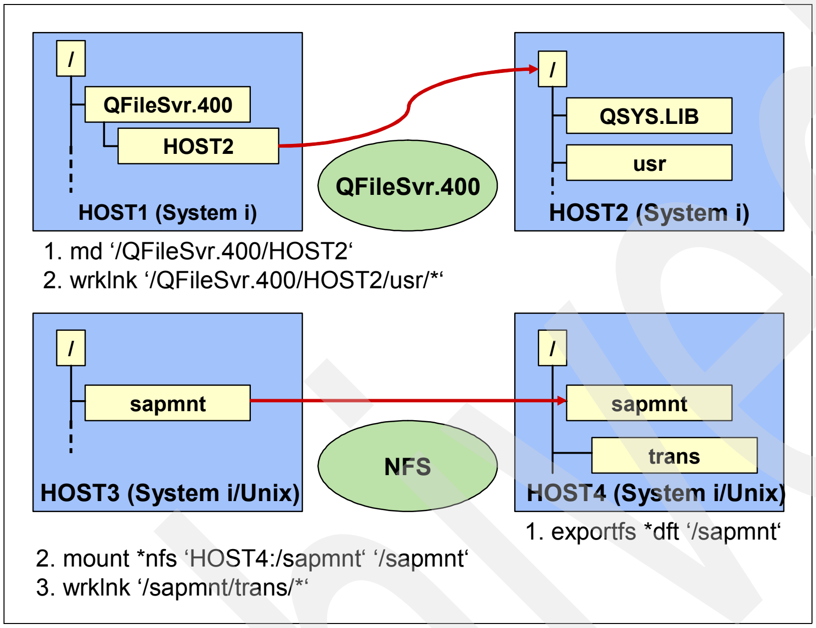

Figure 8-8 Remote file system objects and file systems

The first case shows the access through QFileSvr.400. The /QFileSvr.400 directory is always present on a System i system. In order to access the remote server HOST2, you need to create a directory with the same name as the remote host in the /QFileSvr.400 directory. You can do this with the Make Directory (`md`) command. After this, you can access the root directory and any underlying directories on the remote host by using a path name of **/QFileSvr.400/HOST2/...** if authorized. For this connection type, you must have the same user profile name with the same password on both hosts.

Even though /QFileSvr.400 is persistent, the contents of this directory are not. After an initial program load (IPL), the /QFileSvr.400/*hostname* (sub-) directory does not exist. So it makes sense to create this /QFileSvr.400/*hostname* (sub-) directory after each IPL automatically within the i5/OS QSTRUP start program. 17.3.3, "Starting SAP applications during an IPL" on page 309 describes this procedure.

The second case shows an NFS access. This allows access to directories on a UNIX system from a System i server. The administrator on the target system (HOST4) first needs to export a directory to allow remote access. Then the administrator of the source system (HOST3) can mount the directory over an existing directory on the source system. After that, you can access all subdirectories of the exported directory on the target system if authorized. Unless public access is granted, you need the same user ID on the source and target system.

SAP applications and the Integrated File System

The Integrated File System provides stream file support that is required by SAP applications. The SAP database uses native System i database files located in an i5/OS library (or SQL collection).

Figure 8-9 shows a part of the SAP directory structure on System i models. The directory structure starts with the root (/) file system.

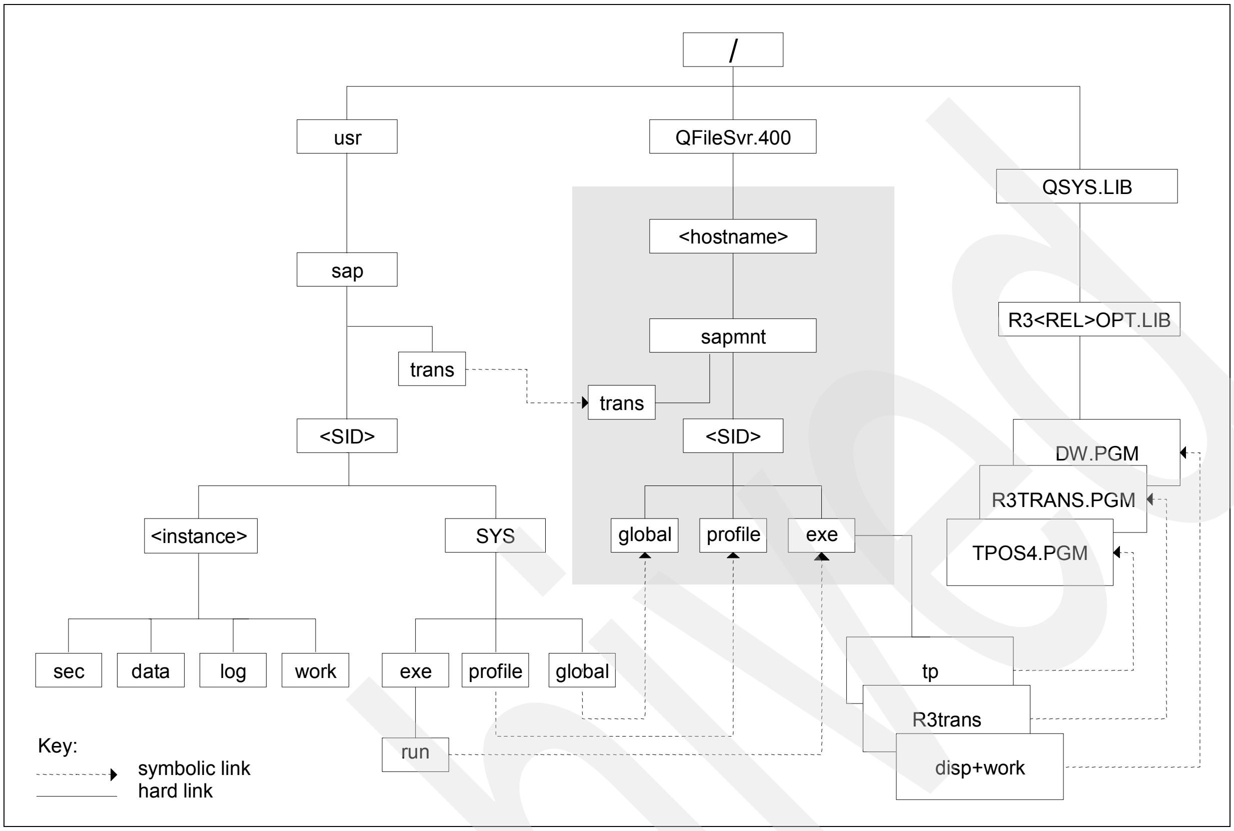

Figure 8-9 SAP Integrated File System directory structure

The Integrated File System directory for SAP systems consists of several file systems. Figure 8-9 shows the following three:

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

The QSYS.LIB file system is shown on the right-hand side of the directory tree. It contains the i5/OS library structure, which is used by the SAP database, the executable SAP kernel programs, work management objects, and so on.

The QFileSvr.400 file system, shown in the middle of Figure 8-12 on page 71, allows access to information about the remote System i server. The /QFileSvr.400 directory tree contains subdirectories that represent all other System i host systems that can be accessed by the local System i server.

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 points to /QFileSvr.400/*hostname*/sapmnt/*SID*/profile. The latter is where the information is actually stored. /QFileSvr.400/*hostname* identifies the System i server where the information is located. The *hostname* parameter can be the name of any System i system in the network.

> **Note:** You can change the default soft links to improve performance in certain situations. One such case is documented in *SAP note 68487*.

The /sapmnt/trans directory (shown in the shaded box in Figure 8-9 on page 66) is used by the SAP Transport Management System (TMS) to transport the SAP system objects between different systems. For more information about the Transport Management System, refer to 14.4, “Global transport directory” on page 189.

The SAP system continues to run if the connection to the system where this directory resides is lost. However, the Transport Management System cannot be used, because it needs the /QFileSvr.400/*hostname*/sapmnt/trans directory.

In an SAP landscape with multiple System i servers, the sapmnt/trans directory is used on only one of the System i systems. Each System i server contains its own /sapmnt/trans directory. Use the /QFileSvr.400 directory on all System i servers in the SAP system landscape to reference the /sapmnt/trans directory of the designated System i system.

The /usr directory, shown on the left-hand side of Figure 8-9 on page 66, contains all stream files used by SAP applications, including the profiles and SAP log files.

> **Note:** In Figure 8-9, several kernel programs, such as Disp+Work (DW), are shown in QSYS.LIB. This is only valid prior to 6.40. In 6.40 and later, DW is moved to the Integrated File System. There is a program object named DW in the kernel library which is just a wrapper to call the PowerPC executable.

Other significant directories are:

- /usr/sap/*SID*/sys/profile

  This contains all the profiles that are necessary to control the SAP system. It is normally linked to /QFileSvr.400/*hostname*/sapmnt/*SID*/profile.

- /usr/sap/*SID*/SYS/exe/run

  This contains references to the SAP kernel executables in the QSYS.LIB file system.

  With Web AS 6.40 and later, this directory contains PowerPC executables.

- /usr/sap/*SID*/*instance_name*

  This directory contains the data, log, and work subdirectories that are used to store SAP job logs and logs of work processes for that instance. One exists for each defined instance. This subdirectory contains the SAP Personal Security Environment files, which are used for running digital signatures. The /sec subdirectory is also found here.

- /usr/sap/trans/config/*SID*

  This contains the *hostname_system number*.cfg. and system.cfg central configuration files.

- /sapmnt

  Global SAP directory shared as *sapmnt* on the global host of the SAP system.

These configuration files contain information about all the SAP systems and the configured instances. They are used to manage the installation of the SAP systems. These configuration files are required to make changes to the installation configuration. Link every System i server in your landscape to the same configuration file.

Figure 8-10 and Figure 8-11 on page 69 show other views to the Integrated File System for an SAP implementation:

1. Integrated File System structures for the SAP Web Application Server, ABAP stack
2. Integrated File System structures for the Web Application Server, ABAP and Java

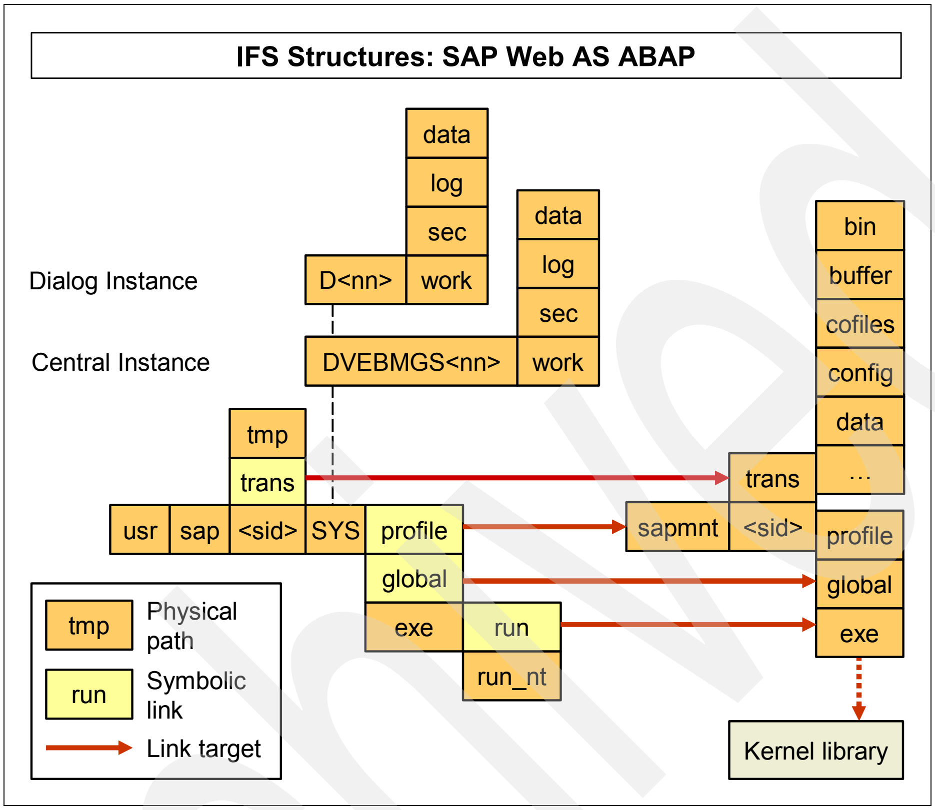

Figure 8-10 Integrated File System structures for the SAP Web Application Server, ABAP stack

Figure 8-10 shows the directories that are used in an ABAP-only installation. Note that different SAP releases can show different structures

The usr and sapmnt directories start in the root. All other directories are setup underneath /usr/sap or /sapmnt. The directory /usr/sap/trans and the subdirectories underneath /usr/sap/*sid*/SYS are symbolic links, so they can point to a common directory in a distributed environment. For example, in a 3-tier installation all remote application servers can share a common profile directory.

Usually the /usr/sap/trans directory is shared between the development, test, and production systems. The kernel library cannot be accessed remotely but needs to reside on the same system as the application server.

Prior to kernel release 6.40, most executables where objects of type *PGM or *SRVPGM are stored in the kernel library. Starting with 6.40, a lot of the executables are moved into directory /sapmnt/*sid*/exe, but there are still some executables left in the kernel library.

The subdirectory /usr/sap/*sid*/SYS/exe/run_nt is used for the executables of a Windows application server for System i models.

Figure 8-11 on page 69 shows some key paths that are used in a combined ABAP and Java SAP WebAS. In comparison to the ABAP-only Figure 8-10, some paths are left out so that the structure fits on one page.

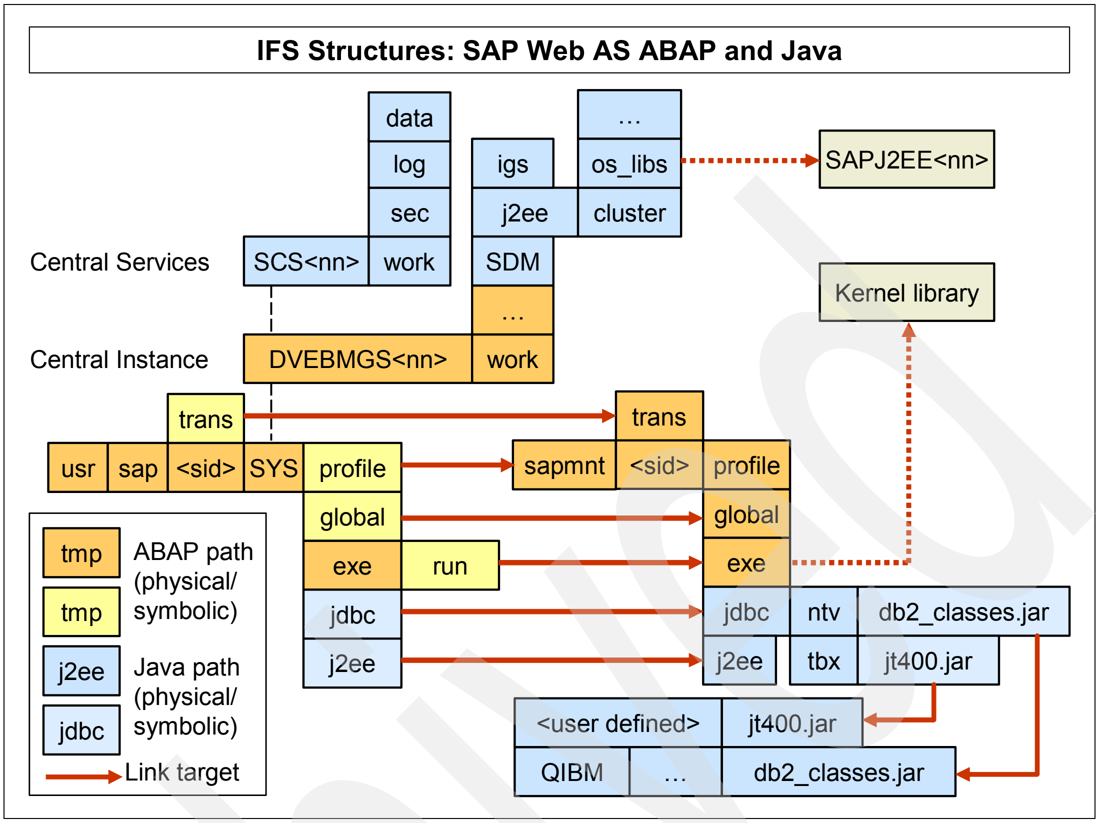

Figure 8-11 Integrated File System structures for the Web Application Server, ABAP and Java

When using Java services, you have a Central Services instance that provides message server and enqueue functions for Java. The J2EE engine itself runs in the ABAP central instance or a dialog instance. The symbolic link /usr/sap/*sid*/SYS/jdbc points to the JDBC drivers.

There are two JDBC drivers available:

- The Native driver (db2_classes.jar)
- The Toolbox or JTOpen driver (jt400.jar)

While the Native driver is always in a fixed location (/QIBM/ProdData/OS400/Java400/ext/db2_classes.jar), the Toolbox or JTOpen driver (jt400.jar) should be downloaded and put to /sapmnt/*sid*/jdbc/tbx/jt400.jar.

Remark: With the initial shipment of Web AS 6.40 the Toolbox or JTOpen driver could be downloaded to any user-defined location. This is no more recommended.

to any user-defined location. The Java instance uses another i5/OS library with executables, which is called SAPJ2EE*nn* (*nn* being the instance number). The path /usr/sap/*sid*/DVEBMGS*nn*/j2ee/os_libs contains native executables and symbolic links pointing to objects in that library.

The following list shows the most relevant Integrated File System paths for Java:

- /usr/sap/*sid*/DVEBMGS*nn*/SDM

  This has data for the Software Deployment Manager (SDM).
- /usr/sap/*sid*/DVEBMGS*nn*/igs

  This has data for the Internet Graphic Server.
- /usr/sap/*SID*/SYS/exe/run

  This contains PowerPC executables from the SAP kernel plus links to objects remaining in the SAP kernel library.
- /usr/sap/*SID*/*instance_name*

  This contains all J2EE data for that instance including administration tools for Java (Admin Tool, Config Tool), and log files, and in case of the central instance, the SDM.
- /usr/sap/*SID*/SCSxx

  This is the Java central services instance directory (Java enqueue, and message server).
- /usr/sap/*SID*/SYS/jdbc

  This contains links to the JDBC drivers.

i5/OS commands to work with the Integrated File System

Use the i5/OS Work Link (WRKLNK) command to view the details of the directory structures on the System i server. The WRKLNK command contains parameters that allow you to specify the level of detail you require to be displayed from the Integrated File System.

The Work Link SAP (WRKLNKSAP) command, contained in the SAP kernel library, also offers options for changing, copying, deleting, displaying, renaming, moving, printing, and working with stream files using numbered options.

Refer to “SAP Integrated File System structure” on page 384 for information about the Integrated File System relative to an SAP system.

For more information about the functions of the QFileSvr.400 file system, refer to *Integrated File System Introduction*, SC41-5711.

8.2.5 Language support

i5/OS allows you to use multiple national languages in parallel on one server. This is accomplished by splitting up all applications into objects that are language independent (for example, programs and service programs), and language dependent objects that contain text elements, and need to be translated (for example commands, message files and data files). IBM products for the System i product line, such as the operating system and additional licensed products, support this concept.

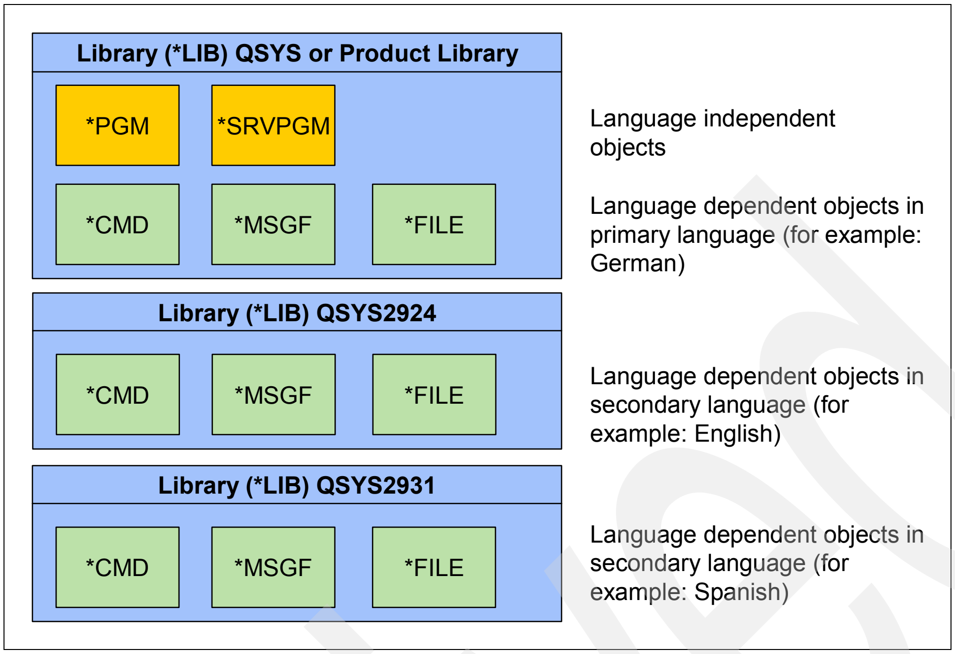

Figure 8-12 Object and file system language support

When you install a System i model for the first time, you should select a primary language. Choose a language your administrators are most familiar with, or English in a multi-national environment. Language-independent objects and the language-dependent objects of the primary language are contained in the QSYS library for the operating system, or in the product library for a licensed program.

In addition, you can install one or more secondary languages. Each language has a four-digit code, such as 2924 for English, 2929 for German, 2931 for Spanish, or 2928 for French. The language-dependent objects of each secondary language are stored in QSYS*nnnn* libraries with *nnnn* being replaced by the four-digit language code.

In order to see screens in a chosen secondary language, the QSYS*nnnn* library must be in the job library list in a position prior to QSYS. You can control this system-wide through the system library list system value (QSYSLIBL), subsystem-wide through the subsystem description, or in a specific job with the CHGSYSLIBL command.

Refer to Chapter 6, “Encoding data used in SAP systems” on page 35 for considerations for encoding data in other languages.

8.2.6 TCP/IP

i5/OS supports the TCP protocol. The Configure TCP/IP (CFGTCP) CL command displays TCP configuration options. Interfaces, routing, host table, services, and domain information are maintained through this interface.

For detailed information, refer to the iSeries Information Center, or the following books:

- *Cool Title About the AS/400 and Internet*, SG24-4815
- *V4 TCP/IP for AS/400: More Cool Things Than Ever*, SG24-5190
- *V5 TCP/IP Applications on the IBM @server iSeries Server*, SG24-6321

8.2.7 Security and authorization

Authority is controlled at the i5/OS user profile level. You can assign user profiles to predefined classes. You can also use group profiles to ease administration. You can give group authority to objects, and then add user profiles to the group to adopt these authorities.

You can set up the i5/OS system running SAP applications to use a single sign-on concept. In this concept, a user can sign on to the network once and is logged to the SAP application by the same action. Refer to *SAP note 138498* for more information.

Refer to Chapter 10, “Users and authorities” on page 115, and the IBM Information Center at:

`http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp`

8.3 System i software products for an SAP implementation

This section discusses the names, features, and handling of IBM licensed programs, licensed program options, and other software products for System i models. Refer to the related IBM product documentation for more details on the Web at:

`http://www-03.ibm.com/servers/eserver/iseries/software/`

In the following sections we discuss the IBM software and licence programs especially needed for an SAP system installation.

8.3.1 Software product names and options

Each IBM licensed program is named with a seven-digit product identifier, such as 5722-SS1 for the Version 5 operating system. Most related products in i5/OS Version 5 have a 5722 prefix. A 5769 prefix is used for OS400 Version 4 software products. A 5799 prefix is used for special purpose products.

The level of a software product is specified as a combination of the version, release, and modification. Licensed program products can have a different level than the operating system that they are installed on. However, there are dependencies. Many licensed products need to be updated to a new release at the same time that the base operating system is upgraded to a later release.

Many products come with multiple options. The *BASE option is always installed. Additional options are optional. Some additional options are free of charge, some require a license fee. Software products usually specify a separate ordering code (feature 5050) to indicate language-dependent objects that are translated. You can install multiple languages for one product.

The LIC does not have language objects but consists only of programming code.

Table 8-1 illustrates the names and option specifications of a few System i licensed programs.

Table 8-1 Name and options of select System i licensed programs

Product ID	Option	Feature		Version	Description
5722999	*BASE	5050	*CODE	V5R4M0	Licensed Internal Code
5722SS1	*BASE	5050	*CODE	V5R4M0	IBM i5/OS
5722SS1	*BASE	2924	*LNG	V5R4M0	IBM i5/OS
5722SS1	1	5050	*CODE	V5R4M0	Extended Base Support

Product ID	Option	Feature		Version	Description
5722SS1	1	2924	*LNG	V5R4M0	Extended Base Support
5722AF1	*BASE	5050	*CODE	V5R2M0	IBM Advanced Function Printing™ Utilities for iSeries (AFP™ Utilities)
5722AF1	*BASE	2924	*LNG	V5R2M0	IBM Advanced Function Printing Utilities for iSeries (AFP Utilities)

The official SAP installation documentation describes which System i licensed programs may be necessary for a specific SAP system installation. The documentation can be found at:

http://www.service.sap.com/instguides

8.3.2 Work with licensed programs

The Work with Licensed Programs (GO LICPGM) menu allows you to display installed software products, install new software products, delete installed software products, and save software products for distributing to other systems.

Figure 8-13 shows you the options of the Work with Licensed Programs menu with the command “**go licpgm**”. Figure 8-13 shows all **go licpgm** options.

```
                                    Manual Install
   1. Install all

Preparation
   5. Prepare for install

Licensed Programs
  10. Display installed licensed programs
  11. Install licensed programs
  12. Delete licensed programs
  13. Save licensed programs

Secondary Languages
  20. Display installed secondary languages
  21. Install secondary languages
  22. Delete secondary languages

Software Agreements
  30. Work with software agreements
  31. Prepare software agreements for transfer of ownership

Redistribution
  40. Create distribution media
  41. Work with installation profiles

Completion Status
  50. Display log for messages

Related Commands
  70. Save and restore commands
  71. Program temporary fix commands
  72. Licensed commands
```

Figure 8-13 Options of the Work with Licensed Programs menu

To work with products that are not in the list, use the following commands:

- DSPSFWRSC - show all installed products
- RSTLICPGM - install a new product
- DLTLICPGM - remove a product
- SAVLICPGM - save a product for distribution

8.3.3 System i software product licensing

Some software products require license keys to enforce a user-based licensing concept. Licensing can be based on:

1. Concurrent users, that is, users that use the product at the same time.
2. Registered users, that is, users that have been registered to the product.
3. The number of processors.

You can have an unlimited number of users, or you can limit the number of users if only few of your users are using a specific software. There is a 70 days grace period after the software is installed during which you can use the software without a valid license. After that, you have to install a valid license key to continue working with the software.

The Work with License Information (WRKLICINF) command is the significant command for license administration. Use this command to see the current number of users and the peak number (for usage type concurrent). This information can help to determine the actual number of licenses you require. You can also use this menu to add the license key received from IBM.

8.3.4 Required IBM software products for an SAP implementation

The SAP installation guide provides a list of required software products and options for your SAP system installation. Table 8-2 shows, for example, that you need additional software (“product IDs”) in addition to the base set if you need Unicode or Java for your SAP implementation.

> **Note:** This information is based on an SAP NetWeaver ´04 implementation at the time of this writing.

Table 8-2 IBM software products required for an SAP system (ABAP and Java)

Product ID	Option	Description	Required for SAP component
5722999	*BASE	Licensed Internal Code	WAS ABAP Non-Unicode
5722SS1	*BASE	i5/OS	WAS ABAP Non-Unicode
5722SS1	1	Extended Base Support	WAS ABAP Non-Unicode
5722SS1	2	Online Information	WAS ABAP Non-Unicode
5722SS1	3	Extended Base Directory Support	WAS ABAP Non-Unicode

Product ID	Option	Description	Required for SAP component
5722SS1	12	Host Servers	WAS ABAP Non-Unicode
5722SS1	13	System Openness Includes	WAS ABAP
5722SS1	21	Extended NLS Support	WAS ABAP
5722SS1	30	Qshell	WAS ABAP
5722SS1	33	Portable App Solution Environment (i5/OS PASE)	WAS ABAP
5722SS1	39	International Components for Unicode	WAS ABAP Unicode
5722 AC3	*BASE	IBM Cryptographic Access Provider 128-bit for iSeries (included in i5/OS at V5R4)	WAS ABAP Java
5722JC1	*BASE	IBM Toolbox for Java	WAS Java
5722JV1	*BASE	IBM Developer Kit for Java	WAS Java
5722JV1	5	Java Developer Kit 1.3	WAS Java
5722JV1	6	Java Developer Kit 1.4	WAS Java

Refer to Chapter 11, “Installation overview” on page 133 for an introduction installation, planning and preparation activities for an SAP system installation.

Many of these products or options are free of charge. However, not all are installed automatically with the i5/OS initial installation.

Table 8-3 identifies additional IBM software products that are useful for the customer.

Table 8-3 Additional IBM software products useful for SAP system environment (ABAP)

Product ID	Option	Description	Required for SAP component
5722WDS	21	iSeries Tools - Application Development	Programming Development Manager (PDM)
5722ST1	*BASE	IBM DB2 Query Manager and SQL Development Kit for iSeries	SQL Tools (STRSQL) SQL Precompiler
5722PT1	*BASE	IBM Performance Tools for iSeries	i5/OS Performance Analysis Tools
5722BR1	*BASE	IBM Backup Recovery and Media Services for i5/OS (BRMS)	Administration of tape libraries and backup policies

8.3.5 Useful and optional software products for SAP systems

The products listed in Table 8-4 are useful for your administration business or in special environments like the Business Information Warehouse application. They are not required for an SAP system installation.

Table 8-4 Useful and optional software products for SAP systems

Product ID	Option	Description	Useful for
5722SS1	26	DB2 Symmetric Multiprocessing	Database performance: Parallel data retrieval, parallel index build, table partitioning (BW)
5722SS1	27	DB2 Multisystem	
5722XE1	*BASE	IBM @server iSeries Access for Windows	Useful for access to the System i server by graphical user interface (GUI). Also for 5250 datastream applications or 5250 emulation. The last one for maintaining the System i server from the PC.
5722DG1	*BASE	IBM HTTP Server for iSeries	HTTP server for i5/OS includes also the IBM HTTP Server (powered by Apache) which you can use for the environment of the SAP Web AS Java.

8.3.6 Program temporary fixes

Keep your operating system and other software current with fixes, or program temporary fixes (PTFs). Fixes play an important part in your system's maintenance strategy. Your maintenance strategy can reduce system downtime, add functionality, and provide optimal availability.

Periodically, problems are discovered in System i programs. IBM issues a fix (also known as a PTF) to correct the problem. Multiple fixes are bundled together to form a cumulative PTF package (CUM tape), which contains certain recommended fixes. Typically, you should install cumulative PTF packages quarterly in dynamic environments, and less frequently in stable ones. You should also consider cumulative PTF packages when you make major hardware or software changes to your environment. Its presence is seen in CL commands such as:

- Display PTF (DSPPTF)
- Send PTF Order (SNDPTFORD)
- Display PTF Cover Letter (DSPPTFCVR)

Fixes, fix groups, cumulative packages, and high-impact pervasive (HIPER) fixes play an important part in your system's maintenance strategy. Your maintenance strategy can reduce system downtime, add functionality, and provide optimal availability.

IBM provides you with a variety of fixes such as:

- Cumulative program temporary fix (PTF) packages
- Fix groups
- High Impact PERvasive (HIPER) fixes

This section introduces you to the different kinds of fixes and discusses when each type is available. Each fix has its own purpose as follows:

- Single fixes

  Single PTFs (fixes) are applied to correct specific reported problems.

- Cumulative PTF packages

  Cumulative program temporary fix packages contain fixes for a given release of the operating system. The fixes contained in a cumulative PTF package are for the LIC and most licensed programs that can be installed on that operating system release.

- Fix groups

  A PTF group or fix group is a name that is used to order and manage a group of logically related PTFs.

- Service packs

  A service pack is a collection of code fixes (not PTFs) for iSeries Access for Windows products that is contained in a single PTF.

In i5/OS, go to the PTF menu using the GO PTF command to see the tasks associated with the PTFs, as illustrated in Figure 8-14.

```
PTF                          Program Temporary Fix
                                                             System:   ITSO-SYS1
Select one of the following:

     1. Load a program temporary fix
     2. Apply a program temporary fix
     3. Copy a program temporary fix
     4. Remove a program temporary fix
     5. Display a program temporary fix
     6. Order a program temporary fix
     7. Install a program temporary fix from a list
     8. Install program temporary fix package
     9. Copy PTF Cover Letter
    10. Display PTF Cover Letter
    11. Delete a program temporary fix
    12. Work with PTF groups
    13. Copy PTF group
                                                                        More...
Selection or command
===>

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=Information Assistant
F16=AS/400 Main menu
```

Figure 8-14 The i5/OS PTF menu

With option 70 you see the related commands for this task.

> **Note:** iSeries Navigator includes *Management Central*, a technology for performing system management tasks across one or more servers at the same time. Management Central provides wizards that simplify fix management. The wizards allow you to easily send, install, and uninstall fixes on multiple systems. You can also use the compare and update wizard to compare a model system to multiple target systems to find missing or extra fixes.

You can find more details about handling of PTFs with i5/OS or with the iSeries Navigator in the IBM information center on the Web using the keyword “PTF” at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp

Install the necessary PTFs for any additional components or license programs installed on your system.

When you apply a software patch for the DW of the SAP kernel, it also installs a copy of the Information Authorized Program Analysis Reports (Info APARs) into the /usr/sap/trans/config Integrated File System directory and then checks if the required PTFs are installed. Use the SAP CHKR3PTF command to perform this function from the SAP kernel.

Additionally, there are group PTFs available that you can call with the WRKPTFGRP command, for example, for the functions listed in Table 8-5.

Table 8-5 For SAP relevant i5/OS PTF groups

PTF Group	Text
SF99530	Cumulative PTF Package C5207530
SF99529	Group HIPER
SF99503	DB2 UDB for iSeries
SF99298	Electronic Service Agent™
SF99269	Java
SF99185	530 Backup Recovery Solutions
SF99139	Performance Tools

You can also use the group PTFs for the following:

- iSeries Access, especially for the Operations Console and iSeries Navigator
- Logical partitioning (LPAR)
- Cross-site mirroring (XSM) for System i clustering
- Hardware Management Console (HMC)

SAP note 83292 lists IBM Informational Authorized Program Analysis Reports (APARs) for various i5/OS releases.

Note: It is very important to have your system on an actual state with all necessary PTFs. Be sure to install the SAP system related PTFs given in the release-specific Informational APAR identified in *SAP note 83292*.

Refer to 19.3.2, “Frequency of PTF application” on page 333 for further information about software maintenance.

8.4 Additional tools for SAP systems

In the past, there were some IBM tools available and delivered with PTFs in the QGPTOOLS library for maintaining SAP applications on System i models. You should use the tools as commands from the QGPTOOLS library after implementing them, as described in *SAP note 68732*. The commands are:

- EDTF
- DSPSTMF
- DSPF
- RCLSPACE
- SQLUTIL

However, since i5/OS V5R3, only the SQLUTIL tool is delivered in the QGPTOOLS PTF. See 8.4.1, “SQLUTIL” on page 79.

the following tools are no longer delivered separately:

- EDTF - now integrated in QSYS
- DSPSTMF - not available anymore, use DSPF instead
- DSPF - now integrated in QSYS
- RCLSPACE - not available anymore

Many System i customers use the following tools because they are available from upgrades from earlier releases on their system:

- EDTF

  The EDTF command allows you to edit an Integrated File System stream file or a database physical file member. It can also be used for editing CL programs. The F4 prompt function is not available.

- DSPSTMF

  Display Stream File (DSPSTMF) was designed to allow you to display stream files within the Integrated File System.

  This is not necessary anymore. Now you can use the DSPF command.

- 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.

- RCLSPACE

  The RCLSPACE command should no longer be used on any system. Use RGZPFM if it is necessary to reorganize a database file.

8.4.1 SQLUTIL

You can use the SQLUTIL command to issue SQL commands interactively. This tool is useful if the i5/OS 5722ST1 product is not installed.

SQLUTIL is available as a PTF, as follows:

- With PTF SI22455 for i5/OS V5R4
- With PTF SI14800 for i5/OS V5R3

After importing the PTFs, the QGPL library contains a backup file named QGPTOOLS. Create a library named QGPTOOLS, and save the objects from the backup file to this library. Add the QGPTOOLS library to the QUSRLIBL system value. Do not save the objects back to the QUSRSYS library as this replaces objects that are already in the library.

SAP note 68732 explains this procedure in further detail. *SAP note 68732* contains, in its attachment, a fix for a known problem (SQL0104 (“Token <...> invalid. Valid tokens: <...>”) with SQLUTIL which sometimes appears.

The SQL utility uses i5/OS native database APIs. For more information about DB2 UDB for i5/OS SQL, refer to the iSeries Information Center and keyword “SQL Concepts”:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp?topic=/sqlp/rbafysqlconcepts.htm

Refer to 8.4.3, “Interactive SQL” on page 81 for a discussion of the STRSQL command and especially to Figure 8-15 on page 82 to compare the functions offered from Interactive SQL (“STRSQL”) with those from SQLUTIL.

Table 8-6 describes the parameters of SQLUTIL.

Table 8-6 The SQLUTIL tool

Parameter	Value	Description
Output	 * *PRINT *OUTFILE	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. Output is displayed. Output is printed with the job's spooled output Output is directed to the database file specified on the OUTFILE parameter.
COMMIT	 *NONE *CHG *CS *ALL *RR	Specify the commitment isolation level for the level of work. No commitment control Uncommitted read Cursor stability Read stability Repeatable read
NAMING	 *SYS *SQL	Specify the naming convention you want to use. Use system naming convention. 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.

Note: The F4 prompt function is not available when using SQLUTIL.

8.4.2 Other tools

Other tools for SAP systems include:

- DSPCURSQL

  This displays the most recently executed SQL statement of a System i job.

  Refer to *SAP note 497608* for more information.

- SAVR3SYS

  This is a composite SAP tool intended to save (backup) a total SAP system. It is often used, but SAP no longer recommends it as of operating system OS/400 V5R1.

  Refer to *SAP note 202593* for more information.

- CHKXDA

  This is used to monitor QXDARECVR jobs that have lost the job that initiated them.

Refer to 20.4, “The CHKXDA SAP tool” on page 349 for a description of the tool and program listing. The CHKXDA tool is also described in *SAP note 450351*.

If you have the latest kernel patches applied, you no longer require the CHKXDA tool.

In the SAP kernel, there are also a lot of tools or commands such as:

- APYR3FIX - Apply an SAP fix for a kernel component
- APYR3KRN - Apply an SAP kernel for an SAP system
- CHKR3PTF - Check PTFs
- CRTSAPUSR - Create an SAP user with multiple options
- PRTSTMF - Print stream file
- RRM - Recursive remove (within Integrated File System)
- WRKLNKSAP - Work with object links. It has more options, as well as other options than WRKLNK.
- WRKPID - Work with process ID (from SAP transaction SM50)

Chapter 12, “The SAP kernel on the System i server” on page 137 describes the kernel commands in more detail.

8.4.3 Interactive SQL

Sometimes it is helpful to enter SQL statements directly, for example, to map long SQL names to short i5/OS names. There are two tools available to enter SQL statements in a 5250 emulation stream:

- STRSQL
- SQLUTIL

The STRSQL command is only available if DB2 Query Manager and SQL Development Kit (product number 5722-ST1) is installed on the System i system.

The STRSQL command allows you to prompt SQL statements similar to the i5/OS F4 command prompt. Prompting also presents a list of available columns after specifying a table, so that it is quite easy to build the complete SQL statement. In addition, you can set session attributes, such as the default commit level, and save the session when leaving it. Figure 8-15 on page 82 illustrates this concept.

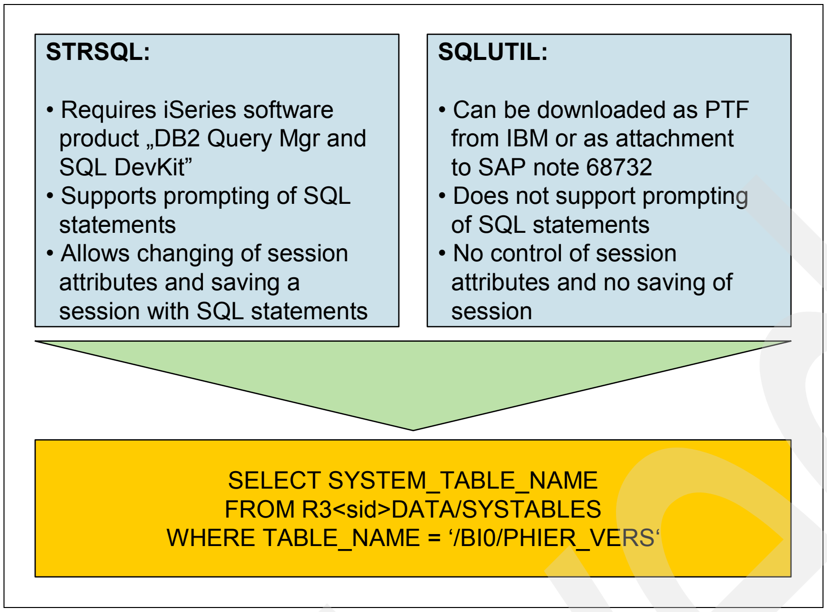

Figure 8-15 Interactive SQL user interface, STRSQL and SQLUTIL

If you have saved a session and you execute the STRSQL command again, you can retrieve the SQL statements from the previous session by pressing the **F9** function key.

Another alternative is the SQLUTIL tool, as discussed in 8.4.1, “SQLUTIL” on page 79. Download it from IBM as a PTF or from SAP as an attachment to *SAP note 68732*. The SAP note contains the PTF number to download, which depends on the operating system release. After entering SQLUTIL, a screen appears where you can enter your SQL statement. There is no F4 prompt support, so you need to know the precise syntax of the statement and the column names. Session attributes cannot be saved and you cannot save your session. By pressing the F9 function key you can retrieve only the last statement in the current session.

8.5 i5/OS commands

i5/OS commands consist of an abbreviated “verb” and an “object”, such as CRT, DSP, and DLT for the verbs, and LIB, JRN, PGM, USRPRF, JRNRCV for objects. The objects supported with each command are listed through the prompting level of each WRK command parameter. Further actions can be taken on each element in the list.

Press **F4** to prompt each parameter to see the parameter options and keywords. Read the help text of the commands to learn about the commands. If you know only the beginning of the command, enter the beginning of the command followed by an asterisk (*), for example, `WRKJ*`. This brings up all the commands that start with the characters WRKJ, such as WRKJRN, WRKJRNA, and WRKJRNRCV.

The GO command brings up a menu (object type *MENU), such as LICPGM, PTF, or MAIN.

The following tables list important commands for a system administrator. Table 8-7 on page 83 lists i5/OS commands commonly used to manage objects.

Table 8-7 i5/OS object management commands

Object management commands	
CRTxxx	Create object, for example CRTUSRPRF, CRTLIB, CRTJRN
DSPxxx	Display object, for example DSPUSRPRF, DSPFD, DSPFFD
WRKxxx	Work with object, for example WRKLIB, WRKF, WRKJRNA
DLTxxx	Delete object, for example DLTUSRPRF, DLTLIB, DLTF
MOVOBJ	Move object into another library
RNMOBJ	Rename object
DSPOBJD	Display object description
CRTDUPOBJ	Create duplicate object (copy object)

Table 8-8 lists commands used with Integrated File System objects.

Table 8-8 i5/OS Integrated File System commands

Integrated File System commands	
WRKLNK	Work with links; similar to Windows: dir or UNIX: ls
ADDLNK	Add link; create symbolic (soft) or hard link
RMVLNK	Remove link; removing last hard link deletes object
MKDIR/MD	Make directory
CHGCURDIR/CD	Change current directory
RNM/MOV/CPY	Rename object or move object or copy object
EDTF	Edit file; simple editor, like Notepad on Windows

Table 8-9 lists i5/OS system control commands that are useful to administer System i resources.

Table 8-9 i5/OS system control commands

System control commands	
WRKACTJOB	Work with active jobs; like Windows task manager or ps -ef
WRKSYSSTS	Work with system status; CPU usage, paging, disk filling
WRKDSKSTS	Work with disk status; disk usage and filling per drive
WRKSYSVAL	Work with system values; control system properties
PWRDWNSYS	Power down the system; with or without restart
SIGNOFF	End interactive job

Table 8-10 on page 84 lists i5/OS system control menus that are useful to administer System i resources.

Table 8-10 i5/OS system control menus

System control menus	
GO LICPGM	Work with licensed programs; install or uninstall
GO PTF	Work with program temporary fixes (code corrections)

Table 8-11 lists commands useful for a System i system administrator in a communications environment.

Table 8-11 i5/OS communication configuration commands

Communication configuration commands	
WRKHDWRSC	Work with hardware resources; entry point to create line descriptions for an Ethernet fine
CFGTCP	Configure TCP/IP menu; important options: 1 Work with TCP/IP interfaces: own IP address, line description, 2 Work with TCP/IP routes: how to get to other systems 10 Work with TCP/IP host table entries: local host table 12 Change TCP/IP domain information: own system name

8.6 i5/OS Portable Application Solutions Environment

The IBM licensed program “i5/OS - Portable Application Solution Environment” (i5/OS PASE) is required for SAP system installations that are based on Web Application Server 6.40 and later. i5/OS PASE is available at no charge as part of i5/OS.

This section introduces i5/OS PASE and explains some aspects of using i5/OS PASE together with an SAP system. The SAP kernel code makes extensive use of the i5/OS PASE with SAP kernel release 6.40 and later. Apart from the requirement to install i5/OS PASE, i5/OS PASE is largely transparent to the user and requires only a minimal amount of disk storage.

8.6.1 i5/OS PASE on System i models

i5/OS PASE is an integrated runtime program that provides simplified porting of applications to i5/OS. i5/OS PASE complements and expands the System i solution portfolio by focusing on rapidly porting applications to the System i platform.

i5/OS PASE exploits the System i processor’s ability to switch between i5/OS and AIX runtime modes within an i5/OS job. This allows applications deployed using i5/OS PASE to run directly on System i models and take full advantage of i5/OS services such as file systems, security, and DB2 UDB for iSeries.

i5/OS PASE is not an operating system. Applications deployed using i5/OS PASE run in a normal i5/OS job and are managed using standard i5/OS operations. Serviceability, backup and restore, and other administrative tasks are performed using standard i5/OS operation and system management facilities.

i5/OS PASE applications can be fully integrated with other System i applications, for example, an Enterprise Resource Planning (ERP) application implemented in ILE, a WebSphere application written in Java, or Lotus Domino. A suite of applications can run together in a job mix or be separated into their own LPARs, depending on the performance and scheduling requirements of the client.

i5/OS PASE programs can now also launch the System i integrated Java virtual machine (JVM™).

You can find further information about i5/OS PASE at:

http://www-1.ibm.com/servers/enable/site/porting/iseries/pase

And for more details, use the keyword "PASE" in the IBM information center on the Web at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp

8.6.2 i5/OS PASE with SAP systems

i5/OS PASE is the latest of several integrated runtime environments that are used in the implementation of SAP applications on System i models since 1994. i5/OS PASE supports the direct processing of PowerPC machine instructions, and does not emulate anything.

SAP kernel source code is sent through a common compiler to generate PowerPC executable instruction streams for both i5/OS and AIX.

> **Note:** *SAP note 863803* describes that SAP kernel release 6.40 and later, that runs on System i models, has changed from using only ILE programs, to using a hybrid kernel that uses ILE and a i5/OS PASE program.

SAP applications run in i5/OS PASE and are integrated with the i5/OS Integrated File System and DB2 UDB for iSeries. These SAP applications can call (and be called by) Java and ILE applications. In general, those programs can take advantage of all aspects of the i5/OS operations environment, such as security, message handling, communication, and backup and recovery.

When SAP applications run in i5/OS, the SAP applications:

- Run as an i5/OS job
- Use i5/OS system functions, such as the file systems, security, and sockets
- Use the i5/OS integrated database

Install i5/OS PASE as Option 33 of the Operating System (5722-SS1), when moving to SAP kernel release 6.40. i5/OS PASE is a no-charge option beginning with OS/400 V5R2. It is available on the standard set of operating system installation CDs.

i5/OS PASE includes a subset of the locales provided by AIX, supporting both 32-bit and 64-bit applications. i5/OS PASE locales are installed as symbolic links in directory /usr/lib/nls/loc. The i5/OS PASE product installs only the locale objects that are associated with the language features that are installed on i5/OS. Installing additional locales is an optional installation step. Order and install additional i5/OS language features if you need locales that are not included with the language features.

The full name of any i5/OS PASE locale includes a code set name, which equates to the Coded Character Set Identifier (CCSID) shown in the table. Some locales also have a short name that exclude the code set part of the name. Any locale with a name ending in @euro uses the euro as the currency symbol. Most i5/OS PASE locales are shipped with i5/OS language feature codes. Only locales in the base *CODE load and locales for installed language feature codes exist on a particular i5/OS system.

For information about i5/OS PASE Globalization and i5/OS PASE locales, use the keywords "PASE Globalization" and "PASE locales", respectively, in the iSeries Information Center at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp?topic=/rzalf/rzalfintro.htm

8.7 Special operating system settings for SAP applications

This section describes special i5/OS settings unique for SAP system installations.

8.7.1 Multiple SAP systems on one server

For platforms other than the System i family, we recommend that you have dedicated servers for each SAP production system.

For System i installations, by using LPAR systems, you can separate your production systems on one System i server in a very comfortable way. The intention is to separate the production system from side effects of other applications. A separate server or LPAR has the advantage of the physical separation of dedicated hardware resources. Maintenance and backup functions, as well as performance influences are avoided as much as possible.

Use separate logical partitions to dynamically move resources between the SAP system installations on the different partitions. Beginning with i5/OS V5R3, dynamic LPAR with uncapped processor resource partitioning is supported. Other partitions can use the free capacity if one partition does not use the assigned processor resources.

You can install more than one SAP system in a partition. This allows a direct sharing of hardware resources between the SAP systems. This also affects the response times of the other systems if one partition does not use the assigned processor resources.

By default, all SAP subsystems are assigned to the *BASE Pool. Place the databases of the SAP system in separate user ASPs to reduce disk space management considerations. Create a shared pool for each SAP system to separate main memory usage. Assign memory to the IBM subsystem (R3_nn) of this SAP instance (nn).

In an SAP system, different types of work processes run alongside one another and compete for system resources, such as CPU time. In general, optimal response times for online users are more critical than response times for spool processes and processes running in the background.

Chapter 5, “System i virtualization” on page 31 provides more information about LPAR.

8.7.2 Adjusting the run priority of SAP work processes

There are two options to adjust the priority for executing work processes:

- Adjusting the priority in the job description (JOBD)
- Using relative priority classes

The priority value specified in JOBD must be between 15 and 80. If this is not the case, a warning is given and the priority specified in JOBD is set to 20 (for priority less than 15), or to 80 (for priority greater than 80).

Adjust the run priority of the jobs in an SAP system according to *SAP note 45335*, especially if there are performance problems in the online mode, high response times for batch work, or spool sessions that run simultaneously.

Before you can use the new priority classes, you must activate the Dynamic Priority Scheduler (QDYNPTYSCD) system value by setting the value to **1**. An IPL is required to make the change take effect.

i5/OS allows you to change job priority. By default, the SAP processes dispatcher, gateway, message server (and icman - SAP internet connection manager) have the execution priority of **12**, and all other SAP work processes have an execution priority of **20**, as defined in the job class description. To help optimize performance for the production system, job priority can be controlled via the job class description. Specify a priority of **40** for the test system, and set priority **20** for production.

Four priority classes are available to define the priority of an SAP instance:

1. HH - Very high priority

   The value of the priority of execution from the job class description is reduced by eight during the runtime of these processes.

2. H - High priority

   The value of the priority of execution from the job class description is reduced by four during the runtime of these processes.

3. M - Normal priority

   The value of the priority of execution from the job class description is not reduced during the runtime of these processes.

4. L - Low priority

   The value of the priority of execution from the job class description is increased by four during the runtime of these processes.

You can then specify priority classes for the processes Update (UPD), Batch (BTC), and Spool (SPO). All the other processes have a fixed relative priority. The critical processes Message Server (MS), Dispatcher (DISP), Enqueue (ENQ) and Gateway (GW) are assigned priority class HH, for example. The user cannot change this.

Table 8-12 displays which type of priority classes are permitted for which work processes. The work processes in italics can be changed by the user.

Table 8-12 Priority classes for the SAP work processes

Priority class ID	Work process
HH	MS, DISP, GW, ENQ
H	UPD
M	UPD, DIA, UPD2, BTC, SPO
L	BTC, SPO

The Dispatcher, Message Server and Enqueue processes (known as jobs to i5/OS installations) are assigned to class HH with dialog and update2 processes of class M. This cannot be changed by the customer within the SAP system.

So the critical SAP work processes Dispatcher, Message Server and Enqueue processes automatically receive a higher priority. All SAP work processes have a default priority, as specified in the job class description and described at the beginning of this chapter.

You can assign batch processes and spool processes to class M or class L. By default, they belong to class M. The assignment of the update, batch and spool processes can be changed by the customer within the SAP system.

You can adjust the execution priority of the SAP work processes by SAP transaction **RZ10** or **RZ11** in the SAP instance profile:

- Set the following parameters in the SAP instance profile to determine the priority of the respective processes as follows. M is the default.

```
rdisp/prio/upd = [ H | M ]
rdisp/prio/btc = [ M | L ]
rdisp/prio/spo = [ M | L ]
```

- Set the instance profile parameter to control the Gateway priority as follows:

```
rdisp/prio/gwrd to M or L
```

The following feature is available from Web AS 6.40 patch level 119 onward.

- To avoid other SAP systems being affected by the CPU resources required when starting an SAP system, use the instance profile parameter **rdisp/prio/wp/start = [L|LL|LLL]**. In this case, the priority of the work processes are set to [24|28|32] during the startup phase. After the system is up the priority is set to the default values or to the values defined by the rdisp/prio-parameters given above.

8.8 Java

Different versions of the Java Development Kit (JDK™) for iSeries are installed as options of IBM Developer Kit for Java (5722-JV1) licensed program. JDK 1.3 is installed as option 5, and JDK 1.4, is installed as option 6. Both are currently used by SAP applications.

8.9 Java database connectivity

Java on the System i server supports two Java database connectivity drivers:

- Native Driver

  The Native Driver can be used when the Java server and the database are on the same System i server in the same partition. It provides better performance than the Toolbox Driver.

- The Toolbox Driver

  The Toolbox Driver can be used on any platform that connects to a System i database.

You can find a detailed description of the two driver types and usage considerations in *SAP note 654800*. *SAP note 809693* explains the driver locations. *SAP note 826449* explains the necessary changes to switch JDBC drivers in an existing SAP WebAS installation.

9

SAP storage and database considerations

The System i server is an advanced and complex server capable of running different types of applications simultaneously, creating a demand for versatile storage. System i models can serve as storage hosts for integrating Windows servers, for Linux partitions, AIX partitions, and Domino, Java, and WebSphere application requirements, all at the same time and usually from the same or multiple i5/OS storage pools. Capacity is a must for any storage solution on System i models that support traditional DB2 UDB online transaction and batch applications, as well as other native System i applications, such as image and file serving.

This chapter discusses System i hardware disk configuration, internal and external storage, disk space, and storage functionality. It also discusses the use of journals for SAP databases and applications in a typical SAP system installation. Refer to “Journal management” on page 377 for an introduction to journal concepts.

9.1 Hardware disk configuration

Disk units are assigned to a disk pool on a storage unit basis. The system treats each storage unit within a disk unit as a separate unit of auxiliary storage. When a new disk unit is attached to the system, the system initially treats each storage unit within it as thought it were not configured. You can add these unconfigured storage units to either the system disk pool, basic disk pool, or independent disk pool of your choosing.

When you add an unconfigured storage unit to a disk pool, the system assigns a unit number to the storage unit. The unit number can be used instead of the serial number and address. The same unit number is used for a specific storage unit even if you connect the disk unit to the system in a different way.

The storage unit that is addressed by the system as unit 1 is always used by the system to store Licensed Internal Code and data areas. The amount of storage that is used on unit 1 is quite large and varies depending on the configuration of your system. Unit 1 contains a limited amount of user data. Because unit 1 contains the initial programs and data that are used during an IPL of the system, it is also known as the *load source unit*.

When you store data in a disk pool, the data is stored along all disk units that belong to the disk pool. The system reserves a fixed amount of storage on units other than unit 1. The size of this reserved area is 1.08 MB per unit, reducing the space available on each unit by that amount.

An SAP application can easily utilize several hundred gigabytes and in large instances, terabytes of disk space. A hardware disk configuration must take into account the following aspects:

- The expected amount of data for all SAP applications and other applications running on the system or in a partition. Consider the monthly or yearly growth of data. Also consider any additional SAP systems and other data requirements which are planned to be added to the same machine or partition.
- The number of disk arms the data requires depends on the size of a single disk drive. You must determine whether to install more disk drives with a lower disk size capacity, or fewer disk drives with a higher disk size capacity. The first solution is more expensive and can require an extra expansion unit, while the second solution can result in fewer disk arms, which can then have a negative influence on I/O performance.
- Implementing data protection, such as RAID-1, RAID-5, or RAID-6 (new with i5/OS V5R4) is in part dependent on the I/O processor (IOP) and disk controller implementation selected.

  RAID-6 protects against the simultaneous failure of two disk units. Logically, the capacity of two disk units is dedicated to storing parity data. In practice, parity data is spread among multiple disk units. The minimum number of disk units in a parity set is four and the maximum is eighteen. When a RAID-6 parity set is started, all the disk units contain parity data.
- The number of RAID sets configured. There are different options for setting up a RAID-5 configuration that influence the amount of installed data and performance capabilities, for example, performance options such as capacity, balance, or RAID across a Small Computer System Interface (SCSI) configuration. This configuration option can be set through Management Central.
- The frequency of data access and utilization of data. For optimal performance, appropriate and efficient access to the data stored on disks is required. Consider the characteristics and number of “channels” to the disks, when planning for disk capacity, for example fibre

channels. The application load (workload) can be calculated by the I/O rate, especially the I/O rate per disk.

- Determine the data requirements for applications outside of the SAP application. Add this data requirement to your storage considerations.

9.1.1 Internal and external disks

Application data on System i models is stored on disk devices (disk drives), commonly referred to as “disks”. Disks are attached to the System i server internally or externally.

Note: Disk units are also referred to as Direct Access Storage Device (DASD).

Internal disks

Internal disks are located within the system unit of the System i server or within a direct attached expansion unit, so that the operating system has direct access to the data on the disk. Disk drives can utilize a limited number of slots in the system unit and a limited number of slots in the expansion unit.

Typically, the location of the disk drive for the load source is in a predefined disk slot in the system unit.

Note: As of October 2005, IBM System i5 and *eServer* i5 models support a load source disk unit that is located in an IBM TotalStorage® subsystem. This enables a direct boot from a storage area network (SAN) attached IBM TotalStorage subsystem offering, such as the IBM TotalStorage ESS, DS6000™, and DS8000™ families.

The enablement of a SAN-based boot support for the load source device is available with the installation of a #2847-Fibre Channel i5/OS Load Source Boot for SAN feature.

There are specific slot restrictions on where IOPs or disk controllers are located for the system unit and expansion towers, as well as placement rules and considerations. There is a limit to the number of disk drives that can be placed in the system unit. Furthermore, some disk controllers in the system unit provide better performance than others, or can be restricted on where they are placed (expansion unit versus system unit). Consider the performance relative to the disk configuration in the system unit and the RAID requirements.

Disks that are located in an expansion unit are attached to the System i server with a special cable. The expansion units support different disk capacity configurations, for example, a maximum of 8 disks, 12 disks, or 45 disks. Expansion units can be concatenated, enabling a wide range of internal disks connected to the System i server.

See *PCI Placement Rules for the IBM @server i5 and iSeries Servers with i5/OS V5R3*, REDP-4011 for a discussion of placement considerations.

Each System i model has a limit on the number of disks it supports. The maximum amount of disk capacity is a typical specification documented for each System i model. Refer to the summary chart in the *IBM @server i5 and iSeries System Handbook: IBM i5/OS Version 5 Release 3 October 2004*, GA19-5486 or *IBM System i5 Handbook IBM i5/OS Version 5 Release 4 January 2006*, SG24-7486, to determine the configuration maximums for the number of disk arms and disk capacity supported per model.

Note: The maximum amount of disk capacity referenced is based on a calculation of the product of the maximum number of disk drives supported in the model and highest capacity disk drives supported.

High capacity disk drives are allowable for an SAP installation as long as there are sufficient disk arms to ensure good performance. Should you consider replacing your current disk units with a fewer number of higher capacity drives, be aware that reducing the number of disk units (arms) in your environment can lead to performance issues.

A typical SAP installation needs hundreds of gigabytes (GB) for storing data. The System i server allows up to 381 terabytes (TB) of internal disk capacity with i5/OS V5R4.

Internal disks are accessed by an IOP or disk controllers. It is important to consider the performance of the disk controller and the controller's cache capabilities for optimum performance. And different IOPs support different RAID implementations. Not every disk controller supports RAID-5.

External Storage

External disks are located in an external storage subsystem (not in the external tower of an expansion unit). The use of external disk storage becomes more important in an SAP system environment with the advancement of external storage functionality and performance in System i environments as compared to early external storage products.

External storage can provide more flexibility and options. External disks provide high-availability functions that are convenient, especially for SAP customers, for example, FlashCopy® and Metro Mirror.

The primary disadvantage of external disks is that there are more administration and planning tasks to do because you have to manage a separate (storage) server with the external disk drives. Adding external capacity requires careful planing. You also have to control the Storage Area Network (SAN) located between the external disk storage server and the System i system unit.

Performance of external disks

External disks can be configured in such a way that performance is not worse than using internal disks. The appropriate Fibre Channel configuration and suitable external main storage for buffering and caching can provide optimum performance.

Mirrored protection

Mirrored protection is a software availability function that protects data from being lost because of failure or because of damage to a disk-related component. Data is protected because the system keeps two copies of data on two separate disk units.

When a disk-related component fails, the system can continue to operate without interruption by using the mirrored copy of the data until the failed component is repaired.

When you start mirrored protection or add disk units to a disk pool that has mirrored protection, the system creates mirrored pairs using disk units that have similar capacities. The overall goal is to protect as many disk-related components as possible. To provide maximum hardware redundancy and protection, the system attempts to pair disk units that are attached to separate IOAs, IOPs, buses, and expansion units.

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 System i model and is a part of the Licensed Internal Code.

Remote mirroring support allows you to have one mirrored unit within a mirrored pair at the local site, and the second mirrored unit at a remote site. For some systems, standard disk unit mirroring is the best choice; for others, remote disk unit mirroring provides important additional capabilities. You must evaluate the uses and needs of your system, consider the advantages and disadvantages of each type of mirroring support, and decide which is best for you.

Device parity or mirroring

The level of mirrored protection determines whether the system keeps running when different levels of hardware fail. The level of protection is the amount of duplicate disk-related hardware that you have. The more mirrored pairs that have higher levels of protection, the more often your system is usable when disk related hardware fails.

You may decide that a lower level of protection is more cost effective for your system than a higher level. When determining what level of protection is adequate, you should consider the relative advantages of each level of protection, such as:

- Levels of protection

  The level of mirrored protection determines whether the system keeps running when different levels of hardware fail.

- Disk-unit level protection

  Mirrored protection always provides disk-unit level protection because the storage units are duplicated.

- Input/output bus level protection
- Input/output adapter level protection
- Input/output processor level protection
- Bus-level protection
- Expansion-unit level protection

> **Note:** Device parity protection is not a substitute for a backup and recovery strategy.
>
> Device parity protection can prevent your system from stopping when certain types of failures occur. It can speed up your recovery process for certain types of failures. But device parity protection does not protect you from many types of failures, such as a site disaster or an operator or programmer error. It does not protect against system outages that are caused by failures in other disk-related hardware (such as I/O adapter (IOA)s, disk I/O processors (IOP), or a system bus).
>
> You should protect all the disk units on your system with either device parity protection or work with mirrored protection. This prevents the loss of information when disk failure occurs. In many cases, you can keep your system operational while a disk unit is being repaired or replaced.

9.2 System i integrated storage

The System i product line offers a sophisticated, high-performance, automated, integrated storage solution. System i integrated storage offers large capacities (up to 288 TB) through multiple, high-speed controllers, scalable, jumbo write cache (757 MB of nonvolatile storage

(NVS) write cache for every 15 disk units), RAID and mirroring performance configuration options, as well as 35 GB, 70 GB, or 140 GB 15 K rpm disk options.

An i5/OS integrated storage management implementation can eliminate the need for specialized storage skills, such as storage subsystem, SAN, and maintenance management. A fully automated storage management design facilitates the functions that include automatic distribution of data (data spreading) across disk pools for excellent performance, minimal operator involvement, complete graphical storage management tools, program temporary fix (PTF) management, and concurrent maintenance.

Integrated i5/OS disk storage offers RAID data protection with several performance options (performance, capacity, balance, and RAID across SCSI) and complete data, non-volatile shared (NVS) cache and I/O path storage redundancy via i5/OS disk mirroring (for example, RAID-10).

The benefits of an integrating a storage management solution with i5/OS include:

- Automatic balance of storage and spread of data
- Optimized RAID-5 configuration with the flexibility to exploit built-in system level mirroring (similar to RAID-1) for parallel reads and writes
- Extensive write cache capabilities
- Self-optimized workload input and output (I/O) through self-caching architecture, with expert cache enhancing it further
- Innovative SAN-like storage management through virtual I/O hosting to Linux applications that is important for a 3-tier SAP installation with a Linux application server
- AIX and Windows server environments which are important for a 3-tier SAP installation with a Windows application server

Unlike SAN, which offers a great amount of flexibility when it comes to attaching storage at varying distances using Fibre Channel, the internal storage is required to stay within the limits of the SCSI or high speed loop (HSL) boundaries (also known as RIO-G). For optimum performance, HSL-2/RIO-G cable lengths should not exceed 15 meters.

9.3 IBM TotalStorage solutions

IBM TotalStorage solutions bring new and advanced storage capabilities to System i models. It allows more storage consolidation and flexibility in an enterprise environment, such as multiserver connectivity, fully redundant hardware including NVS cache, RAID-5 or RAID-10 protection, Copy Service solutions, and advanced storage allocation capabilities. In many client environments, it is a strategic direction that all servers utilize SAN-based storage. This can be done with an Enterprise Storage Server® (ESS) or TotalStorage Server implementation, which are designed to meet these needs.

When an IBM TotalStorage subsystem is attached to an System i model, it becomes an automatic extension of the overall storage pool and the i5/OS automated storage management capabilities. The IBM TotalStorage devices also benefit from the i5/OS automated data distribution, graphical storage management, availability capabilities, integrated backups, and the ability to support all types of System i applications. In fact, Copy Service functions exploit the i5/OS solutions such as Independent Auxiliary Storage Pools (IASP) to enable business continuity solutions and increase overall system availability.

Note: The #2847 PCI IOP for SAN Load Source and its capability to support external load source boot functions is supported only on IBM System i5 servers and only with a Hardware Management Console (HMC) attached to the IBM ESS Model 800, DS6000, and DS8000 series. Earlier ESS models such as E10, E20, F10, and F20 do not support the placement of external bootable load source disk units.

Consolidation of storage begins with compatibility. The IBM TotalStorage DS family supports a broad array of IBM and non-IBM server platforms, including IBM z/OS, z/VM®, i5/OS, OS/400, and AIX 5L operating systems, as well as Linux, HP-UX, Sun™ Solaris™, Novell NetWare, UNIX, and Microsoft® Windows environments. Consequently, you are free to choose your preferred vendors and run the applications required to meet your enterprise's requirements, while extending your previous IT investments along with SAP applications.

Virtualization greatly assists storage asset consolidation. Virtualization software solutions are designed to logically combine separate physical storage systems into a single, virtual storage pool, thereby offering dramatic opportunities to help reduce the total cost of ownership (TCO).

The IBM TotalStorage DS family supports enterprise class data backup and disaster recovery capabilities. As part of the IBM TotalStorage Resiliency family of software, IBM TotalStorage FlashCopy point-in-time copy capabilities back up data in the background, while providing you with nearly instant access to information about both the source and target volumes. Metro Mirror, Global Mirror, and Global Copy capabilities allow the creation of duplicate copies of the application data at remote sites.

Systems retain frequently accessed or high-value data in one storage server and archive less valuable information in a less costly storage server. This allows systems such as the IBM TotalStorage DS family to help improve the management of information according to its business value from the moment of its creation to the moment of its disposal.

In order to correctly plan the implementation of the ESS Model 800, DS6000, or DS8000 series with IBM System i family, you must plan for the solution that you want to implement. The solutions cover a range of options, based on the application and overall business continuity requirements.

Figure 9-1 on page 96 outlines the planning steps required to determine the "right" storage configuration for your business requirements.

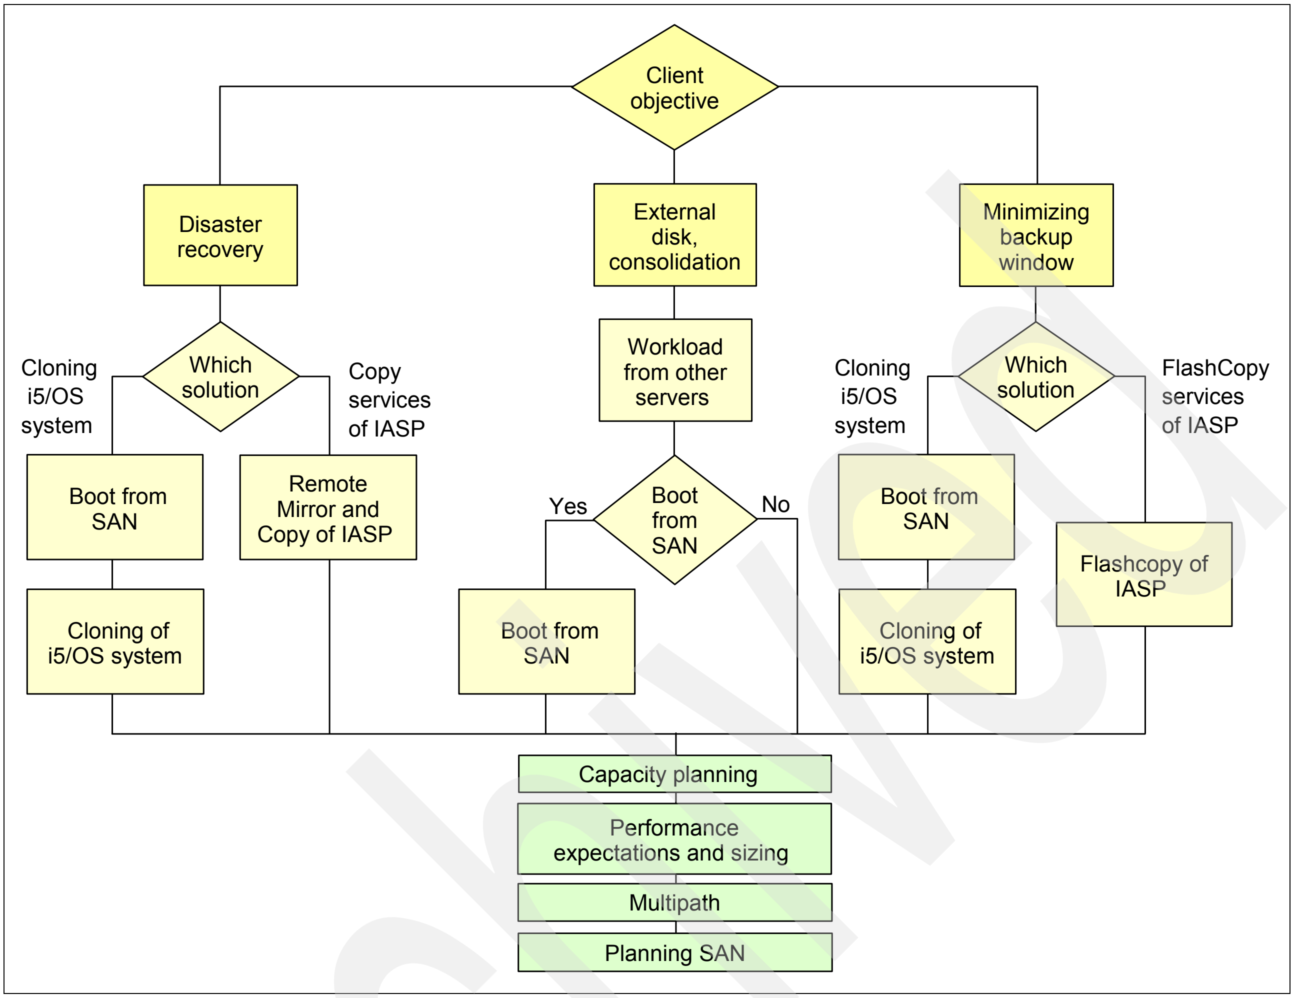

Figure 9-1 Planning for storage

Understand the following success factors when planning your storage solution:

- Evaluate the supported hardware configurations.
- Understand the minimum software and firmware requirements for i5/OS, HMC, system firmware, and microcode for the ESS Model 800, DS6000, and DS8000 series.
- Understand additional implementation considerations such as multipath I/O, redundancy, and port setup on the storage subsystem.

Each solution has its own benefits and considerations. See *iSeries and IBM TotalStorage: A Guide to Implementing External Disk on @server i5*, SG24-7120 for more information.

9.3.1 Fibre Channels for external disks in an SAP system environment

External storage is connected to System i server with a Fibre Channel Adapter. It is critical to provide enough Fibre Channel Adapters in the system if you are connecting to external storage. If there are too few adapters, the link between the server and the external disk will be a performance bottleneck. For a description of how to plan, configure, install, and run external disks and storage attached with Fibre Channels, refer to *iSeries in Storage Area Networks A Guide to Implementing FC Disk and Tape with iSeries*, SG24-6220.

The information includes the following details:

- The fundamentals of internal and (especially) external disk storage
- An introduction to SAN and the iSeries server
- Making the right connections to SAN
- Planning for external disks
- Creating iSeries external storage
- Recommendations for volume size
- Multipath for iSeries

9.4 Planning for disk storage in an SAP installation

To help prepare the necessary disk drives before you install an SAP application on a System i configuration, review the following questions:

- Will internal or an external disk be installed?

  When internal disks are used, distinguish between their location in the system unit or in a System i expansion unit.

- What amount of disk space is sized for each SAP application as well as for any additional non-SAP applications?

  Consider all potential SAP system environments, such as production, customizing, test, quality assurance, sandbox, backup, cluster, and training.

  Include monthly data growth and the growth of data coming from additional users, additional functions, or additional applications within the planned life cycle of the disk drives in the calculation. Remember that performance significant degrades if the disk utilization is higher than 80%.

- How many disks are used and what is the capacity of each disk?

  Disks with high capacity need fewer disk slots, which means fewer expansion units are required. Disks with higher capacity are more expensive per GB compared to disks with lower capacity.

  Also consider the rotation per minute (rpm) of the disks. Disks with 15k rpm are more suitable for production SAP system environments than disks with a lower rpm. Avoid using disk devices with different rpm (different architecture) within any Auxiliary Storage Pool (ASP).

  If fewer disk drives with higher capacity represent the required disk space capacity, check the number of disk arms. For optimal performance, use disk devices with the same capacity within one ASP.

  Only disk drives of the same size can be placed in one RAID set. At least three disk device are required for a RAID-5 set.

  > **Note:** Avoid using disks with different capacities of more than one architecture within one ASP. For example, you can combine 17.5 GB disks with 35 GB disks, but not with 70 GB disks.

- Which disk controller is used and how many disk controllers are required?

  Consider the type and number of disk controllers you can use. Performance characteristics differ. For example, some controllers perform about seven times faster than others.

  Define the characteristics of a RAID-5 set and the type and number of disks assigned to the set. The configuration also affects the selection of the controller.

9.5 Basic principles of SAP databases and SAP systems

In this section the relationship and the handling of the databases and their journals of standard and special SAP installations is discussed, including:

- SAP databases and SAP journals in standard installations
- SAP databases with shared or common SAP journals
- Multiple Components in one database (MCOD) and their journals

9.5.1 SAP databases and SAP journals in standard installations

Figure 9-2 illustrates a typical implementation of an SAP installation with a ABAP and a Java instance:

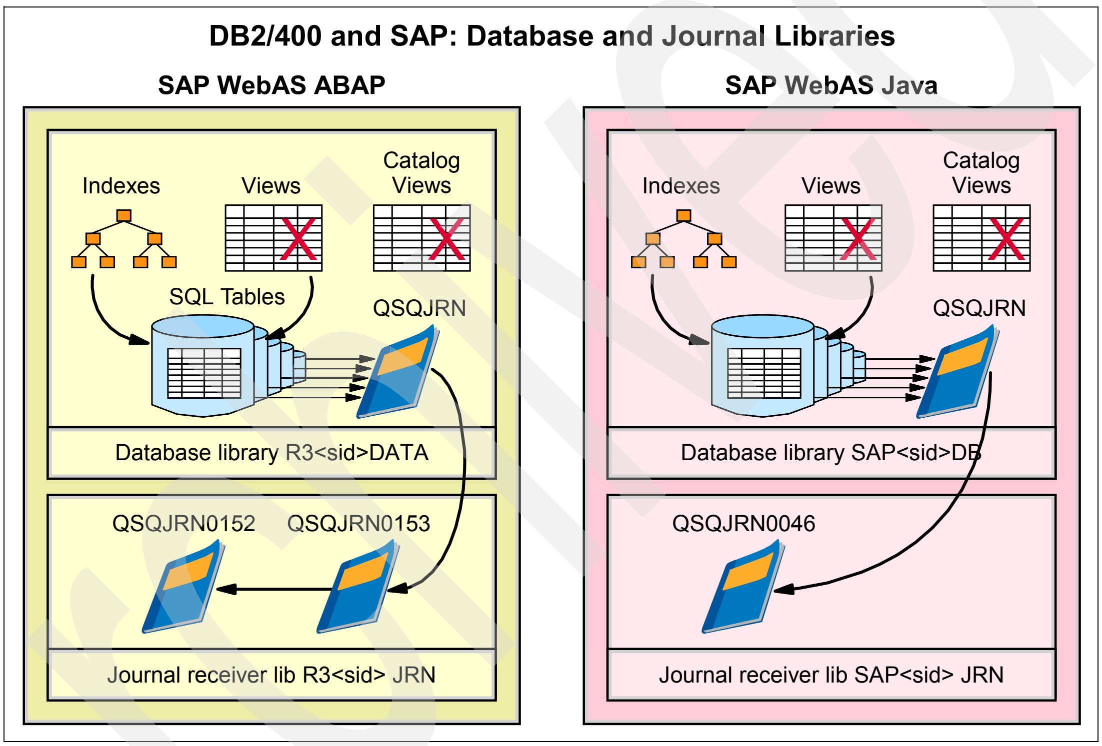

Figure 9-2 SAP ABAP and Java database and SAP journals in a standard SAP installation

SAP WebAS ABAP and SAP WebAS Java store their data in different SQL schemas. With DB2 UDB for iSeries, a schema is implemented by a library. There is a R3*SID*DATA database library for SAP WebAS ABAP and an SAP*SID*DB database library for SAP WebAS Java. Substitute the SAP system ID for *SID*. With a combined installation (SAP WebAS with both ABAP and Java), you have both libraries.

The database libraries contain all tables, indexes and views of the schema. In addition, they contain the catalog views (SYSTABLES, SYSINDEXES, for example) and the QSQJRN journal. All tables in the database library are journaled to that journal. The actual entries of the journaling function are written to the journal receivers, which are located in different libraries. This is done to make it easier to store the journal receivers in a separate ASP. The journal receiver library for SAP WebAS ABAP is called R3*SID*JRN. The journal receiver library for SAP WebAS Java is called SAP*SID*JRN.

9.5.2 SAP databases with shared or common SAP journals

As of i5/OS V5R3, you have the option to journal both the ABAP database library and the Java database library into the same journal. See Figure 9-3.

A shared journal can be useful to administer the ABAP and the Java instance of one SAP application (for example, SAP ECC). Additionally you can use a shared journal for multiple SAP applications like SAP ECC and SAP BI. In both cases, you can handle the procedure for Roll Forward or Roll Backward of the journal receiver easier, because you have not to synchronize the starting points, ending points, or both of two different journals.

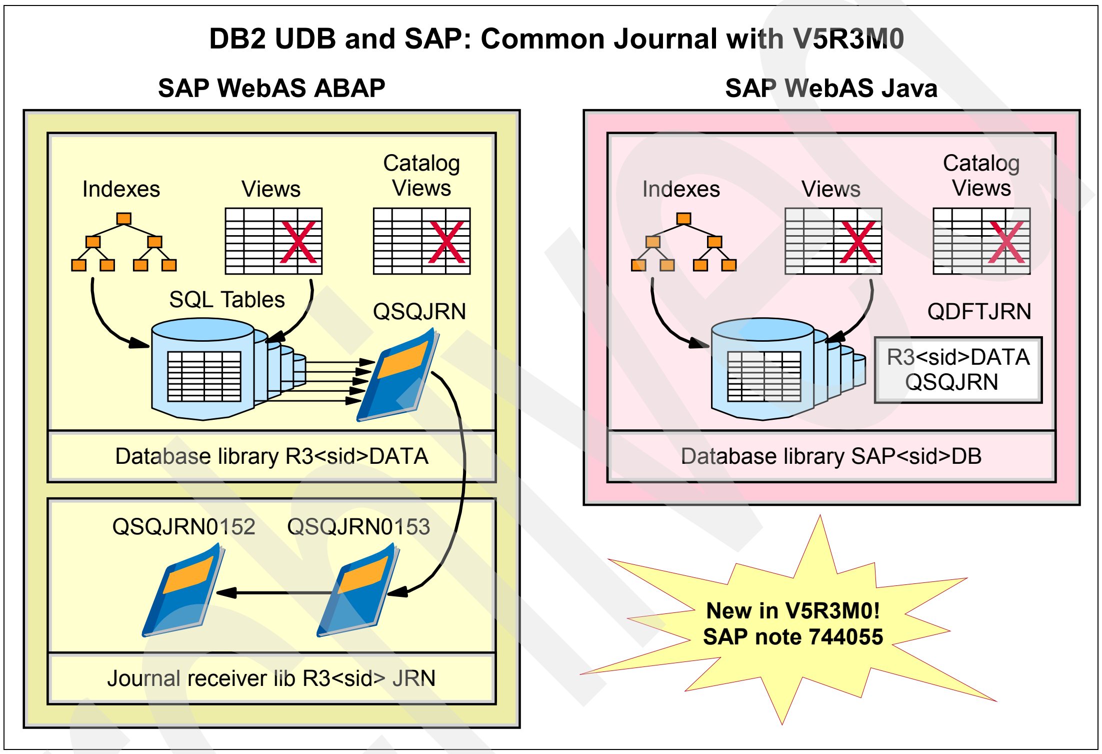

Figure 9-3 SAP databases with shared or common journals

The implementation of a shared or common journal has two advantages:

1. You only need to backup and maintain one set of journal receivers.
2. In case of a recovery, journal changes can be applied on both the ABAP and the Java server to exactly the same point in time. If two independent journals are used, you must apply the changes separately and can only be applied to a specific point in time (the granularity is one second). In this scenario, it is possible that a ABAP transaction is considered complete and restored, while a related transaction in Java (that was committed a fraction of a second later) is considered incomplete and rolled back.

All data libraries that are to use a shared journal must be contained in the same Auxiliary Storage Pool (ASP) on the same System i server.

There is a limit to the number of objects a journal can handle. The maximum depends on the level of the operating system:

- 10,000,000 for i5/OS V5R4
- 250,000 for i5/OS V5R3

The number of internal objects cannot be precisely determined in all cases. The significant “conversion rate” for SAP applications are listed in Table 9-1.

Table 9-1 Conversion rates between System i objects and internal objects; limits of a journal

System i object	Internal object
Table, physical file	2
Primary key	1 *
Index	1 - 2 *
View	1
Objects	1
SQL sequence	1
Data area, data queue	1
Other	0
* Depends on whether system managed access path protection (SMAPP) is used.	

It is mostly a 1:1 conversion. However, for each table or physical file you need two internal objects, so you are to double the number of the related tables to get the number of the entries in the journal. If in doubt, contact your local IBM Support Center.

For more detailed informations refer to the journal limits topic in the IBM Information Center at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp

As the number of objects in the journals increases, the time to restore objects also increases, even if only parts of objects are to be restored.

> **Note:** The time it takes **APYJRNCHG** or **APYJTRNCHGX** to apply changes increases in proportion to the number of tables for which changes have been journaled for. This is not very significant for ABAP and Java as a unit, but should be considered when thinking about MCOD with two and more SAP systems involved. The same consideration applies to **RMVJRNCHG**.

For a discussion, refer to the “Journal limits” topic in the Information Center at: http://publib.boulder.ibm.com/pubs/html/as400/infocenter.html.

Refer to the Information Center for i5/OS operating system limits at: http://publib.boulder.ibm.com/pubs/html/as400/infocenter.html

And refer to *OS/400 Maximum Capacities - V5R2*, REDP-0204 for more operating system limits.

Setting up a shared journal

To convert an existing configuration, for example, to reconfigure an existing installation consisting of two data libraries (R3*SID*DATA and SAP*SID*DB), proceed as follows:

1. Select which libraries should contain the shared journal (referred to as the primary library). R3*SID*DATA is selected for our example.
2. Download the latest AS4FIXFILE patch from the SAP Service Marketplace.

3. Shut down all relevant SAP systems.
4. In all secondary libraries, create a data area called QDFTJRN, which contains the name of the new standard journal. Ensure that:
   a. The first 10 letters contain the name of the primary library
   b. The second 10 letters contain the name of the new default journal
   c. The next 5 letters contain the value “*FILE”

   Example:

   ```
   CRTDTAARA DTAARA(SAPSIDDB/QDFTJRN) TYPE(*CHAR) LEN(30) +
      VALUE('R3SIDDATA  QSQJRN     *FILE     ')
   ```

   Make sure that *SID*OWNER *USE has rights to this object.
5. Call AS4FIXFILE to convert the journals for existing objects into the new journal. For example:

   ```
   AS4FIXFILE DBLIB(SAPSIDDB) ENDJRN(*YES) STRJRN(*YES) FILES(*ALL)
   ```
6. Delete the journals and journal receivers of the secondary libraries from the system. They are no longer required.

The activation of this feature is explained in SAP note 744055. For more information, see “Auxiliary storage pools” on page 378.

9.5.3 Multiple components in one database and their journals

The term Multiple Components on One Database (MCOD) stands for a type of SAP system installation used on some databases. An SAP system can now be installed into an existing SAP system database, so that you have multiple SAP components in one single database. Since database administration is costly and difficult, some savings can be obtained on other databases by using this application work-around to force the installation of multiple components into the same database.

While the installation of multiple SAP systems in one database requires special processing on other databases, it is not a significant difference for single installations on System i models, because the integrated i5/OS database is already capable of handling multiple schemas and allowing them to be journaled in a consistent manner. Each component has its own R3*SID*DATA and R3*SID*JRN library, residing in the same database. If you are running at i5/OS V5R3 or later, all data can be journaled to the same journal, as shown in Figure 9-4 on page 102.

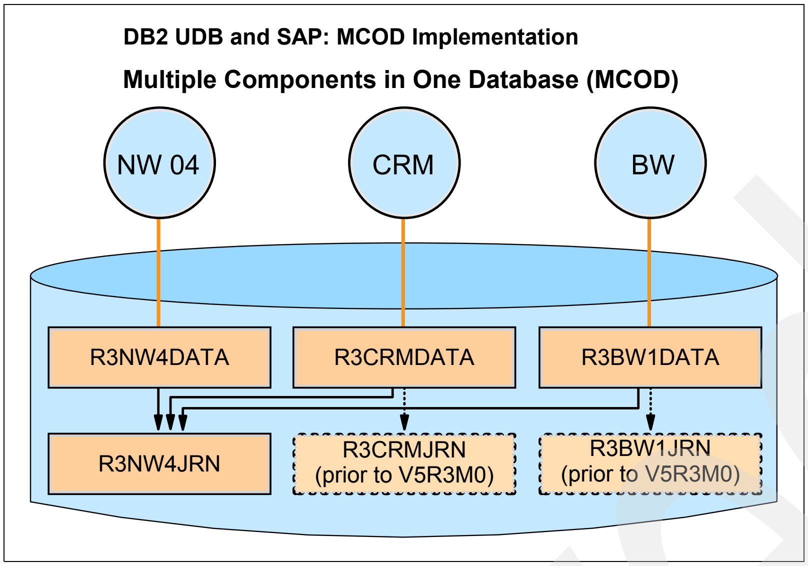

Figure 9-4 SAP databases in a MCOD implementation

When sharing a common journal, there are some limits concerning the number of the objects and more. See 9.5.2, "SAP databases with shared or common SAP journals" on page 99 for more details.

More information about the MCOD concept can be found at:

http://service.sap.com/mcod

9.6 SAP applications and journal management

In the SAP installation guides, SAP recommends the use of a separate ASP for the journal receiver. It is necessary to highlight this recommendation since you have to decide where you want to store the journal receivers for the SAP installation.

What follows is an introduction to the SAP journaling mechanism for i5/OS solutions. The fundamentals of journaling are described in "Journal management" on page 377. See "Auxiliary storage pools" on page 378 for further information about ASPs.

9.6.1 The SAP journal and SAP journal receiver

SAP applications are designed to work strictly with commit control cycles. Therefore, the i5/OS *JRN journal and *JRNRCV journal receiver objects are required.

A journal in i5/OS is named QSQJRN with an object type of *JRN. A journal is related to a collection or to say it better because of the shared journal feature there is no 1:1 relationship anymore, each collection has got a associated default journal which is part of the local SQL catalog for a collection. It is created in the catalog with the CREATE COLLECTION command (for example in the SAP installation procedure).

The following figures show the options and parameters for the following significant i5/OS commands.

- Create Journal receiver - **CRTJRNRCV**, as shown in Figure 9-5 on page 103.

- Create Journal - **CRTJRN**, as shown in Figure 9-6 on page 104 and Figure 9-7, “The Create Journal CRTJRN command (2/2)” on page 104.
- Change Journal - **CHGJRN**, as shown in Figure 9-8 on page 105 and Figure 9-9 on page 105.

```
                     Create Journal Receiver (CRTJRNRCV)

Type choices, press Enter.

Journal receiver . . . . . . . . . JRNRCV
  Library  . . . . . . . . . . .                   *CURLIB
ASP number . . . . . . . . . . . . ASP            *LIBASP
Journal receiver threshold . . . . THRESHOLD      1500000
Text 'description' . . . . . . . . TEXT           *BLANK


                          Additional Parameters

Preferred storage unit . . . . . . UNIT           *ANY
Authority  . . . . . . . . . . . . AUT            *LIBCRTAUT





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

Figure 9-5   The Create Journal Receiver command CRTJRNRCV

```
                              Create Journal (CRTJRN)

 Type choices, press Enter.

 Journal  . . . . . . . . . . . . JRN
   Library  . . . . . . . . . . .                  *CURLIB
 Journal receiver . . . . . . . . JRNRCV
   Library  . . . . . . . . . . .                  *LIBL

                                                   *LIBL
 ASP number . . . . . . . . . . . ASP            *LIBASP
 Journal message queue  . . . . . MSGQ           QSYSOPR
   Library  . . . . . . . . . . .                  *LIBL
 Manage receivers . . . . . . . . MNGRCV         *SYSTEM
 Delete receivers . . . . . . . . DLTRCV         *NO
 Receiver size options  . . . . . RCVSIZOPT      *MAXOPT1
                          + for more values
 Minimize entry specific data . . MINENTDTA      *NONE

 Journal caching  . . . . . . . . JRNCACHE       *NO
                                                                      More...
 F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
 F13=How to use this display       F24=More keys
 Parameter JRN required.                                                    +
```

Figure 9-6 The Create Journal CRTJRN command (1/2)

```
                              Create Journal (CRTJRN)

 Type choices, press Enter.

 Manage receiver delay time . . . MNGRCVDLY      10
 Delete receiver delay time . . . DLTRCVDLY      10
 Fixed length data  . . . . . . . FIXLENDTA      *JOBUSRPGM
                          + for more values
 Text 'description' . . . . . . . TEXT           *BLANK

                                                                       Bottom
 F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
 F13=How to use this display       F24=More keys
 Messages pending on other displays.
```

Figure 9-7 The Create Journal CRTJRN command (2/2)

```
                           Change Journal (CHGJRN)

Type choices, press Enter.

Journal  . . . . . . . . . . . . . JRN
  Library  . . . . . . . . . . .                   *LIBL
Journal receiver:                JRNRCV
  Journal receiver . . . . . . .                 *SAME
    Library  . . . . . . . . . .
  Journal receiver . . . . . . .
    Library  . . . . . . . . . .
Sequence option  . . . . . . . . . SEQOPT         *CONT
Journal message queue  . . . . . MSGQ           *SAME
  Library  . . . . . . . . . . .
Manage receivers . . . . . . . . MNGRCV         *SAME
Delete receivers . . . . . . . . DLTRCV         *SAME
Receiver size options  . . . . . RCVSIZOPT      *SAME
                          + for more values
Journal state  . . . . . . . . . JRNSTATE       *SAME

                                                                      More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
Parameter JRN required.
```

Figure 9-8 The Change Journal command CHGJRN (1/2)

```
                           Change Journal (CHGJRN)

Type choices, press Enter.

Minimize entry specific data . . MINENTDTA      *SAME

Journal caching  . . . . . . . . JRNCACHE       *SAME
Manage receiver delay time . . . MNGRCVDLY      *SAME
Delete receiver delay time . . . DLTRCVDLY      *SAME
Fixed length data  . . . . . . . FIXLENDTA      *JOBUSRPGM
                          + for more values
Text 'description' . . . . . . . TEXT           *SAME

                                                                       Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
Messages pending on other displays.
```

Figure 9-9 The Change Journal command CHGJRN (2/2)

For more details, use the <**F1**> Online Help for these commands within i5/OS.

Different journals can exist in i5/OS, including the journals for the operating system itself located for example in QRECOVERY, QUSRSYS or QSYS2 library. There can be two different journal receivers in two different places, for example, libraries, attached to the same journal.

The SAP journal is located in the SAP database in the R3*SID*DATA library for an ABAP instance and in SAP*SID*DB for a Java instance.

A journal identifies the time the data is written to the journal receiver: the data before the data is changed (*BEFORE) or the data before and after the data change (*BOTH). A *BOTH option provides forward and backward recovery and is used in an SAP implementations.

The journal itself does not contain application data. The journal refers to journal receivers which are i5/OS objects of the type *JRNRCV, which contain the data and which can be located in a different library. The default location of the SAP ABAP journal receivers is the R3*SID*JRN library. The SAP Java journal receivers are stored in SAP*SID*JRN library. The default naming convention of the SAP journal receivers is QSQJRNnnnn, assigned after the completion of the SAP system installation, where nnnn is a consecutive number beginning with 0001.

A specific threshold size (in kilo bytes (KB)) for each object is defined for the journal receiver. When this size is reached, a new journal receiver must be attached to the journal. Either the operating system can manage the detach of the old and attach of the new receiver, or it can be done manually. If the system generates the new journal receiver its name `QSQJRN`*`nnnn`* is increased by 1 (`QSQJRN`*`nnnn`*`+1`). If a threshold value for the journal receiver is specified, management of the receivers can be automated.

You must also decide if you should delete the old and detached journal or keep it in the R3*SID*JRN library.

Usually the SAP installation program sets up the journal automatically so you do not need to be concerned about the settings of the SAP journal and the SAP journal receiver. If manual operations are required, it is helpful to be familiar with the significant parameters.

Set the Manage receivers parameter (MNGRCV) of the journal to *SYSTEM, so that the system generates a new journal receiver automatically when the threshold of the current one is reached. The new journal receiver is then attached. The system detaches the old one automatically. If the Delete receivers parameter (DLTRCV) is set to *NO, the detached journal receivers are kept, so that they can be saved. If point-in-time recovery is not needed (for example in a test system), you can also set this parameter to *YES. The system then deletes detached receivers automatically, so that you do not need to manage the space they occupy. The “Receiver Size Options” parameter (RCVSIZOPT) must have a value of *MAXOPT2 to allow journal entries for Long Object Management (LOB) type columns, as they are used by SAP applications. The value *RMVINTENT is recommended to keep the receiver sizes smaller.

Also consider:

- By default, SAP applications journal only tables. Indexes are not journaled.
- The user who writes the journal receiver is the owner of the SAP work process, so it is ***SIDnn***. Therefore, you cannot directly determine the journal receiver changes made by the user to the database.

9.6.2 Default SAP application *JRN and *JRNRCV settings

The SAP installation procedure provides the following default settings of the journal and the journal receiver:

- Journal QSQJRN in library R3*SID*DATA

  CRTJRN JRN(R3*SID*DATA/QSQJRN):

  – One journal receiver

    Option: **JRNRCV(R3*SID*JRN/QSQJRN0001)**

  – Operating system management of journal receiver

    Option: **MNGRCV(*SYSTEM)**

  – Minimize entry specific data

    Option: **MINENTDTA(*FILE)**

  – Keep journal receiver after journal switch

    Option: **DLTRCV(*NO)**

  – Another option for the size of the journal receiver:

    Option **RCVSIZOPT(RMVINTENT *MAXOPT2)**

- Journal receiver in library R3*SID*JRN:

  CRTJRNRCV JRNRCV(R3*SID*JRN/QSQJRN001):

  – Which ASP to write journal receiver

    Option: **ASP**

  – Threshold of the journal receiver

    Option: **THRESHOLD(200000)**

9.6.3 Manual procedure to set up journal logging

Under exceptional circumstances, you may need to manually change the journal configuration of an SAP installation. This can be true, for example, if customers run a high availability solution that requires different settings.

> **Note:** SAP installations, upgrades, migrations, or system copies that use standard SAP tools, such as SAPINST or R3SETUP, do not require manual activities. They implement the default settings.

First, decide how the data and journals are to be distributed on your system, in particular, across the available ASPs. Additionally, decide whether to use a shared journal for multiple databases.

Procedure to set up the journal logging manually

This description assumes that the data to be logged is already in the corresponding data libraries in the correct ASP (R3*SID*DATA or SAP*SID*DB). The SAP system is stopped and the data is saved if necessary. Thus, there are no dependencies on obsolete journal recipients. Perform the following steps:

- Log in as user *SID*OFR on the System i server.
- Clean up journals that are in R3*SID*DATA. Use the following command to check journal status:

  ```
  WRKJRNA JRN(R3SIDDATA/QSQJRN)
  ```

- If necessary, follow the subsequent steps to remove the journal and obsolete journal recipients from the system:

a. Use the following command to ensure that you have current backup copies for all objects that are logged against R3*SID*DATA/QSQJRN:

```
WRKJRNA JRN(R3SIDDATA/QSQJRN)
```

With SAP Release 7.0, or if the statements in Note 744055 are implemented, R3*SID*DATA/QSQJRN can be the default journal for objects outside of R3*SID*DATA, for example, for objects in SAP*SID*DB when you use a “shared journal”.

b. Use the following command to terminate any table log:

```
ENDJRNPF FILE(*ALL) JRN(R3SIDDATA/QSQJRN)
```

c. If the journal exists, use the following command to remove it:

```
DLTJRN JRN(R3SIDDATA/QSQJRN)
```

d. Use the following command to remove the obsolete journal recipients:

```
DLTJRNRCV JRNRCV(R3SIDJRN/QSQJRN*)
```

e. Remove the corresponding library using the following command:

```
DLTLIB LIB(R3SIDJRN)
```

- Create the R3*SID*JRN library in the desired target ASP with the correct authorizations using the following command:

```
CRTLIB LIB(R3SIDJRN) ASP(asp) AUT(*EXCLUDE) CRTAUT(*EXCLUDE)
```

- Create the initial journal recipient using the following command:

```
CRTJRNRCV JRNRCV(R3SIDJRN/QSQJRN0001) ASP(*LIBASP) THRESHOLD(200000)
```

- Then create the journal itself in R3*SID*DATA using the following command:

```
CRTJRN JRN(R3SIDDATA/QSQJRN) JRNRCV(R3SIDJRN/QSQJRN0001) MNGRCV(*SYSTEM) DLTRCV(*NO)
RCVSIZOPT(*RMVINTENT *MAXOPT2) MINENTDTA(*FILE)
```

- Start the journal log using the following command:

```
AS4FIXFILE DBLIB(R3SIDDATA) ENDJRN(*YES) STRJRN(*YES) FILES(*ALL)
```

If a Java part exists in the SAP installation, manage it as follows:

- Treat the Java part as a special case if the installation is a 6.40 SAP installation on OS/400 V5R2. Proceed in the same way with the Java portion that includes the Java data library SAP*SID*DB, the journal SAP*SID*DB/QSQJRN, and the journal recipient library SAP*SID*JRN, to log objects in SAP*SID*DB separately against SAP*SID*DB/QSQJRN.
- In all other cases, implement the recommends described in SAP note 744055 so that objects in SAP*SID*DB are also logged against R3*SID*DATA/QSQJRN. In this case, a separate journal recipient library named SAP*SID*JRN is not required. You can create a special data area in SAP*SID*DB instead, using the following command:

```
CRTDTAARA DTAARA(SAPSIDDB/QDFTJRN) TYPE(*CHAR) LEN(30) VALUE('R3SIDDATA QSQJRN *FILE')
```

- Use the following command to start the journal log:

```
AS4FIXFILE DBLIB(SAPSIDDB) STRJRN(*YES) FILES(*ALL)
```

For more details about this procedure see SAP note 828354.

9.6.4 Journal switch

A journal switch can be done manually. That is, you can manually close the old journal receiver and attach a new journal receiver with the option *GEN within the CHGJRN command.

If journal receivers reach the defined size in normal SAP operation, in standard SAP installations the system automatically creates a new journal receiver and attaches it to the journal. The old journal receiver is detached from the journal, but remains on the system in the R3>SID>JRN library.

You can save or delete these detached journal receivers so that the disk space they use is available again, avoiding an overflow of the ASP.

See 9.7.4, "ASP overflow" on page 113 for a discussion of ASP overflow.

9.6.5 Reduction in journal receiver size

An SAP system can send a large quantity of information to the journal receiver, depending on the activities, because all changes to the database have to be maintained. In each case, a complete before image and after image of the record being changed is saved, so that you can carry out both a rollback and a forward recovery.

Using the Edit Recovery for Access Paths command (EDTRCYAP), you can display a list of the access path recovery times for the system and the ASP. The lower this value is set, the more information the system saves to meet the specified time.

The EDTRCYAP command is a very simple and comfortable tool. It is illustrated in Figure 9-10.

```
                        Edit Recovery for Access Paths                ITSO-SYS1
                                                             29.05.06  19:58:19
 Estimated system access path recovery time  . . . :             57   Minutes
 Total not eligible recovery time  . . . . . . . . :              0   Minutes
 Total disk storage used . . . . . . . . . . . . . :          0,204   MB
 % of disk storage used  . . . . . . . . . . . . . :          0,000

 Type changes, press Enter.
   System access path recovery time  . . .   60          *SYSDFT, *NONE, *MIN,
                                                           *OFF, Recovery time
   Include access paths  . . . . . . . . .   *ALL        *ALL, *ELIGIBLE

             ------Access Path Recovery Time------  -----Disk Storage Used-----
 ASP               Target          Estimated             Megabytes     ASP %

   (No user ASP configured or information not available)




                                                                         Bottom
 F3=Exit   F4=Prompt   F5=Refresh        F12=Cancel
 F13=Display not eligible access paths   F14=Display protected access paths
```

Figure 9-10 The i5/OS command EDTRCYAP (Edit Recovery for Access Paths)

The Edit Recovery for Access Paths display shows a list of access path recovery times for the system and for auxiliary storage pools (ASP) that are currently active on the system. The information shown reflects the current target and estimated access path recovery times.

Access path and access path recovery information is shown also for all auxiliary storage pools (ASPs) if they are active and if system-managed access-path protection has not been turned off (the system access path recovery time value is not *OFF). The ASP information is not shown when the system ASP (ASP 1) is the only ASP that is active.

On this list, you can type changes for the target access path recovery times for the system and for the ASPs. You can also view updated access path recovery status information. Where necessary, adjust the system access path recovery time with the command EDTRCYAP.

The access path changes because of system-managed access-path protection (SMAPP) are journaled to the same journal receiver as table changes unless the RCVSIZOPT parameter is used. For more information about SMAPP, see the IBM Information Center at:
http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp

Use the RCVSIZOPT parameter of the CHGJRN command to limit the amount of permanently saved data. Enter the following command for this purpose:

```
CHGJRN JRN(R3SIDDATA/QSQJRN) JRNRCV(*GEN) RCVSIZOPT(*RMVINTENT)
```

There are three factors that influence the journal receiver size

- System access path recovery time (in minutes) of the **EDTRCYAP** command
- Receive size option set to *RMVINTENT of the **CRTJRN** or **CHGJRN** command
- Minimize entry specific data set to *FILE of the **CRTJRN** or **CHGJRN** command

For more details see the <**F1**> Online Help to the related commands.

9.7 ASP considerations for SAP journal management

Typically, the SAP journal is stored in the SAP database in R3*SID*DATA for the Advanced Business Application Programming (ABAP) instance and in SAP*SID*DB for the Java stack. In typical installations, these databases are stored in the System ASP (SYSBAS) which is the ASP 1.

If you implement SAP applications with an IBM cluster solution (cross-site mirroring (XSM)), the SAP databases are located in the Independent ASP and not in SYSBAS.

The SAP journal receivers are stored in the R3*SID*JRN library for the ABAP and in the SAP*SID*JRN library for the Java part. If you use the Java environment normally, implement a shared journal for the ABAP and the Java database with only one R3*SID*JRN library for all journal receivers. Therefore, in the following sections, we discuss only the journal receiver library named R3*SID*JRN.

9.7.1 Using a separate ASP for the journal receiver

In the SAP installation guides, SAP recommends the setup of a separate ASP (ASP2) for the journal receiver. The following statements help clarify the reasons to install the journal receiver in a separate ASP:

In the SAP installation guides, SAP recommends the setup of a separate ASP (ASP2) for the journal receiver. The following section clarifies this statement.

- Better availability of the journal receiver when using a separate ASP
- Better performance with mirrored disks (RAID-1) for the journal receiver library R3*SID*JRN or SAP*SID*JRN.

> **Note:** Remember that “mirroring” on a System i is not limited to the disk units (RAID-1) but to other devices like IOP, bus ore more.

The SAP databases are normally protected with device parity (RAID-5) not with mirrored devices (RAID-1).

If the underlying journal receiver of the SAP journal is not available, the SAP system immediately encounters errors.

The availability of the journal receivers is improved when stored into a separate ASP. It is more likely to lose two disks in a large RAID-5 array for the application data rather than in a small separate ASP for the journal receiver. In a mirrored disk implementation, both related disks must fail in order for a commit control function to be lost.

The use of a separate ASP is a relatively expensive feature. Disk space must be separated into parts. The advantage of the evenly distributed data over all disks is lost. When you divide the disks for the application data and journaling, consider a separate ASP for the journal receiver and also the following factors:

- The size of the disk drives
- The number of disks and number of disk arms for performance reasons
- The type of data protection required (RAID-1, RAID-5, RAID-6)
- The location to place the disks (system unit, expansion unit, external disks)
- The type of controller

9.7.2 General recommendations

The primary reason to implement journal receivers in a separate ASP is to help guarantee the availability of the journal receiver. Other recommendations are:

- There is no need to set up a separate ASP2 for the journal receiver if you have a high availability (HA) solution for your production database.
- Only one separate ASP2 is required for all journal receivers of the different databases if you have multiple production SAP systems on one machine or one LPAR setup.

  It is not advisable to install one separate ASP for each production SAP system. Therefore, you do not have ASP2, ASP3, and so on, on one server or in one logical partition.

  > **Important:** Ensure that your system is sized effectively for journal receivers so that you do not run into an ASP overflow. You must have enough disk arms for performance reasons for the journal receiver.

- For customizing, test, and quality assurance, and other non-productive SAP systems, you do not need a separate ASP for the journal receiver. Implement all journal receivers of all non-productive SAP systems on one machine or logical partition in the system ASP.

  See “Auxiliary storage pools” on page 378 for a further discussion.

  > **Note:** Independent ASP cannot overflow. The SAP system, as well as the System i server, stops if the Independent ASP reaches its limit. Consider this if you use Independent ASP for the *JRNRCV.

For more information, see SAP note 654801.

9.7.3 Sizing the SAP journal receiver

The required disk size for the journal receiver depends on multiple factors, especially the following:

- The strategy of backing up the journal receiver

  This helps to determine when you can delete and when you can save the detached journal receiver

- The amount of data that is written to the journal receiver

If you save and delete the detached journal receivers on a daily period, then you have to spare journal receiver data for only one day.

Estimate the required disk space for the SAP journal receiver

There is no formula to calculate the amount of data for the situation listed previously. However, consider the following list when estimating the size of the SAP journal receiver. Take the maximum for each item:

- As a rule of thumb, count approximately one twentieth to one tenth of the database size as the needed disk space for the journal receiver.
- Take into account the daily growth of the database. Calculate about five to ten times the monthly growth of the database for the reserved journal receiver.
- The daily activity of the database is the main influence on the size of the journal receiver. Calculate approximately 100 MB to 200 MB per active user to get an approximate estimate of the size for the journal receiver.

Mass procedures such as client copy, large batch input, long-running batch jobs, and other data-consuming batch jobs are not calculated. Reserve appropriate disk space in the journal receiver for such unpredictable situations.

Note: The factor for mass procedures is not to be underestimated. Critical situations can occur due to the following factors:

A journal receiver can technically only be removed from the system when its data doesn't participate in any open transaction any more. The high degree of parallelism and long running transactions typical for conversion processes can require that journal receivers are temporarily assigned, so give the reserved journal receiver disk space a buffer.

As an example for the previous criteria, an SAP installation with a 400 GB database and a monthly growth of approximately 5 GB used by approximately 200 active users requires a journal receiver size of approximately 20 GB to 50 GB. The required disk space is calculated as follows (see the bullets from the section above):

- For the first bullet: 1/20 to 1/10 of 400 GB = 20 GB to 40 GB
- For the second bullet: 5 times to 10 times of 5 GB = 25 GB to 50 GB
- For the third bullet: 0.2 GB multiplied by 100 to 200 users = 20 GB to 40 GB

The 40 GB to 50 GB figure is the net amount of the disk size for this example. This can be maintained by a single disk unit. However, you should not configure only one unit for journal receivers because that single unit will become a bottleneck. Use a minimum of two mirrored units or three units in a RAID configuration. Your transaction volume will determine whether this is sufficient or whether you require more disk arms for the journal receiver ASP.

9.7.4 ASP overflow

When a user ASP reaches its limit (ASP threshold), the journal receiver in that ASP overflows into the system ASP (ASP1). An error message is displayed in the QSYSOPR message queue, but the SAP system does not stop (as it is the case with other databases). If the system ASP fills, the system stops. This error message remains even when you correct the overflow situation and the journal receiver is again in ASP2.

A manual initial program load (IPL) of the server can allow an operator to logon and check for objects that can be deleted to free up enough temporary storage to continue.

The procedure to recover from an overflowed ASP is as follows:

1. Use the CRTJRNRCV command to create a new (temporary) journal receiver in a library other than R3*SID*JRN, for example, in the database R3*SID*DATA.
2. Use the CHGJRN command to attach the SAP journal to this new (temporary) journal receiver.
3. Use the SAVOBJ or SAVLIB command to save all detached journal receiver in ASP2 or the library R3*SID*JRN.
4. Use WRKJRNA command to delete all journal receivers of the SAP journal QSQJRN in the database library R3*SID*DATA except the last receiver, which is the new (temporary) journal receiver.

   Click the function key **F15** and select option 4 (delete).
5. Use the DLTLIB command to delete the journal receiver library R3*SID*JRN.
6. Signed on as *SID*OFR, use CRTLIB to create a new journal receiver library in ASP2 (CHGOBJOWN to R3OWNER).
7. As *SID*OFR, use CRTJRNRCV command to create a new journal receiver in this new journal receiver library with the known attributes (CHGOBJOWN to R3OWNER).
8. Attach the journal to this new journal receiver (journal switch) using the CHGJRN JRN command.
9. Use SAVOBJ or WRKJRNA command to save and delete the temporary journal receiver from step 1. Then click the function key **F15** and select option 4 (delete).

   Save and delete all the journal receivers, except the last.

10

Users and authorities

This chapter describes user authorization relative to SAP implementations on the System i server. It will help you to understand authorization problems or problems with user settings and file or object permissions.

This information is based on *SAP notes 173579, 834218,* and *46063*.

10.1 Authorization concepts

The following sections introduce the SAP authorization concepts to the reader. The objects and user profiles handled by SAP system environments together with their usage, attributes and authorities are explained.

10.1.1 Objects, owners, and users

The SAP installation procedure creates i5/OS objects and assigns these objects to a specific owner, to specific owner groups and assigns other attributes to them. When SAP system environments are implemented, normally there is no need to change these objects relative to the authority configuration.

It is helpful to understand the authority attributes of these i5/OS objects if you run in authority problems and you have to correct these settings.

> **Note:** Take care to have the right authorities, group profiles and other *USRPRF attributes when you work as an SAP system administrator. In most cases the `<SID>OFR`, `<SID>ADM` or a user profile with the same attributes is the correct user profile. You can avoid a lot of problems if you mostly use the `<SID>OFR` user profile.
>
> In most cases the i5/OS `QSECOFR` user profile is not the right user profile when you work as an SAP system administrator.
>
> Take care if you use an i5/OS user profile that is not created by an SAP application, including the IBM-supplied `QSECOFR` user profile.
>
> Always use the right user profile when you work on System i servers in the SAP system environment.

By design, it should not be necessary to change the attributes of user profiles that are created during the SAP installation. For example, authority problems can be caused by incorrect ownership of programs in the kernel library. Some programs in the kernel library need to adopt authority of the `QSECOFR` user profile. Sign on as `QSECOFR` and execute the **`FIXR3OWNS`** command in the kernel library to ensure proper authority settings.

> **Note:** Do not modify the i5/OS user profiles created by the SAP system.
>
> Do not give the `<SID>OFR SPCAUT(*ALLOBJ)` authority, even for resolving problems. If you change the `<SID>OFR` user profile you can experience more sophisticated problems. Use other means to correct your problem.
>
> In critical situations where your support organizations do not react fast enough it can be a valid *short-term work-around* to grant `*ALLOBJ` special authority to the `<SID>OFR`.

SAP objects

SAP objects include all libraries starting with the characters **`R3<SID>*`**, as well as the libraries named **`R3400`**, **`R3SYS`**, **`R3WRK<instance>`** and the **`R3<rel>OPT`** kernel libraries.

Beginning with the mySAP Technology application and its Internet Architecture in 1999, the SAP NetWeaver application with its Enterprise Services Architecture (ESA) in 2003, and the total mySAP Business Suite application, there are a growing number of libraries which contain SAP objects. Libraries starting with the characters “R3” or “SAP” contain SAP objects.

In the Integrated File System, the SAP objects are located in the subdirectories of `/usr/sap/*` or `/sapmnt/*.` For more details, refer to the Integrated File System directory structure in the iSeries Information Center at:

http://publib.boulder.ibm/infocenter/iseries/v5r4/index.jsp

Owner of SAP objects

All SAP traditional objects (besides the kernel objects) belong to **<*SID*>OWNER**. The kernel objects which are also traditional objects too belong to **R3OWNER**. **<*SID*>OWNER** owns the database objects. SAP Integrated File System objects have the primary group owner of **<*SID*>GROUP**.

The following programs are exceptions. These objects are owned by the **QSECOFR** user profile:

- APYR3KRN
- CHGR3SYSLE
- CRTR3INST
- CRTR3SYS
- CRTSAPUSR
- DBINCONDB4
- LODR3KRN
- LODR3SDK
- RMVR3KRN
- STRRMTDB

The following programs are owned by the **QSECOFR** user profile in SAP Release 46D:

- CHKOBJAUT
- CHKR3PTF
- SAVR3OBJS
- STRSAP

The following programs are owned by the **QSECOFR** user profile in SAP Release 45B:

- CHK04PTF
- CHKR3PTF
- SAPOSCOL

> **Note:** Be aware that the objects owned by QSECOFR can differ between the SAP releases.
>
> You always get the correct settings when preceding the FIXR3OWNS command, see “Programs that correct authorizations:” on page 118.

The characteristics of SAP user profiles are as follows:

- **<*SID*>OFR** SAP user profile

  <SID>OFR is member of the group defined by the profiles <***SID*>OWNER** and **R3OWNER.** Consequently, **<*SID*>OFR** is authorized for all SAP libraries. The **<*SID*>OFR** user is not authorized for data libraries of other SAP system IDs, since the object owner of a database library corresponds to **<*SID*>OWNER**.

  <*SID*>OFR is also a member of the groups defined by the profiles **<*SID*>GROUP** and **R3GROUP**. It is therefore authorized for all Integrated File System objects of this SAP system ID.

  <*SID*>OFR is created by the SAP installation procedure with ***JOBCTL** and ***SAVSYS** as special authorizations.

- <***SID***>**OPR** SAP user profile

  This user is authorized to start and stop the SAP system and to initiate a backup of the SAP system. But he cannot look at the data in the database or Integrated File System.

  The ***SID***>**OPR** user profile has the group profiles named **<*SID*>GROUP** and **R3GROUP** and consequently is authorized for all Integrated File System objects. **<*SID*>OPR** has the special authorizations of ***JOBCTL** and ***SAVSYS**. **<*SID*>OPR** does not have **<*SID*>OWNER** as a group profile and is therefore only authorized for the following objects. These objects are protected by an authorization list named **R3ADMAUTL**:

  - R3TASKS
  - R3INLPGM
  - R3MNUOPT
  - STARTSAP
  - STRSAP
  - STRRMTDB
  - STOPSAP
  - AS4STOPSAP
  - SAPOSCOL
  - SAPROUTER
  - SAPROUTAS4

- ***<SID><instance>*** SAP user profile

  The ***<SID><instance>*** user profile has <***SID***>**GROUP**, <***SID***>**OWNER**, **R3OWNER** and **R3GROUP** as group profiles. It is explicitly authorized for the **R3<*sid*>DATA** data library. ***<SID><instance>*** has the special authorizations of ***JOBCTL** and ***SAVSYS**.

10.1.2 Solutions for some authority situations

This section describes some procedures and i5/OS commands used to implement and correct the authority attributes, as well as the primary group or the passwords of some SAP objects.

- Correct the primary group:

  All directories, all subdirectories and all stream files must have a primary group of <***SID***>**GROUP**. Only global SAP directories, subdirectories and stream files should have **R3GROUP** as primary group. Use the following command to make this change:

  CHGPGP OBJ('<*obj*>') NEWPGP(<SID>GROUP) DTAAUT(*RWX) OBJAUT(*ALL)

- Synchronization of user passwords:

  If authorization problems occur with accesses via '/QFileSvr.400', check whether the accessing user profile has the same password on the remote system. Use the following command to make this change:

  CHGUSRPRF USRPRF(<*user*>) PASSWORD() STATUS(*ENABLED)

- Programs that correct authorizations:

  SAP command **FIXR3OWNS**

  The SAP **FIXR3OWNS** command is used to change the owners of the executables (programs of the kernel library) in an SAP system. You must have QSECOFR authorization to run it.

  Figure 10-1 on page 119 shows you an i5/OS screen of the **FIXR3OWNS** command.

```
                         Fix SAP owners (FIXR3OWNS)

Type choices, press Enter.

Kernel library . . . . . . . . . LIB
Object . . . . . . . . . . . . . OBJ            *ALL

























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

Figure 10-1 The SAP command FIXR3OWNS

The syntax is as follows:

```
FIXR3OWNS LIB(<library>) OBJ(<generic name>)
```

Or

```
CALL PGM(FIXR3OWNS) PARM(<library> <generic name>)
```

- Kernel library (LIB) *or <library>*

  The library of the kernel to be fixed, for example **R3640AOPT**

- Object (OBJ) or <*generic name*>

  You can specify single objects to be fixed. Usually, this is ***ALL**.

SAP command **AS4FIXFILE**

The **AS4FIXFILE** command corrects object ownership in the data library, if necessary. The runtime can take up to four hours. Use the following parameters, as shown in Figure 10-2.

```
                   Fix file attributes in library (AS4FIXFILE)

 Type choices, press Enter.

 Name of the database library . . DBLIB
 Change object owner  . . . . . . CHGOWN         *NO
 End journaling . . . . . . . . . ENDJRN         *NO
 Start journaling . . . . . . . . STRJRN         *YES
 Change file attributes . . . . . CHGFILE        *NO
 New owner  . . . . . . . . . . . NEWOWN         *SIDOWNER
 Selection of Files . . . . . . . FILES          *ALL
                           + for more values









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

Figure 10-2 The SAP command AS4FIXFILE

The syntax is as follows:

```
AS4FIXFILE DBLIB(<library>) CHGOWN(*NO | *YES) ENDJRN(*NO | *YES) STRJRN(*YES | *NO)
   CHGFILE(*NO | *YES) NEWOWN(*SIDOWNER | <owner>) FILES(*ALL | <generic>)
```

Or

```
CALL PGM(AS4FIXFILE) PARM( ... )
```

- **DBLIB**: This parameter describes the data library which should be processed.
- **CHGOWN**: This parameter allows changes of owners of all tables, indexes, and views. The default value is ***SIDOWNER**, which affects a change on the user profile **<*SID*>OWNER**. Other entries lead to changes to the specified user profile.
- **ENDJRN**: This parameter ends the journalizing of all tables of the database. The SAP system cannot be started when it is in this status.
- **STRJRN**: This parameter starts the journalizing of all tables of the database. Use this option if you are not sure if all the tables are being journaled.
- **CHGFILE**: This parameter changes all tables so that deleted records can be used again (**REUSEDLT(*YES))**. The wait time for all tables, indexes and views is set to 120 seconds.
- **NEWOWN:** This parameter allows you to specify a new owner for all tables, indexes and views for the selected files in the selected database library. Take with care.
- **FILES:** This parameter allows you to specify a special selection of files in the specified database library. Take with care.

10.1.3 The i5/OS concept of adopted authority

Sometimes a user needs different authorities to an object or an application, depending on the situation. A solution for this temporary gain of authority is to use adopted authority.

Identify and document all programs that are created to use adopted authority, and review the code for any activities that could be exploited by the program users. As a part of regular system monitoring, review the system for new programs that adopt authority. These new programs need to be reviewed and documented.

Object ownership

Each object is assigned an owner when it is created. The owner is either the user who creates the object or the group profile if the member user profile has specified that the group profile should be the owner of the object. When you create an object, the owner is given all the object and data authorities to the object.

The owner of an object always has all the authority for the object unless any or all authority is removed specifically. As an object owner, you can choose to remove some specific authority as a precautionary measure. For example, if a file exists that contains critical information, you can remove your object existence authority to prevent yourself from accidentally deleting the file. However, as object owner, you can grant any object authority to yourself at any time.

Ownership of an object can be transferred from one user to another. Ownership can be transferred to an individual user profile or a group profile. A group profile can own objects whether or not the group has members.

When changing an object's owner, you have the option to keep or revoke the former owner's authority. To transfer ownership, any user (including the object's present owner) must have the following types of authority:

- Object existence authority for the object (except authorization list)
- Object operational and object existence authorities if the object is a file, library, or subsystem description
- All object (*ALLOBJ) special authority or ownership if the object is an authorization list
- Add authority for the new owner's user profile
- Delete authority for the present owner's user profile
- All object (*ALLOBJ) and security administrator (*SECADM) special authorities to change the object owner of a program or an Structured Query Language (SQL) package that adopts its owner's authority
- *USE authority to the auxiliary storage pool device if one is specified

Object ownership is used as a management tool by System i models. The owner profile for an object contains a list of all users who have private authority to the object. This information is used to build displays for editing or viewing object authority.

Adopted authority

Certain programs or commands called by a user can require a higher level of authority (for the duration of the command) than is normally available to that user. Adopted authority provides a means for handling this situation. Adopted authority allows a user to temporarily gain the authority of the owner of a program (in addition to the user's own authorities) while that program is running. This provides a method to give a user additional access to objects, without requiring direct authority to objects.

To set the adopted authority to i5/OS program you can do a CHGPGM with the parameter USEADPAUT *YES as shown in Figure 10-3 on page 122:

```
                            Change Program (CHGPGM)

 Type choices, press Enter.

 Program  . . . . . . . . . . . . PGM
   Library  . . . . . . . . . . .                     *USRLIBL
 Optimize program . . . . . . . . OPTIMIZE       *SAME
 User profile . . . . . . . . . . USRPRF         *SAME
 Use adopted authority  . . . . . USEADPAUT      *YES <== set for adopted authority
 Remove observable info . . . . . RMVOBS         *SAME
                           + for more values
 Enable performance collection:   ENBPFRCOL
   Collection level . . . . . . .                *SAME
   Procedures . . . . . . . . . .
 Profiling data . . . . . . . . . PRFDTA         *SAME
 Teraspace  . . . . . . . . . . . TERASPACE      *SAME
 Force program re-creation  . . . FRCCRT         *NO
 Text 'description' . . . . . . . TEXT           *SAME


                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
 Parameter PGM required.
```

Figure 10-3 The i5/OS CHGPGM command to change the USEADPAUT parameter

To check the authority to an object, the i5/OS searches for the authority according to a schema. For detailed information about the authorization search sequence, see the *iSeries Security Reference*, SC41-5302.

In general, the authorization is searched as shown in Figure 10-4 on page 123 and as explained in the following sequence.

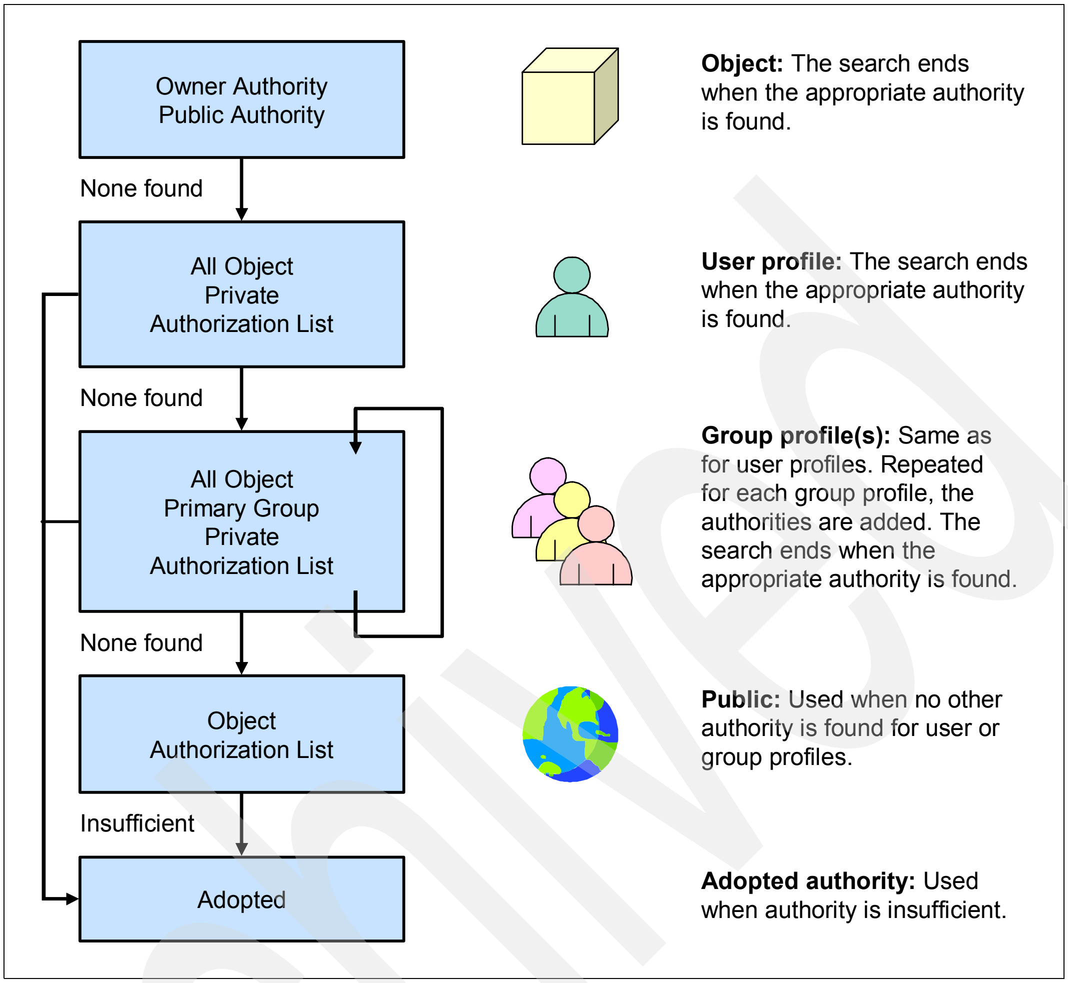

Figure 10-4 i5/OS authority search order

1. An object's authority is checked in the following order:
   a. No private authority for the object
   b. Not secured by an authorization list
   c. Sufficient owner authority
   d. No primary group for the object
   e. Sufficient public authority

   If this is true, the user is authorized. Otherwise, continue with the next step.
2. The user's authority is checked in the following order:
   a. If the user has *ALLOBJ special authority then the user is authorized.
   b. If the user's specific authority to the object is sufficient

      If the authority is not sufficient, the checking stops; continue with step 6.
   c. If the user has no specific authority to the object continue with step 3.
   d. If the user's specific authority on the authorization list securing the object is sufficient

      If the authority is not sufficient, the checking stops; continue with step 6.
   e. If the user has no authority on the authorization list securing the object continue with step 3.

If 2b or 2d is true, the user is authorized.

3. The group authority is checked in the following order, for each group profile:
   a. If the group profile has *ALLOBJ special authority.

      If this is true the user is authorized.

   b. If the group profile’s specific authority to the object is sufficient.

      If the authority is not sufficient the checking stops, continue with step 6.

   c. If the group profile has no specific authority to the object.

      Continue with step 4.

   d. If the group profile’s authority on the authorization list securing the object is sufficient.

      If the authority is not sufficient, the checking stops; continue with step 6.

   e. If the group profile has no authority on the authorization list securing the object. Continue with step 4.

   If step 3b and 3d are true, the user is authorized.

4. Public authority is checked in the following order:
   a. If the public authority specified for the object is sufficient
   b. If the public authority for the authorization list securing the object is sufficient

      If this is true the user is authorized. If authority checking is being performed for the current user in control of the job or thread, continue with the next step. Otherwise, go to step 6.

5. The program is checked for adopted authority:
   a. Is adopted authority used?
   b. Is the program owner’s authority sufficient?

      If this is true, the user is authorized. Otherwise continue with the next step.

6. The user is not authorized and one or more of the following actions happen:
   – A message is sent to the user or program.
   – The program fails.
   – An AF entry is written to the audit journal.

For further information about security, refer to *WebSphere Application Server Network Deployment V6: High Availability Solutions*, SG24-6688.

10.2 i5/OS system users with SAP applications

SAP user profiles serve different purposes and control access to different operating system objects which are used by SAP applications.

There are many types of operating system objects. In general terms, i5/OS objects can be categorized as:

- Library objects (*LIB, *FILE, *PGM...)
- Integrated File System objects (files and directories)

This section describes what purpose the SAP user profiles serve and which objects are allowed to access.

10.2.1 SAP system administrator

Figure 10-5 illustrates the relationship of the SAP i5/OS system users with their group profiles, the databases and the Integrated File System.

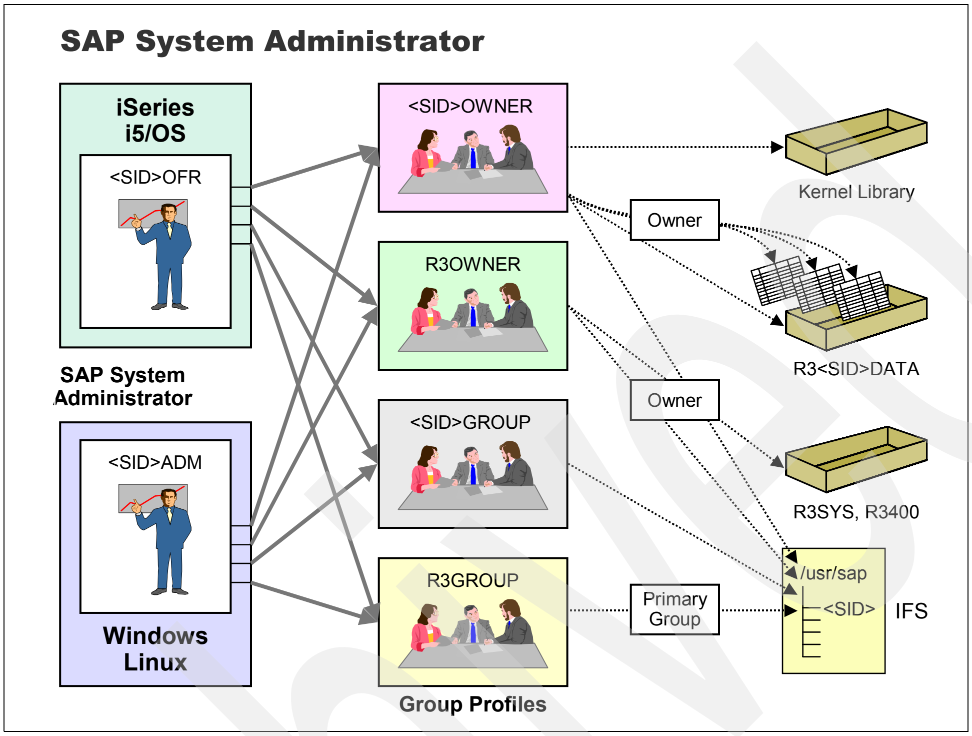

Figure 10-5 SAP System Administrator

The user ID of the SAP system administrator at operating system is <SID>OFR for System i configurations with i5/OS or <SID>ADM for Windows or Unix application servers. The <SID>ADM user profile also exists at an i5/OS level, but when working with an System i installation, <SID>OFR is usually preferred. In the future, the <SID>OFR user profile can be replaced by <SID>ADM to enable a common naming convention across platforms.

Both user profiles have four group profiles: R3OWNER, R3GROUP, <SID>OWNER, and <SID>GROUP. This allows access to all objects to which the four group profiles are authorized.

All objects that can be shared between multiple SAP systems are owned by R3OWNER, for example the libraries R3SYS and R3400 or the directory /usr/sap. Objects that are specific to one SAP system are owned by <SID>OWNER, for example the database library R3<SID>DATA or the directory /usr/sap/<SID> and all directories and files underneath. Integrated File System objects also have a primary group profile, which is comparable to group authority in Unix.

There is a primary group profile R3GROUP for Integrated File System objects that are shared between SAP systems and a primary group profile <SID>GROUP for all Integrated File System objects that belong to one SAP system.

The initial password for user profiles <SID>OFR and <SID>ADM is “SAPOFR”.

> **Note:** Change the passwords of the SAP standard user profiles immediately after installing an SAP system.

10.2.2 SAP system operator

Figure 10-6 shows the authority settings for the SAP system operator user profile.

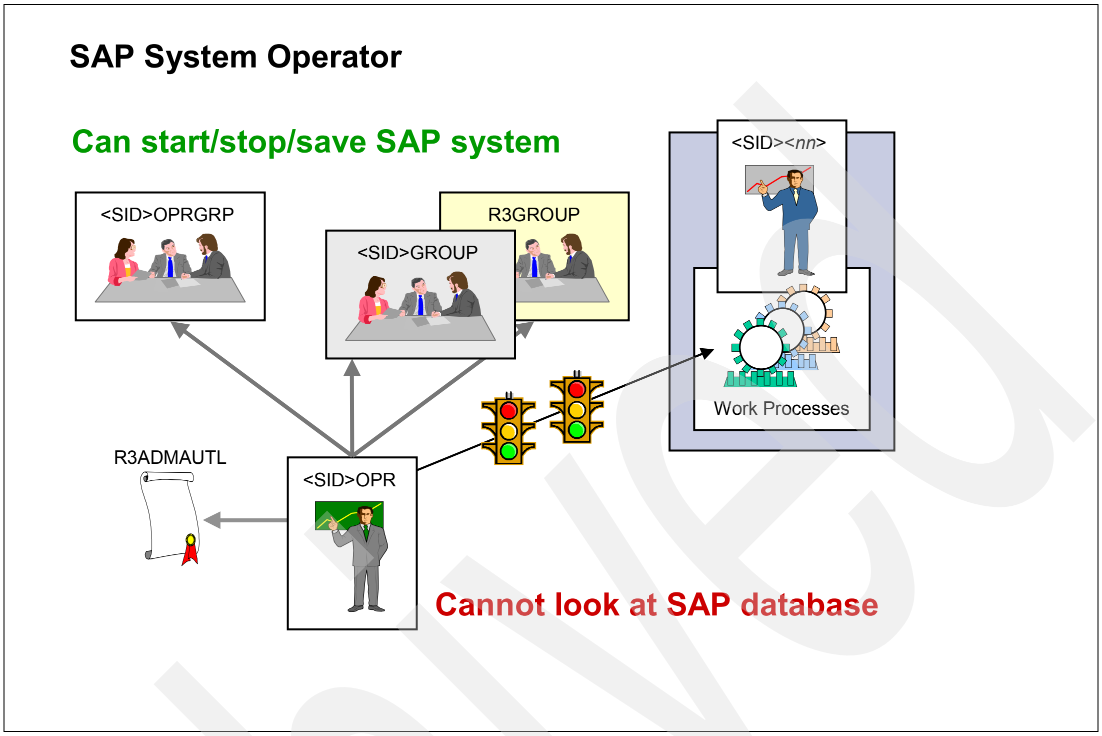

Figure 10-6 SAP system operator

The <SID>OPR user profile is designed to allow someone to start, stop, or save the SAP system, but not to look at any of the SAP system data through native methods (like STRSQL or SQLUTIL). Note that this user does not have R3OWNER or <SID>OWNER as group profiles. The necessary authority to objects in the kernel is obtained through the group profile <SID>OPRGRP and the authorization list R3ADMAUTL. User profile <SID>OPR has group profiles R3GROUP and <SID>GROUP, so it can access all Integrated File System data including log files.

In earlier SAP releases, the default password for user profile <SID>OPR is "SAPOFR". In newer releases it is *NONE, so that no one can sign on with this user until a password has been actively set by QSECOFR or a similar user.

10.2.3 ABAP and JDBC user profiles

There are additional i5/OS user profiles for Java instances on the Web Application Server. The processes of the SAP WebAS ABAP server and the SAP WebAS Java server on System i configurations run under user profiles that are named <SID><nn>. <SID> is a three digit SAP system ID and <nn> is the instance number. These user profiles have the group profiles <SID>OWNER, <SID>GROUP, R3OWNER, and R3GROUP with all their authorities. These user profiles must be enabled and have a valid password.

SAP ABAP and JDBC user profiles are illustrated in Figure 10-7.

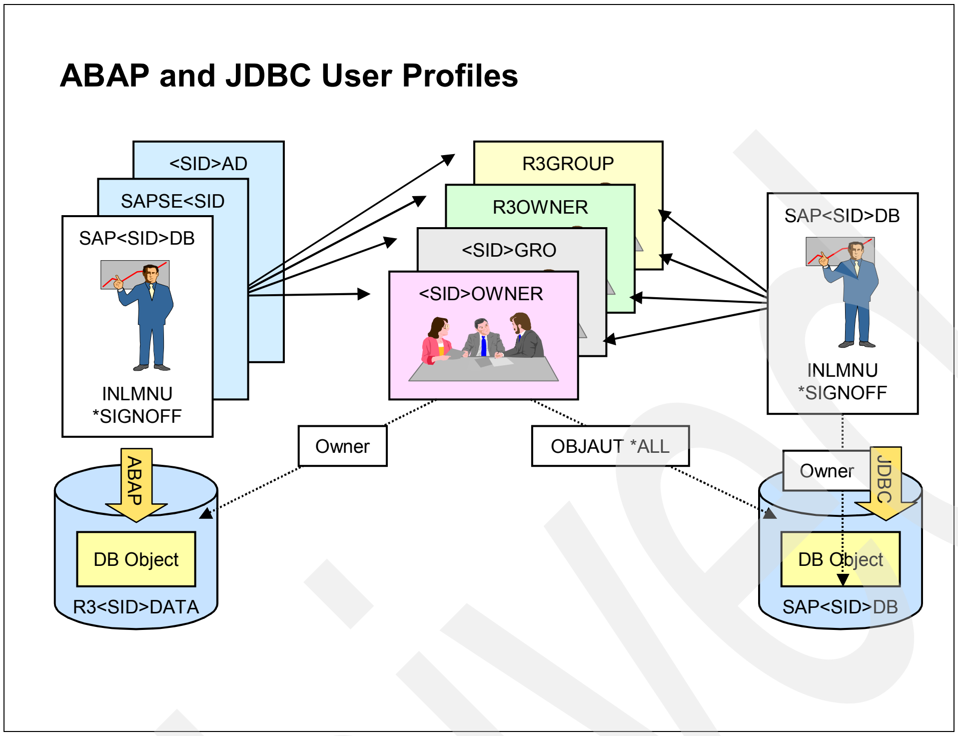

Figure 10-7 SAP ABAP and JDBC user profiles

The default password after an installation is SAP<nn>PWD. Because the initial menu is set as INLMNU(*SIGNOFF), you cannot sign on with that user profile to an interactive session. On Windows and Linux application servers, the work processes run with the users <SID>ADM or SAPse<SID>. These user profiles have the same group profiles as <SID><nn>.

On a System i SAP WebAS ABAP server, database access is done with user profile <SID><nn>. The database access of a Windows application server prior to release 6.40 is done with user profile <SID>ADM. Starting with 6.40, the database access is done by a user named SAPSE<SID>. The database access from a Linux application server is done with user <SID>ADM. Objects that are created by these users are owned by <SID>OWNER, so that the instance users <SID><nn> or SAPSE<SID> as well as the administrators <SID>OFR or <SID>ADM can access the data. On an SAP WebAS Java server, the access is done through JDBC with user profile SAP<SID>DB.

Because of the SQL conventions in JDBC, objects that are created by this user must also be owned by SAP<SID>DB. In order to allow <SID>OFR or <SID>ADM to access the data, an additional object authority *ALL is granted for user <SID>OWNER for each object that is created by SAP<SID>DB.

10.2.4 Characteristics of the SAP system user profiles on i5/OS

On any given SAP system, some user profiles do not exist or can have different settings than described in this chapter. This can be because the system evolved through upgrade procedures. Deviations introduced by earlier installation procedures can exist as well.

Note: Update the user profiles according to the information provided in this section.

Obtain the latest patch for **CRTSAPUSR** for your kernel release. Then create or update the user profiles and execute the command explained about a specific user profile described below. Check the user profile settings afterwards with respect to the attached user-settings list in Figure 10-8.

User/Param	R3GROUP	R3OWNER	<SID>OPRGRP	<SID>OPR	SHD<SID>DB	SAP<SID>DB	SAPSE<SID>	<SID>GROUP	<SID>OWNER	<SID><nn>	<SID>OFR	<SID>ADM
Password	*NONE	*NONE	*NONE	*NONE	sap	sap	sapofr	*NONE	*NONE	SAP<NN>pwd	sapofr	sapofr
Status	*ENABLED	*ENABLED	*DISABLED	*DISABLED	*ENABLED	*ENABLED	*ENABLED	*ENABLED	*ENABLED	*ENABLED	*ENABLED	*ENABLED
USRCLS	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR	*PGMR
SPCAUT	*NONE	*NONE	*NONE	*JOBCTL *SAVSYS	*NONE	*NONE	*NONE	*NONE	*NONE	*JOBCTL *SAVSYS	"*JOBCTL *SAVSYS"	"*JOBCTL *SAVSYS"
INLPGM	*NONE	*NONE	*NONE	R3INLPGM	*NONE	*NONE	*NONE	*NONE	*NONE	*NONE	R3INLPGM	R3INLPGM
INLMNU	*SIGNOFF	*SIGNOFF	*SIGNOFF	R3TASHS	*SIGNOFF	*SIGNOFF	*SIGNOFF	*SIGNOFF	*SIGNOFF	*SIGNOFF	R3MAIN	R3MAIN
LMTCPB	*YES	*YES	*YES	*NO	*YES	*YES	*YES	*YES	*YES	*YES	*NO	*NO
PWDEXPITV	*DEFAULT	*DEFAULT	*DEFAULT	*DEFAULT	*NOMAX	*NOMAX	*NOMAX	*DEFAULT	*DEFAULT	*NOMAX	*DEFAULT	*DEFAULT
GRPPRF	*NONE	*NONE	*NONE	<SID>OPRGRP	<SID>OWNER	<SID>OWNER	<SID>OWNER	*NONE	*NONE	<SID>OWNER	<SID>OWNER	<SID>OWNER
OWNER	*USRPRF	*USRPRF	*USRPRF	*USRPRF	*USRPRF	*USRPRF	*GRPPRF	*GRPPRF	*GRPPRF	*GRPPRF	*GRPPRF	*GRPPRF
SUBGRPPRF	*NONE	*NONE	*NONE	<SID>GROUP R3GROUP	<SID>GROUP R3OWNER R3GROUP	<SID>GROUP R3OWNER R3GROUP	<SID>GROUP R3OWNER R3GROUP	*NONE	*NONE	<SID>GROUP R3OWNER R3GROUP	<SID>GROUP R3OWNER R3GROUP	<SID>GROUP R3OWNER R3GROUP

Figure 10-8 Operating system user's settings

Characteristics of the user profiles are:

- **<*SID*>OFR** and **<*SID*>ADM**

  These user profiles are the SAP system administration users. The SAP system administrator should log in and perform its tasks with these user profiles. The settings of both user profiles are the same.

  <*SID*>OFR is the historical name for the i5/OS administration user while **<*SID*>ADM** is the SAP standard name. As heterogeneous environments become more and more important, the tendency is to use only **<*SID*>ADM**. You can use the **<*SID*>ADM** user profile instead of the **<*SID*>OFR** user profile.

  Users have **<*SID*>OWNER** as a group profile and are therefore allowed to access the database library. Furthermore, they have the supplemental group profiles **<*SID*>GROUP**, **R3OWNER** and **R3GROUP** and thus are allowed to access system specific Integrated File System objects, SAP specific library objects including the SAP kernel library, and SAP system specific Integrated File System objects.

  Use the following command to create the **<*SID*>ADM** user profile:

  CRTSAPUSR USER(*ADM) SID(<*SID*>)

  Use the following command to create the **<*SID*>OFR** user profile:

  CRTSAPUSR USER(*OFR) SID(<*SID*>)

- <SID><nn>

  This is the user profile under which an i5/OS instance **<nn>** runs. It has to exist and have the same password on the database host, as well as the instance host. The user profile is created with the parameter PASSWORD(*NONE) so the user <*SID*><*nn*> cannot sign on.

  The **<*SID*><*nn*>** user profile has **<*SID*>OWNER** as group profile and is therefore allowed to access the database library. Furthermore, it has the additional group profiles named **<*SID*>GROUP**, **R3OWNER** and **R3GROUP** and thus is allowed to access system specific Integrated File System objects, SAP system specific library objects, including the SAP kernel library and SAP system specific Integrated File System objects.

Use the following command to create the ***<SID><nn>*** user:

```
CRTSAPUSR USER(*SIDINST) SID(<SID>) INSTANCE(<nn>)
```

- <SID>OWNER

  The ***<SID>*****OWNER** user profile owns the SAP system specific library objects, specifically the SAP ABAP database objects. It serves as the group profile for many user profiles which work with the ABAP database. The password is set to ***NONE** to prevent users logging on with this profile.

  Use the following command to create the ***<SID>*****OWNER** user profile:

  ```
  CRTSAPUSR USER(*SIDOWNER) SID(<SID>)
  ```

- <SID>GROUP

  The ***<SID>*****GROUP** user profile is the group to which the SAP system specific Integrated File System objects are assigned. It serves as a supplemental group for many user profiles which work with SAP system specific Integrated File System objects. The password is set to ***NONE** to prevent users logging on with this profile.

  Use the following command to create the ***<SID>*****GROUP** user profile:

  ```
  CRTSAPUSR USER(*SIDGROUP) SID(<SID>)
  ```

- SAPSE<SID>

  The **SAPSE*****<SID>*** user profile was introduced starting with SAP NetWeaver ´04 SR1 as the user for SAP dialog instances running on Microsoft Windows. It is required on the SAP ABAP database host. Windows instances are no longer allowed to run as ***<sid>*****adm**, but use **SAPSe*****<SID>*** which has less authority. The passwords of **SAPSe*****<SID>*** on Windows and of **SAPSE*****<SID>*** on the i5/OS database host have to match. The user must have a password but the possibilities to log on should be restricted.

  The user has ***<SID>*****OWNER** as group profile and is therefore allowed to access the database library. Furthermore, it has the supplemental group profiles ***<SID>*****GROUP**, **R3OWNER** and **R3GROUP** and thus is allowed to access system specific Integrated File System objects, SAP specific library objects including the SAP kernel library and SAP system specific Integrated File System objects.

  Use the following command to create the **SAPSE*****<SID>*** user profile:

  ```
  CRTSAPUSR USER(*SAPSESID) SID(<SID>)
  ```

- SAP<SID>DB and SHD<SID>DB

  These users are required for the SAP J2EE database (J2EE upgrade).

  SAP<SID>DB and SHD<SID>DB user profiles own the objects in the SAP J2EE database SAP<SID>DB. Since the user is used to connect via JDBC, a password is required. Restrict the log on possibilities.

  These user profiles have ***<SID>*****OWNER** as a group profile and are therefore allowed to access the ABAP database library. Furthermore, the supplemental group profiles named **<SID>GROUP**, **R3OWNER** and **R3GROUP** are associated, and thus are allowed to access system specific Integrated File System objects, SAP specific library objects including the SAP kernel library and SAP specific Integrated File System objects.

  Use the following command to create the **SAP*****<SID>*****DB** user profile:

  ```
  CRTSAPUSR USER(*J2EEDB) SID(<SID>)
  ```

  Use the following command to create the **SHD*****<SID>*****DB** user profile:

  ```
  CRTSAPUSR USER(*J2EESHD) SID(<SID>)
  ```

- <SID>OPR

  The **<SID>OPR** profile is an operator user profile with less authority than **<SID>ADM**. It has no access to the database library and only limited access to the kernel. In newer installations, the profile is delivered in a disabled status. Disable **<SID>OPR** if it is used.

  This user profile has **<SID>OPRGRP** as a group profile and has therefore limited access to the kernel. Furthermore, it has additional group profiles of **<SID>GROUP** and **R3GROUP** and thus is allowed to access system and SAP specific Integrated File System objects.

  Use the following command to create the **<SID>OPR** user profile:

  ```
  CRTSAPUSR USER(*OPR) SID(<SID>)
  ```

- <SID>OPRGRP

  The **<SID>OPRGRP** user profile is a group profile for **<SID>OPR** which controls **<SID>OPR**'s access to some (kernel) objects. In newer installations, it is delivered in a disabled status. Disable **<SID>OPR** if it is not used.

  Use the following command to create the **<SID>OPRGRP** user profile:

  ```
  CRTSAPUSR USER(*SIDOPRGRP) SID(<SID>)
  ```

- R3OWNER

  The **R3OWNER** user profile owns the SAP specific library objects, most importantly SAP kernel objects. It serves as an additional group for many user profiles use the SAP kernel. The profile password is set to ***NONE** to restrict users from logging on with it.

  Use the following command to create the **R3OWNER** user profile:

  ```
  CRTSAPUSR USER(*OWNER)
  ```

- R3GROUP

  The **R3GROUP** user profile is a group profile to which the SAP specific Integrated File System objects are assigned. It serves as an additional group profile for user profiles which work with SAP specific Integrated File System objects. The profile password is set to ***NONE** to restrict users from logging on with it.

  Use the following command to create the **R3GROUP** user profile:

  ```
  CRTSAPUSR USER(*GROUP)
  ```

  Figure 10-8 on page 128 summarizes the settings and authorities for all operating system user profiles, such as:

  – Password
  – Status
  – User class (USRCLS)
  – Special authorities (SPCAUT)
  – Initial program (INLPGM)
  – Initial menu (INLMNU)
  – Limited capabilities (LMTCPB)
  – Password expiration interval (PWDEXPITV)
  – Group profile (GRPPRF)
  – Owner
  – Supplementary group profile (SUPGRPPRF)

> **Note:** Refer to SAP note 834218 for authoritative information about System i user concept for SAP systems based on 6.xx technology or later. For earlier releases, see SAP notes 46063 or 173579.

10.3 Environment variables

i5/OS environment variables must be set properly to effectively run SAP applications on the System i server. This section describes this process.

Environment variables can be used to control the behavior of applications or operating system functions. SAP applications use environment variables to query the name of the installation (the system ID and instance number), as well as the database type and the name of the database host. In addition, environment variables can be used by support personnel to control trace level information in the case of diagnosing errors. Support personnel (SAP or IBM) direct you to what values to set specific variables to.

As illustrated in Figure 10-9, environment variables on i5/OS can be specified with the commands Add Environment Variables (ADDENVVAR), Change Environment Variables (CHGENVVAR), or Work with Environment Variables (WRKENVVAR).

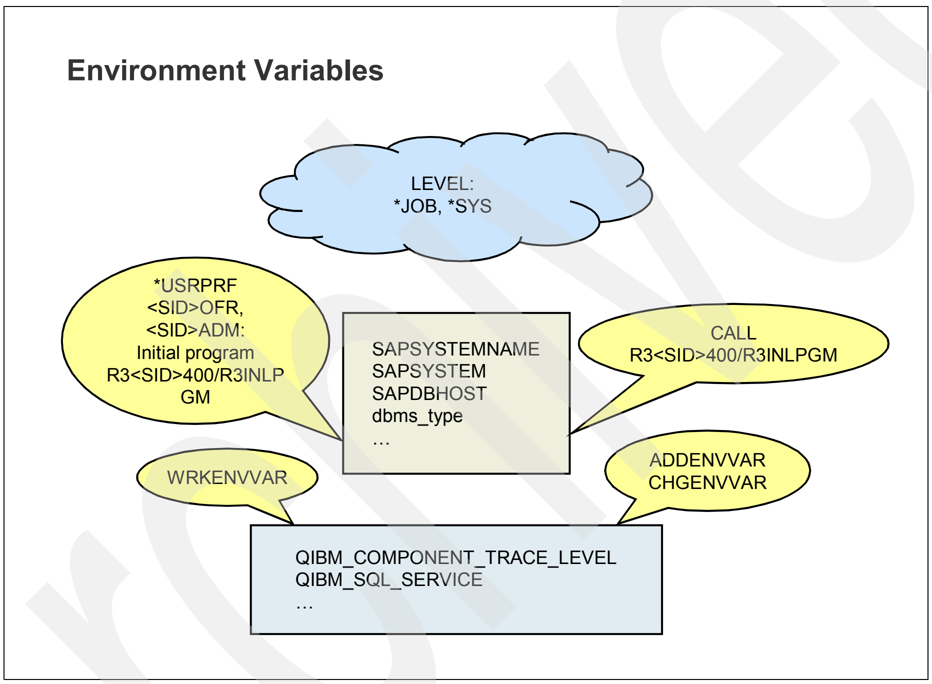

Figure 10-9 SAP system environment variables within SAP system environments and i5/OS

Environment variables can be specified at the job (*JOB) or system level (*SYS). When environment variables are specified at the system level, they are valid after a restart of the system. Environment variables at a job level are typically inherited by jobs that are spawned from the job. The name and values of environment variables are case-sensitive.

The SAP program named R3INLPGM in the library R3<SID>400 sets the necessary environment variables for an SAP system. This program is executed automatically when the <SID>OFR or <SID>ADM user signs on because it is specified as an initial program in the user profile.

You can also execute the program manually by calling it if you need the SAP system environment within a session of another user profile.

11

Installation overview

The installation topic is highly dependent on the SAP component and release level concerned. Detailed SAP documentation is available for each and every situation.

Refer to *Installing SAP Applications on IBM System i Models: An Overview*, REDP-4189 for information about planning, maintenance, and operational aspects of SAP applications that are based on SAP NetWeaver ´04 components. It offers information about system copies, a different type of installation.

11.1 SAP documentation

Always use the official Installation Guides, which you can download at:

http://www.service.sap.com/instguides

Note that you will use multiple guides. For example, if you install a complete SAP Web AS release 6.40 SR1, you need four documents: "Planning and Preparation" and "Installation and Post-Installation" for both ABAP and Java. The Installation Guides refer to individual SAP notes that have the latest information, as well as the IBM Informational APARs with current information about the PTFs necessary for your system.

Refer to the SAP notes at:

http://service.sap.com/notes

Refer to the IBM Information APARs at:

http://www-03.ibm.com/servers/eserver/iseries/service/erp/support.html

11.2 IBM Redpaper

A detailed discussion about installation of SAP NetWeaver components and solutions in System i has been created in a separate IBM Redpaper, *Installing SAP Applications on IBM System i Models: An Overview*, REDP-4189. This document gives a comprehensive overview that can be used as a checklist. It complements the SAP documentation and can be found at:

http://www.redbooks.ibm.com/abstracts/redp4189.html?Open

11.3 Set up an SAP system by system copy

There are two different classes to perform an SAP system copy:

- Homogeneous system copy
- Heterogeneous system copy

A *homogeneous* system copy refers to the copying of a specific SAP System to another SAP System, where the source and target operating system and databases are the same type. For example, a copy of an SAP system on i5/OS and DB2 for i5/OS to create a new SAP system is homogeneous if the target is also on i5/OS and DB2 for i5/OS. If you are copying between two i5 partitions, the PTF levels of partitions do not necessarily have to be the same.

If the target operating system or the target database type is different than that of the source, the copy is said to be *heterogeneous*. For example, a copy of an SAP system on Windows and Oracle to an SAP system on i5/OS and DB2 for i5/OS would be a heterogeneous system copy. Migration is another term used for a heterogeneous system copy.

The *SAPSYSTEM-ID* can remain the same, or be changed, when you do a homogeneous or a heterogeneous system copy.

The official SAP Installation Guides for both the homogeneous and the heterogeneous system copy for ABAP as well for Java are found in the SAP Marketplace at:

http://service.sap.com/instguides

Follow the navigation path of: *SAP NetWeaver → Release 04 → Installation → SAP Web AS → SAP Web AS 6.40 SR1 and Related Documentation*

SAP provides an overview and more detailed information about system copies, their functions and procedures on the following Web site:

http://service.sap.com/systemcopy

11.4 Differences between 2-tier and 3-tier landscapes

The main concepts and management involved in 2-tier versus 3-tier systems and the biggest differences between these two classes are as follows:

- 3-tier systems invariably display higher database response times than 2-tier systems. This is due to the network overhead between the application and database servers.
- Set the following parameters in a 3-tier landscape in either the default profile or in all instance profiles:

  ```
  rdisp/bufrefmode = sendon,exeauto
  rdisp/bufreftime = 120
  ```

  Refer to SAP note 14754 for more information.
- The majority of instance parameters must have the same setting in all instances. There are a few instance-specific exceptions, for example:

  ```
  INSTANCE_NAME
  SAPSYSTEM
  rdisp/wp_no_xxx
  ```

  Pay special attention to this point when additional instances are added to an existing system.
- If you intend to use logon load balancing, then implement it as described in SAP note 26317.
- Ensure that host names are correctly defined everywhere, as SAP note 60252 describes (for example, upper and lower case characters).
- Make sure that the EDRSQL server is active on the database server and on all System i application servers. If required, start the EDRSQL server with the following command:

  ```
  STRTCPSVR SERVER(*EDRSQL)
  ```
- The TCP/IP send and receive buffers on System i servers must be at least one MB (1048576 bytes) in size, as described in SAP note 92589. Restart TCP/IP afterwards in order to activate any changes.
- Make sure that the system values on all System i servers are defined in accordance with SAP note 428855.
- Because of the higher database response times, the SAP buffers have more significance in a 3-tier system than in a 2-tier system. For this reason, avoid nightly system restarts on 3-tier systems. SAP note 202593 describes how an online backup can be used to avoid stopping the system in order to take a backup.
- If TCP/IP is used on a 3-tier system, then the CHKXDA tool described in SAP note 450351 should also be installed and permanently activated on System i models. See 20.4, “The CHKXDA SAP tool” on page 349 for information about the CHKXDA tool.
- If the database server can be addressed via different host names, then configure the local instances in such a way that the host name matches that found with CFGTCP Option 12 (case-sensitive). Otherwise the work processes run via Unix-Domain-Sockets or TCP/IP, which is somewhat slower than a local connection.

  You can spot such incorrect configurations in the developer trace. If the fields named dbjobname (mostly QXDARECVR SAP database shadow job) and wpjobname contain different values on a DB-server instance, then you have an example of the situation described above.

This is the normal situation from release 6.10 onward if the following profile parameter has been set:

```
dbs/db4/allow_cancel = 1
```

- In spite of the correct implementation of the points given above, if the system performance is still unsatisfactory, then check the sizing of the whole system.

  The database server, in particular, can be too small if a central instance is installed on the same server. Contact your IBM account manager in order to arrange for a sizing check. In the meanwhile, you can reduce the load on the database server by shifting some other processes to the application server (for example batch jobs or updates).

See 3.3, “SAP system dimension” on page 20 for a discussion of 2-tier and 3-tier landscapes.

11.5 Upgrade guides

An SAP release upgrade is similar to an installation in that it requires special planning and consideration before starting. When upgrading from older SAP releases to SAP NetWeaver components, there may also be significant additions to the existing infrastructure as there are during an installation. The following is a list of the guides for upgrades:

- *Upgrade Master Guide*

  The Upgrade Master Guide is the starting point for upgrading the business scenarios of an SAP application. It provides scenario-specific descriptions of the preparation, execution, and follow-up of an upgrade. It also refers to other documents, such as the component upgrade guides and SAP notes.

  Instead of an Upgrade Master Guide, there can be several business scenario upgrade guides or a solution upgrade guide for previous shipments of the business scenarios of an SAP application. This guide is located in the SAP Service Marketplace at:

  http://service.sap.com/instguides

- *Component Upgrade Guide*

  The Component Upgrade Guide describes the technical upgrade of an SAP component, taking into account the combinations of operating systems and databases. It does not describe any business-related configuration. This guide is also located in the SAP Service Marketplace at:

  http://service.sap.com/instguides

12

The SAP kernel on the System i server

The SAP kernel is the interface between the SAP system (the database and application servers) and the platform-specific operating system and database.

SAP provides specific kernels with the following classification:

- SAP basis system (4.6D, 6.40 and so on)
- Operating system (i5/OS, OS/390®, AIX, HP-UX, Solaris, Linux and so on)
- Database (DB2 UDB for i5/OS, DB2/390, DB2 Universal Database™ (UDB), Oracle, SAP DB, MSSQLSVR and so on)
- Codepage ((Extended Binary Coded Decimal Interchange Code (EBCDIC), American Standard Code for Information Interchange (ASCII), 32-bit, 64-bit, Unicode)
- SAP Maintenance strategy (EXT = Extended Maintenance with OS/DB conflicts, COM = SAP components 3.1I and 4.0, no extension = else)

  For details about the extended maintenance of the SAP kernels and the general SAP maintenance strategy, refer to Figure 12.2.3 on page 145.

Therefore, the kernel provides the infrastructure for the complete SAP Web AS environment, including client/server technology, a platform-independent database interface, Advanced Business Application Programming (ABAP) compiler and runtime, and more components.

The SAP kernel is a part of the Basis Component (BC) of an SAP system or application. Therefore, the release cycle of the kernel is that of the basis system, rather than the release cycle of the application. For example, the kernel for SAP release 6.20 is 6.40.

Beginning with SAP NetWeaver, the kernel belongs to the Web Application Server (Web AS) which is the basis of SAP applications such as mySAP Enterprise Resource Planning (mySAP ERP) or mySAP Customer Relationship Management (mySAP CRM). For example, kernel 6.40 is used for mySAP ERP Central Component (mySAP ERP Central Component) 5.0 and other SAP applications.

In this chapter we discuss the SAP kernel for:

- SAP Web AS 6.40
- IBM i5/OS V5R4
- IBM DB2 UDB for iSeries i5/OS V5R4
- GLS and Unicode

12.1 Introduction to SAP kernels

An i5/OS library and Integrated File System objects (streamfiles and links) represent the SAP kernel of an SAP system on System i models. This section discusses these options and provides recommendations concerning the SAP kernel libraries.

> **Note:** The fundamental changes offered with kernel Release 6.40 are described in this section. Kernels prior to Release 6.40 are mentioned if it is important for the SAP releases which are still in maintenance, for example SAP R/3 4.6C.

It is useful to understand the various options when installing kernel libraries, especially if more than one SAP system is installed on an i5/OS system.

The SAP kernel is often called a *downward compatible kernel* because SAP basis release 4.6B, 4.6C, and 4.6D all use the 4.6D kernel, and SAP basis release 6.10, 6.20, and 6.40 all use the 6.40 kernel.

12.1.1 What is a kernel?

As of Release 6.40, you can see a kernel from the view of a running system as a collection of Integrated Language Environment® (ILE) objects and Integrated File System executable files in the /usr/sap/*SID*/SYS/exe/run directory. This dependency leads to modified recommendations and procedures, both of which are described in this section.

From an installation and transport perspective, a kernel on an i5/OS system can be considered as a library containing a different object types such as:

- Programs (*PGM)
- Service programs (*SRVPGM)
- Commands (*CMD)
- Physical files (*FILE PF)
- Data areas (*DTAARA)
- Message files (*MSGF)
- Menus (*MNU)
- Panel groups (*PNLGRP)
- Locales (*LOCAle)

See the *IBM Information Center* if you are interested in background information about i5/OS object types:

http://publib.boulder.ibm/infocenter/iseries/v5r4/index.jsp

The kernel library contains programs that are optimized, and therefore, offer the best performance possible. So sometimes this kernel is referred to as an *optimized kernel*. Before

4.6C, there was also a kernel available for debugging activities called the *debug kernel*. It is not available now.

In the pre-6.40 releases, the SAP system basically ran using links pointing from an Integrated File System directory to a single library object. With Release 6.40, the system now exploits objects directly located in the Integrated File System.

In the directory named /usr/sap/*SID*/SYS/exe/run which is a link to /sapmnt/*SID*/exe, there are links to the program objects in the library as well as binary executables. The content of the directory is used by SAP applications on i5/OS, for example, when starting programs.

For Release 6.40 kernel and later, the database file named O4AFFILE is available in the kernel library. Unlike previous ILE objects in the kernel library this database file contains program objects for the i5/OS Portable Application Solutions Environment (i5/OS PASE) programming model. You can use the **APYR3KRN** command to unpack the program objects.

As of kernel 6.40 there are objects in the i5/OS (kernel) library as well as Integrated File System objects (in /usr/sap/<SID>/SYS/exe/run) that belong to the kernel. These kernel objects are stored in the O4AFFILE file and are placed in the Integrated File System when the kernel is applied. The **APYR3KRN** with OPTION(*INCR) (incremental) adds or replaces newer objects in the Integrated File System.

Refer to SAP note 751132 for instructions on how to import the kernel into the /usr/sap/*SID*/SYS/exe/run directory.

12.1.2 Naming the kernel library

On a System i model, the SAP kernel is represented by the following:

- An i5/OS library where all the objects of the SAP kernel are stored
- Some files that are stored within the Integrated File System

The name of the files in the Integrated File System and the objects in the i5/OS library are fixed and cannot be changed. The name of the i5/OS kernel library is not fixed. There are several options to name the kernel library, as follows:

- Option A

  The name of the kernel library is freely definable. Names are suggested according to the template R3*REL*OPT, where *REL* stands for the SAP release. For example, 640 relates to the library name R3640OPT. Additionally, you can add an *x* between the release and the OPT, where *x* can be A (for GLS) or U (for Unicode). For example, R3640UOPT for the Release 6.40 Unicode kernel.

- Option B

  Template R3*SID*OPT is an alternative, but cannot be adhered to in an upgrade, because an SAP *SID* (system ID) temporarily has two kernels available. One kernel is from the source release, and the second kernel is from the target release. Include the release (*REL)* in the name (R3*SID*O*REL)*, for the optimized kernel, for example R3PRDO640. Or you can add *x* between *SID* and the release, where *x* can be A (for GLS) or U (for Unicode). For example, R3PRDU640 for the 6.40 Unicode kernel in an SAP system with the system ID PRD.

- Option C

  Another approach for naming the kernel is R3*RELxpatchlevel*, with statements as illustrated in Example 12-1 on page 140. For example, R3640A0098 is the kernel and the library name for a 6.40 GLS kernel with patch level 98 for disp+work (DW).

Example 12-1 R3RELxpatchlevel library names

```
REL           =   Kernel release (for example 46D or 640)
x             =   type:   E for EBCDIC
                          A for ASCII
                          U for UNICODE
patchlevel    =   4-digit actual patch level for the disp+work (DW) kernel component
```

The procedure to update the kernel with actual patches is straight forward. Each kernel (library) with a new patch-level gets a new name so you can distinguish the different kernel libraries. So a fast return to the old kernel with the former patch-level in its name is possible. Moreover, the name of the kernel library is seen as the actual patch-level of the disp+work (DW) kernel component.

On the other hand, with each DW kernel patch, a new name is used for the SAP kernel. This is not useful if you have to adapt some interfaces where you need to refer to the kernel library name. However, we recommend that you do not refer directly to the library name of an object, but work with library lists. Set the special *SID* library list with the following command:

```
CALL R3SID400/R3INLPGM
```

There are different type of kernel patches. Refer to 12.4, “Kernel 6.40 content” on page 171.

Standard names for SAP kernels

Many SAP tools for kernel processing require a standard name, for example, the parameters for **SAVLIB** options in the **LODR3KRN** command or the **APYR3FIX** command. These names follow a pattern such as:

- R3*rel*OPT: when downloading a kernel using the **LODR3KRN** command
- GEN*rel*OPT: when importing patches using the **APYR3FIX** command

More rules apply for the abbreviation *rel* depending on the SAP release. For releases 31I, 40B and 45B, the only type is EBCDIC. The release name is therefore used without being changed.

As of SAP Release 46D, as well as EBCDIC, GLS is a variant. The release name is therefore completed with the letter E or A. As of SAP Release 6.20, there is only GLS or Unicode. In this case, add the letter A or U.

Some examples of kernel names are as follows:

- R331IOPT: for **LODR3KRN** in SAP Release 31I
- GEN40BOPT: for **APYR3FIX** in SAP Release 40B
- R346DEOPT: for **LODR3KRN** in SAP Release 46D EBCDIC
- GEN46DAOPT: for **APYR3FIX** in SAP Release 46D GLS
- R3620AOPT: for **LODR3KRN** in SAP Release 6.20 GLS
- GEN640UOPT: for **APYR3FIX** in SAP Release 6.40 Unicode
- R3*rel*OPT: when downloading a kernel using the **LODR3KRN** command
- GEN*rel*OPT: when importing patches using the **APYR3FIX** command

12.1.3 Using separate kernels for each SAP system

As of Release 6.40 of the SAP kernel, SAP does not recommend using a shared kernel. By default in Release 6.40, executable files are put into /sapmnt/*sid*/exe, and are therefore, specific to a *sid*. You cannot reuse a kernel for other systems. Consequently, you should keep the ILE parts of a kernel separate in different appropriately named ILE *LIB objects. SAP does not recommend keeping all the Integrated File System objects in the same directory, independent of the *sid*.

In the case you use one kernel for several SAP systems, major changes occur when you run the **APYR3KRN** command because the kernel consists of both ILE and Integrated File System objects. **APYR3FIX** is explained in 12.3.7, “Applying the SAP kernel patches” on page 164. **APYR3KRN** is explained in 12.3.2, “Loading an SAP kernel” on page 149 and 12.3.3, “Applying an SAP kernel” on page 152.

For example, the two SAP systems, PRD and TST share a kernel library. New corrections need to be imported. In previous releases, only links to the kernel library are available in the directory named /usr/sap/*SID*/SYS/exe/run. This means that all corrections are automatically displayed and are consistent in all SAP systems. This is changed in kernel Release 6.40.

For a kernel to be consistent, run the **APYR3KRN** command after running the **APYR3FIX** command, so that the Integrated File System objects are unpacked into the /usr/sap/*SID*/SYS/exe/run directory. Only then is the kernel consistent again. As in our example, if the PRD and TST systems use the same kernel library, the kernel for the PRD system is no longer consistent after running the **APYR3KRN** command for the TST system.

Perform the following steps to use one kernel for two or more SAP systems:

1. Stop all SAP instances that use the old, shared kernel library.
2. Import all corrections into the kernel library using the **APYR3FIX** command.
3. Use the **APYR3KRN** command once per SAP system that shares the updated kernel in order to activate the kernel.
4. Start the SAP systems.

However, we prefer to keep things separate.

> **Note:** Keeping the kernels separate is enforced beginning with the 7.10 kernel when SAP plans to prescribe the name of the ILE library containing the ILE kernel parts.

For an installation as well as for upgrades, you can reuse existing kernel libraries from the same release. Consequently, several SAP systems can use the same kernel. Refer to “Option A” and “Option C” on page 139.

From Release 6.40 onward, the changes do not take effect immediately as a substitute of the kernel library, and will not automatically lead to an update of the Integrated File System-content of the exe/run directory.

The time and disk space you need for an additional kernel is compensated by the fact that you can import and test program corrections for each SAP system without interfering with another SAP system. When you upgrade the systems, the new kernel is imported in each instance and the old kernel is deleted, as there is no other SAP system that uses the old kernel.

12.1.4 Other libraries that include runtime objects (J2EE and RFC)

Table 12-1 on page 142 shows other libraries, besides the kernel library, that provide connections between the SAP system and the i5 Operating System (i5/OS) and database.

Table 12-1 Additional libraries for kernel component patches

Library name *	Library type	Definition
SAPJ2EE*nn* *	Java library	Provides runtime objects that are used by Java to access native libraries and operating system functionality. Used in a Java instance.
R3*REL*RFC R3*REL*ARFC	Library for RFC SDK the same for GLS	Software Development Kit that provides Remote Function Call (RFC) functions (service programs and modules) for use in high-level language programs (C++ and Report Program Generator (RPG))
* *nn* = SAP system number *REL* = SAP release		

Note: The library SAPJ2EE<nn> is going away in Web AS 7.00.

Installation of the SAP Remote Function Call Software Development Kit

To install the SAP Remote Function Call Software Development Kit (RFC-SDK) with i5/OS, refer to SAP note 450870.

The RFC-SDK for the different i5/OS operating system versions can be located as follows:

- SAP Basis Release 4.6D:

  On the presentation CD in file /SDK/OS400/*i5/OS-Release*/R3RFC

- SAP Web AS Release 6.20:

  On the kernel CD1 in the file /OS400/AS400/*i5/OS-Release*/R3RFC

- SAP Web AS Release 6.40:

  In the Software Download Center under SAP RFC SDK

Note: *i5/OS-Release* stands for the version of the i5/OS operating system in which the SDK was generated.

The RFC SDK is updated regularly and is available on the Software Distribution Center at:

http://service.sap.com/swdc

The RFC SDK is provided as a backup file for all SAP releases.

If the relevant CD is not available, download the R3RFC patch from the SAP Software Distribution Center. This patch contains program fixes and updates. The Software Development Kit is not delivered on CD for Release 6.40.

Note: The i5/OS generated version of the RFC-SDK in the SAP Software Distribution Center must not correspond to your current i5/OS version. You can also use an older version of the RFC-SDK in your SAP implementation.

After installation, the RFC-SDK is stored in the System i library R3*REL*RFC. The README subfile in the file R3*REL*RFC/RFC contains the necessary documentation.

Installing the RFC-SDK for SAP Basis Release 4.6D

To install the RFC-SDK in the presentation CD, follow these steps on the System i server:

1. Place the presentation CD in the CD drive.
2. Log in with the user name QSECOFR.
3. Enter the following i5/OS commands:
   - `CRTSAVF FILE(QTEMP/R3RFC)`
   - `CPYFRMSTMF FROMSTMF('/QOPT/CD-NUM/SDK/OS400/i5/OS-Release/R3RFC') TOMBR('/QSYS.LIB/QTEMP.LIB/R3RFC.FILE') MBROPT(*ADD)`
   - `CRTLIB LIB(R346DRFC)`
   - `RSTOBJ OBJ(*ALL) SAVLIB(R346DERFC) DEV(*SAVF) SAVF(QTEMP/R3RFC) RSTLIB(R346DRFC)`

   Replace `i5/OS-Release` with the version of your i5/OS operating system. Replace `CD-NUM` with the actual CD number of the presentation CD.

Installing the RFC-SDK for Web SAP AS release 6.40

To install the RFC SDK from the kernel CD1, perform the following steps:

1. Place the kernel CD1 into the free CD drive of your Windows computer.

   Alternatively, you can also place the CD into the System i server CD drive, or copy the file /OS400/AS400/*OS400_RELEASE*/R3RFC.SAR to your Integrated File System on the System i server.
2. Log into the System i server as QSECOFR.
3. Follow these steps and enter the following i5/OS commands:
   - `ADDLIBLE Kernel_library`

     Replace *Kernel_library* with the name of your kernel library, for example, R3620AOPT.

     Replace *path_to_R3RFC.SAR* with the real path name of the R3RFC.SAR file.

     > **Note:** If you have placed the kernel CD1 in the CD drive of your Windows personal computer, release the CD drive with release name `D`. To access the kernel CD1 via QNTC, use the following path name:
     >
     > `/QNTC/<(>Windows_computer/D/OS400/AS400/<(>OS400_RELEASE/`

     For more information, refer to the chapters *Sharing the Windows CD Drive*, *Connecting to iSeries Using QNTC*, and *Preparing the CDs* in the SAP document *SAP Web Application Server 6.20 Installation: IBM eServer iSeries.*
   - `SAPCAR '-xvf path_to_R3RFC.SAR/R3RFC.SAR -R /QSYS.LIB/QTEMP.LIB -flat R3RFC.FILE'`
   - `CRTLIB LIB(R3620RFC)`
   - `RSTOBJ OBJ(*ALL) SAVLIB(R3620ARFC) DEV(*SAVF) SAVF(QTEMP/R3RFC) MBROPT(*ALL) ALWOBJDIF(*ALL) RSTLIB(R3620RFC)`

To install the RFC-SDK for SAP Web AS release 6.40, perform the following steps:

1. Download the RFC.SVF file from the SAP Software Download Center.
2. Transfer the file RFC, which is a save file, to your System i server.
3. Install the SDK with following commands:

   a. `CRTLIB LIB(R3640RFC)`

b. `RSTOBJ OBJ(*ALL) SAVLIB(R3640ERFC) DEV(*SAVF) SAVF(`*`savf-lib`*`/RFC) MBROPT(*ALL) ALWOBJDIF(*ALL) RSTLIB(R3640RFC)`

12.2 SAP kernel, operating system and database release

Ensure that you are using the right SAP kernel version for your operating system and database version.

> **Note:** SAP recommends that you use the most recently released SAP kernel. This minimizes the risk of running into an SAP kernel upgrade constraint when problems arise that you are not able to fix with your current kernel.
>
> However, in production SAP systems, it is hard to follow this recommendation. A reasonable compromise is to implement new patches for your SAP Kernel about four times a year.

The SAP Service Marketplace gives you information about possible combinations.

Note that not all the combinations listed in this note have been released for all mySAP components. Some mySAP components such as BW, APO, BBP, CRM and so on, are based on SAP Basis and SAP kernel technology, but might nonetheless have their own database and operating system (OS) platform release and additional restrictions.

You can find additional information about this, as well as information regarding scheduled releases at either of the following Web sites:

- `http://service.sap.com/platforms` (Platform and Technology Info Center)
- `http://service.sap.com/pam` (Product Availability Matrix (PAM))

> **Note:** Ensure that the operating system and database versions used on your SAP servers are fully supported by the respective manufacturer. Otherwise, if problems occur during a required database or operating system upgrade, finding a solution can be unnecessarily delayed.
>
> If you receive the message "`release xxx is not supported with this kernel`" from Transaction SICK, this indicates either of the following problems:
>
> - The database or operating system is no longer supported by the manufacturer.
> - You require a current SAP kernel patch that includes a newly certified version.

Refer to the following SAP notes for details on the SAP kernels:

- SAP note 156557: *"Released operating systems SAP Kernel 4.6x DB2/400"*
- SAP note 410783: *"Released operating systems SAP Kernel 6.x DB2/400"*

> **Important:** Use the right SAP kernel for your environment.

An explanation of the functionality and handling of SAP kernels from SAP installations on Linux on POWER is beyond the scope of this redbook.

You can find more information about this by using the keywords "Linux on POWER" in the SAP Service Marketplace at:

`http://service.sap.com/notes`

12.2.1 Released operating systems for SAP kernel 4.6x and DB2 UDB for iSeries

SAP note 156557 describes the database and operating system releases which are supported with the SAP kernel 4.6x.

Because basis release 4.6D enters special maintenance at the end of 2006, new installations based on basis release 4.6D are not expected.

> **Note:** The downward-compatible 4.6D kernel is used for SAP basis releases 4.6B, 4.6C, and 4.6D.

12.2.2 Released operating systems for SAP kernel 6.x and DB2 UDB for iSeries

The SAP note 410783 provides information about database and operating system releases that are released with the SAP kernel 6.x and 7.x and for DB2 UDB for iSeries.

> **Note:** The downward-compatible 6.40 kernel is used for SAP basis releases 6.10, 6.20, and 6.40.

The tables in this note are essential references for 3-tier landscapes with an additional application server. The tables indicate the operating system releases supported with SAP kernels 6.xx and 7.xx for the following:

- i5/OS and OS/400
- Linux on Power
- Windows 32-bit
- Windows 64-bit

You can find more information about Unicode and Global Language Support (GLS) in Chapter 6, “Encoding data used in SAP systems” .

12.2.3 Maintenance for SAP kernels

We learned that an SAP system and the underlying SAP kernel have different SAP release strategies. So it can happen that the end of the maintenance for an SAP kernel comes before the end of maintenance for an SAP application that uses it.

SAP provides extended maintenance for certain products and releases after the mainstream maintenance phase as part of various maintenance models. You can find more details about the general SAP maintenance strategy by selecting **SAP’s Release and Maintenance Strategy** in the SAP Service Marketplace at:

http://service.sap.com/releasestrategy

SAP communicates the actual maintenance situation of all their applications and products in the SAP Service Marketplace in the SAP Product Availability Matrix (PAM) at:

http://service.sa.com/pam

In some cases, the end of maintenance for certain kernels communicated by SAP (for example, in the PAM) can come before the end of maintenance for the SAP products that use these kernels.

This may seem like a contradiction. As a matter of principle, you can assume that SAP provides *successor versions* of the kernel available in time, as long as at least one of the SAP applications based on them is still maintained.

The end of maintenance for SAP kernels (such as 4.6D or 4.6D_EXT) is determined by the end of maintenance for the OS or database combinations with which these kernels are produced and released. Since these do not generally correspond to the end of maintenance for the SAP applications that use them, SAP builds new versions of the kernels from time to time, which then use earlier OS or database combinations. In this sense, the 4.6D_EXT kernel is a *successor version* of the 4.6D kernel, for example. If SAP no longer maintains the 4.6D kernel, you must use the 46D_EXT kernel as an alternative. The same applies for all other SAP kernel releases.

As a result, if you are using a more recent kernel version, as part of SAP's extended maintenance, you may have to upgrade your operating system and database software.

The kernel maintenance strategy is further described in SAP note 787302.

The SAP support strategy after the end of the mainstream or extended maintenance is also described in SAP note 52505.

As a result of extended support for certain SAP components, SAP also provides extended support for the SAP kernels on which these components are based. However, in some cases, extended support of the SAP kernels is only possible in combination with a database release upgrade or operating system release upgrade (DB/OS). This situation arises, for example, because the partner no longer supports a certain DB/OS release for this period. As a result, the SAP kernel with the higher DB/OS release has to be recompiled. In order to be able to distinguish between this new SAP kernel and the old one, the new kernel is called an *EXT kernel*, for example, 4.6D_EXT. For an overview of the various DB/OS combinations, select **Strategies → Extended Maintenance, OS/DB - Kernel Matrix** in the SAP Service Marketplace at:

http://service.sap.com/platforms

> **Note:** The red fields in the matrix displayed under this path mean that you have a DB/OS combination for which at least one partner no longer provides support after 2004. If you do not upgrade to a supported DB/OS combination with the corresponding SAP kernel, your system might no longer be supported if an error occurs.
>
> Yellow and red fields mean also that there is no EXT kernel for the corresponding DB/OS combination. SAP recommends that you upgrade to a DB/OS combination with a green field.

The none-EXT SAP kernels for SAP release 3.1 and 4.0B are also known as *COM kernels* with COM derived from components.

Some of the differences between the EXT and the COM kernel are the minimum release requirement of OS/400 V5R1 for the EXT kernels as compared to the earlier OS/400 releases for the COM kernel, as well as the Teraspace-2 architecture that was implemented in the 4.6D EXT kernel as discussed in SAP note 611946.

Maintenance Extension Support expires at the end of 2004. After this, you may have to carry out an immediate DB/OS upgrade. If an error occurs in the future, you may have to carry out an SAP upgrade in order to update your system to a supported DB/OS combination.

This behavior and the difference between EXT and COM kernel are described in SAP note 663811.

12.3 Managing the SAP kernel

This section discusses managing the SAP kernel when:

- Upgrading the i5/OS operating system
- Loading an SAP kernel
- Applying an SAP kernel
- Patching an SAP kernel

The managing phases of an SAP kernel are divided into the sub-phases that load and apply the kernel, as follows:

1. An i5/OS load to put the kernel (from CD or from the Web) on the machine in the i5/OS environment
2. Make the kernel visible and usable to the SAP system by copying it to the directory `/usr/sap/`*`SID`*`/SYS/exe`

12.3.1 SAP kernel when planning an operating system upgrade

Prior to upgrading the i5/OS system, ensure that there is a compatible SAP kernel and SAP application version for the new i5/OS release available. Refer to the following SAP notes:

- SAP note 68440 for general information
- SAP note 156557 when using SAP release 4.6x
- SAP note 410783 when using SAP release 6.xx

You cannot upgrade to an i5/OS release that is not certified by SAP.

Ensure that the i5/OS release being upgraded to supports the SAP application version and SAP kernel to be installed.

As of January 2006, SAP supports the following:

- 3.1I kernel (EXT, EBCDIC)
- 4.0B kernel (EXT, EBCDIC)
- 4.5B kernel (EXT, EBCDIC)
- 4.6D EXT kernel (EXT, GLS, EBCDIC)
- 6.40 kernel (GLS, Unicode)
- 7.0 kernel (GLS, Unicode)

Download the kernels and related patches from the SAP Software Download Center in the SAP Service Marketplace at:

`https://websmp102.sap-ag.de/~SAPIDP/002006825000000234912001E`

For an Informational Authorized Program Analysis Report (Info APAR) for each i5/OS release that lists the recommended and required cumulative program temporary fix (PTF) level, individual PTF numbers, and the group PTFs for all relevant operating system licensed programs for the installation of the current kernel patch levels, refer to the Web at:

`http://www-03.ibm.com/servers/eserver/iseries/service/erp/support.html`

> **Note:** Fill in the notification form at the Web site listed above in order to be notified when a change is made to the information APAR.

After upgrading i5/OS to the target release, you may not need to import a new kernel if the kernel in use meets the required criteria. However, we recommend that you use the latest available patch level of the kernel to help ensure the SAP system operates normally.

For new i5/OS releases, there can be a small delay between when the i5/OS release becomes generally available, and when this i5/OS release receives its formal SAP certification. Once an operating system release becomes available and is certificated by SAP, *release patches* are made available for the kernels that are already delivered.

Refer to Table 12-2 to determine the required SAP kernel patches needed prior to upgrading to a new i5/OS operating system level.

Table 12-2 SAP kernel patch level for i5/OS and OS/400 releases

	V4R4	V4R5	V5R1	V5R2	V5R3	V5R4
3.x	351	524	604	684	742 EXT	768 EXT
4.0B	478	712	840	9534	1017 EXT	1056 EXT
4.6D	0	0	628	1278	1850 EXT	2200 EXT
6.40	—	—	—	0	19	117
7.00	—	—	—	—	0	47

Note: In January 2006, i5/OS V5R4, i5/OS V5R3, and OS/400 V5R2 are the supported operating system release levels. It has been announced that program support for OS/400 V5R2 will be withdrawn on 30 April 2007. OS/400 V5R2 has been withdrawn from IBM marketing on 01 October 2005.

Perform the following when upgrading to a new release:

- Obtain kernel patches from the download center at the SAP Service Marketplace on the Web at:

  `https://websmp102.sap-ag.de/~SAPIDP/002006825000000234912001E`

  SAP periodically releases kernel CDs that package kernel patches together. Depending on the size and number of patches, it may be more efficient if you import the kernel patches from a CD.
- Apply all the PTFs from the current information APAR after an operating system upgrade. Refer to SAP note 83292 for more information.
- Check the SAP Service Marketplace for your new operating system for known problems on the Web at:

  `http://service.sap.com/notes`

  Note: Use the terms **VxRyM0** and **OS upgrade** to search for i5/OS-specific SAP notes in the SAP Service Marketplace. Refer to SAP note 743113 for i5/OS V5R3M0.
- Complete the i5/OS release upgrade and import the compatible kernel before upgrading to the next SAP level. In some cases, the SAP upgrade is automatically performed with an additional kernel exchange.

Mixed i5/OS release with SAP systems in a 3-tier environment

An SAP kernel on Windows and Linux platforms is available for application servers that are connected to an i5/OS database server.

The operating system releases can differ between the individual application servers. However, the SAP kernel must be at the same release on all application servers. This regulation allows a *rolling upgrade*, that is, you can run the upgrade between the database and the application servers, and among the application servers, separately.

You can see all operating system releases that are available for a specific SAP kernel and a database server in the SAP notes 156554 and 410783.

12.3.2 Loading an SAP kernel

This section describes the procedure to load an SAP kernel. The description assumes that you have an SAP release 6.40 kernel.

Refer to SAP note 410783 to check whether the 6.40 kernel supports your current database release. To load a 6.40 SAP kernel from the kernel CD, install **Option 33 - Portable Application Solutions Environment (i5/OS PASE)** of i5/OS on your System i server. Use the **DSPSFWRSC** command or the **GO LICPGM** command, and then select option **10** to verify its installation.

> **Note:** After installing additional licensed programs, you need to install PTFs from the information APAR again, including the cum package and group PTFs.
>
> This is especially important for i5/OS PASE.

To use a double-byte language, such as Chinese, Japanese or Korean in your SAP system, refer to SAP note 807912 for more details.

To load the new kernel library from CD, use the **LODR3KRN** command as Figure 12-1 shows.

```
                        Load SAP Kernel Library (LODR3KRN)

 Type choices, press Enter.

 Saved library  . . . . . . . . . SAVLIB
 Device . . . . . . . . . . . . . DEV
 Save file  . . . . . . . . . . . SAVF
   Library  . . . . . . . . . . .                  *LIBL
 Directory  . . . . . . . . . . . STMFDIR



 Kernel library name  . . . . . . KRNLIB         *SAVLIB
 Apply kernel library . . . . . . APYR3KRN       *NO






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

Figure 12-1 The Load SAP kernel Library (LODR3KRN) command

The parameters for the **LODR3KRN** command include:

- SAVLIB

  Specify the standard release-specific library name like the kernel is delivered from SAP. Refer to Standard names for SAP kernels.

- DEV

  This specifies the name of the device used for the restore operation. If you use standard SAP kernel CDs, specify ***STMF**. If you use special devices, the device name must already be known on the system by a device description. For example:

  `*STMF`, `*SAVF`, *`optical_device_name`*, *`tape_device_name`*

- STMFDIR

  This specifies the path name of the directory containing the stream files used to restore the saved data. Valid stream files are SAPCAR archives.

- KRNLIB

  Specify the name of the library to which you want to restore.

- APYR3KRN

  This specifies whether or not the loaded kernel library should be applied to an SAP System after loading the kernel (***YES** or ***NO**).

Standard names for SAP kernels and their delivery procedure

Many SAP tools for kernel processing require a standard name (for example the parameter SAVLIB with the **LODR3KRN** command or the **APYR3FIX** command). These names follow a simple pattern:

- R3*rel*OPT: when downloading a kernel using the **LODR3KRN** command
- GEN*rel*OPT: when importing patches using the **APYR3FIX** command

For the abbreviation *rel*, there are more rules applicable depending on the SAP release. As of SAP Release 46D, as well as EBCDIC, there is also the GLS variant. The release name is therefore completed with the letter E or A. As of SAP Release 6.20, there is only GLS or Unicode. Here you should add the letter A or U.

Examples include:

- R346DEOPT: for **LODR3KRN** in SAP Release46D EBCDIC
- GEN46DAOPT: for **APYR3FIX** in SAP Release46D GLS
- R3620AOPT: for **LODR3KRN** in SAP Release 6.20 GLS
- GEN640UOPT: for **APYR3FIX** in SAP Release 6.40 Unicode

The installation tool LODSAPKRN

SAP delivers the library LOKKRN for installing a kernel which you can restore on your System i server using the **LODRUN** command. Use the **LODSAPKRN** command to load a kernel.

If the kernel is delivered on DVD, then place the System i kernel CD in the subdirectory /K01 of the DVD and perform the following steps:

1. Log in as QSECOFR and load the latest tools from CD into the LODKRN library as follows:

   ```
   LODRUN DEV(*OPT) DIR('/OS400/AS400/TOOLS')
   ```

2. Log in as *SID*OFR.
3. Use the new tools to import the 6.40 kernel as follows:

   ```
   ADDLIBLE LODKRN
   ```

```
LODSAPKRN SAVLIB(R3640AOPT) DEV(*STMF) KRNOSREL(V5R2M0) MNTPNT('/QOPT/VOLID')
   APYR3KRN(*NO) KRNLIB(kernel library)
```

Replace *VOLID* with the name of the CD in the MNTPNT parameter.

Loading an SAP kernel from a support package stack

The installation of support package stacks on System i models is different from that on other platforms. To load an SAP kernel from a support package stack, perform the following steps:

1. Download the R3OPT.SAR file for the support package stack for SAP kernel Release 6.40. It contains all the necessary programs for running an SAP system with a 6.40 kernel.

   The file named R3OPT.SAR is updated on a regular basis. Its constituent parts are taken from the current production release, and thus have components that are optimally adjusted to one another.

   In Figure 12-2 you see the SAP entry point from the SAP Software Download Center (SWDC) on the SAP Service Marketplace for the Download of the SAP NetWeaver ´04 Support Package Stack.

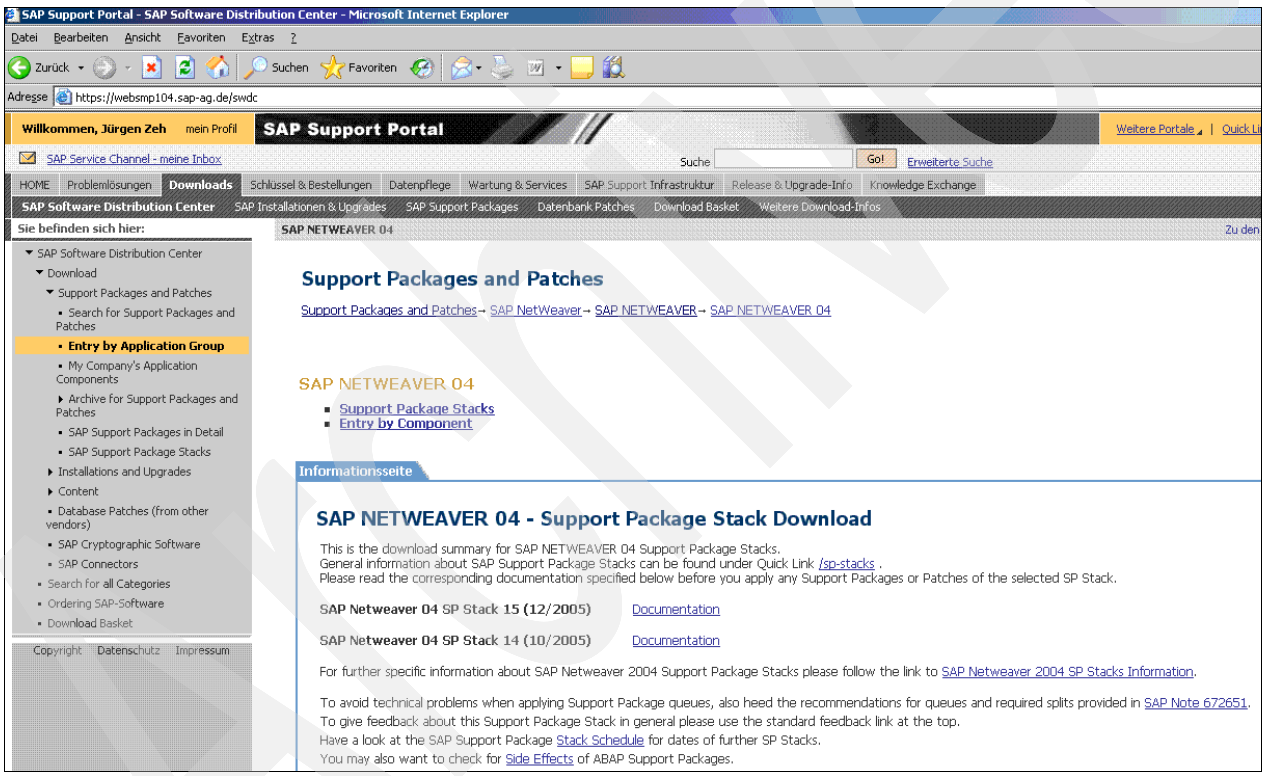

Figure 12-2 SAP Software Download Center (“swdc”) on the SAP Service Marketplace

2. Map an appropriate Integrated File System share drive if it is available via the Management Central file share option, or download to a Windows directory and then use file transfer protocol (FTP) to transfer into the /tmp directory in binary mode.
3. Download the R3OPT.SAR file to the Integrated File System of your System i server. Then use the following commands to load the library from the file:

   ```
   ADDLIBLE LIB(old_krn_name)
   LODR3KRN SAVLIB(R3640xOPT) DEV(*STMF) STMFDIR('/tmp') KRNLIB(new_krn_name)
   ```

   The file name R3OPT.SAR is assumed implicitly, which means that the system searches for the file R3OPT.SAR in the directory /tmp. Use the x with the letter **A** as a placeholder for an GLS kernel and **U** for a Unicode kernel. Give the target library a name of your choice, with a parameter name of KRNLIB.

4. Use the following commands if the R3OPT file is already in the save file format:

```
ADDLIBLE LIB(<old_krn_name>)
LODR3KRN SAVLIB(R3640<x>OPT) DEV(*SAVF) SAVF(savf_lib/R3OPT) KRNLIB(new_krn_name)
```

5. Enter the following commands as the user *SID*OFR `to` activate the kernel that has been loaded:

```
STOPSAP
APYR3KRN SID(SID) KRNLIB(new_krn_name)
STARTSAP
```

Note: The procedure to load an SAP kernel from a support package stack can change. Refer to the current SAP notes 732453 (for 4.6D) and 904977 (for 6.40).

12.3.3 Applying an SAP kernel

Perform the following steps to apply an SAP release 6.40 kernel:

1. Log in to the System i server as *SID*OFR.
2. Stop the SAP system.
3. Remove the old kernel from the SAP system using the following command:

```
RMVR3KRN SID(SID)
```

4. Remove the old kernel from library list using the following command:

```
RMVLIBLE LIB(old_kernel)
```

5. Add the new kernel library to the library list so that it is available for commands using the following command:

```
ADDLIBLE LIB(new_kernel)
```

6. Apply the new kernel using the following command:

```
APYR3KRN SID(SID) KRNLIB(new_kernel)
```

7. Log out from the System i server and log in again as `SIDOFR`.
8. Restart the SAP system.

Refer to SAP note 751132 for more information.

Perform the following steps to apply a 6.40 kernel for multiple SAP systems:

1. Log in to the System i server as *SID*OFR.
2. Stop the SAP system.
3. Remove the old kernel from the SAP system using the following command:

```
RMVR3KRN SID(SID)
```

4. Remove the old kernel from library list using the following command:

```
RMVLIBLE LIB(kernel)
```

5. Repeat step 1 to step 4 for each system that shares the 6.40 kernel in order to keep the name of the kernel library after the switch.
6. Rename the old library *kernel* to *kernel*`X`, adding a suffix of `X` to the name.
7. Rename the library *new_kernel* to *kernel*.

Repeat the following steps 8 to 11 for every SAP system that shares this kernel:

8. Add the new kernel library to the library list so that it is available using the following command:

   `ADDLIBLE LIB(`*`kernel`*`)`

9. Apply the new kernel as follows:

   `APYR3KRN SID(`*`SID`*`) KRNLIB(`*`kernel`*`)`

10. Log off from the System i server and log in again as *SID*OFR.
11. Restart the SAP system.

New features of the APYR3KRN command in SAP release 6.40

The SAP kernel from 6.40 onward consists of Integrated File System-based files, as well as i5/OS objects. These Integrated File System-based files are stored in a database file inside the kernel library. The **APYR3KRN** command unpacks the files to the Integrated File System and creates links to programs in the Integrated File System.

You can use the Apply R3 kernel (**APYR3KRN**) command to test the consistency of both the SAP system and the kernel library that you need to apply patches to. Running the **APYR3KRN** command for a standard kernel apply can take 10 to 15 minutes.

Run the **APYR3KRN** command to install patches to a kernel using the **APYR3FIX** command in order to exchange Integrated File System files that have been patched.

You can run the **APYR3KRN** command in two modes:

- In the FULL mode (the default), **APYR3KRN** deletes the executable directory (/usr/sap/*SID*/sys/exe/run by default) and rebuilds it from scratch. The **RMVR3KRN** command is performed implicitly. Do not save data in the executable directory because the content of the executable directory is discarded.
- Use the incremental mode (`MODE(*INCR)`) in case there is only a minor (non-DW) patch applied to the kernel library. This mode does not require the SAP system to shut down. Integrated File System-based files which were changed during the patch install are exchanged.

When in doubt, or if the incremental mode does not work because of object locks, proceed with the FULL mode.

> **Note:** The **APYR3KRN** command does not exist in Web AS 7.00. The function is replaced in 7.00 with the **APYSAP** command.

Configuration directory

The files in the /usr/sap/trans/config/*SID* directory are created during the installation and are used by a few tools. However, the path is fixed and you cannot adjust it by changing the location for the DIR_TRANS transport directory.

There are cases (mainly for distributed installations with several hosts) when the files cannot be located or changed consistently. **APYR3KRN** then issues a warning if the configuration files cannot be updated. This does not affect the running of your SAP system. Create a local copy of the files as a work-around.

12.3.4 Updating a Java kernel

For information about updating a kernel in a 6.40 environment that contains Java, refer to SAP note 760162: *Loading a new J2EE kernel or Java Startup Framework*.

12.3.5 SAP kernel patches

In this section we explain, as an overview, the concept of how to maintain the SAP kernel with patches for the single-kernel components. In the following sections, we show you in more detail how this concept is realized.

When you install an SAP system, the related kernel is delivered (as a whole) and installed with all its kernel components. After that subsequent kernel errors, or kernel components errors, must be corrected in individual executables (disp+work in particular). These individual executables are also called *kernel patches*. To correct a kernel error you must download one or multiple kernel patches from the SAP Service Marketplace and import them into your existing kernel (library).

Figure 12-3 illustrates the concepts of maintaining the SAP kernel components. On SAP systems these components are *binary patches*.

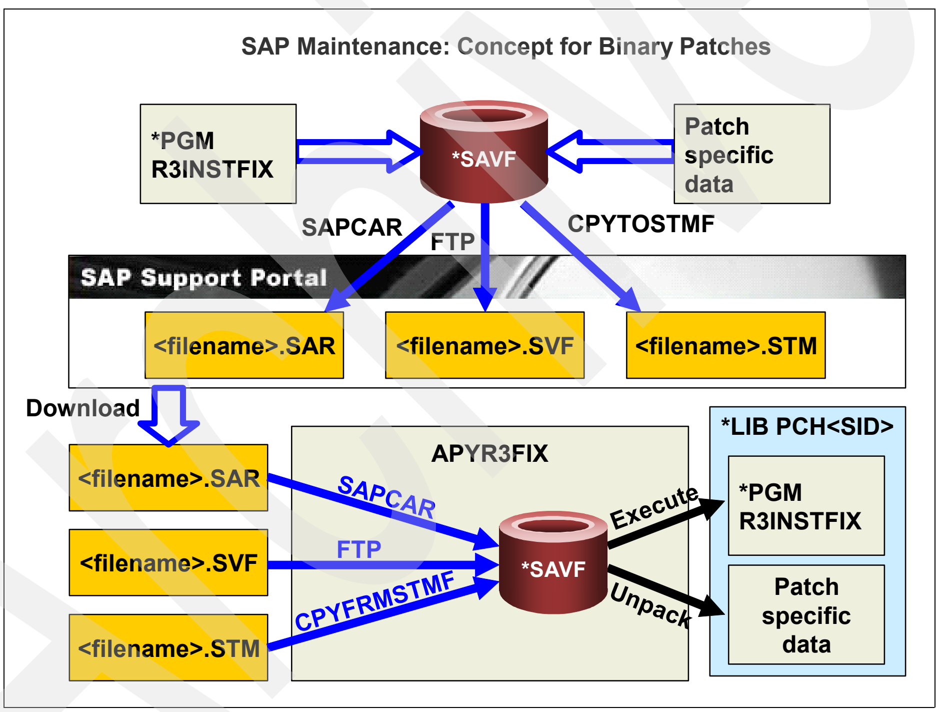

Figure 12-3 SAP kernel component concept ("Binary Patches")

Binary patches for SAP Web AS on System i servers are packaged in save files. You must convert the save file into a flat file in order to make it downloadable from the SAP Service Marketplace at:

http://service.sap.com/patches

Flat files for i5/OS are known as *stream files*.

Depending on the SAP release and kernel component, there are three different ways to convert the save file object which is identified by the file name extension:

- SAR when the SAPCAR tool is used

- SVF when FTP is used
- STM when CPYTOSTMF is used

In order to apply a patch, you must first download the flat file from the SAP Service Marketplace to a directory that can be reached from the System i server, usually a shared drive with the PC that you are downloading with. Log in as *SID*OFR and use the **APYR3FIX** command to unpack the patch. This tool performs the appropriate action to convert the stream file into a save file and restores the contents of the save file to a temporary library. You can specify your own library here, otherwise the **APYR3FIY** tool creates a library named PCH*SID* and restores the objects to it. Then execute the **R3INSTFIX** program to install patch-specific data to the kernel library.

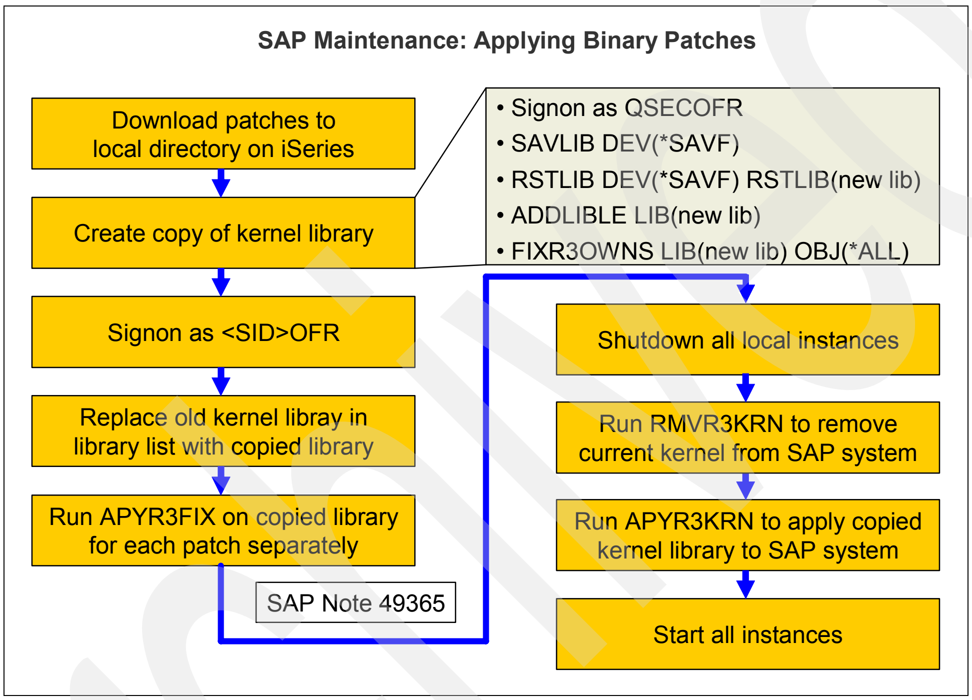

Figure 12-4 Applying SAP kernel components ("Binary Patches")

When you apply kernel patches (also known as "binary patches"), you should operate on a copy of a kernel library. For certain patches, such as DW or LIB_DBSL, this is a requirement. We recommend it for the others as well. To create a copy, log in as QSECOFR or as a user with similar authorities, because some program authorities are lost during the save and restore operations and need to be restored with the **FIXR3OWNS** command. The **FIXR3OWNS** command is available from release 4.6D onward. Prior to 6.20, you must run the program named FIXR3OWNS with two parameters, the first of which is the name of your library and the second of which is either *ALL or the name of the object to be processed.

The activities shown on the left side of Figure 12-4 can be done at any time while the SAP system is active. You need to stop the SAP system only when running the **APYR3KRN** command. The procedure for applying binary patches on System i servers is described in SAP note 49365, which covers all supported releases.

Use the **SAVLIB** command and the **RSTLIB** command to copy the kernel library to the other systems if you want to distribute the kernel library with the new patches to other systems in your landscape. Then use the **FIXR3OWNS** command to correct the authorities. The kernel must be applied before it can be used. Use the **APYR3KRN** command to apply it. This command copies all necessary objects from the kernel library into directory /usr/sap/*SID*/SYS/exe/run and creates symbolic links to the objects in the kernel library, as required.

SAP note 751131 describes how to apply a 6.40 kernel patch for SAP running with i5/OS, beyond what is documented in SAP note 49365.

Starting with kernel release 6.40, the kernel consists of ILE objects and Integrated File System objects which need to be at a consistent level. Refer to SAP note 49701 for more details.

Before a kernel is applied to an SAP system, the Integrated File System objects are stored in a database file in the kernel library. These files are extracted when you use the **APYR3KRN** command to apply a kernel to an SAP system. As a result, run the **APYR3KRN** command after you run the **APYR3FIX** command (or a series of calls to **APYR3FIX**), as described in "Making a kernel change with an updated new kernel" on page 170.

SAP support package stacks

SAP provides support packages for ABAP and Java to correct their applications. The handling of support packages is not explained in this book because it is a platform independent task and described in the SAP Service Marketplace, for example, at:

http://service.sap.com/swdc

The SAP support package strategy is extended for certain product versions to Support Package Stacks (SP Stacks). This strategy supports the actual way in which the majority of customers apply Support Packages (SPs). Quality and service can be further improved, and therefore total cost of ownership can be reduced.

The increasing diversity of components within a product version calls for greater transparency when it comes to applying support packages and patches. It also requires clear guidelines on the recommended or permitted combinations of these. For this purpose, a new SP Stack is compiled for each product version included in the new strategy (usually quarterly). This SP Stack contains the optimal combination of support package and patch statuses for the individual components (also kernel patches) at the given time.

The SP Stacks are sets of support packages and patches for the respective product version that must be used in the given combination. SAP recommends regular application of these SP Stacks. The technology for applying support packages and patches does not change. You should see the SP Stacks as an entity in themselves. Customers must heed the minimum requirements and dependencies between individual components, and apply the support packages and patches specified in the SP Stack together.

Reducing the wide range of support package combinations theoretically possible, to practical and sensible combinations has numerous advantages for customers, including:

- An enhanced quality of single support packages, since other components involved meet specific minimum requirements, and thus, corrections are less complex and of a higher quality.
- SAP's test activities can focus more on the required combinations (SP Stacks), which means that the quality and compatibility of these combinations can be improved.
- Download pages are specifically tailored to SP Stacks. This makes it simpler to download the required support packages and patches.
- Application instructions are specifically created for SP Stacks. This reduces the time and effort needed to apply corresponding support packages and patches.
- Knowledge of possible restrictions and the transparency of side effects and their solutions is generally better within the context of SP Stacks compared to support packages and patches that are combined randomly. Potential problems can be preempted or solved faster, and operating costs can ultimately be reduced.

SP Stacks support the demands by customers with productive applications for regular support package and patch recommendations at minimized cost of ownership. For customers in upgrade or implementation projects, (higher) recommendations that meet the given requirements can also apply other than those that are specified in the last SP Stack.

The introduction of SP Stacks leads to reduced complexity, increased quality, greater transparency, and easier execution of planned maintenance. In operative business, risks are reduced even more and problem resolution is accelerated. SP Stacks take the reduction of the total cost of ownership one step further.

Figure 12-5 shows you the entry point to the SAP support package stacks in the SAP Service Marketplace with the “Quick Link” /sp-stacks:

http://service.sap.com/sp-stacks

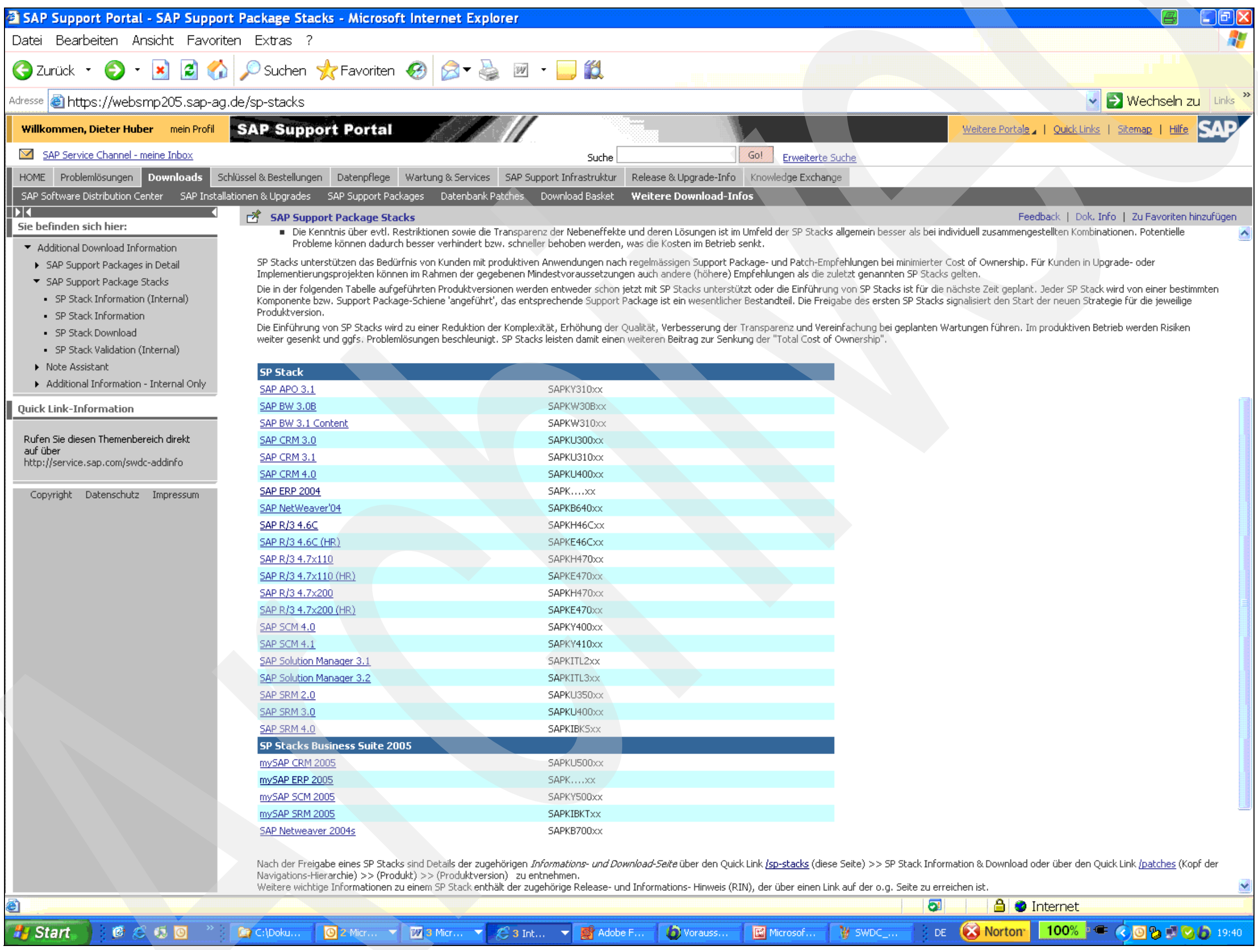

Figure 12-5 SAP support package stacks

The use of SP Stacks reduces the complexity and increases the quality and transparency of easy-of-use execution for planned maintenance. It reduces the risk of impacting operations of a business and accelerates problem resolution.

Nevertheless, at times, you need to upgrade some components (kernel and applications) prior to installing a specific SP stack, particularly the kernel or kernel components.

You can find more information about this on the Web at:

http://service.sap.com/sp-stacks

12.3.6 Downloading SAP kernel patches

You can download the SAP kernel patches and other SAP patches by selecting **Download → Entry by Application Group** on the left side of the SAP Service Marketplace Software Distribution Center on the Web at:

http://service.sap.com/swdc

You can directly download them by selecting from the site:

http://service.sap.com/patches

Figure 12-6 shows you the entry point to the SAP kernel patches where you can select the next level *Entry by Application Group.*

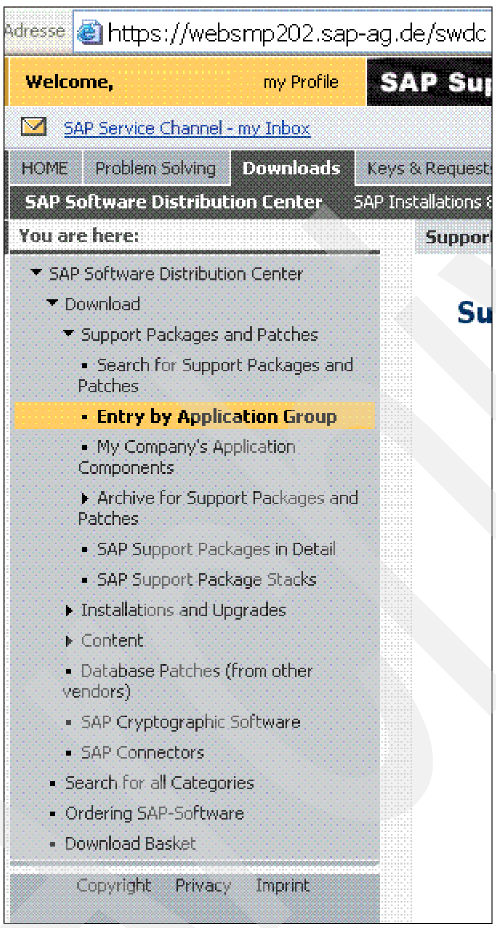

Figure 12-6 The SAP Service Marketplace Software Distribution Center (SWDC)

By selecting *Entry by Application Group* from the figure above you come to Figure 12-7 that shows you an overview of all SAP support packages and patches grouped by the application group.

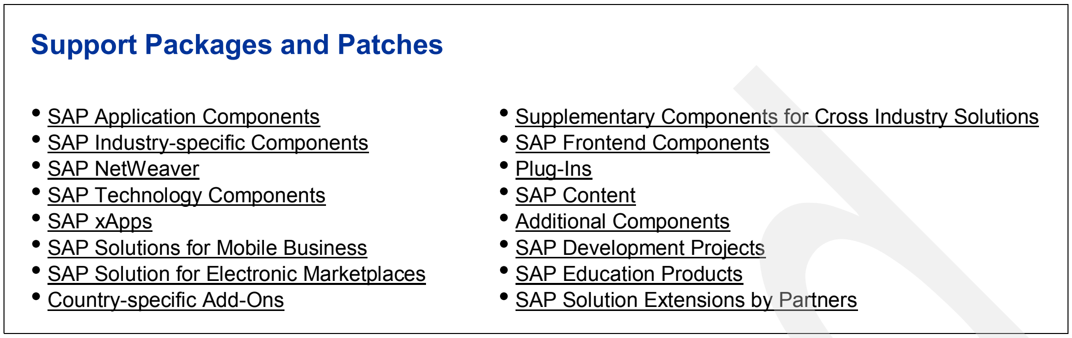

Figure 12-7 Overview SAP support packages and patches grouped by applications

> **Note:** You can find the SAP kernel patches under the application group "Additional Components".

When you open the "Additional Components" group, you will see the SAP kernel group and further options, such as:

- SAP kernel 32-bit
- SAP kernel 32-bit Unicode
- SAP kernel 64-bit
- SAP kernel 64-bit Unicode

For Unicode, only the following are available:

- SAP kernel 6.40 64-bit Unicode
- SAP kernel 7.00 64-bit Unicode

The SAP kernel patches for System i models are located under the 64-bit kernels as shown in Figure 12-8.

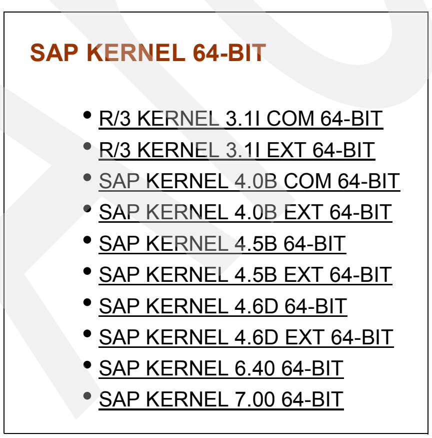

Figure 12-8 The SAP 64-bit kernels

Here, we select the SAP kernel 64-bit, not Unicode.

In the example shown in Figure 12-9, you see the complete tree for the selection of the SAP kernel 6.40, 64-bit, non-Unicode.

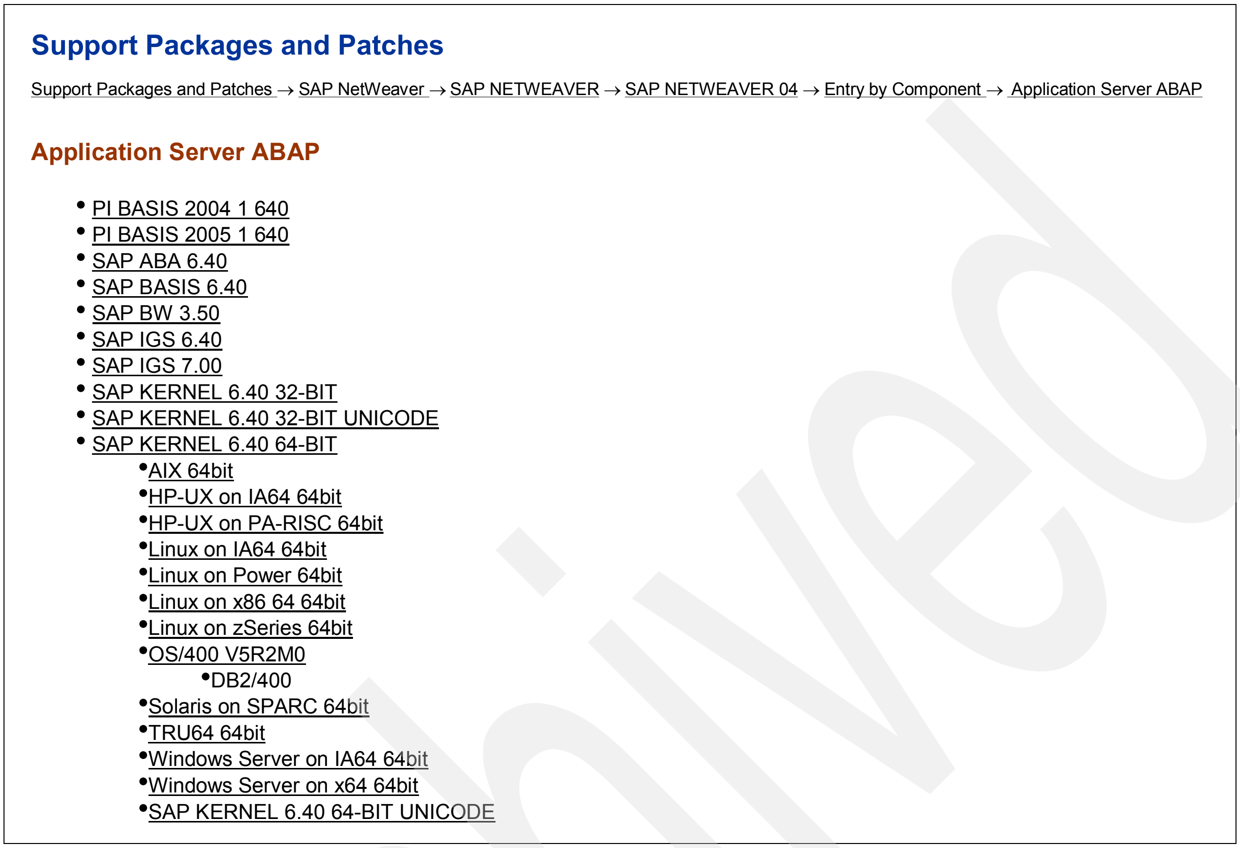

Figure 12-9 Selection of 6.40 GLS kernel patches from the SAP "SWDC"

We select the **DB2/400** line from the OS/400 V5R2 kernel and see a display as shown in Figure 12-10 on page 161.

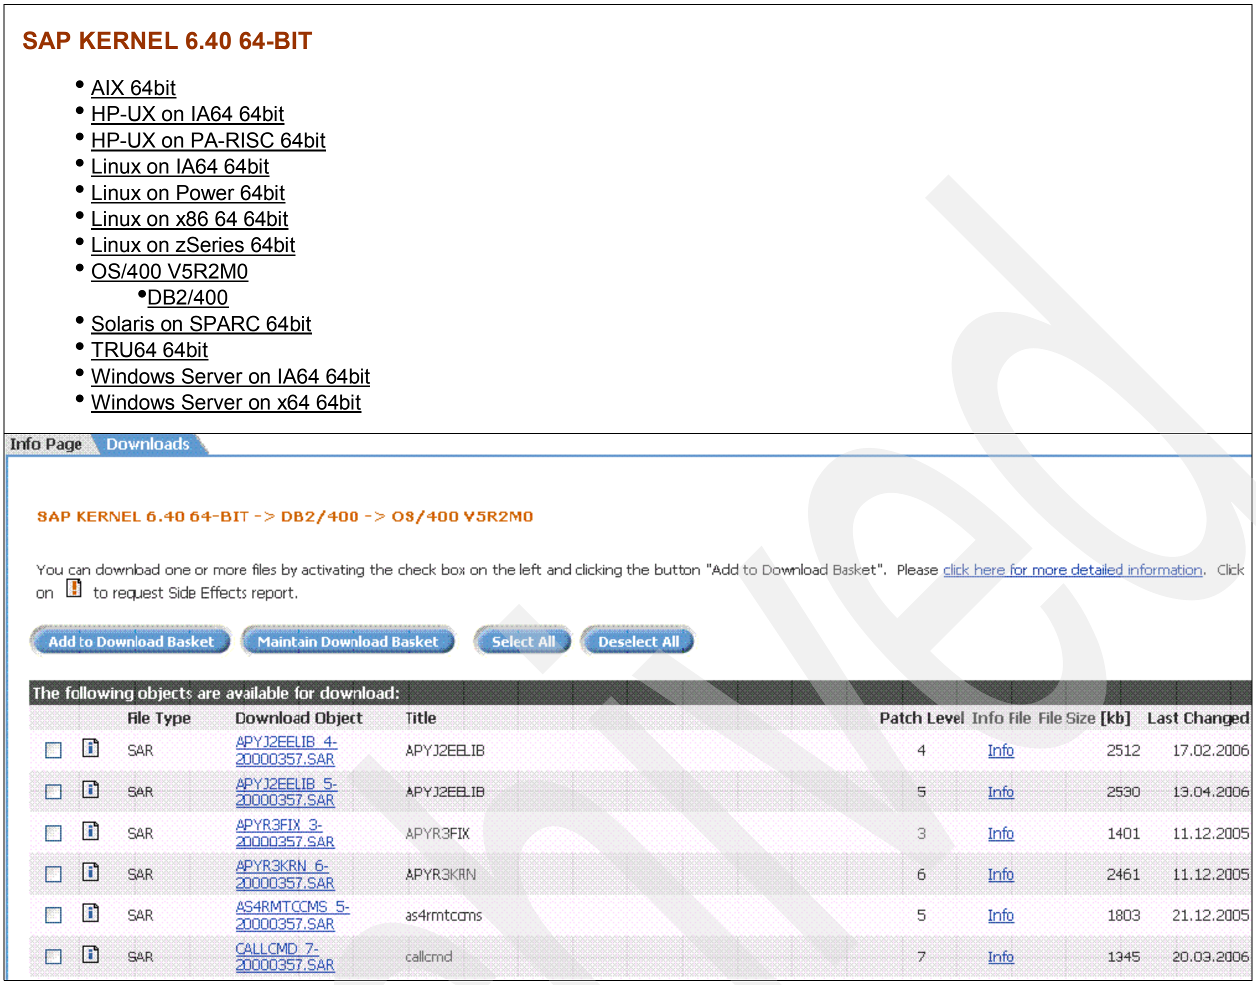

Figure 12-10 6.40 ASCII kernel patches (extract) from SAP Software Distribution Center

To ease the download process, SAP provides a Java-based download manager.

Add patches in the download basket for your *SAP Service Marketplace User* (formerly called OSS-User or SAPNET-User). Select the patches that need to be applied from the list and perform the download.

If the SAP Solution Manager is not available, refer to “SAP Download Manager” on page 164 for instructions on how to download it.

The R3CPIC and R3RFC patches are Software Development Kits. You can install them in separate libraries. Do not install them in the kernel library.

> **Tip:** Rename the files for an easier application of the patches. For example, rename the patch file for the component TP from “`TP_105-20000357.SAR`” to “`TP.SAR`” where `105` represents the patch level of the TP program. Refer to Figure 12-11 for an illustration.

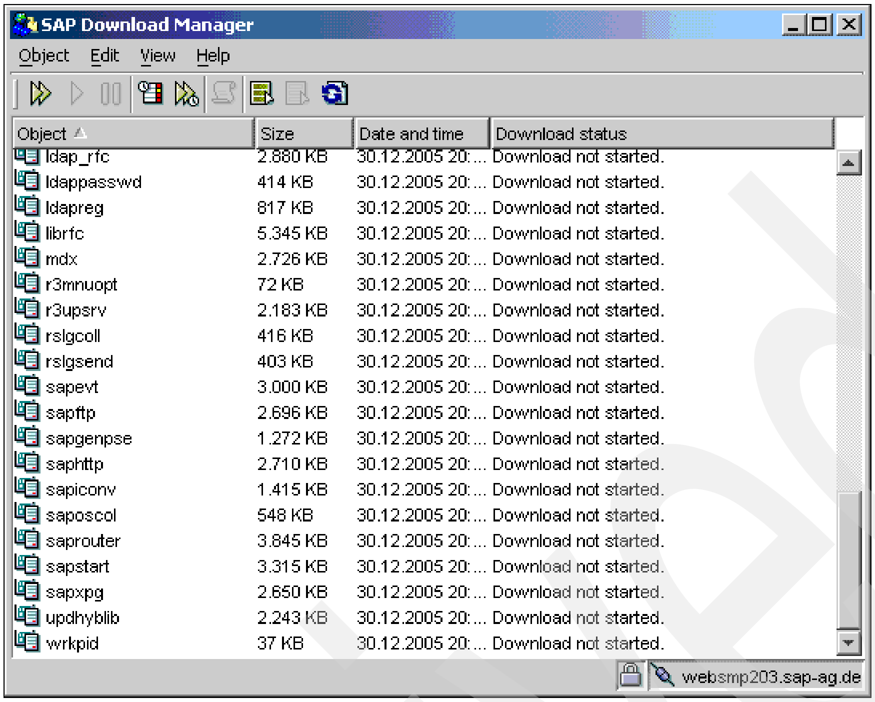

Figure 12-11 6.40 GLS kernel patches in the basket of the SAP Download Manager

Patch formats

The patches of the individual programs are “*name*.SAR” SAR or CAR archives. In the SAP Software Distribution Center (SWDC), the patch level and further information are displayed for every patch (select the **Info** button on the left to get a link to the documentation). Always use the patch with the highest version number.

From 4.6D onward, the following archives contain the complete SAP kernel:

- SAPEXE.SAR: contains all database independent executables and libraries
- SAPEXEDB.SAR: contains database specific executables

You always need both archives to update a kernel.

For regular system maintenance, use the newest SAPEXE / SAPEXEDB packages (kernel stacks).

In addition to the disp+work kernel runtime, a number of programs are contained in the dw.SAR CAR archives. Because of existing dependencies, always use it in consistent versions. The link `Info` in the Software Center has the same contents as is shown when calling disp+work -V. The link contains information about the SAP and database version and the kernel patch level.

With basis release 4.6C, an SAPCAR archive tool was imported, which describes archives in a changed format. To differentiate between the archive formats, SAP uses the .SAR file name extension for the this archives.

Use `SAPCAR -xvf` *`name`*`.SAR` to unpack the SAR archives.

> **Note:** The SAP kernel patches are normally located in .SAR archive files.

You can only unpack the SAR archive using the SAPCAR tool. You can unpack the CAR archives with the CAR and SAPCAR tools. Refer to SAP note 212876: *The new SAPCAR archiving tool*.

Use `CAR -xvf name.CAR` to unpack the CAR archives. You can find the CAR tool for UNIX at /usr/sap/*SAPSID*/sys/exe/run/CAR, for Windows at %homedrive%%homedir%\INSTALL\car.exe, or on the kernel CD in the platform-specific directory.

> **Note:** You can use .CAR archive files in relation with Windows or Linux in a 3-tier landscape.

Earlier patch formats include the following:

- Binary patches for i5/OS implementations are saved in save files. To import them, they need to exist in a save file in your System i configuration. You can support and use other file formats as well via the SAP Service Marketplace. These have the following file extensions:
  - .SVF

    This is a save file that was transferred via ftp from a System i server to the SAP Service Marketplace. Those files can again be transferred via ftp into a save file.
  - .STM

    This is a save file that was copied to a stream file into the Integrated File System using CPYTOSTMF. Those files need to be transferred into a save file using CPYFRMSTMF.

> **Note:** The FTP server named sapserv(x) is no longer used for the general SAP kernel patches. However, you can still download non-SAP patches and projects from sapserv(x).
>
> The FTP servers sapserv(x) contain files without extension and patch level addition. These are save files that were transferred via FTP from a System i server to the server. Thus, they correspond to the patch files .SVF on the SAP Service Marketplace.

Other relevant facts include:

- As of Release 6.20, a so-called Unicode kernel (*directory SAP KERNEL 6.20 32/64-bit Unicode*) is offered for download. This kernel is exclusively suited for Unicode systems. You must not use it for normal non-Unicode systems under any circumstances.
- The kernel for Release 6.40 is downward-compatible with 6.x so that you can use the kernel of the current release to eliminate kernel errors without having to upgrade the SAP system itself to the new release. Use the current 6.40 patch kernel for troubleshooting because patches for lower 6.20 kernels only remain available for a short time.

  It is important that each patch always matches a particular kernel release. For example, you cannot use individual 6.40 executables together with older kernel versions. Therefore, if you want to import a 6.40 patch and are still working with an older kernel, you have to exchange the kernel beforehand.

  Refer to SAP note 664679: *Installing a 6.40 kernel into a system with 6.10/6.20 Web AS*.

For more details about the download procedure of the kernel patches, refer to SAP note 330793 and SAP note 19466.

SAP Download Manager

The Java-based Download Manager automates the download of a larger number of files. It provides the following functions:

- A download basket for the administration of the files provided for the download
- Support of Microsoft Windows and UNIX-platforms
- Downloads can be scheduled as background jobs at night or over the weekend.
- Interrupted downloads can be restarted again at the same location.

You can download the Download Manager from:

http://service.sap.com/download-basket

Select **Get Download Manager** in the navigation area to get detailed instructions on how to download, install, or configure the SAP download manager on a Windows (or Linux) PC. Refer to Figure 12-12.

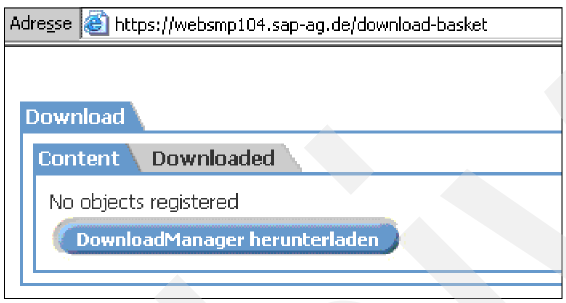

Figure 12-12 Download Manager from the SAP "SWDC"

12.3.7 Applying the SAP kernel patches

After loading the SAP kernel patches from the SAP Software Distribution Center (SWDC), you can apply them to your kernel library. You can do this with the SAP tool APYR3FIX. Before we explain the functionality and the handling of the APYR3FIX, we will show you how you can determine the patch level of the most important kernel components.

Determining latest version of the kernel, disp+work, DBSL, TP and R3trans tools

SAP recommends deploying the latest version of the kernel components, also known as kernel *tools*.

Determine the version of the kernel, DBSL, TP, and R3trans kernel components in your system as follows:

1. To determine the patch level of the kernel (this is to determine the patch level of the main component disp+work (DW)), perform the following steps:
   a. Select **system** → **status** → **other**, **kernel info** on the menu bar.
   b. Click **Release Info** in the SM51 transaction.
   c. Check the developer trace file for `patchno`

   In Figure 11-13 you see the value "105" for this `patchno`.

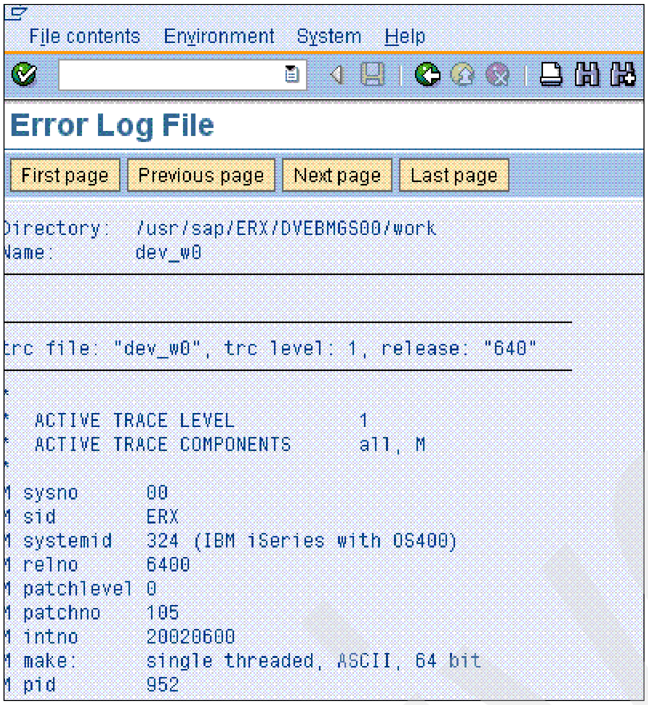

Figure 12-13 Extract from the developer trace

The DBSL level is listed in the beginning of the ST11 developer trace file. For example:

```
'Version of library '.../run/dbdb4slib' is "46D.00",
patchlevel (0.638).'
```

In this example the DBSL patch level is 638 with a 46D version.

2. Enter the following command to determine the level of the TP command:

   CALL PGM(*kernel library*/tpos4) PARM('-V')

3. Enter the following command to determine the level of the R3trans command:

   CALL PGM(*kernel library*/R3trans) PARM('-V')

Importing the kernel patches

Before you can import a patch, you should to download it via the Internet. The file can be stored directly on your System i server in the Integrated File System. For this purpose, set up a share which can be written to on your System i server. Specify this share name when saving (for example, `\\myas400\patches`).

Carry out this import using the **APYR3FIX** command. **APYR3FIX** supports the file formats described above and can generate a save file from them. It is important that the patch file you download can be found by **APYR3FIX**.

APYR3FIX restores the saved objects from a save file into library PCH*sid*. Patch files always contain the program R3INSTFIX which is called by **APYR3FIX**, and then imports the patches.

> **Note:** The import of a patch must be carried out by *sid*OFR.

APYR3FIX asks for the name of a kernel library into which you want the patch imported. The library must be the only kernel library in the library list. Otherwise **APYR3FIX** terminates.

Use the central command, **APYR3FIX** command, to apply a kernel patch. Figure 12-13 on page 165 and Figure 12-14 display the screens of the **APYR3FIX** command.

```
                             Apply R/3 Fix (APYR3FIX)

 Type choices, press Enter.

 SAP system ID  . . . . . . . . . SID
 From save file . . . . . . . . . SAVF
   Library  . . . . . . . . . . .                  *LIBL
 Saved library  . . . . . . . . . SAVLIB
 Kernel library to be processed   KRNLIB         *CURRENT
 Executable Path To Be Updated  . EXEPATH        *SYSEXERUN

 Install Program  . . . . . . . . INSPGM         R3INSTFIX
 Get save file from stream file   GETSAVF        *YES
 Stream file host . . . . . . . . FROMHOST       *LOCAL







                                                                      More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
 Parameter SID required.                                                    +
```

Figure 12-14 APYR3FIX command (1 of 2)

```
                             Apply R/3 Fix (APYR3FIX)

 Type choices, press Enter.

 Stream file directory  . . . . . FROMDIR        '/tmp'



 Stream file name . . . . . . . . FROMFILE       *SAR



 Only allow SIDOFR to patch . . . CHKSIDOFR      *YES

                           Additional Parameters

 Remote user ID . . . . . . . . . RMTUSER        *DEFAULT

 Remote password  . . . . . . . . RMTPWD

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

Figure 12-15 APYR3FIX command (2 of 2)

When using the **APYR3FIX** command, the following parameters are required:

- SAP system ID (SID):

  Enter the system ID that you want to apply the patch for. This parameter is required.

- From save file and Library (SAVF):

  This is the name of the save file that the patch is extracted to. It must be the same name as the patch, that is the stream file name without the patch number and the suffix, for example DW in the example above. If the save file does not exist, it is created. Otherwise the content is overwritten.

  .SAR files are unpacked by **APYR3FIX** in batch. Therefore you can specify any library except QTEMP.

- Saved library (SAVLIB):

  This is the name of the library where the patch was saved from. This name depends on the kernel release and code page, for example:

  – GEN46DAOPT for 46D GLS
  – GEN620AOPT for 620 GLS
  – GN620UOPT for 620 Unicode
  – GEN640AOPT for 640 GLS
  – GEN640UOPT for 640 Unicode

  For more information, refer to SAP note 49365.

- Kernel library to be processed (KRNKLIB):

  Enter the name of the library into which you want to import the patch. *CURRENT is the kernel library currently used by your SAP system.

  If you are working on a copy of the kernel library (which we recommend), specify the name of your library here.

- Executable Path To Be Updated (EXEPATH):

  Use this with the default value *SYSEXERUN unless you are told differently by SAP.

- Install Program (INSPGM):

  Use this with the default value R3INSTFIX unless you are told differently by SAP.

- Get save file from stream file (GETSAVF):

  Specify ***NO** if the save file you entered for parameter SAVF is already contained in your System i server and already contains the patch.

  Specify ***YES** if the patch file is in the Integrated File System or on another machine. In this case, **APYR3FIX** creates or overwrites the save file specified for parameter SAVF.

- Stream file host (FROMHOST):

  This parameter is only required if you specified **GETSAVF(*YES)**.

  Specify ***LOCAL** if the patch file is on your System i server in the Integrated File System.

  Specify the name of the machine where the patch file is stored.

- Stream file directory (FROMDIR):

  This is the directory that you download the patch file to.

- Stream file name (FROMFILE):

  We do not recommend using the default value (*SAR) or other special values (*SVF, *STM). The default values requires the save file specified above to have the same name as the stream file, without the extension (for example DW.SAR). Instead, specify the stream file name as it was downloaded from the SAP Service Marketplace.

- Only allow SIDOFR to patch (CHKSIDOFR):

  Specify ***YES** if you only want to allow the *SID*OFR to apply the kernel patches. This helps avoid authorization problems.
- Remote user ID / Remote password (RMTUSER/RMTPWD):

  Enter the login data to be able to fetch the patch files from another machine.

Use the default values unless otherwise instructed by SAP.

Additional considerations when applying a 6.40 kernel patch include the following items. They are further described in SAP note 751131:

- Starting with kernel release 6.40, the kernel consists of ILE objects and Integrated File System objects which needs to be consistent. (Refer to SAP note 49701.) Before a kernel is applied to an SAP system, those Integrated File System objects are stored in a database file in the kernel library. When applying a kernel to an SAP system with **APYR3KRN**, those files are extracted. As a result, you have to run the **APYR3KRN** command after running the **APYR3FIX** command (or a series of calls to **APYR3FIX**).
- The **APYR3FIX** command has three new parameters (EXEPATH, INSPGM and CHKSIDOFR). Always use them with the default values unless you are told differently by SAP.
- SAP recommends the following procedure for handling kernels starting with 6.40 and how to apply patches to these kernels:
  - Use one kernel per SAP system.
  - When you have to apply patches to the system, make a copy of the kernel library. Then apply all patches using the **APYR3FIX** command to the copy of the kernel library. Once this is done, use the **APYR3KRN** command to apply the copy of the kernel library to the SAP system. In this case, use the option **MODE(*FULL)** of **APYR3KRN**, which is the default.
  - In some cases, you can apply a patch to a running SAP system (for example, TP or R3TRANS). In these cases, you can apply the patch directly to the kernel library using the **APYR3FIX** command, and then use the **APYR3KRN** command with the option **MODE(*INCR)** to activate the patch for the SAP system. This process might fail if you use this procedure for objects which are used for a running an SAP system (for example, disp+work). In this case, you have to stop the SAP system and run **APYR3KRN** again.

    Note: Refer to SAP note 49365 for more information about the **APYR3FIX** command.

Examples using APYR3FIX

This section has five practical examples of how to use the **APYR3FIX** command.

Example 1 - Download a kernel patch

Use the **APYR3FIX** command to download a kernel patch from sapserv(*x*) and import it into your System i server. Example 12-2 uses AS001 as the system name.

Example 12-2 Using the APYR3FIX command to download a kernel patch from sapserv(x)

```
APYR3FIX SID(...)
         SAVF(QGPL/DW)
         SAVLIB(GEN....OPT)
         KRNLIB(.......)
         GETSAVF(*YES)
         FROMHOST(sapservX)
```

```
        FROMDIR('general/R3server/patches/...')
        FROMFILE(*SAVF)
```

Example 2: Import a kernel patch

Use this save file to import the patch into other application servers (AS002) as shown in Example 12-3.

Example 12-3 Importing the patch into other application servers

```
APYR3FIX SID(...)
        SAVF(QTEMP/DW)
        SAVLIB(GEN....OPT)
        KRNLIB(.......)
        GETSAVF(*YES)
        FROMHOST(AS001)
        FROMDIR(QGPL)
        FROMFILE(*SAVF)
```

Example 3: Repeat the import

Using APYR3FIX if you want to repeat an import as shown in Example 12-4.

Example 12-4 Repeating the import

```
APYR3FIX SID(...)
        SAVF(QTEMP/DW)
        SAVLIB(GEN....OPT)
        KRNLIB(.......)
        GETSAVF(*NO)
```

Example 4: Kernel patch as .SAR

A kernel patch is contained as .SAR file in the directory /patches620 in the Integrated File System on your System i server (AS001). Refer to Example 12-5.

Example 12-5 Kernel patch contained as .SAR file in the directory /patches620

```
APYR3FIX SID(...)
        SAVF(QGPL/DW)
        SAVLIB(GEN....OPT)
        KRNLIB(.......)
        GETSAVF(*YES)
        FROMHOST(*LOCAL)
        FROMDIR('/patches620')
        FROMFILE(*SAR)
```

Do not specify QTEMP as the library for the **SAVF** command since SAPCAR is run as a background job.

Example 5: Kernel patch as .STM

A kernel patch is stored as .STM file in a PATCH share on your PC PC001 and the System i file system /QNTC is used to access the PC shares as Example 12-6 shows.

Example 12-6 Kernel patch stored as .STM file in a PATCH share

```
APYR3FIX SID(...)
        SAVF(QGPL/DW)
        SAVLIB(GEN....OPT)
        KRNLIB(.......)
        GETSAVF(*YES)
```

```
FROMHOST(*LOCAL)
FROMDIR('/QNTC/PC001/PATCH')
FROMFILE(*STM)
```

Fixing the R3QTEMP not found problem

You can receive the following message:

```
Object XXX in library R3QTEMP not found.
```

Due to the switch from an old **APYR3FIX** version, the following error might occur:

APYR3FIX reimports the objects required for the patch from the save file into library PCH*sid* and wants the R3INSTFIX program to take control. However, this program has been created for the earlier version of **APYR3FIX** exclusively and expects that the reimported objects are in the R3QTEMP library.

To rectify this, perform the following steps:

1. If the R3QTEMP library does not exist, then use **CRTLIB** to create it.
2. Empty the library R3QTEMP as follows:

   ```
   CLRLIB R3QTEMP
   ```

3. Run **APYR3FIX** until it terminates.
4. Copy all objects from library PCH*sid* to the R3QTEMP library.
5. Restart **APYR3FIX**.

Making a kernel change with an updated new kernel

Perform the following steps to create a copy of the kernel library:

1. Use the Save Library (**SAVLIB**) command to save the kernel library into a save file.
2. Use the Load R3 kernel (**LODR3KRN**) command to reload the kernel library under a different name.

   Avoid using other copy methods, for example, the Restore Library (**RSTLIB)** command or the Copy Library (**CPYLIB**) command. Also, do also not use the **CPYF** command or the **CRTDUPOBJ** command, as these methods do not retain object authorizations.

3. Correct authority problems by logging in with the QSECOFR user profile and use the Fix R3 Ownership (**FIXR3OWNS**) command to fix the authority as follows:

   ```
   FIXR3OWNS LIB(new_lib) OBJ(*ALL)
   ```

   Refer to SAP note 148208 for a more detailed description.

4. Use the **APYR3FIX** command to import all corrections to the copy of the library. Enter the name of the copy in the KRNLIB parameter of the **APYR3FIX** command.

   The **APYR3FIX** command uses the **UPDPGM** command or **UPDSRVPGM** command to transfer the ILE object corrections into the kernel library. i5/OS PASE parts are exchanged in the O4AFFILE by using CHGO4AFREC.

5. Stop the SAP system and use the copy for the other SAP system.
6. Use the Remove R3 kernel (**RMVR3KRN**) command to remove the original kernel.
7. Delete or rename the original and give the copy the required name. Use the Rename Object command (**RNMOBJ**) to rename the object.
8. Use the **APYR3KRN** command to activate the copied, corrected kernel. Use the default option of **MODE(*FULL)** for **APYR3KRN**.

9. Restart the SAP system.

This total process “Making a kernel change with an updated new kernel” typically takes a couple of minutes, maybe even an hour or more. The time duration mainly depends how long it takes to import all corrections by the **APYR3FIX** command (see step 4). So it make sense to do all these **APYR3FIX** activities not to the active kernel but to a copy of it. This means you prepare a new shadow kernel while the SAP system is still running. However, the kernel switch to activate the new kernel must be done when the SAP system is down. The following sequence takes only a few minutes:

```
Stop SAP → RMVR3KRN → RNMOBJ → APYR3KRN → Start SAP
```

When doing this your SAP system must be shut down for a longer time.

Instead of applying the patches to a copy of the kernel, you can also directly patch the original kernel by performing the following steps:

1. Stop the SAP system, since the transfer can take a long time and the kernel library cannot be used by other jobs on a System i server during this time.
2. Import corrections to the original kernel by using the **APYR3FIX** command. Enter the name of the current kernel in the KRNLIB parameter of the **APYR3FIX** command.
3. Use the **APYR3KRN** command to activate the corrected kernel.

Use this procedure only if you do not have a problem with shutting down your SAP system for a longer period of time.

Other remarks and recommendations

SAP recommends the following procedure to handle kernels and how to apply patches to those kernels. These procedures apply to level 6.40 and later:

- Use one kernel per SAP system.
- In some cases, a patch can be applied to an operational SAP system (for example, TP or R3TRANS). In this case, use the **APYR3FIX** command to apply the patch directly to the kernel library. Then use the **APYR3KRN** command with the `MODE(*INCR)` option to activate the patch for the SAP system.

  This process can fail if the procedure is used for objects which are used for a production SAP system (for example, disp+work). If this occurs, stop the SAP system and run the **APYR3KRN** command again.
- The kernel patches are separated into several functional areas to ensure that the patches are not too large, and to allow an installation of only the patches required for your implementation.
- Install all available kernel patches at the same time for a more consistent and accurate installation. Do this on a regular basis, perhaps two or four times a year.
- Single kernel component patches are available for SAP. Single components can be upgraded and set to a specific component patch level, if it is not possible to apply all patches at the same time. You can upgrade some components without stopping the system.

12.4 Kernel 6.40 content

There are different types of kernel patches. We recommend taking the most recent version of each part when you are performing kernel maintenance. If you choose to update a minimal set, the most important ones are as follows:

- DW: for the disp+work kernel component
- LIB_DBSL: for the database interface kernel component
- tp and R3trans: for the transport program and the transport management system

Refer to the SAP Service Marketplace to find the kernel patch at either of the following Web sites:

- http://service.sap.com/swdc
- http://service.sap.com/patches

13

SAP instance profile parameters

Each SAP system has one *default profile* that every instance shares. Every instance has its own *start profile* and an *instance profile*. When the SAP system evaluates a profile parameter it does it in the following sequence:

- SAP kernel (default parameter setting)
- Default profile
- Start profile (only for starting and stopping the instance)
- Instance profile

Most parameter settings are performed in the instance profile. In the default profile, add only the parameters that affect all instances. The 3-tier installations have more than one SAP instance. So typically, in 3-tier installations, you can summarize a few SAP instance profile parameters in the SAP system default profile.

This chapter lists the SAP system instance profile parameters specific to an i5/OS implementation. The parameter values of the default profile are described. Also, recommended parameter settings are given for starting and customizing SAP, managing memory, and more.

13.1 SAP instance profile parameters for i5/OS installations

Table 13-1 identifies the instance profiles for an i5/OS implementation.

Table 13-1 SAP instance profile parameters for an i5/OS implementation

Parameter	Default	Parameter settings and description
as4/dbmon/enable	ON	OFF → dbmon function is not used ON → dbmon is automatically started for all work processes
as4/dbmon/memory_size	20	Maximum memory size used by DBMON to keep the data in the main memory (in MB) 0 = no maximum
as4/evistage2suport	0	Set to **1** when using enhanced vector index (EVI) stage 2 support with a BW system. See *SAP note 501572*.
dbms/type	db4	Database type (Default profile setting)
dbs/db4/allow_cancel	0 1	On i5/OS solutions, the work process or the database shadow process are masked during critical phases in the database code so that external signals cannot interrupt the operation. This, however, makes it temporarily impossible for the user to cancel the operation. Database operations can be cancelled in masked processes with dbs/db4/allow_cancel = 1. Default value for 3-tier See *SAP note 484548* and *SAP note 494086*.
dbs/db4/alternate_qaqqinilib	<not specified>	Specify the library where a file QAQQINI exists that affects those statements that are used with the optimizer hint ALTERNATE_QAQQINILIB. See also parameter: dbs/db4/qaqqinilib See *SAP note 820325*, *SAP note 485420*, *SAP note 501572*, and related SAP notes.
dbs/db4/da_cache_size	100	SAP NetWeaver Installation System i server with Linux: In case you did not use the attached control.xml files, be sure the instance profile SID_Dnn_host of your application server contains the profile parameter setting dbs/db4/da_cache_size = 0. (ABAP Only and ABAP+Java) See *SAP note 875679*.
dbs/db4/dbsl_trace	OFF	Activates a DBSL trace for SAP applications on DB2/400 to trace calls to the database See *SAP note 395365*.
dbs/db4/install_no_cmtctl	OFF	Switch off commitment control for the installation.
dbs/db4/max_hostvars	2000	Maximum number of host variables for SQL statements Default is 2000.
dbs/db4/nqe_optimize_method	O	DBSL provides a new instance profile parameter which can be used to configure the behavior of the default connection to the SAP database as well as the service connection: dbs/db4/nqe_optimize_method = [F/A/O], (default: O) Possible values are F, A, and O, where O is the default. See *SAP note 705888* and *SAP note 898203*.
dbs/db4/opticonnect	OFF	The parameter specifies whether IBM Opticonnect is used. ON means that Opticonnect is used. The default value is OFF.

Parameter	Default	Parameter settings and description
dbs/db4/qaqqinilib	QUSRSYS	Specify the library where a file QAQQINI exists that effects the whole SAP system (only to be used when directed to do so by SAP or IBM support). See also parameter: dbs/db4/alternate_qaqqinilib See *SAP note 820325* and related SAP notes.
dbs/db4/rdbname		Name of the Relational Database Directory Entry associated with the Independent Auxiliary Storage Pool (IASP). Use the WRKRDBDIRE RDB(*ALL) command to determine the name. See *SAP note 568820* and *SAP note 874028* for installation in an iASP.
dbs/db4/reopen	1	Allow reopen of access paths. See *SAP note 051651*.
dbs/db4/use_hints	-1	Allow database hints. -1 = Use database vendor defaults. See developer trace: dev_w0, search for 'use_hints' for the current active value. For >= 640: -1 means: abap hints: ON, upto limit: <*big number*> For <= 46D: -1 means: abap hints: OFF, upto limit: <*big number*> The upto limit controls when OPTIMIZE FOR n ROWS is passed to the database for ABAP statements that contains the UP TO n ROWS hint specification. Any n less than the upto limit is passed down. A <*big number*> means always pass OPTIMIZE FOR n ROWS down to the database.
em/as4/max_free_list	-1	Number of released segments administered at their default size. This parameter is only active, if extended memory on the System i configuration is implemented using single level storage and not TeraSpace (all releases prior to kernel 4.6A).
install/umask	007	Authority for IFS files created by the SAP system
rdisp/PG_MAXFS	32768	Size of the SAP paging file in 8KB blocks. The paging file is used to store extracts and for Export to Memory. The size must be sufficient. It is not a problem if it is set too large. i5/OS: Set to **131072** for release 6.xx 65536 and earlier.
rdisp/PG_SHM	32768	Size of the paging buffer in 8KB blocks. A paging buffer is not required for new memory management because internal tables and lists do not reside within SAP paging. Set **Value = 0** to save swap space. i5/OS: Set to **131072** for release 6.xx 65536 and earlier.
rdisp/ROLL_MAXFS	32768	Size of the roll file in 8KB blocks. The roll file is used to store part of the user contexts from the roll area. The size must be sufficient. It is not a problem if it is set too large. i5/OS: Set to **131072** for release 6.xx 65536 and earlier.
rdisp/ROLL_SHM	32768	Size of the roll buffer in 8KB blocks. The roll buffer is used as a cache for the roll file. To ensure fast context change, the roll buffer should provide enough space for the roll areas of all user contexts. Set for a minimum time for context change. i5/OS: Set to **131072** for release 6.xx 65536 and earlier.
ztta/as4/roll_shm_cpp_16MB	5	Size of System i C++ administration area (pieces of 16MB). This parameter is no longer used in *newer* kernel patch levels, so remove it from your instance profile. See *SAP note 376827*.

Recommendations for the parameter settings depend on the system load and also on the available physical memory, disk and processor resources. In configurations where you use dynamic logical partition (LPAR), you can change the number of processors and the main memory during production when the SAP system starts. SAP profile parameters do not typically have to be changed, only the QBASACTLVL i5/OS system value and the size of the share pools need to be adjusted.

Configuration to connect an SAP database to and from external databases

Some database specific parameters are connection based and can also be specified in the dbcon table. See *SAP note 146624*, which describes the possibility of accessing data with EXEC SQL that is not on the SAP database server but on another System i server.

See *SAP note 445872*, which describes two scenarios if you want to access data residing on a System i server from any Windows or Linux PPC application server:

- Your Windows and LinuxPPC application server is using a System i server for the primary database server.
- Your Windows and LinuxPPC application server is not using a System i server for the primary database server.

> **Note:** After the SAP kernel release 640, the System i installations use UNIX_STD memory parameters. Documentation for the parameters and their settings can be found in *SAP note 33576*.

13.2 SAP instance profile parameter values for i5/OS installations

The recommendations listed in Table 13-2 are based on *SAP note 103747*, *SAP note 121625*, *SAP note 110157*, and *SAP note 119754*, as well as the results seen in transaction ST02. In addition, earlier SAP notes such as *SAP note 44695*, *SAP note 97497*, and *SAP note 121808* are taken into consideration for SAP releases 6.xx.

The values in Table 13-2 are for SAP release 6.40. For large systems, you can significantly increase some of these values.

We make some additional remarks or give you some parameter value proposals in the "Parameter settings and description" column. In any case you find a description of the parameter in transaction **RZ11**. For special remarks to a parameter you can look for SAP notes with the parameter as a key-word.

Table 13-2 Typical SAP instance profile parameters for i5/OS installations

Parameter	6.40 values	Parameter settings and description
abap/buffersize	600,000	Maximum 2,000,000
abap/fieldexit	YES	
abap/programs	1,000	
abap/unicode_check	ON	

Parameter	6.40 values	Parameter settings and description
as4/dbmon/enable	0 or 1	0 = if LPAR or server resources are minimal 1 = otherwise
auth/rfc_authority_check	0	In test systems
bdc_iob_size	7,900	for batch input See *SAP note 011842*.
dbs/db4/opticonnect	0	
dbs/db4/use_hints	1	See *SAP note 531337*.
em/as4/max_free_list	200	
em/initial_size_MB	16,384	Maximum 32,768 See *SAP note 371029*.
em/global_area_MB	255	Maximum 255
em/stat_log_size_MB	50	
em/stat_log_timeout	600	
enque/table_size	32, 768	Maximum 65,536
gw/max_conn	2, 000	
gw/max_overflow_size	25, 000,000	Maximum 999,000,000
gw/max_shm_req	400	
gw/max_sys	2,000	
install/umask	000	Authorities for IFS files created from SAP applications
login/system_client		Set value
rdisp/appc_ca_blk_no	1,000	Maximum 2,000 See *SAP note 384971* and *SAP note 79478*.
rdisp/wp_auto_restart	86,400	For automatic restart of work processes: Time in seconds, after which a work process restarts Default: 0 (==> no restart)
rdisp/wp_abap_restart	70,000	For automatic restart of work processes: Time in seconds, after which a work process, which exceeded the abap heap limit, restarts Default: 0 (==> work process is restarted as soon as possible)
rdisp/restartable_wp	DIA, BTC, UPD, UP2, BTC, SPO	List of the work process types that should restart. This could be a list of the types DIA, BTC, UPD, UP2, SPO, ENQ, separated with ",". Default: DIA, BTC, UPD, UP2, BTC, SPO
rdisp/autothtime	102	For analysis only
rdisp/elem_per_queue	2,000	Maximum 5,000
rdisp/max_comm_entries	2,000	

Parameter	6.40 values	Parameter settings and description
rdisp/PG_MAXFS	131,072	
rdisp/PG_SHM	131,072	
rdisp/rfc_max_own_used_wp	50	Analog to transaction RZ12 *
rdisp/rfc_min_wait_dia_wp	3	Analog to transaction RZ12 *
rdisp/ROLL_MAXFS	131, 072	
rdisp/ROLL_SHM	131, 072	
rdisp/tm_max_no	2,000	
rdisp/TRACE_LOGGING	ON, 500 m	Only for analysis
rdisp/vbreorg	1	Reorganize with system start
rdisp/wp_ca_blk_no	300	
rdisp/wp_no_spo	3	
rdisp/wp_no_vb	3	about number WP (DIA + BTC) / 2
rec/client	ALL	In production system
rsdb/cua/buffersize	20,000	Maximum 40,000
rsdb/max_blocking_factor	20	See *SAP note 531337*.
rsdb/min_blocking_factor	10	See *SAP note 531337*.
rsdb/ntab/entrycount	50,000	In production systems about 72,000 See *SAP note 126210*.
rsdb/ntab/ftabsize	48,000	In production systems about 68,000 See *SAP note 126210*.
rsdb/ntab/irbdsize		
rsdb/ntab/sntabsize		Is allowed > 3,000
rsdb/obj/buffersize	20,000	
rsdb/obj/max_objects	20,000	
rspo/local_print/guitimeout	0	
rspo/wp_no_spo_Fro_max	2	Printing in parallel with access method *F*, if more than 1 spool work process
rtbb/buffer_length	30,000	
rtbb/max_tables	3,000	
sap/bufdir_entries	20,000	Maximum 30,000
zcsa/db_max_buftab	20,000	Maximum 30,000L
zcsa/presentation_buffer_area	16,773,119	Maximum 50,000,000
zcsa/table_buffer_area	129,000,000	
ztta/dynpro_area	800,000	See *SAP note 167229*.
ztta/max_memreq_MB	2,047	Default

Parameter	6.40 values	Parameter settings and description
ztta/parameter_area	8,000	Maximum 16,000 See *SAP note 3467*.
ztta/roll_area	16,773,120	Not used
ztta/roll_first	1	See *SAP note 103747*.
rdisp/wp_auto_restart	86,400	Do not use See *SAP note 182207*.
rdisp/wp_abap_restart	70,000	Do not use See *SAP note 182207*.
# ztta/as4/roll_shm_cpp_16MB	64	Do not use See *SAP note 371029*.

For Business Information Warehouse (BW) systems, see *SAP note 307077*, *SAP note 425593*, *SAP note 501572*, *SAP note 541508*, *SAP note 568139*, *SAP note 597312*, *SAP note 601109*, *SAP note 601110*, and *SAP note 641336* (Successor EVI-Stage2).

> **Note:** After the SAP kernel release of 640, SAP applications installed on System i models use UNIX_STD memory parameters. Documentation for the parameters and their settings can be found in *SAP note 33576*.

You can check an instance profile for consistency using the SAP Instance Profile Check Tool as follows:

- Logon as <SID>OFR or <SID>ADM
- Start the i5/OS PASE Shell within i5/OS:
  ```
  call qp2term
  ```
- Call the function:
  ```
  sappfpar pf=/usr/sap/<SID>/sys/profile/<SID>_DVEBMGS<nn>_<host> check
  sappfpar pf=/usr/sap/<SID>/sys/profile/<SID>_D<nn>_<host> check
  ```
  Example:
  ```
  sappfpar pf=/usr/sap/PRS/sys/profile/ERX_DVEBMGS00_ibas003 check
  ```

Adapt the RFC quotas in SAP transaction RZ12 in accordance to the adapted profile parameters beginning with *rdisp/rfc**

Refer to *SAP note 74141* for more information.

14

Software logistics

SAP programs and system configurations can be customized to better suit customer requirements. The SAP system can record these changes in the Change and Transport System (CTS). Within an SAP system, the Transport Organizer (TO) manages the customizing changes in *transport requests*. The Transport Management System (TMS) transports these customizing changes from the customizing system through the system landscape into the production system.

Transporting is the act of copying change requests from one client or one SAP system to another. The transport is accomplished by using the Transport Organizer, the Transport Management System, as well as the commands **`tp`** and **`R3trans`**.

Change requests contain program changes, customizing, or other configuration changes or data. Changes made to repository objects, such as table definitions, or Advanced Business Application Programming (ABAP) reports, or the so-called *development modifications* that the CTS manages, apply to all clients. These objects do not need to be transported between different clients in one SAP system. They must be transported from one SAP system to another.

The Change and Transport System helps to organize the customizing and development of objects when implementing the SAP applications. It also helps them to adapt to the customers requirements, and to transport changes from one SAP system to another. The contents of tables or other business objects on a higher level, such as authority definitions, can also be transported from one client to another and from one SAP system to another.

For more information about the Change and Transport System, the Transport Management System, or the Transport Organizer visit the SAP Help Portal at:

`http://help.sap.com`

14.1 SAP system landscape, clients, and roles

When you log on to an SAP system, you log on to a particular client of the system. Any activities performed in the system are executed in one client. Consider which clients you need for which activities when planning for your SAP system landscape. Assign activities to be performed in a client to give each client a particular role.

Each SAP system landscape requires a client where customizing settings, and possibly ABAP Workbench developments can be made. This adapts the SAP software for your own business needs. This client is known as the *Customizing and Development Client*, or *Customizing Client*. We use the abbreviation *CUST* to represent this client in this redbook.

Test the customizing settings and workbench developments extensively for errors before using them in a production environment. Faulty settings can seriously disrupt production operations, and lead to the loss of production information.

The integrated nature of the various SAP applications means that there are many dependencies between the different customizing settings. The accuracy of the settings can only be guaranteed with extensive testing, as even an experienced customizing developer may not discover these dependencies immediately. The client where these tests are made is called the *Quality Assurance Client*. We use the abbreviation *QTST* to represent this client in this redbook.

A separate client is required for *production* use of the SAP system. It is essential that no customized settings, workbench developments, or tests are made in this client so that the client can be used without disruption. This client is known as the *Production Client*, and we use the abbreviation *PROD* to represent this client in this redbook.

The central clients that exist in every SAP system landscape are as follows:

- CUST: customizing client
- QTST: test client
- PROD: production client

Create all customizing settings in a single customizing client. Then, use the CTS to transport them to the other clients. Do not create customization settings or workbench development in the quality assurance or production clients. Create appropriate client settings in the SCC4 transaction.

In 14.2, “SAP system landscapes” on page 183 we show you that:

- A client is always part of an SAP system.
- An SAP system is always part of the SAP system landscape.

You can set up other clients for other tasks in addition to these central clients. Each extra client takes up additional system resources (main memory and database space), and needs to be administered. For example, you need to set up and administer access authorization for the users, and distribute any changes to other clients with CTS.

Note: Consider the advantages and disadvantages before setting up other clients.

Other SAP clients and roles

The following are the roles of additional clients that can be utilized in an SAP system.

- Development test client (TEST)

  Use this client to test customization settings and workbench developments before releasing change requests. Developers can create test application data for realistic tests in this client. You can remove errors in the customizing client when you discover them.

  A development test client is set up in the same SAP system as the customizing client. This means that any change made to cross-client data in the customizing client is immediately visible in the development test client.

  Changes to client-specific data are copied from the customizing client to the development test client using a special client copy function, transaction SCC1. The client copy function uses the unreleased change requests from the customizing client to do this.

  The development test client is set in such a way that changes cannot be made to customizing data and repository objects.

- Prototype or sandbox client (SAND)

  Use this client to test any client-specific customized settings. Settings that you want to keep are then entered in the customizing client. Changes cannot be made to cross-client customizing data or repository objects in the prototype client. This is in order to prevent conflicts between the prototype client settings and real settings in the customizing client. The CTS does not record changes made to client-specific customizing data, and does not transport them from the prototype client.

- Training client (TRNG)

  Set up a training client to prepare end users for new functions that are to be transported into the Production Client. New functions can be used in this client with specially-created application data. The attributes of the TRNG client are set in such a way that a user cannot make changes to customizing data and to repository objects.

Additional clients can be used for special test purposes, to test data migration, to save the actual state of a client (customization and data), for a central user administration, and more.

14.2 SAP system landscapes

The number of SAP systems in a landscape are determined by the complexity of the implementation and the requirements of the organization. A three-system landscape provides a development environment, enables an environment for sufficient testing and quality assurances, and creates a separate production system. The appropriate system must be installed, and clients created to establish a three-system landscape. Once established, use a change request to *bundle* the changes made in the development system. Then, use the transport system to propagate the changes over the SAP system landscape.

This chapter discusses SAP three-system landscapes, two-system landscapes, as well as SAP one-system landscapes. See Chapter 3, “SAP system landscapes” on page 17 for a definition and introduction of landscape types, and the criteria explaining how to select a landscape for your business.

14.2.1 SAP system environments

In an SAP system landscape there are *environments* for different activities and roles, for example:

- Sandbox
- Development
- Customizing

- Test
  - Function test
  - Module test
  - Integration test
- Quality assurance
- Training
- Production

There can be even more *environments*, for example, to separate customizing and development environments for global and for multiple regional activities. Additional test system environments can be valuable for validating, data migration testing, and more. Moreover, even if there is a close data exchange between the systems, there can be multiple customizing or production environments, separated by country, location, or according to different security reasons (for example, human resources).

An SAP system environment can be represented by a:

- Computer (system, *machine* or *box*)
- Logical partition (LPAR)
- SAP system
- SAP client

There are a lot of different combinations *how* and with *which* types of *environments* you can implement and realize an SAP system landscape using computers, partitions, SAP systems, and SAP clients.

When planning for an SAP system landscape, you must size, configure, taking into consideration availability and economic factors in order to implement solid solutions that meet the required SAP system environments.

> **Note:** The SAP development system (DEV) and the SAP customizing system (CUS) are typically used as the same SAP system, because development and customizing is typically done in the same SAP system. The labels DEV and CUS can be considered synonymous.
>
> The SAP quality assurance system (QAS) and the SAP productive system (PRD) are distinguished separately from the DEV or CUS system.
>
> The term *test* system (TST) is not portrayed consistently. Sometimes it is used as a synonym for the CUS and DEV systems and sometimes it is used as a synonym for the QAS system.

The name of an SAP system is the so-called SAP system identifier or SAP system ID. Typically you see the abbreviation of SID, SAP SID, or SAP SYS-ID, or something similar. The customer is free to assign an SAP system name of his choice, keeping in mind the following rules:

- The SID must be unique throughout your organization.
- The SID must consist of exactly three alphanumeric characters.
- The SID must contain only uppercase letters.
- The SID must have a letter for the first character.
- The SID does not include any of the following reserved IDs:
  `ADD, ALL, AND, ANY, ASC, COM, DBA, END, EPS, FOR, GID, IBM, INT, KEY, LOG, MON, NIX, NOT, OFF, OMS, RAW, ROW, SAP, SET, SGA, SHG, SID, SQL, SYS, TMP, UID, USR, VAR`

14.2.2 One-system landscape

A one-system landscape contains of all central clients in a single SAP system. SAP does not recommend this configuration for a production system. Joint usage of hardware resources and cross-client data places serious restrictions on how a single system operates. In particular, after the system is used productively you can no longer develop in it, unless you stop production operations for the development and test phases.

You can use one-system landscapes to isolate test and training systems.

14.2.3 Two-system landscape

A two-system landscape is an alternative solution for smaller SAP implementations where little workbench development takes place. The two-system landscape does not include a separate quality assurance system. Normally the quality assurance client is also implemented in the DEV development system.

The disadvantage of a two-system landscape is that cross-client data is used in both the customizing and quality assurance clients. Therefore, any change that is made to cross-client data in the customizing client can affect the tests in the quality assurance client.

Transports from the customizing client are not guaranteed to be complete. Even if all tests in the quality assurance client are successful, errors can still occur after the transport into the production client. The changes being made to cross-client data in another transport request that has not yet been transported is the cause for this problem.

There are restrictions when you want to refresh the Customizing Client, or the Quality Assurance Client from the Production Client. In a three-system landscape, the QTST or the QAS is refreshed by means of a (remote) client copy or a (homogenous) system copy. In a two-system landscape both methods are critical, such as:

- It is unknown with a (remote) client copy whether the copied environment is consistent with the existing client-independent objects.
- Settings in the DEV system are overwritten with a (homogenous) system copy, including all customizing and development settings with its version management and documentation.

14.2.4 Three-system landscape

For the purpose of this discussion, the following naming conventions of the SAP Enterprise Resource Planning (ERP) systems are used:

- DEV (development system)
- TST (test system)
- PRD (production system)

The TST system may also be known as a Quality Assurance System (QAS).

There are many ways to distribute a three SAP system landscape on System i configurations, such as DEV, TST, and PRD. One way is to have all three landscapes on a single System i server with two logical partitions defined: one for DEV and TST, and a second partition for PRD. Another way is to have each system run on a separate System i server. The third option is to run any two SAP systems on one System i server and to run a third SAP system on another System i server.

These options are all the same at the conceptual level. These options are explored in more detail in Part 1, “SAP applications on System i models with i5/OS” on page 1 and especially in Chapter 3, “SAP system landscapes” on page 17.

SAP recommends you to isolate your SAP production system from your test and development systems. The following are several of the significant disadvantages of running a production system on the same System i server as the development or test system:

- Development activity can 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, impacts production users.
- Training classes run in a development or testing client can degrade performance in the production system.
- The i5/OS Programming Temporary Fix (PTF) impacts all environments: development, test, and production. Test all changes to the production system in a non-production environment first before applying them to the production system.

Each SAP system in the landscape is installed separately. A directory named /sapmnt/trans used by the Transport Management System is located on only one of the System i servers in the landscape. This key directory is referred to as the *global transport directory*. It consists of all the transport files that contain the changes that are transported between the SAP systems.

Typically, the DEV SAP system is installed first so that the link is created from /usr/sap/trans directory on this System i server to /sapmnt/trans directory. The Change SAP Shared Location (CHGR3SHLOC) command can be used to specify the server that contains the global transport directory when additional SAP systems, such as TST and PRD are installed at a later date on different System i servers. CHGR3SHLOC creates links using the *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. PRD is installed on host PRDSYS. With DEVSYS specified as the host name for the global transport directory, links are created after the CHGR3SHLOC command is run for TST and PRD, as Table 14-1 shows.

Table 14-1 Links for the global transport directory in the DEV system before relocation

SAP systems	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

Table 14-2 shows the links after the relocation of the global transport directory from the DEV to the PRD machine.

Table 14-2 Links for the global transport directory in the PRD system after relocation

SAP systems	Directory	Link
DEV	/usr/sap/trans	/QFileSvr.400/PRDSYS/sapmnt/trans
TST	/usr/sap/trans	/QFileSvr.400/PRDSYS/sapmnt/trans
PRD	/usr/sap/trans	/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:** Use the /usr/sap/trans directory on the production system. Since PRD might not exist at the beginning of an SAP implementation project (it is first installed in the DEV system), the /usr/sap/trans directory can be moved to the PRD system at a later time.
>
> You gain an advantage with this migration. First, you can get better performance for the import phase of an SAP transport because the transport request is physically on the same machine as the PRD system so the import must not be done via /QFileSvr.400 but rather done directly. Second, we can assume that the PRD machine has a higher availability than the DEV machine. Therefore, the common /usr/sap/trans directory also has a higher availability as it would be when physically staying on the DEV machine.

14.3 Manage clients after the installation of an SAP system

When an SAP system is installed, three clients are included by default: 000, 001, and 066. Client 000 is the SAP delivery client and 001 is the SAP standard client. Client 066 is for the SAP EarlyWatch Service and is not available for customer use.

Do not delete the reference client 000 and do not use clients 000 or 001 for customizing or development activities. Copy the standard SAP client 001 to a new client and use the new client for customizing and development activities.

Perform this copy using the SAP (local) client copy tool with the help of the SAP transaction SCCL. Check the result of the client copy with SAP transaction SCC3.

Each SAP client, including the CUST, QTST, and PROD, have client specific settings or attributes. We explain these settings in 14.5.1, “Client attributes” on page 197. You can alter these with the SAP transaction SCC4.

After the SAP system installation, client 001 is an exact copy from client 000. Do not modify client 000 with customer specific data. Client 001 can be used as a reference or for customer purposes.

The restrictions when selecting a client number are:

- Choose any three digits other than 000 or 066
- Use only digits (no letters or other special characters)

Figure 14-1 on page 188 illustrates using a copy of the standard SAP client 001 as client 200. Client 200 is where the customizing and development activities can take place.

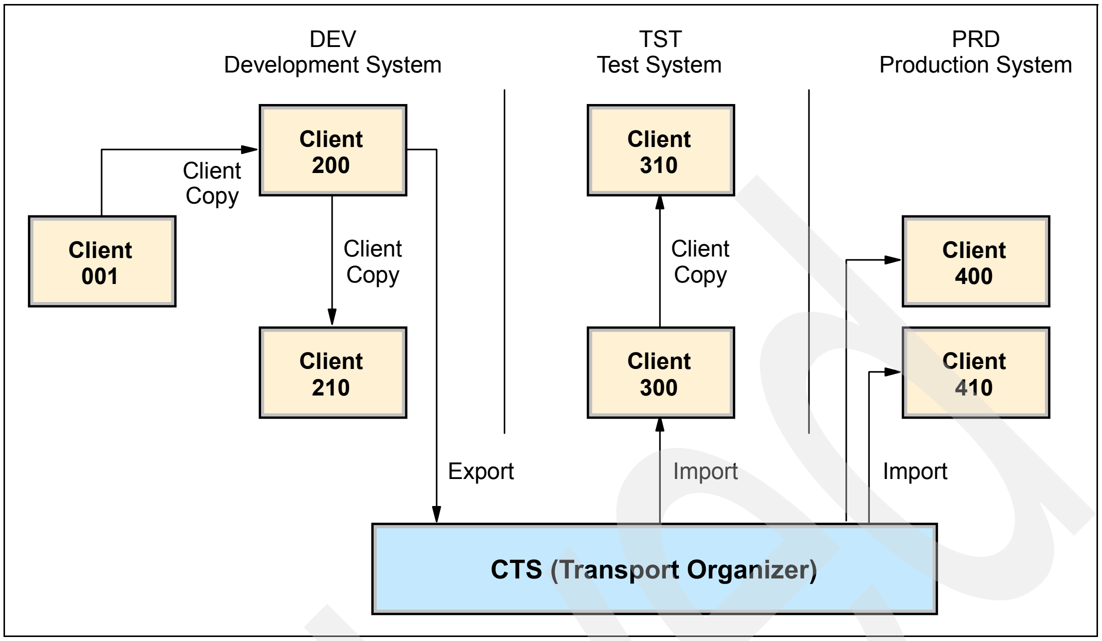

Figure 14-1 Client structures in a three-system landscape

Client 210 is created as a sandbox or work area that is regularly refreshed from client 200. The change requests generated in client 200 are exported and then imported using TMS to a clean client 300 in the TST. The TST is used for continuous testing upon completion of a development phase.

Typically client 300 is copied to a new client 310 in the TST and used for the training of end users. The changes are imported to client 400 and client 410 in the PRD system for preproduction integration testing when testing is verified. Further changes are imported only into client 400, the production client, when production begins. The 410 client can then be deleted.

Specify appropriate client change settings for each client and also for the system as a whole in the system change options. These settings define the various types of objects that can be modified. They also define if the changes made in a particular client are to be saved in a change request so that it can enable them to be transported.

The Transport Management System provides a central administration point for transports, called the *transport domain controller*. It defines the SAP systems and the transport paths between them.

The terminology *transport group* is used to refer to those SAP systems that share a common transport directory. *Transport domain* refers to one or more transport groups. See Figure 14-2 on page 189 for an example with the mySAP systems defined as follows:

- D01 - ERP Development System (domain controller for DOMAIN_D01)
- Q01 - ERP Quality Assurance System
- P01 - ERP Productive System
- BD1 - BW Development System (other domain controller for DOMAIN_BD1)
- BQ1 - BW Quality Assurance System
- BP1 - BW Productive System
- SM1 - Solution Manager

Figure 14-2 on page 189 shows the transport domain for a complex SAP system landscape.

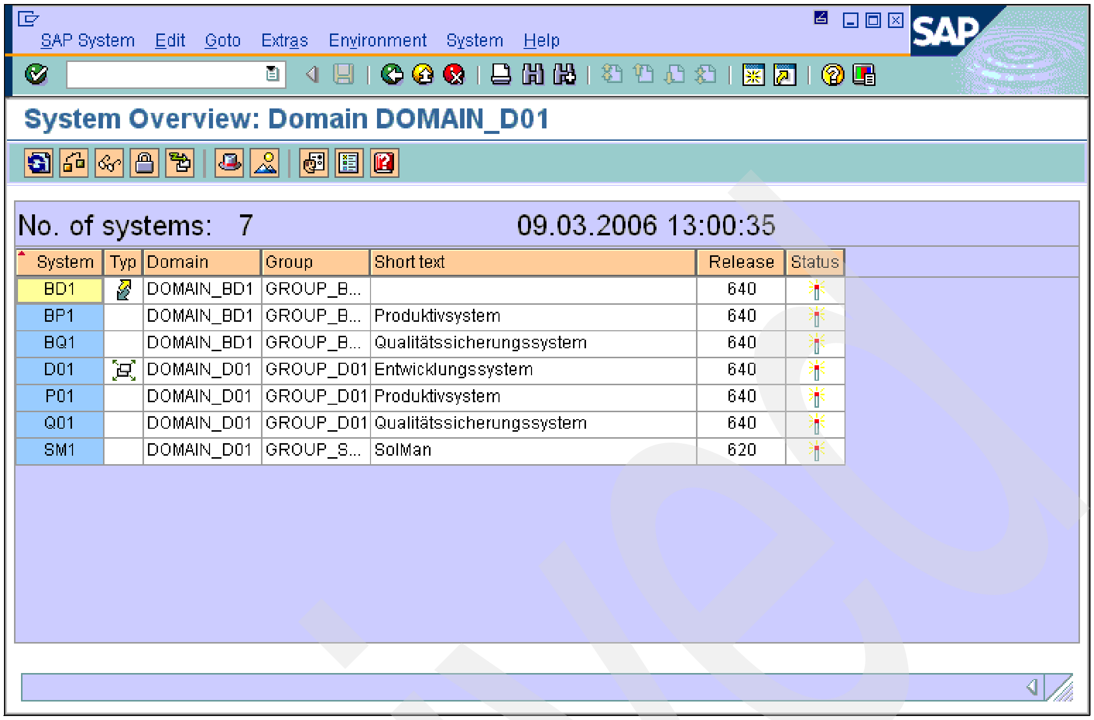

Figure 14-2 Transport domain for a complex SAP system landscape

> **Note:** i5/OS supports the network file system (NFS) as part of the standard operating system. This enables System i servers to share stream file systems with non-i5/OS such as AIX. This makes it easier to transport changes between the System i server and other systems.

14.4 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 system to another SAP system. The /usr/sap/trans directory is a symbolic link to one physical directory called */sapmnt/trans*. The /usr/sap/trans directory also contains a config subdirectory that contains the Informational authorized program analysis report (APAR) or also called *infoapar* files.

After SAP Web Application Server 6.40, all SAP system definitions that are formerly stored in this /config subdirectory are now completely covered by the SAP profiles in the path /usr/sap/SID/SYS/profiles as it is for other platforms.

> **Note:** See 8.2.4, “The Integrated File System” on page 63 for more information about the System i specific Integrated File System and the SAP specific directory structure.

If you have all three SAP systems on a single System i server, each SAP system's /usr/sap/trans directory has a symbolic link to the /sapmnt/trans directory on that system.

If you have multiple System i servers for three SAP systems, you have the /sapmnt/trans global directory on one of the System i servers, for example, the System i server for DEV. Other System i servers access this shared directory, which is referred to as the global transport directory, through the *QFileSvr.400* file system. See Figure 14-3 for reference.

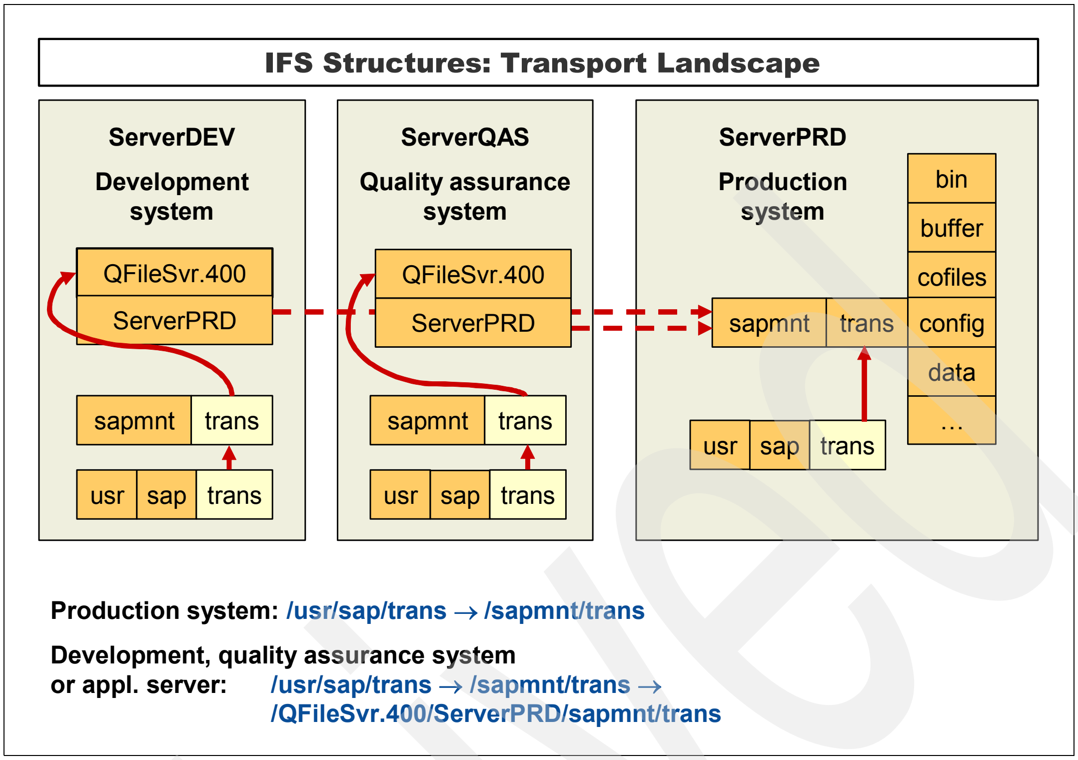

Figure 14-3 The SAP transport landscape

In a transport landscape, it can be useful to share one common directory between the development, quality assurance, and production systems. In this case, the data is physically located on the production system (in our example, ServerPRD), because that server is usually the most stable environment, and imports are usually twice as fast if the transport directory is accessed locally. Directory /usr/sap/trans is a symbolic link in all three systems.

> **Note:** In a homogeneous system landscape with remote Windows or Unix or Linux server there is no /sapmnt/trans directory on the remote hosts. The communication and data exchange is done with the /usr/sap/trans directory.

On the production system, it simply points to /sapmnt/trans. On the other two servers it accesses the remote server through the QFileSvr.400 remote file system.

The QFileSvr.400 file system needs the following to work correctly:

1. The host server *FILE must be started with the command STRHOSTSVR SERVER(*FILE). This is usually done automatically in the system startup program (QSTRUP).
2. Add the target host as a link under QFileSvr.400 with the following command:

   ```
   MD DIR('/QFileSvr.400/<target host>')
   ```

   Instead of MD you can also use MKDIR.

   This link is lost when you restart the server, so it is recommended to add this command to the QSTRUP system startup program.

Note: After an initial program load (IPL), these directories no longer exist. Therefore, add the MKDIR or MD commands to the i5/OS startup program (QSTRUP) of the System i server.

In 17.3.2, “QSTRUP in an SAP system environment” on page 307, you can find an explanation of how to maintain the QSTRUP program so that the /QFileSvr.400/<*hostname*> directory is not lost after an IPL.

3. A user profile with the same name and password must exist on the target host. This is described in *SAP note 67213*.

Use the Make Directory Command MKDIR or MD to create the QFileSvr.400 subdirectory for the host names of each System i server for which access is required.

The R3SETUP or SAPINST SAP installation tools have a menu option for changing the location of the /sapmnt/trans directory that executes the CHGR3SHLOC command. Enter the TCP/IP host name of the System i server with the global transport directory. Figure 14-4 shows an example of using the /sapmnt/trans shared directory with multiple System i servers in a three-system landscape.

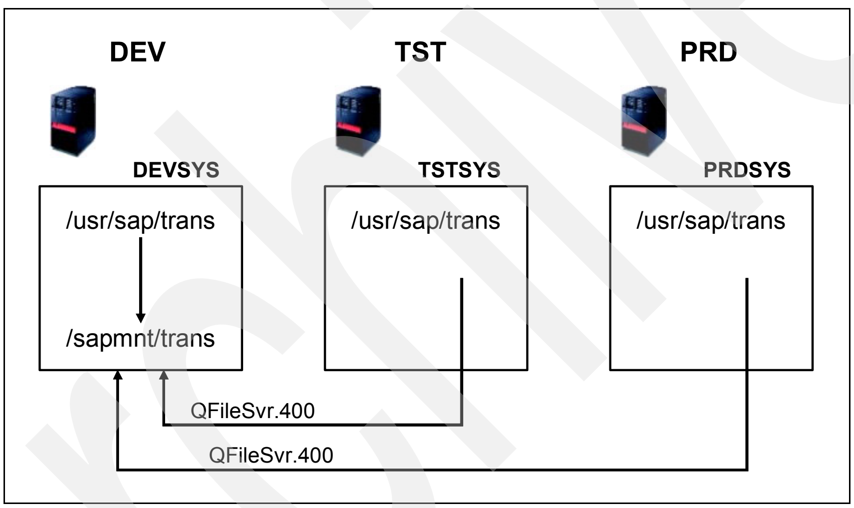

Figure 14-4 /sapmnt/trans shared directory

In this example, the /sapmnt/trans directory exists on only one System i server. A symbolic link named /usr/sap/trans points to /sapmnt/trans directory on the DEV system. It is recommended to move the /sapmnt/trans directory to the PRD system once the PRD system starts up.

To further understand how to work with the QFileSvr.400 system, refer to *Integrated File System Introduction*, SC41-5711.

Important: The <SID>OFR and <SID>nn i5/OS user profiles, where <SID> is the SAP System Identifier and nn is the instance number, must exist on the servers that share the common transport directory. Use the CRTSAPUSR *OFR and CRTSAPUSR *SIDINST commands to create them.

The passwords of QSECOFR, <SID>OFR, and <SID>nn user profiles 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 *SAP note 67213* for requirements of transports between SAP systems on different System i configurations.

We recommend that you locate the global transport directory on the database server in 3-tier system landscapes when you have separate database and application System i servers.

14.4.1 Setup the Transport Management System

The Transport Management System must be set up after the installation of an SAP system. From SAP Release 4.5A onwards, this is performed using the STMS Transaction. STMS automatically updates the /usr/sap/trans/bin/TP_DOMAIN_DEV.PFL file, where DOMAIN_DEV is the name of the transport domain, with the required parameters for each system, as Example 14-1 shows.

Example 14-1 TP_DOMAIN_D01.PFL (extract)

```
Browse : /usr/sap/trans/bin/TP_DOMAIN_D01.PFL
 Record :       1   of      70 by  14          Column :    1     76 by  79
 Control :

....+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....
 ************Beginning of data*************
#TMS:0009:DOMAIN_D01
#
#----------------------------------------------------------------------------
#Caution !
#This file was generated by the Transport Management System.
#Do not change this file using a text editor.
#For further information please see the online documentation of
#transaction STMS.
#If this file was destroyed, it can be regenerated in the TMS.
#To do this, log on to the domain controller (system D01) and call
#Transaction STMS. In the System Overview, you can distribute the TMS
#configuration. This regenerates the file.
#STMS -> Overview -> Systems -> Extras -> Distribute TMS Configuration
#----------------------------------------------------------------------------
#
COFILELIFETIME      = 365
DATALIFETIME        = 365
LOGLIFETIME         = 365
NEW_SAPNAMES        = TRUE
OLDDATALIFETIME     = 180
TRANSDIR            = /usr/sap/trans
#
D01/CTC             = 1
D01/DBHOST          = sapqe101
D01/DBNAME          = d01
D01/DBTYPE          = db4
```

```
D01/NBUFFORM       = 1
D01/TP_VERSION     = 266
#
P01/CTC            = 1
P01/DBHOST         = sappr101
P01/DBNAME         = p01
P01/DBTYPE         = db4
P01/NBUFFORM       = 1
P01/TP_VERSION     = 266
#
Q01/CTC            = 1
Q01/DBHOST         = sapqe101
Q01/DBNAME         = q01
Q01/DBTYPE         = db4
Q01/NBUFFORM       = 1
Q01/TP_VERSION     = 266
#
SM1/DBHOST         = sapqe101
SM1/DBNAME         = sm1
SM1/DBTYPE         = db4
SM1/NBUFFORM       = 1
SM1/TP_VERSION     = 266
#
#----------------------------------------------------------------------------
```

Prior to SAP Release 4.5A, the TP command used with the change and transport system requires parameters within the global transport parameter (TPPARAM) file. On these systems, 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 information that Example 14-2 shows.

To edit the TPPARAM file, issue the following command:

```
EDTF '/usr/sap/trans/bin/TPPARAM'
```

See Example 14-2 for a sample of the TPPARAM file.

Example 14-2 TPPARAM file listing

```
##################################################
# 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
```

DEV, TST, and PRD represent respectively a <SID> in Example 14-2. DEVSYS, TSTSYS, and PRDSYS are the TCP/IP host names of the database server on which the <SID> runs.

After the release 4.5B, the TPPARAM is no longer maintained. Instead, the file TP_DOMAIN_DEV.PFL is used. Set up a link to TP_DOMAIN_DEV.PFL to ensure that the data does not run out of sync. This is useful because every TP from the operating system level where *PF=..* is not explicitly used, uses TPPARAM. Therefore, as of release 4.5B, perform the following actions:

```
DEL /usr/sap/trans/bin/TPPARAM
ADDLNK OBJ('/usr/sap/trans/bin/TP_DOMAIN_DEV.PFL')
NEWLNK('/usr/sap/trans/bin/TPPARAM')
```

The SAP WRKLNKSAP command provides options for changing, copying, removing, displaying, renaming, moving, printing, and working with stream files. See Example 14-3.

Example 14-3 Work with Object Links (WRKLNKSAP)

```
                        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  10=Locks  11=Move   12=Work with

 Opt   Object link          Type      Owner               Size  Date
       .                    *DIR      R3OWNER            24576  Feb 21 11:43
       ..                   *DIR      R3OWNER            40960  Feb 27 12:42
       tpparam_instSM1      *STMF     L123504              145  Oct 17  2005
       trans.log            *STMF     QO1OFR             30811  Sep 29  2005
       DOMAIN.CFG           *STMF     D0102                700  Feb 28 10:28
       TMS_TEST.BD1         *STMF     D0102                  0  Jan 02 12:41
       TMS_TEST.SM1         *STMF     SM120                 31  Feb 21 11:43
       TP_DOMAIN_D01.BAK    *STMF     D0102               1998  Feb 28 10:28
       TP_DOMAIN_D01.PFL    *STMF     D0102               1998  Feb 28 10:28
       TP_DOMAIN_ID1.PFL    *STMF     ID130                872  Oct 12  2005
       TPPARAM              *STMF     D0102               1998  Feb 28 10:28






                                                                          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
```

14.4.2 Define the transport domain and transport routes

The Transport Management System is used to define the SAP systems that are part of the transport domain. To create the transport domain, log in to client 000 as user `DDIC` and choose **Tools** → **Administration** → **Transports** → **Transport Management System**, or transaction code **STMS**.

The basic setting and standard configuration is 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. You can define systems that are not yet installed as *virtual* systems initially and then change them to real systems later.

Figure 14-5 on page 195 and Figure 14-6 on page 196 show the setup of TMS after the installation of the development system for a three-system landscape. Figure 14-5 on page 195 shows the graphical representation.

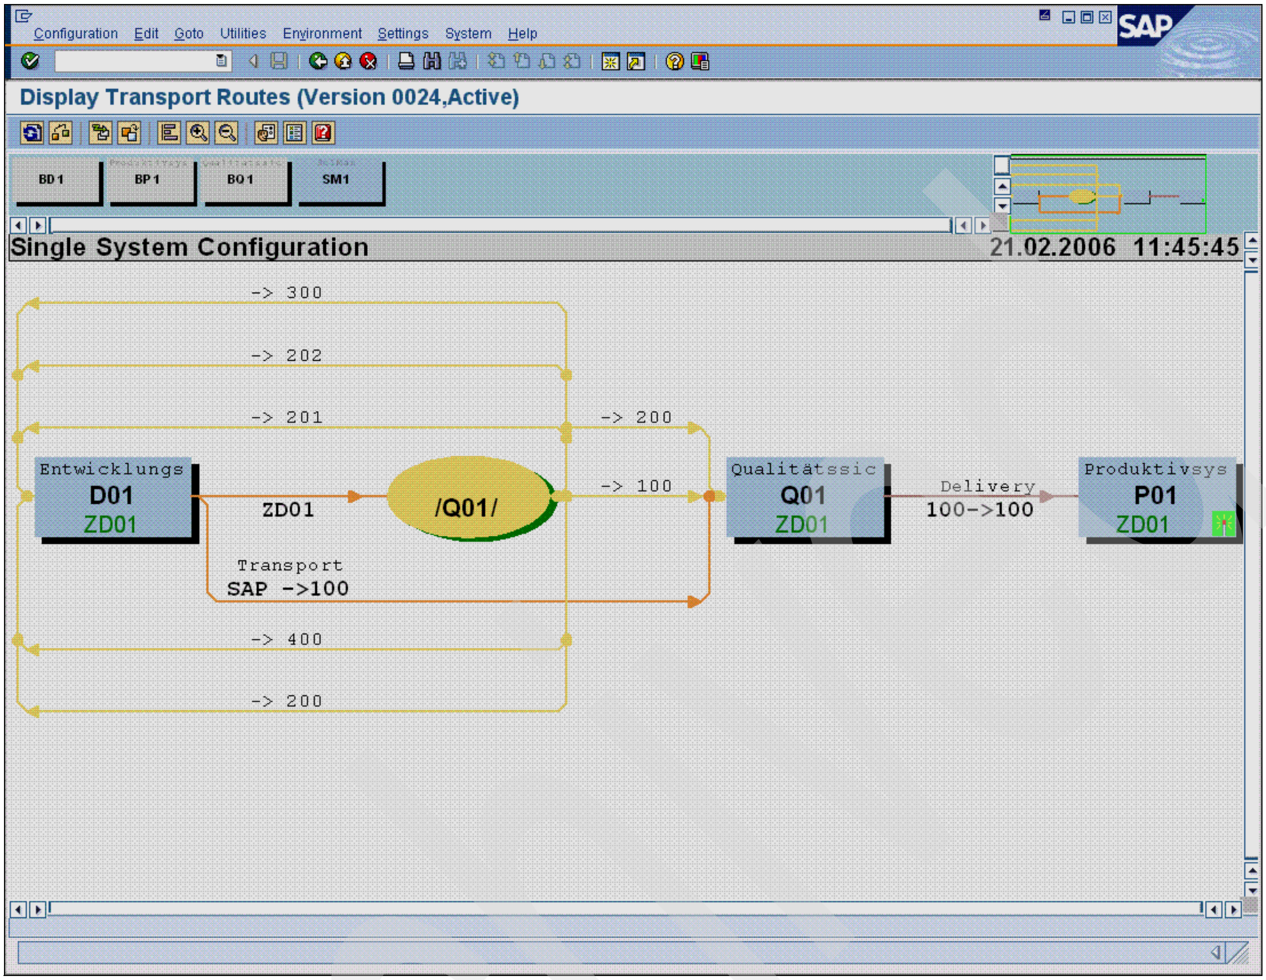

Figure 14-5 TMS setup of three-system landscape showing transport routes (graphical)

Figure 14-6 on page 196 shows the hierarchical representation.

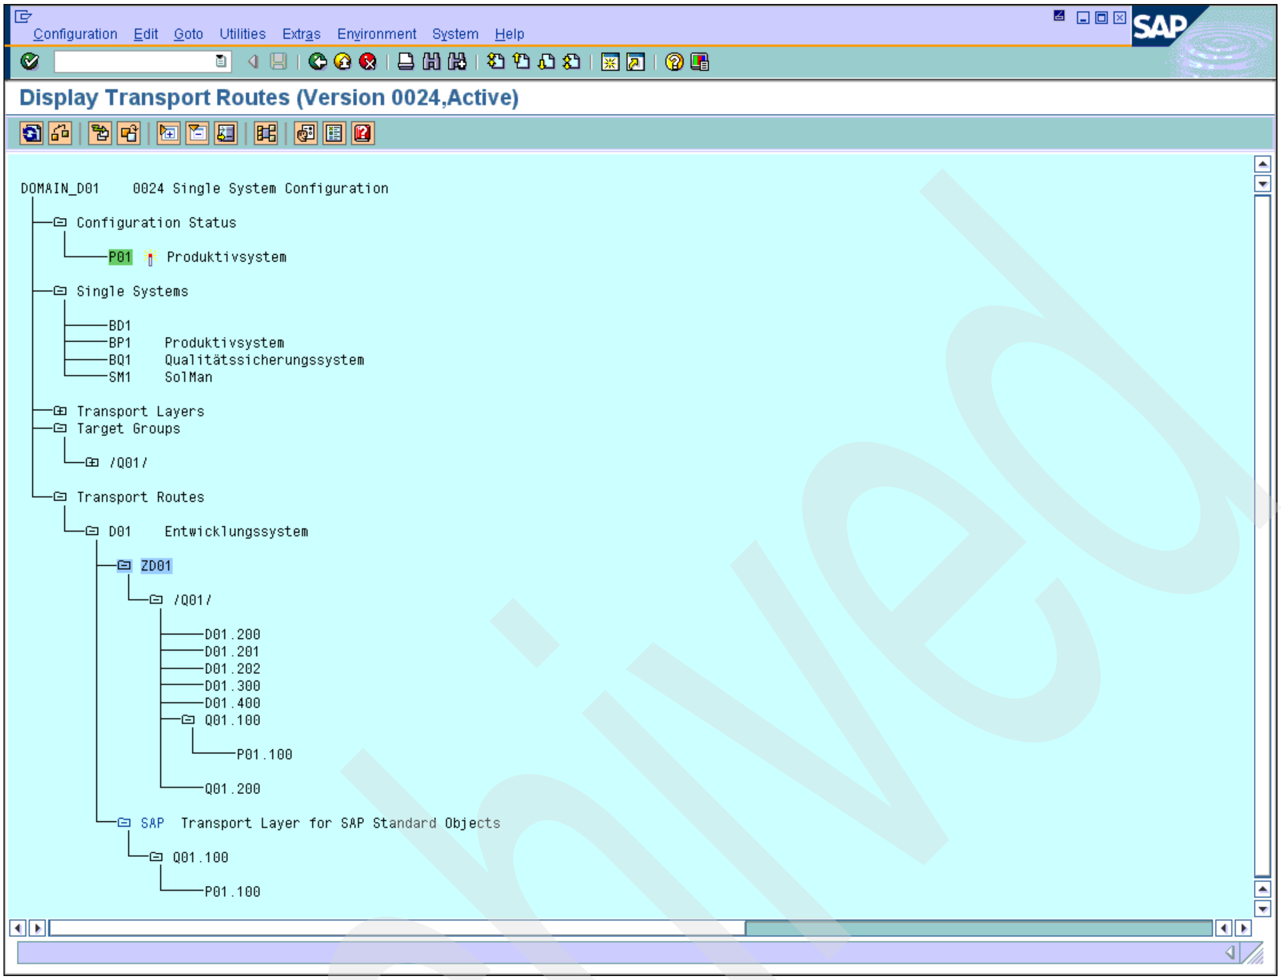

Figure 14-6 TMS setup of three-system landscape showing transport routes (hierarchical)

After installing TST and PRD, you can include them into the transport domain via RFC connections to the transport domain controller. The changed configuration is then redistributed to all SAP systems in the transport domain using an RFC from the domain controller.

The Transport Management System setup prior to SAP Release 4.6A creates the following table entries:

- Known SAP systems in the landscape (table TSYST)

  This defines all of the systems in the system landscape. You can create SAP systems that are yet to be installed as *virtual* systems in TMS and fully define them at a later time.

- Delivery routes or paths (table TASYS)

  This defines those systems for which transports are delivered automatically after successfully importing them into the consolidation system.

- Consolidation routes or paths (table TWSYS)

  This allows the transporting of 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.

TMS includes an entry for a delivery system with a <SID> of SAP to accept the changes and imports from the SAP system directly, such as via SAP Support Packages. Support Packages are referred to as *hot packages* and *legal change patches (LCPs)* on releases prior to release 4.5.

From SAP Release 4.6A onwards, the four tables are replaced by the tables:

- System List (table TCESYST)
- Deliveries (table TCEDELI)
- Transport Layers (table TCETRAL)

These tables contain a field for the Version number that allows you to store and retrieve different TMS configurations as needed. The TMS version information is stored in table TCEVERS.

> **Note:** Do not maintain the SAP tables TCESYST, TCEDELI, TCETRAL, TCEVERS manually.

14.5 Customization and development process in an SAP system

Customizing is an SAP term that means making appropriate changes to the SAP-supplied system that parallels a company's needs. The objects that these changes to the development system affect are moved to the test and production systems through the Transport Management System (TMS).

The term *objects* can 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 SAP repository.

All objects and data within an SAP system 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. The client dependent tables have the first field named MANDT. This is a key field that designates an entry to a specific client. Client independent objects cross all clients (a change to this table affects all clients in the system).

14.5.1 Client attributes

Client attributes allow development efforts in each client to be controlled and are assigned for each client in an SAP system. To check client attributes, select transaction **SCC4**. Or select from the SAP Easy Access Start menu (*SESSION_MANAGER*) **Tools** → **Administration** → **Administration** → **Client Administration** → **Client Maintenance**.

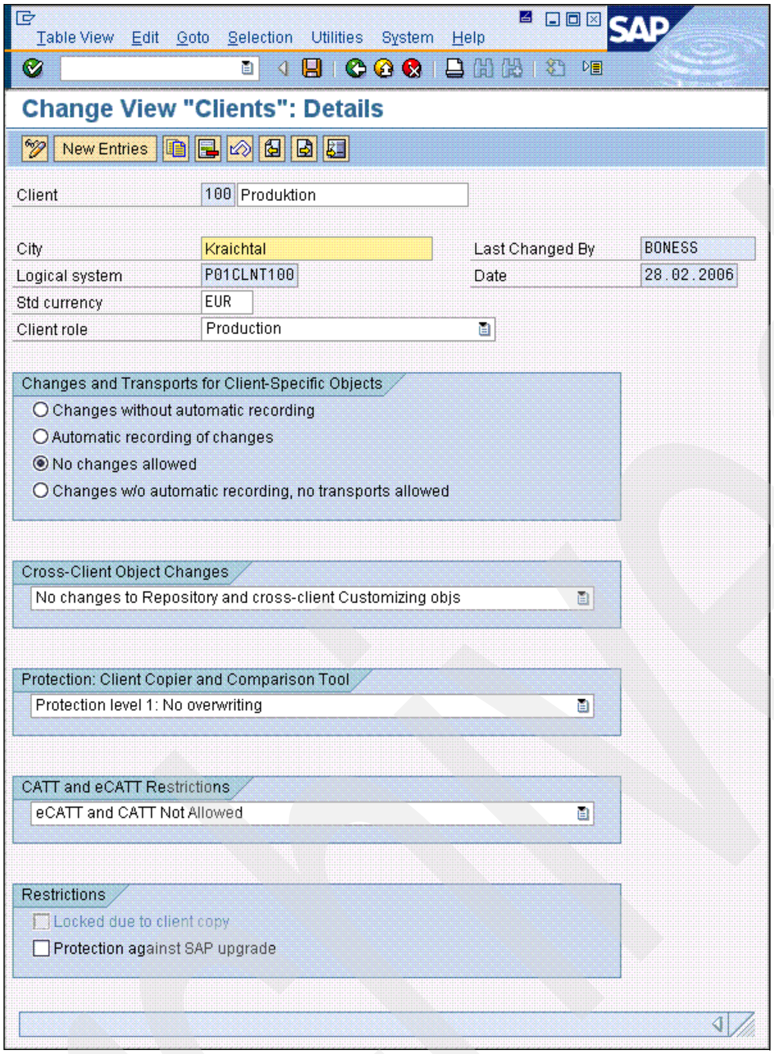

Figure 14-7 Client-dependent attributes

A client can have different roles in an SAP system:

- Production
- Test
- Customizing
- Demonstration
- Training and education
- SAP reference

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 you to change customizing table entries (the default). Changes to customizing and table entries can be included in a task manually or through table view selection.

- Automatic recording of changes

  Allows you to change customizing table contents and automatically includes it in a task.
- No changes allowed

  This display-only mode prevents users from making changes to customizing entries.
- Changes without automatic recording, no transports allowed

  You can change customizing table entries, but they cannot be included in a change request. You can use this option to isolate a demonstration or training client.

Customizing transactions are used to maintain customizing using the Implementation Guide (IMG). System parameters and table views, for example, 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 system and then production systems.

Change requests are numbered in the format <SID>K9xxxxx, where <SID> is a three-character system identifier and xxxxx is an SAP system supplied number. For example, DEVK900065. The tasks under this change request usually follow sequentially. That is, DEVK900066 is the change request number and DEVK900067 is the task. You can use a change request to create more than one task.

At the start of a development project, the project leader can create a change request manually and assign the project team members to it. On the other hand, SAP applications can 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.

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 **SE01**, **SE09**, or **SE10**.

Tasks are the lowest in hierarchy level of the transport system. Key reference to tasks is what records all saved customizing and configuration changes. Each task is owned by an individual user master record. The task contains the actual transport objects that can consist of table entries or objects such as SAPscript forms. The owner of the task or a transport system administrator 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.

You can view the results of a transport in the definition window or transport requests if you choose the menu option **Goto** → **Transport Logs**.

14.5.2 Transport change requests

Once a change request is released and exported, it needs to be transported into the test system. Transporting is accomplished via the transaction STMS, which calls the SAP system TP command in the kernel library. The TP command can also be entered manually on the i5/OS command line, although this method is not recommended. STMS ensures that the steps in exporting and importing objects are performed in the correct order.

When a change request is released, the objects it contains are extracted from the database of 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 14-3 are created.

Table 14-3 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. In the import phase, the requests in the transport buffer are applied to the test system. At the same time, the requests are added to the buffer of the production system. Once the change is tested in the test system, it can then be imported into the production system.

14.6 Handling of a repository object change request

For the handling of a repository object change request, the system names in the landscape are DEV, TST, and PRD.

The following sections show 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
- Import the ABAP program into the test system

14.6.1 The SAP Software Change Registration (SSCR)

If you want to create or modify SAP objects within the SAP repository you must register in the SAP Service Marketplace:

- A developer who is allowed to do these changes

  A developer need to be registered *only once for each SAP installation number*. An SAP installation number is related to an SAP application like mySAP ERP or SAP BW. Within an SAP installation number you can run one productive SAP system together with four non-productive SAP systems, in sum five SAP systems.

  You need a developer SSCR key in any case when you create, change or delete an object from the SAP repository. It doesn’t matter if you do this with SAP specific objects or with objects of (your own) customer namespace.

- The SAP object that is to be created or changed

  An object SSCR is necessary *once* when you want to modify an SAP specific object belonging to an SAP installation number, meaning an object of the SAP repository that belongs to the SAP namespace. After this registration each registered developer can work with this object.

Note: See the SAP Online Help `http://help.sap.com` or the SAP library with the keyword "*SAP namespace*" to check the SAP namespace.

Register the SAP user ID as a Developer in the SAP Service Marketplace at:

`http://service.sap.com/SSCR`

On this page there are the following functions:

- The registration itself
- List of objects that you register
- List of developers that you register

Additionally you can retrieve a user documentation.

Note: If you are not able to evoke this Web site you might not have the correct user authorities in the SAP Marketplace. Here you need the authorization *Register SSCR keys*.

The SAP Service Marketplace customer IDs are maintained within the path from the menu: **Data Administration** → **User Data** → **Maintain User Data**.

In the next window you can select **Register developer** or **Register object**. This registration depends on your SAP installation number that SAP gives you. Figure 14-8 shows you the input fields for registering an SAP object.

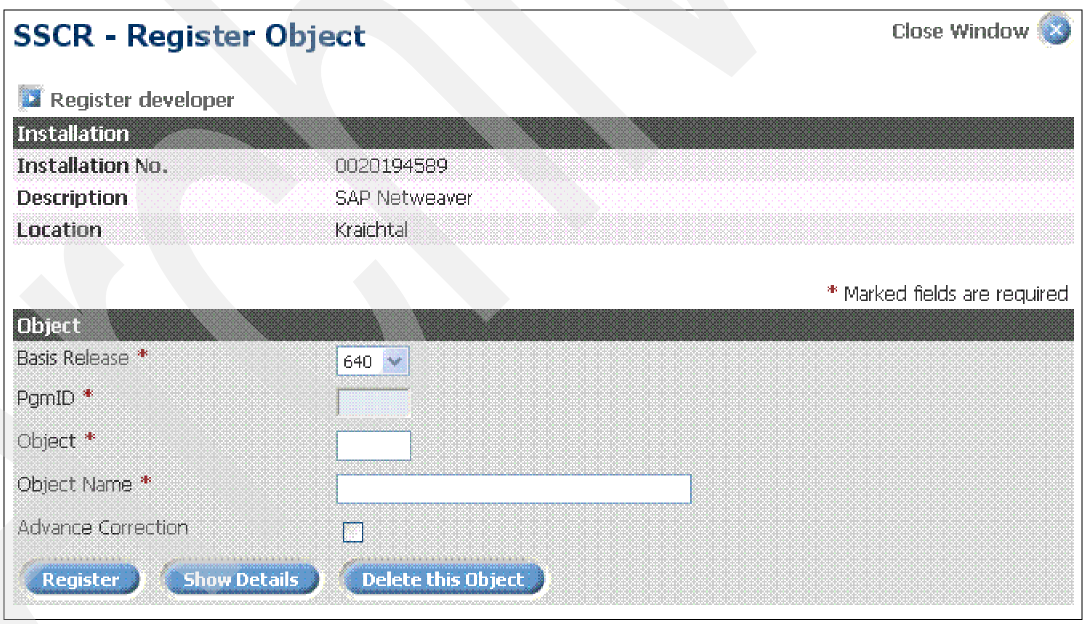

Figure 14-8 The SAP Software Change Registration (SSCR) for an SAP object

You enter the following input fields for your SAP installation:

- Basis Release

  Enter your Web Application Server release, for example 640. There is a match code for this input field.

- PgmID, meaning program ID

  You get the program ID from your SAP application when you are requested to register an object. There is a match code for this field.
- Object, meaning object type

  You get the object (type) from your SAP application when you are requested to register an object. There is a match code for this field.
- Object name

  Enter the object name according to the SAP naming conventions.
- Advance correction

  This is the only input field that is not required. Use it for documentation purposes when you do an advance correction of an SAP standard object. You see this marker in the list of the registered objects.

The SAP Software Change Registration (SSCR) is a procedure, valid from Release 3.0A onward, which registers all manual changes to SAP sources and SAP dictionary objects. SAP application match codes and tuning measures such as the setup of database indices and buffers are excluded from the registration. The availability and reliability of productive SAP installations is further improved with this procedure.

SAP note 86161 explains how to register developers and objects.

SAP note 27532 explains the SAP Software Change Registration (SSCR).

14.6.2 Creating a repository object change request

Complete the following steps to create a repository object change:

1. Click **Tools** → **ABAP Workbench** → **Development** → **ABAP Editor**, or call transaction **SE38**.

   On the ABAP Editor window, give the program a name (for example, ZTMSPGM) and click the **Create** button. The ABAP Program Attributes window appears.
2. When prompted, enter the developer key that SAP assigned to you. You are not prompted to enter the developer key each time you make a change or create a new object.
3. 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.
4. Use transaction SE80 to create a development class. You must create at least one development class before change requests can be created. Select the appropriate development class and click **Save**. The Change Request Query window is shown.
5. Click the **Create request** button. The Create Request window opens. Enter a short description, and click **Save**. You are returned to the Change Request Query window. The request number is automatically created. Press **Enter**.
6. The ABAP Program window opens. Click the **Source code** button.
7. 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 ztmspgm
   write: / 'This is a test of SAP TMS'
   ```

14.6.3 Transport Organizer

You can reach the Transport Organizer using transaction SE09 or SE10, as illustrated in Figure 14-9.

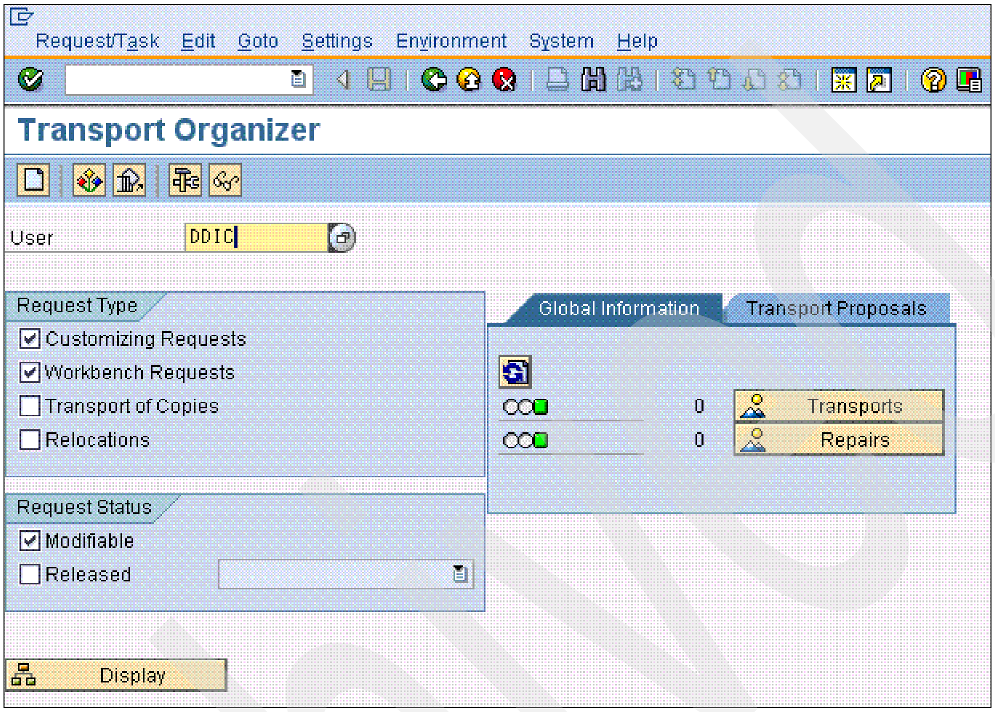

Figure 14-9 SAP transport organizer starting window (transaction SE10 or SE09)

To reach the Transport Organizer, perform the following steps:

1. Type transaction `SE10`. This shows the Transport Organizer window.

   The Modifiable and Released options are both selected by default. To select only requests that are not 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 shown. You can only release tasks or change requests that your ID owns.
2. Click the **plus icon** next to the change request number that was created in the previous steps. The task number 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 the objects are now being exported appears. Press **Enter**.
6. You are returned to the Transport Organizer where the change request is reflected under the *Released* category. To view the export logs and 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.

14.6.4 Strategies for importing requests

There are three methods available to control the import of change requests into the consolidation and delivery systems, depending on your system landscape and the requirements of the project team. 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 library under **Change and Transport System - Overview** → **Transport strategy in the CTS** for a more detailed explanation of the different options available. And refer to the SAP documentation online help under **Transport Management System** → **Performing Transports** for more details on how to proceed in each case.

14.6.5 Workflow strategy for imports

Use the *transport workflow* to automate your transport process. It is an efficient method of transporting a selected number of requests into a group of transport targets, and uses clearly defined approval steps to ensure the quality of your target systems.

If you want to use a *workflow oriented strategy* you must define a TMS workflow engine. See Figure 14-10.

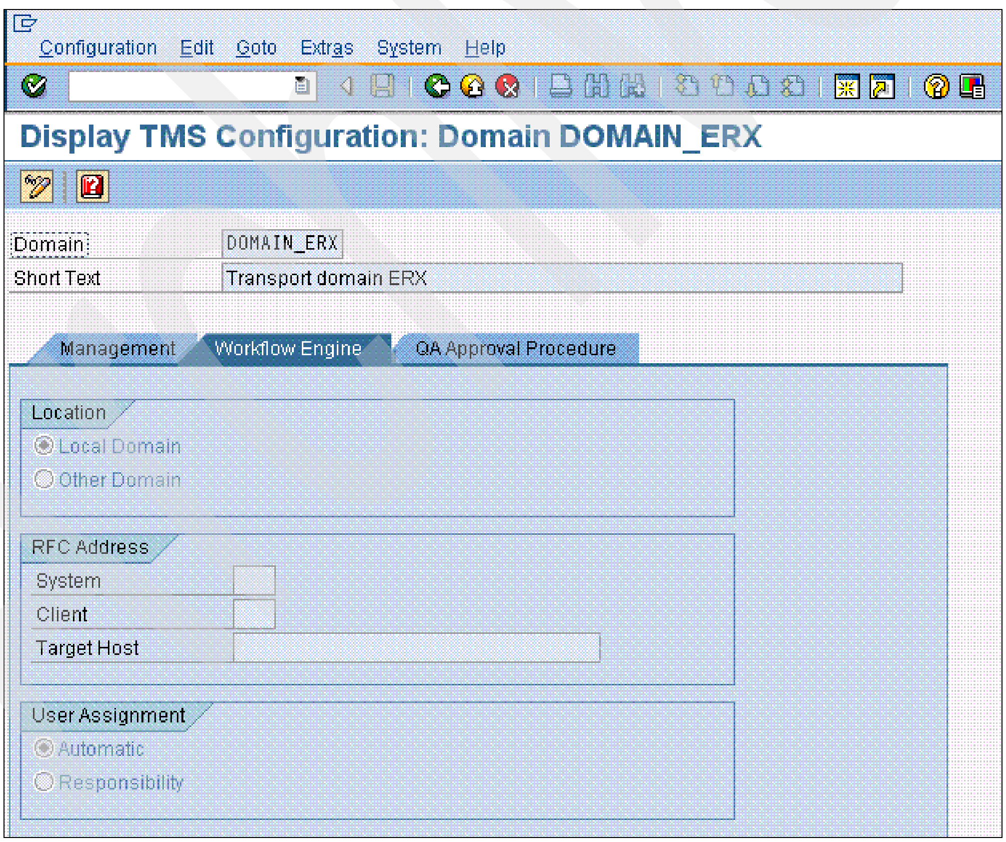

Figure 14-10 TMS Workflow Engine

The prerequisites for the TMS Workflow Engine are that:

- Transport administration has configured the transport workflow.
- You have a user in the Workflow Engine system/client.

After you have released the requests, you must create a transport proposal in the Transport Organizer. To do this, specify the requests, and the target systems/clients into which you want to transport them. You can also specify an import time, import options, and notes for the transport administration.

> **Note:** If you have selected the transport workflow as a transport strategy, the transport proposal appears directly after the selected requests are released. Choose to trigger the physical release of the requests.

When you create a transport proposal, you place it in the TMS work list for the transport administration. The transport administration can approve the transport proposal and it is then imported automatically into the target systems. The transport proposal appears again in your transport proposal inbox for you to confirm. Confirm it to complete the proposal, or apply to have it transported into other systems.

However, transport administration can also reject the transport proposal. It then appears in your transport proposal inbox again, possibly accompanied by a note. You can then modify the transport proposal and place it in the workflow work list of the transport administration.

For more details refer to the SAP online documentation or directly from the SAP Help Portal under *Transport Workflow* in the Section Transport Organizer (BC-CTC-ORG):

http://help.sap.com

14.6.6 Use STMS to import single requests

Release and export the change request. Then to import the change request into the TST, follow these steps:

1. Sign in to the SAP target system.
2. Call Transaction **STMS**.
3. Select the **truck** symbol. The import overview appears.
4. Position the cursor on the SAP system to import requests. Choose the **glasses** symbol. The import queue of the selected SAP system appears.
5. Position the cursor on the request that you want to import.
6. Choose the **truck** symbol.

   The dialog box Import Transport Request appears. The box displays the request you chose and the import target.
7. Select your choices for settings, and choose **Continue**.
8. The Start Import dialog box appears. Check and confirm the information in this dialog box. The box shows the system and the target client that you selected for the import.
9. After the import is finished, check the import history for the return codes of each of the import steps. From the import history, go to the transport logs to check the reason for any errors, select **Goto** → **History** → **Import History**.

To control the import of a change request there is a selection screen within the transaction STMS with three tabs. The next three figures show you examples on how to use these tabs:

- Figure 14-11 on page 206 shows importing a change request with the *Date* selection.
- Figure 14-12 on page 206 shows importing a change request with the *Execution* selection.
- Figure 14-13 on page 206 shows importing a change request with the *Options* selection.

These selection screens also represent the same *settings* as that of step 7 listed previously.

Importing a change request with the *Date* selection as Figure 14-11 shows.

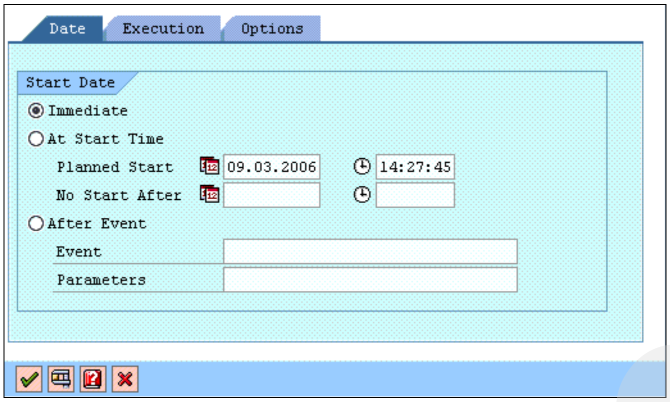

Figure 14-11 Import a change request: Date selection

Importing a change request with the *Execution* selection as Figure 14-12 shows.

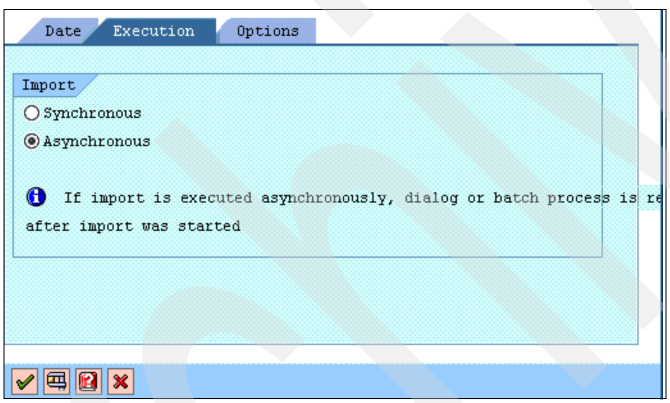

Figure 14-12 Import a change request: Execution selection

Importing a change request with the *Options* selection as Figure 14-13 shows.

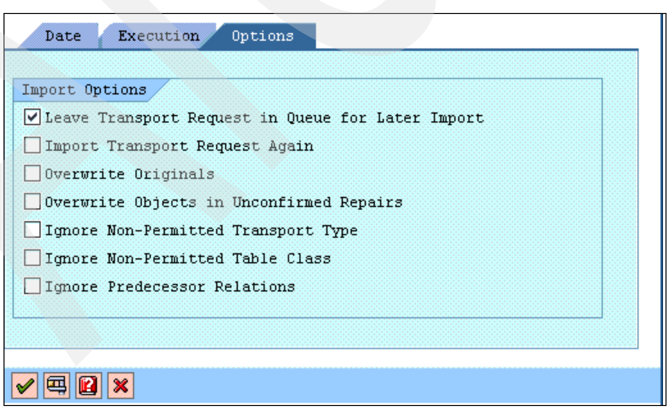

Figure 14-13 Import a change request: Options selection

The *Options* selection are also known as the *unconditional modes* that can be set explicitly and manually with the **tp** command, as described in 14.6.7, "Use TP to import single requests" on page 207.

14.6.7 Use TP to import single requests

Follow these steps to import a single request into the TST using TP:

1. Sign in to the System i server of the target system (TST) as <SID>OFR (TSTOFR in this example). Make sure the kernel library for the SAP implementation is in your library list (maybe R3xxxOPT where xxx is the SAP release). This is the case if you log in under the <SID>OFR user profile.
2. Use the **`Change Current Directory`** or **`Change Directory`** command and type the following to change the current directory to the location of the TP parameter file TPPARAM:

   ```
   cd '/usr/sap/trans/bin'
   ```

3. Alternatively, select the **Change and Transport System** options from the SAP main menu. Then change the current directory to /usr/sap/trans/bin.

   > **Note:** You can control the program tp using the transport profile. Each time tp is started, it must know the location of the global parameter file. Always use tp with its profile parameter. The syntax is:
   >
   > ```
   > tp 'pf=<path_for_your_TP_DOMAIN_SID.PFL' <command>
   > ```
   >
   > For example:
   >
   > ```
   > tp 'pf=/usr/sap/trans/bin/TP_DOMAIN_PRD.PFL import all PRD'
   > ```
   >
   > If the transport profile is not specified in this way, tp searches for the transport profile in the current directory.

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 the command prompt.

5. Type the following command, where xxxxx is the change request number:

   ```
   tp 'import DEVK9xxxxx TST'
   ```

6. You will receive an indication that the TP import completed successfully.
7. Go back to SAP transaction **SE10**. Select the box labeled **Released**. Press **Enter**.
8. The change request is listed in the Released category. Click the change request and select **Goto** → **Transports Logs** from the pull-down menu view the log.

This completes a release and import of a change request.

> **Note:** Use the TP only in special situations.
>
> For a detailed description of the use of the TP, see the SAP library or the SAP Help Portal under *Transport Tools* or *BC-CTS-TLS* at:
>
> http://help.sap.com

Unconditional modes

Unconditional modes are intended for special requirements when you want to ignore certain rules of the Transport Organizer and Transport Management System. An unconditional mode is represented by a digit between 0 and 9. To use the unconditional mode when you call the TP program, enter a `U` in the command line and then add the required digits. For example:

```
tp import T11k904711 P11 U06
```

Table 14-4 lists the individual modes or unconditional modes.

Table 14-4 The tp individual or unconditional modes

Individual or unconditional mode	Type	Text / Activity
0	Overtaker	Import from the buffer without deleting and set unconditional mode 1 in the buffer to trigger another import
1	During export	Ignore the incorrect status of the command file
	During import	Ignore the fact that the change request was already imported
2	During export	Do not expand the selection with TADIR brackets
	During import	Overwrite original objects
3	During import	Overwrite system-dependent objects
6	During import	Overwrite objects in unconfirmed repairs
8	During import	Ignore restrictions based on table classification
9	During import	Ignore the fact that the system is locked for this type of transport

14.7 Transport objects between SAP systems on different host systems

In the following example, a program is exported and released for transport (using SAP transaction SE09) on the machine with a host name of SYSNM001 and an SAP source system name of R01. The transport request named R01K900240 contains the co-files member named K900240.R01and the data member named 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.

14.7.1 Homogeneous environment

In a homogeneous environment (with an System i source and System i target), the /sapmnt/trans directory is on one machine. The /usr/sap/trans symbolic link and the QFileSvr.400 make the directory contents available for the other machine:

```
(/usr/sap/trans --> /QFileSvr.400/SYSNM001/sapmnt/trans)
```

In this case, the symbolic link /usr/sap/trans on the source system points to the Integrated File System directory /sapmnt/trans of the target system with the host name SYSNM001.

14.7.2 Heterogeneous environments

In all SAP system landscapes with all platforms it is not necessary that the transport directory is shared between the SAP systems.

In homogeneous i5/OS environments the QFileSvr.400 normally shares the transport directory. In heterogeneous environments with mixed platforms like i5/OS, Windows, Linux, or other, the sharing of the transport directory can be done with NFS (Network File System). However this is not mandatory so we distinguish in the following between NFS connected and no NFS connected machines.

When using SAP systems on Windows, NFS is not required for transporting between the SAP systems. System i NetServer™ or System i QNTC file systems can also be used.

NFS connected heterogeneous environment

In a heterogeneous environment where NFS is available (for example i5/OS source and AIX target), one machine has the /usr/sap/trans directory and exports this directory to the other machine.

> **Attention:** Windows user names can be longer than ten bytes, an invalid length for i5/OS objects. Therefore, a guest user profile is required to obtain the authority from Windows.

In the following example, the source machine is named SYSNM001:

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 functions are shown based on i5/OS commands.

In the case of a heterogeneous environment, replace the non-i5/OS side with the import command which that particular system supports. If the global transport directory is located on the i5/OS system, proceed as follows:

1. Use the EXPORTFS command to export all subdirectories of /sapmnt/trans/ with the *Data file code page *ASCII* option, except for the directory /data:

```
HOSTOPT(SYSNM002 *ASCII *ASCII)
```

2. Use the EXPORTFS command to export the /data subdirectory as BINARY:

```
HOSTOPT(SYSNM002 *BINARY *ASCII)
```

> **Note:** NFS can also be used in homogeneous System i environments. However, we recommend the use of QFileSvr.400, because it performs better.

For more information about i5/OS Network File System support, see *V5R3 Network File System Support*, SC41-5714.

> **Note:** In heterogeneous system landscapes you have to take care that the operating system user ID (UIDs) and group ID (GID) match between the systems. See *SAP note 818091* for an example with a Linux application server.

Non-NFS connected heterogeneous environment

In heterogeneous environments where the NFS function is not available to connect both file systems (for example, Windows and i5/OS), the Windows machine and the System i server each have a /usr/sap/trans directory. You must transport the files in the /usr/sap/trans/cofiles and /usr/sap/trans/data directories to the i5/OS integrated file system using FTP.

Use the following steps to sign in to the System i server and move the exported data from the Windows NT® system to the i5/OS integrated file system:

1. Sign in to System i 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 (Integrated File System name format).
4. Change the local and remote current directory.
5. FTP the co-files member (sub command get).
6. FTP the data member (sub command get, binary transfer type)

   See Example 14-4 for the FTP data member.

Example 14-4 FTP data member to the Integrated File System

```
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 co-files member is transported using ASCII - EBCDIC translation. The data member is transported in binary format.

7. Fix authority after FTP using:

   ```
   CHGAUT OBJ('/usr/sap/trans/cofiles/K900240.CUS') USER(R3GROUP)
   DTAAUT(*RWX) OBJAUT(*OBJMGT *OBJREF)
   ```

> **Note:** When you use ftp you often make a file transfer with an i5/OS save file (*SAVF). To later transfer the contents of the *SAVF into the Integrated File System you can use the i5/OS command Copy to Copy to Stream File (CPYTOSTMF). Take care of its usage:
>
> - The CPYTOSTMF command must use the ENDLINFMT(*LF) parameter for text files (ASCII).
> - The CPYTOSTMF command must use the ENDLINFMT(*FIXED) parameter for the data files (binary).
>
> You also have to distinguish these different file types (ASCII and binary), if you use the CPYTOSTMF command, when you transfer the SAP co-files (in ASCII format) and the SAP data files (in binary format) of the Transport and Management System (TMS).

For more information about the FTP command, see the *TCP/IP Configuration and Reference*, SC41-5420.

14.8 Test the Transport Management System

To test the Transport Management System, export a user-developed program to the SAP development system (DEV). A transport request number (DEVK9xxxxx) is received. Sign in to the System i server containing the SAP test system (TST), and issue the following commands using user ID <SID>OFR, where <SID> is the target SAP system.

```
cd '/usr/sap/trans/bin'
tp 'ADDTOBUFFER DEVK9xxxxx <SID>'
tp 'import DEVK9xxxxx <SID>'
```

DEVK9xxxxx is the transport request number (for example, DEVK900240), and <SID> is the system identifier of the target system (for example, TST).

> **Note:** Authorize the user ID performing the TP command 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 i5/OS 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
```

15

The SAP system of printing

The SAP spool system is separate from i5/OS spool support. This chapter describes the fundamentals of SAP system printing, gives a background of i5/OS print support and terminology, and provides several examples of printing definitions. It is written from the perspective of an SAP application user.

15.1 Introduction to printing with SAP applications

SAP provides an internal spool support interface to ensure a consistent printing interface across the diverse platforms supported by SAP. To-be-printed data that is produced by an application running under SAP goes into an SAP-internal spooled file. The data in the SAP spooled file is not in an altogether suitable format for printing. The data is ready for printing after SAP transforms it into an SAP-compatible final form and takes one of the following two actions:

- Sends the to-be-printed data to the System i server where it is stored in a spooled file, and managed and printed using standard i5/OS printing facilities.
- Sends the to-be-printed data to a remote print server using the standard Berkeley LPD protocol, or a special SAP LPD protocol for those systems that do not support the standard Berkeley protocol. The remote print server where it is printed might or might not be another System i server.

Figure 15-1 illustrates the flow of data from its origin within an SAP application to its arrival on an i5/OS output queue.

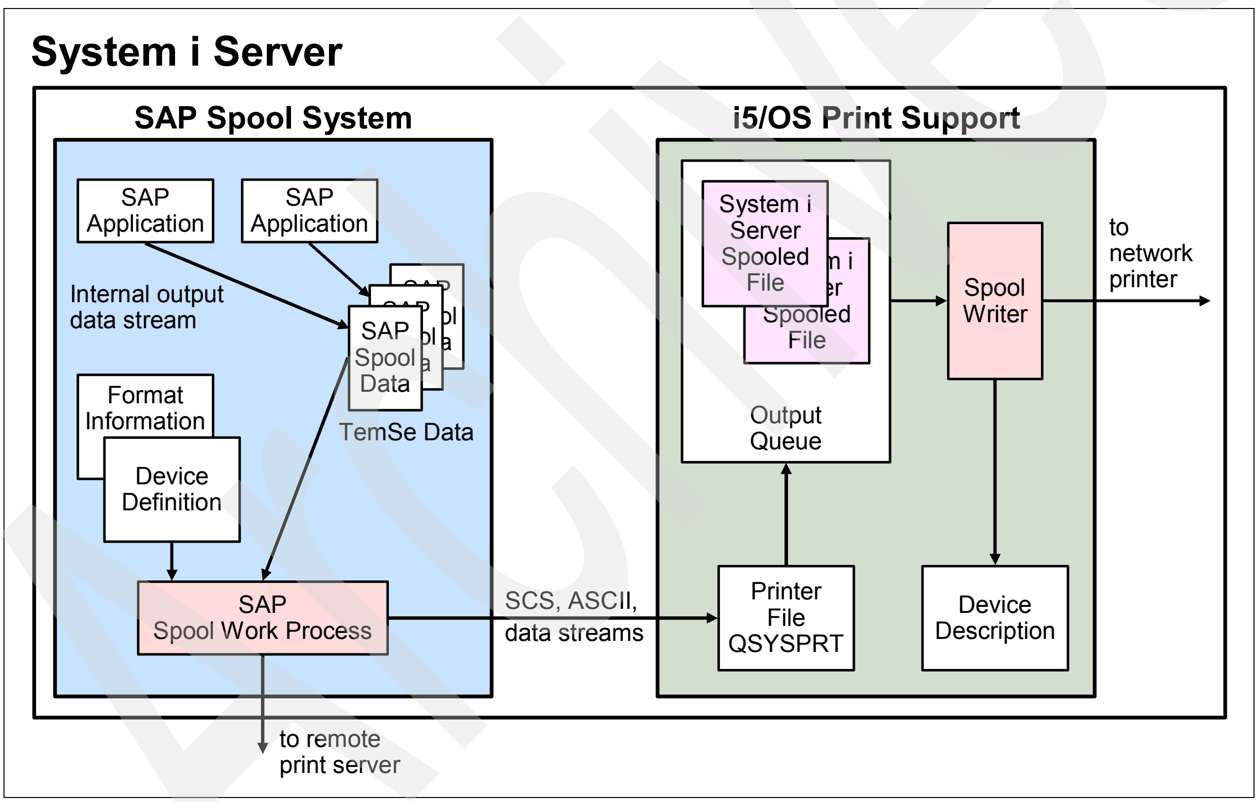

Figure 15-1 Flow of SAP output data on a System i server

Printed output from SAP applications can consist of a simple list or complex documents that require advanced printer functions. There are several ways to request output to be printed from SAP applications. The most common ways are to:

- Click a print button on an SAP GUI display, which causes a simple list to be printed
- Request the SAP programming editor to print an ABAP source file
- Run an ABAP report
- Produce a document from SAPscript

In any case, the SAP system does not actually print the data. The SAP system presents the final output data stream to the i5/OS operating system which then is responsible for its final printing. See 15.2.1, "Overview of i5/OS spool support" on page 216 for an overview of i5/OS spooling.

The SAP spool system requires a definition of print resources and capabilities. Making a printer available to an SAP application is a two-step process, as follows:

1. Configure the printer on the System i server using the normal System i configuration support. This makes the printer known to the System i server, but not to the SAP system.
2. Using the support provided by the SAP system, add a new printer to the set that is known by the SAP system. The printer is then available for use by SAP applications (using transaction SPAD). So with transaction SPAD you define an SAP device type (general spoken: a "printer") with its attributes.

All printing methods supported by i5/OS are available to an SAP application, including SNA Character Stream (SCS), Intelligent Printer Data Stream™ (IPDS™) and American Standard Code for Information Interchange (ASCII). The printers can be:

- Direct attached
- LAN attached
- Remote output queues

A detailed description of how printers are defined in an SAP system is described in 15.6, "Configure new SAP system device" on page 244. For a description of the different printer models and attachment methods supported by the System i server, refer to *IBM AS/400 Printing V*, SG24-2160. For more information about Advanced Function Printing (AFP) data streams, refer to 15.4.2, "AFP printing" on page 222.

The SAP spool system provides its own mechanisms to define the proper format of print data. The SAP spool system does not recognize or use data description specifications (DDS), externally described data, or printer files.

While it is in an internal SAP spooled file, output data is encoded in a character set that is not necessarily the same as the one supported by the printer. SAP 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 15.5.7, "SAP characters and character sets" on page 243.

SAP spooled files are managed using the SAP facilities, not from i5/OS facilities. With SAP system features, you can:

- Display a list of outstanding spooled requests
- Display the content of a spooled file
- Delete spooled requests
- Print spooled requests

Printing a spooled file causes the SAP spool system to transform the data and give it to the System i server for actual printing. The data enters an output queue once the data is received by i5/OS. A printer prints the data as directed by the spool writer.

Note: The i5/OS-supplied QSYSPRT system printer file is used by SAP applications on System i servers. The SAP application overwrites this printer file to match the printer customizing within the SAP system.

Managing SAP spooled files is discussed in 15.8, "Manage SAP spooled and output requests" on page 247.

15.2 Spool support architecture

This section describes the SAP spool system and includes an overview of i5/OS spooling support in order to understand the architecture and information flow in the SAP spool system. Also discussed is the spool work process and where you can find the SAP spool data.

15.2.1 Overview of i5/OS spool support

When a spool request is printed in the SAP system, the spool work process considers the device type and access method of the output device that is to process the request. When an output device is configured for access method C in transaction SPAD, the name of an output queue (object type: *OUTQ) on the System i server is specified. Based on a QSYSPRT (printer file object type: *FILE, file attribute: PRTF), the spool work process creates a spooled file in that output queue. If a writer job is started for that output queue, the spooled file is processed and sent to the attached printer. The printer can be somewhere in the network, as shown in Figure 15-1 on page 214, or it can be physically attached to the System i server. If the printer is locally attached to the System i server, you see a printer writer job instead of a remote writer job.

Refer to 15.4, "Printing with i5/OS systems" on page 221 for further information.

15.2.2 Overview of the SAP spool system

When a request is made to print a document in an SAP system, a spool request is created. It consists of administrative information (origin, date, author name, logical printer), which is stored in the SAP database, and the data to be printed, which is stored in a repository called the temporary sequential database (TemSe). The SAP spool system uses generic representations of printer formatting commands and the SAP internal character set to represent the characters to be printed.

The preparation of a spool request for printing is an output request. In an SAP system, you can either print immediately (an output request is generated immediately) or delay printing (the spool request is not immediately sent as an output request). A spool request can correspond to several output requests which can be sent to different output devices.

The access method defines how the printer is accessed. Commonly used access methods are:

C	Direct operating system call
F	Printing on Front End Computer
S	Print Using SAP Protocol
U	Print Using Berkeley Protocol
L	Print Locally Using LP/LPR

Access method L is no longer supported with i5/OS, as described in SAP note 693802.

Figure 15-2 on page 217 illustrates the SAP printing flow.

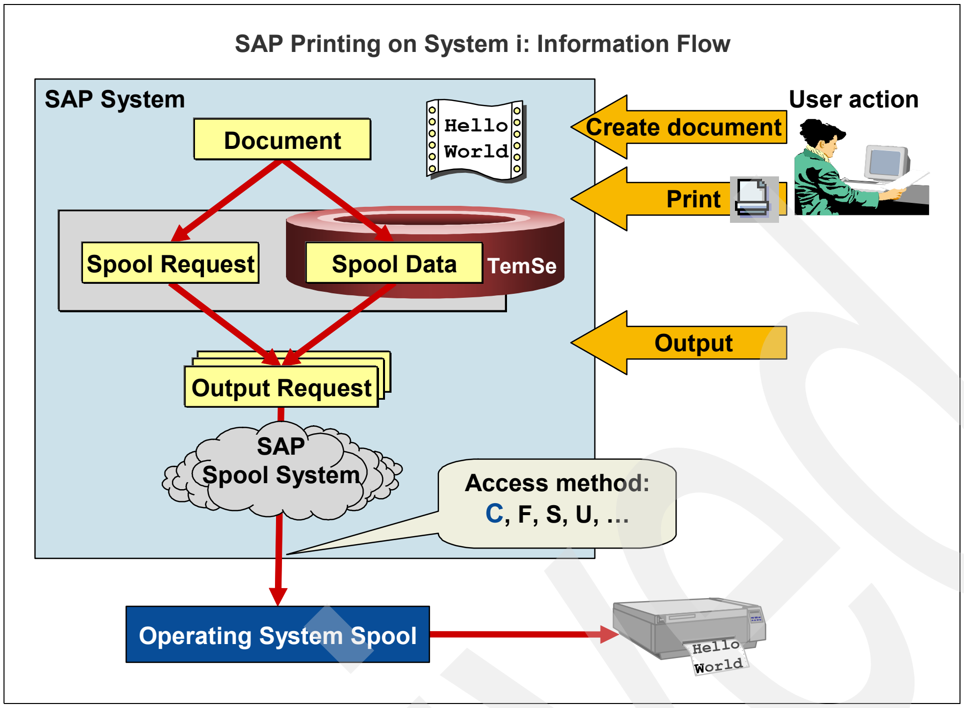

Figure 15-2 SAP printing with i5/OS: Information flow

15.2.3 TemSe data

The set of spooled data held internal to the SAP system 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 with i5/OS configurations: in the SAP database or in the Integrated File System. 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.

Depending on which system the spool data is stored, you might want to run the spool work process on the same machine.

For more information, refer to SAP note 20176.

15.2.4 The spool work process

An SAP 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 i5/OS spooled file on the appropriate output queue or an external print server. Whether there is a spool work process or not 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 defined for this instance. There can be multiple processes. Use transaction SM51 to see which instances have spool work processes. If a spool work process is not running, data cannot be printed, but SAP applications can still place data into the TemSe data. See 16.1.3, “Monitoring an SAP server: Transaction SM51” on page 256 for a further discussion of SM51.

Note: Beginning with SAP Release 4.0A, it is possible to have more than one spool work process within an SAP instance. For more information, refer to SAP note 107899.

In a 2-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 3-tier configuration, there are two options for placing the spool work process:

- On the database server
- On the application server

Each approach is discussed in this section.

Place the spool work process on the database server

The advantage of placing the spool work process on the database server is that spool work processes have local access to the TemSe data. Because it is local, the access to the TemSe data is fast due to the reduced traffic in 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 i5/OS spool support when the data is finally printed. Moreover, if the printed material is needed at the application server, the i5/OS spool support might 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 15-3 illustrates the placement of the spool work processes on the database server.

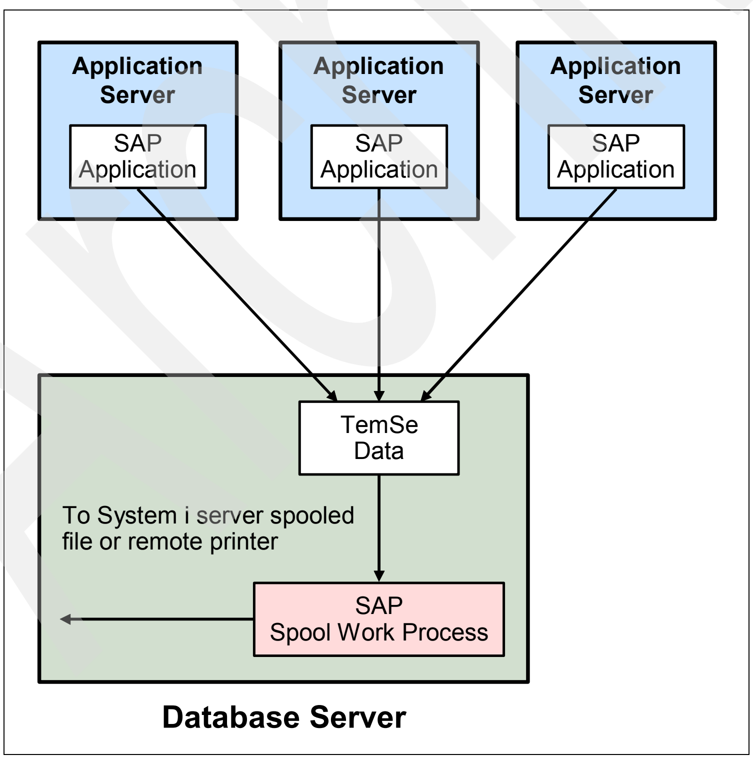

Figure 15-3 Spool work process on the database server

Place 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 SAP applications, 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 Integrated File System, it is the system that contains '/usr/sap/<SID>/global'. If the TemSe data must flow two ways, the SAP 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 these types of connection, there should be no performance problem due to network traffic. If the TemSe data is kept in the Integrated File System file, the data flow is across a TCP/IP connection where performance might be a consideration (unless a high speed connection is used).

Figure 15-4 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.

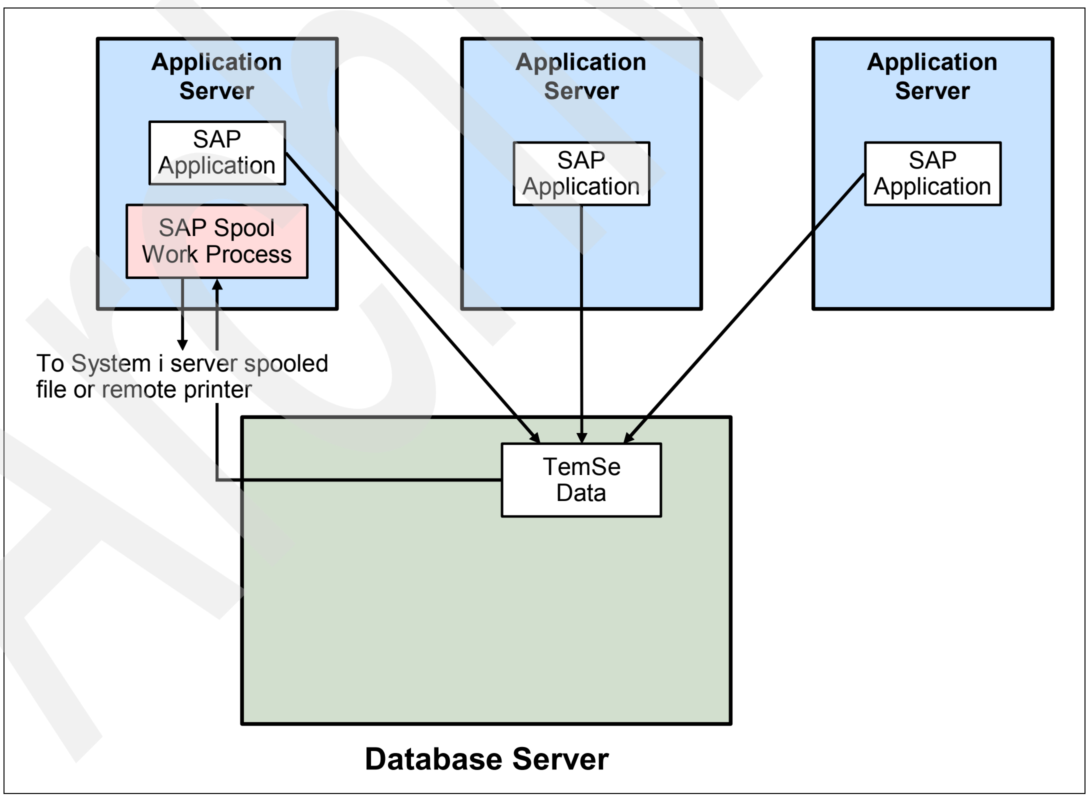

Figure 15-4 Spool work process on the application server

15.3 Spool request and 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 15-5 shows the print flow in an SAP system.

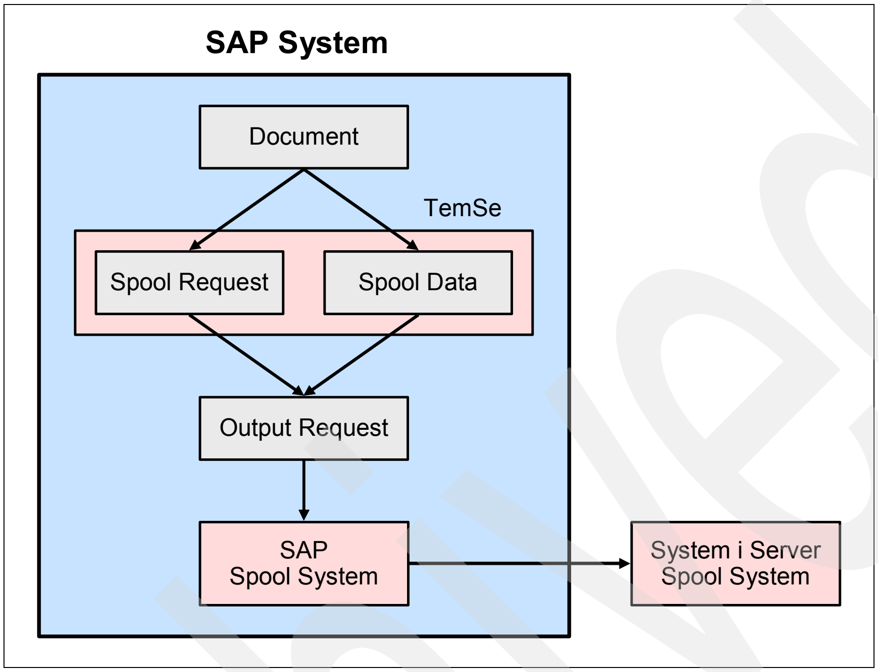

Figure 15-5 SAP spool request and output request

When a document is printed, the SAP system generates a spool request and places the spool data stored in the TemSe. The spool request and the spool data make up a output request. One spool request can generate several output requests. Specify "print immediately" to combine the two-step process into one.

> **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 System i server if the operating system spooler is involved.

A spool request has two parts:

1. Administrative information stored in the SAP database (for example, the origin, date, author name, and name of the logical printer).
2. The data to be printed is stored in a repository called *TemSe data*. The SAP spool system uses generic representations of printer formatting commands and SAP 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 was 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.

15.4 Printing with i5/OS systems

To understand printing on i5/OS systems we first describe the supported data streams. Then we explain how to manage printers and output queues with i5/OS.

15.4.1 i5/OS printer data streams

System i models support three EBCDIC-based printer data streams:

- SCS is a simple, text based data stream that allows only a few formatting options such as bold, underline or characters per inch (CPI) values.
- IPDS allows including graphics or bar codes and selecting different fonts from the printer.
- AFP is the most advanced data stream and allows downloading software fonts and including predefined overlays or form definitions.

These data streams are downward compatible so that an AFP printer can also print SCS data stream.

Traditional applications can also print on ASCII-based printers by using the Host Print Transform functionality. However, the number of supported data streams is limited, and the conversion of the data stream needs resources (CPU and memory) on the System i model. So direct-attached ASCII-based printer are not so popular but rather as network printer attached to a Linux or Windows or another printer server.

Figure 15-6 illustrates the different printing data streams.

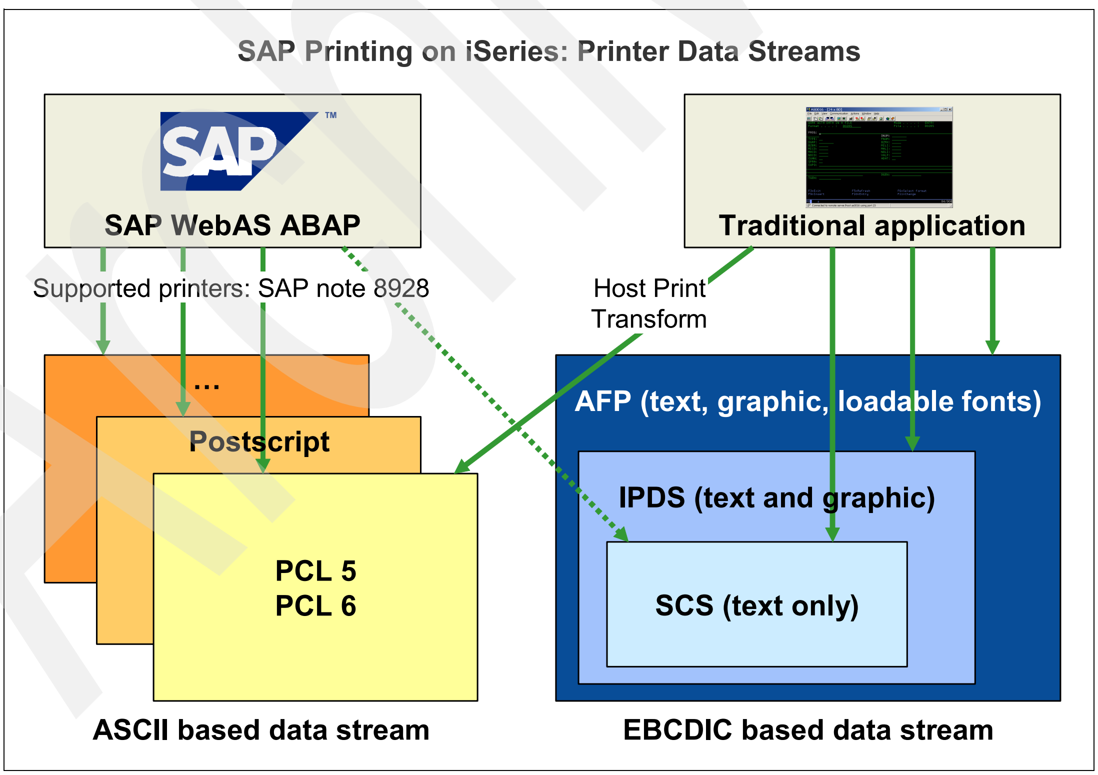

Figure 15-6 EBCDIC and ASCII based data stream

The SAP WebAS ABAP server supports a wide variety of ASCII based data streams. The list of supported printers or data streams can be found in SAP note 8928. This note is updated as new printers are introduced to the market and supported by SAP. You can also create SCS data stream (with its restrictions) by using the device type IBMSCS.

15.4.2 AFP printing

AFP (Advanced Function Presentation™) is the standard page output format on the System i platform. This section covers AFP printing considerations from SAP.

AFP printing from SAP was supported by the SAP AFP module in AFP PrintSuite (5798-AF2 and 5798-AF3) This product is not supported by IBM as of September 30, 2004. This affects users with IBM AFP printers, or users with double byte SCS printers.

AFP PrintSuite consisted of following components:

- Advanced Print Utility (APU) for iSeries: An unsupported version of this product can be downloaded
- Page Printer Formatting Aid (PPFA) for iSeries: This product is replaced by Infoprint® Designer for iSeries
- AFP Toolbox for OS/400: An unsupported version of this product can be downloaded
- SAP R/3 Print SP

As consequence of this withdrawal, printing with SAP access method 'L' will no longer be supported by SAP on the System i platform. However, access method 'L' will still work if you have previously installed the IBM AFP PrintSuite SAP module. Note that new SAP printing functionality (for example, Smartforms) may be limited or not supported.

Access Method 'Z', a hardcoded implementation of Access Method 'L', has not been supported by SAP since release 4.0B. It will continue to function as long as the IBM AFP PrintSuite SAP module is installed on the iSeries.

Printing to double byte character set (DBCS) SCS printers is unsupported as it also used the IBM AFP PrintSuite SAP module.

For detailed information see SAP note 693802.

There are several alternatives if IBM PrintSuite SAP support is no longer being used:

- If your printer hardware supports ASCII-based data streams such as PostScript or PCL-5, it may be possible for you to use Access Methods C or U.
- If you require a non-ASCII data stream and your AFP print images do not contain complex graphics (such as letterheads or barcodes), try using Access Method C with SAP printer driver “SAPSCS”. Note the much of the advanced functionality provided with AFP would not be available with this approach.
- There are some new composition and printing tools that provide various levels of integration with SAP output data. See the next section for more information.

Tools and connectors for AFP printing

There are several approaches to support AFP printing from SAP output.

Mapping

Mapping is a document composition system that can import data from a number of different sources, redesign the data into documents, and produce various print data streams (AFP included). Mapping could be used in conjunction with SAP applications in a couple of ways:

1. Retrieving txt (scs) spool files from SAP and use Mapping to layout the pages again and produce the print data stream.
2. Use the MapReport module to read in SAP data in RDI or XML format and use this imported data for format documents and produce the required print data stream.

StreamServe

StreamServe is business communication middleware (IBM product 5639-DD9) that provides an intelligent connection between application data and business documents. StreamServe has high-function, certified connectors to SAP data. Any SAP output format (SAPGOF, RDI, XSF, ABAP) can be enhanced and distributed via both print and other distribution channels (i.e., Web, e-mail, fax).

StreamServe can reside on an in-board Windows or AIX partition and connect to SAP output data. StreamServe could also be installed on an outboard Windows or AIX server and connect to SAP iSeries data.

For more information, access this Web site and select "Software":

http://www.ibm.com/printers

SAP2AFP converter

SAP note 65224 describes the handling with the converter SAP2AFP:

SAP2AFP is a transform program from IBM that converts SAP output to AFP format. In the SAP system, print data is passed on to the external converter via so-called "Spool Exits" (also called "Access method Z"). The device types IBMAFP and IBMAFP3 are available in SAP for this.

If there are printing problems that cannot be traced back to incorrect print data in the SAP application spool (for example, incorrect OTF syntax, incorrect list format), SAP cannot assist further with these problems.

The format spool exit ("Access method Z and device types IBMAFP, and so on) is no longer supported as of Release 4.0A and is replaced by the SAPGOF format (device type SAPGOF). In principle, SAPGOF delivers the same information as the old format and is officially documented. You find the SAPGOF documentation under

http://sapnet.sap.com/output

Important note: In the meantime, IBM is providing a converter solution based on the SAPGOF data format for printing SAP documents on AFP printers.

SAP2AFP is an IBM product. Contact IBM for support/error correction in concerning the product SAP2AFP.

15.4.3 i5/OS printer commands

The following list contains commands that are commonly used on the System i models for managing printers and printed output.

- Create Printer Device Description (CRTDEVPRT)

  This command creates a device description of the printer. You can indicate how the printer is attached, the type of printer to be attached, the model number, whether the device is varied on at IPL, and other optional parameters that depend on the specific models chosen.

- Create Output Queue (CRTOUTQ)

  An output queue is where spooled files are placed prior to actually being printed. With the CRTOUTQ command, you can order the spooled file entries on an output queue and identify that any spooled files that come to this output queue should be sent to a remote system for actual printing. Other selections are available depending on the parameters and values specified.

- 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 (STRRMTWTR)

  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 is 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. To see all writers, use the command:

  ```
  WRKWTR WTR(*ALL)
  ```

  WRKWTR 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

- 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.

Note: Use of each command requires proper authority.

For a more complete description of these printer commands, refer to the *Printer Device Programming*, SC41-5713, and to the CL Reference in the *iSeries Information Center.*

15.4.4 i5/OS printer configuration

The following printer configurations are supported by i5/OS:

- IPDS printers

  Refer to 9.10 AFP printing with SAP and 9.11.1 LAN-attached IPS printers in the *IBM Redbook SG24-4672 Implementing SAP on OS/400,* for information about how to configure and attach IPDS printers.

- ASCII printers

  i5/OS can route printed output to LAN ASCII printers using TCP/IP. This includes IBM printers, as well as non-IBM ASCII printers with an appropriate network attachment.

  There are two methods to print directly to these printers:

  – *Send TCP/IP spooled file*: This is the i5/OS 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 command, the print file is simply sent. Print is not managed by i5/OS.

The transmission has the following characteristics:

- The entire file is sent.
- Some printers ignore the control file that is sent along, losing job control.
- It is a one-way transmission. There is no control of the data stream, no comfortable feedback of the print status and no good error recovery.
- The entire spooled file is resent when the transmission errors.
- Most printer i5/OS file parameters are inoperable.
- Cancelling the printing yields unpredictable results.

– *PJL driver*: The Printer Job Language (PJL) driver increases the network printing function beyond what is provided by the Send TCP/IP spooled file commands. With the PJL driver, a printer device description is created, unlike the case with the Send TCP/IP spooled file command. This facilitates a dialog between i5/OS 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.

Prior to configuring the printer:

- The TCP/IP network must already be set up on the System i server.
- Install the latest cumulative PTFs on your system to help avoid problems.

Create a printer device description

The printer must support PJL and Printer Control Language or Printer Command Language (PCL) to be LAN-attached with a PJL driver.

Complete the following steps to configure a LAN-attached ASCII printer with PJL drivers:

1. To create the device description for your printer, type the Create Device Description Printer (CRTDEVPRT) command on any command line and press F4 (Prompt). A screen appears, as shown in Figure 15-7 on page 226 through Figure 15-14 on page 229:

```
                     Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Device description . . . . . . . DEVD         > NPLAN
 Device class . . . . . . . . . . DEVCLS       > *LAN
 Device type  . . . . . . . . . . TYPE         > 3812
 Device model . . . . . . . . . . MODEL        > 1
 LAN attachment . . . . . . . . . LANATTACH    > *IP
 Switched line list . . . . . . . SWTLINLST
                           + for more values
 LAN remote adapter address . . . ADPTADR
 Adapter type . . . . . . . . . . ADPTTYPE       *INTERNAL
 Adapter connection type  . . . . ADPTCNNTYP     *PARALLEL
 Emulated twinaxial device  . . . EMLDEV
 Advanced function printing . . . AFP            *NO
 AFP attachment . . . . . . . . . AFPATTACH      *WSC
 Port number  . . . . . . . . . . PORT         > 2501
 Switch setting . . . . . . . . . SWTSET
 Local location address . . . . . LOCADR
                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 15-7   CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 1 of 8)

```
                     Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Auxiliary printer  . . . . . . . AUXPRT
 Emulating ASCII device . . . . . EMLASCII       *NO
 Physical attachment  . . . . . . ATTACH
 Online at IPL  . . . . . . . . . ONLINE       > *YES
 Attached controller  . . . . . . CTL
 Language type  . . . . . . . . . LNGTYPE        *SYSVAL
 Print quality  . . . . . . . . . PRTQLTY        *STD
 Font:                            FONT
   Identifier . . . . . . . . . .              > 11
   Point size . . . . . . . . . .                *NONE
 Form feed  . . . . . . . . . . . FORMFEED     > *AUTOCUT
 Separator drawer . . . . . . . . SEPDRAWER      *FILE
 Separator program  . . . . . . . SEPPGM         *NONE
   Library  . . . . . . . . . . .
 Number of drawers  . . . . . . . NBRDRAWER
 Printer error message  . . . . . PRTERRMSG      *INQ
                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 15-8   CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 2 of 8)

```
                   Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Message queue  . . . . . . . . . MSGQ         > *SYSOPR
   Library  . . . . . . . . . . .
 Maximum length of request unit   MAXLENRU       *CALC
 Application type . . . . . . . . APPTYPE        *NONE
 Activation timer . . . . . . . . ACTTMR       > 170
 Inactivity timer . . . . . . . . INACTTMR       *ATTACH
 SNA pass-through device desc . . SNPTDEV        *NONE
 SNA pass-through group name  . . SNPTGRP        *NONE
 Host signon/logon command  . . . LOGON          *NONE



 Pacing . . . . . . . . . . . . . PACING         7
 Line speed . . . . . . . . . . . LINESPEED      *TYPE
 Word length  . . . . . . . . . . WORDLEN        *TYPE
 Type of parity . . . . . . . . . PARITY         *NONE
                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 15-9 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 3 of 8)

```
                   Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Stop bits  . . . . . . . . . . . STOPBITS       1
 Host print transform . . . . . . TRANSFORM    > *YES
 Manufacturer type and model  . . MFRTYPMDL    > *IBM4317
 Paper source 1 . . . . . . . . . PPRSRC1      > *A4
 Paper source 2 . . . . . . . . . PPRSRC2      > *A4
 Envelope source  . . . . . . . . ENVELOPE     > *C5
 ASCII code page 899 support  . . ASCII899       *NO
 Image configuration  . . . . . . IMGCFG       > *IMGA02
 Maximum pending requests . . . . MAXPNDRQS      6
 Print while converting . . . . . PRTCVT         *YES
 Print request timer  . . . . . . PRTRQSTMR      *NOMAX
 Form definition  . . . . . . . . FORMDF         F1C10110
   Library  . . . . . . . . . . .                  *LIBL
 Character identifier:            CHRID
   Graphic character set  . . . .                *SYSVAL
   Code page  . . . . . . . . . .
                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 15-10 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 4 of 8)

```
                     Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Remote location:                 RMTLOCNAME
   Name or address  . . . . . . .              > '9.9.99.99'



 Local location . . . . . . . . . LCLLOCNAME     *NETATR
 Mode . . . . . . . . . . . . . . MODE           QSPWTR
 DBCS feature:                    IGCFEAT
   Device features  . . . . . . .
   Last code point  . . . . . . .
 User-defined options . . . . . . USRDFNOPT      *NONE
                           + for more values
 User-defined object:             USRDFNOBJ
   Object . . . . . . . . . . . .                *NONE
     Library  . . . . . . . . . .
   Object type  . . . . . . . . .
                                                                       More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 15-11 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 5 of 8)

```
                     Create Device Desc (Printer) (CRTDEVPRT)

 Type choices, press Enter.

 Data transform program . . . . . USRDTATFM      *NONE
   Library  . . . . . . . . . . .
 User-defined driver program  . . USRDRVPGM      *NONE
   Library  . . . . . . . . . . .
 System driver program  . . . . . SYSDRVPGM    > *IBMPJLDRV
 Secure connection  . . . . . . . SECURECNN      *NO
 Validation list  . . . . . . . . VLDL           *NONE
   Library  . . . . . . . . . . .









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

Figure 15-12 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 6 of 8)

```
                   Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Publishing information:         PUBLISHINF
  Support duplex . . . . . . . .                *UNKNOWN
  Support color  . . . . . . . .                *UNKNOWN
  Pages per minute black . . . .                *UNKNOWN
  Pages per minute color . . . .                *UNKNOWN
  Location . . . . . . . . . . .                *BLANK

  Data streams supported . . . .                *UNKNOWN
                          + for more values
Text 'description' . . . . . . . TEXT        > 'DEVD for NPLAN'

Dependent location name  . . . . DEPLOCNAME     *NONE

                           Additional Parameters

Remote network identifier  . . . RMTNETID       *NETATR
                                                                       More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 15-13 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 7 of 8)

```
                   Create Device Desc (Printer) (CRTDEVPRT)

Type choices, press Enter.

Workstation customizing object   WSCST          *NONE
  Library  . . . . . . . . . . .
Authority  . . . . . . . . . . . AUT            *CHANGE

























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

Figure 15-14 CRTDEVPRT for LAN-attached ASCII printer using the PJL driver (Part 8 of 8)

Create a remote output queue

Figure 15-15 on page 230 through Figure 15-18 on page 231 show the parameters for creating a remote output queue:

```
                         Create Output Queue (CRTOUTQ)

Type choices, press Enter.

Output queue . . . . . . . . . . . OUTQ         > PPA6
  Library  . . . . . . . . . . .                >   QUSRLIBL
Maximum spooled file size:         MAXPAGES
  Number of pages  . . . . . . .                  *NONE
  Starting time  . . . . . . . .
  Ending time  . . . . . . . . .
                          + for more values
Order of files on queue  . . . . SEQ              *FIFO
Remote system  . . . . . . . . . RMTSYS         > DEWDFPRO5



Remote printer queue . . . . . . RMTPRTQ        > *PPA6



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

Figure 15-15 CRTOUTQ command (part 1 of 4)

```
                         Create Output Queue (CRTOUTQ)

Type choices, press Enter.

Writers to autostart . . . . . . AUTOSTRWTR   > 1
Queue for writer messages  . . . MSGQ           QSYSOPR
  Library  . . . . . . . . . . .                  *LIBL
Connection type  . . . . . . . . CNNTYPE      > *IP
Destination type . . . . . . . . DESTTYPE     > *OTHER
Host print transform . . . . . . TRANSFORM      *YES
User data transform  . . . . . . USRDTATFM      *NONE
  Library  . . . . . . . . . . .
Manufacturer type and model  . . MFRTYPMDL    > *INFOPRINT1352
Workstation customizing object   WSCST          *NONE
  Library  . . . . . . . . . . .
Image configuration  . . . . . . IMGCFG         *NONE
Internet address . . . . . . . . INTNETADR
VM/MVS class . . . . . . . . . . CLASS          A
Forms Control Buffer . . . . . . FCB            *NONE

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

Figure 15-16 CRTOUTQ command (part 2 of 4)

```
                          Create Output Queue (CRTOUTQ)

Type choices, press Enter.

Destination options  . . . . . . DESTOPT        *NONE


Print separator page . . . . . . SEPPAGE        *YES
User defined option  . . . . . . USRDFNOPT      *NONE
                          + for more values
User defined object:             USRDFNOBJ
  Object . . . . . . . . . . . .                *NONE
    Library  . . . . . . . . . .
  Object type  . . . . . . . . .
User driver program  . . . . . . USRDRVPGM      *NONE
  Library  . . . . . . . . . . .
Spooled file ASP . . . . . . . . SPLFASP        *SYSTEM
Text 'description' . . . . . . . TEXT           *BLANK



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

Figure 15-17 CRTOUTQ command (part 3 of 4)

```
                          Create Output Queue (CRTOUTQ)

Type choices, press Enter.


                            Additional Parameters

Display any file . . . . . . . . DSPDTA         *NO
Job separators . . . . . . . . . JOBSEP         0
Operator controlled  . . . . . . OPRCTL         *YES
Data queue . . . . . . . . . . . DTAQ           *NONE
  Library  . . . . . . . . . . .
Authority to check . . . . . . . AUTCHK         *OWNER
Authority  . . . . . . . . . . . AUT            *USE







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

Figure 15-18 CRTOUTQ command (part 4 of 4)

Important parameters include:

- Output queue and library: Define the name of your output queue. Specify a library that is in the library list of all jobs on the system, such as QGPL.

- Remote system and Remote printer queue: Remote system is the name of a network printer or a print server. The default *NONE is used for a local printer or an output queue that has no printer assigned. Specify the printer name on the print server in the Remote printer queue parameter.
- Writers to autostart: If the output queue actually represents a printer, it is recommended to specify a writer job to be started automatically at system startup.
- Connection type and Destination type: Specify *IP and *OTHER, as shown in Figure 15-16 on page 230.
- Host print transform: Leave the default *YES to allow EBCDIC-based spooled file data to be printed on an ASCII based printer.
- Manufacturer type and model: This parameter is used when Host print transform is set to yes. Function key F4 provides a list of supported printer models.
- Workstation customizing object and library: If the printer type and model is not found in the list of supported types and models, create your own conversion object using the command CRTWSCST.
- Print separator page: The default for this parameter is set to *YES. Change to *NO to allow SAP create the separator page.
- Creating a remote output queue is probably the easiest way to configure a network printer for the System i server. In Figure 15-19, two modifications are described that can be useful in some cases.

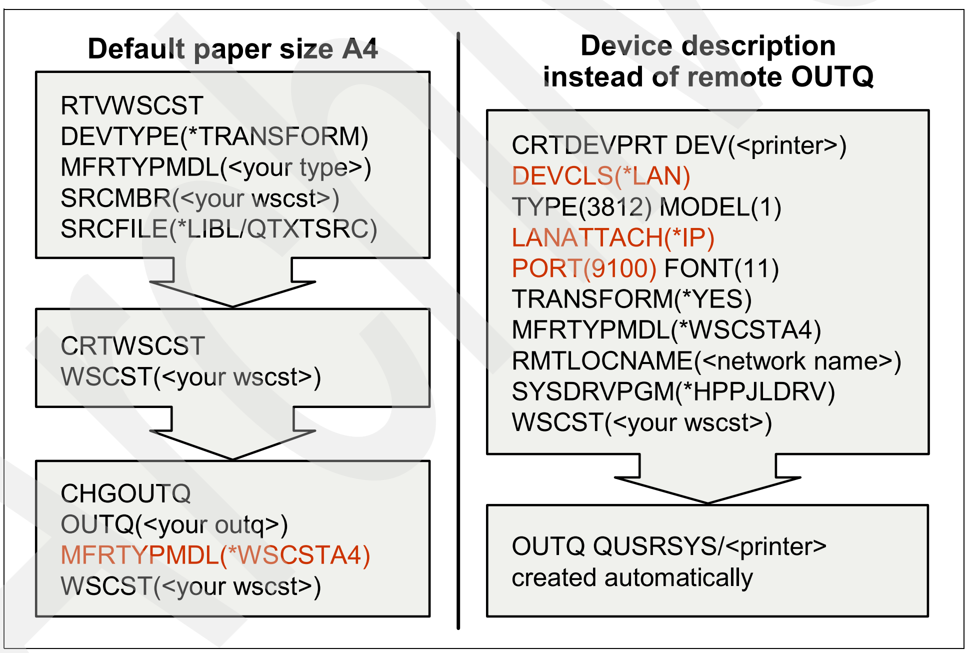

Figure 15-19 Create a WSTCST or a device description

Per default, the i5/OS printer drivers assume that your printer has letter format paper loaded. If you print to a printer that has A4 size paper loaded, you can see a message at the printer telling you to load letter size paper. You can ignore this message, but it is annoying.

To solve this problem, you can create your own workstation customizing object (type *WSCST). Specify *WSCSTA4 in the “manufacturer type and model“ (MFRTYPMDL) parameter as shown in the left part of the slide. First you need to retrieve the source of an existing type, such as *INFOPRINT1352, to a member in a source file QTXTSRC. Then compile the source without changes to create an object of type *WSCST. Specify this object in your remote output queue and set the MFRTYPMDL parameter to *WSCSTA4.

- Instead of a remote output queue, you can also create a device description as shown in the right part of Figure 15-19 on page 232. The advantage is a better feedback handling in case of errors.

  When printing native i5/OS data, you can also specify a page range to print, for example to restart your printout at a certain page after a paper jam. The slide shows an example configuration. Depending on your printer type you might have to use different port numbers (parameter PORT) or system driver program (parameter SYSDRVPGM). When you create a printer device description, an output queue with the same name is created automatically in library QUSRSYS.

Work with an i5/OS output queue

You can use the i5/OS WRKOUTQ command to see spooled file entries in an output queue. An example is shown in Figure 15-20:

```
                              Work with Output Queue

Queue:   PPA6           Library:   QUSRSYS        Status:   RLS/WTR

Type options, press Enter.
  1=Send   2=Change   3=Hold   4=Delete   5=Display   6=Release   7=Messages
  8=Attributes        9=Work with printing status

Opt  File        User        User Data   Sts    Pages   Copies  Form Type   Pty
     QSYSPRT     PO1OWNER    0000011766  HLD       1*      1    DINA4        5
     QSYSPRT     PO1OWNER    0000011766  HLD       1*      1    DINA4        5
     QSYSPRT     PO1OWNER    0000011766  HLD       1*      1    DINA4        5
     QSYSPRT     PO1OWNER    0000011921  HLD       1*      1    DINA4        5







                                                                         Bottom
Parameters for options 1, 2, 3 or command
===>
F3=Exit   F11=View 2   F12=Cancel   F20=Writers   F22=Printers
F24=More keys
```

Figure 15-20 The i5/OS WRKOUTQ command

The name of the output queue is specified as "host printer name" when selecting access method C. (See 15.2.2, "Overview of the SAP spool system" on page 216 to understand method C.) In order to see some entries in the OUTQ, the output queue is held with the command HLDOUTQ. Use the RLSOUTQ command to release the output queue. The status is displayed in the upper right part of the screen. WTR indicates that a writer job is started for that output queue. The user data for each spooled file in the output queue corresponds with the spool request number in the SAP system, so you can map SAP spool requests with entries in the output queue.

Options to process the spooled files include hold (option 3), delete (option 4), and display the attributes (option 8). Option 5 (display) works only when the data stream is EBCDIC-based, which is true with a IBMSCS device type is specified in the SAP system.

If the print requests on a printer with access method C are completed in the SAP system without error but do not appear on your printer, review the WRKOUTQ command display. Verify in the upper right part of the screen that a writer job is started and that the output queue is released.

15.5 Administration of the SAP spool system

This section explains how to navigate to the correct window and function in order to create a new spool administration task. It also describes the main components of a printer definition and how to administer the SAP spool system, such as how to setup:

- Output devices
- Device types
- Format types
- Print controls
- Character sets

Use the SPAD transaction to define resources to the SAP spool system. This transaction takes you to the spool administration window.

15.5.1 Components of a printer definition

Figure 15-21 illustrates the components that make up a printer's definition.

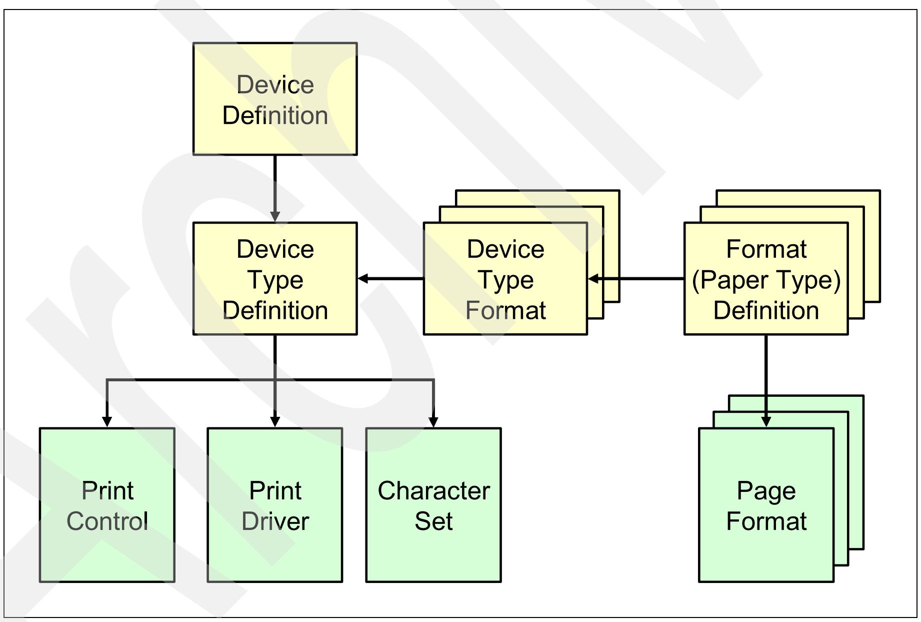

Figure 15-21 Components of a printer definition

An SAP 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, the device type is linked to three components, which include:

- *Print driver*: This function is used to format the final printer output stream. Print drivers are predefined by the SAP system. Choose from the available drivers.
- *Character set*: Select the character set supported by the printer.
- *Print controls*: This parameter maps the generic printer controls of the SAP system to the specific control characters of the printer.

A device type and its components do not provide information about how data is to be formatted for the printer. This information is provided by a format (the type of paper). A format is linked to 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 which specifies printer initialization information. Since a device type is capable of processing more than one data format, there can 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.

15.5.2 Output devices

To define a new printer to the SAP spool system, complete the following steps:

1. Select the Output Devices category from the SPAD transaction. The next window that appears lists all of the current output devices that are already defined to the SAP system.
2. Click the Create or the Create using template button to define a new printer device.
3. Fill in the following parameters on the next window:
   - *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 the SAP system uses to access the printer.
   - *Device type*: This parameter is the SAP method of identifying the model of the device. The SAP spool system provides a number of predefined device types from which you can 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. Click in the input field and the list box that appears to view the list of device types already defined. 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 can create a new one. Device types are discussed in 15.5.4, “SAP device types” on page 239.
   - *Spool server*: This parameter identifies the location of the spool work process that processes any data that is directed to this printer. A list of names appears that identifies where the various spool work processes are located. Click in the input field and then in the list box that appears to determine which spool work processes are available. 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 SAP system ID, and the third is the instance number. The name is the same as 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 the SAP system. When locked from the SAP system, 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 the SAP system, the device can still be used outside of that environment.

– *Host spool access method*: This parameter tells the spool work process what to do with the final output data stream. To send the final output to an i5/OS spooled file, choose one of the following possible values:

  - **C**: This option allows the use of output queues. The SAP system sends the final output data stream to a spooled file. This is mainly used for SCS data stream (SAP device type ZIBMSCS) or ASCII data sent to a remote printer via a remote writer (data stream in the spooled file *USERASCII). You can use either locally attached printers or remote output queues.
  - **F**: Front-end printing allows you to print on locally attached front-end PCs. The i5/OS spool system is not involved.
  - **Z**: The output is placed in an Integrated File System stream file, and then the Convert Printer Data (CVTPRTDTA) command (belonging to the PrintSuite, option 4) is run to print AFPDS. The access method Z is no longer supported, but can still be used. Access method L supersedes Z.

    Access method Z is no longer supported in i5/OS.
  - **L**: This method is more flexible than access method Z. L also uses the CVTPRTDTA command.

    Access method L is no longer supported in i5/OS.
  - **U**: The output is sent via TCP/IP directly to the printer using the Berkeley LPD protocol without going through an output queue. Since you can setup more than one spool work process there are less or no problems with locked spool work process for quite a while. The i5/OS spool system is not involved in this case. Access method U is now more than an alternative to access method C.

    This method is not recommended because problems with the printer can lock up the spool work process for quite a while. The i5/OS spool system is not involved in this case.
  - **S**: The S value causes the SAP system to send the final output to another host rather than putting it into a spooled file. The protocol used is SAP's private protocol. The receiving host must run SAP's transfer program. This method sends print data to a host in a network using TCP/IP that does not support the Berkeley LPD protocol. The i5/OS spool system is not involved.
  - **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 i5/OS spool system is not involved.

– *Host printer*: The name of the printer at operating system level. If an i5/OS output queue is used to receive the data that is produced by the spool work process, enter the name of the output queue. The SAP system does not create output queues. Ensure the output queue exists prior to any attempt by the SAP system to use it.

  For front-end printing (access method F), specify __DEFAULT to access the default printer on the front-end PC.

  Enter the name of the printer on the remote host. This is LPT1 if SAPLPD is used.

– *Destination host*: This field appears after pressing Enter if the U or S access method is used. 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 the SAP system to query the output status from the host spooler. This parameter is valid for the Z and L methods.

15.5.3 Configure Access Method C

Access method C uses the i5/OS specific spooling system. Access method U sends the print data directly to the printer. C and U are the most common method way to define an SAP printer.

An output queue (object of type *OUTQ) is required on the System i server when configuring a printer through access method C. You can see in Figure 15-15 on page 230 until Figure 15-18 on page 231, how that output queue is created by using the i5/OS CRTOUTQ command.

In Figure 15-22 the name of the output queue is specified as host printer. The other configuration elements are the same as for other access methods and platforms.

Figure 15-22 illustrates how to configure Access Method C.

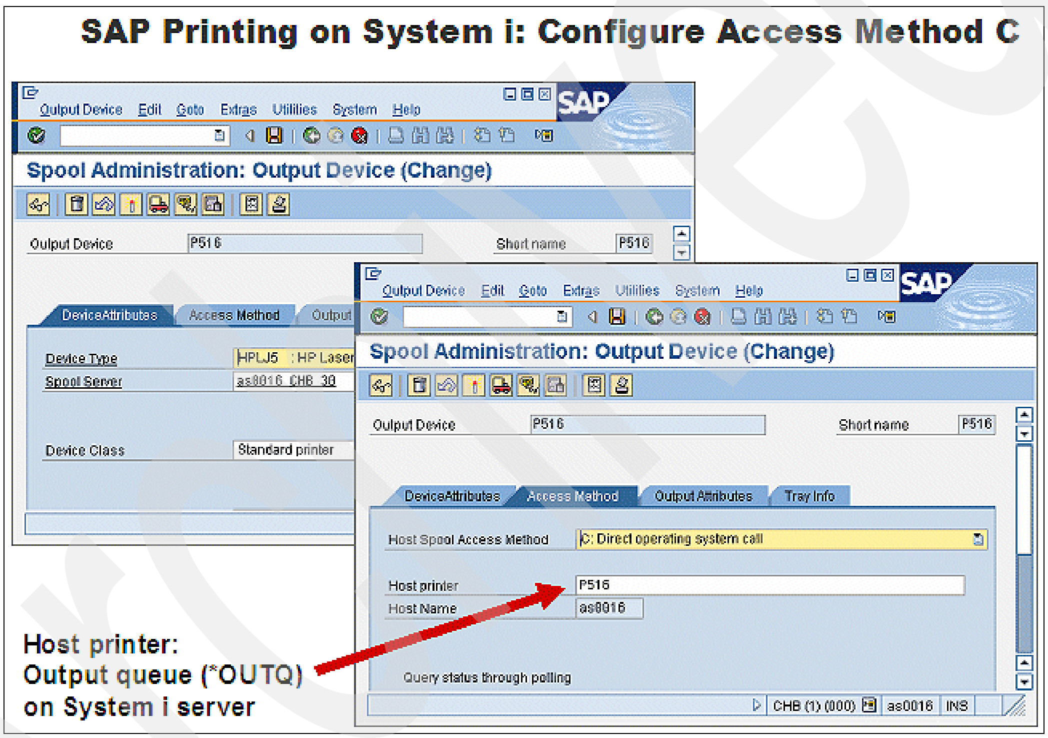

Figure 15-22 Configure access method "C"

Comparing access method C and access method U

When configuring a printer in the network for SAP applications on a System i model, you must decide on an access method for that printer. See Figure 15-23 on page 238 to understand the relationship.

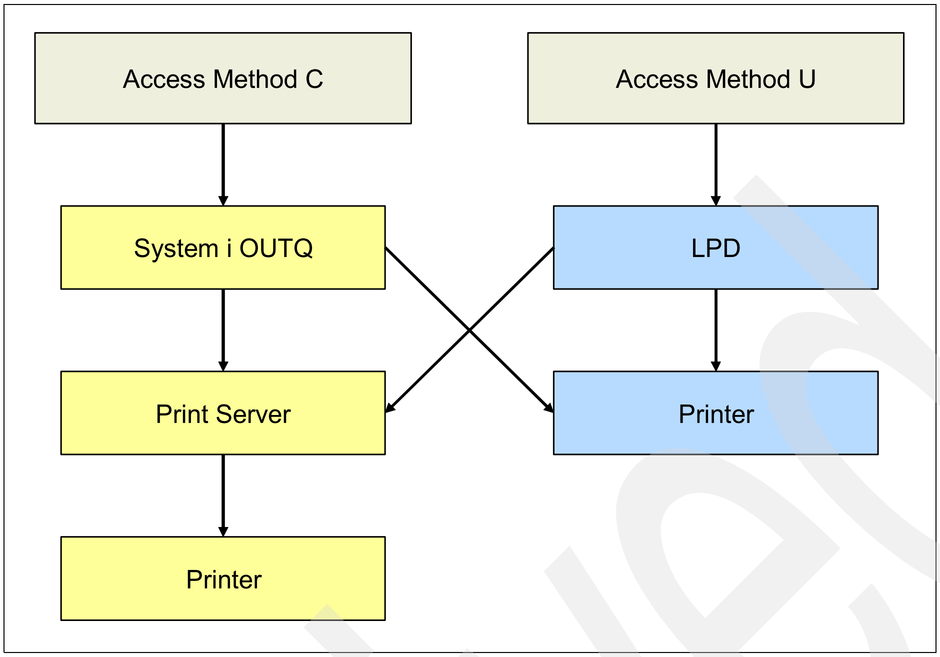

Figure 15-23 Access method C and U

Table 15-1 lists the advantages for access methods C and U.

Table 15-1 Access method C or U

Access Method	C	U
Technique	Print to output queue	Print through LPD protocol
Advantages	Sharing printer between SAP and non-SAP applications Spool work process is not blocked if printer is not available	Direct feedback from printer about success of printout Single point of maintenance in transaction SPAD

From an administration standpoint, access method U is easier because you do not need to watch the output queue and writer job for potential errors. Also, when choosing access method C the spool request is shown as “completed” in transaction SP01 as soon as the spooled file is created in the output queue, not when it has been sent successfully to the printer. Conversely, only access method C allows the sharing of a printer between SAP WebAS and other applications. And when a print server is not available, for example because it is switched off, the request blocks the spool work process for some time before it times out. With access method C, the spooled file is sent to the output queue and the spool work process is available again without waiting for the printout to complete.

Instead of using an output queue and access method C, you can also access a network printer through access method U directly, as on any other platform. Using access method U has the following advantages over access method C:

- You get direct feedback from the printer in your list of spool requests (transaction SP01).
- You do not have duplicate maintenance by checking both the SAP spooling system and the output queues and writers on i5/OS.

If a printer that is accessed through access method U is temporarily not available, the spool work process can wait some time before it assumes a timeout. In order to avoid bottlenecks in the SAP spooling system, configure sufficient spool work processes to process spool requests even if a spool work process is waiting for feedback from a non-available printer.

15.5.4 SAP device types

To create a new device type for SAP applications, 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. Select a device type and click 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.

An existing SAP standard device type can be modified instead of creating a new one. However, these changes can be lost at the next release upgrade. Create a new device type using the template of an existing device type. A reference to the standard device type is kept in the new device type for SAP Release 4.6B and later. The benefit is that changes made to the standard device types (most likely by SAP) immediately affect all device types that reference that one.

The IBM printers and device types currently supported by SAP are referenced in SAP note 8928.

Device types SAPWIN and SWIN

Device types `SAPWIN and SWIN` are described as follows:

- `SAPWIN` is a generic device type for printers (or also fax devices) linked to PCs running under MS Windows by means of the SAP program SAPLPD. Windows printer drivers are used and the character set is ISO 8859-1. Barcode printing from the SAP system is possible with the additional installation of a third-party Dynamic Link Library (DLL), but is not supported in the standard system. OCR-A and OCR-B fonts are possible with an appropriate TrueType font. See SAP note 25344 and 48803 for more information.

  SAPWIN is used to:

  - Print proportional fonts, lines and boxes in SAPscript
  - Print black and white as well as color Tagged Image File Format (TIFF) graphics (with the 32-bit SAPlpd on MS Windows only).

  See Note 39031 for more information. See page 240 for a definition of SAPscript.

- `SWIN` is a generic device type for printing and faxing using MS Windows with SAPlpd. It is available as of Release 3.1H. Device type SWIN differs from SAPWIN in the layout of X_... formats for ABAP/4 list printing. The font size and line spacing are adjusted, particularly for color list printing.

  In addition, SAPWIN contains an expanded list of SAPscript printer fonts for current releases, for example, Greek and Russian.

- SAP note 114426 gives a good overview about all aspects concerning "Frontend printing" as a collective note. This note describes the general mode of operation of frontend printing and all related problems known to date.

  Frontend printing means that the print data is sent to the computer on which the user has started SAPGUI. The data should be output on the default printer of the user, if possible.

As of Release 4.6A, the list output of the frontend printer can be selected in the printing dialog box, for SAPScript only as of Release 4.6. The name of the frontend printer is restricted to 40 characters.

Owing to its mode of operation, frontend printing is subject to the following restrictions:

– Frontend printing is suitable for the output of a small number of documents. In particular, this restriction can occur with a slow GUI connection. Frontend printing does not replace other access methods.

– Background use is not possible. Since there is no connection to a frontend in background processing, frontend printing is not possible here either. In some releases, when you define print parameters for a background job, the entry of a frontend printer can be prevented or it can be permitted, but no output results from the batch.

Definition of a frontend printer

You create a frontend printer with access method 'F' in Transaction SPAD. You do not have to enter a formatting server or a destination host.

The easiest way to define the frontend printer is with the Device type SAPWIN (or SWIN as of Release 4.0) and the host printer name __DEFAULT (put 2 underscores before the word DEFAULT). However, this is only possible if the SAPLPD program in the subdirectory SAPLPD of the SAPGUI is installed on the frontends and no other LPD (for example, TCP/IP Print service) is running at the same time on port 515.

SAPLPD should be used for frontend printing for the following reasons:

– The Device type SAPWIN/SWIN can be used for all printers.

– The Printer name __DEFAULT automatically uses the default printer.

– SAPLPD is started automatically if no LPD is running on the computer.

To learn more about how to use an SAPLPD, refer to SAP note 114426.

Windows Terminal Server (WTS)

To print with the Windows Terminal Server (WTS), refer to SAP note 150533.

Attributes for a new device type

Fill in the following input fields when creating a new device type:

- *Device type*: The name by which the new device type is identified. Start the name with the character Y or Z to avoid a conflict with the names of the predefined device types.
- *Name*: A field that allows you to enter descriptive information about the device type. For example, the name and model of the related printer can be entered.
- Version: A field that can 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. The previous used special driver to convert from OTF to AFP (the CVTPRTDTA tool) is no longer available on i5/OS. 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 defined for 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 used, but must be filled in. The same ID can be used 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. 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, unprintable characters can be produced. If the character does not have the same hexadecimal value, a different character than intended can print.

Create a unique character set if no other set meets your needs. See 15.5.7, "SAP characters and character sets" on page 243 for more information.

15.5.5 Paper types

A format (paper type) groups formatting information for a printed page. The formatting information is not given directly in the format, but is a page format that is linked to the 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. Values are filled in 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 the name with the letter Y or Z to distinguish those provided by SAP.
   - *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. Any page format is valid for this format if you use the value ANY.
   - *Orientation*: Use this field to specify the orientation of the printed page. Both check boxes for portrait and landscape can be selected since this is used for documentation only.
   - *Comment*: Enter a descriptive comment to document the format.

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. Click the **Create** button to create a new one.
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. 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.

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. A window appears where the control characters can be edited for the 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. 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. Refer to the *BC-CCM-PRN SAP Printing Guide in the SAP Online Help* to understand what to specify.
5. Double-click the button of the particular action to edit the control characters for an action. This shows a window where you can enter one or more lines of information. Enter the exact sequence of escape characters that are needed to accomplish the desired effect at the printer.

   The SAP programming editor processes this window. The control character sequences needed can be found in the printer manual for the specific printer being used. As a result, there is a language syntax that you use for entering the information. This syntax is described in the *BC-CCM SAP Printing Guide in the SAP Online Help* at:

   http://help.sap.com

   Enter the attributes for the device type format after you enter the information for each action. From the window that lists the actions, click the **Attributes** tab.
6. 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.

15.5.6 Print controls

Actual print controls are placed in the data when 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 can see, for example, that to set printing at 8 characters per inch (cpi), the symbolic name CI008 is used. A symbolic name is 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 PrCtls**. A window appears that allows you to select either all print controls for every device type, or a print control and device type. A list is shown where each print control table is identified by the name of the associated device type if you leave the defaults. A print control table does not exist independently of its associated device type, nor do two device types share the same print control table.

Complete the following steps to see the print controls that are currently assigned to the device type:

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, including a field that holds the control character sequence. Note the following considerations:

- 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. The control characters are documented in the technical manual for the specific printer.

Define a new print control table when a new device type is created. 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. If you add a new print control name, 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.

15.5.7 SAP characters and character sets

SAP 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.

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.

A character set is used to specify the exact encoding for printing a character at a printer. 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 SAP 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. Click the **Edit character set** button to select one from the list and see the mapping for a character set.

The character sets predefined by SAP are extensive. It is unlikely that you need to define your own characters and character sets. Even if you need to define a new device type, you can use one of the predefined character sets.

SAP has provisions for you to:

- Add new characters to the master list
- Create a new character set

There 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-CCM SAP Printing Guide* in the SAP Online Help at:

http://help.sap.com

15.6 Configure new SAP system device

Sometimes it is necessary to add a new device type to your SAP system because the existing printer device drivers do not match the function of your printer. Before a device type or printer device in SAP is added, the printer must be identified to the System i configuration.

Complete the following steps to add a new device type:

1. Enter the SPAD (Spool administration) transaction.

   When the SPAD Initial window appears, click the **Full administration** button, click the **Device Types** tab, and click the **Device types** button.

2. Click the **Change** button and look for an existing device type in the list where the control sequences resemble your new printer. This is your template device.

   > **Note:** Most of today's impact printers have the ability to emulate an HP laser, an Epson, or a Lexmark printer. You can use one of those printers as a template even if your new printer is an impact printer. The same analogy can be applied to laser printers.
   >
   > Many IBM SCS printers have no escape control sequences. Therefore, these printers must be controlled by i5/OS printer files or directly by their own settings from the control panel.

Select the printer and click the **Create using template** to create the new device type. The new name should start with a letter Y or Z according to the SAP naming convention for customer objects.

3. Click the **Save** button to save the new device type.

   Go back to the spool administration window and click the **Print controls** button. A list of existing standard print controls is shown.

4. Click the **Create using template** button to create new standard print controls based on an existing standard print control. The Spool Administration: Copy Standard PrtCtrls window appears. Enter the name of the new standard print control and a description in the Comment field. In the Print control status box, choose which of the three print types this new print control can be used with.

   > **Note:** Use the standard form CIxxx and LIxxx as the name of the new standard print control, where *xxx* is CPI or LPI if the printer functionality falls into the categories of characters per inch or lines per inch. If the functionality does not fit into one of the standard categories, type a five-character code that is meaningful.

5. Click the **Save** button to save the new print control and go back to the spool administration window. Click the **Device types** button. Select the device type just created, and click the **List of implemented print controls** button. A list of print controls is shown that was copied from your template device. These print controls are currently assigned to your device type.

   To add a new print control, click the **Standard print controls**. All standard print controls that are available are displayed. Double-click the print control just created. Then the Spool Administration: Edit Print Controls window appears.

6. Enter the needed control character sequence according to the device 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 shows a list of the implemented formats for the selected device type.

8. Double-click the format to modify the printer initialization. The Maintain Format window appears.

9. The most frequently modified format action is printer initialization. Double-click **Printer initialization**. Then the Printer init window appears. Enter the commands 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 “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.” on page 242.

10. It can be helpful to also add some or all of the printer initialization control characters to the *Start of page* format, especially for large output requests. This somewhat increases the time needed for the print output. The benefit is that this initializes the printer at the beginning of every page. This helps to prevent displaced print output.

11. If any special formats and page format are needed, follow the instructions in “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." on page 242.

Refer to SAP note 17054 for more information about how to copy or create a device type.

15.7 Barcode printing definitions

The following are hints on how to handle barcode printing:

- Printer with barcode printing support

  If your printer supports barcode printing, 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 the printout.

  – SBS (barcode suffix) is the control sequence following the variable barcode information.

  Follow these steps to create a printer with barcode printing support:

  a. Call the SPAD transaction.

  b. Click the **Full administration** button. Click the **Device Types** tab and then click the **Print Controls** button.

  c. Select the **SBP** entry from the list. Click the **Create using template** button to create the barcode prefix.

  d. Click the **Save** button to save the data.

  e. Select the **SBS** entry from the list, and use **Create using template** to create the barcode suffix.

- Define new barcode print controls

  If barcode print controls are defined for a device type, associate them with an existing system barcode or define a new system barcode to associate them with.

  To do this, complete the following steps:

  a. Call the SE73 (SAPscript Font Maintenance) transaction.

  b. Select **System barcode** to define a system barcode.

  c. Click the **Change** button. On the next window, click the **Create** button and fill in the needed information.

  d. To associate the 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 the device type. If the device type you use as a template has defined barcodes, these are listed in a new window.

  e. To add either standard barcodes or your own previously defined barcode, click the **Create** button and select a desired barcode. Enter the associated print controls SBPxx (barcode prefix) and SBSxx (barcode suffix).

  f. Save the entry and go back to the list of printer barcodes.

  g. Select the device type. Then, click the **Check print controls** button to see the listing.

  h. The print controls for barcode printing are now associated to the device type.

- Create an output device

  To create the output device, perform the following steps:

  a. Call the **SPAD** transaction.

  b. Click the **Output devices** button and the Spool Administration: List of Output Devices window.

  c. Click the **Change** button and then click the **Create** button.

  d. Specify the output device and the short name. Select the device type and the spool server from the list box.

  e. Specify the Host spool access method, and enter the Host printer, which is the name of the output queue with i5/OS.

15.8 Manage SAP spooled and output requests

This section discusses how to manage SAP spooled files and the actions you can take in this process. Spooled files are not managed directly, but rather through the spool requests. By working with a spool request, you indirectly manage the TemSe data.

15.8.1 Attributes of a spool request

Managing SAP spool and output requests requires that you navigate to the Output controller window. To do this from the SAP main menu, call the SP01 transaction.

Once the window appears, you can display a list of spool or output requests that are in the SAP spool system. Click the **Spool requests** tab to list the spool request. Enter the criteria for identifying the spooled requests.

Different criteria can be used to identify specific spooled requests. One or more of the following values can be specified:

- **Spool request number**: The number assigned by SAP to the spool request.
- **Created by:** The name of the SAP 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:** SAP spool authorization.
- **Output device**: Spool requests for a specified output device. The device is identified by the SAP 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 SAP system name. It must be a valid RFC destination. If the field is left empty, all the spool requests for all the systems in table ALCONSEG are listed. Use SAP transaction RZ20 to maintain this list.

Once you enter the criteria, a list appears of those spool requests that meet the criteria. For each entry selected, attributes about that entry are shown, including the 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 the criteria for identifying the output requests. The selection criteria is basically the same as the above list.

> **Note:** Enter a valid RFC destination in the System name parameter to display spool or output requests from other SAP systems.

15.8.2 Spool request actions

The spool request actions available include to:

- Display the content of the spool request in character or hexadecimal format.
- Change attributes of a spool request. Attributes that can be changed include:
  - The title, the name and the user name
  - The output device to print the data
  - The format used 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, whether or not it is printed
  - The number of copies to print
  - The priority at which the spool request is processed for printing
  - The status of the spool request after printing which controls whether it is deleted after successfully printing
- 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 for more information.

  > **Note:** It is important to keep the size consumed by the spool requests to a minimum.

For details about any one action, refer to the *SAP Online Help* at `http://help.sap.com`

15.9 Resolving printing problems

Sometimes spooled files do not print and it can be difficult to determine why. This section points out some common causes for files that do not print if the i5/OS spool system is involved. See Chapter 20, “Problem determination and management” on page 339 for general problem determination techniques.

Follow these steps to resolve specific printing related problems:

- Problem: Spooled file is not generated

  Action: Try the following actions:

- Search for the output request using the SP01 transaction. Make sure the status is “Compl.” (completed).
- Look at the corresponding joblog on the System i server. Use transaction SM50 to determine the process ID of the spool work process 1. Use the **`WRKPID`** (Work with SAP process ID) command to display the joblog of the spool work process. Select the proper spool work process in the correct instance.
- The output queue is not in the library list of the spool work process

  The system might have placed it into the default output queue. Search for the spooled file using the command:

  ```
  WRKSPLF SELECT(<SID>OWNER)
  ```

- Problem: Spool work process or the spool server is not active

  Action: Make sure that the spool work process can be accessed by the application if a spool work process is not in the instance or a spool server is defined on another system.

- Problem: A spooled file is in the output queue, but it is not printing

  Action: Try these actions:

  - Use the `WRKOUTQ` command to check whether the output queue is released. If not, release the output queue using the `RLSOUTQ` command.
  - Use the `WRKWTR` command to verify that the printer device has been started.
  - Check whether the printer that is attached to the output queue can handle the spooled file type.
  - It is possible that the i5/OS spool system tried to print the spooled file, but found some type of error. Depending where the messages for this printer go, there can be a spooled error message indicating that the i5/OS spooled support encountered a problem. There can be a problem with i5/OS spooled support. Check the SAPNet - SAP Frontend (formerly known as OSS) to see if this is a known problem.

- Problem: The print requests on a printer with access method C are completed in SAP without error but do not appear on your printer

  Action: Review the WRKOUTQ command display. Verify in the upper right part of the screen that a writer job is started and that the output queue is released.

15.10 Printing documentation references

For more detailed information about SAP printing than what is provided in this chapter, refer to:

- *BC-CCM SAP Printing Guide* in the SAP Online Help at:

  `http://help.sap.com/saphelp_nw04/helpdata/en/d9/4a8eb751ea11d189570000e829fbbd/content.htm`

  This publication contains information about all aspects of printing from SAP.

- *BC-ABA ABAP Programming*

  See the BC-ABA ABAP topic at the SAP Help Portal:

  `http://sap.help.com`

  Refer to this publication for information about using ABAP programs to print.

- *BC-SRV-SCR SAPscript*

  See the BC_SRV_SCR SAPscript topic at the SAP Help Portal at:

  http://sap.help.com

  Refer to this publication for information about producing documents through SAPscript.

- *ADM525 - SAP NetWeaver on DB2 UDB for iSeries: Database Administration*

  This class teaches concepts about printing on SAP.

- SAP Marketplace for Output Management Systems

  To make an output device that is served by an external output management system (OMS) accessible for SAP printing and or faxing.

  http://service.sap.com/output

16

SAP system administration

This chapter summarizes the activities that you should do to help ensure your SAP system operates smoothly. These include:

- Checking whether the system is up and running
- Checking whether the required functionality is ready
- Checking, analyzing and correcting errors and other situations that do not appear normal
- Performing proactive tasks to help the functionality in the future
- Optimizing the system

In the following sections, we distinguish between monitoring tasks and other administration tasks.

Monitoring tasks are those where the system itself provides you with data and information you have to check. Before the system can do this, you first have to implement and adjust some monitoring programs, monitoring tasks, or monitoring methods. To set up *advanced monitoring*, you should define and implement individual variables and threshing values. When these threshing values are reached or exceeded, a trigger is alerted to inform you about the related event. This alert can be a system message, an e-mail, a Short Message Service (SMS) or some other method of information. Therefore, monitoring tasks are caused automatically in periods, or by events.

Administration tasks are those where a person, for example, a system administrator, works manually on the system. Often, the system administrator has a checklist or of tasks they have to do. They choose when to work on the system to perform the administration tasks. Administration tasks must be done for special situations which cannot be automated. If you do not have a monitoring system, you will have to perform a lot of repetitive and tedious tasks manually. However, you cannot be sure that problems will not arise after you have completed checking the system.

Try to set up a monitoring system to ensure that no problems arise.

16.1 SAP system monitoring

This section describes some important SAP transactions for the SAP system administrator.

An SAP system consists of many software and hardware components that contribute to the overall availability and performance of your SAP installation. As illustrated in Figure 16-1, these components include the operating system, the CPU, the physical memory, and the disk devices. The next layer of major components are the database, SAP buffers, and SAP services (for example, dialog, update, enqueue, spool).

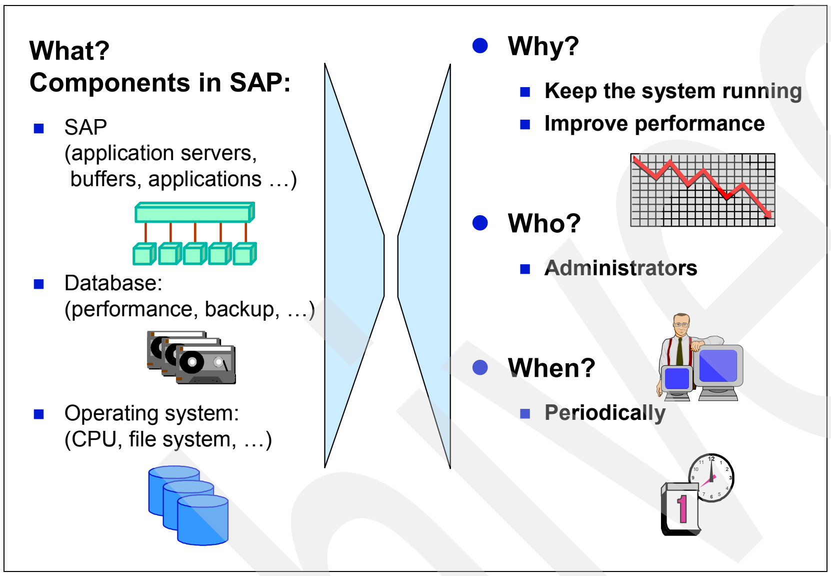

Figure 16-1 Monitoring the SAP system

The administrator should monitor each of these components regularly.

The main goals of system monitoring are to:

- Keep the system running.
- Analyze and correct errors.
- Improve performance, and thus increase user acceptance of SAP.

System monitoring can be performed by the following people depending on their area of responsibility:

- SAP system administrators

  They are responsible for assuring the performance of the SAP system.
- Database administrators

  They are responsible for assuring the consistency of the database and for restoring the database if a database inconsistency or database loss occurs.
- Operating system administrators

  They are responsible for providing physical storage media.

In some SAP systems, you have to log in to the database separately after logging in to the operating system. This is not necessary for System i implementations.

Note: On System i models with its integrated database, you need not distinguish between a database administrator and an operating system administrator. i5/OS implementations do not need an extra database administrator, as recommended by SAP for other platforms, or even an additional operating system administrator.

In general, especially in large installations, the system administrator performs the necessary administration tasks (including the database and operating system-specific tasks) on System i models.

You do not require specialized skills to run SAP applications on the System i platform, in the i5/OS operating system, or in the DB2 UDB for i5/OS database.

For a System i system administrator, the administration of an SAP system on the System i installation is a normal and relatively easy job.

Consider a ratio of about 20:80 for System i tasks to SAP basic system tasks. You can classify administration tasks as:

- Tasks to administer the System i server
- Tasks to administer the SAP basic system

The SAP system should be monitored regularly, at least once a day. SAP recommends more frequent monitoring, depending on the size of the installation.

16.1.1 Monitoring session information

On the bottom right of an SAP session in the SAP GUI front-end window, beside the triangle, there is a drop-down menu as Figure 16-2 shows.

Figure 16-2 Drop-down menu on the SAP GUI front-end window

When you click this menu, you will see the SAP session information as Figure 16-3 shows.

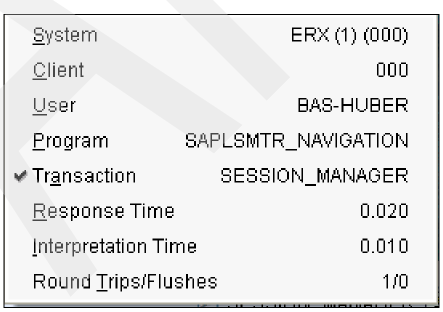

Figure 16-3 SAP session information

Following is a description of some of the fields you see in Figure 16-3. The others are self-explanatory.

- System information

  For example, ERX(1) means that the SAP System ID is ERX, the session or modus is 1 in client 000.

- Response times

  This is the time that is spent on network communication from the moment data is sent to the server to the moment the response from the server arrives.

- Interpretation times

  Interpretation time begins after data has arrived from the server. It includes the time taken for the parsing of the data and the distribution to the SAP GUI elements.

- Round trips

  Before SAP GUI sends data to the server, it locks the user interface. In many cases, it does not unlock the interface once data arrives from the server, but immediately sends a new request to the server. Controls use this technology to load the data they need for visualization. This field offers a count of these token switches between SAP GUI and the server.

- Flushes

  This counts the number of flushes in the automation queue during server communication.

16.1.2 Monitoring SAP information

When you select **System** → **Status** in each SAP menu bar, you can see a lot of current and significant SAP system information.

Figure 16-4 on page 255 and Figure 16-5 on page 255 illustrate this.

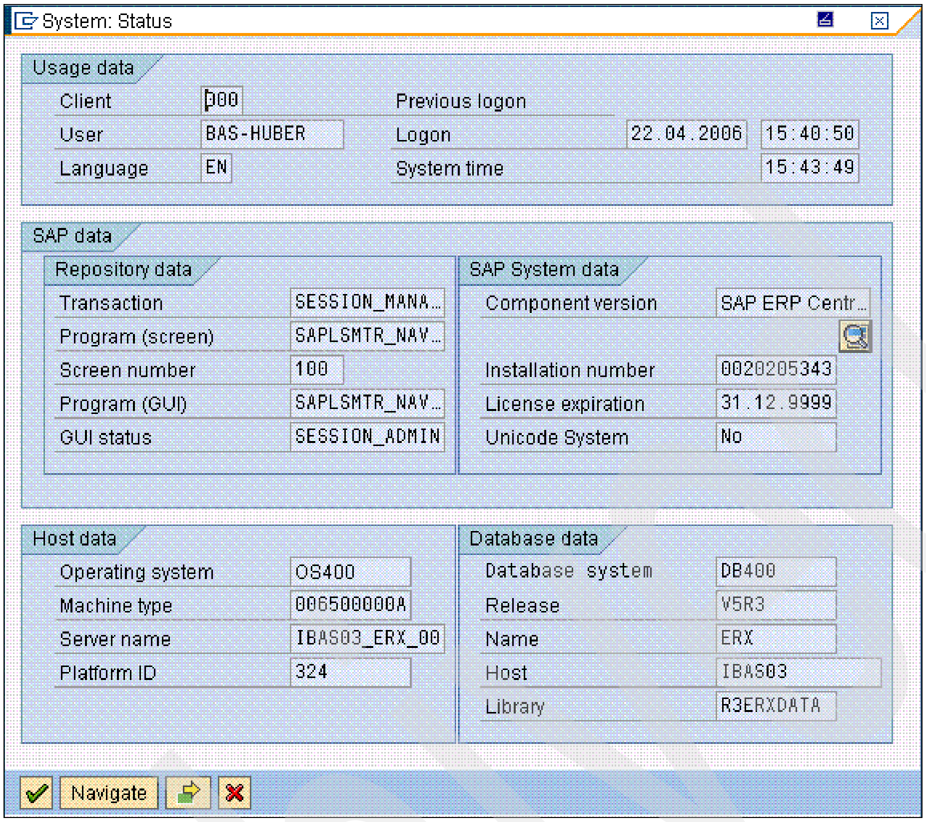

Figure 16-4 The SAP system status (1/2)

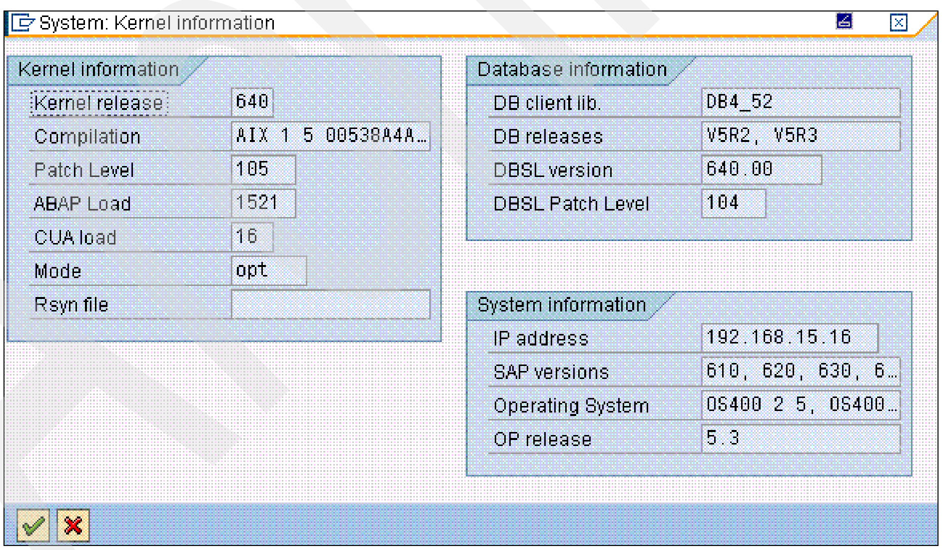

Figure 16-5 The SAP system status (2/2)

Some of the most significant information is as follows:

- Usage data

  This contains previous login information.

- Current SAP data

This contains the following information:

- The version or release of your current SAP component
- SAP installation number of your SAP application
- Unicode or non-Unicode installation

- Host data

  This contains the following information:

  - Operating system: still called OS400 on SAP systems
  - System i serial number: called Machine type on SAP systems
  - Database system: called DB400 on SAP systems
  - Database release (also the i5/OS release, which is the same as the DB2 UDB for i5/OS release)

- Kernel information

  This is the kernel release and the patch level of the Disp+Work (DW) kernel component.

- Database information

  This is the database release and the patch level of the lib_dbsl kernel component.

- System information

  This contains the following information:

  - IP address
  - Operating system release

16.1.3 Monitoring an SAP server: Transaction SM51

Transaction SM51 provides you with an overview of available servers.

You can use this transaction to:

- Examine the processes of the server you are logged in to
- Display the users of the system
- Display the Secure Network Communication (SNC) status
- Display the system log
- Display release information
  - Display kernel information such as:
    - Kernel release
    - Kernel patch level, more precisely the patch level of the DW kernel component
    - Patch level of the lib_dbsl kernel component

      Refer to Chapter 12, “The SAP kernel on the System i server” on page 137 for more information about the kernel components, especially DW and lib_dbsl.
  - Supported environments for that kernel
    - SAP database releases
    - i5/OS and OS/400 releases
    - DB2 UDB for i5/OS releases

      On System i models, the supported operating system release and the database release are the same.
  - The text of the imported kernel patches

- Display details to the application server

  These details include:

  - Application server names
  - Host name
  - Types of the installed work processes
  - IP address of the host name
  - Port number of the host name

- Dynamically switch to another server

16.1.4 Monitoring SAP work processes: Transactions SM50 and SM66

SAP transactions SM50 and SM66 provide an overview of the work processes on the application server you are currently logged on to, or the whole SAP system.

These screens display the type of work process, the status, CPU time consumption, actual action, current user, and client. From this screen, you can stop and start a work process, and debug and trace a work process.

16.1.5 Monitoring an SAP user: Transactions SM04 and AL08

With SAP transactions SM04 and AL08, you can also get an overview of users on a specific server, or of all the users in the SAP system. Detail information is displayed for the user about the location, name, session, modes, and more information about the user data.

From here you can perform actions such as deleting the user or the session, tracing user activities, and more.

16.1.6 Monitoring asynchronous update and update processing: Transaction SM13

To display terminated updates, call Transaction **SM13**.

With Transaction SM13, information is displayed that indicates:

- Whether updating is active
- What the update modules are
- Error information
- Data that should have been updated, but has not due to the following reasons:
  - An error occurred which prevented the data from being updated
  - The update process to update the data has not been run

You can perform the following actions such as:

- Posting the asynchronous update task
- Testing the update task
- Starting the update task (again)

An update in SAP systems is usually executed asynchronously. This means that the update information is written into a set of tables (such as VBMOD, VBHDR, VBDATA and so on), and updated from these tables when the transaction has finished.

You can perform time critical updates (known as *V1 updates*) and non-time critical updates (known as *V2 updates*) in several steps, for example, a statistics update. In its application programs, SAP defines what V1 and V2 processes are. V*x* is an abbreviation for “version x”.

An update is triggered at the end of an SAP transaction by a COMMIT WORK operation. This is done explicitly in the program coding, or implicitly at the end of an SAP transaction by default.

The update work process reads the data from the VB* tables and executes the updates in the corresponding tables of the SAP database. Once the update is successfully completed, the entries in the VB* tables and in the lock table are deleted.

If an error occurs, the lock table entry is deleted. The entries in the VB* tables are not deleted. The user is notified immediately by an express mail entry. Depending on the business data object, the entry can be reposted.

SAP does not recommend “posting” or repeating the transaction when using Transaction SM13 for V1 updates. We recommend that you restart the transaction after you have solved the problem that caused the original update error.

SAP recommends that you post aborted V2 updates using Transaction SM13, but only if the corresponding V1 update is successful.

> **Note:** You can implement an asynchronous update with the Advanced Business Application Programming (ABAP) key word `Call function .... in update task`.You can also update the database synchronously inside a dialog service. Use the synchronous update, otherwise known as a *hard* update, for special cases only.

16.1.7 Monitoring SAP locks: Transaction SM12

To display the current SAP object locks, call **Transaction SM12**. This transaction provides information about the following:

- The user that has initiated the lock
- The client of this user
- The table that is locked
- The lock argument

In this transaction, you will get information about:

- What the locked tables are
- Who the lock holder is
- What the lock type is
- What the lock area is

You can start activities to:

- Delete the locks.
- Test an enqueue.

After investigating and analyzing the errors associated with the user holding the lock, you can delete the lock entry. Before deleting the lock entry, consider the following:

- Is the user holding the lock actually using the transaction?
- Did the user start the update transaction and then leave the computer or workstation?
- Is it a long-running transaction with several locks?
- Is the lock not released due to a terminated update?

You cannot use locking mechanisms on relational database systems for complicated business objects, such as delivery orders located in several tables. To coordinate parallel access to this business object data in an SAP system, use the enqueue work process as Figure 16-6 shows.

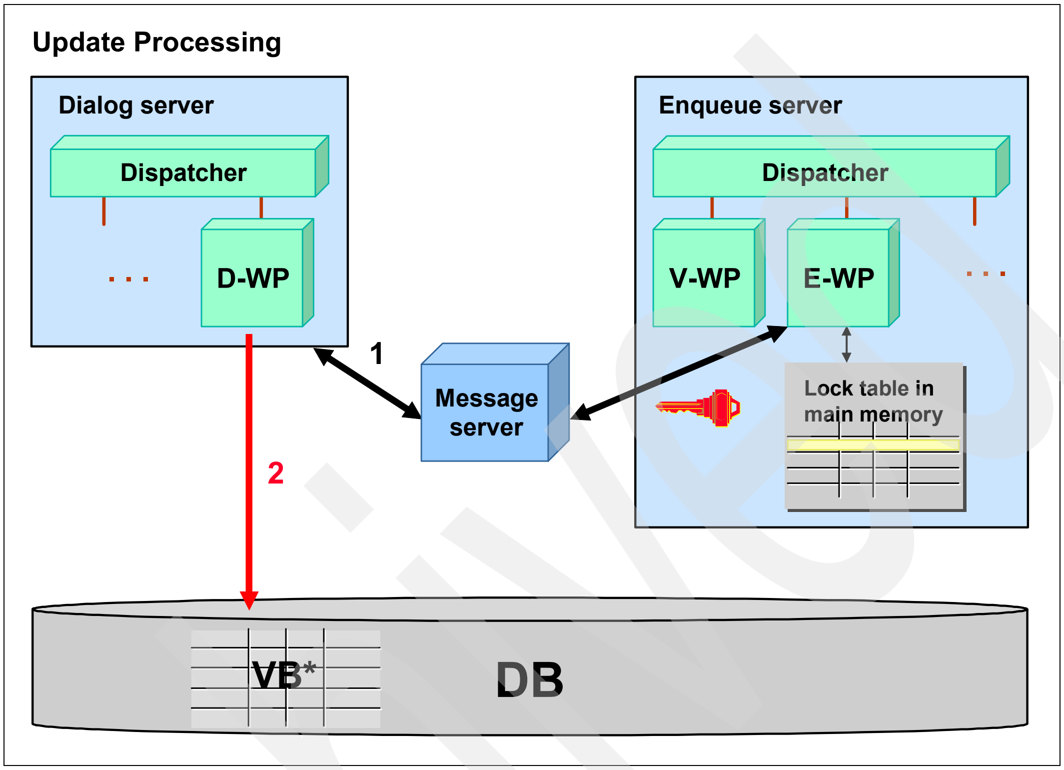

Figure 16-6 Update processing

The locks (=enqueues) are handled by the enqueue work process in a lock table, in the main memory on the server where the enqueue process is running. When a lock is requested, the lock table is checked for an entry of this data set. If there is already an entry of this type, the request is rejected and the user notified. If the dialog work process and enqueue work process are not located on the same server, communication is enabled by the message server.

For every business-related data object like a sales order, a material, a user profile, a printer or similar entities, a lock object can be defined in the ABAP dictionary. Such a business object is not a specific SAP object type, however, the lock object is an SAP specific object type. The customer name range for a lock object must start with the letters EY or EZ.

A lock object can have a shared lock mode (mode S) for read access, or an exclusive lock (mode E) for write access.

16.1.8 Monitoring ABAP dumps: Transaction ST22

If you receive an error message in the SAP system log, or if you see a terminated update in the update service analysis transaction, check for dumps using the dump analysis transaction, **Transaction ST22**. Alternatively, select **Tools** → **Administration** → **Monitoring** → **Dump analysis**.

The ST22 transaction allows you to analyze short dumps from the current and previous day.

The dump analysis function shows you:

- What happened where
- What the termination point in the program code is
- What the environment data at the time of termination is
- What you can do
- How to correct the error

The dump analysis function also provides information you can use for searching in SAP Service Marketplace, as well as information about:

- The system environment
- Users and transactions

You can analyze the following information:

- The date, time, user, and client of the transaction
- Contents of system and data fields
- Contents of internal tables and application tables

16.1.9 Monitoring the SAP system log: Transaction SM21

Call transaction **SM21** that shows you the SAP system log. Alternatively, you can call the transaction **SM21** with the SAP menu by selecting **Tools** → **Administration** → **Monitor** → **SM21 System Log**.

In the RZ20 - CCMS Monitor Sets, you can invoke a monitor template which shows you a lot of information for every application server. All application servers have the same structure, divided into the following parts, for example:

- BasisSystem
- Database
- CCMS (Computing Center Management System)
- Background
- Communication
- Spool
- Security
- TransportSystem
- BatchInput
- Application
- Customer
- Miscellaneous
- SyslogFreq

For more details about these monitor systems within RZ20, refer to the SAP Online Help or the SAP Help from the Help Portal on the Web at:

http://help.sap.com

The SAP system log provides information about these different parts of your SAP system. Use SAP transaction **SM21** to directly access the system log (syslog). The Syslog displays data showing the history of your SAP system.

The system log contains information about the SAP System, which is categorized into problem types, for example:

- MS - Message Server
- DP - Dispatcher
- RD - Gateway reader
- DIA - Dialog work process

- BTC - Batch work process
- SPO - Spool work process
- UPD - Update work process
- UP2 - Update V2 work process
- ENQ - Enqueue work process

There are also the following classes which you can select from within the system log:

- S

  Start and stop the SAP system and the work processes.

  Operation mode switches
- W

  The rollbacks performed
- K

  Kernel program errors
- T

  ABAP transaction errors resulting in short dumps

The displayed information shows the time stamp, client, user, transaction code, and some text. The work process which caused the message is also displayed.

Note: You must monitor the SAP system log separately for each application server.

Figure 16-7 shows you the SAP system log from the SAP Computing Center Management System (CCMS) which is started by transaction **RZ20**.

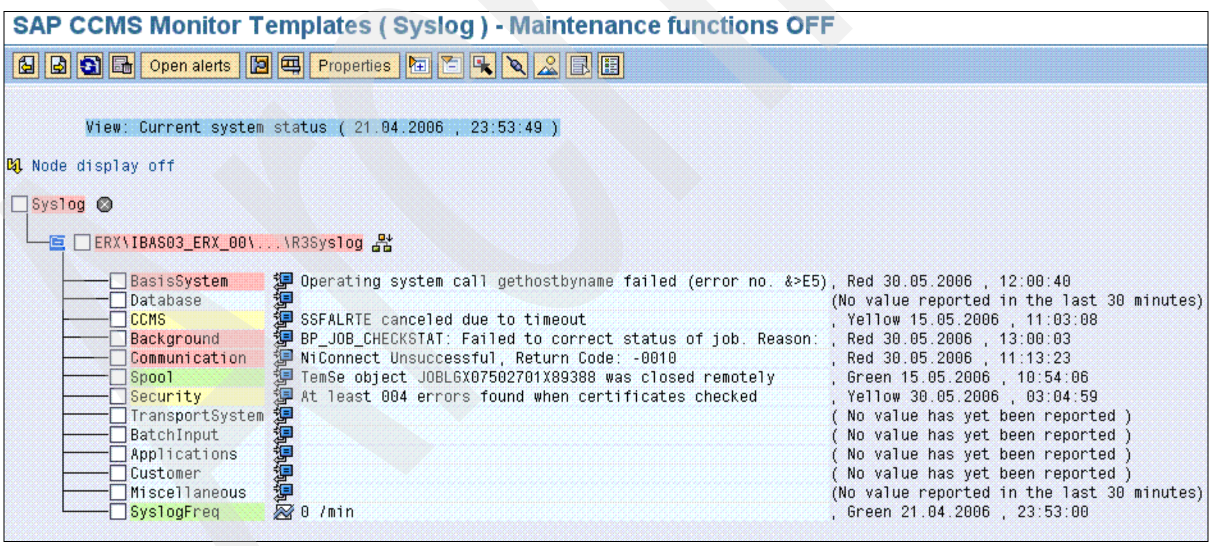

Figure 16-7 The CCMS Monitor templates (Syslog) from transaction RZ20

Here you see additional classes for the SAP syslog information.

16.2 SAP system administration tasks

There are numerous administration tasks within the SAP system environment. In this book, we provide you with a list from the SAP Help portal where you can see more details for each administration task that is of interest to you. Additionally, we show you the activities you should perform before shutting down an SAP system.

16.2.1 Global administration tasks

The SAP Web site provides a good explanation about the meaning and what to do for each of these administration tasks. The lists shown in Figure 16-8 on page 263, Figure 16-9 on page 264, and Figure 16-10 on page 265 show extracts from the Technical Operations Manual in the SAP library which is available in the SAP Help Web site at:

http://help.sap.com/saphelp_nw04/helpdata/en/start.htm

The SAP operational manuals from the other SAP applications, for example, from mySAP ERP ECC 5.0 in the SAP Help Portal, explain the additional tasks that are available. They explain the activities you should perform within the SAP system.

These tasks focus more on the application business view rather than the technical view, so they are only referenced. Most administration tasks are platform-independent.

> **Note:** You should consider the administration, monitoring and performance tasks within an SAP system and the underlying platform, operating system, and database as one entity.

Figure 16-8 on page 263 shows you the ABAP-specific administration tasks for the Web AS.

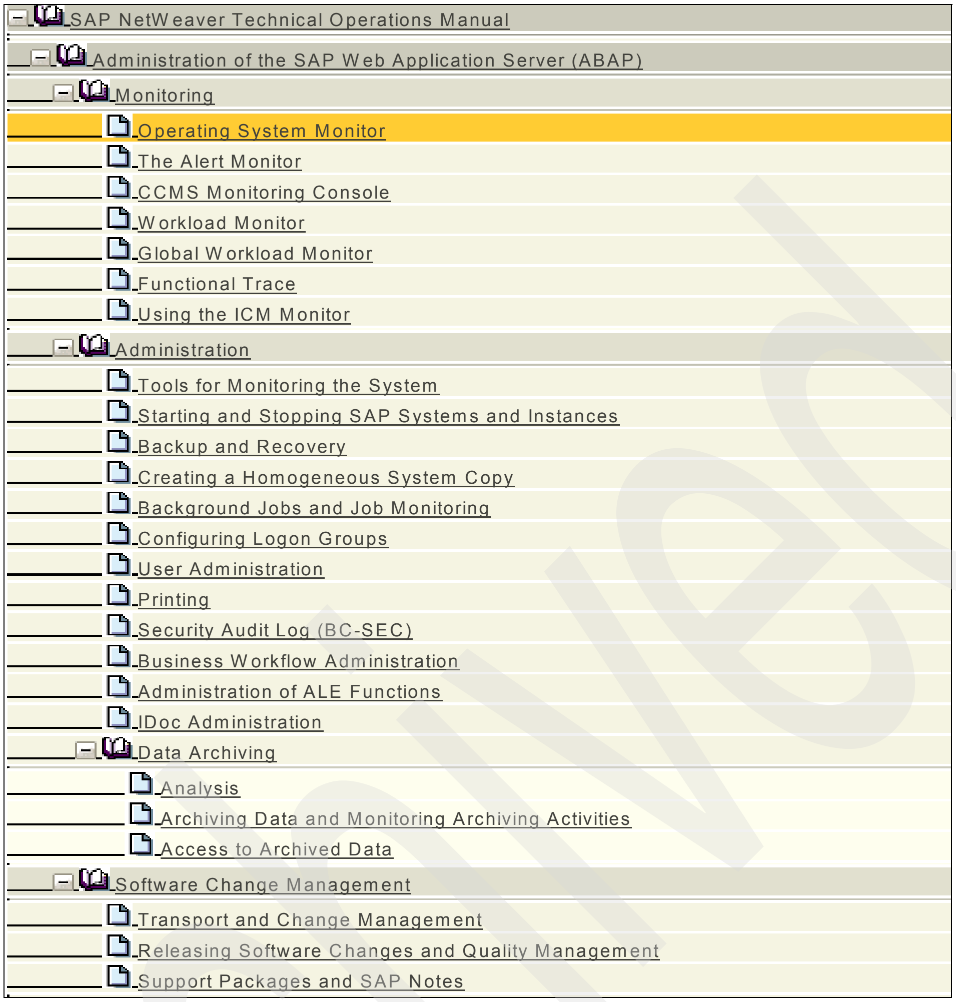

Figure 16-8 SAP NetWeaver Technical Operations Manual part 1 of 3 - ABAP

Figure 16-9 on page 264 shows you the Java-specific administration tasks for the Web AS.

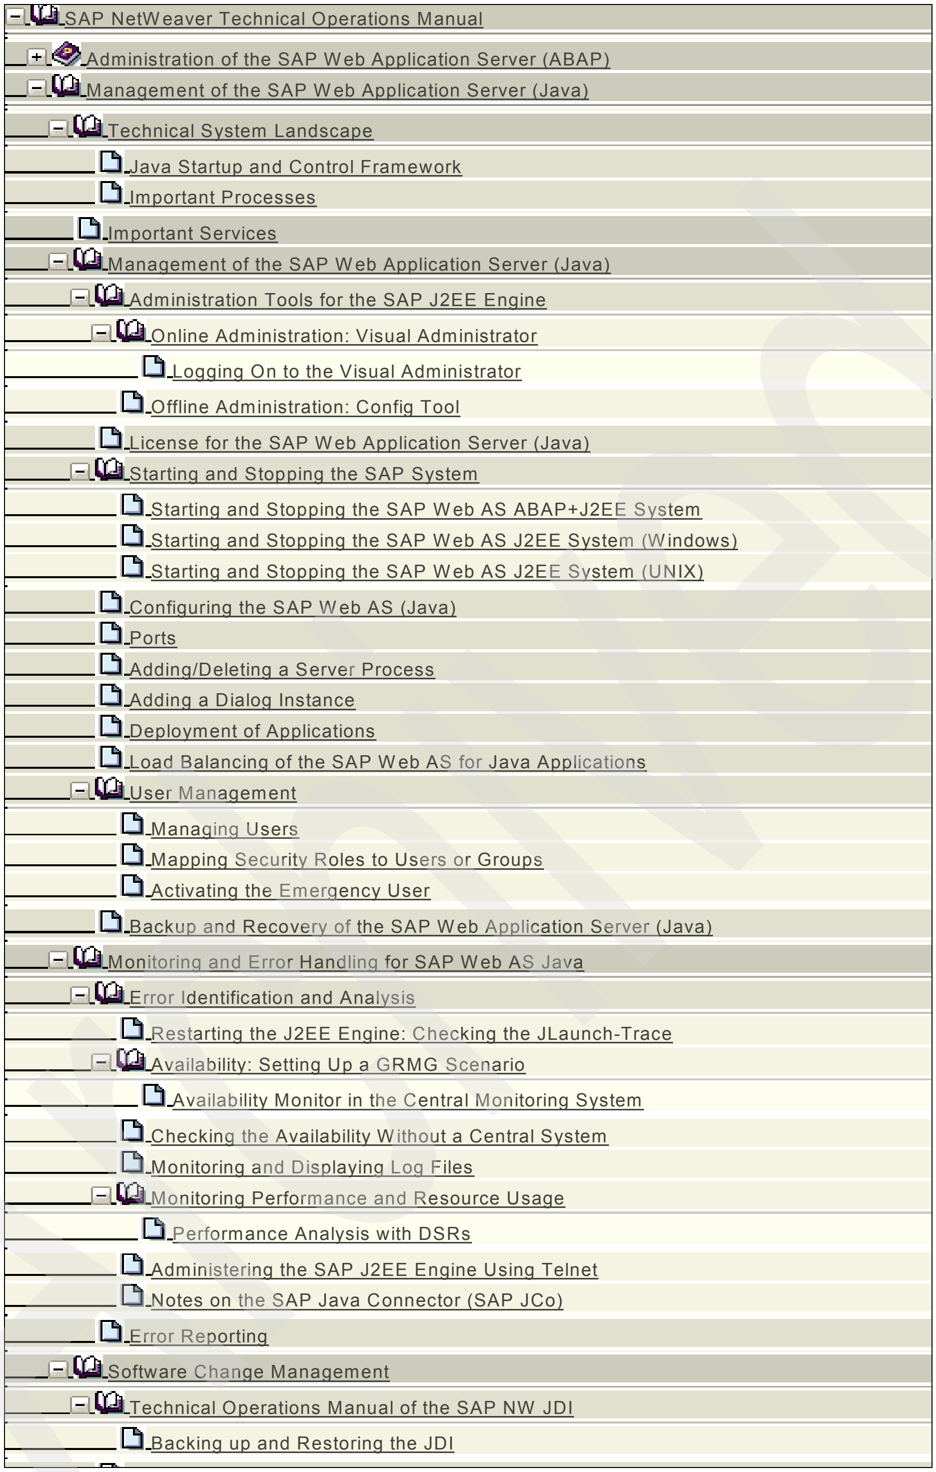

Figure 16-9 SAP NetWeaver Technical Operations Manual (part 2 of 3 - Java)

Figure 16-10 on page 265 shows you administration tasks for other components for the Web AS.

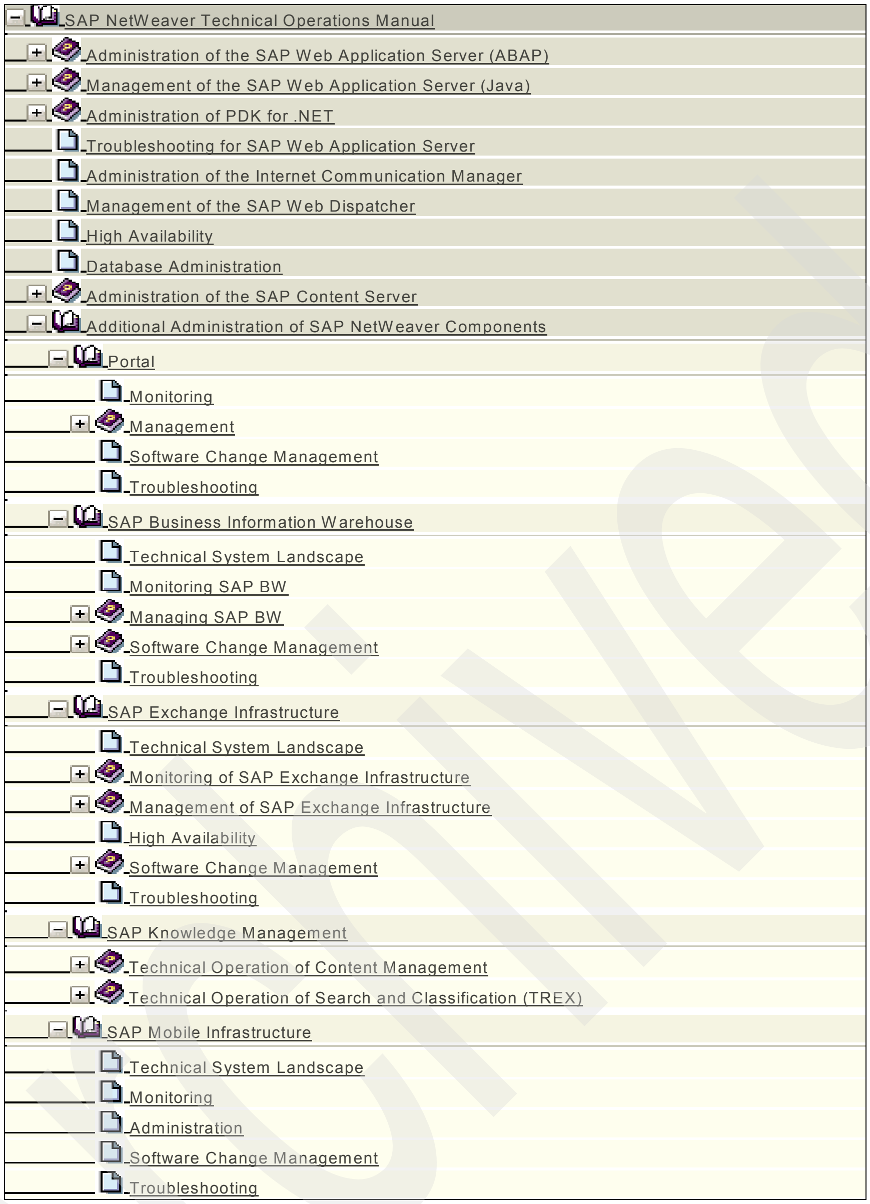

Figure 16-10 SAP NetWeaver Technical Operations Manual (part 3 of 3 - other components)

16.2.2 Activities before stopping an SAP system

Before the SAP system is stopped, the SAP system administrator should check the state of the SAP system by calling the corresponding SAP transaction:

- Job Overview

  Check if any background jobs from any application server are active or have been triggered externally. Use Transaction **SM37**, or choose **System** → **Services** → **Jobs** → **Job overview**.

- Process Overview

  Check if the background work process named BTC is active in any application server. Use SAP Transaction **SM50**, or choose **Tools** → **Administration** → **Computing Center** → **Management System** → **Control** → **All work processes**.

- Batch Input Overview

  Check if any batch input jobs are running. Use SAP Transaction **SM35** or choose **System → Services → Batch input → Edit → Overview**.

- Update records

  Check if any update records are open. Use SAP Transaction **SM13**. When the system is stopped, these records are rolled back and set to an init status. At startup, the records are processed again.

In each case, as SAP system administrator, when performing these steps, you should decide whether to interrupt the job or wait until the job is finished.

Prior to shutting down the SAP system, inform users by setting a system message through Transaction SM02. Before shutting down the system, use Transaction **SM04** to check whether any users are still logged on, and ask them to log off. The SAP system administrator and administrators of external systems should keep each other informed about data transfers between their respective systems.

16.2.3 SAP profiles

This section describes how the SAP profiles are assigned and maintained.

Assigning SAP profiles

In order to provide a stable startup procedure, the parameter read sequence (also known as the *parameter replace sequence*) is defined during startup as shown in Figure 16-11.

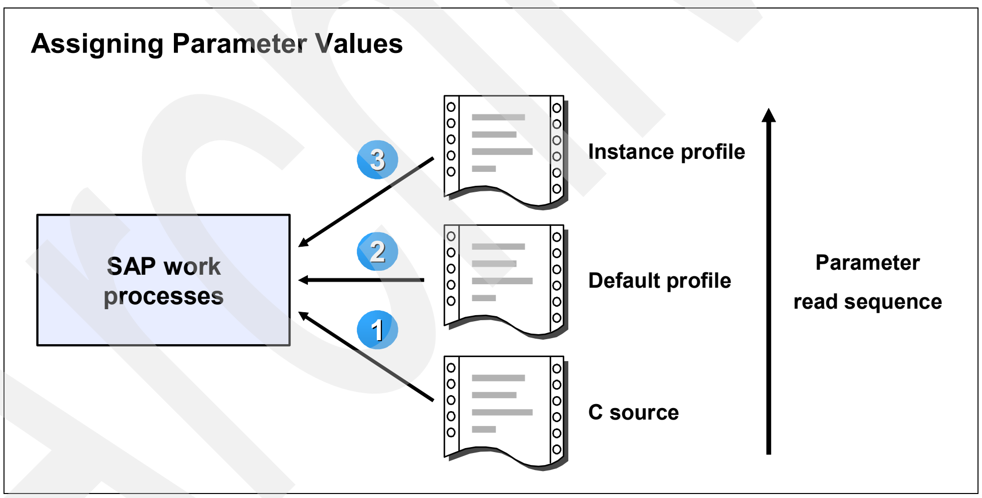

Figure 16-11 Sequence of assigning the SAP profile parameters

The sequence of assigning the SAP profile parameters is as follows:

1. SAP processes read the appropriate parameters from a C program source in the SAP kernel.
2. The default profile /usr/sap/*SID*/SYS/profile/DEFAULT.PFL is read. Profile values already defined in the C source are replaced with the values in the default profile.
3. The instance profile /usr/sap/*SID*/SYS/profile/*SID_INSTANCE_hostname* is read. Profile values already defined in the default profile or in the C source are replaced with the values defined in the instance profile.

This procedure ensures that system parameter values reflect the instance profile and the values in the default profile and the C source.

To display the replace sequence for a particular parameter, execute report **RSPFPAR** in Transaction SE38 or SA38.

The SAP kernel (disp+work) reads the default and instance profile. Therefore, if you change one of these profiles, you must restart the respective SAP instance.

Maintaining the SAP profiles

You should always perform the maintenance of the profile parameters within the SAP system with the transactions RZ10 or RZ11. Only in exceptional situations can you perform the changes on the operating system level in the Integrated File System with i5/OS commands, such as EDTF or STRPDM.

To change profiles belonging to different application servers from a single screen, the profiles have to be stored in the database. However, after installation, profiles are stored at an operating system level. To import them into the SAP system, from the Profile Maintenance window, select **Utilities** → **Import profiles** → **Of active servers**. All three types of profiles are imported and a first check on parameters is performed. The profiles are automatically saved in the SAP database and activated by being written back to the operating system level. If you import single profile files or create profiles, you have to check, save, and activate the profiles manually.

Figure 16-12 shows the processes mentioned above, in other words, how and where to change the SAP profiles:

- Import
- Change
- Save
- Activate

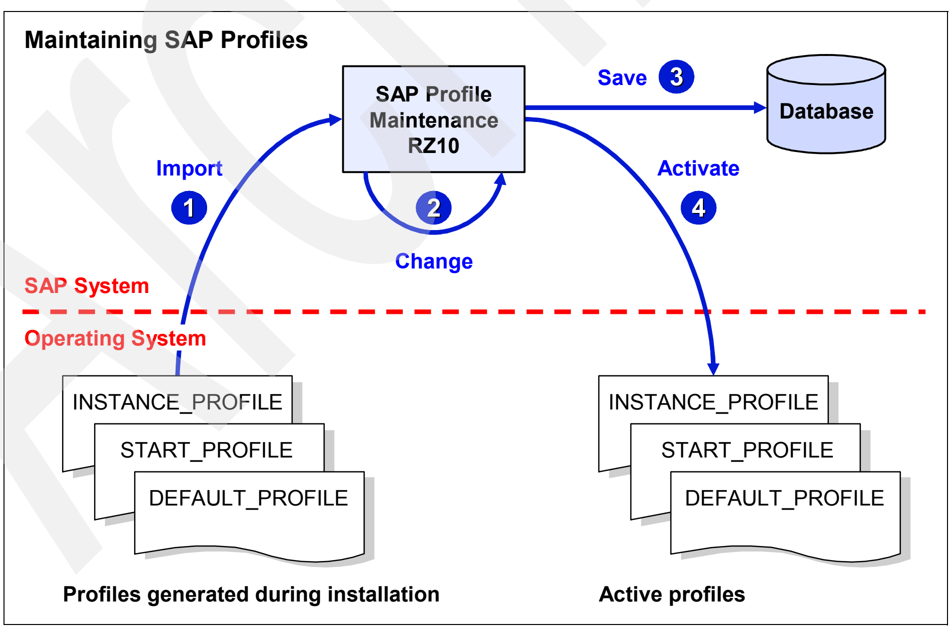

Figure 16-12 Maintaining SAP profiles

You can also create and maintain several profiles in the database under the same name by assigning different version numbers to different files. Each time you save an altered profile a separate version is created. The database thus contains mirrored operating system profile files, old versions, modification histories, and parameter documentation.

The SAP application server is always started using the profile file at the operating system level. From an SAP point of view, a profile consists of two logical parts: entries in database tables and an operating system file that resides in the global profile directory. To activate a profile, write it to the operating system level and restart the SAP system. If other profile files exist with the same name, a confirmation message appears when you activate the profiles from the database to allow the previous files to be overwritten. Additionally, a backup file (.bak) is written.

When activating the profile, a header is inserted in the operating system file containing the name of the profile, the user, who last modified the profile, and the date and time of the change. You can only activate the most recent version of a profile.

16.2.4 SAP background processing

This section provides an overview of the SAP background processing and which tasks to perform.

What is an SAP background job

Background jobs are collections of one or more ABAP programs, external programs, or external commands that run sequentially, without user intervention. Background jobs can:

- Process routine tasks automatically
- Process data collected from core business applications (the existing system)
- Be scheduled to run dynamically based on events occurring in and out of the SAP system
- Process mass data loads at times of low system activity, without affecting online transaction processing

Figure 16-13 on page 269 shows you that you can define background jobs by specifying the following:

- jobname
- jobclass (A = highest, B, C = lowest)
- starting time
- target host (server where the batchjob should be executed)

These attributes for each job are located in a job-scheduling table.

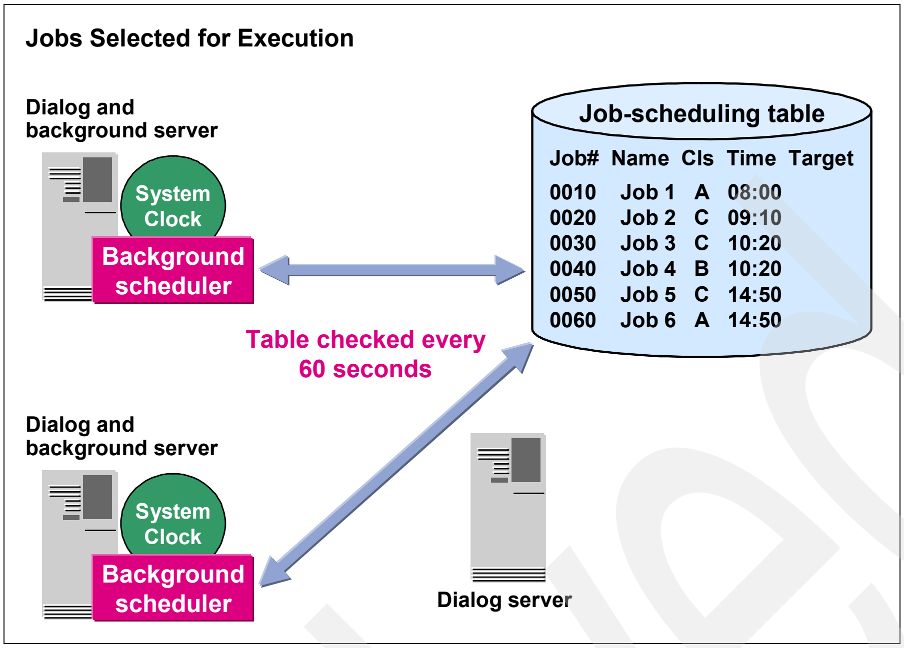

Figure 16-13 SAP background jobs selected for execution

Every server with defined background work processes has a background scheduler active on it. The background scheduler checks the job-scheduling table periodically. Set the rdisp/btctime instance profile parameter to specify the interval between periodic checks. The default value is 60 seconds. Refer to Chapter 13, "SAP instance profile parameters" on page 173 for a discussion of SAP instances.

You can schedule jobs as Immediate, Date/Time, or Event-based. If Immediate or Event-based jobs cannot be executed due to lack of resources, they are placed in the job scheduling table and processed as Date/Time scheduled jobs.

After activation, the background scheduler checks the job-scheduling table, selects jobs that have not yet been executed and have a start time and date already passed, or are scheduled to start immediately. Then the background scheduler gives the control of the process flow to the Dispatcher which assigns a BTC work process.

Configuring a background job

To define a background job, perform the following steps:

1. Select **Tools** → **CCMS** → **Jobs** → **Define**.
2. Specify a job name and job class.
3. Send the spool requests to an SAPoffice user by performing the following steps:
   a. Choose Spool List Recipient.
   b. Enter one of the following in the Recipient field:
      - A user's SAPoffice mail name
      - SAPoffice distribution list
      - SAP user ID
      - An external electronic mail address
   c. Activate any mailing options you want to use, for example, Express, which notifies the recipient as soon as a spool list is available.
   d. Select **Copy** to save the recipient.

All spool requests generated by the job are now sent to this recipient.

4. Select **Steps**, then save the step information to specify the step information.
5. Specify the scheduling information, by selecting **Start Date**, then save it. Save the job data.

Reviewing the background job for errors

To review the background joblog for errors, access **Job Overview** (transaction SM37) and follow the steps as shown in screen that is displayed.

1. Select the job to be viewed.
2. Select **Job log**.
3. Review the joblog for error messages.

If the job shows errors, use Transaction **SM21** to review the SAP system Log (SAPLOG) for details.

Job monitoring - Job overview

The information displayed on the Job Overview screen depends on the selections you have made from the main selection screen (SM37).

When you have jobs scheduled waiting on an event, remember to select the field or **start after event** on the main selection screen. When you have jobs that have not been assigned a start date, select Jobs **without start date** in the main selection screen. When you have jobs that have been scheduled based on a previous job, select Jobs **with previous job** on the main selection screen. From the job overview screen you can also change the status of scheduled jobs.

16.2.5 SAP and DB2 UDB for iSeries data base administration in CCMS

There are only a few tasks involved in DB2 UDB for i5/OS database administration, some of which you can perform using Computing Center Management System (CCMS), as shown in Table 16-1 and Table 16-2. The only required task to keep your system running is to periodically check to see that the system has available disk space to hold the database. A periodic database backup is always recommended. Performance monitoring and tuning are handled with only minimal efforts for most customers, but you may choose to spend more time and resource to optimize your performance, especially if you are developing your own database files and queries.

Table 16-1 General database administration tasks

Area of administration	Can be performed using
Data backup	Operating system
Performance optimization	SAP System (CCMS), operating system
Database check (space)	SAP System (CCMS)
Reorganization	Operating system

Table 16-2 CCMS specific database administration tasks

Area of administration	Can be performed in CCMS using
Performance monitoring	Database monitor (transaction DB4COCKPIT) Alert Monitor (Transaction RZ20)

Area of administration	Can be performed in CCMS using
Database check	State on disk (Transaction DB02)

Computing Center Management System (CCMS) is described in 16.3, “SAP Computing Center Management System (CCMS)” on page 272.

Many functions are performed outside the SAP system, that is, using tools supplied by IBM. Refer to Chapter 25, “Performance concepts and monitoring” on page 457 for a introduction to some of the performance-related operating system and SAP tasks.

16.2.6 Maintaining users and authorities

There are several different users residing outside the SAP system, such as the operating system (OS) user, the database (DB) users, and administrator users. Administrator users (for example, *sid*ofr) and database users (for example, sapr3) must log in at the operating system level as Figure 16-14 illustrates.

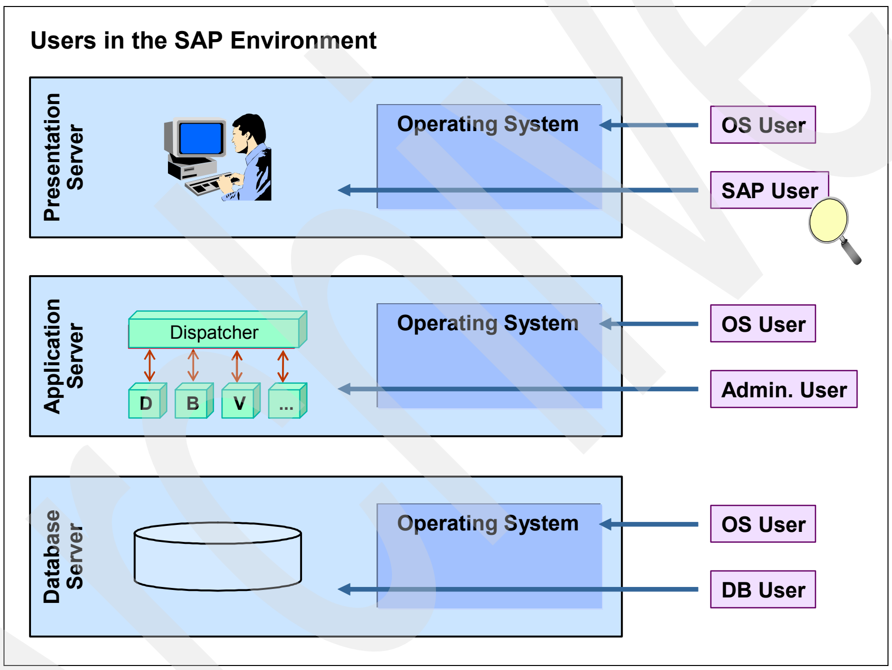

Figure 16-14 Users in the SAP system environment

The SAP user is only known within the SAP system, and it cannot be used to log on to the operating system or the database. Therefore, before accessing the SAP system as an end user, you must first log on to your own operating system.

> **Note:** There is no separate database user on System i models. It is the same system administrator user for the maintenance of i5/OS and the database DB2 UDB for i5/OS.
>
> Refer to 16.1, “SAP system monitoring” on page 252 for more details.

To access SAP data and functions, users must have the appropriate authorization. For example, an SAP user responsible for maintaining user master records has the authorization required to perform tasks in that area, but is not authorized to post a sales document in sales and distribution.

The handling and maintaining of SAP users and the authorities is an important task for the SAP system administrator. Refer to Chapter 10, “Users and authorities” on page 115. For platform independent considerations, refer to the official SAP documentation in the SAP library at:

`http://help.sap.com/saphelp_nw04`

16.3 SAP Computing Center Management System (CCMS)

The SAP CCMS is a monitoring tool which automatically informs you with alerts for the situations you have set up on an alert monitor. Thus, the system, not the administrator, performs the regular checks. This is more efficient than performing the checks manually at a specific time only.

16.3.1 CCMS overview

CCMS is an integral part of the SAP application base. CCMS provides tools for managing the following:

- The SAP system and performance
- Database and archiving
- Workload
- Output
- Security

CCMS can be used to analyze and distribute client workloads and report on resource consumption for system components. CCMS also provides graphical monitors and management utilities.

CCMS provides 24-hour unattended system management functions from within the SAP system through operation modes and instances. CCMS allows you to monitor, control, and configure SAP and provides functions for:

- Profile maintenance
- Unattended 24-hour system management using operation modes, instance definitions, and scheduling
- Starting and stopping instances
- Processing and controlling background jobs, and scheduling database backups
- Automatic reporting of system alerts
- Dynamic logon load balancing
- System and network monitoring and analysis

16.3.2 Setting up CCMS

CCMS must be set up correctly for your environment. Consistency and accuracy of CCMS functioning depend primarily on how it is initially configured by the customer.

Perform the following steps to set up CCMS:

- Maintain SAP profiles. Normally, you import the profiles of all active servers after installation. They are then automatically saved to the database and activated.
- Define at least one operation mode.
- Generate instance definitions for the instances created during SAP system installation.

- Assign instance definitions to operation modes, if necessary, and adapt the work process distribution.
- Define the timetable for normal operation for a full 24-hour cycle.

If operation modes, instances, or timetables are not correctly defined, the CCMS does not display meaningful data.

For more details, refer to the SAP library at:

`http://help.sap.com`

16.4 System Administration Assistant (SSAA transaction)

The *System Administration Assistant* is a hierarchical tree structure that groups important system administration tasks into logical areas and sorts them according to their periodicity. The tasks and their related documentation are at the ends of the hierarchical tree structure.

The System Administration Assistant simplifies the job of the administrator and other members of the SAP implementation or customizing project by providing a preconfigured view of the administration, development and customizing processes. The administrator (or project team member) does most of their work with the System Administration Assistant.

If you want to add customer-specific tasks to the System Administration Assistant, you need to maintain the structure. In this context, bear in mind the customer-specific structure.

When you display the System Administration Assistant, you also display status information. This information will help you carry out the tasks defined in the structure. You can display detail information by choosing different views.

Before you start working with the System Administration Assistant, you need to make some basic settings. For example, for one-system and two-system landscapes, you must perform a conversion.

16.4.1 Using keyword search to find tasks

The System Administration Assistant (transaction SSAA) contains many tasks that are classified in different categories. You can search for tasks using keywords, in the following manner:

1. Select **Utilities** → **Keyword search**.

   All defined keywords are displayed.

2. You can use generic searches to further reduce the list.
3. Select a keyword (or generic part of a keyword). The list expands and all results are displayed in a different color.

   The keyword does not need to be part of the task name. For example, you can use the keyword `user administration` to find the task `User: Copy user`.

16.4.2 Maintaining the System Administration Assistant

To define an own subtree, the so-called *customer subtree*, below the first node in the System Administration Assistant, add new customer tasks at the end of the list in the following steps:

1. Copy SAP tasks to the customer subtree if necessary.

Tasks inherit some attributes of the father node (such as period, function, operating system dependence, database dependence, and system dependence) if tasks are added or copied. These attributes define the structure of subtrees.

You cannot modify SAP tasks in customer systems.

2. An entry for a customer-specific subtree is offered automatically.

   You can define your own customer subtrees depending on customer namespaces.

16.4.3 Defining a new task

You can define a task by defining the global settings, such as:

- Description

  Each task must have a description.

- Documentation

  You can attach documentation to a task (for example, in WinWord or SAPscript). Make sure that the path of WinWord files is accessible.

- Transaction or function module

  Attach a transaction or function module if you want to start a task directly.

- Period

  Define a period (for example, daily or weekly) for getting the status information.

- Client

  Define the client if some tasks must be executed in a special client. For example, transaction SPAM can only be executed in Client 000.

- Authorization

  Define the authorization (administrator, project leader, project member) for access permission.

Dependencies are defined by attributes such as:

- Function

  The function attribute helps to define a logical structure. You can select a function from administrative function, customizing and development, installation function, and so on.

- System landscape

  You might need to distinguish between one-system landscapes and multi-system landscapes for some tasks. The System Landscape attribute can be set to the following:

  - Documentation dependent

    The Knowledge Warehouse includes different documentation for one-system landscapes and multi-system landscapes.

  - Not in One System landscape

    This indicates that the task is not displayed in a one-system landscape.

- Ready-to-Run

  The Ready-to-Run attribute is set by SAP to define a task for Ready-to-Run systems only.

- Operating System

  Tasks that are not valid for the current operating system are not shown in the tree.

- Database

Tasks that are not valid for the current database are not shown in the tree.

- System

  Specify the system if you only need a task for a specific SAP system.

16.4.4 Configuring a system landscape

Configure the system landscape in the development system and transport the settings to all other systems. Considerations include:

- All systems configured by the Transport Management System (TMS) can be administered by the System Administration Assistant.
- You can set the System Type attribute to Development or Production. The System Type determines which set of tasks (SAP templates) should be copied to the customer subtree.
- The Use attribute can be set to one of the following:
  - Development

    The system for modifying the System Administration Assistant
  - Master

    The central system for administering the system landscape

    It allows remote access to all other systems. Status information is saved both locally and in the master system.
  - Inactive

    The system is not visible from other systems.

    The system can see the other systems.

    Only local administration is possible.
  - Standalone

    The system is not visible from other systems.

    The system cannot see other systems.

    Only local administration is possible.
- The SAP tasks (the administrative functions) for that system type are copied to the customer subtree after configuration, and the system name is replaced.

16.4.5 Accepting system landscape settings

Perform the following activities to accept the system landscape settings or to modify them:

- Accept the system landscape in each system after transport of the modified System Administration Assistant and the configured system landscape. The system prompts you automatically to accept when you call the System Administration Assistant for the first time.
- You can overwrite the configuration if one of the target systems is not yet actively used.
- Select **Utilities** → **Check System landscape** to modify the settings after accepting the system landscape.

16.4.6 Allowing remote access

To administer a multi-tier system landscape, you have to allow remote access from a *leading* SSAA SAP system to the other SAP systems in the landscape in the following way:

1. Overwrite the configured settings for Remote access and Security Popup after accepting the configured system landscape.
2. Allow remote access to a selected system.
3. You can choose to send a security pop-up before remote access if remote access is allowed.
4. Allow remote access from the production system to the development system.
5. Mark one system as the master system if you have a system landscape with three or more systems. Administer all systems from the master system. All status information is saved in both the local system and the master system.

16.5 SAP Solution Manager

The SAP Solution Manager is part of SAP NetWeaver. It is designed to manage your entire SAP application.

The Solution Manager functions can be described in the following five scenarios:

- Implementation

  The SAP Solution Manager gives you access to the implementation tools, content, and methodology useful to implement and optimize your SAP system, from both a functional and technical perspective.

- Solution Monitoring

  The SAP Solution Manager helps you meet the performance and availability expectations with system monitoring which monitors the entire business process across multiple components. The following monitoring functions are available:

  – SAP EarlyWatch Alert
  – System monitoring
  – Interface monitoring
  – Business process monitoring
  – Service level reporting

- Operations

  The operations area of the SAP Solution Manager provides access to SAP Support Services, including remote services, on-site services, self-services, and best practice documents. Based on your system configuration, recommendations for these services are triggered dynamically.

- Support Area

  The service desk of the SAP Solution Manager offers a complete infrastructure for organizing and operating a system-wide support organization at your site. This promotes effective user support in your SAP system landscape every step of the way, from the end-users to your internal support organization and, when necessary, to SAP. Solution Manager Diagnostics provides all functionality to centrally analyze and monitor a complete SAP NetWeaver system landscape.

- Upgrade

  SAP Solution Manager helps you to prepare, structure, and execute your upgrades. It provides a portfolio of methodologies, tools and services to ease upgrades of SAP components.

16.5.1 Monitoring an SAP system

The continually increasing number of information technology (IT) solutions and components is a growing challenge for the administration team in a business computing center. The number of components has increased from what was required with SAP (including SAP instances, database, hardware and operating system) to include an ever-growing range of technologies, including products where SAP is a reseller and not a manufacturer. As a result, IT managers require a monitoring method that is centralized with all the information in one tool, yet can be extended to cover new components.

Instead of the classic system monitoring of individual system components, we now discuss system monitoring, where entire business processes can be monitored as a whole across multiple components. This concept is realized through the following three monitoring sections within the SAP Solution Manager:

- Business Process Monitoring

  Real-time monitoring with a focus on the graphical display of business processes based on the Computer Center Management System
- System Monitoring

  Central System Administration and Real-time monitoring based on the Computer Center Management System
- Service Level Management

  Continuous monitoring based on SAP EarlyWatch Alert

Figure 16-15 shows you the processes by which the Solution Manager can monitor the SAP systems.

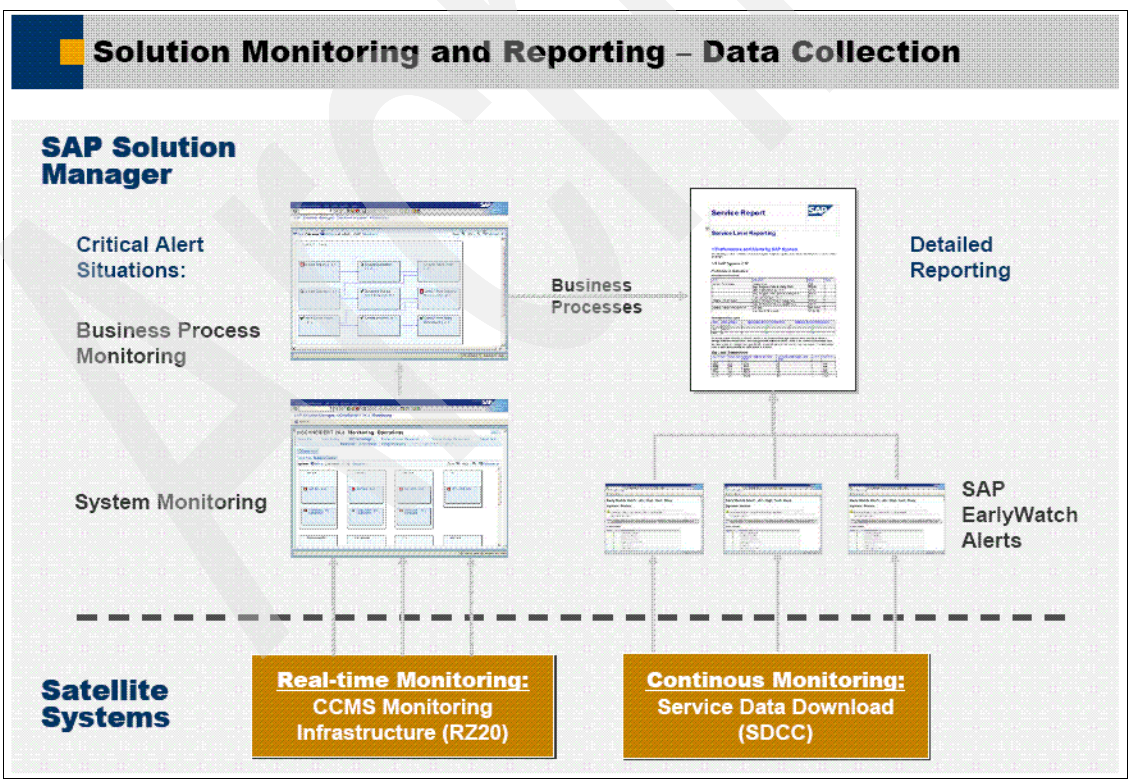

Figure 16-15 SAP Solution Manager - Monitoring, Reporting and Data Collection

There are two general classes of monitoring:

- Continuous monitoring

  Reports inform you on a periodic basis about your SAP systems status and of all your systems (SAP and non-SAP systems) connected to the SAP Solution Manager. In order to get all the information you require, you should define and customize the following:

  – The period of time you require
  – The range of transaction and reports you require
  – The threshold and level of detail you require

- Real-time monitoring

  Here the system should inform you through various methods (such as messages, e-mail or others) when a critical situation arises. In the run-up, you have to define what is a critical situation and when it should be triggered.

Additionally, the SAP Solution Manager gives you a detailed report informing you about the following:

- System information of the connected SAP and non-SAP systems. (For example, you can see the patch levels of all connected SAP systems in these reports.)
- Information from the continuous monitoring
- Information from critical situations of the real-time monitoring

You can use these reports for Service Level Management (SLM) and to set up Service Level Agreements (SLAs).

Monitoring your entire IT landscape efficiently

Managing distributed mySAP.com applications or client-server systems places high demands on the system administrator, demands that increase significantly through the inclusion of the Internet in business processes. Using the Monitoring Architecture of CCMS, SAP provides a flexible and universally applicable infrastructure, with which you can monitor the entire IT landscape centrally, and which quickly and reliably reports problems that occur.

The monitoring architecture is delivered with every SAP Web Application Server with no additional costs and can be used immediately after installation. The architecture runs on all SAP Web Application Servers and can easily be extended with SAP and non-SAP components.

Using alerts

- Alerts are a central element of the monitoring. Alerts quickly and reliably report problems, such as a value exceeding or falling below a particular threshold value, or an IT component remaining inactive for a defined period of time. To improve clarity, alerts are assigned certain colors: yellow for a warning, red for a problem, and a numerical value for the criticality of the problem

There are two products available to display the collected data:

- Measured values and alerts are displayed in the Alert Monitor in various monitors, organized by topic. These monitors display the data in a tree structure in which the highest alert of a subtree is reported upward.
- The SAP Solution Manager is focussed on the technical integration of applications, Software Change Management, and monitoring the most important business processes of the customer.

Analysis and auto-reaction methods

You can assign analysis and auto-reaction methods to the alerts. These contribute to quicker processing of the problem. By double-clicking on an alert, you manually start the analysis method (for example, the job management transaction for a background job that has terminated prematurely). An auto-reaction method, on the other hand, starts automatically as soon as the alert occurs. These include the execution of operating system commands and sending an e-mail or an SMS message to the system administration.

Figure 16-16 shows you the data and process flow of the alert monitoring of the SAP Solution manager. In Figure 16-16, CEN means Central CCMS and ALM means Alert Management.

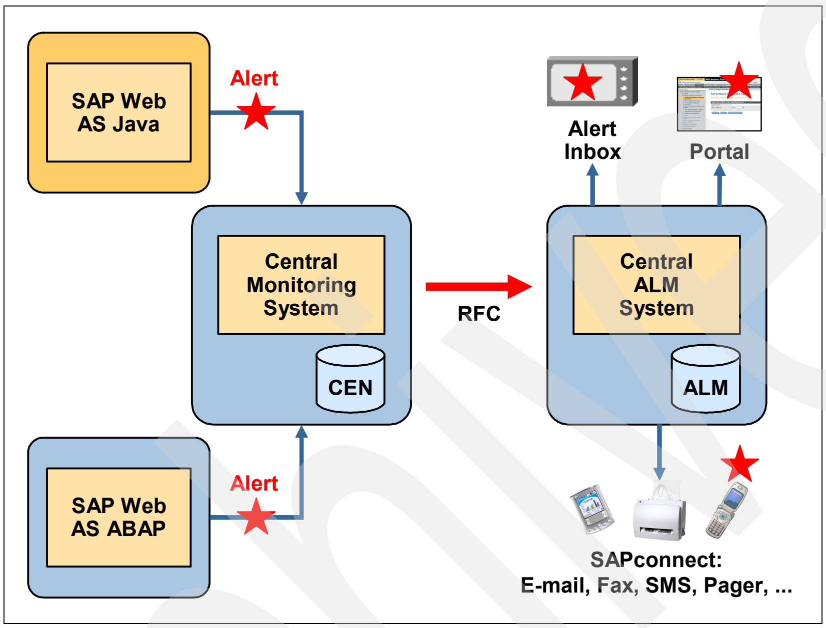

Figure 16-16 SAP Solution Manager - Alert Monitoring

Individual configuration

Threshold values, methods, detailed help, and four extensive monitor sets with monitors for all aspects of system management are predefined in the monitoring architecture and are available in every SAP system. You can also customize all settings to your own requirements and, for example, configure customer-specific monitors so that only the data about your system landscape that is relevant to your needs is displayed.

Communication using agents

Remote monitored components are included using CCMS agents. These have the following properties:

- Connection to the central monitoring system using Remote Function Call (RFC) to provide greater downtime security and universal usability
- Use of a push mechanism to proactively report alerts and data to the central system
- Integration of the SAPOSCOL operating system collector to transfer data for the operating system level, from CPU utilization and fill levels of file systems to details of individual monitored processes
- Connection of systems with SAP Basis 3.X and components on which no SAP Web Application Server is installed
- Monitoring of any log files

SAP note 522878 describes some specifics to CCMS agents and the availability of CCMS agents on the System i server. There are also additional EBCDIC and ASCII considerations. You can download the CCMS agents and tools from the SAP Software Distribution Center at:

http://service.sap.com/swdc

Reporting on landscape-wide performance data

With SAP NetWeaver ´04, highly sophisticated and cost-efficient performance reporting is possible using SAP Business Intelligence (BI). BI features powerful analysis tools for performing correlations and determining trends with flexible and user-definable reporting. Performance and availability data stored in the Central Performance History is transferred to a BI system once a day.

Service Level Management

SLM is the discipline of using proactive methodology and procedures to ensure that adequate levels of service are delivered to all IT users in accordance with business priorities and at an acceptable cost. SLM defines the interface between IT departments and functional departments. It refers to systems and business processes, and involves the customer support desk, SAP, and partners.

SLM introduces clearly defined, measurable goals and communication structures that:

- Improve the IT department's understanding of Business Process Owner (BPO) requirements
- Improve the quality of the IT service
- Clarify the current cost structure
- Facilitate forecasting

With the Service Level Reporting functionality in the SAP Solution Manager you are able measure your key performance indicators (KPIs) of your mySAP application and communicate the agreed service goals. Service Level Reporting is the database to control and prove the keeping of your Service Level Agreements.

SLAs are created between the BPO and the IT organization and consists of the three agreement parts:

- Identify together the Business Process Owner requirements for IT service quality:
  - Availability
  - Performance
  - Correctness
  - Security
- Agree to clear, measurable service goals.
- Agree to a clear communication structure for achieving those goals.

The Service Level Reporting function in the SAP Solution Manager allows you to report on the performance of all systems (that belong to a system landscape), business processes and the support process regularly, for example, on a weekly basis.

By setting clearly defined and verifiable goals, Service Level Reporting enables clear communication between the operators of an e-business solution and the customers of the solution, that is, the owners of the business processes. They can range from server to database operators to network operators and so on). Service Level Reporting reports on how goals are being met within a defined period of time. Requirements include:

- Service Level Reporting is generated once a week and is oriented to various management areas (IT and business department management). For this reason, it is simple and does not need expert knowledge to be set up.
- Service Level Reporting evaluates performance indicators on management level. These action plans can result from recommendations in SAP Support Services, for example, SAP EarlyWatch Check and Solution Management Optimization services. There must be a data connection between Service Level Reporting and these Support Services.
- The action plans are partially configured, yet they can be adapted and expanded to meet the customer's system requirements.
- Service Level Reporting allows the customer to generate forecasts and is connected to the predictive Support Services from SAP, for example, SAP GoingLive Check and SAP GoingLive Functional Upgrade Check.

16.5.2 SAP EarlyWatch Alert

The EarlyWatch Alert is a preventative monitoring method for your SAP application. SAP EarlyWatch Alert is an important part of making sure your core business processes work. It is a monitoring tool that monitors the essential administrative areas of SAP components and keeps you up-to-date on their performance and stability. SAP EarlyWatch Alert runs automatically to keep you informed, so you can handle issues proactively, before they become critical.

SAP EarlyWatch Alert is most effective when you activate it for all SAP components in your system. It is covered by your maintenance agreement with SAP, so you can do so at no extra charge.

Keeping total cost of ownership low and the performance of your SAP system high is a tremendous value to your business, a value delivered by the SAP EarlyWatch Alert. Knowing the status of each SAP component in your system allows you to:

- Greatly minimize the risk of downtime
- React to issues such as bottlenecks before they become critical
- Know what is affecting the performance and stability of your system

The underlying concept of the SAP EarlyWatch Alert is to ensure smooth operation of individual SAP components by taking action proactively, before severe technical problems occur. This monitoring tool is based on experience with thousands of installations giving you a tool to save time, reduce costs, and keep your SAP system running smoothly.

The SAP EarlyWatch Alert is included in your maintenance agreement with SAP. You activate it yourself for each of your SAP components. It is included in the SAP Solution Manager from where you activate it and read the weekly reports. To activate the SAP EarlyWatch Alert, simply follow the instructions given in the SAP Solution Manager Installation Guide.

16.6 Daily administration tasks - sample lists

Table 16-3 on page 282 and Table 16-4 on page 282 show you sample lists for daily administration tasks derived from a real SAP project. You can adapt and use this list for your own purposes.

The daily tasks shown in Table 16-3 on page 282 are from a System i point of view.

Table 16-3 Example daily administration checklist for System i configurations

Task	Transaction or command	How to perform the task
Is there a CPU load on the system? Is the SAP system running?	WRKSYSSTS	Can you logon to the System i server, at least on the system console with the QSECOFR user profile? Is there any CPU utilization?
Is the SAP system up?	WRKACTJOB	Are there jobs with a SEMW status in the System i R3_nn subsystem?
Is TCP/IP running?	WRKACTJOB WRKLIND	Does the i5/OS QSERVER subsystem have at least three jobs (QPWFSERVSD)? Is the ethernet adapter busy with TCP/IP?
Check disk status	WRKDSKSTS	ASP1 should be less than 80% If greater than 90% indicates a critical situation If greater than 95% immediate attention is required ASP2 load should be less than 60% If greater than 60%, immediately save and delete journal receivers If ASP 2 overflows (to ASP 1), save and delete journal receivers from ASP 2 and perform a journal switch
Check System i syslog	DSPMSG QSYSOPR and DSPLOG	Look at answered and not answered messages QHST: Look at “critical” entries (Select time range preceding error in the DSPLOG command)
Is SAProuter running?		Check if the job or SAPROUTER service is running
Check printer and output queues especially for the SAP printer	WRKWTR	Is there any data?
Check the backups		Check carefully Where are the tapes?
Are future backups scheduled correctly?		Check carefully

Table 16-4 Daily SAP system administration checklist (example)

Task	Transaction or command	How to perform the task
Is SAP up?		Can you sign on to the SAP system via SAPLOGON? If not, analyze the situation: WRKLNK to path /usr/sap/*SID*/DVEBMGSnn/work ► File: sapstart.log ► File: dev_ms ► File: dev_disp ► File: dev_ms ► Files: dev_wnnn, nnn= number workprocess

Task	Transaction or command	How to perform the task
Check the work processes	SM50	► Look at the status. The active processes should be `waiting` or `running` ► Are there locks or debugs and traces? ► Are there long running jobs? If yes, compare it with the background job list. ► Are there too many red alerts meaning restart of work processes? If yes, investigate the situation. ► How is the workload? ► Refresh the screen again by pressing **F8.** ► How is the CPU resource? ► Push the CPU button. ► The last tasks of every process should be near to 0:00 seconds.
Check the user list	SM04 AL08	► Are there unknown users? ► Are there unknown work stations? ► How many user and how many sessions are there in total? ► When did the user perform the last action? ► Is the user logged on more than once to the same client? ► How many sessions (external and internal) is the user using?
Check canceled updates	SM13	► Is the update mode running? ► Set user: ***** ► Set update: **all** ► Set the date to 1000 years before. ► Are there error situations? ► Contact the user and discuss each case and decide what to do. ► Are there open V2? ► Are there enough UP2 work processes running? ► Are there open “coll.run” updates (V3)? ► Is there a problem with the job RSM13005? Do not delete updated records without contacting the user. There should be no errors and no problem situations with the updates.

Task	Transaction or command	How to perform the task
Check locks	SM12	Are there old locks (see last entry)? If the lock has a date (no time) entry, the lock is from yesterday or earlier Investigate if there is a job assigned that is still active but which is already cancelled ► Is a lock older than one hour for a dialog process? Contact the user. ► Is there a lock older than two hours for a batch job? Contact the user. Do not delete locks without contacting the user.
Check job list	SM37	Job: * User: * Date: Set to 1000 years earlier ► See if the status of the job is `cancelled`. Investigate and then get in touch with the user. ► Is an active status OK at this time?
Check batch input	SM35	Session: * Date: **today** Created by: * Select: ► incorrect ► locked In processing: ► Is this OK at this time In background: ► Is this OK at this time? Choose the analysis button Contact the user and discuss each situation that you do not understand
Check dumps	ST22	Select the right date. Investigate the list. ► Which time, which user? ► What is the error ID? Are there dumps with the same type? Analyze each dump you do not understand the precise cause for: ► Look at the dump list for hints ► See SAP notes ► Developer traces (dev_wnn, nn=number) – Transaction ST11 Contact the user and discuss each situation that you do not understand.

Task	Transaction or command	How to perform the task
Check syslog	SM21	Date: **today** Time: **00:00:00** Choose all messages Button: **Reread the system log** Are there “red” entries? ► Often it is an known situation ► Analyze any unknown red entry – Dump – SAP OSS – Developer traces
Check spool files	SP01	User: * Date: **today** Client: *empty* Total at the end of the list Number of output requests ► Status: `waiting` ► Status: `error` ► Status: `completed with errors` ► Status: `empty`
Check spool system	SP12	Not necessary every day, only if there is a reason or occasionally Reorganization, consistency
Check workload statistics	ST03	Server *host*, instance nn Select: **today** Press button: **Dialog** Average response time should be less than nnn msec (for example 600) ► average CPU time: about aaa msec (for example 100 - 200 msec) ► average DB time: about bbb msec (for example 200 - 300 msec) ► average GUI, wait, … : about ccc msec (for example 20 - 200 msec) Analyze with: ► detailed menu ► processing time < 2 x CPU time

Task	Transaction or command	How to perform the task
Check buffer	ST02	What is the swapping rate (red fields): ► About 1 - 2000 program buffer swaps are ok ► About 1 - 1000 export / import buffer swaps are ok ► Any other buffer swaps should be avoided Look at extended memory: ► Maximum used memory should always be less than allocated memory ► There should be no heap memory – If there is heap memory perhaps restart SAP or enlarge em/initial_size_MB ► Current memory shows actual situation Be very careful of changing profile parameters. Be sure that you have a backup of the previous file.
Check OS	ST06	Detailed menu: ► Check CPU activation in particular ► CPU activity should not be higher than 60 - 80% for a long time ► View and compare history
Check database	DB02	Detailed menu: ► “Detailed Object Analysis” provides the largest tables – Is it ok? – Is it similar to yesterday? ► “Space Statistics” allows you to see growth rates ► “Missing Indexes” or “Consistency Checks” allow you to find inconsistencies between the database and the dictionary – Are there missing indices ► Check generally the “State on disk” ► Look at other check buttons
Check DB protocols	AL02	Is replaced by RZ20
Check SQL statistics	STAT	Shows and analyze situation of the actual time and earlier ► Select time interval ► Select the number of records Is something unique?
Check CCMS monitor(s)	RZ20	If there is a special monitor set up: ► What about the thresholds? Look at red alerts: ► Especially for the “expert monitor” ► Expand all ► This can give you a lot of informations ► You can execute some check functions

16.7 Useful SAP system administration transactions

Example 16-1 shows a list of useful SAP system administration transactions.

Example 16-1 Useful SAP system administration transactions for System i installations

```
AL03[1] Operating system alert monitor
AL04[1] Monitor call distribution
AL05[1] Monitor current workload
AL08  Users Logged On
AL11  Display SAP Directories
AL12  Display table buffer (Exp. session)
AL13  Display Shared Memory (Expert mode)
AL15  Customize SAPOSCOL destination
AL16[1] Local Alert Monitor for Operat.Syst.
AL17[1] Remote Alert Monitor for Operat. Syst.
AL18[1] Local File System Monitor
AL19[2] Remote File System Monitor
BD64   Maintenance of Distribution Model
BSVW  Linkage Status Update-Workflow Event
CMOD  Enhancements
DB01  Analyze exclusive lock waits
DB02  Analyze tables and indexes
DB11 Maintain Database Integration
DB12  Overview of Backup History
DB15  Data Archiving: Database Tables
DBCO  Database Connection Maintenance
FILE  Cross-Client File Names/Paths
NACE  WFMC: Initial Customizing Screen
OAA1  SAP ArchiveLink: Maint.user st.syst
OAA3  SAP ArchiveLink protocols
OAA4  SAP ArchiveLink applic.maintenance
OAAD  ArchiveLink Administration Documents
OAC2  SAP ArchiveLink: Globaldoc. types
OAC5  SAP ArchiveLink: Bar code entry
OACA  SAP ArchiveLink workflow parameters
OAD0  SAP ArchiveLink: Objectlinks
OAD2  SAP ArchiveLink document classes
OAD3  SAP ArchiveLink: Link tables
OAD4  SAP ArchiveLink: Bar code types
OAD5  SAP ArchiveLink: Customizing Wizard
OADR  SAP ArchiveLink: Print list search
OAM1  SAP ArchiveLink: Monitoring
OAOR  SAP ArchiveLink: Stored documents
OARE  SAP ArchiveLink:St.syst.return codes
OS01  LAN check with ping
OS03  Operating System Parameter changes
OS04  Local System Configuration
OS05  Remote System Configuration
OS06  Local Operating System Activity
OS07  Remote Operating System Activity
OSS1  Logon to Online Service System
OY18  Table history
PFCG  Activity Group
PFUD  Authorization Profile comparison
RZ01  Job Scheduling Monitor
RZ02[1] Network Graphics for SAP Instances
RZ03  Presentation, Control SAP Instances
RZ04  Maintain SAP Instances
RZ06[1] Alerts Thresholds Maintenance
```

RZ08[1] SAP Alert Monitor
RZ10 Maintenance of profile parameters
RZ11 Profile parameter maintenance
RZ12 Maintain RFC Server Group Assignment
RZ20 CCMS Monitoring
RZ21 Customize CCMS Alert Monitor
SA38 ABAP/4 Reporting
SADC Address: Maint. communication types
SALE Display ALE Customizing
SAINT Plug-in Installation
SARI Archive Information System
SARA Archive management
SARP Reporting (Tree Structure): Execute
SART Display Reporting Tree
SB01 Business Navigator - Component View
SC38 Start Report Immediately
SCAT Computer Aided Test Tool
SCC1 Client Copy - Special Selections
SCC2 Client transport
SCC3 Client Copy Log
SCC4 Client administration
SCC5 Client Delete
SCC7 Client Import – Post Processing
SCC8 Client Export
SCC9 Remote Client Copy
SCCL Local Client Copy
SCDO Display Change Document Objects
SCMP View / Table Comparison
SCON SAPconnect - Administration
SCUA Central User Administration: Distribution Model Assignment
SCUG Central User Administration Structure Display
SCUL Central User Administration Log
SCUM Central User Administration Field Selection
SCU0 Table Analyses And Comparison
SCU3 Table History
SD11 Data Modeler
SDBE Explain an SQL Statement
SECR Audit Information System
SE01 Transport and Correction System
SE03 Transport Utilities
SE06 Set up Workbench Organizer
SE07 Transport System Status Display
SE09 Workbench Organizer (Initial Screen)
SE10 Customizing Organizer
SE11 Data Dictionary Maintenance
SE12 Data Dictionary Display
SE13 Maintain Technical Settings (Tables)
SE14 Convert Data Dictionary tables on Database Level
SE15 Repository Info System
SE16 Display Table Content
SE17 Generate Table Display
SE30 ABAP Objects Runtime Analysis
SE32 ABAP Text Element Maintenance
SE33 Context Builder
SE35 ABAP/4 Dialog Modules
SE36 Logical databases
SE37 ABAP Function Modules
SE38 ABAP Editor
SE39 Splitscreen Editor: Program Compare
SE40 MP: Standards Maintentance and Translation

SE41 Menu Painter
SE43 Maintain Area Menu
SE51 Screen Painter
SE54 Generate table view
SE55 Internal table view maintenance call
SE56 internal call: display table view
SE57 internal delete table view call
SE61 Document worklist
SE62 Industry Utilities
SE63 Translation: Initial Screen
SE71 SAPscript form
SE72 SAPscript Styles
SE73 SAPscript font maintenance (revised)
SE74 SAPscript format conversion
SE75 SAPscript Settings
SE76 SAPscript: Form Translation
SE77 SAPscript Translation Styles
SE78 SAPscript: Graphics administration
SE80 Object Navigator
SE81 Application Hierarchy
SE82 Application Hierarchy
SE84 Object Navigator
SE85 ABAP/4 Repository Information System
SE89 Maintain Trees in Information System
SE91 Maintain Messages
SE92 New SysLog Msg Maintenance as of 46A
SE93 Maintain Transaction Codes
SE94 Customer enhancement simulation
SE95 Modification Browser
SEPS SAP Electronic Parcel Service
SERP Reporting: Change Tree Structure
SF01 Client-Specific File Names
SFAW Field Selection Maintenance
SIAC[1] Web Object Administration
SHDB Record Batch Input
SICK Installation Check
SIN1 SAPBPT: Inbox
SINA SAPBPT: Maintain Standard Configuration
SLG0 Application Log: Object Maintenance
SLIN ABAP: Extended Program Check
SM01 Lock Transactions
SM02 System Messages
SM04 User Overview
SM12 Display and Delete Locks
SM13 Display Update Records
SM14 Update Program Administration
SM21 System log
SM28 Installation Check
SM29 Model Transfer for Tables
SM30 Call Up View Maintenance
SM31 Table maintenance
SM31 Old Table Maintenance
SM32 Maintain Table Parameter ID TAB
SM33 Display Table Parameter ID TAB
SM34 View cluster maintenance call
SM35 Batch Input Monitoring
SM36 Batch request
SM37 Background job overview
SM38 Queue Maintenance Transaction
SM49 Execute Logical Commands

```
SM50  Work Process Overview
SM51  List of SAP Servers
SM54  TXCOM maintenance
SM55  THOST maintenance
SM56  Number Range Buffer
SM58  Asynchronous RFC Error Log
SM59  RFC Destinations (Display/Maintain)
SM61   Display Control Object List
SM62   Display / Edit Events
SM63  Display/Maintain Operating Mode Sets
SM64  Release of an event
SM65  Background Processing Analysis Tool
SM66  System-wide Work Process Overview
SM69  Display/Maintain Logical Commands
SMGW  Gateway Monitor
SMLG  Maintain Logon Group
SMLT  Language transport utility
SMOD  SAP Enhancement Management
SMT1  Trusted Systems (Display <-> Maint.)
SMT2  Trusting systems (Display <->Maint.)
SMW0  SAP Web Repository
SMX   Display Own Jobs
SNRO  Number Range Objects
SOBJ  Attribute Maintenance Objects
SOLE  OLE Applications
SOPE  Exclude Document Classes
SOTD  SAPoffice: Maintain Object Types
SOY1  SAPoffice: Mass Maintenance Users
SOY2  SAPoffice: Statistics data collection
SOY3  SAPoffice: Statistics Evaluation
SOY4  SAPoffice: Access overview
SOY5  SAPoffice: Inbox overview
SOY6  SAPoffice: Document overview
SOY7  SAPoffice: Folder overview
SOY8  SAPoffice: Mass Archiving
SOY9  SAPoffice: Inbox Reorganization
SOYA  SAPoffice: Change folder owner
SP00  Spool and Relate Area
SP01  Spool Control
SP02  Display output Requests
SP11  TemSe Contents
SP12  TemSe Administration
SPAD  Spool Management
SPAM  SAP Patch Manager (SPAM)
SPAU  Display Modified DE Objects
SPCC  Spool Consistency check
SPDD  Display Modified DDIC objects
SPHA  Telephony administration
SPIC  Spool : Installation Check
SPRM  Current Customizing
SPRO  Customizing
SQ01  SAP Query: Maintain queries
SQ02  SAP Query: Maintain funct. areas
SQ03  SAP Query: Maintain user groups
SQ07  SAP Query: Language comparison
SQVI  QuickViewer
SSAA  System Administration Assistant
SSCA  Appointment Diary: Administration
SRZL  CCMS
SSM1  Session Manager generation call
```

```
ST01  System Trace
ST02  Setups/Tune Buffers
ST03  Performance, SAP Statistics, Workload
ST04 iSeries Performance Monitor Overview (is going to be replaced by DB4COCKPIT)
ST05  SQL Trace
ST06  Operating System Monitor
ST07  Application monitor
ST10  Table Call Statistics
ST11  Display Developer Traces
ST14  Application Analysis
ST22  ABAP Runtime Error Analysis
ST22  ABAP/4 Runtime Error Analysis
ST62  Create industry short texts
STAD   Select Statistical Records
STMS  Transport Management System
STZAC Maintain time zone act.in client
STZAD Disp.time zone activat.in client
SU01  Maintain User
SU01D Display users
SU02  Maintain Authorization Profiles
SU03  Maintain Authorizations
SU05  Maintain Internet Users
SU10  Mass changes to User Master
SU12  Mass Changes to User Master Records
SU2   Maintain User Parameter
SU20  Maintain Authorization Fields
SU21  Maintain Authorization Objects
SU22  Auth. object usage in transactions
SU24  Disables Authorization Checks
SU25  Imports SAP Check Indicators defaults
SU26  Adjust Authorization checks
SU52  Maintain own user parameters
SU53  Display check values
SU56  Analyze User Buffer
SUPC  Profiles for activity groups
SUPO  Maintain Organization Levels
SUIM  Repository Info System
SWDC  Workflow Definition
SXDA  Data Transfer Workbench
TU02  Display Active Parameters
USMM  Customer measurement

1 = replaced by transaction RZ20
2 = replaced by transaction OS07
```

Start the transaction within the SAP application and press the **F1** Help key to see an explanation of each transaction.

Start the transaction by a prefix “/n” in front of each transaction name, for example **/nsu01** for the transaction **SU01** (Maintain User).

17

General topics

This chapter presents background and summary information for topics and procedures not specifically represented elsewhere in this book. It also covers recommendations for daily work with SAP applications with i5/OS systems. We discuss some issues which often arise with SAP applications in System i installations. We focus in the following specific points:

- SAProuter with the saprouttab
- Automated startup and IPL
- Working with SAP media

17.1 Control network access with SAProuter

SAProuter is an SAP program that acts to protect your SAP network against unauthorized access. It provides the means of controlling access to your SAP system. SAProuter acts as an intermediate station (a proxy) in a network connection between SAP systems, or between SAP systems and external networks. SAProuter controls the access to the network (the application level gateway), and is a useful enhancement to an existing firewall system (such as a port filter).

You can find further information in the SAP “SAProuter Documentation” on the following Web sites:

- `http://service.sap.com/saprouter`
- `http://service.sap.com/access-support`

Refer to the illustration in Figure 17-1 to understand the SAProuter as a firewall for the SAP systems.

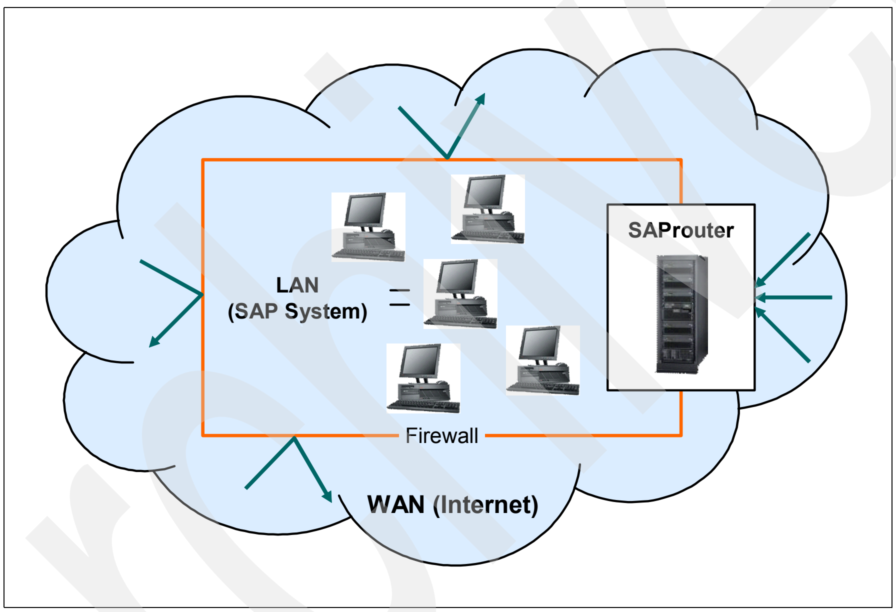

Figure 17-1 SAProuter as a firewall

Figuratively, the firewall forms an impenetrable wall around your network. However, since particular types of connections need to penetrate this wall, an entry through the firewall is required. SAProuter assumes the control of this access.

You should use SAProuter to:

- Control and log the connections to your SAP system, for example, from an SAP service center.
- Set up an indirect connection when programs involved in the connection cannot communicate with each other due to the network configuration, in the event that:
  - Addresses conflict when using non-registered IP addresses
  - Restrictions exist for firewall systems
- Improve network security by means of the following:
  - Using a password to protect your connection and data from unauthorized external access

  – Allowing access from only particular SAProuters
  – Allowing only encrypted connections from a known partner (using the Secure Network Communication (SNC) layer)

- Improve performance and stability by reducing the SAP system load within a local area network (LAN) when communicating with a wide area network (WAN).

Installing the SAProuter

Follow these steps to install the SAProuter on a System i server. Refer to *SAP note 86329* for further information:

1. Import the latest version of sapservX from the general/misc/saprouter directory. Read the associated README file. Import the programs SAProuter and niping into a separate library (for example, SAPROUTER).
2. Log in to the System i server as *SID*OFR.
3. Enter the following commands:
   a. `CRTLIB SAPROUTER`
   b. `CRTSAVF SAPROUTER/SAPROUTER`
   c. `CRTSAVF SAPROUTER/NIPING`
   d. `ftp sapservX`
   – `cd general/misc/saprouter`
   – `lcd SAPROUTER`
   – `bin`
   – `get saprouter.`*Version.Plattform* `SAPROUTER (replace`
   – `get niping.`*Version.Plattform* `NIPING (replace`
   – `quit`
   e. `RSTOBJ OBJ(*ALL) SAVLIB(SAPROUTER) DEV(*SAVF) SAVF(SAPROUTER/SAPROUTER) RSTLIB(SAPROUTER)`
   f. `RSTOBJ OBJ(*ALL) SAVLIB(SAPROUTER) DEV(*SAVF) SAVF(SAPROUTER/NIPING) RSTLIB(SAPROUTER)`

Note: Installing SAProuter without the use of a firewall does not protect your network against access from external networks. You must ensure that all incoming connections go through the SAProuter hole.

Use the SAProuter from the current 6.40 kernel for new installations. Figure 17-2 on page 296 identifies how to find the patch on the SAP Software Distribution Center Web site:

http://service.sap.com/swdc

Select **Search for all Categories** on the left side of the page.

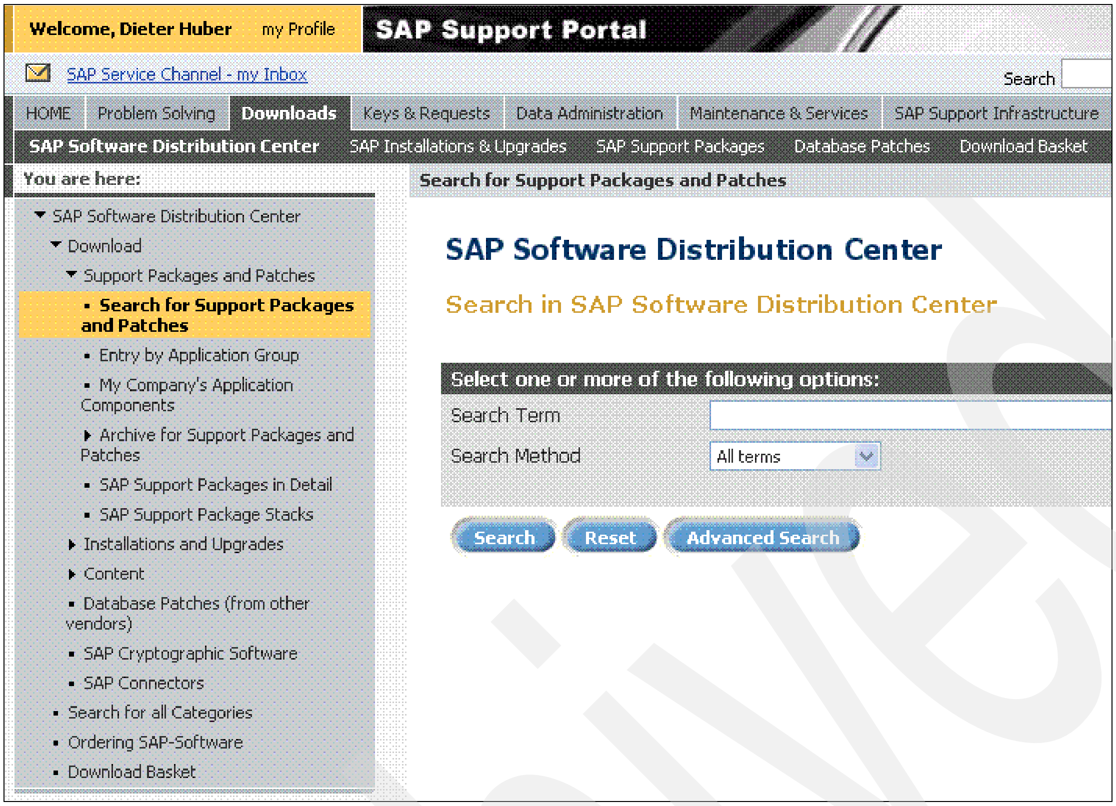

Figure 17-2 How to find and get the SAProuter patch from SAP Service Marketplace

Perform the following steps to apply the current SAProuter patch:

1. Create the **/usr/sap/saprouter** directory.

   Use the Add Link command to create a link to saprouttab in the root directory as follows:

   ```
   ADDLNK OBJ('/usr/sap/saprouter/saprouttab') NEWLNK('/saprouttab')
   ```

2. The corresponding routing table must be maintained by entering the following command:

   ```
   EDTF ’/usr/sap/saprouter/saprouttab’
   ```

   An example of a routing table can be found in this directory on sapservX.

3. Enter the following commands to start the SAProuter:

   a. `ADDLIBLE LIB(SAPROUTER) POSITION(*FIRST)`

      Then, to use `SAPROUTER` from the kernel library, enter the following command:

      ```
      CALL R3sid400/R3INLPGM
      ```

   b. ```
      SBMJOB CMD(CALL PGM(SAPROUTER) PARM('-r' '-R'
      '/usr/sap/saprouter/saprouttab' '-W' '60000')) JOB(SAPROUTER)
      JOBQ(QSYSNOMAX) INLLIBL(*CURRENT) CPYENVVAR(*YES) USER(sidOFR)
      ```

Refer to *SAP note 86329* for more information.

## SAProuter options

You can call the SAProuter with multiple commands.
      ```

Note: Calling saprouter without parameters displays help text, within which a sample saprouttab is contained.

Table 17-1 identifies the SAProuter options.

Table 17-1 SAProuter options

Option	Meaning
Option -s (stop saprouter)	Stop SAProuter
Option -n (new saprouttab)	Reread in the Route Permission Table
Option -t (toggle trace)	Changing the trace level
Option -c*n* (cancel connection n)	Terminate connection *n*
Option -l / -L	Display route information
Option -d (dump buffers)	Write detailed information from the internal buffers to the trace file
Option -f (flush buffers)	Reset internal buffers
Option -a *lib*	Use external library
Option -p	Carry out soft shutdown

Table 17-2 shows an extract of additional saprouter options.

Table 17-2 Additional saprouter options

Option	Meaning	Default
Option -R*saprouttab*	File name and path of the Route Permission Table	./saprouttab
Option -G*logfile*	Name and path of the SAProuter log file	No log file
Option -T*tracefile*	Name and path of the SAProuter trace file	dev_rout in the directory of the SAProuter
Option -V*tracelevel*	Trace level of the SAProuter	1
Option -S *service*	Service (port) on which the SAProuter runs	3299
Option -W *waittimeL*	Timeout for blocking network calls (if there is an error)	5000 msec

Options must be specified in quotation marks under i5/OS. For example, to stop the SAProuter, enter `saprouter '-s'`.

Testing SAProuter basic functions

Test for network problems before using SAProuter. The saprouter and niping programs are required to perform this test, as well as three open windows (shells) on one or more hosts.

Table 17-3 on page 298 shows the test scenario set when using niping.

SAProuter runs in window 1, the server runs in window 2, and the client runs in window 3.

Table 17-3 System i niping test

	Windows 2 (host2)	Window 1 (host1)	Window 3 (host3)
Without SAProuter	`call niping '-s'`		`call niping '-c' '-H' 'host2'`
With SAProuter	`call niping '-s'`	`saprouter '-r'`	`call niping '-c' '-H' '/H/host1/H/host2'`

The SAP program niping is provided by SAP to test the SAP Network Interface (NI) of SAProuter and other SAP connections with the SAP internal communications protocol. Perform the following steps to run the niping test:

1. Enter the following command to start SAProuter in window 1 (on host1):

   ```
   saprouter '-r'
   ```

   This command calls SAProuter without any parameters.

   > **Note:** For a complete list of the SAProuter commands, refer to the SAProuter Options or the online help. Refer to the left side of the *SAProuter Documentation* (an Adobe PDF) at the following Web site:
   >
   > http://service.sap.com/saprouter

2. In window 2 (host2), start the niping test program to emulate a test server. Enter the following command:

   ```
   call niping '-s'
   ```

   For a complete list of the niping commands, refer to the online help. To refer to the online help, enter the key word `niping`.

3. In window 3 (host3), start the test program niping to emulate a client. Enter the following command:

   ```
   call niping '-c' '-H' 'host2'
   ```

   This command tests the connection that is directly between host2 and host3 without using SAProuter.

4. In window 3, start the test program niping again with the following command:

   ```
   call niping '-c' '-H' '/H/host1/H/host2'
   ```

   This command tests the connection using the SAProuter. A host name is interpreted as a route (over one or more SAProuters to the server) if /H/ is added as a prefix to the host name.

In steps 3 and 4, data packages are sent to the server, and the server sends the data packages back. In step 3, the data packages should be sent to the server more frequently, since more process changes take place.

Performing a self test for the local host

To perform a self test for the local host, enter the following command:

```
call niping '-t')
```

A list with function names, parameters, and return codes is displayed. If the self test is successful, the following message appears:

```
*** SELFTEST O.K. ***
```

To see the options provided by niping, enter `niping` without any parameters.

Route Permission Table (saprouttab)

The route permission table contains the host names and port numbers of the predecessor and successor points on the route (from the SAProuter's point of view), as well as the passwords required to set up the connection. This corresponds to a substring of route strings which describe the stations of a connection required between two hosts. It is used to specify which connections are allowed and which are prohibited by SAProuter. It also specifies whether SNC connections are set up and which SNC connections these are.

Standard entries in a route permission table are as follows:

`P/S/D` *`source-host dest-host dest-serv password`*

In these standard entries:

- Permit (`P`) causes SAProuter to set up the connection.
- Secure (`S`) only allows connections with the SAP protocol. Connections with other protocols (such as Transmission Control Protocol (TCP)) are not allowed.
- Deny (`D`) prevents the connection from being set up.
- *source-host* and *dest-host* are represented by a host name or by an IP address.
- *dest-serv* is the IP-service on the destination host which is requested.
- *password* is an arbitrary password.
- You can also add comment lines, which must begin with a number sign character ("#").

There are specific SNC entries in the saprouttab for SNC connections. You can find further information in the SAProuter documentation in the following Web sites:

- `http://service.sap.com/saprouter`
- `http://service.sap.com/access-support`

The following rules apply when the SAProuter evaluates the route permission table:

- First match

  The first entry in the route permission table for which source address, target address, and target port match determines the match.

- No match

  If there is no appropriate entry in the table for a route, the connection is rejected. It behaves as though the last line is "D * * *" .

- Wildcards exception

  A wildcard ("*") can be used with the SAProuter connection using SNC.

You can find further information in the SAP SAProuter Documentation in the following Web sites:

- `http://service.sap.com/saprouter`
- `http://service.sap.com/access-support`

SAProuter connections via SNC

SAProuter connections can use SNC or virtual private network (VPN) connections. This section describes how to configure using SAProuter in order to enable the connection to SAP through the Internet, and receive objects over an SNC connection.

If you want to use the SNC connection over the Internet, then you should connect via sapserv2 to the SAP system instead of using sapserv3, sapserv4, sapserv,5, sapserv6, or sapserv7.

The SAPCRYPTO Service Program (*SRVPGM) and some Integrated File System files in /secude/etc are required for an SNC connection with the System i server. Perform the following steps to configure SAProuter:

1. Download the necessary software components from SAP Service Marketplace.

   Refer to the SNC documentation on the Service Marketplace before downloading the sapcrypto library for i5/OS at the following Web site:

   `http://service.sap.com/saprouter-sncdoc`

   This document generally describes the installation of the sapcrypto library. It is not specific to System i configurations. However, the links and general information can be of interest. Then perform the following steps:

   a. Download the required software from the following Web site:

      `http://service.sap.com/tcs`

   b. Click the **Download Area** and then click **SAP Cryptographic Software**.

      The downloaded file is a CAR file. Use the SAPCAR command to unpack it on your PC. Then upload it to the System i server using the binary FTP command into *SAVF QGPL/CRYPTO.

Follow these steps to restore the contents on the System i server:

1. Logon to the System i server as QSECOFR.
2. Enter the following commands:

   a. `CRTLIB LIB(SAPROUTER)`

   b. `RSTOBJ OBJ(*ALL) SAVLIB(SAPCRYPTO) DEV(*SAVF) SAVF(QGPL/CRYPTO) MBROPT(*ALL) ALWOBJDIF(*ALL) RSTLIB(SAPROUTER)`

      Ignore the CPF3848 `Security or data format changes occurred` message that is displayed.

   c. `GRTOBJAUT OBJ(SAPROUTER/*ALL) OBJTYPE(*ALL) USER(*PUBLIC) AUT(*ALL)`

   d. `MKDIR DIR('/secude')`

   e. `MKDIR DIR('/secude/etc')`

   f. `RST DEV('/qsys.lib/saprouter.lib/secude_etc.file') OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) ALWOBJDIF(*ALL)`

      Ignore the CPD377B `Security changes occurred for 4 objects` message.

   g. `CHGPGP OBJ('/secude/etc') NEWPGP(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)`

   h. `CHGPGP OBJ('/secude/etc/*') NEWPGP(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)`

3. Logon as *SID*OFR and install the ASCII version of SAPROUTER (6.20 kernel) to the library SAPROUTER with the APYR3FIX command.

Follow these steps to create the certificate request:

1. Logon on as *SID*OFR
2. Create the /usr/sap/saprouter directory if it does not exist already with the following commands:

   a. `MKDIR DIR('/usr/sap/saprouter')`

   b. `CD DIR('/usr/sap/saprouter')`

   c. `ADDLNK OBJ('/qsys.lib/saprouter.lib/sapcrypto.srvpgm') NEWLNK('/usr/sap/saprouter/sapcrypto')`

3. Go to the following Web site:

   http://service.sap.com/saprouter-sncadd

   Click **Trust Center Service in Detail**, select **SAProuter Certificates** and click the **Apply Now!** button. Obtain the "Distinguished Name" for your SAProuter from the list of SAProuters registered for your installation.

4. Enter the following command:

```
ADDLIBLE LIB(SAPROUTER)
```

5. As user *SID*OFR, set the environment variable by entering the following command:

```
ADDENVVAR ENVVAR('SECUDIR') VALUE('/usr/sap/saprouter')
```

6. Generate the certificate request by entering the following commands:

```
CALL PGM(SAPGENPSE) PARM('get_pse' '-v' '-noreq' '-p' 'local.pse' 'Your Distinguished
Name')
CALL PGM(SAPGENPSE) PARM('get_pse' '-v' '-onlyreq' '-r' 'certreq' '-p' 'local.pse')
```

   The Distinguished Name can be generated as follows:

   a. Enter the following command:

```
Your Distinguished Name = CN=saprouter-server-name, OU=cust-number, OU=SAProuter,
O=SAP, C=DE
```

      Use single-quotes around the distinguished name, not double-quotes as sometimes, this is specified for NT and other operating systems.

   b. You are prompted twice to enter a PIN. Enter the same personal identification number (PIN) every time this Problem Solving Environment (PSE) is used. Press **Enter** both times the PIN is requested to skip the request. This allows everyone to use this PSE, because it is not secured by a PIN.

7. Display the output file named certreq. Copy and paste the certificate request into the text area of the same form on the SAP Service Marketplace from which you copied the Distinguished Name as follows:

```
EDTF STMF('/usr/sap/saprouter/certreq')
```

8. A certificate signed by the Certification Authority (CA) in the Service Marketplace is received. Cut and paste the text to a local file named /usr/sap/saprouter/srcert. As this signed certificate is quite lengthy, consider creating this file with a PC based text editor and then upload it via NetServer/400.

9. Enter the following command to install the certificate in your saprouter:

```
CALL PGM(SAPGENPSE) PARM('import_own_cert' '-c' 'srcert' '-p' 'local.pse')
```

10. Create the credentials for the SAProuter with the same program:

```
CALL PGM(SAPGENPSE) PARM('seclogin' '-p' 'local.pse')
```

    You can log in to the System i server with any user profile. The "-O"-option is not supported on System i configurations.

11. Check if the certificate has been imported correctly by entering the following command:

```
CALL PGM(SAPGENPSE) PARM('get_my_name' '-v' '-n' 'Issuer')
```

    The name of the issuer should be:

    CN=SAProuter CA, OU=SAProuter, O=SAP, C=DE

12. Check the expiration date of the certificate with the following command:

```
CALL PGM(SAPGENPSE) PARM('get_my_name')
```

    These certificates are valid for one year. It is not possible to use the connection to sapserv2 without a certificate. Create a install a new certificate in a timely manner.

13. Open a customer message for component BC-OP-AS4 when problems are encountered. Describe what has been done and give the output of each commandused.

The following steps must be completed before you can start SAProuter:

1. Setup a new saprouttab, as described in the SNC document located at the Web site:

   http://service.sap.com/saprouter-sncdoc

   a. Edit the saprouttab as follows:

      ```
      EDTF STMF('/usr/sap/saprouter/saprouttab')
      ```

      An example of a saprouttab is shown in Example 17-1. The double quotes are the correct syntax.

Example 17-1 Example saprouttab

```
# SNC-connection from and to SAP
KT “p:CN=sapserv2, OU=SAProuter, O=SAP, C=DE” 194.39.131.34 *
# SNC-connection from SAP to local R/3-System for Support
KP “p:CN=sapserv2, OU=SAProuter, O=SAP, C=DE” * *
# SNC-connection from SAP to telnet in your network
KP “p:CN=sapserv2, OU=SAProuter, O=SAP, C=DE” * 23
# Access from the local Network to SAPNet - R/3 Frontend (OSS)
P * 194.39.131.34 3299
# deny all other connections
D * * *
```

   b. The following environment variable settings are required:

      - `ADDENVVAR ENVVAR('SECUDIR') VALUE('/usr/sap/saprouter')`
      - `ADDENVVAR ENVVAR('SNC_LIB') VALUE('/usr/sap/saprouter/sapcrypto')`

   c. Change the directory, as follows:

      ```
      CD DIR('/usr/sap/saprouter')
      ```

   d. Start the SAProuter, as follows:

      ```
      CALL PGM(SAPROUTER) PARM('-r' '-R' 'saprouttab' '-K' 'p:Your Distinguished Name')
      ```

   Include the three commands above in a Control Language (CL) program to automate this procedure. Run the procedure in batch. Otherwise, a terminal session is blocked by the saprouter.

17.2 Using the IBM Job Scheduler in an SAP system environment

SAP jobs can be started on the System i server from outside the SAP system. Chapter 29, “Data exchange scenarios and examples” on page 621 discusses using the SAP Event (SAPEVT) CL command and the Start Report (STRREPORT) CL command. These commands are SAP-supplied CL commands. In addition, you can use the IBM Job Scheduler to initiate jobs in an SAP system environment.

The following IBM Job Scheduler commands are utilized by the examples described in this section:

- WRKJOBSCDE
- ADDJOBSCDE
- CHGJOBSCDE
- RMVJOBSCDE

- HLDJOBSCDE
- RLSJOBSCDE

The recommended use of parameters for scheduled jobs is encapsulated in the CMD() parameter of the SBMJOB statement, as shown in Example 17-2.

Example 17-2 Recommended schema for IBM Job Scheduler

```
ADDJOBSCDE
JOB(<my_entry_name>)
CMD(SBMJOB CMD(<the_command_to_schedule>) CPYENVVAR(*YES) JOB(<batchjob_name>)
JOBQ(QSYSNOMAX) USER(<sid>OFR))
FRQ(*ONCE) SCDDATE(*CURRENT) SCDDAY(*NONE) SCDTIME(*CURRENT) SAVE(*YES) JOBQ(QSYSNOMAX)
USER(<sid>OFR)
```

Job Scheduler entry to set the SAP system environment

This section describes an example to illustrate the use of the IBM Job Scheduler to invoke SAPROUTER to set the SAP system environment.

Set the environment variables and the library list by executing the SAPROUTER command and statements as shown in Example 17-3.

Example 17-3 Using SAPROUTER with the IBM Job Scheduler

```
ADDLIBLE SAPROUTER /* if you use a separate library instead of the SAP kernel */
CALL R3<sid>400/R3INLPGM
ADDJOBSCDE JOB(SAPROUTER)
CMD(
SBMJOB CMD(CALL PGM(SAPROUTER) PARM('-r' '-R' '/usr/sap/saprouter/saprouttab' '-W' 60000))
CPYENVVAR(*YES) JOB(SAPROUTER) JOBD(R3<sid>400/R3_<instance>) JOBQ(QSYSNOMAX)
USER(<sid>OFR))
FRQ(*ONCE) SCDDATE(*CURRENT) SCDDAY(*NONE) SCDTIME(*CURRENT) SAVE(*YES) JOBQ(QSYSNOMAX)
USER(<sid>OFR) TEXT('Manual Start of SAPROUTER if needed')
```

Hold the scheduling entry if you do not want the job to start immediately as specified by the scheduling parameter when the scheduling entry was created, by entering the following command:

```
HLDJOBSCDE JOB(SAPROUTER) ENTRYNBR(*ONLY)
```

To submit the job once (immediately) and to keep the job scheduling entry definition, enter the following command:

```
CHGJOBSCDE JOB(SAPROUTER) FRQ(*ONCE) SCDDATE(*CURRENT) SCDDAY(*NONE) SCDTIME(*CURRENT)
SAVE(*YES)
```

It is important for every SBMJOB command in the SAP system environment to set the correct library list and environment variables. There are three ways to achieve this:

- Invoke the SBMJOB command when signed on as user *SID*OFR.
- If not signed on as *SID*OFR, invoke the SBMJOB after calling R3*SID*400/R3INLPGM. Add the USER(*SID*OFR) parameter to ensure that access rights are correct while the batch job executes (Integrated File System and SAP database rights), and authority definitions for new objects created by the batchjob.
- Write a Control Language Program (CLP) to create the example program named STRSAPRTR, as illustrated in Example 17-4.

Example 17-4 CLP listing for SAPROUTER procedure

```
PGM PARM(&SID)
    DCL VAR(&SID) TYPE(*CHAR) LEN(3)
    DCL VAR(&LIB) TYPE(*CHAR) LEN(10)
    DCL VAR(&ROUTTAB) TYPE(*CHAR) LEN(80)
    CHGVAR VAR(&LIB) VALUE('R3' *CAT &SID *CAT '400')
    CALL PGM(&LIB/R3INLPGM)
    CHGVAR VAR(&ROUTTAB) VALUE('/usr/sap/saprouter/saprouttab')
CHGLIBL LIBL(QTEMP SAPROUTER QGPL QGPTOOLS)
    MONMSG MSGID(CPF0000)
    CALL PGM(*LIBL/SAPROUTER) PARM('-r' '-R' '/usr/sap/saprouter/SAP-
     ROUTTAB' '-W' '60000')
ENDPGM:

   ENDPGM
```

Run the program CALL PGM(STRSAPRTR) PARM(*sid*) in the CMD() parameter of the Job Scheduling Entry, as follows:

```
ADDJOBSCDE JOB(SAPROUTER) CMD(
SBMJOB CMD(CALL PGM(STRSAPRTR) PARM(<sid>) CPYENVVAR(*YES)) JOB(SAPROUTER)
JOBQ(QSYSNOMAX) USER(<SIDOFR>))
FRQ(*WEEKLY) SCDDAY(*MON) SCDDATE(*NONE) SCDTIME('00:10') OMITDATE('04.01.06')
JOBQ(QSYSNOMAX) USER(<sid>OFR) TEXT('weekly start of SAPROUTER')
```

The library list and environment parameters do not need to be set.

The optional parameters and statements used in this example CLP program are as follows:

- To set the schedule as a weekly for early each Monday morning, specify the following parameter values:

  ```
  FRQ(*WEEKLY) SCDDAY(*MON) SCDDATE(*NONE) SCDTIME('00:10')
  ```

- If you do not want to schedule the job to run on special holidays, enter the following parameter values:

  ```
  OMITDATE('24.12.06') or
  OMITDATE('12/24/06') /*
  ```

  The syntax depends on the date format and date separator of “*/”.

- The job queue QSYSNOMAX is connected to the subsystem QSYSWR and has no maximum number of parallel jobs. Specify this as follows:

  ```
  JOBQ(QSYSNOMAX) /*
  ```

- To run the job in the subsystem of the SAP instance, specify the following, where *xx* is the instance number:

  ```
  JOBQ(R3<sid>400/R3_xx) /* xx = instance number
  ```

Optional parameters for the SBMJOB command include:

- USER(*sid*OFR) for user authority
- JOBD(R3*sid*400/R3_*xx*>) INLLIBL(*JOBD) where *xx* is the instance number.

Job schedule entry to add jobs to the IPL Startup program

You can create a job scheduling entry for every additional command you want to include in the i5/OS IPL Startup program. For example:

Add job schedule entries for all commands for processes that should easily be stopped and restarted without performing an IPL.

Enter the following command for every process in the IPL Startup program to invoke the start of your process during an IPL:

```
CHGJOBSCDE JOB(<my_process>) FRQ(*ONCE) SCDDATE(*CURRENT) SCDDAY(*NONE) SCDTIME(*CURRENT)
SAVE(*YES)
```

Enter the following command to deselect the start of a process IPL, without changing and recompiling your Startup program:

```
HLSJOBSCDE JOB(<my_process>)
```

Job schedule entry for miscellaneous jobs

Additional examples for starting jobs with the Job Scheduler include:

- All periodic jobs, for example daily and weekly backup
- A job entry to save and cleanup journal receivers that you are manually run when a journal ASP runs out of disk space.

17.3 Automated startup

At system startup, you can start subsystems and jobs, including those from your SAP infrastructure.

17.3.1 i5/OS functions

Many of the i5/OS functions can be automatically started when the system starts at Initial Program Load (IPL). The IPL attributes screen shows some of the functions performed during the IPL of a System i server. Use the Display IPL Attributes (DSPIPLA) command to see the attributes of the IPL, The display shows, for example, whether the Transmission Control Protocol over Internet Protocol (TCP/IP) is automatically started. Refer to Figure 17-3 on page 306 for an illustration.

```
                          Change IPL Attributes (CHGIPLA)

 Type choices, press Enter.

 Restart type . . . . . . . . . . RESTART        *SYS
 Keylock position . . . . . . . . KEYLCKPOS      *NORMAL
 Hardware diagnostics . . . . . . HDWDIAG        *MIN
 Compress job tables  . . . . . . CPRJOBTBL      *ALL
 Check job tables . . . . . . . . CHKJOBTBL      *ABNORMAL
 Rebuild product directory  . . . RBDPRDDIR      *NONE
 Mail Server Framework recovery   MSFRCY         *NONE
 Display status . . . . . . . . . DSPSTS         *ALL
 Start TCP/IP . . . . . . . . . . STRTCP         *YES
 Clear job queues . . . . . . . . CLRJOBQ        *NO
 Clear output queues  . . . . . . CLROUTQ        *NO
 Clear incomplete joblogs . . . . CLRINCJOB      *NO
 Start print writers  . . . . . . STRPRTWTR      *YES
 Start to restricted state  . . . STRRSTD        *NO



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

Figure 17-3 i5/OS CHGIPLA command

The attributes of an IPL can be changed with the Change IPL Attributes (IPLA) command.

Alternatively, you can run a program to modify the attributes of the startup program, as described in this section. The startup program runs at each IPL of the partition. The name of the startup program is defined by system value QSTRUPPGM. You can find the default IBM startup program named QSTRUP in the QSYS library.

To retrieve the source used for the IBM IPL startup program, follow these steps:

1. Copy the QSTRUP program from QSYS library into the QGPL library, or in another general purpose library.
2. Change the system value QSTRUPPGM to use the IPL start program to that QSTRUP program you just copied. Until now, this QSTRUP program is the same as the original QSTRUP from QSYS.
3. Now retrieve the source code from the original or the copied QSTRUP from step 1, edit the source code and create a new and modified QSTRUP program to override the copied QSTRUP program from step 1.

> **Note:** Do not change the IBM-supplied QSTRUP program in QSYS. Instead, create a modified version of QSTRUP in another library, for example, QGPL. Otherwise, any modifications to the IBM-supplied QSTRUP program in QSYS are overwritten when an IBM system upgrade is performed.

 – Run the Retrieve CL Source (RTVCLSRC) command to retrieve the source code for QSTRUP.
 – Use Start Source Entry Utility (STRSEU) or Edit File (EDTF) to edit the QSTRUP source code. SEU is part of the Program Development Manager (PDM) program in 5722-WDS.

– Run Create CL Program (CRTCLPGM) to create the CL Program.

Tip: Your modified startup program can include statements to start SAP systems automatically after an IPL. If you are performing an IPL to apply program temporary fixes (PTFs) or to perform hardware maintenance and want to verify the results before starting the SAP instances, you can set system value QSTRUPPGM to `QSTRUP QSYS` for this IPL.

After verifying the system state, you can execute your own startup program from the system console or set the system value back to your modified program and perform another IPL.

The sequence of events in an IPL process is as follows:

- TCP/IP is started during a normal IPL according to the IPL attributes specifications
- The IPL startup program (QSYS/QSTRUP) invoked when the controlling subsystem is started. The controlling subsystem is normally QBASE or QCTL.

A standard System i recommendation is that the QSTRUP program should perform all customer-specific IPL tasks, and then issue a call to PGM(QSYS/QSTRUP). This rule is beneficial if the QSERVER subsystem is also started in your QSTRUP routine. In this case, the start is issued twice, which does not cause a problem other than generating warning messages in the QSYSOPR message queue.

17.3.2 QSTRUP in an SAP system environment

Use the information in this section to customize the QSTRUP IPL startup program to control various events in an SAP system environment. For example:

- Starting the QXDAEDRSQL job automatically

  The QXDAEDRSQL job is a prerequisite for working with an SAP system. It is required for the following tasks:

  – Calling operating system specific interfaces from ABAP, for example, in the Computing Center Management System (CCMS)

  – Secondary connections, for example, in mySAP BW

  – 3-tier installations

  – 2-tier installations with instance profile parameter dbs/db4/allow_cancel = 1

  To set the QXDAEDRSQL job to be started automatically after each IPL, enter the following code into the QSTRUP program:

```
STRTCPSVR SERVER(*EDRSQL)
MONMSG MSGID(CPF0000)
```

- Starting the *DDM TCP/IP Server and Host Servers automatically

  To use the Toolbox JDBC Driver for your J2EE or ABAP+J2EE system, start the DDM TCP/IP server and host servers after each IPL. Enter the following code into the QSTRUP program to start the servers automatically after each IPL:

```
STRTCPSVR SERVER(*DDM)
MONMSG MSGID(CPF0000)
STRHOSTSVR SERVER(*ALL)
MONMSG MSGID(CPF0000)
```

Note: You will require the *DDM TCP/IP server when using the native Java Database Connectivity (JDBC) driver. You need the host servers when using the Toolbox JDBC driver for SAP WebAS Java.

- Starting the NFS-Server in order to access remote file systems

  Enter the following:

  ```
  STRNFSSVR SERVER(*ALL) NBRSVR(5) NBRBIO(5)
  MONMSG MSGID(CPF0000)
  ```

  Export all file systems, for example /etc/exports, as follows:

  ```
  CHGNFSEXP OPTIONS('-A')
  MONMSG MSGID(CPF0000)
  ```

  Refer to 17.3.4, “Concerning Linux Application Server” on page 310 for an example.

- Adding links with /QFileSvr.400 in a system landscape with multiple System i servers, for example, to create a common transport directory as follows:

  ```
  MD DIR('/QFileSvr.400/my_server1')
  MONMSG MSGID(CPF0000)
  MD DIR('/QFileSvr.400/my_server2')
  MONMSG MSGID(CPF0000)
  ```

- Retrieving QSTRUP source statements

  Source statements are not provided for the QSTRUP startup program. However, because *YES was specified for the RTVCLSRC parameter when QSTRUP was created, QSTRUP statements can be retrieved, as follows:

  a. Retrieve the CL source statements by entering the following command:

     ```
     RTVCLSRC PGM(QSYS/QSTRUP) SRCFILE(QGPL/QCLSRC) SRCMBR(QSTRUP)
     ```

     If the source statements are retrieved successfully, the following message is displayed:

     ```
     CL source retrieved for program QSTRUP in QSYS.
     ```

  b. Edit the program with EDTF by entering the following command:

     ```
     EDTF FILE(QGPL/QCLSRC) MBR(QSTRUP)
     ```

     To start the QXAEDRSQL job automatically, enter the following command:

     ```
     STRTCPSVR SERVER(*EDRSQL)
     MONMSG MSGID(CPF0000)
     ```

Press **F2** and then **F3** to exit the editor.

Example 17-5 shows you what a QSTRUP program modification can look like.

Example 17-5 Customer specific entries in QSTRUP

```
..+....2....+....3....+....4....+....5....+....6....+....7....+....8....
DONE:
     QSYS/STRSBS SBSD(QSPL)
     MONMSG MSGID(CPF0000)
/*------------------------------------------------------------------*/
/* Begin Customer specific part for SAP         17.05.2005, D.H. */
/*                                              04.01.2006, J.Z, */
STRTCPSVR  SERVER(*EDRSQL *REXEC *NETSVR)
     MONMSG MSGID(CPF0000 TCP0000)
     STRTCPSVR  SERVER(*DDM)
     MONMSG MSGID(CPF0000 TCP0000)
     STRHOSTSVR SERVER(*ALL)
```

```
     MONMSG MSGID(CPF0000 TCP0000)

/* IFS and/or NFS mount                                              */
   MD DIR('/QFileSvr.400/my_server1')
     MONMSG MSGID(CPF0000)
     MD DIR('/QFileSvr.400/my_server2')
     MONMSG MSGID(CPF0000)

/* Cluster Environment if needed                                     */
     STRTCPSVR  SERVER(*INETD)
   MONMSG MSGID(CPF0000 TCP0000)
/* High Availability Software Solutions if implemented               */

/*   And then the handling of iSeries Cluster if implemented         */
     STRCLUNODE CLUSTER(SAPCLUSTER) NODE(<my_server1>)    /* primary */
     MONMSG MSGID(CPF0000)
     STRCLUNODE CLUSTER(SAPCLUSTER) NODE(<my_server2>)    /* backup */
     MONMSG MSGID(CPF0000)
/* perhaps STRTCPIFC for a switchable IP Address, if not handled */
     /* by an APPLICATION CRG in the takeover IP-Address              */
     STRCRG CLUSTER(SAPCLUSTER) CRG(IASP33)      /* independent ASP */

/* Start SAP System, example with <SID> PRD and instance number 00*/
   SBMJOB CMD(CALL PGM(<MYLIB>/<MYCLPGM>) PARM('PRD' '00')) +
   JOB(STARTSAP) USER(PRDOFR) MSGQ(QSYSOPR)
   MONMSG MSGID(CPF0000)
/* End of Modifications for SAP                    17.05.2005 */
/*                                                 04.01.2006 */
/*------------------------------------------------------------------*/

     QSYS/RTVSYSVAL SYSVAL(QSTRPRTWTR) RTNVAR(&STRWTRS)
     IF COND(&STRWTRS = '0') THEN(GOTO CMDLBL(NOWTRS))
     CALL PGM(QSYS/QWCSWTRS)
     MONMSG MSGID(CPF0000)
NOWTRS:
     RETURN
     CHGVAR VAR(&CPYR) VALUE(&CPYR)
     ENDPGM
```

17.3.3 Starting SAP applications during an IPL

The SAP system can be automatically started during an IPL. To create this program, generate a CL program with the source listed in Example 17-6. Name the CL program for example *mylib/myclpgm* for reuse in the SBMJOB command.

Example 17-6 CL source for mylib/myclpgm

```
PGM PARM(&SID &INST)
DCL  &SID *CHAR 3
DCL  &INST *CHAR 2
DCL  &LIB *CHAR 10
DLYJOB DLY(300)        /* Wait until TCP/IP and QSERVER are running */
/* If your "/usr/sap/<SID>" directory is a link to           */
/* "/QFileSvr.400/<someAS400>/sapmnt/<SID>":                */
MKDIR /QFileSvr.400/<someAS400>
MONMSG CPFA0A0
CHGVAR &LIB ('R3' *CAT &SID *CAT '400')
CALL &LIB/R3INLPGM     /* Sets the library list */
STARTSAP *DB           /* Starts the remote DB server for R/3 */
```

```
STARTSAP &SID &INST    /* Starts the instance */
ENDPGM
```

Add the following statement to the System i startup program. (This is already done in the Example 17-6 on page 309.)

```
SBMJOB CMD(CALL PGM(mylib/myclpgm) PARM('PRD' '00')) JOB(STARTSAP) USER(PRDOFR)
MSGQ(QSYSOPR)
```

This example assumes that the system ID of *SID* is PRD and the instance number is 00.

4. Create the CL program from the source code after the source statements are changed in program QSTRUP by entering the following:

   ```
   CRTCLPGM PGM(QGPL/QSTRUP) SRCFILE(QGPL/QCLSRC) SRCMBR(QSTRUP)
   ```

   If the program is created successfully, you will receive this message:

   ```
   Program QSTRUP created in library QGPL.
   ```

   To make sure that your adjusted startup program (and not the default startup program) is started after an IPL, change the QSTRUPPGM system value to point to the modified startup program and library.

17.3.4 Concerning Linux Application Server

If you have Linux Application Server in your SAP system landscape, you must start the Network File System (NFS) server and export the related file systems named /etc/exports.

The best way to do it is to include it in the following QSTRUP start program:

```
QSTRUP:
...
/* Start the NFS-Server to really export the filesystems */
STRNFSSVR SERVER(*ALL) NBRSVR(5) NBRBIO(5)
MONMSG MSGID(CPF0000)

/* Export all filesystems of /etc/exports */
CHGNFSEXP OPTIONS('-A')
MONMSG MSGID(CPF0000)
...
```

Include it with these associated export files in /etc/exports in the SAP system environment:

- /sapmnt/<SID>/exe ANON=-1 ROOT=<LinuxHost>
- /sapmnt/<SID>/profile ANON=-1 ROOT=<LinuxHost>
- /sapmnt/<SID>/global ANON=-1 ROOT=<LinuxHost>
- /sapmnt/<SID>/j2ee ANON=-1 ROOT=<LinuxHost>
- /sapmnt/<SID>/jdbc ANON=-1 ROOT=<LinuxHost>

17.3.5 Automatically start and stop multiple instances and the application server

Normally, you should start all application servers when the database instance starts.

If all instances are on the same host as the central instance, start the application servers by entering the following command:

```
STARTSAP SID(SID) INSTANCE(*ALL)
```

Stop the application servers by entering the following command:

```
STOPSAP SID(SID) INSTANCE(*ALL)
```

However, this procedure does not work unless all instances are on the same host. It also does not work if you have Windows or Linux application servers.

To start and stop additional System i application servers at the same time, consider the following:

- The R3RMTDB job (prior to OS/400 V4R4) or QXDAEDRSQL (as of OS/400 V4R4) must be active on all System i servers before the startup.
- Enhance the start-up profile of the central instance. In this example, the host name of the application server is APPSRV1, the SAP system ID is PRD, the instance number is 01, the start entries in the SAP start profile are consecutively numbered up to 05, and the stop entries are numbered up to 02.
  - Enter an additional startup program for each instance. Add the following statement at the end of the start-up program:

    ```
    Start_Program_06 = APPSRV1 STR_PRD_01
    ```

  - Enter an additional stop program for each instance:

    ```
    Stop_Program_03 = APPSRV1 STP_PRD_01
    ```

  When you start the central instance, these entries cause a job to be started on APPSRV1 as follows:

  ```
  SBMJOB CMD(CALL STR_PRD_01)
  ```

  To stop the job, enter the following command:

  ```
  SBMJOB CMD(CALL STP_PRD_01)
  ```

  You can set the library list in these jobs by entering the following command:

  ```
  CALL PGM(R3PRD400/R3INLPGM)
  ```

  - Create the two CL programs listed in Example 17-7 on the System i server named APPSRV1.

Example 17-7 Two CL programs to start and stop an SAP system

```
CL program STR_PRD_01:
PGM
CALL PGM(R3<SID>400/R3INLPGM)
STARTSAP SID(PRD) INSTANCE(01)
ENDPGM

CL program STP_PRD_01:
PGM
CALL PGM(R3<SID>400/R3INLPGM)
STOPSAP SID(PRD) INSTANCE(01)
ENDPGM
```

  - Place the CL programs into a general purpose library, such as QGPL.

Refer to *SAP note 93316* for further information.

> **Note:** For starting and stopping remote Windows application servers along with the System i application servers, refer to SAP note 682281.

17.3.6 Starting Windows AppServer from System i central host on remote systems

The following statements highlight the considerations outlined in *SAP note 682281* to start the Windows application server from the central host on the System i server on remote systems:

1. Start the service “iSeries Access for Windows Remote Command” on each Windows server.
2. Select **Start** → **Run** and type `services.msc`. This opens a new window with a list of services available.
3. Double-click **iSeries Access for Windows Remote Command** to get the property sheet. Change Startup type to **Automatic**.

   This opens the Windows server to rexec connections from remote hosts. You can start Windows application servers from the System i central instance.

   > **Note:** This procedure exposes your Windows server to the risks associated with a rexec demon running on the server. For example, it opens the possibility to allow any user with a password to run commands on that machine.

4. Enhance the start-up profile of the central instance. In the following example, the host name of the windows application server is app01, the SAP system ID is PRD, the instance number is 01, the start entries in the SAP startprofile are consecutively numbered up to 05, and the stop entries up to 02.
5. Enter an additional start-up program for each instance, that is, at the end of the start-up program, as follows:

   ```
   Start_Program_06 = local STR_PRD_01
   ```

6. Enter an additional stop program for each instance as follows:

   ```
   Stop_Program_03 = local STP_PRD_01
   ```

   When you start the central instance as follows, these entries cause a job to be started on the central instance:

   ```
   SBMJOB CMD(CALL STR_PRD_01)
   ```

   You can stop the job as follows:

   ```
   SBMJOB CMD(CALL STP_PRD_01)
   ```

7. Set the library list in these jobs by entering the following command:

   ```
   CALL PGM(R3PRD400/R3INLPGM)
   ```

8. Create the STR_PRD_01 and STP_PRD_01 CL programs on the System i central instance, as shown in Example 17-8 and Example 17-9 on page 313.

Example 17-8 CL program STR_PRD_01

```
PGM
CALL PGM(R3<SID>400/R3INLPGM)
RUNRMTCMD +
     CMD('D:\usr\sap\<SID>\D<INST>\exe\startsap.exe +
        NAME=<SID> NR=<INST> SAPDIAHOST=app01') +
     RMTLOCNAME(app01 *IP) +
     RMTUSER(<SID>ADM) +
RMTPWD('xxxxxxxx')
ENDPGM
```

Example 17-9 CL program STP_PRD_01

```
PGM
CALL PGM(R3<SID>400/R3INLPGM)
RUNRMTCMD
CMD('D:\usr\sap\<SID>\D<INST>\exe\stopsap.exe
        NAME=<SID> NR=<INST> SAPDIAHOST=app01')
     RMTLOCNAME(app01 *IP)
     RMTUSER(<SID>ADM)
     RMTPWD('xxxxxxxx')
ENDPGM
```

"`D:`" is the drive designation for the application server on the Windows Server. Substitute the correct password for the `xxxxxxxx` value.

Use the Work with Spool File (WRKSPLF) command to display the results. Place the CL programs into a general library such as QGPL. The results of running this remote command can be displayed with the WRKSPLF command.

17.4 Copying CDs and DVDs to hard disk

With the installation of an SAP system at 6.20, 6.40 or 7.00, it is very useful to copy the installation CDs and DVDs to hard disk, because for the installation the CD, data needs to be accessible from both the System i server and the PC that runs the SAPInst tool.

There are several ways to copy CD or DVD content to the Integrated File System of a System i configuration. These are as follows:

- Using the operating system CPY command as follows:

  ```
  SBMJOB CMD(CPY OBJ('/QOPT/*') TODIR('/tmp/') SUBTREE(*ALL) REPLACE(*YES) OWNER(*KEEP))
  JOBQ(QCTL) JOB(XCOPY)
  ```

- Using the Q-shell as follows:

  ```
  MKDIR '/tmp/<CD-number>' DTAAUT(*RWX) OBJAUT(*ALL)
  strqsh
  cp -r /QOPT/<CD-number>/*.* /tmp/CD51014013
  <F3>
  ```

  Use the WRKLNK or WRKLNKSAP command to verify that the content is copied successfully.

- Using the CPYDIR command from the SAP kernel.

  The CPYDIR command is not delivered on releases later than 4.6D. Depending on the patch level and operating system release, CPYDIR might not run. Use the CPYDIR command as follows:

  ```
  CPYDIR FROMDIR('/QOPT') TODIR('/tmp/<directory>')
  ```

- Using iSeries Operations Navigator and xcopy (fast)

  Sign on to the System i server with a user profile of QSECOFR or user with *SECADM authority and create a shared drive with read and write access to the /tmp directory. Map the /tmp network drive as X:\ in Windows Explorer. Alternatively, use the DOS command NET USE. Open a command line on the Windows PC and enter the following commands:

  ```
  E:\
  xcopy E:\*.* X:\<CD-Number>\*.* /S/E/V
  exit
  ```

Deleting the copied CDs does not typically work with the RRM SAP kernel command RRM because of the read-only attribute of the files. Use Q-Shell or Windows Explorer to clean up.

- Using iSeries Operations Navigator (OpNav) and xcopy (fast)

  Create a share to the directory /tmp (read/write access) from the OpNav with user profile QSECOFR or similar *SECADM user.

  Map the network drive /tmp as X:\ in the Windows Explorer. Alternatively, you can use the DOS command NET USE.

  Ensure that your Explorer settings show hidden and system files, as follows:

  a. Create the empty target directory.

  b. Mark the complete content of the CD (in its root directory) in the Explorer by pressing **Ctrl+A**, and then **Ctrl+C**.

  c. Copy the content to the target directory by pressing **Ctrl+V**. This normally takes between 5 and 20 minutes per CD.

- Using File Transfer Protocol (FTP)

  FTP can be used as an efficient method to copy CDs or DVDs to the System i disk drive, as follows:

  a. Use the Make Directory command to create the empty target directory for the CD, as follows:

  ```
  MKDIR '/tmp/CD51014013' DTAAUT(*RWX) OBJAUT(*ALL)
  ```

  b. From a Windows command line, call the Make Directory (MKDIR) for every subdirectory of the CD.

  c. From a Windows command line, enter the following:

  ```
  ftp <server>
  user <sid>ofr
  <pwd>
  namefmt 1
  bin
  lcd E:\
  cd /tmp/CD51014013
  prom
  mput *.*
  # shows errors for the subdirectory entries, but
  # copies all files in the directory
  # Repeat this for every subdirectory (and hierarchy level)
  # in the directory tree of the CD, for example install:
  cd E:\install
  lcd /tmp/CD51014013/install
  mput *.*
  quit
  exit
  ```

- For the transfer between two System i servers, enter the following:

  ```
  CRTSAVF QGPL/CD AUT(*ALL) /* on both servers */
  SAV DEV('/QSYS.LIB/QGPL.LIB/CD.FILE) IFS('/tmp/CDs*')
  ... transfer using FTP in binary mode
  (do not forget CD QGPL, LCD QGPL)
  RST DEV('/QSYS.LIB/QGPL.LIB/CD.FILE) IFS('/tmp/CDs*')
  ```

  The save files can now be deleted.

If the authorizations are not correct, use the following commands to grant access for every directory of the CD:

```
CHGPGP OBJ('/tmp/CD51014013/') NEWPGP(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)
CHGPGP OBJ('/tmp/CD51014013/*') NEWPGP(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)
```

Or use the Change Authority command, as follows:

```
CHGAUT OBJ('/tmp/CD51014013/') USER(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)
CHGAUT OBJ('/tmp/CD51014013/*') USER(R3GROUP) DTAAUT(*RWX) OBJAUT(*ALL)
```

- Using image catalogs

  To use image catalogs, enter the following:

```
CRTDEVOPT DEVD(OPTVRT01) RSRCNAME(*VRT) TEXT('Virtual Optical Drive') /* call once */
VRYCFG CFGOBJ(OPTVRT01) CFGTYPE(*DEV) STATUS(*ON) /* call once */
CRTIMGCLG IMGCLG(NW04) DIR('/tmp/NW04') CRTDIR(*YES) TEXT('Image Catalog for NetWeaver
Installation CDs/DVDs')
```

  Follow these steps for every CD or DVD to be processed:

  a. Insert the CD or DVD in the OPT01 DVD drive of the System i server.

  b. Enter the following command:

```
ADDIMGCLGE IMGCLG(NW04) FROMDEV(OPT01) TOFILE(*FROMFILE) LODIMGCLG IMGCLG(NW04)
DEV(OPTVRT01) OPTION(*LOAD)
```

  All CDs of the image catalog are now available and are seen under the /QOPT device.

  The Work with Optical Volume (WRKOPTVOL) command lists all physical and virtual CDs that are loaded. You can administer multiple image catalogs with the Work with Image Catalog (WRKIMGCLG) command. You can perform a cleanup with the following commands:

```
LODIMGCLG IMGCLG(NW04) DEV(OPTVRT01) OPTION(*UNLOAD)
DLTIMGCLG IMGCLG(NW04)
RRM '/tmp/NW041' or DEL '/tmp/NW04/*' and RD '/tmp/NW04'
```

18

Life cycle

This chapter considers the ways that your SAP system landscape and systems can significantly change over time. We focus on the following specific points:

- Growth management
- Available tools to track resource utilization over time
- Common upgrade scenarios

18.1 Growth management

It is safe to assume that your SAP system environment expands over time. More users are added, more data is stored, and more SAP systems are implemented. You also upgrade your existing infrastructure to more current release levels of hardware, operating systems and applications.

Your original system sizing is based on the information known at the time of the estimate. This is a good starting point based on the historical experience of others. However, just as your company is unique, it's growth pattern is unique. Over time, the actual use of your system can prove to be quite different than the original projection which was based on the original assumptions. At this time, your experience and knowledge of company history become a more important predictor of your future business needs.

Therefore, it is always a good idea to pay attention to the usage and growth patterns of your SAP systems. This section discusses some of these considerations.

18.1.1 Storage capacity growth

Expect database tables to grow in size over time. Track your storage growth on disk to make sure the available storage capacity of Auxiliary Storage Pools (ASPs) does not fill up and overflow.

The Work with System Status (WRKSYSSTS) command shows the utilization of the System ASP. The operating system shuts down if the system ASP reaches capacity. Take action immediately when you detect that the system ASP is nearing capacity to avoid an unplanned shutdown.

The QSTGLOWLMT system value specifies the percentage of available storage remaining in the system ASP when the lower limit is reached. The QSTGLOWACN system value specifies the action associated with this lower limit. The default action for QSTGLOWACN is to send a message to the QSYSOPR message queue. QSTGLOWACN can also be changed to use an exit point so that a program of your design can be called to take action to avoid the system shutdown.

The Work with Disk Status (WRKDSKSTS) command shows the utilization of each ASP, and of each disk within the ASP.

Use the Start ASP Balance (STRASPBAL) command to rebalance the data on the disk drives when you notice a large difference in utilization between disks within the same ASP so that all drives are used comparably.

Take the following options to try to reduce the capacity utilization if you notice that capacity utilization of an ASP is getting high:

- Check the temporary storage usage on Work with System Status output display (WRKSYSSTS). Compare the "Current unprotect used" and "Maximum unprotect" fields to the value of the "System ASP" field.

  Temporary storage is used by the system and by SAP applications. The "Current unprotect used" value increases rapidly when SAP instances start. Once the system is up and has been used for a short time, the value should increase much less rapidly.

  The "Maximum unprotect" field shows the high water mark for the "Current unprotect used" field. Increase the amount of disk storage or contact SAP to see if there is a memory leak in the application if these temporary storage values continue to increase rapidly, or become a significant portion of the System ASP. Check and record the measured values from time to time so that you are able to detect a significant change in their normal values.

The SAP system command Display Temporary Storage (DSPTMPSTG) shows the amount of temporary storage used by each work process and the shared memory of an instance. In the case that you are using an excessive amount of temporary storage, you can stop the SAP instance to release all of the temporary storage it is using. Then you can restart the instance. The temporary storage usage should be much less than it was. Continue to monitor the growth. If the problem reoccurs, contact SAP service.

- Use the i5/OS Retrieve Disk Information command (RTVDSKINF) to collect size information for all objects on the system if temporary storage usage is not a problem. Submit the command to batch. Use the Print Disk Information command (PRTDSKINF) to get a list of the largest libraries and objects when RTVDSKINF is finished. Use this process to determine if there are objects that are no longer needed on the system.
- Use the SAP DB02 transaction to find tables that should be compressed to remove deleted rows. The space that the rows of a table occupy is not actually freed when rows are deleted from a table. Instead, the rows are marked as deleted so that when new rows are added to the table, the same space can be reoccupied. New rows are added to the end of the table if there are no open rows.

  Certain SAP system operations, such as Client Delete or Archiving, can leave a large number of deleted rows in a table that are not filled for some time. Use the Reorganize Physical File Member command (RGZPFM) to reorganize these files to reclaim the space held by the deleted rows. Run the DB02 SAP transaction to find candidate tables. Check the “File statistic analysis date” field to make sure the information is current. Choose the “Deleted Row Analysis” button. The resulting report is ordered by the number of deleted rows. Review the “%deleted” and “Deleted bytes” columns of this report to identify candidates for reorganization. Note that reorganizations are not usually necessary unless huge amounts of data are deleted by the Client Delete or Archiving operations.

  It can also be helpful to reorganize large tables that contain Large Object (LOB) fields that are updated often, even if they do not have a significant number of deleted rows. Examples of these tables are REPOLOAD and DYNPLOAD. You can determine how often to reorganize tables and which tables are most important to reorganize on a monthly, quarterly or yearly basis with periodic examination of DB02 output.

  It takes time to reorganize a table. Use the RGZPFM command with the KEYFILE(*PRIMARY) parameter to order the file as it is reorganized. This can have a performance benefit since SAP system queries typically access data in primary key order.

  See SAP note 84081 for more information related to RGZPFM in an SAP context.
- Consider SAP Archiving if you still need to reclaim some storage space. SAP Archiving can significantly reduce the database size by moving SAP transactional data and objects from a period of time in the past, to storage that is external to the current database. Archiving can also improve performance of the system as a whole since it physically removes data from the database and therefore removes it from consideration by the database during queries. Remember to check DB02 for deleted rows and a possible table reorganization after a significant amount of archiving is done.
- Add more disk hardware to the system if archiving is not possible or additional disk capacity is yet required. Adding disk can often be done while the system is up and running without disrupting operations, however, be sure to use the STRASPBAL command immediately to spread the data in the ASP evenly across all of the existing and newly installed disk units.

See Chapter 9, “SAP storage and database considerations” on page 89 for a discussion of disk topics. See Chapter 22, “Availability, backup and recovery concepts” on page 373 for more information about managing the availability of disk.

18.1.2 Workload growth

Over time, you can add more users to a system, start using more functions within a given system or perform more background processing. Each of these occurrences place more demand on the CPU resources of the system, and can also require additional memory and disk resources.

Consider the following:

- Application instance profile tuning can be necessary as users are added to the system. See 26.1.1, “Tuning an SAP system” on page 539.
- You might want to add a separate application instance. Multiple instances can be used, even in a 2-tier configuration. Use the instances to separate work logically, or to give one instance priority to resources (CPU priority and memory pools) over another. Make sure sufficient memory is available to absorb the additional overhead if you add an instance, and ensure there is available disk capacity for additional temporary storage use.
- You might eventually outgrow the transactional throughput that can be provided by a single i5/OS partition. This requires switching to a 3-tier environment, where another partition or server is used to provide some or all of the application server resources. Keep data intensive operations in an instance on the database server to reduce network lag and an increase in response time. See Chapter 3, “SAP system landscapes” on page 17 for a discussion of two and three tier landscapes.

18.2 Tools to track resource utilization

Resource usage can be tracked with several tools and services in an SAP system and i5/OS implementation. Three tools are outlined in this section.

- IBM Insight

  The IBM Insight for SAP utility program and its subsequent analysis process and report are designed to provide a high level and convenient workload analysis for an in production SAP system complex. The analysis includes the actual active user counts, machine utilizations, user and module load distributions, dialog counts, information about batch and reporting usage, system information, and database information.

  The IBM Insight for SAP utility program is packaged as an all-in-one Microsoft Windows 2000 SP3 and Microsoft Windows XP install image. It is ready for installation on a Windows/Intel® system capable of communicating with the production SAP complex. Documentation is included with this software. IBM Insight for SAP is validated to run against SAP production complexes on System i servers.

  Install the utility on your PC to use IBM Insight. Set up the initial communication parameters, ensure authorization and access, and then begin the recording session. The software continues to record performance data from the SAP complex, using SAP's RFC functionality until the session is ended. At this point, the collected data is forwarded via e-mail to IBM for reduction, analysis, and customer report production. Send customer collected data sets, as well as requests for further information or support, to: eSizings@us.ibm.com.

  Collect three days worth of data during a period of the month with reasonably high usage. Only the data from the production system is analyzed.

  Refer to 7.2.5, “Size upgrades” on page 46 for more information about sizing. Refer to 27.1.1, “IBM Insight” on page 568 for more information about IBM Insight, and the following Web site for instructions on how to download the Insight tool:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS381

Consider the following complementary services:

- SAP EarlyWatch. The SAP EarlyWatch Check is a service from SAP that analyzes the components of your SAP application, your operating system and database to determine how to optimize performance and keep the total cost of ownership to a minimum. It can also be used to identify usage patterns and determine where growth is occurring in your SAP systems.

  Refer to 27.1, "SAP EarlyWatch" on page 566 for more information about SAP EarlyWatch service.

- Performance Management (PM) (Performance Tools (5722-PT1) for analysis. These tools track System i utilization in detail and can be helpful to identify hardware bottlenecks. Refer to Chapter 27, "Performance management tools and services" on page 565 for more information.

18.3 System i resource upgrades

Remember that CPU, memory (sometimes referred to as main storage) and disk all work together to achieve a balanced system. A bottleneck can move from one part of the system to another if a significant increase is made for one resource without accounting for resource utilization in other areas.

You might have hardware that is out of date or is nearing the end of a leasing term. Use IBM Insight to collect your current system utilization. Then use the this information to determine an appropriate size system when you consider a new hardware configuration.

Perform a complete save and restore when updating to a new set of hardware. Or install the operating system and then perform a homogeneous system copy procedure.

See 11.3, "Set up an SAP system by system copy" on page 134 for a discussion of homogeneous system copies.

This section further discusses typical hardware configuration scenarios.

18.3.1 Adding memory or disk resources to a partition

Adding memory resources to a partition is almost guaranteed to provide some benefit because they increase the amount of data that can be quickly accessed by the processor without requiring a disk operation. The objective is to allow sufficient memory so there is minimum paging activity in a given subsystem.

Adding more I/O adapters and disk units can also be beneficial because any disk load that does exist can be serviced by more resources. Use STRASPBAL to rebalance the disks in the ASP when new drives are added.

18.3.2 Adding CPU resources to a partition

There are several ways to add CPU resources to a given partition. Before adding CPU resources, consider whether you have enough memory and disk arms to support the additional CPU resources

You can increase the available processor resources in a partition by permanently adding physical CPUs by way of a hardware upgrade.

Another way to increase the amount of hardware resources is to take advantage of the processor virtualization functionality offered by your System i configuration. It is possible to manually move CPU resources from one partition to another.

Still another way is by utilizing System i Capacity Upgrade on Demand (CUoD) features. Depending upon the model and the number of CPUs that you are currently entitled to use, there can be additional physical CPUs in your System i that are present but not active. CUoD software keys allow you to activate some or all of the available processors for a defined period of time and then you can allocate them for use by a partition on the system.

Finally, System i hypervisor technology also allows CPU virtualization. This involves automated dynamic sharing of a pool of processors among two or more logical partitions. Guidelines regarding minimum and maximum CPU entitlements can be defined by the system administrator and the hypervisor enforce them automatically. This can be a very effective way to utilize the full capacity of your server and to automate the needs of peak workload times on your system.

A virtualization study paper is available that outlines the use of the shared processor pool in a virtualized SAP system environment. Refer to the *SAP Solutions on IBM eServer™ iSeries* whitepaper at:

http://www.ibm.com/servers/eserver/iseries/perfmgmt/pdf/sapvirti5.pdf

Be aware that the SAPOSCOL and ST06 behavior does not seem consistent when using the shared processor pool because values greater than 100% utilization can be reported. See 27.3.1, “SAP OS Collector” on page 577 and 20.6.10, “SAPOSCOL” on page 360 for information about SAPOSCOL. See SAP note 787127 for further details.

18.3.3 Instance profile tuning changes

Revisit SAP application tuning whenever hardware changes are made, particularly to CPU or memory. For example, if CPUs are added, you might be able to add more work processes to the overall system workload. Consider increasing the buffer settings or the number of tables that are buffered when memory is added.

18.4 Software release updates

When operating system or SAP application software is upgraded to a new release, take the opportunity to reverify the recommended size and capacity of the CPU, memory and disk capacity of your hardware model, and functions of an operating system version. Changing to a later software version can be enough to trigger a performance problem if a serious problem is imminent or your system is currently undersized.

Refer to 7.2.5, “Size upgrades” on page 46 when considering an upgrade.

18.4.1 i5/OS release upgrade

i5/OS release upgrades usually do not require additional hardware resources. However, changing the operating system can expose a problem with an undersized system.

Normally the development and test partitions are upgraded before the production partition is upgraded. Remember that data is often accessed between partitions, either through QFileSvr.400 to the shared transport directory, or between a 3-tier application server and the database server. Therefore, consider applying current PTFs to the partitions that are not updated to a newer operating system release so that the PTF level of these partitions is compatible with other partitions. Running different operating system release levels in your network has less risk of incompatibility if the PTF levels are near the same level.

The recommended steps when upgrading software are as follows:

1. Verify that installed hardware is still supported.
2. Upgrade the kernel level prior to the operating system release upgrade. See SAP note 68440 for release and patch level combinations by kernel level.
3. Refer to the Information Authorized Program Analysis Requests (APARs) to get a list of recommended fixes. See 19.3.1, “IBM Informational APARs for SAP customers” on page 332.
4. Check SAP notes for the i5/OS release you are upgrading to. For example, specific information regarding i5/OS V5R3 is described in SAP note 743113.

   Consider any planned SAP release upgrades as well. Some SAP applications require a specific i5/OS level as a minimum.
5. Install any secondary languages that were installed previously. See SAP note 805447.
6. Tune individual queries, as necessary.

18.4.2 i5/OS database upgrade

The database version on System i servers is upgraded when you upgrade i5/OS to the latest release. A separate database upgrades is not required on System i configurations because the database (DB2 UDB for i5/OS) is integrated into the i5/OS operating system.

18.4.3 SAP release upgrade

The complexity of the SAP upgrade process increases with the number of modifications you have made to the SAP standard installation. The upgrade process is least difficult when you have not made changes to the SAP standard. In any case, there can be many application changes made by SAP in a particular area in a particular release that can significantly change the internal processing, data structures or even the user interface. The opportunity for customizing and the potential for change as a result of the release upgrade highlight the need to test and plan for any SAP release upgrade.

Note that a later SAP release level often requires additional hardware resources as compared to the prior release. SAP sizing guidelines for the releases you are coming from and going to should be reviewed.

Depending on your skill level with SAP systems, consider bringing in an experienced consultant to perform the upgrade. Carefully review the upgrade guide and all SAP notes prior to planning and attempting an upgrade.

There are few differences between the SAP release upgrade procedures on the System i platform and other platforms. However, if you are upgrading from a release based on 4.6C or prior Web Application Server technology using the EBCDIC code page, a codepage conversion must be performed when upgrading to 6.x or a later Web Application Server technology, as described in 18.7.2, “Code Page Conversion (CPC)” on page 325.

The following list gives an overview of technical considerations for an SAP release upgrade.

- Apply the i5/OS PTFs as listed in the Informational APAR prior to performing an SAP release upgrade.
- Consider upgrading to the latest support packages for the new SAP release

- For release upgrades:
  - A codepage conversion might be required
  - Resize the system. SAP releases typically require additional hardware resources.
  - The Central Instance must run on the Database Host during the upgrade process, which can require you to make configuration changes if it is currently running elsewhere in the system.

See the upgrade guide for information about SAP system upgrades. It can be found at:

http://service.sap.com/instguides

Click the product name, and then the release level.

18.5 Add SAP systems

You might decide to add a new SAP system to an existing partition. You can decide to copy a system from somewhere else by using the homogeneous or heterogeneous system copy procedures. You might also decide to add a new system to an existing partition by using standard installation tools.

Be aware that additional resources are used when you start an instance for a new system when it is added. Additional memory for each additional instance that is running within a partition is recommended. In addition to memory, sufficient disk space for the system is required. You might need additional disk units if they are already busy with the current workload. If you already have a separate ASP for journal receivers, make sure that it is capable of absorbing the additional workload from a capacity and arm utilization perspective.

See 11.3, “Set up an SAP system by system copy” on page 134 for more information.

18.6 Adding languages to SAP applications

If you choose to install additional languages, additional storage is required. Refer to the language guide for more information, which is found at

http://service.sap.com/instguides

Click the product name, and then the release level.

Use the Create SAP Locales (CRTSAPLCL) command after the language is imported if you are using a GLS system and kernel release of 6.20 or earlier.

The CRTSAPLCL command is not necessary with kernel release 6.40 and later because single byte locales are included. Install double-byte language support for i5/OS for the particular language you want if you want to use double-byte languages within SAP applications. See SAP note 807912 for details.

18.7 Migration and code page conversion

This section discusses the differences between migration and Code Page Conversion (CPC).

Migration is what one does to change hardware, operating system and database platforms. It implies a complete unload of the source system, and a reload on the target system. A CPC is simply a conversion of columns in database files from one codepage to another within the same database. For systems in the Latin-1 codepage, a CPC does not imply any complete unload or reload to a flat file format, nor does it imply a platform change.

An introductory discussion on this topic is found in Chapter 6, “Encoding data used in SAP systems” on page 35.

18.7.1 Migration

Migration refers to the process of migrating an SAP application from one operating system and database platform combination to another. This is also referred to as a heterogeneous system copy.

SAP requires that a certified migration consultant assists with migrations. SAP also offers an SAP OS/DB Migration Check service. This offering prepares systems for migration from another platform to the i5/OS operating system and database, and helps ensure a smooth operation on the new System i platform.

Note: Contact the IBM SAP Competence Center at isicc@de.ibm.com for assistance in finding a migration consultant in your area that has experience with migrating systems to System i models.

Use the 6.40 migration tools if you are already at the 6.40 kernel level or later. Download the homogeneous and heterogeneous System Copy Guide from the Service Marketplace at:

`http://service.sap.com/instguides`

The System Copy manual explains the prerequisites, the step-by-step procedure, and the post-processing tasks.

The migration procedure involves a complete unload of the database to an SAP system specific stream file format. These files are moved to the target platform, then the SAP tools are used again to move the data from the stream files to the target database. This load is very similar to how the SAP install tools load the database from the SAP installation DVDs.

SAP also has migration tools for releases prior to 6.40.

18.7.2 Code Page Conversion (CPC)

SAP systems with an i5/OS configuration are delivered with an EBCDIC implementation through SAP Basis release 4.6D. SAP applications based on the 6.x and later Web Application Server technology are available in the Global Language Solution (GLS) version or in a Unicode version. They are not available in EBCDIC. If you presently use an EBCDIC based solution and want to upgrade to a current SAP release level, a Codepage Conversion (CPC) is required.

The Code Page Conversion process is a fully supported offering from SAP. SAP requires that the process be performed by someone that has attended the SAP course UABC92, i5/OS Codepage Conversion. There are no prerequisites for the course, so you can choose to attend and then perform your own CPC. An SAP certified OS/DB migration consultant may also run the CPC.

The default codepage for i5/OS is EBCDIC. However, Integrated File System files created by SAP systems are stored in the same codepage as the application server. When looking at SAP Integrated File System files from an i5/OS interactive job, the i5/OS Display File (DSPF) and Edit File (EDTF) commands automatically convert the files to a readable format regardless of the codepage in which they are stored.

GLS

The GLS codepage originated with the availability of level 4.6C and is available for Enterprise Central Component (ECC) 5.0 and most other current applications that use the 6.x kernel. Some SAP applications are available in only a full Unicode implementation.

The GLS version features an ASCII application server with Unicode (UCS-2) tagged database files. It is sometimes referred to as the “ASCII version”. A full Unicode version is also available which uses Unicode for both the application server and the database server.

GLS works with more different languages than EBCDIC does. GLS can handle double-byte Asian languages. However, there are limitations on what combinations of code pages are allowed to be used together. For example, west European, east European, and Asian. Since the GLS version is ASCII based, GLS application servers connected to your i5 database can run on i5/OS, Windows, or Linux on POWER servers.

The CPC takes an SAP database at the 4.6B or 4.6C EBCDIC release level and converts the code pages of specific fields in the database to the UCS2 format. The bulk of the process is performed as a series of simple statements in SQL. It takes advantage of multiple processors on your machine by running in parallel jobs. For Latin-1 systems, it does not require a complete unload of the database to a flat file format, so it does not require a reload. For systems in other code pages, an export and then an import are required.

The resulting database is approximately 50-70% larger than the original database since a large number of fields are converted from single byte (EBCDIC) type to a double byte (UCS2) type. Note that the extra capacity requirement does not cause an extra disk expense line of 50-70% because higher capacity disk units are usually only marginally more expensive than lower capacity drives and that a single new disk unit of larger capacity might be able to replace the capacity of multiple older disk units if extra or replacement disk units are needed. Multiple disk units might be needed to provide enough input/output bandwidth. Remember that disk drive performance is also dependent upon the number of disk arms and the number and type of IO adapters used in a configuration. Size the system accordingly.

See 7.2.3, “Work with Quicksizer output” on page 44 for additional considerations when converting code pages.

> **Note:** Documentation for the Code Page Conversion is available on the SAP Service Marketplace. See SAP note 500950 for the latest information.

Unicode

Also known as “full Unicode”, Unicode is available for SAP ECC 5.0 and all other solutions built with the 6.x kernel. SAP ships both the GLS and Unicode DVD media in the installation package. You must decide which version you prefer to use.

A Unicode solution offers the most flexibility for using multiple languages in a single system. Like GLS, it can use i5/OS, Windows or Linux on POWER as application server platforms connected to an i5/OS database.

In a Unicode implementation, both the application server and database are in Unicode (UCS2). Additional memory is required for the application infrastructure because data within the application server is now represented with two bytes per character, as compared to GLS and EBCDIC application servers which represent data within the application server with one byte per character.

Conversion is required to go from GLS to Unicode. Some additional database fields are converted to UCS2 format in a Unicode solution as compared to the GLS, so there is a little database growth. The complexity of converting from GLS to Unicode varies depending on whether you are using the Latin-1 codepage or not. If you are using Latin-1, there is a simplified “In-Place Procedure” described in SAP note 800791. If you are not using the Latin-1 codepage, a conversion through Export/Import is required.

An SAP Codepage Conversion Check is available to support codepage conversion from EBCDIC to GLS formats, and any necessary release upgrades.

Through a combination of release upgrades and CPC operations, it is possible to move from a 4.6B or prior EBCDIC system to a later SAP release level running with a Unicode application server and database.

Refer to 7.2.3, “Work with Quicksizer output” on page 44 for sizing considerations.

A technical description of GLS and Unicode

Encoding characters into bit patterns has been a matter of concern for the information technology (I/T) industry since its inception.

While IBM promoted the Extended Binary Coded Decimal Interchange Code (EBCDIC) encoding on the System z™ and System i configurations, the American Standard Code for Information Interchange (ASCII) encoding was developed from teletype encoding. ASCII is used on UNIX and Windows-based systems.

> **Note:** The sort sequence of EBCDIC and ASCII encoding is different.

There are 256 different combinations if one byte is used to encode a character. These combinations are not sufficient when you take into consideration the different languages and the need for control characters. This leads to the concept of code pages for different languages, for example, a code page for Latin-2. Asian languages typically have more than 256 characters. More than a one-byte encoding scheme is required.

A specific bit pattern can represent different characters in different code pages. Figure 18-1 on page 329 provides an example for this fact. If you want to have a code page, which might be of use for more than one language, you might choose the following approach. All characters which are commonly used in all of these languages are defined. Only a subset of the characters used in just one of these languages have a bit pattern assigned. A code page which serves this purpose is called a blended code page. It can be used for more than one language, but you might not be able to have a bit representation of all characters and symbols used in these languages.

Creating a universal encoding scheme that provides definitions for all of the characters in living and dead languages can solve these limitations. The Unicode Consortium was formed to come up with a suitable solution. The consortium came up with the following standards:

- A unique definition of a code point for each possible character and symbol
- Three different encoding schemes for all code points:
  - UTF-8, which uses one byte as a building block. Encoding a code point requires up to six bytes.
  - UTF-16, which uses two bytes as a building block. Encoding a code point requires either two or four bytes.
  - UTF-32, which uses four bytes as a building block. Every code point can be encoded in one building block.
- An algorithm to convert from one encoding scheme to another encoding scheme

We could observe the following:

- UTF-8 is very compact for normal ASCII encoding. Most common characters require one byte. However, encoding a character from an Asian language requires more than two bytes.
- UTF-16 uses only two bytes to encode any character from the living languages. However, the sort sequence of characters differs from that of UTF-8.
- If only living languages are concerned, UTF-16 would require only two bytes. If we can restrict the definition domain in the sense that 2 bytes encoding of UTF-16 is sufficient, we could use a more simple implementation. An International Organization for Standardization (ISO) norm named UCS-2 defines this type of implementation.
- On the other side UCS-2 is a proper subset of UTF-16. Although UCS-2 shares all the features of a Unicode encoding scheme, it cannot encode all the defined Unicode code points. It can, however, enclose all the characters that can be used with all the living languages in a commercial application. That is, you can use UCS-2 for Chinese or Japanese symbols, but not for the symbols used in Egypt 2,000 years ago.
- An SAP system uses UTF-16 internally for its Unicode implementation.
- i5/OS implements a UCS-2 encoding. It is used to store data in the database to support applications that use ASCII encoding, for example, the SAP systems encoded in ASCII. This is called a global language system (GLS).

> **Note:** SAP applications up to SAP BASIS 4.6D on system, support EBCDIC encoding. SAP BASIS 4.6 C supports ASCII or EBCDIC encoding. NetWeaver-based systems support either UCS-2 or ASCII encoding.

The results of the Unicode Consortium are published in *The Unicode Standard, Version 4.0* by The Unicode Consortium, ISBN-0321185781.

Encoding schemes in SAP systems on System i

Different versions of SAP systems have various encoding schemes, as follows:

- SAP versions prior to SAP BASIS 4.6 are in EBCDIC
- SAP version BASIS 4.6 are in EBCDIC or ASCII (GLS)
- SAP version BASIS 6.20 systems are in ASCII or Unicode, that is SAP Enterprise or SAP ECC or later.

Figure 18-1 on page 329 illustrates the translation process for different encodings.

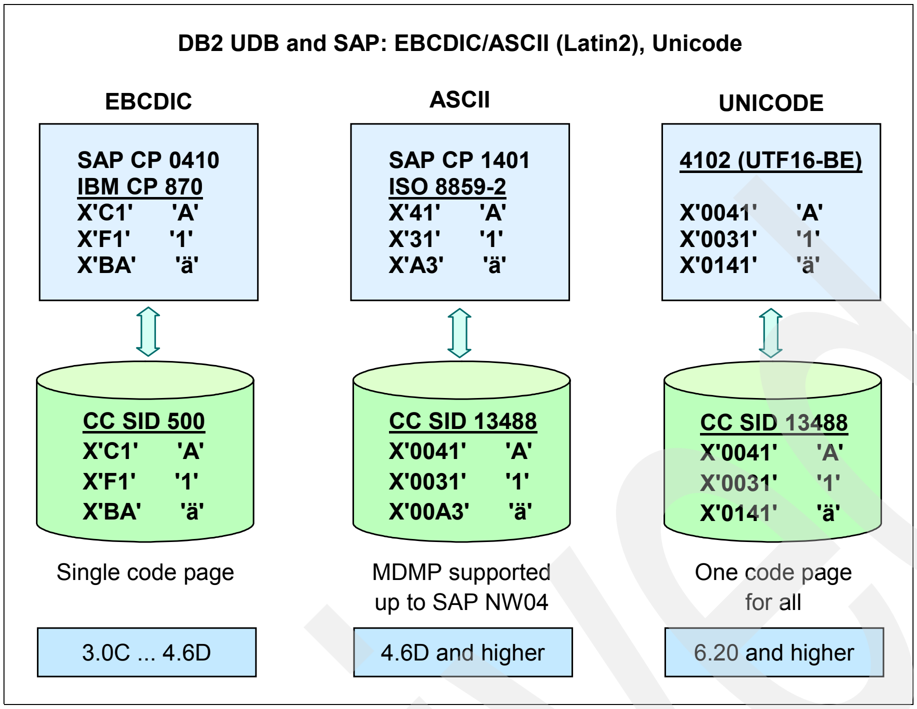

Figure 18-1 Code pages for SAP systems running on i5/OS - Source SAP

This illustration is based around three characters: “A”, “1”, and “.”. The encoding in the database is represented as systems of red disk devices. The coded character set identifier (CCSID) defines the code page and the characteristics for the translation process to other CCSIDs. CCSID 870 is for the EBCDIC, and 13488 is for the GLS and Unicode system. 13488 is the UCS-2 CCSID.

Some of the other details illustrated in Figure 18-1 are:

- On the application side, for EBCDIC, the code page is 870 for IBM and 0410 for SAP. The model characters preserve the encoding.
- For the ASCII system, the interpretation of the encoding is given by the codepage 1401 for SAP especially for the character X’A3’, which would have a different interpretation in the ISO 8859-2 code page and the SAP codepage 1401.
- The mapping of the ASCII codepoint X’A3’ to the representation on the database is done via the CCSID 13488.
- For the Unicode system with the mapping to UTF-16, the respective code pages for SAP are 4102 if running on System i models or 4103 if the application server runs on Windows. The character encoding is preserved.
- ASCII encoding requires 70% more disk space compared to EBCDIC encoding. Unicode requires more central processing unit (CPU) and random access memory (RAM) resources.

Different code pages and CCSIDs are important factors to consider when exchanging SAP data with programs written in traditional programming languages. Using traditional programming languages to manipulate data residing in an SAP system requires that you follow certain rules. For an explanation of the rules and to see coding examples, refer to Chapter 29, “Data exchange scenarios and examples” on page 621. A Report Program Generator (RPG) example is found on 29.2.1, “RPG/400 example” on page 625.

For more information about CCSID and globalization, refer to the following Web site:

http://www-03.ibm.com/servers/eserver/iseries/software/globalization/

19

Problem avoidance

This chapter presents ways that you can help yourself avoid problems and, should problems arise, solve them faster. We focus on the following specific points:

- How to avoid hardware and software errors
- Where to look for information about potential issues
- What to do to help get timely service

19.1 Prevent instead of react

Your SAP system environment constantly develops and changes. Users do not do the same things in the same way each day. The number of rows in a database changes day by day. Your system is different today than what it was yesterday, even if you do not make significant changes, like adding systems or performing maintenance. Many changes go unnoticed. Are you close to an SAP application limit? Is a hardware failure imminent?

A business can suffer if problems with the system are allowed to linger. Successful businesses effectively handle problems and move forward. You can avoid many problems by taking preventative action. Enlist the support of SAP and IBM to help make sure that your system continues to run smoothly, so that your business can remain successful.

Your decision to implement SAP applications on the System i platform simplifies your system maintenance strategy. With a System i implementation, you can call one place to resolve any issue related to the hardware, the operating system, or a database situation. Some SAP platform combinations require as many as three or four vendors.

In this chapter, some of the components involved in a preventive approach to problem management is discussed.

SAP and IBM resources to help in problem resolution are summarized in Appendix A, “Support for SAP applications” on page 711.

19.2 Hardware

Your System i server has excellent built-in automated hardware error detection. A component can start to show signs of trouble. The System i server logs the error, and in many cases can “phone home” to IBM to request service. You are notified through the QSYSOPR message queue and the Work with Problem (WRKPRB) command. Many hardware components can be replaced while the system is running with little or no impact to your end users.

Diagnose hardware issues immediately.

Further information regarding hardware support can be found at:

http://www.ibm.com/servers/eserver/support/iseries/

19.3 IBM software

A list of the IBM software installed on your system can be viewed by using the command:

```
GO LICPGM
```

A list of the IBM software fixes applied to your system can be viewed by using the command:

```
DSPPTF
```

Press Enter to progress from page to page to view the information for each IBM software product.

19.3.1 IBM Informational APARs for SAP customers

A Programming Temporary Fix (PTF) is a software fix to an IBM program. IBM provides sets of PTFs in Cumulative PTF packages that cover all aspects of i5/OS. Group PTFs are sets of

PTFs that relate to a certain area of i5/OS, for example database or Java. A special group of PTFs is called the Hiper group. This group contains PTFs that correct problems of a high impact or pervasive nature (highly pervasive).

Lists of i5/OS fixes that are recommended for SAP customers are maintained by IBM in a set of Informational Authorized Program Analysis Requests (Info APAR). There is an Info APAR for SAP applications on System i models for each i5/OS release.

The Info APAR for "SAP on System i5" contains the current cumulative PTF level, the most critical group PTFs for SAP, and a number of individual PTFs that are not included in any group at a given time. The info APAR also lists downloadable IBM code that can be useful in an SAP system environment running as a System i solution, including tools, the JDBC Toolbox, and database client code for non-i5/OS application servers.

Table 19-1 lists Info APAR numbers for the latest i5/OS levels.

Table 19-1 Info APARs for i5/OS and OS/400

Operating system release	Informational APAR number
i5/OS V5R4	II14126
i5/OS V5R3M5	II14125
i5/OS V5R3M0	II13868
OS/400 V5R2	II13337
OS/400 V5R1	II12833

These Info APARs are updated frequently to add fixes that are known to be of benefit to SAP applications.

19.3.2 Frequency of PTF application

It is important to apply PTFs on a regular basis. The frequency of PTF application depends on the requirements and preferences of a particular environment. You might have more freedom to apply PTFs to test environments than to update your production environment. This is a benefit to establishing a test environment, so that fixes to your environment can be tested prior to the affected program reentering production. See 19.5, "Testing" on page 334.

Some PTFs that are listed in the Info APAR are for unique situations that might not be applicable to your environment. It is not necessary to apply PTFs each time the Info APAR for SAP changes. It is recommended to include all of the PTFs that are listed in the Info APAR when you apply maintenance, because they are known to be helpful in SAP system environments.

Consult the IBM Preventative Service Planning (PSP) information at the following Web site when planning your PTF maintenance strategy:

http://www-912.ibm.com/s_dir/sline003.nsf/sline003home

The PSP "bucket" contains links to cumulative and group PTF information, including the scheduled date for the next cumulative PTF package.

At a minimum, it is highly recommended to keep up-to-date on the Hiper PTF Group, since this group contains the most critical fixes for the system.

19.3.3 IBM software support

Make sure to obtain an IBM Software Support contract that corresponds to the needs of your business. SAP systems are usually critical to business operations. Note that SAP Support does not provide service for any IBM Software. Various IBM Software Support contracts are available, each with different terms and conditions. Offerings can vary by geographical region.

IBM resources to help in problem resolution are summarized in Appendix A, “Support for SAP applications” on page 711.

19.4 SAP software corrections

All SAP software fixes are obtained from the SAP support portal at:

http://www.service.sap.com

SAP application code corrections are delivered as Support Packages. Corrections to the infrastructure code are delivered as Binary Patches. Both of these types of fixes can be downloaded independently from the SAP support portal.

In 2003, the SAP Support Package strategy was extended for certain product versions to Support Package Stacks (SP Stacks). An SP Stack contains an optimal combination of Support Package and patches for the individual components. Use the SP Stacks to reduce the risk of incompatibility between SAP elements.

19.4.1 SAP Support Marketplace for Notes and Hot News

SAP offers TopNotes and HotNews through the SAP Support Marketplace portal. This Web site for SAP service is available to all SAP customers. A user ID and password that is supplied by SAP is required to use the service. You can access the SAP Support Marketplace at:

http://www.service.sap.com

To navigate to the TopNotes and HotNews topics, choose the “Problem Solving” tab from the main Service portal page. These features allow you to select products and components of interest so that you can identify the most critical notes. For example, the BC-DB-DB4 component is concerned with specific notes about the i5/OS database platform.

> **Tip:** Read the relevant SAP notes carefully before you take action. If there appears to be a conflict between notes, contact SAP for clarification.

19.5 Testing

The optimal test environment for avoiding problems is an exact copy of your production environment. IBM and SAP both have significant test cases to try to identify issues before they leave the development lab. However, your environment is unique to your system. No one can test your production environment as well as you can.

Increasing the similarities between the test and production system reduce the severity and the number of surprises that can happen once changes are made to the production system. Is the hardware similar? Are the data sets the same size? Do they have the same content? Are instance profiles the same?

Effective workload simulation is also important. Computer systems behave differently under higher loads. A handful of users and a handful of data is a test, but this might not be adequate to simulate a production system running with many more users and much more data.

One commonly overlooked testing effort is in the area of disaster recovery and high availability. Develop a documented procedure to recover your SAP systems from media in the event your environment experiences a catastrophe. Plan regular failover testing of your primary system to your backup system that includes running on the backup system for some time before swapping back again. This test is a good insurance policy.

Refer to Chapter 24, “Disaster recovery and high availability” on page 433 for further information about availability and recovery strategies.

19.6 History

Maintain a history of all changes that are made to your system. This helps with problem determination in the event something goes wrong in the system, because you are able to explain to your consultant, or an SAP or IBM Service representative what changed, when it changed and what the situation looked like before the error.

Keep documentation on the date and identifier for each of the following maintenance items:

- SAP notes applied
- SAP Support Package applied
- SAP Binary patches applied
- Operating system cumulative, group and individual PTFs that are applied
- Database client changes on non-i5/OS application servers

Keep a log of the historical performance information so that you can tell what performance was like before a change.

The following are useful sources of information to review past performance characteristics:

- Early Watch Alert Analysis
- IBM Insight data
- i5/OS performance monitoring data

In addition to knowing performance characteristics are on a normal day, it can also be useful to note any special peak workloads like month or year end processing.

See Chapter 25, “Performance concepts and monitoring” on page 457 for information about monitoring system performance.

19.7 Avoid more problems once one is identified

You might not be able to avoid every problem, despite everyone’s best efforts. However, the choices you take can help get your environment back to normal again as quickly as possible. To help avoid obstacles in the service process, consider the following process:

- Do the first search yourself

  If you have easily recognizable symptoms for a problem like an error identifier or some specific text string, perform an online search of the SAP notes and the IBM PTF cover letters. You might quickly find a match and the solution.

- Open a problem record

  See Chapter 20, “Problem determination and management” on page 339 to understand what is involved in reporting a problem. Understanding what happens when managing problems can make the process more efficient.

  > **Note:** Effective communication between the client and support personnel is critical to manage problems in an efficient manner. Getting the right information to the right person is the most critical component of resolving an issue.It is beneficial to draw a picture of your SAP and System i implementation so that the IBM or SAP service representatives who assist in problem management can better understand the environment involved.

  Problems are described to SAP using an SAP Message, and to IBM using a Problem Management Record (PMR). Each and every problem is assigned a reference number.

  When you first report your problem to a service representative, their first course of action is to determine if it is a new problem or if it has been reported before. Be as specific as possible in the initial problem description. Include facts like error codes, the SAP and IBM program names involved in the failure, and statement numbers of the programs whenever they are available. These facts can be effective keywords for service representatives or developers to use when searching for the solution to the problem. Include the i5/OS release level and the number of the Cumulative PTF package that is applied, the Hiper Group and Database Group PTF levels, the SAP release, hot package levels and the kernel level of your SAP and System i installation in your report.

  When opening your problem, consider mentioning to your IBM Service Representative that a Remote Technical Assistance Information Network (Retain) item #1695 exists to provide a template for collecting basic information related to SAP applications on i5/OS.

  With good information, it can be quick and easy to determine which PTF or SAP fix is required to solve a problem.

  While waiting for a response from service, it can be useful to obtain the latest PTFs. In the event the problem has already been identified and a fix is available, the required PTF has already been downloaded. If the fix is not known, you might eventually need the interim fix anyway, so this effort is not counter productive.

  There is no guarantee that a resolution to a related problem will resolve your problem. In this situation, you can try the suggested fix before diagnosing the problem further.

- Work with official service organizations to resolve your problem

  Additional information is required if the problem appears to be new or is a unique issue. Answer the following questions to help to speed up problem diagnosis:

  - Can the problem be reproduced? If yes, how?
  - Does the problem occur all of the time or intermittently?
  - Are connections to your system available to allow access for remote service?
  - Are the connections open and are the userids and passwords known to the service personnel?

  A unique problem scenario can be described only with specific information that is obtained from your system. You must provide service access to the environment and be able to describe the problem to the service representative. IBM or SAP can provide fixes to unique issues only when they have sufficient evidence to diagnose the problem.

  Some problems are more difficult than others to reproduce or detect, and can take more time.

Note: Usage of internet forums can be good for sharing best practices and other types of information, but they are not always the best place for you to try to diagnose a particular problem.

It can be tempting to post a problem situation in the hopes that the problem can be recognized and resolved. If you encounter a unique SAP system or IBM issue, no one on the internet is in a position to perform the required level of investigation on your system other than SAP or IBM.

Keep in mind that non-IBM and non-SAP posted replies can be wrong, and even damaging to your system. Limiting your problem management to this type of unofficial source leaves the situation “unreported” in the eyes of SAP or IBM.

19.8 Education

There are several places to find information about education sources about SAP implementations with the System i platform. Keep up with current information to help make effective operations and management choices, and to help avoid problems.

Consider the following education options:

- Technical courses

  SAP and IBM offer a technical education course for individuals who plan, implement, operate, or manage SAP applications with i5/OS called “SAP NetWeaver on DB2 UDB for iSeries: Database Administration”. The target audience for this course is an experienced technical individual, with knowledge of either basic SAP systems or System i architecture. The course discusses many aspects, ranging from basic SAP system and System i concepts, to installation, operation, database management, backup and restore, including performance management, and tips for basic troubleshooting. The course is a combination of lectures and hands-on sessions.

  This course is part of SAPs education curriculum and certification path for basis consultants.

  Consult the IBM or SAP course catalog for dates and locations. The SAP course number is ADM525 and the IBM course number is AS525.

- International SAP IBM Competence Center

  The International SAP IBM Competence Center (ISICC) has a team of personnel dedicated to working on SAP implementations with i5/OS. The ISICC is the central point for IBM, SAP, and customers in all aspects of SAP projects. To contact the ISICC, send an e-mail to isicc@de.ibm.com or visit:

  http://www.ibm-sap.com

- IBM briefings and events

  A customized briefing in the Rochester Executive Briefing Center in Rochester, Minnesota or in the International SAP IBM Competence Center in Walldorf, Germany can be arranged through your partner or IBM representative.

  Informational meetings for System i customers with SAP implementations are periodically sponsored by IBM and SAP. Contact the ISICC at isicc@de.ibm.com to see if any such events are available for you to attend.

- Internet user groups

  The user group community can be an effective way to ask questions, share information and, in general, learn about what other SAP clients are doing with an i5/OS and SAP application. Information gained can help you avoid and better manage problems in the future.

IBM and SAP sometimes post information to internet user groups in order to reach as many clients as possible in a timely manner.

Independently validate anything read in the user group community before using it in a professional manner within your business. Be cautious. It has been known to find errors, unintentional or intentional falsehoods, or omissions of critical details posted to some user forums.

SAP and IBM resources, including education courses and online information to help in problem resolution are summarized in Appendix A, “Support for SAP applications” on page 711.

20

Problem determination and management

This chapter is intended to help you diagnose and manage problems that may be encountered when running an SAP application with i5/OS on an IBM System i model. Refer to 15.9, “Resolving printing problems” on page 248, for help in resolving printer problems, and Chapter 19, “Problem avoidance” on page 331 for a discussion of how to prevent problems, or manage them effectively. Chapter 21, “Correctable scenarios” on page 365 may also interest you.

20.1 Implementing SAP applications on i5/OS

This section describes how most SAP applications that are based upon the SAP Web Application Server technology are implemented as a System i solution. It also describes how to find the server jobs on the operating system (OS) level that are used to run the Web Application Server.

The SAP application is started with the STARTSAP command and ended with the STOPSAP command. The STARTSAP command has the Default Spool File (DLTSPLF) parameter set to a default of *YES. This causes all spooled files that are created during the previous run to be deleted. These spool files are helpful for program diagnosis. If the SAP system is stopped and restarted for debugging, use the following command to start the SAP system so that the spool files remain available:

```
STARTSAP DLTSPLF(*NO)
```

20.1.1 Job structure

When you run the following Work with Active Jobs command, (WRKACTJOB SBS(R3_*<nn>*), where *<nn>* is the instance number, the display shows all of the SAP jobs for that instance running on the System i server in a particular subsystem. These jobs correspond to the following SAP processes:

- Startup job
- Message server
- Dispatcher
- Gateway
- Work processes

The number of each type of work process 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 System i model also runs the SAP OS collector (SAPOSCOL) process.

The STARTSAP command first starts the SAP subsystem (instance) and submits a background job to this subsystem that runs the SAP start program (SAPSTART). The SAPSTART program starts the SAP services and processes associated with each instance, as shown in Figure 20-1 on page 341. Instances are described in Chapter 13, "SAP instance profile parameters" on page 173.

```
 Opt  Job         User        Type     -----Status-----  Function
      DISP_WORK   ERP03       BATCHI   ACTIVE            PGM-disp+work
      GWRD        ERP03       BATCHI   ACTIVE            PGM-gwrd
      ICMAN       ERP03       BATCHI   ACTIVE            PGM-icman
      IGS400      ERP03       BATCHI   ACTIVE            PGM-igswd_mt
      IGS400      ERP03       BATCHI   ACTIVE            PGM-igsmux_mt
      IGS400      ERP03       BATCHI   ACTIVE            PGM-igspw_mt
      IGS400      ERP03       BATCHI   ACTIVE            PGM-igspw_mt
      MSG_SERVER  ERP03       BATCHI   ACTIVE            PGM-MSG_SERVER
      RSLGCOLL    ERP03       BATCHI   ACTIVE            PGM-rslgcoll
      RSLGSEND    ERP03       BATCHI   ACTIVE            PGM-rslgsend
      SAPSTART    ERP03       BATCH    ACTIVE            PGM-sapstart
      WATCHDOG    ERP03       BATCHI   ACTIVE            PGM-disp+work
      WP00        ERP03       BATCHI   ACTIVE            PGM-disp+work
      WP01        ERP03       BATCHI   ACTIVE            PGM-disp+work
      WP02        ERP03       BATCHI   ACTIVE            PGM-disp+work
      WP03        ERP03       BATCHI   ACTIVE            PGM-disp+work
```

Figure 20-1 SAP jobs in the subsystem

By default, the STOPSAP command does not end the subsystem job or the SAP performance collector job. To end these jobs, specify the ENDSBS(*YES) parameter when issuing the STOPSAP command.

Job status

Table 20-1 lists all jobs in the subsystem for the SAP central instance. The Initial status field shows the status of the jobs after the instance has been started.

Table 20-1 Jobs, job type, and status

Subsystem/job	Description	Job type	Initial status
R3_<instance>	Subsystem job	SBS	DEQW
DISP_WORK	Dispatcher	BCI	SELW
GWRD	Gateway (reader)	BCI	SELW
MSG_SERVER	Message server	BCI	SELW
SAPSTART	SAPSTART program	BCH	EVTW
SAPOSCOL	SAP performance collector	BCI	SIGW
WRK<nn>	Work processes	BCI	SEMW

You should see the QXDAEDRSQL job in the QSYSWRK subsystem. The listener (Port 4402) for remote database requests from application servers are started by the **STRTCPSVR SERVER(*EDRSQL)** or the **STARTSAP SID(*DB)** command.

If the system is running in a 3-tier environment with Opticonnect, the database server has jobs named QSQnnnnnn in the QSOC subsystem that correspond to the work process jobs on the application server.

20.2 Working with job logs

Every SAP work process has a corresponding job. And every job has an associated job log. The job log for a job records the messages that are sent between the SAP 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 can be written to a spooled file where it can be viewed or printed.

This section explains how you can find the corresponding job and its job log on a System i configuration based on the SAP work process.

20.2.1 Changing the job attributes

Whether a spooled file is generated or not when a job ends depends on the job attributes and the end code of the job. Because the SAP kernel monitors for many messages, an end code of 0 is received even if an error has occurred. To receive a job log in any event, change the job descriptions used by the SAP work processes using the command:

```
CHGJOBD JOBD(R3400/R3_<nn>) LOG(4 00 *SECLVL)
```

This change takes affect after the next restart of the SAP system.

During the installation or upgrade of an SAP system, the install jobs inherit the attributes from the job that started the installation. To make sure that a spooled file is 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 sets the job attributes automatically. To force a job log in this situation, press F21 (User Window) from the R3UP screen to get a command line and enter the command:

```
CHGJOB LOG(4 00 *SECLVL)
```

Leave the window by pressing F12 and continue with the upgrade.

20.2.2 Work process overview

Use the SAP transaction SM50 (Work process overview) to display the SAP work processes, as shown in Figure 20-2 on page 343.

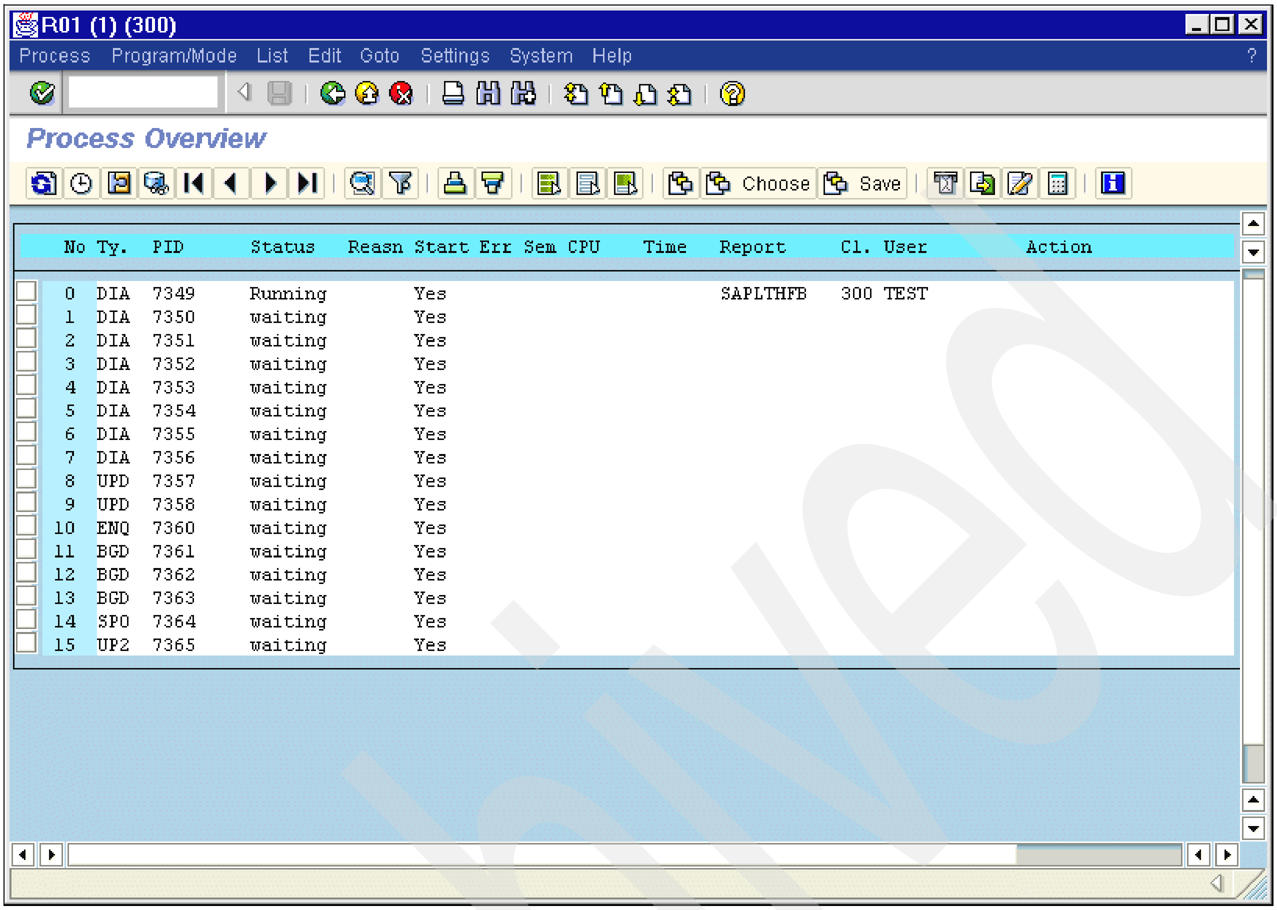

Figure 20-2 Work process overview (SM50)

Use the SAP dispatcher monitor (DPMON) if access to the SAP system through an SAP GUI is not available to process transaction SM50. Run the following command to see the same information from a 5250 screen:

```
CALL PGM(DPMON) PARM('pf=/usr/sap/<SID>/SYS/profile/<instance profile>')
```

Figure 20-3 on page 344 shows the output of this 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 20-3 Work process table (DPMON)

There are three different ways to find the job log of a dialog work process on the System i model. In this example, we want to display the job log of work process number 0.

20.2.3 The WRKPID command

Note the process ID (PID) of the job from transaction SM50. In this example the PID is 7349. Make sure that the SAP kernel 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 SAP work process to the job. The WRKPID PID(7349) command is used in our example. This command runs the WRKJOB command for the corresponding job, as shown in Figure 20-4 on page 345.

```
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 20-4 WRKPID command

Press Enter and type option **10** to display the job log of the job.

20.2.4 The database lock monitor (DB01)

Note the work process number from transaction SM50. In our example it is 0. Follow these steps:

1. Call transaction DB01 (Database lock monitor).
2. Select the instance, and click the **Execute** button. A window as shown in Figure 20-5 on page 346 appears.

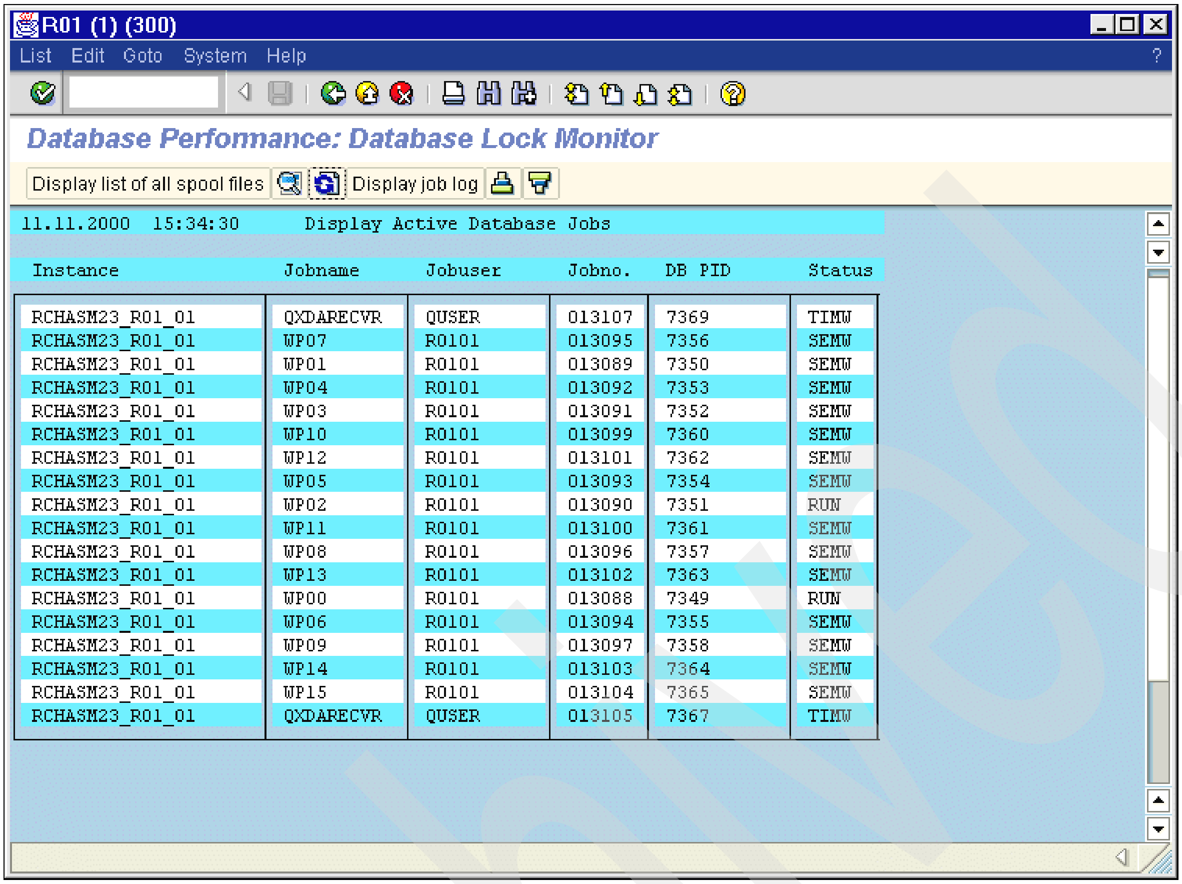

Instance	Jobname	Jobuser	Jobno.	DB PID	Status
RCHASM23_R01_01	QXDARECVR	QUSER	013107	7369	TIMW
RCHASM23_R01_01	WP07	R0101	013095	7356	SEMW
RCHASM23_R01_01	WP01	R0101	013089	7350	SEMW
RCHASM23_R01_01	WP04	R0101	013092	7353	SEMW
RCHASM23_R01_01	WP03	R0101	013091	7352	SEMW
RCHASM23_R01_01	WP10	R0101	013099	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 20-5 Database Lock Monitor (DB01)

3. Select WP00 and click the **Display job log** button.

> **Note:** Select System → List → Save → Local file to save the job log from this window. Or enter %pc. Always use the unconverted format as this is useful in the event you have to forward the job log to support personnel.

20.2.5 The WRKACTJOB command

With SAP kernel release 4.6 and later, SAP work processes can be mapped to the job by using the WRKACTJOB command. In our example, the work process number from transaction SM50 is 0. Follow this process:

1. Enter the command:

   ```
   WRKACTJOB SBS(R3_<nn>)
   ```

   Refer to Figure 20-1 on page 341 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.

20.2.6 Printing and locating the job log

If the job found by these steps is still active, you can use the Display Joblog command to obtain a spool file of the job log:

```
DSPJOBLOG JOB(qualified job name) OUTPUT(*PRINT)
```

Use the WRKSPLF 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 review the spooled file QPJOBLOG.

If problems occur during installation or upgrade of an SAP application and you cannot identify a certain job by the above steps, you can scan through all spooled files generated by the user profiles *<SID>*OFR (for startup, installation, or upgrade problems) and *<SID><nn>*. Replace *<SID>* with the SAP system ID and *<nn>* with the instance number.

Some files are named QPRINT in the list of spooled files. Some SAP functions send output to the standard output device (STDOUT) or the standard error device (STDERR). Since all of the jobs under the SAP 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.

20.2.7 SAP instance profiles

To locate the SAP profiles, use the i5/OS 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 SAP:

- Default profile

  ```
  DEFAULT.PFL
  ```

  The default profile indicates basic information about the SAP system, such as the system name, database host name. It also contains defaults for all instances in an SAP 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 specifies the SAP parameter overrides for the instance, including the number of work processes to be started.

Usually these profiles should be maintained through SAP transaction RZ10 (Profile Maintenance). In emergency cases (for example, if the SAP system does not 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.

20.3 Trace and log files

The following are SAP and IBM traces and log files that can be used for problem analysis.

- System log

  The main source of error information in an SAP system is the system log. The system log can be displayed with transaction SM21. In case of errors, the log shows the error message, information about where in the application the error occurred, 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. Reviewing the log entries using a 5250 session can be useful. In general, the entries are unreadable, unless you are familiar with the encoding scheme.

- Short dumps

  When an SAP 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 ST22 transaction

  The short dump has detailed information about the failing source statement, including the contents of some variables. The SQL statement shown in the short dump is written in ABAP syntax, which is different from the statement syntax that is executed by the i5/OS database code.

- Developer traces

  SAP trace files can be viewed using transaction ST11. 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.

  A number of files are seen 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.

  The trace level is set with the rdisp/TRACE=*level* parameter. The trace level is effective for all processes after the next time the SAP system starts.

  You can set the trace level for work processes in the work process overview named SM50. The trace level for users can be set in the user overview.

  To select the desired level and the components to be traced, use the SM50 transaction by selecting **Process** → **Trace** → **Components** in the menu bar. You can set the trace level and the components to be traced. If the files become too large before the error is logged, restart the trace in the same menu path with a “Reset files” parameter.

  The trace level can be set as follows:

  – 0: Nothing is recorded. The trace file is closed.
  – 1: Only errors are logged. This is the default setting.
  – 2: Internal activities are logged.
  – 3: Internal activities and data are logged.

  Be sure to return the trace level to a value of 0 once the required data is captured. Tracing produces some performance overhead on the system and the higher trace levels can quickly create a significant amount of data in the trace files.

- SQL trace

  The SQL trace can be switched on to capture and log the SQL statements being passed to the database management system (DB2 UDB for iSeries) in a trace file. Turn off the SQL trace function once the report you are interested in has completed. Click the “Show Trace” button in ST05 to display the trace file.

- Startup log

  If SAP does not start correctly or not at all, review the startup log that can be found in '/usr/sap/<SID>/<instance>/work/sapstart.log' file.

- Upgrade logs

  The upgrade to a higher SAP database release is done in many phases. Each phase writes its own log into the '/usr/sap/put/log' directory. A list of the upgrade phases and the logs created during each phase was provided in an appendix of the SAP Upgrade Manual until SAP R/3 Enterprise. With SAP R/3 enterprise, the R3UP program places the list in a subdirectory of ‘/usr/sap/put’ during the initial execution of the program.

- Transport logs

  The '/usr/sap/trans/log' directory contains logs from operations, such as:

  – Customer transports in the system landscape

  – Import/export of SAP corrections

  – Client import/export/copy

- XDA trace

  The trace can be activated in three different levels for one specific job, for all jobs in an SAP 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. Contact your local support to run and analyze the XDA trace.

20.4 The CHKXDA SAP tool

When you install a 3-tier system with QXDARECVR and TCP/IP with, database shadow processing performance problems can occur on 3-tier SAP systems with high CPU-utilization on the database server. Jobs with the names QXDARECVR or APIAnnnnn have a consistently high CPU-usage, even though the SAP system is not significantly used or the specially allocated work process is not active.

> **Note:** The DBSL interface allows 2-tier database access when using virtual host names. Therefore, the use of XDARECVR database shadow work processes are no longer necessary.

SAP work-processes can end during daily operations, due to time-outs during online processing, for example. Sometimes the corresponding QXDARECVR database shadow-process does not end at the same time. QXDARECVR stays active, and therefore, continues to use system resources without being attached to a work-process. The system determines that the work-process is missing when the current Structured Query Language (SQL) command completes.

There is an SAP tool named CHKXDA which you can use to end a database shadow process. It belongs to work processes which have already been ended by the system. Download the CHKXDA tool which is an attachment in the *SAP note 450351*.

Start the tool with the following command:

```
CALL CHKXDA/CHKXDASM
```

This procedure starts the CHKXDAJOBS job in the background. The job runs in the QCTL subsystem and checks for duplicate QXDARECVR jobs on the system every 30 seconds.

The CHKXDA tool should always be started on the database server by a user with DSPJOBLOG and ENDJOB authorizations. The program checks the QXDARECVR jobs in intervals to see if there is more than one server job for each work process (WP). When the tool finds more than one, the first job is ended because it is no longer used.

The default wait time is 30 seconds. In between checks, you can change the wait-time in the CHKXDA/DELAY data area, if necessary. Independent of this setting, you can see a delay of “DLY-59” against the CHKXDAJOBS job. This is required to ensure that data consistency and is not reduced.

The QXDARECVR jobs are in the QSYSWRK subsystem, and also in the QSOC subsystem if OPTICONNECT is active.

The WP trace shows the QXDARECVR job that handles the database connection for the work process, for example, an entry for a secondary connection (used by the Advanced Business Application Programming (ABAP) application programming interface (API) call). For example:

```
M Tue Jan  3 00:50:00 2006
M  AS400 API connected to ‘iSeries’. DBjob: (04851/QUSER    /QXDARECVR)
```

In a 3-tier installation, you get the qualified shadow process information from the “DbSl EGO Structure” in the trace file, for example:

```
C     Connect type                = TCPIP SOCKETS
C     pid                         = 2438071
C     dbjobname                   = QXDARECVR
C     dbjobuser                   = QUSER
C     dbjobno                     = 256155
C     wpjobname                   = WP00
C     wpjobuser                   = CHB31
C     wpjobno                     = 767716
```

The job queue used by program CHKXDASM is entered in the CHKXDA/JOBQ data area. The default is the QCTL job queue.

The currently active QXDARECVR jobs are held in the XDAJOBS table. You can display the current contents by entering the following command:

```
RUNQRY QRYFILE((CHKXDA/XDAJOBS))
```

To establish which QXDARECVR jobs belong to which work processes, and on which application server, look for entries which contain the value `*NONE` for a WP number. These jobs are not database shadow processes, they are other database jobs. If the CHKXDAJOBS monitor job finds more than one QXDARECVR job for a work process on a particular application server, the job with the lower job number is stopped automatically, since it is no longer attached to an active work process.

Refer to SAP note 450351 for download and installation instructions.

In earlier releases, when users cancelled a long-running database operation, only the work process is restarted. The database shadow process continues to process the database request and uses a significant amount of resources. The CHKXDA tool can identify the orphan jobs and end them in order to free up the resources.

As of kernel 6.20, the SAP database interface uses a cancel function that can also interrupt long-running database operations in the shadow process.

Note: The SAP system itself is now enabled to automatically end database jobs (QXDARECV and APInnnnnn). The CHKXDA tool is now considered an alternative to standard processing.

The CHKXDAJOBS tool is delivered and supported as-is.

20.5 Special considerations for SAP Java

The SAP Java environment continues to evolve. At the current time, the best place to get the latest information is from SAP note 717376 - Recommended settings for Web AS Java Engine. This note is updated with critical settings and recommendations for successful implementations of SAP Java solutions on i5/OS.

The SAP Java environment adds a new aspect to the typical Web Application Server environment. For the most part, it is similar to the Web Application Server, as follows:

- Job structure with Java

  The job structure of the SAP J2EE environment is similar to that of the Web AS.

  When running using the Java Toolbox (usually in a 3-tier environment), QZDAINIT, QZDASOINIT, QZDASSINIT jobs can be seen in the QSYSWRK subsystem where you normally see the QXDAEDRSQL jobs in an ABAP environment.

  When running with the native JDBC driver, database access occurs through the QSQSRVR jobs in the QSYSWRK subsystem.

  Job attributes can be changed to specify particular job logging options or priorities by using the CHGJOB command in the same way as for Web AS jobs.

- SAP profile settings for Java

  There are some specific settings in the SAP instance profile that pertain to the Java environment. The profile is the same one used by the Web AS.

- Developer traces for Java

  All trace files are located under the ‘/usr/sap/<SID>/<instance>/work’ path as in the Web AS environment. SAP transaction ST11 is not available in J2EE. The SAP Visual Administration tool is used instead, and from within this tool, the LogViewer is used.

- Startup log for Java

  If the J2EE does not start correctly or at all, look at the startup log found in ‘/usr/sap/<SID>/<instance>/work/sapstart.log’

20.6 Problem analysis

This section describes what needs to be considered when a problem occurs.

20.6.1 Where to look first

Perform the following checks in the order given:

1. PTFs: IBM maintains Informational APARs to keep track of the PTFs that are found to be useful for running SAP applications as a System i solution. Apply all of the PTFs listed in the corresponding Informational APAR. Table 19-1 on page 333 in 19.3.1, “IBM Informational APARs for SAP customers” on page 332 lists the Informational APARs based on the i5/OS release.

2. Kernel patch level: Keep at the most recent patch level for the SAP kernel. To determine the installed patch level, use the SM51 transaction. Position the cursor on the line with your server, and click the Release Information button. Information about how to install kernel patches can be found in SAP note 49365. See 16.1.3, “Monitoring an SAP server: Transaction SM51” on page 256 for a further discussion of SM51.
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 System i model, like the example shown in Figure 20-6.

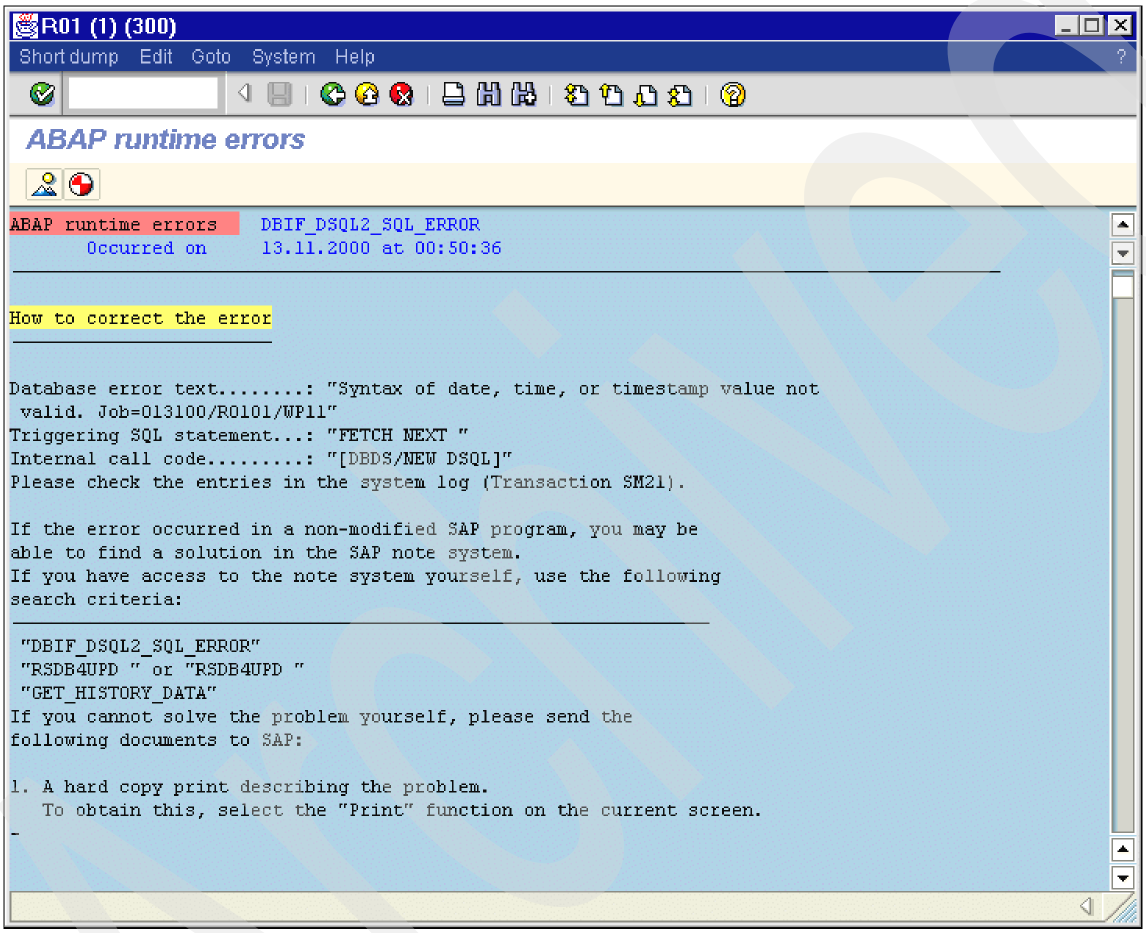

Figure 20-6 Short dump

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 date and 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 can identify the type of the problem by looking at the SAP trace and log files, the job logs, and the system configuration. The following sections describe the data that can be reviewed or collected when preparing to report the problem.

Further information about problem determination tools can be found in Chapter 19, “Problem avoidance” on page 331.

20.6.2 Database error

In many cases, the end user receives an SAP 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 occurred. 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 determine if messages in the job log are related to the problem seen by the 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...”, search for the term “SQL0913”. First, 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, look for escape-type messages around the time of the error.

> **Note:** Use the DLTR3PKG command to delete the SQL packages after the following events:
>
> - Kernel patch installation
> - PTF installation

Some message types can generally 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 the SAP system, 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, collect the following information:

- SAP system log
- SAP short system dump
- Developer trace of the SAP work process
- Job log of the corresponding job (Refer to 20.2, “Working with job logs” on page 342.)
- SAP SQL trace (ST05) to identify the failing SQL statement if the problem is reproducible

20.6.3 Program exceptions

Use transaction ST22 or SM21 to view and administer SAP termination messages. These transactions show you the Process-ID (PID) of the SAP work process.

In case of a program exception (for example, jobs terminating with a MCH3601 or MCH0601 error), perform the following steps:

- Start transaction ST22 or SM21 to view and administer SAP termination messages. These transactions also show the Process-ID (PID) of the SAP work process.
- Run the WRKPID command to find the associated i5/OS job
- DSPJOBLOG shows the i5/OS termination message
- Press F1 to display additional message information
- Press F9 to display and analyze message details. The 'From program' information gives a first indication of where the problem is primarily reported. Note the following:
  - Program names starting with the letter 'Q' should primarily be reported to IBM
  - Program names starting with a character different than 'Q' should primarily be reported to SAP.

This helps to localize the problem more precisely and to directly address the responsible support team either from SAP or IBM.

20.6.4 Performance problems

For all kinds of performance problems, you should ask yourself the question: "What has changed since the performance degraded?" Think about recent changes in the system configuration, network, an increased number of users, kernel patch, support packages, PTFs, modifications, SAP release, or i5/OS release.

If you cannot solve the performance problem with this analysis, contact SAP and describe the problem as precisely as possible.

For initial consideration when diagnosing performance problems, consider the following.

- System-wide performance

  If the general performance of your SAP system seems worse than you might expect, the following lists identifies some areas to look for bottlenecks in the system-wide performance.

  - For i5/OS, check:
    - Sizing: Make sure the sizing of the System i configuration is still suitable.
    - 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 disks are equably busy. If the usage on one disk is significantly different, use CL commands TRCASPBAL and STRASPBAL TYPE(*USAGE) to balance the usage and the capacity of the disks within an ASP.
    - Interactive load: The interactive load can be too high on a server model. Use the WRKSYSACT command from the Performance Tools (5722-PT1) licensed

program, if installed, to check if the CFINT<nn> tasks are consuming most of CPU time.

- JVM: Check that the garbage collector is properly configured, as described in SAP note 717376.

- For SAP, check:
  - Trace level: Set the trace level for the developer traces to 1.
  - DBMON: Use the ST04 transaction to check whether the database monitor is running. Note that even the new memory-based database monitor needs system resources.
  - Buffering: Check the buffering quality with transaction ST02.
  - **Workload**: Transaction ST03 shows some information about the current workload and response times.

For more information about how to analyze performance problems, refer to Chapter 19, “Problem avoidance” on page 331. Refer to 26.1, “Tuning System i models” on page 532, for information about i5/OS and SAP logs and tools available for diagnosing problems relative to performance.

- Transaction-based

  The performance of a particular transaction can be poor. Use the following steps to collect all data that is necessary for a detailed analysis if the problem is reproducible:

  a. Complete these steps from the SAP GUI session where you want to run the poor performing transaction:
     i. Go to transaction SE38, enter `RSTRC000` for the report, and click the Execute button.
     ii. 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.
  b. Using transaction SM50 from another SAP GUI session, you should see the work process with a status of “Stopped” and a reason of “Lock” for the first session. The PID is needed for the session for the next step.
  c. From i5/OS, complete these steps:
     i. Enter the command, WRKPID PID(<nnn>)

        Here *<nnn>* is the PID for the stopped work process. This displays the associated job for the work process that needs to be debugged.
     ii. 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.
     iii. Start the remote service operation with the command:

          STRSRVJOB JOB(<jobname>)
     iv. Put the job into debug mode using the command:

         STRDBG UPDPROD(*YES)
  d. From the second SAP GUI session, start the SQL trace for the user using transaction ST05.
  e. From the first SAP GUI session, run the poor performing transaction. It can run slower than before because there are two levels of tracing running.
  f. Once the transaction completes, end the SQL trace with transaction ST05 from the second SAP GUI session.

g. From the first SAP GUI session, use SE38 to unlock the work process by to run report RSTRC000 again. This time, remove the “X” in the Keep work process box and click the Change button. You see a message stating that the work process is now unlocked.

h. From i5/OS,

   i. End the debug mode with the ENDDBG command

   ii. End the remote service operation with the ENDSRVJOB command.

   iii. 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.

i. Provide the job log to the service organization that is looking into the problem.

20.6.5 Integrated File System problems

An SAP system uses the Integrated File System and especially the /QFileSvr.400 file system. If problems occur, check if a directory for each system in the SAP system landscape exists under /QFileSvr.400. Also verify that the following command has been run:

```
STRHOSTSVR SERVER(*FILE)
```

Make sure port 8473 is in state “Listen”. Also verify that the instance user profile <SID><nn> is enabled, and uses the same password with identical authorities and identical User-ID (UID) on all systems.

20.6.6 Printing problems

Refer to 15.9, “Resolving printing problems” on page 248 for more information.

20.6.7 Damaged objects on the System i server

Damaged database objects have different effects on System i configurations. Usually, the damage is observed when the data is backed up or when an SAP transaction terminates with one of the following error messages:

- `MCH0601 - Space offset...outside current limit for object <name>`
- `MCH1604 - System object <name> damaged ...`
- `MCH1668 - System object <name> partially damaged ...`
- `CPF3285 - Damage found on file <name> in library`
- `CPF3819 - Damaged object <type> <name> not restored`
- `CPF8111 - Damage on member <name> file <name>`

Objects can become damaged on the System i server after hardware errors occur, the system terminates, there are errors in the database source code, and other factors. The log file of the licensed internal code (LIC log) contains information that can help find the cause of the error.

Recovering from damaged objects may be simple or difficult, depending on the type and number of objects that are damaged and the operations active at the time of the damage.

Contact IBM Support for help in analyzing the log to determine the cause of the damage. You can also need SAP Service to help determine if it is safe to recover an individual application table or if it must be recovered with other data to ensure consistency across the database.

If a large number of objects are damaged (due to the failure of multiple hard disk drives within the same Redundant Array of Independent Disks (RAID) set, for example), you may need to restore the entire system from the latest system backup. Refer to SAP note 825473 *“Recovery after loss of data”* for instructions if you are advised to recover in this manner by the IBM or SAP Service organizations.

20.6.8 Unable to start the SAP application

When the SAP application is unable to start, the problem is typically due to errors made in the setup and startup of the SAP application, not errors in the SAP application.

Consider the following.

- TCP/IP and server support

  Ensure that TCP/IP support is started. 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. Ping the host using the IP address. If this job is not running, start TCP/IP using the STRTCP command.

  Use the i5/OS 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 QPWFSERVSO should be running in the QSERVER subsystem.

- 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 the SAP application 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 the host names specified in the SAP 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. The first host name for an address should be the host name itself. The second entry for an address should be the host name with the domain suffix (<hostname>.<domainname>). 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.

- Remote file system authority

  In a 3-tier environment, the application server uses the remote file system to read or write Integrated File System stream files to the database server. It is required that the i5/OS user profile are 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 i5/OS user profile for an application instance must exist on the database server. The user profile for an instance has the name <sid>*<nn>*. If, for example, system PRD instance 00 exists on the database server and instance 01 exists on an application server, the application server must have a user profile named PRD01. The database server must have user profiles named PRD00 and PRD01. If the standard installation procedure is used, this is done automatically. If the standard installation procedure is not used, use the CRTSAPUSR command to create the i5/OS user profiles on the database server that are found on the application server.

Refer to "Remote file systems" on page 65 for a discussion of remote file systems.

- 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 determine what mutex is causing the wait situation. This is useful information for SAP to identify the root cause of the problem. Try terminating the subsystem and restarting the SAP system once to see if this situation remains the same.

- Review the job logs of user *SIDnn* and the user starting the SAP system (for example, *SID*OFR).

  Use the following commands to review the QPJOBLOG file of user *SIDnn* and *SID*OFR:

  ```
  WRKUSRJOB USER(SIDnn)
  WRKUSRJOB USER(SIDOFR)
  ```

  Select option **8** to see the spool files, or option **5** and then option **4** to review the QPJOBLOG file.

  Another way to see the spool files is with the following commands:

  ```
  WRKSPLF SELECT(SIDnn)
  WRKSPLF SELECT(SIDOFR)
  ```

- Review the output of standard i5/OS output files

  Review the QPRINT spool files of user *SIDnn*.

  Note: By default the logging level in the job description R3*SID*400/R3_*INST* is LOG(4 00 *NOLIST), so that you usually do not get a job log spooled.

  In order to get a job log, issue the following before the STARTSAP command:

  ```
  CHGJOBD JOBD(R3<SID>400/R3_<INST>) LOG(4 00 *SECLVL)
  ```

- Developer trace output

  By default SAP applications write what is referred to as *developer trace files* for the dispatcher, the message server, the gateway server, each work process, and for more processes.

  Change to the developer trace directory by entering the following command:

  ```
  cd '/usr/sap/<SID>/DVEBMGS<INST>/work'
  ```

  Use the Display File (DSPF) command to display the files with the following parameters:

  – dev_disp(Dispatcher) for the dispatcher
  – dev_w0(1st work process) for the first work process
  – sapstart.log (starter) for the starter job log
  – dev_rd(Gateway)for the gateway log

If you cannot detect or solve the problem with the job log, spool file, or developer trace information, create a customer message and report the problem to SAP.

20.6.9 Inability to stop a job or work process

Ending an SAP work process is a step-by-step approach. If one step does not end the job, try the next step after about two minutes, as follows:

1. Look for jobs in a *COMMIT or *ROLLBACK status when viewed from WRKACTJOB. These jobs are trying to recover to a point where they can end cleanly. If you terminate them abnormally, they start the recovery process again at IPL.
2. Use the SM37 SAP transaction to end the batch job.
3. Process Cancel with core from SAP Transaction SM50.

   To generate a dump by cancelling a process, use the **cancel with core** SAP function from the menu of transaction SM50. You will see the dump in transaction ST22 with the runtime error.
4. Process Cancel without core from SAP transaction SM50.

   With this SAP function, you can cancel the job without generating a dump.
5. If the SAP work process still does not stop, find the process ID from SAP transaction SM50 and retrieve the System i process (*user/jobname/number*) from it with the System i SAP kernel command. Use the following commands:

   a. `WRKPID PID(process-ID from SM50).` The user is typically *SID*OFR.

   b. `ENDJOB JOB(number/user/jobname) OPTION(*CNTRLD)`

   c. `ENDJOB JOB(number/user/jobname) OPTION(*IMMED)`

   d. For SAP kernel 640 and greater, to send an i5/OS PASE signal, enter the following commands:

   ```
   CALL QP2SHELL PARM('/QOpenSys/usr/bin/kill' '-TERM' '<process-ID from SM50>')
   CALL QP2SHELL PARM('/QOpenSys/usr/bin/kill' '-KILL' '<process-ID from SM50>')
   ```

   e. For all releases, to send an ILE signal enter the following command:

   ```
   KILL PID(<process-ID from SM50>) SIGNAL(*SIGTERM)
   ```

   KILL is an SAP-provided *CMD.

   Alternatively, you can combine step d and e with the following steps:

   i. Identify the process ID (PID) of the work process, for example, by looking at the developer trace file.

   ii. Enter the QP2TERM command.

   iii. Enter the following command:

   ```
   kill -9 pid
   ```

   f. Enter the following command:

   ```
   ENDJOBABN JOB(number/user/jobname)
   ```

   > **Tip:** Only run the End Job Abnormal (ENDJOBABN) command about ten minutes after issuing the ENDJOB *IMMED command.

6. Before IPLing the machine, Open a Problem Management Record (PMR) with IBM.

Refer to 20.7, “Reporting the problem” on page 360 to understand how to report a problem. IBM Service may be able to assist you to recover without an IPL or whether an IPL is the correct action to take and how to perform this operation safely.

Never shut down the partition by unplugging the machine.

20.6.10 SAPOSCOL

Normally only one SAPOSCOL job should run in a partition, servicing all systems. However, there are instances where multiple SAPOSCOL versions may exist within a partition because of SAP release or codepage differences.

SAP has changed the performance data collection for all SAP application releases, so that all kernels from before Version 630 (Releases 61I, 40B, 45B, 46D, and 620) collect their performance data in one format. Kernels that are based on Version 640 (Releases 640 and 620 with the 640 downward-compatible kernel) collect their performance data in a new format and data structure. Therefore, if you have a combination of 6.40 and older kernels, two SAPOSCOL jobs can be active on a System i model, collecting performance data independently of each other.

If you have difficulty with transaction ST0,6 we recommend that you import the latest patches for the kernel, the SAPOSCOL and the RFCOSCOL for consistency. Refer to SAP note 371786 for an example.

Refer to the related SAP notes listed in Table 20-2.

Table 20-2 SAP notes for SAPOSCOL

SAP note	Summary
17852	iSeries: Authority Problems in SAPOSCOL
637174	SAPOSCOL cannot access Libraries of different SAP systems
850218	iSeries: SAPOSCOL can’t be stopped from the menu
708136	iSeries: No monitor data in ST06
889292	iSeries: SRVPGMe LIBICUDATA, LIBICUI18N, LIBICUUC are missing
12103	Contents of the TCOLL table
874480	Entries in table TCOLL and domain COLL_RNAME
787127	iSeries: ST06 CPU statistics (additional information)

20.7 Reporting the problem

The following information is useful to provide in any problem description regardless of whether you are initially reporting the problem to SAP or to IBM. Getting a good understanding of the current software levels can help SAP and IBM solve the problem faster. Table 20-3 on page 361 shows the information that needs to be collected in any case, regardless of the type of problem that is encountered.

Table 20-3 General information to be collected

Type of Information		How to find it	Example
SAP	SAP application level	SPAM → Display all support packages	SAP_BASIS SAPKB64009 SAP_ABA SAPKA64009
	SAP kernel release	SM51 → Release notes	640
	SAP kernel patch level	SM51 → Release notes	104
	SID, instance and client	-	P01, 00, 300
	SAP system landscape	DSPR3SYS <SID>	Instance numbers, hostnames, roles
i5/OS	System name	CFGTCP opt 12	AS23
	System model number	DSPSYSVAL QMODEL	550
	System serial number	DSPSYSVAL QSRLNBR	100ABCD
	Cumulative PTF package level	DSPPTF	C4209530
	Group PTF levels for database and Java	WRKPTFGRP	SF99530 - 9

20.7.1 SAP system environment

Always provide the following information about the SAP system where the error occurred:

- SAP system landscape
- SAP application: ERP, BW, CRM ...
- Test/QA or productive system
- SAP database release
- SAP database codepage (EBCDIC, GLS, Unicode)
- SAP database in system ASP or IASP
- SAP kernel release
- SAP kernel patch level
- SID, instance, and client
- 2-tier or 3-tier installation

  For 3-tier, provide the operating system of the application server (i5/OS, Windows 32/64 bit, Linux)

To determine the kernel patch level, follow these steps:

1. Call transaction SM51, and select the server name.
2. Click the Release info button. Make note of the kernel patch number and DBSL patch number under the SAP Kernel information section.
3. If SAP GUI is not available, sign on to i5/OS with <SID>OFR and use the i5/OS command:

```
DSPDTAARA DTAARA(KERNEL)
```

20.7.2 i5/OS environment

Provide the following information about the i5/OS environment:

- System name: The system name of the System i model that identifies the system.

- System model number: Every system has a system model number.
- System serial number: IBM uses the system serial number to open a Problem Management Record (PMR) on your behalf.
- Cumulative PTF package and database fix package level: This PTF package contains a collection of PTFs for specific licensed program products and DB2 UDB for iSeries.

  Enter the command:

  `WRKPTFGRP`

  Press F11to see all installed PTF groups, as shown in Figure 20-7:

```
                              Work with PTF Groups
                                                              System:   ITSO-SYS1
 Type options, press Enter.
    4=Delete   5=Display   6=Print   8=Display special handling PTFs
    9=Display related PTF groups

  Opt  PTF Group             Text
       SF99530               CUMULATIVE PTF PACKAGE C5207530
       SF99529               GROUP HIPER
       SF99503               DB2 UDB FOR ISERIES
       SF99295               WEBSPHERE MQ FOR ISERIES - VERSION 5, RELEASE 3.
       SF99285               WEBSPHERE APP SERVER V5.1 (BASE/DEV. EDITION)
       SF99269               JAVA
       SF99173               IBM BUSINESS SOLUTIONS
       SF99148               WAS - ADVANCED 4.0.7 5733WA4
       SF99099               IBM HTTP SERVER FOR ISERIES




                                                                          Bottom
  F3=Exit   F6=Print   F11=Display status   F12=Cancel
  F22=Display entire field
```

Figure 20-7 WRKPTFGRP, F11 - Display descriptions

20.7.3 Saving spooled files

There are two ways to save all problem-related spooled files (such as job logs, etc.) so that they can be delivered in a manageable way to IBM or SAP:

- Use iSeries Navigator within iSeries Access. Open the Work Management functions, choose Output Queues, choose the spool files of interest, then right-click and choose the Export menu tab.
- CL commands: If you prefer to use CL commands on the 5250 interface to save spool files to a physical file, and then save the physical file to a save file, enter the following commands in the order presented:
  a. CRTLIB LIB(<library>) TYPE(*TEST)
  b. CRTPF FILE(<library>/<file>) RCDLEN(132) MBR(*NONE) MAXMBRS(*NOMAX) SIZE(*NOMAX)
  c. CPYSPLF FILE(<spooled file name>) TOFILE(<library>/<file>) JOB(<qualified job name>) SPLNBR(<spooled file number>) TOMBR(<spooled file name>)
  d. CRTSAVF FILE(<library>/<save file>)

e. SAVOBJ OBJ(<file>) LIB(<library>) DEV(*SAVF) OBJTYPE(*FILE) SAVF(<library>/<save file>) DTACPR(*HIGH)

20.7.4 Reporting the problem to SAP

You may report issues to SAP by using the SAP Service Marketplace at:

http://www.service.sap.com/message

Information that you have collected (joblogs, developer traces, spool files) can be attached to the message. It is recommended that you zip the files into a single archive file before attaching it. If larger amounts of data are requested by SAP, they can be submitted through a separate interface at:

https://sapmats.wdf.sap.corp/

Keep all relevant information directly attached to the message in order to keep an immediate and tangible link between the data and the message to which it relates.

20.8 Additional information

Additional information is described in Information APARs, IBM Web sites, and SAP notes, as follows:

- IBM Information APARs
  - II12632: Terminology: An overview of terms used with SAP implementation, for example the job and object names used in an SAP system
  - II12633: General debug information and how to find job logs
  - II12634: Typical errors and tips on how to handle typical errors
  - II12635: Collecting documentation
- IBM Web sites
  - Support Line Knowledge Base: http://www.ibm.com/servers/eserver/iseries/service/erp/support.html
- 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	Tips about applying SAP fixes (APYR3FIX)
107116	Information about the termination of a work process
121625	Buffer sizes
717376	Recommended settings for Web AS Java Engine

See Appendix A, “Support for SAP applications” on page 711 for further information about obtaining support for IBM and SAP implementations.

21

Correctable scenarios

This chapter describes some scenarios that have been found and corrected at client sites. If you notice that a particular scenario describes your current environment, follow the links within each scenario to the appropriate section in this book to understand how to avoid or resolve the situation. Seventeen scenarios are provided with these topics:

- "Scenario 1: Scale up within a partition" on page 366
- "Scenario 2: Give journal receiver ASPs enough resources" on page 366
- "Scenario 3: Place database intensive batch work in a 2-tier instance" on page 366
- "Scenario 4: Revalidate hardware sizings when the environment changes" on page 366
- "Scenario 5: Test your recovery scenario" on page 366
- "Scenario 6: Optimize custom ABAP coding" on page 367
- "Scenario 7: Minimize the use of "FOR ALL ENTRIES" statements in ABAP" on page 367
- "Scenario 8: Perform a shutdown and cleanup only when necessary" on page 367
- "Scenario 9: Let i5/OS jobs complete ROLLBACK" on page 367
- "Scenario 10: Understand your system when things are going right" on page 368
- "Scenario 11: Avoid problems in the service process" on page 368
- "Scenario 12: Understand that your situation may not match others" on page 368
- "Scenario 13: Be cautious with Internet advice" on page 368
- "Scenario 14: EBCDIC transitions to GLS or Unicode have a tool" on page 369
- "Scenario 15: Be cautious with platform comparisons and evaluations" on page 369
- "Scenario 16: Authorize R3OWNER to all SAP libraries" on page 369
- "Scenario 17: Install latest patches for SAPOSCOL and all SAP releases" on page 369

Scenario 1: Scale up within a partition

- **Situation:** A large System i model is configured with multiple logical partitions (LPARs) so that i5/OS application servers can be created in a 3-tier configuration.
- **Problem**: Logical partitions should never be introduced for scaling purposes. A 3-tier configuration introduces server jobs and network overhead that degrade response time.
- **Solution**: A partition can be scaled up to 32 processors. Use 2-tier i5/OS application servers whenever possible to achieve the best response times and to take advantage of System i simplicity. Create another 2-tier application instance within the partition to separate work or users. Define a separate logical partition if you use a Linux on POWER application server. See Chapter 3, “SAP system landscapes” on page 17.

Scenario 2: Give journal receiver ASPs enough resources

- **Situation:** Multiple SAP systems are on a partition and all journal receivers are on a second ASP. The journal receivers are saved to save files in that ASP as they are detached. The ASP has enough capacity because there are a few large disk units.
- **Problem**: Journal performance and system throughput can be degraded if too few disk units are allocated to ASP2.
- **Solution**: Provide enough disk units to a second ASP to avoid excessive disk unit utilization, or move the journal receivers to the system ASP. See 9.7.1, “Using a separate ASP for the journal receiver” on page 110.

Scenario 3: Place database intensive batch work in a 2-tier instance

- **Situation:** You have an under utilized 3-tier application server and have high batch requirements.
- **Problem**: Batch jobs usually involve a significant amount of database access.
- **Solution**: Run batch jobs in a 2-tier instance on the database server whenever possible to reduce the amount of data transfer over the network through server jobs. See 11.4, “Differences between 2-tier and 3-tier landscapes” on page 135.

Scenario 4: Revalidate hardware sizings when the environment changes

- **Situation:** You are upgrading an SAP installation to a later release level, or are adding a significant amount of work to an existing system.
- **Problem**: The original system sizing is based on common assumptions that may or may not still apply for the present-day or future environment.
- **Solution**: Use the IBM Insight tool to determine your hardware utilization. Use this data to revalidate your sizing and to project future requirements. Although your system may be performing satisfactorily at the moment, a significant change can extend the required resources beyond what you currently have. Revalidate the hardware sizings when entering production, upgrading to new releases, or when significant workload is added. See 18.2, “Tools to track resource utilization” on page 320.

Scenario 5: Test your recovery scenario

- **Situation:** You have just implemented a high availability or disaster recovery solution, or both.
- **Problem**: You have not tested recovery. You will discover if your recovery solution works the first time you use it.
- **Solution**: At a minimum, test your procedure immediately after implementation. It is also advisable to retest when you change high availability solutions, the operating system, or

SAP release levels. See Chapter 22, “Availability, backup and recovery concepts” on page 373.

Scenario 6: Optimize custom ABAP coding

- **Situation:** You have created a large amount of customized ABAP code. You use a lot of internal tables that select large amounts of data from the database, then process these internal tables to find the few rows that are to be used.
- **Problem**: ABAP code can be poorly written. Internal table processing is expensive. Large internal tables use a lot of memory. Database queries that return large numbers of rows are often implemented by the optimizer as table scans.
- **Solution**: Use the ABAP workbench tools to scan and optimize ABAP coding. Use SQL selection to minimize the number of rows and columns returned to internal tables. Use the SQL trace utility to optimize database queries, and SQL performance tools to identify which tables should have additional SQL indexes. See 26.2, “Analyzing and optimizing single transactions” on page 554.

Scenario 7: Minimize the use of “FOR ALL ENTRIES” statements in ABAP

- **Situation:** Your ABAP coding uses “FOR ALL ENTRIES” statements with large internal tables.
- **Problem**: “FOR ALL ENTRIES” is just one ABAP statement. Each entry in the internal table it applies to creates several WHERE clauses in the SQL statement that is passed to the database optimizer.
- **Solution**: Minimize the internal table sizes whenever possible. See 25.1.3, “Impact of database growth on System i performance” on page 459 and 26.1.2, “Database tuning” on page 545.

Scenario 8: Perform a shutdown and cleanup only when necessary

- **Situation:** You end all SAP instances, delete all SQL packages, run Reclaim Storage (RCLSTG) and Initial Program Load (IPL) i5/OS on a regular basis.
- **Problem**: These operations are often unnecessary.
- **Solution**: Perform these operations only when necessary. Some customers have run for months without shutting SAP instances down. If anything, use the instance profile setting to restart work processes periodically in an automated way. Do not use the Reclaim Storage command (RCLSTG) unless you are sure your system cross reference files are corrupted, or if recommended to do so by an IBM service representative. Selectively use the Reorganize Physical file Member command (RGZPFM) or the Start Auxiliary Storage Pool Balance command (STRASPBAL) to optimize disk layout. An IPL should only be necessary if temporary storage is consuming an excessive portion of the system Auxiliary Storage Pool (ASP), or as required for some PTF applications.

Scenario 9: Let i5/OS jobs complete ROLLBACK

- **Situation:** i5/OS jobs are in ROLLBACK status for a long period of time. The job consumes very little CPU and you are considering IPLing the machine.
- **Problem**: An error has occurred and the rollback status for the job indicates that recovery is already in progress. It may be very input/output (IO) intensive, so CPU utilization may not be high. An IPL in this state forces the system to start the rollback processing over from the beginning during the IPL.

- **Solution**: Check for disk activity before concluding that the job is not responding. Wait for the job to complete the rollback recovery. Check with IBM service before IPLing in this state.

Scenario 10: Understand your system when things are going right

- **Situation:** You notice poor performance, but you cannot identify the source. You are certain that nothing has changed.
- **Problem**: Your system is always changing, whether you realize it or not. Users behave differently, file contents and sizes continually change.
- **Solution**: Check your performance metrics when users are satisfied to understand what “normal” performance is. This helps you identify specific areas to focus on if problems occur. Log all changes so that each of them can be evaluated as a possible cause. See 18.2, “Tools to track resource utilization” on page 320.

Scenario 11: Avoid problems in the service process

- **Situation:** You have an issue and you need to contact IBM or SAP service. You are busy so you do not provide much information other than to say “SAP isn’t working.”
- **Problem**: It will take at least one more iteration between you and service before your problem is understood.
- **Solution**: Provide basic data when submitting a problem to IBM or SAP, such as the current release of the system and SAP, patch levels and any symptoms or error codes that are available so that service can search the known issues. Many reported problems are duplicates of known issues that are already solved.

  If you have already opened a problem with IBM, let SAP know the Problem Management Record (PMR) number. If you have opened a problem with SAP, let IBM know the SAP message number. Attach (complete) SAP developer traces and IBM joblogs if it is possible to identify in which work-process the problem occurred. If opening an SAP message, open SAP GUI and telnet service connections before reporting the problem, and provide logon information in the secure area of the message.

  If you have Windows application servers or are having trouble with SAPInst, Windows Terminal Server service connections are a great advantage. See 19.7, “Avoid more problems once one is identified” on page 335.

Scenario 12: Understand that your situation may not match others

- **Situation:** You know another customer running SAP applications on the System i platform. You share information about performance data or configuration options.
- **Problem**: Very few SAP installations have the same environment or requirements. What is right for one customer, may not be effective for another.
- **Solution**: If differences are noted, it does not necessarily mean that one or the other is incorrect. Take hardware, software, workload differences into account if comparisons are made. With so many variables, you may find that it is not useful to compare.

Scenario 13: Be cautious with Internet advice

- **Situation:** You read it on an user group forum related to SAP applications on System i models. It must be true!
- **Problem**: Opinions, correct or otherwise, are freely available. Contexts are often not complete and emotions may be involved in any particular posting.
- **Solution**: Be careful to verify facts and relevance before taking action on your own system.

Scenario 14: EBCDIC transitions to GLS or Unicode have a tool

- **Situation:** You are on an older release level using the EBCDIC code page. You want to go to a new release of SAP and believe that you must do a full migration process to get there.
- **Problem**: Migration tools are only necessary when coming to the System i solution from another platform.
- **Solution**: Use the Code Page Conversion Tool to do an in-place conversion at the 4.6 release level before upgrading beyond 4.6. This tool changes the data from EBCDIC to UNICODE without having do perform a complex unload/reload migration procedure. A second, but even easier data conversion step is needed to move from a GLS version to the full Unicode version. See 18.7.1, “Migration” on page 325.

Scenario 15: Be cautious with platform comparisons and evaluations

- **Situation:** You are facing a major upgrade and must evaluate other platforms. A consultant that also offers services for another platform says that another platform is better and less expensive.
- **Problem**: Initial offers often only include hardware acquisition costs.
- **Solution**: Evaluate based on Total Cost of Ownership (TCO) concerns, using verified TCO studies. Add in the database license cost, other 3rd party software, service and maintenance fees for each, consulting to get it to work together, consulting to get it to perform, additional networking infrastructure for the server farm and every other aspect of a complete implementation that you can think of. At the end of the day, the hardware acquisition costs are a small piece of the TCO story. Be aware that most other platforms require more consulting services and that they may be offered by the person recommending a change to one of those other platforms. See page *7* as a reference to TCO and industry reports.

Scenario 16: Authorize R3OWNER to all SAP libraries

- **Situation:** You want to run SAPOSCOL on a System i configuration. There is more than one SAP system installed on that machine.
- **Problem**: Starting SAPOSCOL fails. In the SAPOSCOL log file (/usr/sap/tmp/dev_coll) you see a CPF9820 error message indicating SAPOSCOL cannot access libraries of different SAP systems.

  SAPOSCOL is not able to access every library with the standard default security settings.
- **Solution:** For every SAP system installed there is a library with the name R3<SID>400. SAPOSCOL needs access to this library. SAPOSCOL runs under the user R3OWNER, but the R3<SID>400 libraries belong to user <SID>OWNER).

  Log on as QSECOFR or similar. Enter the following command:

  ```
  ADDAUTLE AUTL(R3ADMAUTL) USER(R3OWNER) AUT(*ALL)
  ```

  This command adds the user R3OWNER to the R3-specific authorization list R3ADMAUTL, thus granting it all rights to all SAP libraries.

Scenario 17: Install latest patches for SAPOSCOL and all SAP releases

- **Situation:** You have multiple SAP application servers and use SAPOSCOL to collect SAP trace data.
- **Problem**: One or more of the following fault situations:
  - When you select Transaction ST06, the system displays a message, “SAPOSCOL not active” or “SAPOSCOL not running”, even though WRKACTJOB tells you that a SAPOSCOL job is active; or

– Performance data is not displayed in Transaction ST06.

– Work processes terminate because they cannot connect to a shared memory segment.

- **Solution:** Special considerations apply if more systems run in one partition mixed with SAP kernels prior to kernel version 6.40 and version 6.40 or after. You must use a SAPOSCOL program level consistent with each kernel version. Install the latest patches for all SAP releases. See 20.6.10, “SAPOSCOL” on page 360 and 27.3.1, “SAP OS Collector” on page 577.

Part 3

Availability, backup and recovery

The integrity and consistency of an SAP system database is protected by the SAP 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. A high availability or disaster recovery concept optimizes the capability to withstand planned or unplanned outages and to reach predefined service level agreements. The impact of the failure depends on the duration of interruption as well as the importance of the applications impacted by the failure.

Chapter 22, “Availability, backup and recovery concepts” on page 373 discusses the terminology, availability options, and the features available for backing up your SAP system. Chapter 23, “Backup and recovery” on page 391 and Chapter 24, “Disaster recovery and high availability” on page 433 outline the standard System i backup, recovery, and availability facilities available. These chapters also provide examples about backup and recovery facilities and procedures in the SAP system environment.

You can find full, detailed coverage of the subjects in:

- *Backup and Recovery,* SC41-5304

- SAP online documentation on the SAP Help Portal at:

  `http://help.sap.com/`

- iSeries Information Center at:

  `http://publib.boulder.ibm.com/infocenter/iseries/v5r3/ic2924/index.htm`

Refer to these documents to plan backup, recovery, and system availability to protect and maintain crucial business functions.

22

Availability, backup and recovery concepts

This chapter describes the techniques you can use to provide the appropriate availability for your SAP application running on a System i server.

Before you determine the best availability solution for your business, it is important to understand some basic availability terminology:

- **Outage**: This is 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 program temporary fixes (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.

> **Note:** SAP database upgrades always require downtime. You can minimize scheduled downtimes for maintenance purposes, but you cannot completely avoid it.

It is important to understand which activities or events, scheduled or unscheduled, can influence the availability of your system, as the following sections discuss.

22.1 Scheduled outages

Scheduled outages are required to maintain your system. During this time, your SAP system is not available on this system. To minimize scheduled outages and optimize continuous operations, you must implement an effective backup system.

Activities that cause scheduled outages include the following:

- Hardware maintenance
  - Processor
  - DASD (Might not require you to shut down the SAP system)
  - Network (Might not require you to shut down the SAP system)
- Operating system maintenance
  - i5/OS release upgrade
  - Applying PTFs [*]
  - Save system
- SAP system maintenance
  - Upgrade to a new database release.
  - Install a new SAP kernel release.
  - Install support packages or patches for the SAP system.
- Offline backup, as necessary to perform the following:
  - Save an entire system.
  - Save all SAP stream files located in the Integrated File System.

    A shutdown of the SAP system is only needed in special cases, such as when tables contain a lot of deleted rows that cannot be reused and cause performance or space problems. For more information about how to reorganize SAP tables, refer to SAP note 84081.

22.2 Unscheduled outages

The following unscheduled outages influence system availability, each of them requiring a specific solution that covers your demands:

- Complete system loss

  A fire, flood, or other natural disaster can destroy your entire system. You should have a complete set of save tapes and documentation stored off site at a secure, accessible location to rebuild your entire system.
- System failure

  A system failure means that some part of your system hardware fails, other than the disk unit subsystems. 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 can be partially updated.
  - Access paths for files can be incomplete.
  - Objects that are in use can become damaged.
  - Relationships between files can be partially validated.

  When you restart initial program load (IPL), your system after the failed component is repaired, the system analyzes the possible damage, rebuilds or recovers access paths, tries to verify the file relationships, and attempts to synchronize files to transaction boundaries. The first IPL after the system ends abnormally can take a longer time than a normal IPL.

- Power failure

  A loss of power causes your system to end abnormally. You can experience the same type of problems that occur with a system failure. Either implement the feature called system power control network (SPCN) or use an uninterruptable 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 System i server a very productive system to program and to manage. However, it 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 and, in most 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 is not anticipated by the software developers. A program error can cause incorrect information in some of your data files and database tables. 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. Correct or restore the data that has been damaged when this type of error occurs, or rollback the journal entries of your database.

22.3 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 on the business must be minimized. This is why the System i architects continuously emphasize on the improvement of recovery times.

22.3.1 Hardware solutions

This section describes the available hardware solutions for managing availability, including system power control network, uninterruptable power supply, device parity protection, mirrored protection, and dual systems.

Table 22-1 on page 376 shows a comparison of the primary hardware configuration availability solutions.

Table 22-1 Hardware availability solutions

Impact	Availability	Performance	Cost efficiency
High	Dual systems	Mirrored protection	Dual systems
	Mirrored protection	Dual systems	Mirrored protection
Low	Device parity protection	Device parity protection	Device parity protection

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 after a power loss.

Uninterruptable power supply

This 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 prevents 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 the following:

- Technology similar to the Redundant Array of Independent Disks (RAID-5) technique
- Redundant power
- Concurrent maintenance for single disk failures
- Concurrent maintenance for power supply failures for the disk array subsystem

Note: Device parity protection should be the minimum disk protection for the ASP on which the SAP database library is located. This is usually the system ASP.

In case journal receivers for the SAP database are in a separate user ASP, do not include the disks for the journal receiver user ASP in the same parity set used for the database library.

Mirrored protection

If a disk failure occurs, mirrored protection is prevents 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 System i product line 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 the following:

- 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

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 methods for maintaining data on the secondary system are either through the use of data replication software which transmits the journal entries from the primary system to the backup system, or by installing the SAP database into an independent auxiliary storage pool (IASP) cluster, which allows you to switch the SAP database from one system to another.

Table 22-1 on page 376 positions device parity protection, mirrored protection and dual system concepts from a performance and cost point of view.

22.3.2 Software solutions

This section describes the available software solutions for managing availability, including access path protection, journals, and ASPs. Refer to the *Backup and Recovery Guide*, SC41-5304 for more information about each topic.

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 might have to rebuild the access paths before you can use the files again. This is a time-consuming process. It can take many hours to perform an IPL on large busy System i servers that end abnormally.

System-managed access-path protection (SMAPP) provides a simple method to reduce the recovery time after an abnormal system end. SMAPP manages the required environment for you. SMAPP is turned on by default.

Control the settings with the Edit Recovery Access Path command (EDTRCYAP). Access path changes are journaled to the same journal as the table changes if a table is journaled, and SMAPP chooses its related access paths for implicit journaling. The journal entries for access path changes are needed only for a faster recovery after an abnormal IPL. They are not needed for a recovery after a loss of data. Select the **RCVSIZOPT**(***RMVINTENT**) journal parameter to write access path-related journal entries into separate internal spaces so that the journal receivers do not grow too fast. The access path changes are journaled to an internal journal if SMAPP chooses access paths for journaling and the underlying table is not journaled.

Journal management

Use journal management to recover the changes to database files or other objects since the last complete backup. Use a journal to define what files and access paths to protect with journal management. This is 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 journaled. This includes changes to database files, changes to other journaled 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 can update more than one database file.

Journaling uses two object types:

- journals (*JRN)
- journal receivers (*JRNRCV)

The journal acts like a funnel. All database table adds, changes, and deletes are received by the journal. It then writes them to the journal receivers.

SAP journals the database tables automatically to a Structured Query Language (SQL) collection, but it does not journal the access paths explicitly. The journal object itself is located in the SAP database, where the journal receiver objects are stored in the journal receiver library.

The library names for SAP Web Application Server Advanced Business Application Programming (ABAP) systems are:

- R3*SID*DATA for the database
- R3*SID*JRN for the journal receiver

The library names for SAP Web Application Server Java systems are:

- SAP*SID*DB for the database
- SAP*SID*JRN for the journal receiver

Note: Without a dual system approach for high availability and permanent logical replication of the database files, it is recommended to create a user ASP to provide dedicated resources for journal receiver objects in order to reduce the loss of data due to a disk media failure.

See 9.6, "SAP applications and journal management" on page 102 for further information about journals.

Auxiliary storage pools

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. 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 can 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.

Three types of ASPs are available on the system:

- The system ASP (ASP1)

  ASP1 is always configured and contains the LIC, licensed programs, system libraries, and temporary system workspace.

- A basic user ASP (ASP2-ASP32)
- Independent user ASPs (IASPs, ASP33-ASP64)

An IASP can be defined as a switchable one, which means it can be switched among System i servers. There are different options to run a switchable IASP implementation. These are:

– Switchable IASP in a high-speed link (HSL) loop

  This allows you to move an IASP from one System i server to another.

– Switchable IASP with cross-site mirroring (XSM)

  The mirroring is done by replicating data segments from the primary system to a secondary one containing the mirrored disks.

– Peer-to-Peer Remote Copy (PPRC) data replication service with the Rochester installation toolkit

  This option requires IBM Enterprise Storage Server (ESS) or data server systems with PPRC license program to replicate the switchable IASP from the primary site to a backup storage system.

> **Note:** You require 5722-SS1 Option 41 of i5/OS (HA switchable resources) to implement a switchable IASP.

ASPs are further described in 8.1.5, “Auxiliary Storage Pool” on page 53. IASPs are described in 8.1.6, “Independent ASP” on page 53.

22.3.3 Backup solutions

This section shows the available backup solutions for the System i server. The Backup Recovery and Media Services (BRMS) and Tivoli® Storage Manager solutions are discussed.

Backup Recovery and Media Services

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 5722-BR1 offers a set of functions for defining and performing the following tasks:

- Backup, recovery, archiving, retrieval, and media management
- Policy-driven full function tape media management software for a single or multiple networked System i servers
- Auto archive and recall hierarchical storage management (HSM) dynamic retrieval

You can use BRMS in conjunction with the job scheduler (5722-JS1) to provide a very robust and flexible unattended automated backup strategy. We recommend a save strategy using multiple tape drives operating concurrently when performing unattended backups on systems.

See 17.2, “Using the IBM Job Scheduler in an SAP system environment” on page 302.

Tivoli Storage Manager

You can use Tivoli Storage Manager for iSeries server (formerly known as ADSM/400) to protect data on your workstations and local area network (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 a System i server. It can back up data from a variety of workstation platforms. You can use the BRMS program to back up System i user data to any Tivoli Storage Manager when the System i server is in a client or 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.

22.3.4 Availability solutions

Applying PTFs and i5/OS release upgrades have an impact on system availability in the SAP system environment.

Waiting for an extended time until their systems are available after such maintenance is difficult for some customers. Implement logical partitions or an effective dual systems approach to minimize this situation.

> **Note:** The dual systems approach is the most secure way to meet 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 either the journaling capabilities of the System i server along with independent software options to reproduce the duplicate system over a communications link, or the switchable IASP technology of i5/OS. Switch users to the duplicate system in the event of a planned outage. Processing continues until the scheduled outage is completed on the master system and journaled transactions are applied.

Obtain the necessary license keys from SAP for both the production and backup machines to use SAP applications in a dual systems configuration, as described in SAP note 94998. SAP returns the license keys once this information is provided. Install the SAP system license key as described in SAP note 870871. The SAP system cannot be used on that machine until the license keys are obtained for a server.

Refer to Chapter 24, “Disaster recovery and high availability” on page 433 for more information.

22.4 Backup and recovery

Figure 22-1 on page 381 and Figure 22-2 on page 382 show the commands and menu options that you can use to save parts of the system or the entire system. Review these figures to see which parts of your system are saved when you use the options from the SAVE menu if you choose to use a simple save strategy.

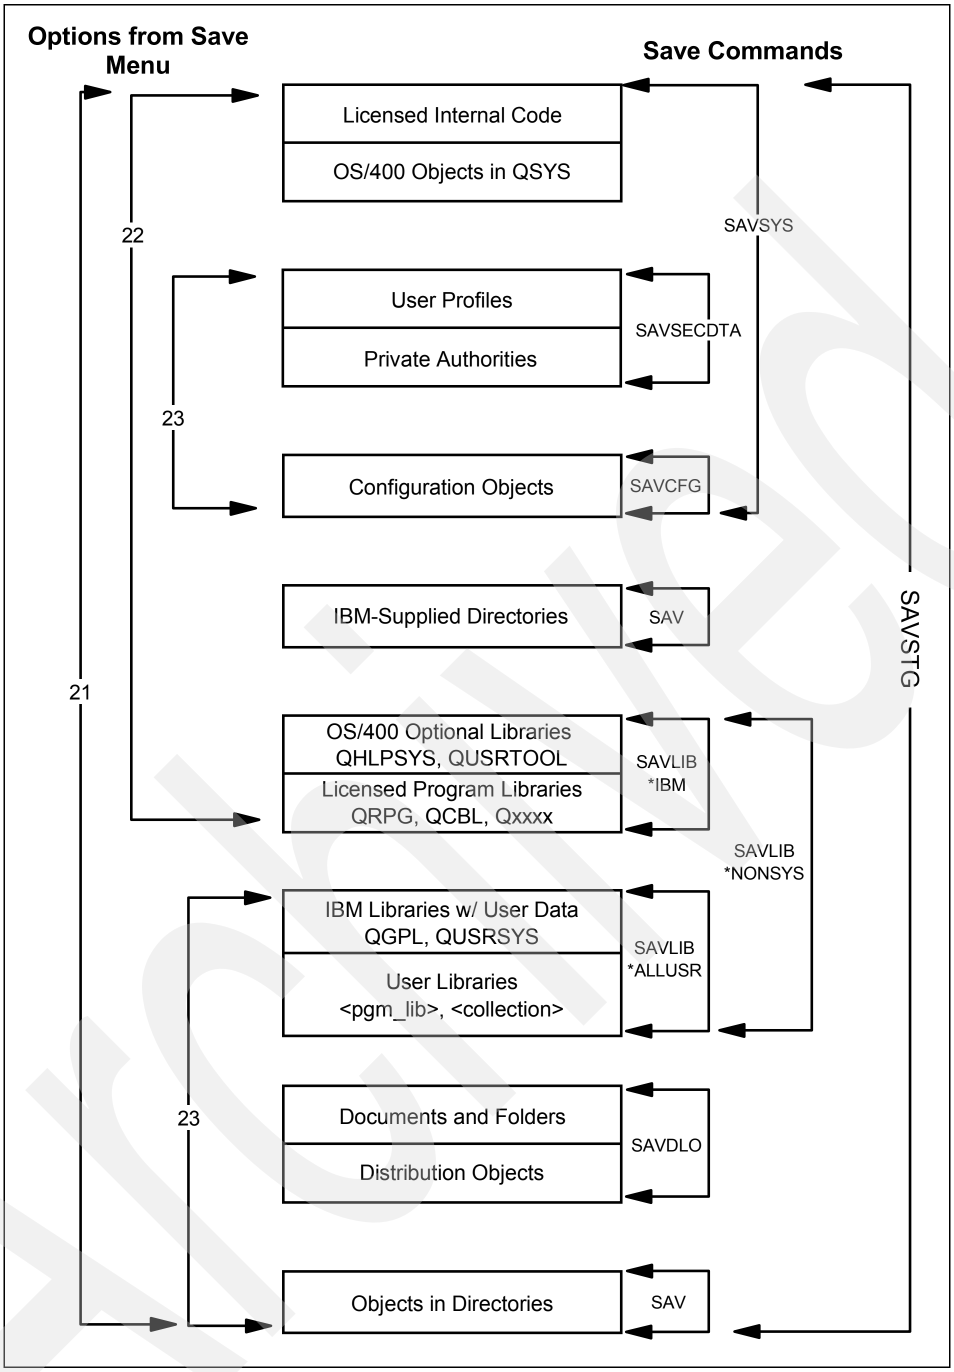

Figure 22-1 Save commands and Save menu options

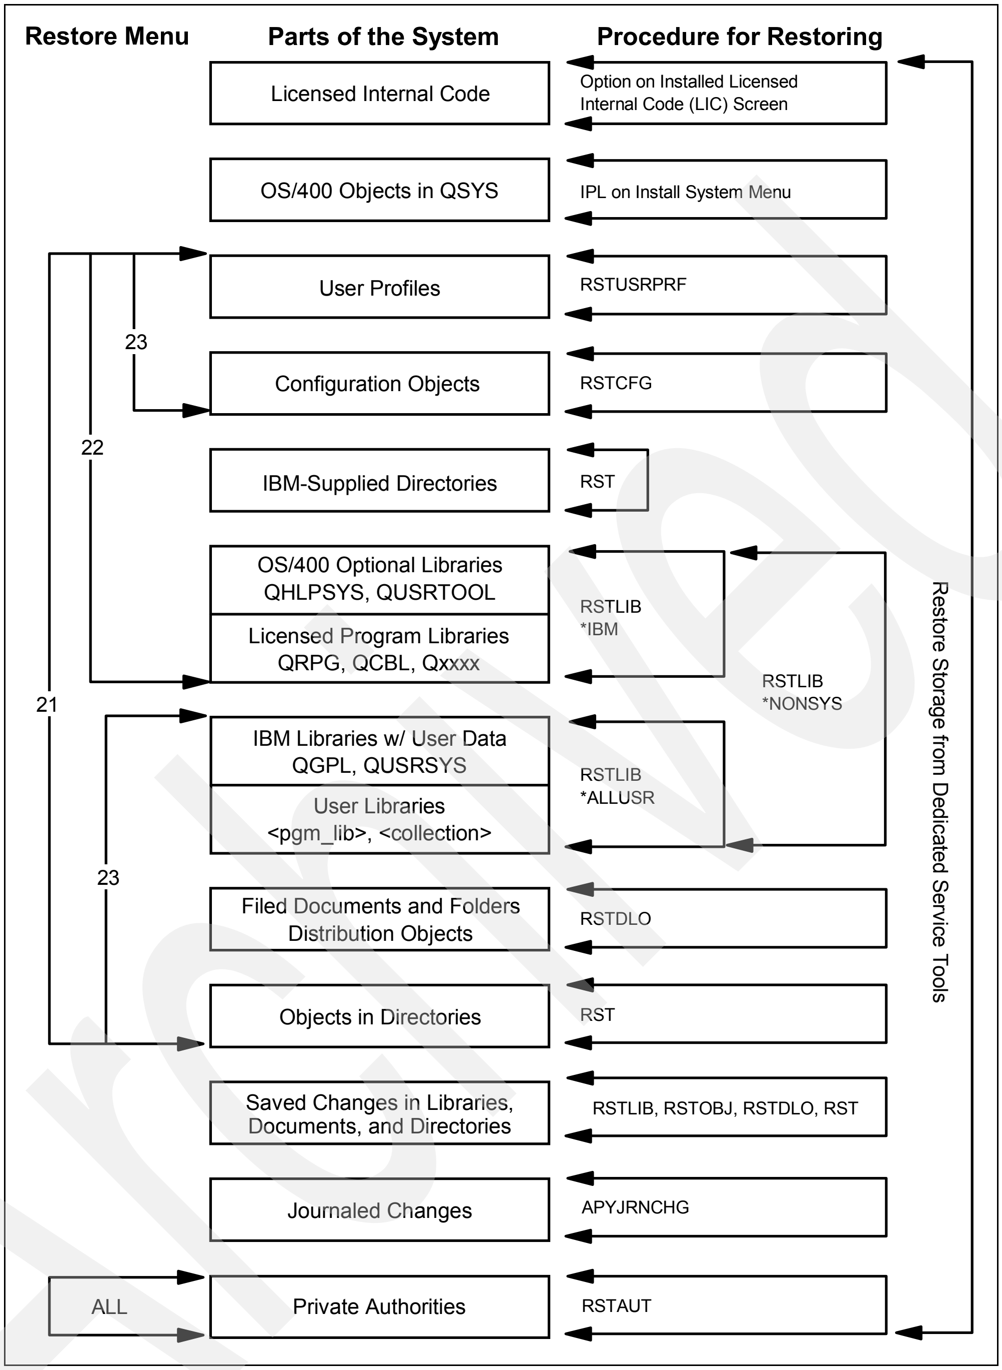

Figure 22-2 Restore commands and Restore menu options

22.4.1 i5/OS save and restore commands

All i5/OS save and restore commands are provided in the base i5/OS operating system. Use IBM BRMS if you need to backup and save media management.

Save and restore functions can be accessed in 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: Saves document library objects

- SAVLIB: Saves library
- SAVOBJ: Saves objects
- SAVSECDTA: Saves security data
- SAVSTG: Saves the whole content of auxiliary storage (except unused and temporary)
- SAVCFG: Saves configuration
- SAVSYS: Saves system (copy of the Licensed Internal Code)
- SAVCHGOBJ: Saves changed objects
- SAV: Saves 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: Restores document library objects
- RSTLIB: Restores library
- RSTOBJ: Restores objects
- RSTAUT: Restores user authorities to objects
- RSTCFG: Restores system configuration
- RSTUSRPRF: Restores user profiles and authority tables
- RST: Restores 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: There is no Restore System command or a Restore Changed Object command as this is done as a dedicated service tools function.

This chapter does not explain each command in detail. There are additional commands available that are not listed or described here. You can find detailed information about the save and restore commands in *Backup and Recovery*, SC41-5304, or *CL Reference*, SC41-5722.

Save commands requiring a dedicated system

A dedicated system (restricted state) on the System i server has only the console and system jobs running. No users or other applications can use the system.

The following IBM commands require a dedicated system:

- SAVSYS
- SAVSTG
- SAVLIB LIB(*NONSYS)

Save commands not requiring a dedicated system

The following IBM commands do not require a dedicated system (restricted state):

- SAV
- SAVCFG
- SAVCHGOBJ OBJ(*ALL) LIB(*library-name*)
- SAVCHGOBJ OBJ(*ALL) LIB(*ALLUSR)
- SAVDLO
- SAVLIB LIB(*IBM)
- SAVLIB LIB(*ALLUSR)
- SAVOBJ OBJ(*library-name/object-name*)
- SAVSECDTA

22.4.2 Save strategies

This section provides an overview of the objects that you need to save to ensure that you can restore your data in a consistent state.

SAP system libraries

Table 22-2 shows the System i libraries that are part of the SAP product. The libraries listed are those based on the SAP database ECC 5.0 and kernel release 6.40.

Table 22-2 System i libraries included in the SAP product

Library	Description	Initial number of objects	Initial size (approximate)
R3*REL*OPT	Library for optimized executables (example: 6.40 Unicode)	373	5 GB
R3*SID*DATA	Database library for: - SAP R/3 Enterprise 4.7) - WebAS ABAP 6.40 Unicode	 50,169 19,080	 55.8 GB [a] 17.5 GB
SAP*SID*DB	Database library for: SAP WebAS Java 6.40 Enterprise Portal 6.0	 485 727	 291 MB 775 MB
R3*SID*JRN SAP*SID*JRN	Journal receiver libraries for SAP WebAS ABAP and Java systems	1	_[a]
R3*SID*400	Library for work management objects	57	32 MB initial size 100-500 MB later
R3400	Library for non-*SID* specific objects	143	458 MB [c]
R3WRKnn	Internal SAP library	0	80 KB
R3*SID*xxxxx	SQL package library		_[c]

a) This value does not include customer data.
b) The amount depends on the activity on the system and the frequency with which the journal receivers are backed up.
c) The R3400 library contains objects that are not SID specific, including the OS collector data.
d) The R3SIDxxxxx libraries that contain SQL packages are dynamically created when you run the SAP application. Therefore, the number and size of libraries varies.

In Table 22-2, consider the following facts:

- *SID* is the SAP system ID (for example “P01”)
- *REL* is the SAP kernel release (for example “6.40” for the 6.40 kernel)
- *nn* is the instance number

Note: The values shown in Table 22-2 are approximations. They change depending on the SAP release and changes within an SAP release.

SAP Integrated File System structure

This section shows the Integrated File System structure of SAP systems. Figure 22-3 on page 385 illustrates the Integrated File System structure for pure ABAP environments.

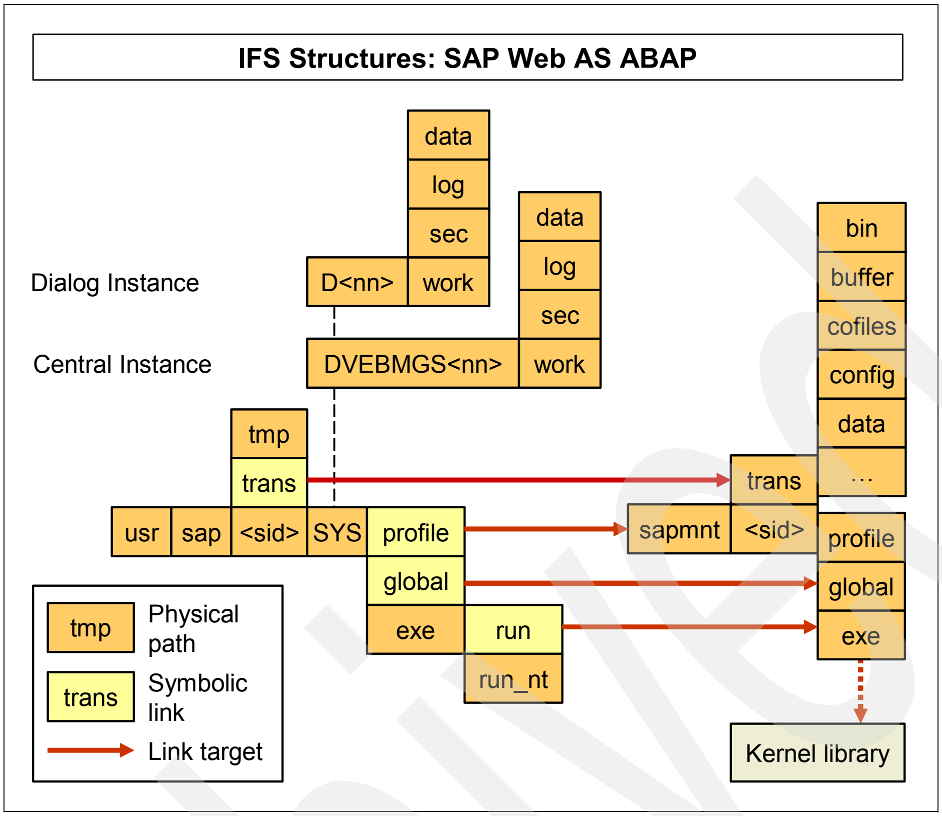

Figure 22-3 SAP WebAS ABAP system - directory structure

Figure 22-4 illustrates the Integrated File System structure in an ABAP plus Java environment.

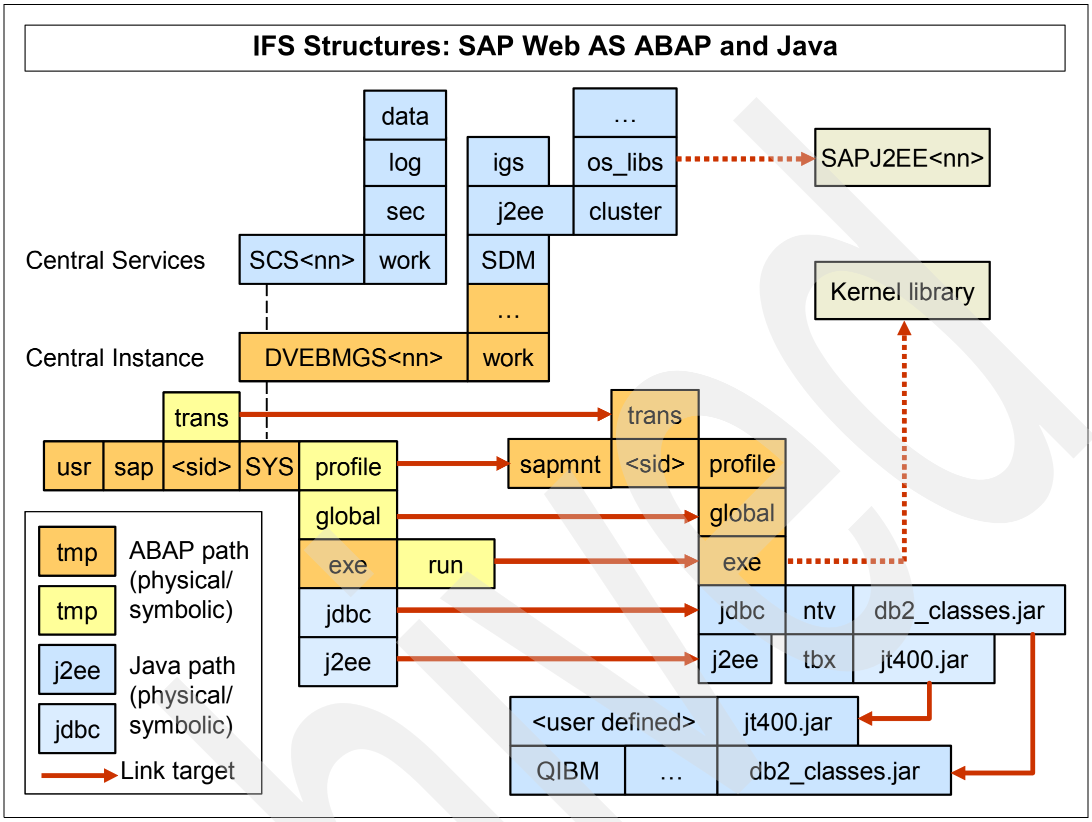

Figure 22-4 SAP WebAS ABAP and Java system - directory structure

Both environments have identical root-directory entries as follows:

- (/user/sap/)
- /sapmnt/

These directory structures allow an identical save-procedure for Integrated File System objects.

SAP objects to be saved

Save the SAP objects as shown in Table 22-3.

Table 22-3 When to save an object

Object	Frequency
R3*SID*DATA, SAP*SID*DB	Daily
R3*SID*JRN, SAP*SID*JRN	At least daily. Make sure the user ASP does not overflow.
R3*REL*OPT	After installing new levels of the kernel
R3*SID*400	After installing new levels of the kernel
R3400	After installing new levels of the kernel
R3*REL*CPIC	After installing new levels of the kernel
R3*REL*RFC	After installing new levels of the kernel
Integrated File System objects (refer to Table 22-4)	Daily

The R3*SID*xxxxx library does not need to be saved. SQL package libraries and SQL packages are created automatically.

Table 22-4 shows what to save in the Integrated File System.

Table 22-4 What to save in the Integrated File System

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.

Table 22-5 shows the recommended backup methods.

Table 22-5 Recommended backup methods

Object	Weekly backup	Daily backup
R/3*SID*DATA SAP*SID*DB	Offline (optional) See “Saving the entire system” on page 394	Online with ’SAVE ACTIVE (*SYNCLIB)’ function in SAVLIB-command: - wait time: saves database in a consisting state, needs to find a database-wide checkpoint. - *NOCMTBDY: saves database in a non-consistent state, needs journal receiver to complete the restore. Refer to “Saving the SAP database” on page 396
R3*SID*JRN SAP*SID*JRN		Online (SAVDLTRCV) Refer to 23.1.5, “Saving journal receivers” on page 405
Integrated File System objects		Online (SAV) Refer to “Saving the SAP Integrated File System objects” on page 423
R3*REL*OPT R3*SID*400 R3400 R3*REL*RFC (optional) R3*REL*CPIC (optional) R3SYS		---

22.4.3 Backup considerations

Refer to this section before saving any objects.

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: This includes the operating system (i5/OS), user profiles, authorities, configurations, system, and network values.
- Application software: This is 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: These contain the organization's information.

22.4.4 Recommendations prior to backup

Consider the following recommendations before you 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 IBM 359x or LTO tape drives. They are high speed, high capacity, and very reliable tape drives.

- 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 on to 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 application programming interfaces (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. Refer to the *System API Reference*, SC41-5801, for more information about these APIs.

  – Using the Backup and Recovery Media Services 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 SAP database and journal receivers to provide a better protection.

- Stopping the SAP system before offline backup

  End the SAP system (central and secondary instances) before you perform an offline backup.

- Symbolic links

  Saving symbolic links in the Integrated File System saves 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.

- Online backup

  For online backup use the i5/OS SAVLIB command with the Save Active function:

  – Use *SYNCLIB with system wait time for object lock (Save-While-Active) as your primary choice as long as your backup-job can find a database-wide synchronization point for a consistent backup.
  – Use *SYNCLIB with *NOCMTBDY parameter (Rapid Checkpoint Save-While-Active) when your backup-job cannot find a database-wide checkpoint at any time. Save the associated journal-receivers as shown in the job log to complete the backup.

- SAVR3SYS, DSCR3SYS, RCNR3SYS

  As of SAP Release 6.10 or i5/OS Release V5R3 these commands are no longer supported. For more information, refer to SAP note 202593.

- 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

  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 Integrated File System objects online because the backup always fails.

- Update history option

  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 Integrated File System objects online

  It is currently not possible to save the all Integrated File System objects of an SAP system online because, for example, developer trace files (/usr/sap/*SID*/DVEBMGS*nn*/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 kernel release (6.40).

  You have to either save the Integrated File System objects of an SAP system 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 should read the job log to check 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. Retain the information for disaster recovery.

23

Backup and recovery

This chapter outlines the standard System i backup and recovery facilities available. It also provides examples of backup and recovery facilities and procedures in the SAP system environment.

For in-depth information about these subjects, refer to:

- *Backup and Recovery,* SC41-5304
- SAP online documentation on the SAP Help Portal:

  `http://help.sap.com/`

- iSeries Information Center at:

  `http://publib.boulder.ibm.com/infocenter/iseries/v5r3/ic2924/index.htm`

These documents will help you plan backup, recovery, and system availability to protect and maintain crucial business functions.

23.1 Backup

Backup and recovery options for SAP systems on a System i model are managed using native System i server facilities. Backup Recovery and Media Services (BRMS) can provide automated tape backup and archive operations, recovery services, and tape media management to treat tape as an extension of direct access storage device (DASD) from the application point of view.

i5/OS provides many backup facilities that you can select 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. This section is for your general guidance for developing a suitable save strategy. It does not provide detailed information about save strategy and procedure. For more complete information, refer to *Backup and Recovery*, SC41-5304.

23.1.1 Backup methods

There are three methods of saving the SAP database on a System i model:

- Offline backup

  This means saving the objects while either the SAP system is down, or the System i server is in a restricted state, which implies that the SAP system is down.

- Online backup with Save-While-Active function

  This method saves objects in an SAP database library while they are in use by other jobs. At the beginning of the backup, all the objects in the SAP database library need to reach a checkpoint together, where all open commit cycles are successfully completed. After having reached this status, all these objects are saved in a consistent state in relationship to each other.

- Online backup with Rapid Checkpoint Save-While-Active function

  This function is available with i5/OS V5R3 onward and allows you to save the SAP database in a non-consistent state, regardless of whether it is in use by another job or not. Objects in the SAP database can reach checkpoints at different times and may not be in a consistent state in relationship to each other.

  To achieve data consistency, you also need to backup all appending journal receivers as shown in the job log. This function cannot save objects that are in use by uncommitted transactions that include Data Definition Language (DDL) statements such as ALTER TABLE, and so on. Run the APYJRNCHGX command in order to apply DDL statements without interruptions. For more details, refer to “Forward recovery” on page 419.

  The usage of this statement does not occur very often in an SAP system. For example, this can happen if transport requests are imported into a system. Therefore, we recommend that you avoid SAP system backups at the same time.

  Note: Use Save-While-Active as your primary choice for online backup. Consider Rapid Checkpoint Save-While-Active in the event that a database-wide checkpoint for the backup job cannot be found.

Table 23-1 lists the advantages and disadvantages of each method.

Table 23-1 Methods of saving the database

Method	Advantages	Disadvantages
Offline backup	► You do not have to manage synchronization. ► The fastest way to save the data.	► The SAP system needs to be down for a relatively long time. ► The content of the SAP system buffers is lost.
Online backup with Save-While-Active	► SAP system stays up and running. ► You do not lose the contents of the SAP system buffers.	The save operation with Data Manipulation Language (DML) or DDL may not be successful if the jobs do not commit their work within the specified amount of time.
Online backup with Rapid Checkpoint Save-While-Active	► SAP system stays up and running. ► Contents of the SAP system buffers are not lost. ► You do not need to find a database-wide synchronization point.	► The save operation with DDL statements may not be successful if the jobs do not commit their work within the specified amount of time. ► To achieve data consistency, you also need to backup (and apply in the event of a Restore) all related journal receivers.

Note: If you write your own save or backup programs using the **`save-while-active`** command and the system is in a dedicated mode, the save-while-active processing is ignored.

23.1.2 Initializing the tape

Before you begin with the backup operation, perform the following 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 23-1 on page 394. If your media density does not comply with your tape drive, you might have to specify a different density. Specify a volume identifier or take the existing one (*MOUNTED).

Figure 23-1 shows how to initialize a tape.

```
                           Initialize Tape (INZTAP)

 Type choices, press Enter.

 Device . . . . . . . . . . . . .   TAP01         Name
 New volume identifier  . . . . .   SAP           Character value, *NONE...
 New owner identifier . . . . . .   *BLANK
 Volume identifier  . . . . . . .   *MOUNTED      Character value, *MOUNTED
 Check for active files . . . . .   *YES          *YES, *NO, *FIRST
 Tape density . . . . . . . . . .   *CTGTYPE      *DEVTYPE, *CTGTYPE, *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 23-1 Initialize tape (INZTAP)

23.1.3 Offline backup

You can select the backup options either from the save menu (GO SAVE) or a command line entry.

Saving the entire system

This task saves the entire System i server, including all SAP system objects. This must be done after the initial installation of an SAP system is complete. We recommend entire system saves before or after major system changes such as the following:

- Hardware addition or hardware removals
- Extensive system configuration changes
- Operating system upgrades
- Major program temporary fix (PTF) installation (for example, applying cumulative PTFs)
- Application software additions, removals, and upgrades

> **Note:** This option requires a dedicated system or restricted state. This means that no subsystem is available, except the controlling subsystem (for example, QCTL) which only supports 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 initial program load (IPL) startup program, you must issue it manually to start the SAP system.

To perform backup operation, follow these steps:

1. Stop the SAP system using STOPSAP command.
2. Sign on from the console as QSECOFR.

3. Go to the save menu by using the GO SAVE command and page down. Select option **21** to save the entire system.
4. A screen appears which describes the selected option in detail.
5. Press **Enter** again to specify save options. You also can use the system reply list to perform an unattended backup operation.
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.

Saving all user data

Option 23 from the save menu saves all SAP 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)

This option does not save the Licensed Internal Code (LIC) or the operating system, but it saves all other data that is included in option 21, that is, save the entire system.

We recommend that you *do not* use this option because it does not reduce the backup time as much as compared to option 21. You cannot restore an older version of the Licensed Internal Code or the operating system which can be helpful in case of a software failure that came along with a PTF.

Saving the SAP database

To perform the backup operation (in our example, for the WebAS ABAP database), perform the following steps:

1. Stop the SAP system using the STOPSAP command.
2. Sign on as QSECOFR or *SID*OFR.
3. Save the SAP database by selecting option **2**, **Save libraries**, from the save menu, or prompt the SAVLIB command as Figure 23-2 shows.

```
                            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
Starting library . . . . . . . .   *FIRST        Name, *FIRST
Use optimum block  . . . . . . .   *YES          *YES, *NO

                           Additional Parameters

Target release . . . . . . . . .   *CURRENT      *CURRENT, *PRV, V5R1M0...
                                                                    More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 23-2 SAVLIB: Save the SAP database (Part 1 of 3)

4. Specify ***YES** for Object pre-check and ***YES** for Save access paths as shown in Figure 23-3.

```
                              Save Library (SAVLIB)

 Type choices, press Enter.

 Update history . . . . . . . . .   *YES          *YES, *NO
 Clear  . . . . . . . . . . . . .   *NONE         *NONE, *ALL, *AFTER, *REPLACE
 Object pre-check . . . . . . . .   *yes          *NO, *YES
 Save active  . . . . . . . . . .   *NO           *NO, *LIB, *SYNCLIB, *SYSDFN
 Save active wait time:
   Object locks . . . . . . . . .   120           0-99999, *NOMAX
   Pending record changes . . . .   *LOCKWAIT     0-99999, *LOCKWAIT...
   Other pending changes  . . . .   *LOCKWAIT     0-99999, *LOCKWAIT, *NOMAX
 Save active message queue  . . .   *NONE         Name, *NONE, *WRKSTN
   Library  . . . . . . . . . . .     *LIBL       Name, *LIBL, *CURLIB
 Save access paths  . . . . . . .   *yes          *SYSVAL, *NO, *YES
 Save file data . . . . . . . . .   *YES          *YES, *NO
 Storage  . . . . . . . . . . . .   *KEEP         *KEEP, *FREE
 Data compression . . . . . . . .   *DEV          *DEV, *NO, *YES, *LOW...
 Data compaction  . . . . . . . .   *DEV          *DEV, *NO

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

Figure 23-3 SAVLIB: Save the SAP database (Part 2 of 3)

5. Specify ***PRINT** for the Output parameter as shown in Figure 23-4.

```
                              Save Library (SAVLIB)

 Type choices, press Enter.

 Libraries to omit  . . . . . . .   *NONE         Name, generic*, *NONE
                + for more values
 Objects to omit:
   Object . . . . . . . . . . . .   *NONE         Name, generic*, *NONE, *ALL
     Library  . . . . . . . . . .     *ALL        Name, generic*, *ALL
   Object type  . . . . . . . . .   *ALL          *ALL, *ALRTBL, *BNDDIR...
                + for more values
 ASP device . . . . . . . . . . .   *             Name, *, *SYSBAS, *CURASPGRP
 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 23-4 SAVLIB: Save the SAP library (Part 3 of 3)

6. 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.

23.1.4 Online backup of the SAP database

You can backup the SAP database while your SAP system is up and running. This helps to minimize the downtime of your application, and avoids losing the contents of the SAP system buffers. Due to the additional processing involved, the online backup takes longer to complete. However, to minimize the duration of the backup, it is a good idea to perform a save-while-active during periods of low activity.

There are two possibilities to perform an online backup:

- Save-While-Active
- Rapid Checkpoint Save-While-Active

To avoid unnecessary downtime, make Save-While-Active your primary choice for online backup. Consider Rapid Checkpoint Save-While-Active if you cannot find a database-wide checkpoint for your backup job. For more information about how to position both methods, refer to 23.1.1, “Backup methods” on page 392. For more details, refer to SAP note 825473.

Save-While-Active

Prior to actually saving the objects to the save media, 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. 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, that is, the checkpoint cannot be reached, if the jobs do not commit their work within the specified amount of time. In that case, consider Rapid Checkpoint Save-While-Active as an alternate solution.

When commitment control is active for any job on the system, which is the case for jobs running the SAP instance subsystem, the system performs the following tasks during the save-while-active checkpoint processing:

- Identifies all the 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, it 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. Allow a maximum wait time of 1200 seconds.

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. It is the time the checkpoint image was established. This process consumes a lot of resources.

The save operation fails if no commitment boundary can be reached within the predefined wait time. In this case, all users will have to wait the maximum wait time as specified in the SAVE command.

In our example, the SAP database library (the WebAS ABAP library R3R01DATA) is saved while it is in use by an active SAP system.

To perform the backup operation, perform the following steps:

1. Sign on as QSECOFR or *SID*OFR.
2. Prompt the SAVLIB command, and specify the value as shown in Figure 23-5.

```
                          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
 Starting library . . . . . . .   *FIRST        Name, *FIRST
 Use optimum block  . . . . . .   *YES          *YES, *NO

                           Additional Parameters

 Target release . . . . . . . .   *CURRENT      *CURRENT, *PRV, V5R1M0...
                                                                    More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Figure 23-5 SAVLIB: Save-While-Active (Part 1 of 3)

3. Specify ***YES** for Object pre-check, ***SYNCLIB** for Save active. We recommend that you change the default value for Save active lockwait wait time to 1200 as Figure 23-6 shows.

```
                             Save Library (SAVLIB)

Type choices, press Enter.

Update history . . . . . . . . .   *YES          *YES, *NO
Clear  . . . . . . . . . . . . .   *NONE         *NONE, *ALL, *AFTER, *REPLACE
Object pre-check . . . . . . . . > *YES          *NO, *YES
Save active  . . . . . . . . . . > *SYNCLIB      *NO, *LIB, *SYNCLIB, *SYSDFN
Save active wait time:
  Object locks . . . . . . . . . > 1200          0-99999, *NOMAX
  Pending record changes . . . .   *LOCKWAIT     0-99999, *LOCKWAIT...
  Other pending changes  . . . .   *LOCKWAIT     0-99999, *LOCKWAIT, *NOMAX
Save active message queue  . . .   *NONE         Name, *NONE, *WRKSTN
  Library  . . . . . . . . . . .     *LIBL       Name, *LIBL, *CURLIB
Save access paths  . . . . . . .   *SYSVAL       *SYSVAL, *NO, *YES
Save file data . . . . . . . . .   *YES          *YES, *NO
Storage  . . . . . . . . . . . .   *KEEP         *KEEP, *FREE
Data compression . . . . . . . .   *DEV          *DEV, *NO, *YES, *LOW...
Data compaction  . . . . . . . .   *DEV          *DEV, *NO


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

Figure 23-6 SAVLIB: Save-While-Active (Part 2 of 3)

4. Specify ***PRINT** for Output as Figure 23-7 shows.

```
                             Save Library (SAVLIB)

Type choices, press Enter.

Libraries to omit  . . . . . . .   *NONE         Name, generic*, *NONE
               + for more values
Objects to omit:
  Object . . . . . . . . . . . .   *NONE         Name, generic*, *NONE, *ALL
    Library  . . . . . . . . . .     *ALL        Name, generic*, *ALL
  Object type  . . . . . . . . .   *ALL          *ALL, *ALRTBL, *BNDDIR...
               + for more values
ASP device . . . . . . . . . . .   *             Name, *, *SYSBAS, *CURASPGRP
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 23-7 SAVLIB: Save While Active (Part 3 of 3)

5. Check the job log and the save listing to be sure that all the objects could be saved together even if the object pre-check was used.

Rapid Checkpoint Save-While-Active

The Rapid Checkpoint processing saves objects without requiring transactions with pending record changes to reach a commitment boundary. Therefore, objects can be saved with partial transactions.

> **Note:** If DDL instructions are executed during a save with the value *NOCMTBDY, then the affected work process sends a query message `CPA8351` to the QSYSOPR message queue. The work process remains in status MSGW until the message is answered. The reply to the message can be `C` for Cancel or `R` for Repeat (after waiting). In order to prevent operator replies being necessary during a save, an entry with reply `R` should be made for the message ID CPA8351 in the following system reply list:
>
> `ADDRPYLE SEQNBR(xxxx) MSGID(CPA8351) RPY('R')`

For more information about Rapid Checkpoint Save-While-Active, refer to *Improve Whole System Backups with the New Save-While-Active Function*, REDP-7200.

To perform the backup operation, perform the following steps:

1. Sign on as QSECOFR or *SID*OFR.
2. Prompt the SAVLIB command, and specify the value as shown in Figure 23-8.

```
                              Save Library (SAVLIB)

 Type choices, press Enter.

 Library  . . . . . . . . . . . LIB          > R3R01DATA
                          + for more values
 Device . . . . . . . . . . . . DEV          > TAP01
                          + for more values
 Volume identifier  . . . . . . VOL            *MOUNTED

                          + for more values

 Sequence number  . . . . . . . SEQNBR         *END
 Label  . . . . . . . . . . . . LABEL          *LIB
 File expiration date . . . . . EXPDATE        *PERM
 End of media option  . . . . . ENDOPT         *REWIND
 Starting library . . . . . . . STRLIB         *FIRST
 Use optimum block  . . . . . . USEOPTBLK      *YES


                                                                    More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
 Already at top of area.
```

Figure 23-8 SAVLIB: Rapid Checkpoint Save-While-Active (Part 1 of 3)

3. Specify ***YES** for Object pre-check, ***SYNCLIB** for Save active, and ***NOCMTBDY** for saving objects without requiring transactions with pending record changes to reach a commitment boundary as Figure 23-9 shows.

```
                              Save Library (SAVLIB)

Type choices, press Enter.


                           Additional Parameters

Target release . . . . . . . . . TGTRLS         *CURRENT
Update history . . . . . . . . . UPDHST         *YES
Clear  . . . . . . . . . . . . . CLEAR          *NONE
Object pre-check . . . . . . . . PRECHK       > *YES
Save active  . . . . . . . . . . SAVACT       > *SYNCLIB
Save active wait time:           SAVACTWAIT
  Object locks . . . . . . . . .              > 120
  Pending record changes . . . .                *NOCMTBDY
  Other pending changes  . . . .                *LOCKWAIT
Save active message queue  . . . SAVACTMSGQ     *NONE
  Library  . . . . . . . . . . .                  *LIBL
Save access paths  . . . . . . . ACCPTH       > *YES
Save file data . . . . . . . . . SAVFDTA        *YES
                                                                       More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 23-9 SAVLIB: Rapid Checkpoint Save-While-Active (Part 2 of 3)

> **Note:** Other pending changes, such as DDL object level changes, or changes by application programming interface (API), still need to be resolved until the backup can be completed. The DDL object level changes typically happen if you are activating transports in your SAP system.
>
> We recommend that you do not run transports during a backup of your SAP system.

4. Specify an outfile name for Output as shown in Figure 23-10 on page 403.

```
                          Save Library (SAVLIB)                    Level: 2

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
ASP device . . . . . . . . . . .   *             Name, *, *SYSBAS, *CURASPGRP
Output . . . . . . . . . . . . .   *outfile      *NONE, *PRINT, *OUTFILE
File to receive output . . . . .   uscsave       Name
  Library  . . . . . . . . . . .     engelb      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 23-10 SAVLIB: Rapid Checkpoint Save-while-Active (Part 3 of 3)

5. Check the job log and the save listing to ensure that all objects can be saved together even if the object pre-check was used.
6. In order to establish which tables are saved while transactions are pending, and which journal receivers need to be available so that you can recover completely from the backup, you can specify an outfile on the SAVLIB command. This is the USCSAVE file, as shown in Example 23-1. From this file, you can query which tables have pending transactions and which is the oldest journal receiver which has to be saved.

Example 23-1 Specifying an outfile on the SAVLIB command

```
Select SRONAM, SROLIB, SROJRN, SROJTL
FROM ENGELB/USCSAVE
WHERE SROPRT = 1
```

Alternatively, you can use the STRSQL command as Figure 23-11 shows.

```
                              Specify SELECT Statement

 Type SELECT statement information.  Press F4 for a list.

   FROM files . . . . . . . . .   USCSAVE
   SELECT fields  . . . . . . .   SRONAM, SROLIB, SROJRN, SROJRL
   WHERE conditions . . . . . .   SROPRT = '1'
   GROUP BY fields  . . . . . .
   HAVING conditions  . . . . .
   ORDER BY fields  . . . . . .
   FOR UPDATE OF fields . . . .

                                                                     Bottom
 Type choices, press Enter.

   DISTINCT records in result file  . . . . . . . .   N    Y=Yes, N=No
   UNION with another SELECT  . . . . . . . . . . .   N    Y=Yes, N=No
   Specify additional options . . . . . . . . . . .   N    Y=Yes, N=No



 F3=Exit         F4=Prompt   F5=Refresh   F6=Insert line   F9=Specify subquery
 F10=Copy line   F12=Cancel   F14=Delete line   F15=Split line   F24=More keys
```

Figure 23-11 STRSQL: analyzing associated journal receivers

Figure 23-12 shows the results of this.

```
                Display Data

Data width . . . . . . :      46
 Position to line  . . . . .
Shift to column  . . . . . .
....+....1....+....2....+....3....+....4....+.
 Object    Library   Journal   Journal
 Name      Name      Receiver  Receiver
                     Name      Library
 TESTTABLE  R3R01DATA  QSQJRN0001 R3R01JRN
 ********  End of data  ********

Bottom
 F3=Exit   F12=Cancel   F19=Left   F20=Right   F21=Split   F22=Width 80
```

Figure 23-12 Tables with pending transactions and associated journal receivers

In our example, you also need to save the journal receiver QSQJRN001 to achieve a consistent backup.

> **Note:** You need all associated journal receivers to achieve data consistency when you need to backup your system.

23.1.5 Saving journal receivers

The SAP 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.

The following are some considerations when saving journal receivers:

- Separate tapes for journal receivers

  Use separate backup media for the SAP database and journal receivers to provide a better protection.

- Threshold

  The default threshold for SAP 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 auxiliary storage pool (ASP) overflow

  Ensure 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.

  > **Note:** An independent ASP cannot overflow into another ASP. Provide enough disk space to avoid unplanned SAP system downtime in case of an Independent Auxiliary Storage Pool (IASP) overflow.

  ASP overflow is discussed in 9.7.4, “ASP overflow” on page 113. Recovering from an overflow situation is introduced in 23.2.1, “User ASP overflow recovery” on page 412.

- Do not change MNGRCV(*SYSTEM) and DLTRCV(*NO).

  For production systems we recommend that 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.

  > **Note:** It is mandatory to save the journal receivers when you are using Rapid Checkpoint Save-While-Active for online backup of the SAP database.

- Keep the backup tapes.

  Keep the tapes of the journal receiver backup to minimize the amount of data loss in case of an emergency. The journal receiver chain must not be broken. You can reuse the tapes after the SAP database has been saved in a consistent state.

- Do not save attached journal receivers.

  We recommend that you do not save attached journal receivers (showing the status ATTACHED) because they are not fully saved. When you have to restore an incompletely saved journal receiver, the status shows 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 following command:

  ```
  CHGJRN JRN(R3SIDDATA/QSQJRN) JRNRCV(*GEN)
  ```

This command generates a new journal receiver, detaches the current one, and attaches the new one.

Saving manually

This section explains how to save journal receivers individually. To do this, perform the following steps:

1. Enter the following command:

   `WRKJRNA JRN(R3SIDDATA/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 23-13.
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.

```
                          Work with Receiver Directory

 Journal  . . . . . . :   QSQJRN         Library  . . . . . . :   R3USCDATA

 Total size of receivers (in kilobytes)  . . . . . . . . . . . :   3864

 Type options, press Enter.
   4=Delete   8=Display attributes


                                     Attach                        Save
 Opt  Receiver    Library     Number  Date       Status              Date
      QSQJRN1266  R3USCJRN    00001   10/11/05  SAVED               10/11/05
      QSQJRN1267  R3USCJRN    00002   10/11/05  ONLINE              00/00/00
      QSQJRN1268  R3USCJRN    00003   10/11/05  ATTACHED            00/00/00



                                                                         Bottom
 Parameters or command
 ===>
 F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F11=Display size
 F12=Cancel
```

Figure 23-13 WRKJRNA and F15: work with receiver directory

The advantage in using this technique is that each journal receiver is saved only once. You do not have problems with duplicate names and partial receivers if you need to restore. The disadvantage of this technique is that it requires manual effort to determine the names of the journal receivers to be saved.

Automatic save using a CL program

You can use a Command Language (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 message 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.

Automatic save using the SAVDLTRCV command

We recommend that you use the backup method 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.

> **Note:** 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 addition, they can still use the receivers to transfer the data to the target system and control the deletion of the files themselves. In this case, we recommend you use the backup method 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. To use this program, perform the following steps:

1. Before using this program, sign on as *SID*OFR and create a message queue using the following command:

   ```
   CRTMSGQ MSGQ(R3SID400/SAVDLTRCV)
   ```

2. Change the journal to use the messages queue as follows:

   ```
   CHGJRN JRN(R3SIDDATA/QSQJRN) MSGQ(R3SID400/SAVDLTRCV)
   ```

3. Submit a batch job to run this program as follows:

   ```
   SBMJOB CMD(SAVDLTRCV MSGQ(R3SID400/SAVDLTRCV) DEV(TAP01)) JOB(SAVDLTRCV)
   ```

Alternatively, add the following entries to your start profile:

```
_SAVDLTCMD = SAVDLTRCV MSGQ(R3SID400/SAVDLTRCV) DEV(TAP01)
Execute_04 = local SBMJOB CMD($(_SAVDLTCMD)) JOB(SAVDLTRCV)
Stop_Program_04 = local SAVDLTRCVE MSGQ(R3SID400/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 23-2.

Table 23-2 SAVDLTRCV parameters

Parameter	Description
MSGQ	Qualified message queue name as passed by command. Recommended value: R3*SID*400/SAVDLTRCV

Parameter	Description
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)
DLTRTYTIM	If the journal receiver cannot be deleted (CPF7024), the job is delayed some time before retrying 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 following command:

```
SAVDLTRCVE MSGQ(R3SID400/SAVDLTRCV)
```

23.1.6 Saving SAP Integrated File System objects

You can backup the Integrated File System objects of your SAP system whether your system is up and running or not. In case that your SAP system is active, some SAP log and trace files may not be saved because of object locks. But these objects are temporary ones. They are automatically rebuilt when you start your SAP system again. Therefore, they are not mandatory for a consistent backup of an SAP system.

To perform the backup operation follow these steps:

1. Stop your SAP system.
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 23-14.

```
                              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
Name pattern:
  Pattern  . . . . . . . . . . .   '*'

  Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT
              + for more values
Directory subtree  . . . . . . .   *ALL          *ALL, *DIR, *NONE, *OBJ, *STG
Save active  . . . . . . . . . .   *NO           *NO, *YES, *SYNC
                                                                      More...
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Figure 23-14 Save object

4. Enter the Integrated File System paths as shown in Figure 23-15.

```
                              Save Object (SAV)

 Type choices, press Enter.

 Objects:
   Name . . . . . . . . . . . . . . > 'usr/sap/R01'

   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT

   Name . . . . . . . . . . . . . . > '/sapmnt/R01'

   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT

   Name . . . . . . . . . . . . . . > '/sapmnt/trans'

   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT
                + for more values



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

Figure 23-15   OBJ parameter

> **Note:** In case you are running an SAP application in an IASP, you also can specify the directories /*asp*/usr/sap*sid*, /*asp*/sapmnt/*sid* and /*asp*/sapmnt/trans.

5. Press **F9** and **F10** to see the additional parameters as shown in Figure 23-16. Specify ***PRINT** for the Output and ***LEAVE** for the End of media option to save rewind time.

```
                              Save Object (SAV)

Type choices, press Enter.

Name pattern:
  Pattern  . . . . . . . . . . .   '*'

  Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT
               + for more values
Directory subtree  . . . . . . .   *ALL          *ALL, *DIR, *NONE, *OBJ, *STG
Save active  . . . . . . . . . .   *NO           *NO, *YES, *SYNC
Output . . . . . . . . . . . . .   *PRINT

Volume identifier  . . . . . . .   *MOUNTED
               + for more values
Label  . . . . . . . . . . . . .   *GEN
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
                                                                       More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 23-16 Save Object

6. Specify ***YES** for Object pre-check and ***YES** for Update history as shown in Figure 23-17.

```
                              Save Object (SAV)

Type choices, press Enter.

Type of output information . . .   *ALL          *ALL, *ERR, *SUMMARY

                           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, V5R1M0...
Update history . . . . . . . . .   *YES          *NO, *YES, *SYS, *PC

Clear  . . . . . . . . . . . . .   *NONE         *NONE, *ALL, *AFTER, *REPLACE
Data compression . . . . . . . .   *DEV          *YES, *NO, *DEV, *LOW...
                                                                       More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 23-17 Save object, continued

7. Check the job log to make sure that all Integrated File System objects have been saved. You should retain the information for disaster recovery.

23.2 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 can be selected from the standard Restore menu, which can be accessed by entering the GO RESTORE command from an System i 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, up to the previous evening.

23.2.1 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 the message CPI0953 to the QSYSOPR message queue warning you that an ASP is approaching its storage threshold. The system sends the message CPI0954 when the storage threshold is exceeded and the ASP is in overflowed status.

Reset a user ASP in overflowed status as soon as possible. An overflowed ASP affects system performance. It also makes recovery more difficult and can 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.

> **Note:** There is no ASP overflow function for IASPs. In case of an IASP overflow, your SAP system stops.

23.2.2 Restoring the entire system

The following are the major steps and commands to restore the entire system:

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 using the RSTUSRPRF command.
4. Restore the configuration using the RSTCFG command.
5. Restore all libraries using the RSTLIB command.
6. Restore the document library objects using the RSTDLO command.
7. Restore the integrated file system using the RST command.
8. Apply the journal changes using the APYJRNCHGX command. Refer to "Forward recovery" on page 419.
9. Restore public and private authorization using the RSTAUT command.

For more information, refer to the manual *Backup and Recovery,* SC41-5304.

23.2.3 Restoring the SAP system 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.

> **Note:** The sequence of restoring the SAP system environment is very important. Make sure that the journal receiver library R3*SID*JRN is restored before the SAP database, R3*SID*DATA or SAP*SID*DB. If not, a new journal receiver chain is created in the database library.
>
> Do not restore individual files to the SAP database. Such an action can result in an inconsistent database. Do this only under the direction of support personnel.

Perform the following steps to restore the SAP system environment:

1. Stop the SAP system if active.
2. Sign on as QSECOFR.
3. Change your job so the message queue does not wrap when it is full. Use the following command:

   ```
   CHGJOB JOBMSGQFL(*PRTWRAP)
   ```

4. Restore the SAP kernel library and configuration library if necessary and run the FIXR3OWNS command.
5. Restore the SAP Integrated File System structure if necessary using the following command:

```
RST
```

6. Restore the journal receivers for forward recovery.
7. Delete the SAP Advanced Business Application Programming (ABAP) or Java database using the following command:

```
DLTLIB LIB(R3SIDDATA)
DLTLIB LIB(SAPSIDDB)
```

8. Restore the SAP ABAP or Java library with the following command:

```
RSTLIB SAVLIB(R3SIDDATA) DEV(TAP01) MBROPT(*ALL) ALWOBJDIF(*ALL)
RSTLIB SAVLIB(SAPSIDDB) DEV(TAP01) MBROPT(*ALL) ALWOBJDIF(*ALL)
```

Check the job log to make sure that all objects of the database library restored successfully. Monitor CPI3731, if you were using Rapid Checkpoint Save-While-Active to backup the database. There can be database objects being saved with partial transactions as shown in Figure 23-18. In that case you need to apply the associated journal receivers. For more information, refer to 23.2.4, “Recovering the SAP database” on page 415.

```
                          Additional Message Information

 Message ID . . . . . . :   CPI3731       Severity . . . . . . . :   10
 Message type . . . . . :   Information
 Date sent  . . . . . . :   10/28/05      Time sent  . . . . . . :   07:07:13

 Message . . . . :   *FILE TESTTABLE saved with partial transaction.
 Cause . . . . . :   Object TESTTABLE type *FILE in library R3R01DATA was saved
   with one or more partial transactions. If the object is a database file, the
   member name is TESTTABLE.
 Recovery  . . . :
     -- If this is a restore operation, you cannot use this object until you
   apply or remove journaled changes (APYJRNCHG or RMVJRNCHG command) to reach
   commit boundaries. To apply or remove the changes, you will need journal
   QSQJRN in library R3R01DATA, and the chain of journal receivers starting
   with QSQJRN1001 in library R3R01JRN on ASP device *SYSBAS.
     -- If this is a save operation, make sure that you save journal QSQJRN in
   library R3R01DATA and the chain of journal receivers starting with
                                                                        More...
 Press Enter to continue.

 F3=Exit   F6=Print   F9=Display message details   F12=Cancel
 F21=Select assistance level
```

Figure 23-18 DSPJOBLOG: message CPI3731

9. Restore the SAP Integrated File System structure:

```
RST DEV(’/QSYS.LIB/TAP01.DEVD’) OBJ((/usr/sap/*’ *INCLUDE) (’/sapmnt/*’ *INCLUDE))
```

> **Note:** In case you are running SAP in an IASP, you also need to restore the directories /*asp*/usr/sap*sid*, /*asp*/sapmnt/*sid* and /*asp*/sapmnt/trans.

10. Restore the private authorities using the following command:

```
RSTAUT USRPRF(*ALL)
```

11. Sign off and sign on as *SID*OFR.

12. Delete Structured Query Language (SQL) packages with the command:

```
DLTR3PKG SID(SID) PKGTYPE(*ALL
```

23.2.4 Recovering the SAP database

The reasons to recover the SAP database are:

- The database is saved with partial transactions. For example, a Rapid Checkpoint Save-While-Active is used.
- A system failure or power failure that caused damage to the database objects
- A disk failure that caused damage to the database objects
- Program failure or human error that caused damage to the content of the database.

Use the associated journal receivers to recover the database.

Database consistency

The following are the important considerations with regard to database consistency:

- Ensure that commitment transaction boundaries are honored on the apply or remove journaled changes operations by using CMTBDY(*YES) for any APYJRNCHGX or RMVJRNCHG command. Otherwise, the database can be inconsistent.
- Always recover all the database files in the SAP libraries using the APYJRNCHGX parameter FILE(R3*SID*DATA/*ALL) for ABAP and FILE(SAP*SID*DB/*ALL) for Java. Otherwise, the SAP database becomes inconsistent.
- If the recovery process fails, it does not terminate at a commitment boundary. The SAP database is therefore inconsistent.

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 is 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 the *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 ended. Certain illogical conditions, such as a duplicate key in a database file defined as unique, also cause processing to end.

> **Note:** Use OBJERROPT(*CONTINUE) for the APYJRNCHGX command. This enables to continue processing of journal entries for other objects.

DB2 UDB for iSeries offers two ways to recover a database to a specific point in time.

Table 23-3 shows you which recovery procedure to use for each type of outage.

Table 23-3 When to use what recovery procedure

Type of outage	Recommended recovery procedure
Database saved with partial transactions	Forward recovery
System failure or power failure	Forward recovery
Disk failure	Forward recovery
Program failure or human error	Backout recovery, if possible (DDL statements cannot been rolled back)

If you need to recover the database due to a program failure or human error, consider the following aspects:

- Find out at what point in time the database was last in a consistent state.
- When you last backed up the data
- It is faster to perform a backout recovery since you do not have to restore the database.
- It is 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 to go back to a certain database level.

Backout recovery

Depending on the type of damage to the physical file, that is 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 SAP database. Therefore, it is usually the fastest method to reset the SAP database.

Use the Remove Journaled 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 table. For example, assume that an application updated entries incorrectly for a certain period of time. In this case, you can remove the changes from the table until that application first opened the member.

To perform the backout recovery, perform the following steps:

1. Stop the SAP database if active.
2. Sign on as QSECOFR.
3. Make sure that no job locks the SAP database. Use the following command:

   ```
   WRKOBJLCK OBJ(R3SIDDATA) TYPE(*LIB)
   ```

4. Determine the point in time when the SAP database was last in a consistent state. For example, let's assume the database was last all right at about 10/31/05, 18:00:00.
5. Restore the necessary journal receivers.
6. Search for F code journal entries in the corresponding receiver chain using the following command:

   ```
   DSPJRN JRN(R3RO5DATA/QSQJRN) RCVRNG(*CURCHAIN) FROMTIME('10/31/05' '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 check if the process will end. Depending on the number and types of entries found, it is better to use forward recovery. If there aren't any entries, proceed with the next step.

For backout recovery, you should check the journal sequence number for the ending point to remove changes that were journaled. Use the following command to list the commitment boundaries in the 30 minutes prior to the failure:

```
DSPJRN JRN(R3R05DATA/QSQJRN) RCVRNG(*CURCHAIN) FROMTIME('10/28/05' '17:30:00')
TOTIME('10/28/05' '18:00:00') JRNCDE((C))
```

This command shows you a screen similar to the example in Figure 23-19.

```
                             Display Journal Entries

 Journal  . . . . . . :   QSQJRN          Library  . . . . . . :   R3R05DATA

 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 23-19 DSPJRN: Finding the journal sequence number

Perform the following steps:

1. Note the journal sequence number from the last C code and CM type entry. In this case, the sequence number is 5647727.

2. Prompt the RMVJRNCHG command, and enter the values as illustrated in Figure 23-20.

```
                       Remove Journaled Changes (RMVJRNCHG)

 Type choices, press Enter.

 Journal  . . . . . . . . . . . . > QSQJRN        Name
   Library  . . . . . . . . . . . >   R3RO1DATA   Name, *LIBL, *CURLIB
 Journaled file identification:
   Journaled physical file  . . . > *ALL          Name, *ALL
     Library  . . . . . . . . . . >   R3RO1DATA   Name, *LIBL, *CURLIB
   Member . . . . . . . . . . . . > *FIRST        Name, *ALL, *FIRST
                 + for more values
 Objects:
   Object . . . . . . . . . . . .                 Name, *ALL
     Library  . . . . . . . . . .     *LIBL       Name, *LIBL, *CURLIB
   Object type  . . . . . . . . .                 *FILE, *DTAARA
   Member, if data base file  . .   *ALL          Name, *ALL, *FIRST
                 + for more values



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

Figure 23-20 RMVJRNCHG (Part 1 of 2)

3. Specify ***YES** for Commitment boundary as shown in Figure 23-21.

```
                       Remove Journaled Changes (RMVJRNCHG)

 Type choices, press Enter.

 Range of journal receivers:
   Starting journal receiver  . .   *CURRENT      Name, *CURRENT, *LASTSAVE
     Library  . . . . . . . . . .                 Name, *LIBL, *CURLIB
   Ending journal receiver  . . .                 Name
     Library  . . . . . . . . . .                 Name, *LIBL, *CURLIB
 Starting large sequence number     *LAST
 Ending large sequence number . .   5647727
 Fully qualified job name . . . .                 Name
   User . . . . . . . . . . . . .                 Name
   Number . . . . . . . . . . . .                 000000-999999
 Commitment boundary  . . . . . .   *YES          *YES, *NO
 Option . . . . . . . . . . . . .   *NONE         *NONE, *IGNINQMSG
 Object error option  . . . . . .   *CONTINUE     *CONTINUE, *END
 Output . . . . . . . . . . . . .   *NONE         *NONE, *OUTFILE



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

Figure 23-21 RMVJRNCHG (Part 2 of 2)

4. Sign off and sign on as *SID*OFR.
5. Delete SQL packages with the following command:

```
DLTR3PKG SID(SID) PKGTYPE(*ALL
```

6. Start the SAP system.

Forward recovery

If a table becomes damaged or is not usable, you can recover it using the Apply Journaled Changes (APYJRNCHGX) command directly. You must first reestablish the table to a condition that you know is undamaged.

> **Note:** We recommend that you do not use the APYJRNCHG command for forward recovery, because it is not capable of applying object-level journal entries.

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.

For more information about how to find the associated receivers when you are performing a Rapid Checkpoint Save-While-Active, refer to "Rapid Checkpoint Save-While-Active" on page 401 and 23.2.3, "Restoring the SAP system environment" on page 413.

The system applies the changes to the file in the same order as they were originally made. When you use the APYJRNCHGX command, the file cannot be in use by anyone else.

> **Note:** Avoid deleting the QSQJRN journal in the R3*SID*DATA library whenever possible. Otherwise, the restore forces a chain break in the journal receiver. The FROMENT(*LASTSAVE) and TOENT(*LASTRST) parameters may not be possible on the APYJRNCHGX command.

To perform the forward recovery, follow these steps:

1. Stop the SAP system if active.
2. Sign on as QSECOFR.
3. Make sure that no job locks the SAP database. To do this, use the following command:

```
WRKOBJLCK OBJ(R3SIDDATA) TYPE(*LIB)
```

4. Determine the point in time when the SAP database was last in a consistent state. For example, let us assume the database was last OK at about 10/31/05, 18:00:00.
5. Restore the necessary journal receivers.
6. In case you have not lost the journal QSQJRN in database library, perform the following steps:
   a. Lock the journal objects from another job using the following command:

   ```
   WRKJRNA JRN(R3SIDDATA/QSQJRN)
   ```

   b. Clear the database library using the following command:

   ```
   SBMJOB CMD(CLRLIB LIB(R3SIDDATA)) JOB(CLRR3LIB)
   ```

   The journal object can be deleted with this command if it is not locked.

   c. Display the library to make that only the journal remains. Use the following command:

   ```
   DSPLIB LIB(R3SIDDATA)
   ```

7. Restore the SAP database from tape using the command:

```
SBMJOB CMD(RSTLIB SAVLIB(R3SIDDATA) DEV(TAP01) OPTION(*NEW) MBROPT(*ALL) ALWOBJDIF(*ALL))
JOB(RSTR3LIB) JOBMSGQFL(*PRTWRAP)
```

   Check the job log to make sure that all objects of the database library, except for the journal, have been restored successfully. Monitor CPI3731, if you were using Rapid Checkpoint Save-While-Active to backup the database. There can be database objects being saved with partial transactions as shown in Figure 23-18 on page 414. In that case you need to apply the associated journal receivers.

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 23-22:

```
                               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

                              ASP
 Opt  Journal     Library     Device      Text
  9   QSQJRN      R3R01DATA               COLLECTION - created by SQL









                                                                 Bottom
 Command
 ===>
 F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

Figure 23-22 WRKJRN: associate receivers with journal

> **Note:** This function may not run if you have implemented BRMS and the job QBRMNET is active in subsystem Q1ABRMNET. To solve this problem, stop the job and redo the command again.

9. Use the WRKJRNA command, and press **F15** to check whether all receivers have been added to the journal.

10. Type the APYJRNCHGX command and enter the values as shown in Figure 23-23 on page 421:

```
                   Apply Journaled Changes Extend (APYJRNCHGX)

 Type choices, press Enter.

 Journal  . . . . . . . . . . .   QSQJRN        Name
   Library  . . . . . . . . . .     R3R01DATA   Name, *LIBL, *CURLIB
 Journaled file identification:
   Journaled physical file  . . .   *all          *ALL
     Library  . . . . . . . . .     R3R01DATA   Name
   Member . . . . . . . . . . . .   *ALL          *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 large sequence number     *LASTSAVE
 Ending large sequence number . .   *LASTRST



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

Figure 23-23 APYJRNCHGX (Part 1 of 2)

11. Specify the ending date and time and ***YES** for Commitment boundary as shown in Figure 23-24:

```
                   Apply Journaled Changes Extend (APYJRNCHGX)

 Type choices, press Enter.

 Ending date and time:
   Ending date  . . . . . . . . . > '10/31/05'    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          *YES, *NO
 Option . . . . . . . . . . . . .   *NONE         *NONE, *IGNINQMSG
 Object error option  . . . . . .   *CONTINUE     *CONTINUE, *END
 Output . . . . . . . . . . . . .   *NONE         *NONE, *OUTFILE




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

Figure 23-24 APYJRNCHGX (Part 2 of 2)

12. In case you lost the journal object, you must find 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.
13. Sign off and sign on as *SID*OFR.
14. Delete SQL packages with the following command:

```
DLTR3PKG SID(SID) PKGTYPE(*ALL
```

15. Start your SAP system.

23.3 Using BRMS for save and restore

You can use BRMS to save and restore the entire SAP system 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.

Note: Use the *SID*OFR user profile to run the scheduled jobs.

Recovery policy considerations

All SAP database tables are journaled and all SAP transactions are running under commitment control. In our example, the SAP system R01 has both an ABAP and a Java database, R3R01DATA and SAPR01JRN respectively. We specified a shared journal environment, which means that the journal receivers for both databases are stored in the library R3R01JRN.

We recommend that you specify that saved journal receivers for journaled objects are automatically restored when the journaled objects are restored. The journal receivers are not restored if they already exist on the system.

The apply option specifies whether the journal changes are applied to the journaled objects, when the objects are restored. The i5/OS command APYJRNCHG is used to apply the changes, which is not recommended for SAP system environments.

Note: Specify the Apply option, ***NONE**, as shown in Figure 23-25 and set APYJRNCHGX manually if you need to apply the changes.

```
                          Change Recovery Policy                    ITSOSYS1

 Type choices, press Enter.

   Document name generation . . . . . *SAME         *SAME, *NEW
   System resource management . . . . *ALL          *ALL, *NONE, *HDW, *TRA
   Restore into folder  . . . . . . . *SAME

   Apply journaled changes:
     Restore journal receivers . . .: *YES          *YES, *NO
     Apply option  . . . . . . . . .: *NONE         *NONE, *APPLY, *REMOVE
       Ending date . . . . . . . . .: *CURRENT      Date, *CURRENT
       Ending time . . . . . . . . .: *CURRENT      Time, *CURRENT
   Lotus point-in-time:
     Ending date . . . . . . . . . .: *CURRENT      Date, *CURRENT, *FULL
     Ending time . . . . . . . . . .: *CURRENT      Time, *CURRENT, *FULL




 F3=Exit    F4=Prompt    F5=Refresh    F9=System policy
 F12=Cancel
```

Figure 23-25 Define recovery policy

Saving the SAP Integrated File System objects

Use the WRKLBRM command to add a link for saving SAP Integrated File System objects as shown in Figure 23-26.

```
                          Display Link List (DSPLNKLBRM)

 Type choices, press Enter.

 List . . . . . . . . . . . . . . > R3IFS        Character value
 Objects:
   Name . . . . . . . . . . . . . > '/usr/sap/<SID>/DVEBMGS01'
   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT

   Name . . . . . . . . . . . . . > '/sapmnt/<SID>'
   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT

   Name . . . . . . . . . . . . . > '/sapmnt/trans'
   Include or omit  . . . . . . .   *INCLUDE      *INCLUDE, *OMIT
 Directory subtree  . . . . . . .   *ALL          *ALL, *DIR, *NONE, *OBJ
 Text . . . . . . . . . . . . . . > 'SAP IFS Structure'




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

Figure 23-26 LInk definition: BRMS Integrated File System backup

Performing offline backup

To customize an offline backup, use the Work with Backup Control Groups, WRKCTLGBRM, command to create a new control group with option **1**. You can use the example shown in Figure 23-27 to add similar entries to your backup control group.

```
                      Edit Backup Control Group Entries               ITSOSYS1

Group . . . . . . . . . . : SAVSAPOFF
Default activity  . . . . . FFFFFFF
Text  . . . . . . . . . . . Offline Backup of SAP system

Type information, press Enter.

                          Auxiliary   Weekly    Retain  Save      SWA
      Backup      List  Storage     Activity  Object  While     Message
Seq   Items       Type  Pool Device MTWTFSS   Detail  Active    Queue

  10  R3IFS       *LNK  *SYSBAS     *DFTACT   *NO     *NO
  20  R3R01DATA         *SYSBAS     *DFTACT   *ERR    *NO
  30  SAPR01DB          *SYSBAS     *DFTACT   *ERR    *NO
  40  R3640UOPT         *SYSBAS     *DFTACT   *ERR    *NO
  50  R3R01400          *SYSBAS     *DFTACT   *ERR    *NO
  60  R3R01JRN          *SYSBAS     *DFTACT   *YES    *NO


                                                                        Bottom
F3=Exit    F5=Refresh    F10=Change item    F11=Display exits
F12=Cancel    F14=Display client omit status    F24=More keys
```

Figure 23-27 WRKCTLGBRM: BRMS offline backup

Note: Leave the default value *ERR for Retain Object Detail except for the journal library. This does not save object details for the backup items, but keeps the message identifier in case of an error.

Performing online backup with Save-While-active

Refer to Figure 23-28 for the link list definition. Use a similar logic for the Backup Control Group definitions, as shown in Figure 23-28 and Figure 23-31 on page 427.

.

```
                   Display Backup Control Group Entries              ITSOSYS1

 Group . . . . . . . . . . : SAVSAPON
 Default activity  . . . . : FFFFFFF
 Text  . . . . . . . . . . : Online Backup of SAP system

                             Auxiliary   Weekly    Retain  Save      SWA
       Backup      List  Storage     Activity  Object  While     Message
 Seq   Items       Type  Pool Device MTWTFSS   Detail  Active    Queue
   10  R3IFS       *LNK  *SYSBAS     *DFTACT   *NO     *YES      SAVE
   20  R3R01DATA         *SYSBAS     *DFTACT   *ERR    *SYNCLIB  SAVE
   30  SAPR01DB          *SYSBAS     *DFTACT   *ERR    *SYNCLIB  SAVE
   40  R3640UOPT         *SYSBAS     *DFTACT   *ERR    *NO
   50  R3R01400          *SYSBAS     *DFTACT   *ERR    *SYNCLIB  SAVE
   60  R3R01JRN          *SYSBAS     *DFTACT   *YES    *YES      SAVE




                                                                        Bottom
 Press Enter to continue.

 F3=Exit   F11=Display exits   F12=Cancel   F14=Display client omit status
```

Figure 23-28 WRKCTLGBRM: BRMS Save-While-Active

You should specify ***SYNCLIB** for the following:

- R3*SID*DATA database library (ABAP environment)
- SAP*SID*DB database library (Java environment)
- R3*SID*400 library because the memory-based database monitor is using files in that library.

All the other SAP 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 online backup for all the libraries. This means that the tape drive can be used for another save operation, for example, if the tape drive is shared between two systems.

Define a message queue to which save-while-active synchronization messages should be sent. In our example shown in Figure 23-28, we specified the message queue SAVE in the library QUSRBRM.

Leave defaults for backup attributes as shown in Figure 23-29:

```
                       Change Backup Control Group Attributes

Group  . . . . . . . . . . . . . . . . : SAVSAPON

Type information, press Enter.

IPL after backup . . . . . . . . . . . . *BKUPCY      *YES, *NO, *BKUPCY
   How to end  . . . . . . . . . . . . . *BKUPCY      *CNTRLD, *IMMED, *BKUPCY
   Delay time, if *CNTRLD  . . . . . . . *BKUPCY      Seconds, *NOLIMIT
   Restart after power down  . . . . . . *BKUPCY      *YES, *NO, *BKUPCY
   IPL source  . . . . . . . . . . . . . *BKUPCY      *PANEL, A, B, *BKUPCY
Save active wait time:
   Object locks  . . . . . . . . . . . . 120          0-99999, *NOMAX
   Pending record changes  . . . . . . . 120          0-99999, *NOCMTBDY, *NOMAX
   Other pending changes . . . . . . . . 120          0-99999, *NOMAX






                                                                       More...
F3=Exit   F12=Cancel
```

Figure 23-29 WRKCTLGBRM: option 8, change backup control attributes

Performing online backup with a Rapid Checkpoint Save-while-Active

Change the command default of the SAVACTWAIT parameter of the SAVLIB command in order to specify that you are running the Rapid Checkpoint Save-while-Active function as shown in Figure 23-30 on page 427.

```
                   Display Backup Control Group Entries            ITSOSYS1

Group . . . . . . . . . . : SAVSAPONR
Default activity  . . . . : FFFFFFF
Text  . . . . . . . . . . : Online Backup of SAP system

                        Auxiliary   Weekly     Retain  Save       SWA
      Backup      List  Storage     Activity   Object  While      Message
Seq   Items       Type  Pool Device MTWTFSS    Detail  Active     Queue
  10  R3IFS       *LNK  *SYSBAS     *DFTACT    *NO     *YES       SAVE
  20  R3R01DATA         *SYSBAS     *DFTACT    *ERR    *SYNCLIB   SAVE
  30  SAPR01DB          *SYSBAS     *DFTACT    *ERR    *SYNCLIB   SAVE
  40  R3640UOPT         *SYSBAS     *DFTACT    *ERR    *NO
  50  R3R01400          *SYSBAS     *DFTACT    *ERR    *SYNCLIB   SAVE
  60  R3R01JRN          *SYSBAS     *DFTACT    *YES    *YES       SAVE





                                                                        Bottom
Press Enter to continue.

F3=Exit   F11=Display exits   F12=Cancel   F14=Display client omit status
```

Figure 23-30   WRKCTLGBRM: BRMS Rapid Checkpoint Save-while-Active

Change defaults for Backup Control Group Attributes as shown in Figure 23-31.

```
                       Change Backup Control Group Attributes

Group  . . . . . . . . . . . . . . . : SAVSAPONR

Type information, press Enter.

IPL after backup . . . . . . . . . . . . *NO         *YES, *NO, *BKUPCY
   How to end  . . . . . . . . . . . . . *BKUPCY     *CNTRLD, *IMMED, *BKUPCY
   Delay time, if *CNTRLD  . . . . . . . *BKUPCY     Seconds, *NOLIMIT
   Restart after power down  . . . . . . *BKUPCY     *YES, *NO, *BKUPCY
   IPL source  . . . . . . . . . . . . . *BKUPCY     *PANEL, A, B, *BKUPCY
Save active wait time:
   Object locks  . . . . . . . . . . . . 120         0-99999, *NOMAX
   Pending record changes  . . . . . . . *NOCMTBDY   0-99999, *NOCMTBDY, *NOMAX
   Other pending changes . . . . . . . . 120         0-99999, *NOMAX







                                                                       More...
F3=Exit   F12=Cancel
```

Figure 23-31   WRKCTLGBRM: option 8 page 3 of 4 Change Backup Control group attributes

Restoring the SAP databases

This section shows how you can restore your SAP databases for ABAP and Java from your BRMS environment. In our example, we have a common journal environment for both databases. All journal receivers are stored in the library R3R01JRN. The BRMS command WRKMEDIBRM prompts all items being backed up by BRMS. To restore your SAP databases, perform the following steps:

1. Select option **7** to restore the SAP libraries as shown in Figure 23-32.

```
                       Work with Media Information                 ITSOSYS1

 Position to Date . . . . .

 Type options, press Enter.
   2=Change   4=Remove   5=Display   6=Work with media   7=Restore
   9=Work with saved objects

    Saved                            Save   Volume      File    Expiration
Opt Item             Date     Time   Type   Serial     Sequence    Date
    R3R01400       11/23/05 14:06:11 *FULL  E00011            4   12/28/05
    QUSRBRM        11/23/05 14:06:19 *QBRM  E00011            5   12/28/05
    R3R01JRN       11/23/05 14:08:42 *FULL  E00012            1   12/28/05
    QUSRBRM        11/23/05 14:09:00 *QBRM  E00012            2   12/28/05
 7  R3R01DATA      11/30/05 11:40:46 *FULL  E00007            1    1/04/06
 7  SAPR01DB       11/30/05 11:40:46 *FULL  E00007            2    1/04/06
    R3640UOPT      11/30/05 11:41:08 *FULL  E00007            3    1/04/06
    R3R01400       11/30/05 11:41:16 *FULL  E00007            4    1/04/06
 7  R3R01JRN       11/30/05 11:41:23 *FULL  E00007            5    1/04/06
    QUSRBRM        11/30/05 11:41:32 *QBRM  E00007            6    1/04/06
                                                                       Bottom

 F3=Exit   F5=Refresh   F11=Object detail   F12=Cancel
```

Figure 23-32 WRKMEDIBRM: select libraries

2. Press **Enter** to confirm the selection as Figure 23-33 shows.

```
                           Select Recovery Items                     ITSOSYS1

 Type options, press Enter.  Press F16 to select all.
   1=Select   4=Remove   5=Display   7=Specify object


     Saved         Save      Save     Save   Volume     File      Expire  Objects
Opt  Item          Date      Time     Type   Serial   Sequence    Date     Saved
 1   R3R01DATA   11/30/05 11:40:46  *FULL  E00007            1  1/04/06     21
 1   SAPR01DB    11/30/05 11:40:46  *FULL  E00007            2  1/04/06      1
 1   R3R01JRN    11/30/05 11:41:23  *FULL  E00007            5  1/04/06      2










                                                                          Bottom
 F3=Exit   F5=Refresh   F9=Recovery defaults   F11=ASP View   F12=Cancel
```

Figure 23-33 WRKMEDIBRM: confirm selected libraries

3. Press **F9** to validate command defaults as Figure 23-34 shows, as recommended in "Recovery policy considerations" on page 422.

```
                        Restore Command Defaults

Type information, press Enter.

 Document name generation  . . . . . *SAME       *SAME, *NEW
 System resource management  . . . . *ALL        *ALL, *NONE, *HDW, *TRA
 Apply journaled changes:
   Restore journal receivers . . . . *YES        *YES, *NO
   Apply option  . . . . . . . . . . *NONE       *NONE, *APPLY, *REMOVE
     Ending date . . . . . . . . . . *CURRENT    Date, *CURRENT
     Ending time . . . . . . . . . . *CURRENT    Time, *CURRENT
 Lotus point-in-time:
   Ending date . . . . . . . . . . . *CURRENT    Date, *CURRENT, *FULL
   Ending time . . . . . . . . . . . *CURRENT    Time, *CURRENT, *FULL






                                                                         Bottom
 F12=Cancel
```

Figure 23-34 WRKMEDIBRM: restore command defaults

Restoring the entire SAP control group

In order to restore all objects of the entire SAP control group, print the BRMS recovery plan. In our example, we print the recovery plan for the backup control group SAVSAPONR, which describes a Rapid Checkpoint Save-while-active scenario as Figure 23-35 shows.

```
                    Start Recovery using BRM (STRRCYBRM)

Type choices, press Enter.

Option . . . . . . . . . . . . . > *CTLGRP      *SYSTEM, *ALLDLO, *ALLUSR...
Action . . . . . . . . . . . . .   *REPORT      *REPORT, *RESTORE
Time period for recovery:
  Start time and date:
  Beginning time . . . . . . . .   *AVAIL       Time, *AVAIL
  Beginning date . . . . . . . .   *BEGIN       Date, *CURRENT, *BEGIN
  End time and date:
  Ending time  . . . . . . . . .   *AVAIL       Time, *AVAIL
  Ending date  . . . . . . . . .   *END         Date, *CURRENT, *END
Use save files . . . . . . . . .   *YES         *YES, *NO
Control group selection:
  Control group  . . . . . . . .   SAVSAPONR    Name, *SELECT, *NONE...
  Sequence number  . . . . . . .                1-99
              + for more values
Allow duplicate entries  . . . .   *NO          *NO, *YES
Use TSM  . . . . . . . . . . . .   *YES         *YES, *NO
                                                                  More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 23-35 STRRCYBRM: print recovery plan

Follow the instructions of the recovery plan to restore the environment. Figure 23-36 shows an example of the recovery plan (`step 003`), which describes how to recover the user libraries of the control group SAVSAPONR, as shown in Figure 23-30 on page 427.

```
5722BR1 V5R3M0 Recovery Report  SYSTEMXY  11/30/05  11:42:33    Page    1
.
.
STEP 003 : Recover User Libraries
Start date/time__________Stop date/time__________Duration_____________
   You should restore the current version of your user libraries.
   If you are performing a complete system restore, run the following
   command to continue.
      STRRCYBRM OPTION(*RESUME)
   Otherwise, run the following command.
      STRRCYBRM OPTION(*ALLUSR) ACTION(*RESTORE)
   Type your command choice then press "Enter".
.
.

Saved                Save     Save          Sequence  Control    Volume
Item       Type     Date     Time    Saved Number    Group      Identifier
---------- -----  -------- -------- ----- --------  ---------  ---------
R3R01DATA  *FULL  11/30/05 11:40:46   21     1      SAVSAPONR  E00007
SAPR01DB   *FULL  11/30/05 11:40:46    1     2      SAVSAPONR  E00007
R3640UOPT  *FULL  11/30/05 11:41:08    1     3      SAVSAPONR  E00007
R3R01400   *FULL  11/30/05 11:41:16    1     4      SAVSAPONR  E00007
R3R01JRN   *FULL  11/30/05 11:41:23    2     5      SAVSAPONR  E00007
```

Figure 23-36 SAP recovery plan, step 003

Note: To enable BRMS-integrated forward recovery of your journal receivers with the APYJRNCHGX command, perform the following:

- Install PTF SI21665 (V5R3) or SI21664 (V5R4).
- Create a data area as follows:

  `CRTDTAARA (DTAARA(QUSRBRM/Q1AAPYJRNX) TYPE(*CHAR) LEN(1)`

Manual forward recovery for BRMS saves

For manual forward recovery of your journal receivers, run the APYJRNCHGX command for all objects of the libraries R3R01DATA and SAPR01DB, and delete the SQL packages, as explained in “Forward recovery” on page 419.

24

Disaster recovery and high availability

For those customers that commit to the technology of SAP NetWeaver, high availability becomes an important and complex subject. It is important because the availability of the SAP system environment is essential to the day-to-day operations of the enterprise. It is complex because the client/server environment of an SAP system does not lend itself to the traditional solutions found in the centralized environment of the past. You also not only need to be concerned about the availability of one single SAP system, but typically about a set of closely coupled SAP systems, as described in Chapter 3, “SAP system landscapes” on page 17.

For an SAP system environment, the focus of a highly available solution must be at the network level, that is, a focus that supports minimal disruption to any active clients while a failing element of the SAP 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 whether the entire SAP network must come down while the failing component is replaced or whether the clients can remain operational. It is also this procedure that determines whether the backup environment is prepared properly for an SAP system to run or a manual intervention is required to do so.

24.1 Solution discussion

There are three critical components of an SAP system that can affect the availability of an SAP system environment. They are:

- The SAP database
- The central instance
- The critical Integrated File System information

The failure of any one of these components can impact the entire SAP system 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 components, we advise that you keep them on the same System i server. The database and critical Integrated File System information are automatically together. There is, however, no restriction that limits the location of the central instance. It can be located on a server separate from the one that holds the database. If this is done, an SAP implementation for high availability becomes more complex.

There are now two points of failure within the SAP network: the machine that holds the database and the system that is running the central instance. Provisions must be made to recover from failure in either of these two nodes (database or central instance) if SAP high availability is to be maintained.

Note: To optimize the availability of your SAP system environment, install a disaster recovery and a high availability concept.

24.1.1 Disaster recovery

Disaster recovery means that plans are in place recover your SAP system environment off-site from a disaster on a separate server. However, the business operations might tolerate IT-services for some amount of time while disaster recover is performed. In addition, disaster recovery can involve some amount of manual processing since it is assumed to be rarely needed.

Typical disaster recovery environments are:

- To reload the SAP system environment on a separate server
- Implement remote journaling and a manual apply of journal receivers
- Mirrored Sysbas environments

24.1.2 High availability

The requirement for this category is to ensure that critical system resources for your SAP system environment are continuously available during both planned and unplanned outages. Failover processing to a backup server must be fully automated and complete in a short time frame. Normal production is expected to continue with acceptable degradation.

Typical high availability environments are:

- Implementations based on Switchable Independent Auxiliary Storage Pool (IASP)
  - Switchable device environments
  - Cross Site Mirroring (XSM)
  - IBM Storage server and System i Copy Services

- Logical replication environments

24.2 Concepts and basic technologies

This section discusses some components of a high availability or disaster recovery solution, including LPAR, a backup system, clustering and IASP technology.

24.2.1 Logical partitioning (LPAR)

Logical partitioning is the ability to make a single multiprocessing System i server run as though it were two or more independent systems (introduced in OS/400 V4R4). Each logical partition represents a division of resources on a System i 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, IOPs and IOAs. 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 can 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 system landscapes” on page 17, and Chapter 5, “System i virtualization” on page 31.

24.2.2 Backup system

The SAP production system can be replicated to another machine for high availability or disaster recovery reasons. Since the “replicated” machine does not need the same resources as the production machine all of the time (memory, disks, CPUs, and so on), it can be configured into two LPARs, with another LPAR used for testing, for example.

Figure 24-1 on page 436 shows a typical scenario: an SAP production system running on a dedicated server with a replication of the production system on a backup server. This backup server is also used for development and quality assurance. Both environments are running in separate LPARs.

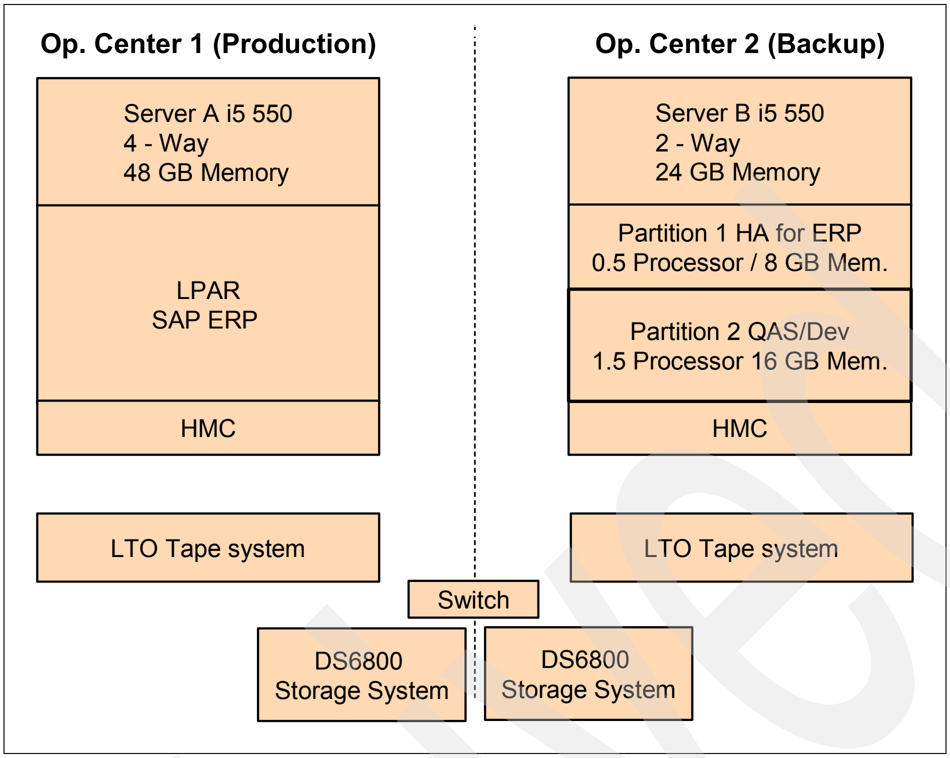

Figure 24-1 SAP production system with additional server

24.2.3 Clustering

Clustering builds the foundation of System i high availability concepts. A cluster is a collection of interconnected complete systems used as a single, unified computing resource. The integrated cluster resource services enable you to define a cluster of systems and the set of resources that should be protected against outages. Cluster resource services detect outage conditions and coordinate automatic movement of critical resources to a backup system. The systems in a cluster can be physically connected via a LAN or high-speed OptiConnect bus.

In the event of a failure, Cluster Resource Services (CRS), which is running on all systems, provides a switch over. This switch causes minimal impact to the end user or applications that are running on a server system. CRS needs to wait for several minutes. In this time frame the SAP database is activated on the backup server.

This can require a database recovery in the event of an unplanned downtime. After that, the SAP system needs to be restarted on the backup server. Finally, the end user can logon again to the system. Data requests are automatically rerouted to the new primary system. You can easily maintain multiple data replicates of the same data.

Clusters can 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.

Figure 24-2 on page 437 shows how SAP applications can be implemented in a System i cluster environment.

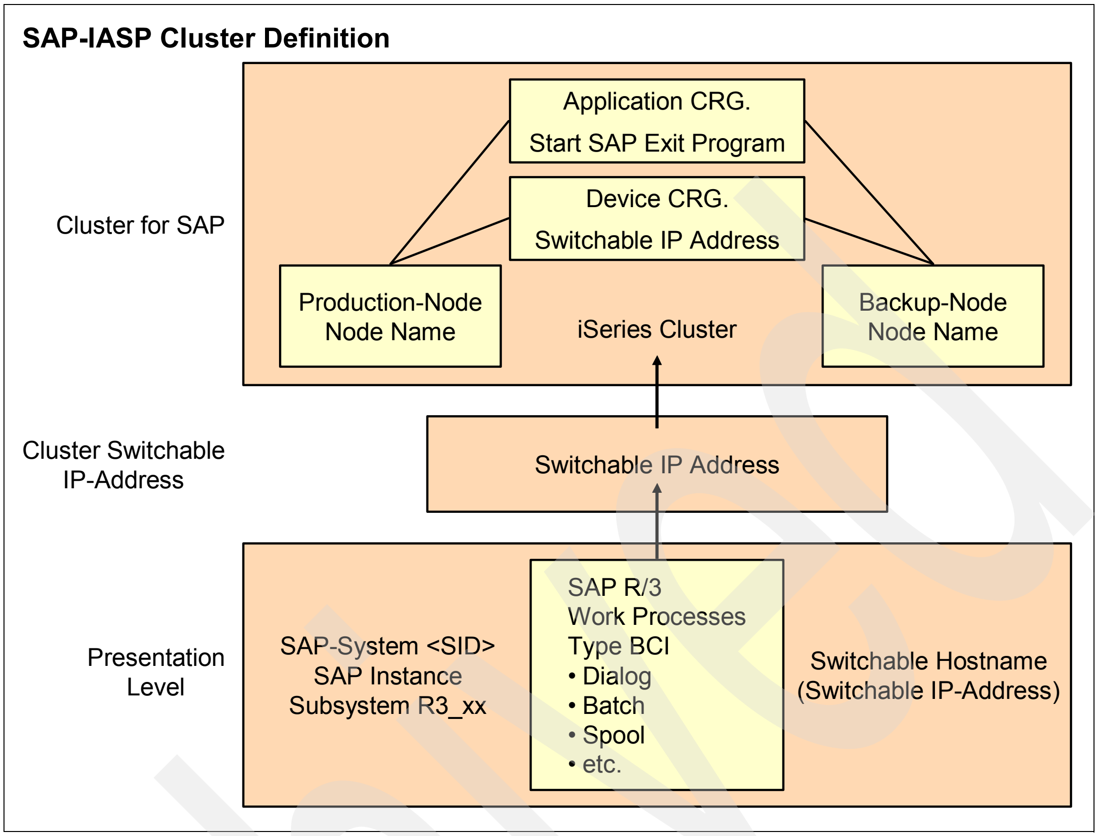

Figure 24-2 SAP applications running in a System i cluster

24.2.4 Independent Auxiliary Storage Pool

Auxiliary storage pools (ASP'S) have been part of the System i architecture since the announcement of the AS/400 in 1988, and of the System/38™ before it. System ASP's and user ASP's enable a division of the total disk storage on the system into logical groups. One or more applications or data objects can be isolated onto one or more ASP's to provide support for improving backup and recovery, performance, and more.

IASPs can be used on a single system or switched between multiple systems or LPAR partitions. When used on a single system (as a standalone IASP), the IASP can be dynamically turned on or off. Figure 24-3 on page 438 shows example for standalone or switchable IASP's being assigned to System i servers.

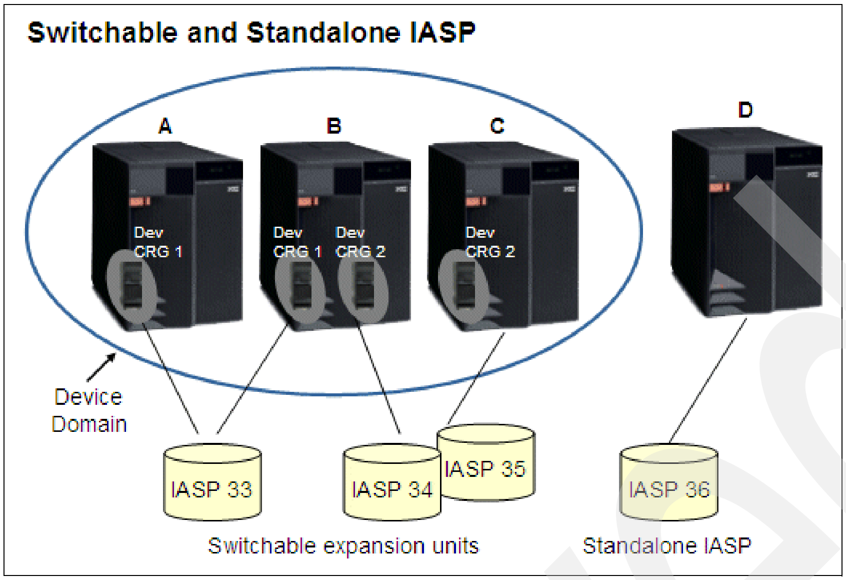

Figure 24-3 Switchable and standalone IASPs

When used across multiple systems, the IASP, and therefore its contents (application and data) can be switched between those systems. The ability to be turned off or on, or to be switched between systems is what differentiates IASPs from regular ASP's and provides IASPs the independent characteristics. The characteristic of an IASP to be varied on or off, or attached and detached, can be done by performing a system IPL.

Clustering support is required to switch IASPs across multiple systems. Cluster management is used to switch the IASP across systems in a cluster. At any one time, the IASP can be used from one of the clustered systems. The IASP cannot be used simultaneously from multiple systems.

For more information about auxiliary storage pools, refer to:

- 8.1.5, "Auxiliary Storage Pool" on page 53
- 9.7, "ASP considerations for SAP journal management" on page 110
- *Clustering and IASP's for Higher Availability on the IBM @server iSeries Server,* SG24-5194
- *IBM @server iSeries Independent ASP's: A Guide to Moving Applications to IASP's,* SG24-6802
- *Independent ASP Performance Study on the IBM @server iSeries Server,* REDP-3771

24.2.5 Mirrored system environments

You can run a disaster recovery solution by permanently duplicating the whole disk environment of a System i server onto another separate disk environment, which also can be assigned to a separate backup system. This can be done by:

- Mirrored protection at the bus level (an i5/OS integrated function), or
- Metro Mirror copy services (for IBM TotalStorage servers only)

In the case of a failover you need to:

1. Assign the disk environment to the backup server
2. Recover the internal Load Source Unit (not necessary for SAN bootable IOPs)
3. Run an i5/OS IPL

This implementation does not affect your installation of an SAP system environment. However, this solution requires manual intervention in case of switchover to the backup server. The recovery time depends on:

- The time needed for assigning the disk environment to the backup server
- The load source recovery process (not needed for bootable IOPs)
- The system IPL
- The database recovery time in case of an unplanned outage.

Figure 24-4 shows a typical landscape for the following kind of mirrored system environments:

- Mirrored internal disks:
  - Improved performance because of Raid-1 disk mirroring
  - Limited distance between production and backup environment (HSL cable length)
  - Load source recovery required in case of switchover

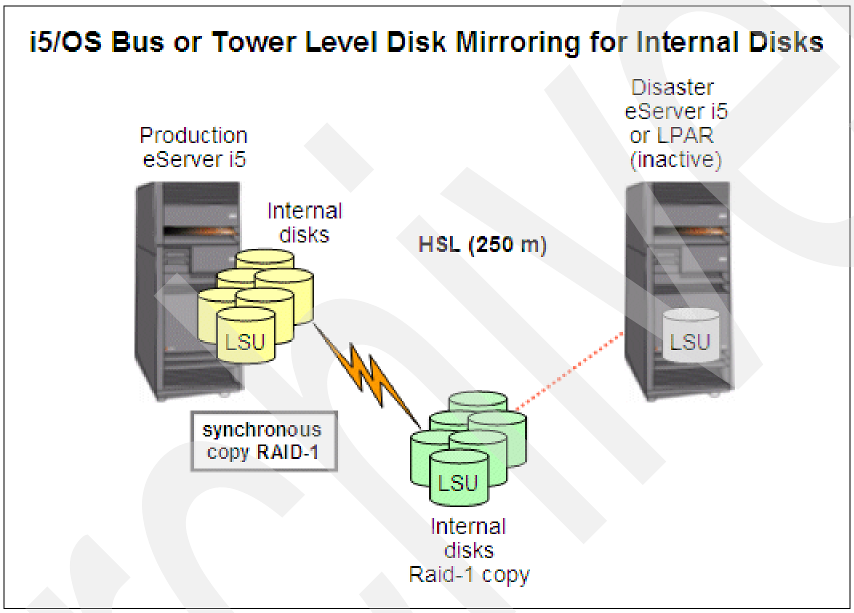

Figure 24-4 Mirrored internal disks (RAID-1)

- Mirrored external disks, RAID-1 (Figure 24-5 on page 440):
  - Improved performance because of RAID-1 disk mirroring
  - Extended distance because of Fibre Channel connection (up to 10 Km)
  - No load source recovery for SAN bootable IOPs required (as of System i servers with i5/OS V5R3M05 or later)

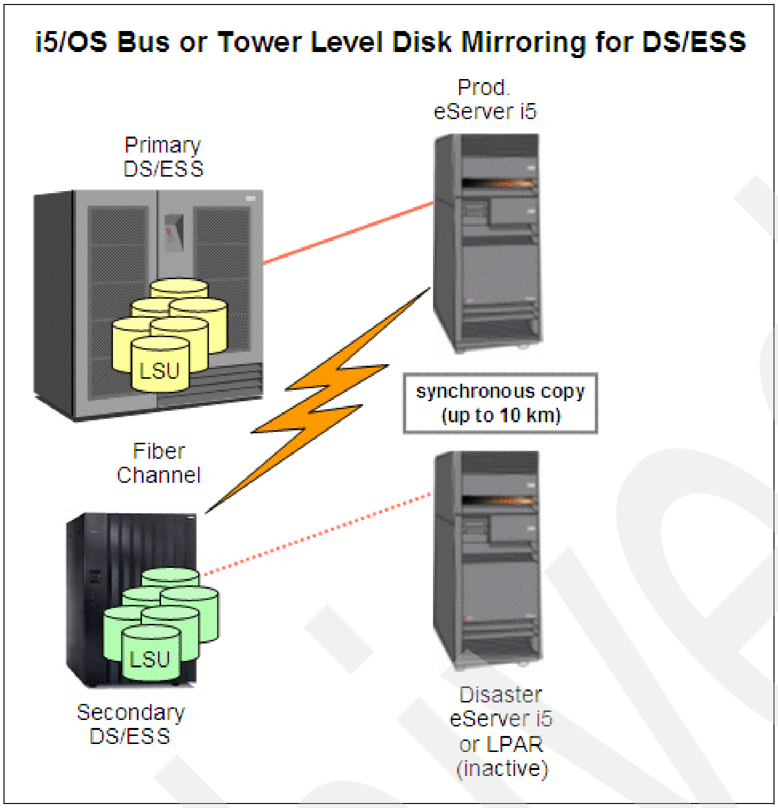

Figure 24-5 Mirrored external disks (RAID-1)

- Mirrored external disks, Metro Mirror (Figure 24-6 on page 441):
  - Synchronous or asynchronous connection (unlimited distance for asynchronous connection)
  - No load source recovery for SAN bootable IOPs required (as of System i servers with i5/OS V5R3M05 or later)

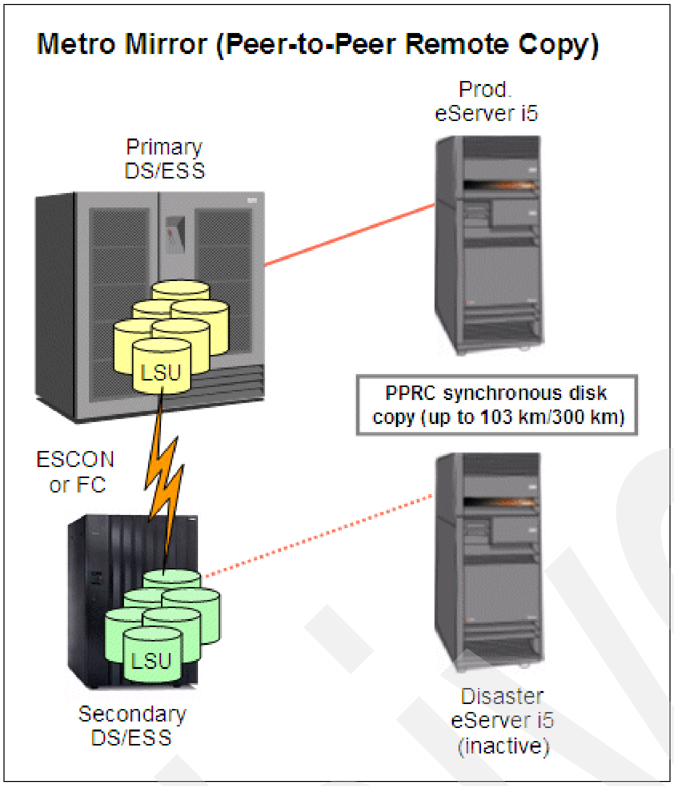

Figure 24-6 Mirrored external disks (Metro Mirror)

24.2.6 Remote journaling

The high availability solutions developed prior to OS/400 V4R2 used local journaling and the Receive Journal Entry (RCVJRNE) command, as shown in Figure 24-7 on page 442.

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 links. 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.

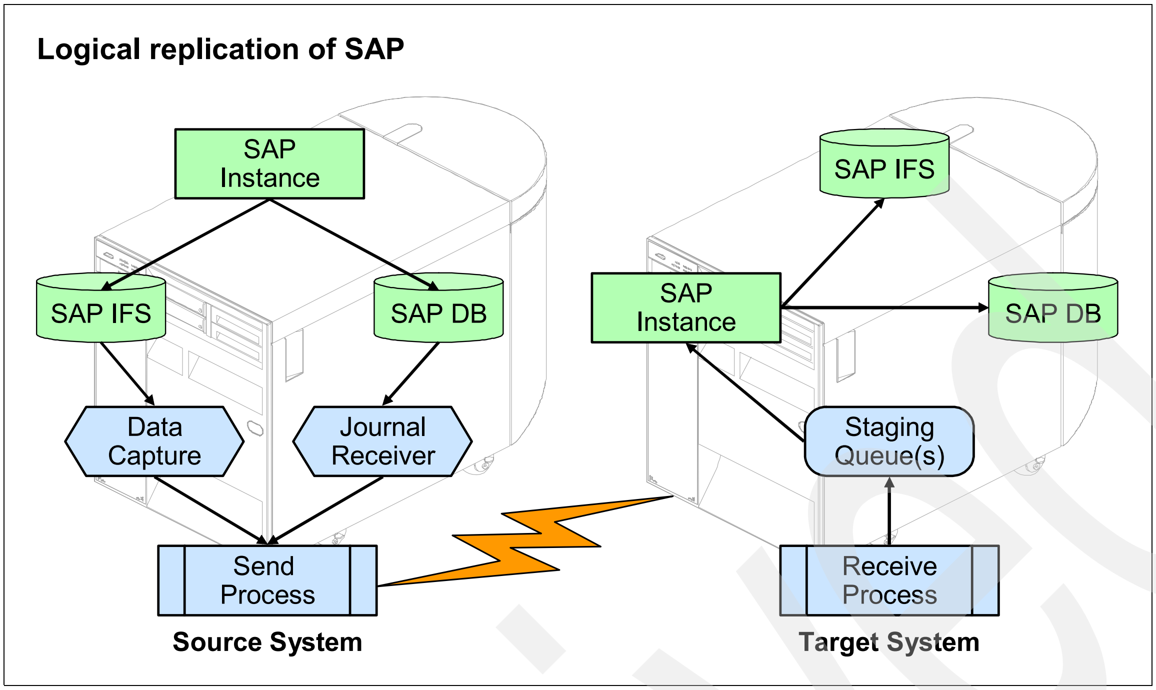

Figure 24-7 Logical replication of SAP

The remote journal function as described in Figure 24-8 provides a much more efficient transport of journal entries than the traditional approach.

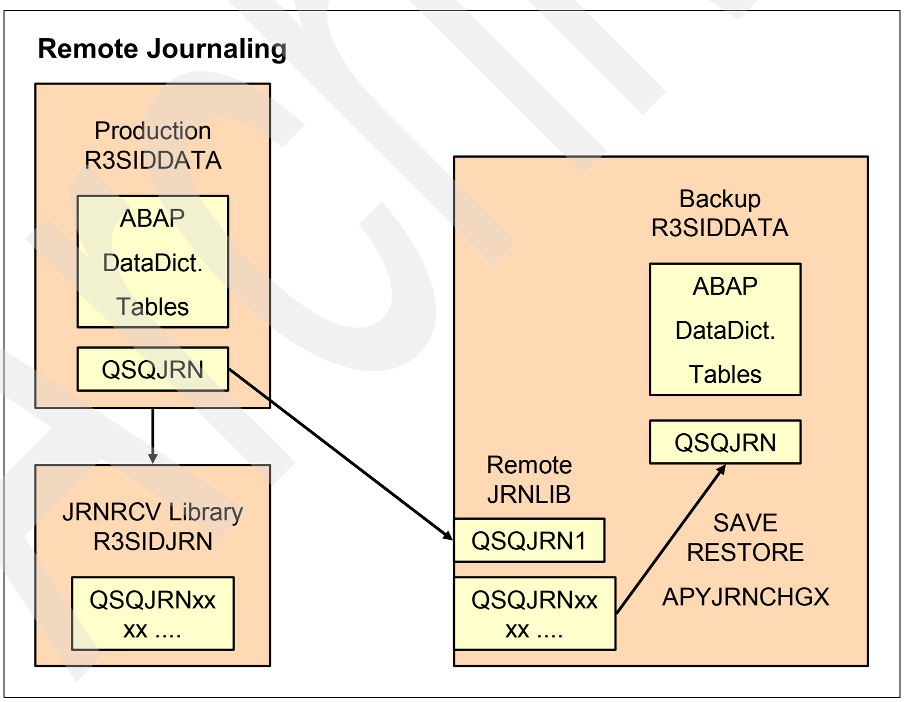

Figure 24-8 Remote journal environment for an SAP system

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.

In a second step you need to connect the journal receivers to the journal of the backup database and apply the changes to synchronize the backup database with the original one.

The high availability solutions available on System i models 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 24-9 shows how a logical database solution can benefit from the i5/OS remote journaling concept.

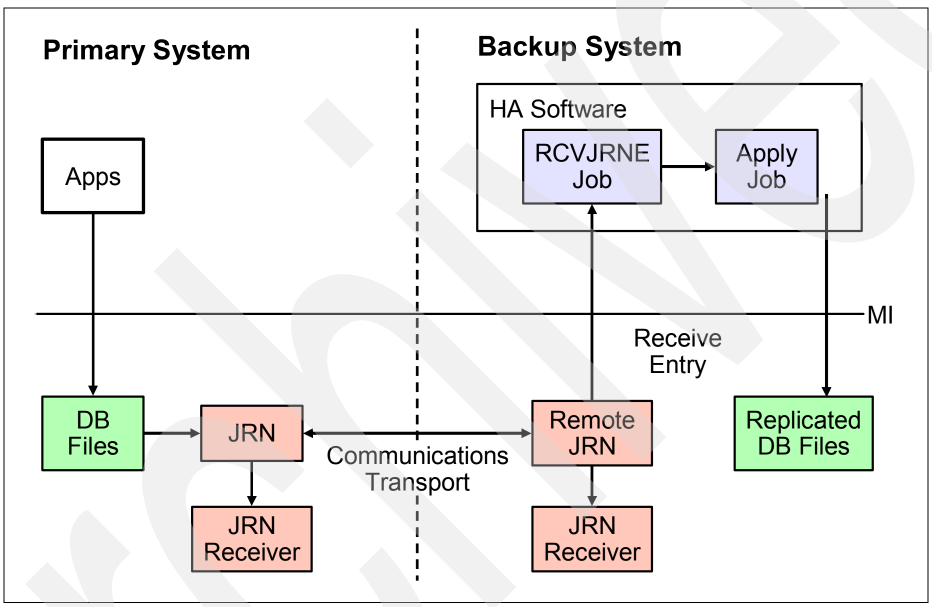

Figure 24-9 With remote journaling

You can consider the journal receivers on the target machine as a replica of the production machine's journal receivers. It is similar to saving the the journal receivers of the production machine and restoring 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 can see a minimal discrepancy in size between the local receiver and the associated remote receiver. Note that you are using two different machines that allocate space for the journal receivers in different operating system environments. This can result in slightly different sizes, but the data in the associated receivers is always the same.

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 is activated.

Synchronous or asynchronous transmission mode

The different delivery modes are described as follows.

- 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 does not lose 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 can have journal entries ahead of those journal entries that are known to the target system. Using this method allows for recovery that can lose a number of journal entries in the event of a failure on the source system. It should have minimal impact to the local system when compared to the synchronous delivery mode.

> **Note:** It is recommended to always use asynchronous mode. Otherwise your SAP system stops if the target system is not available.

Implementing remote journaling

To implement remote journaling for an SAP system, perform the following steps:

1. Set up the communication paths that are used to replicate journal entries between the source and the target systems.
2. Create a Relational Database (RDB) entry that identifies the target system and communication as shown in Figure 24-10 on page 445:

```
                    Display Relational Database Entry Detail

 Relational database  . . . . . . :   ERPSF2
 Remote location:
   Remote location  . . . . . . . :   ERPSF2
     Type . . . . . . . . . . . . :   *IP
   Port number or service name  . :   *DRDA
   Remote authentication method:
     Preferred method . . . . . . :   *ENCRYPTED
     Allow lower authentication . :   *ALWLOWER
 Text . . . . . . . . . . . . . . :


 Relational database type . . . . :   *REMOTE



                                                                         Bottom
 Press Enter to continue.

 F3=Exit   F12=Cancel
 (C) COPYRIGHT IBM CORP. 1980, 2003.
```

Figure 24-10 Relational database entry for remote journaling

3. Create a library for the remote journal (in our example R3R01RJRN) and journal receivers (in our example R3R01JRN) on the backup system:

```
CRTLIB LIB(R3R01DATA1)
CRTLIB LIB(R3R01JRN1)
```

4. Remote journaling is running on DDM based connections without password checking. Therefore, switch off password checking for DDM connections on the target server and start the DDM service again:

```
ENDTCPSRV SERVER(*DDM)
CHGDDMTCPA PWRQD(*NO)
STRTCPSVR SERVER(*DDM)
```

   If you do not want to switch off password check for DDM, specify the user profile for which the server authentication is to be added:

```
ADDSVRAUTE USRPRF(<REMOTEJRN>) SERVER(QDDMSERVER) PASSWORD(<password>)
```

5. Add the remote journal related to the existing SAP journal environment (QSQJRN in SAP database library), as shown in Figure 24-11 on page 446:

```
                         Add Remote Journal (ADDRMTJRN)

Type choices, press Enter.

Relational database  . . . . . . > ERPSF2
Source journal . . . . . . . . . > QSQJRN        Name
  Library  . . . . . . . . . . . >   R3RO1DATA   Name, *LIBL, *CURLIB
Target journal . . . . . . . . . > QSQJRN1       Name, *SRCJRN
  Library  . . . . . . . . . . . >   R3RO1DATA1  Name
Remote receiver library  . . . .   *SRCRCVLIB    Name, *SRCRCVLIB
Remote journal type  . . . . . .   *TYPE1        *TYPE1, *TYPE2
Journal message queue  . . . . .   QSYSOPR       Name
  Library  . . . . . . . . . . .     QSYS        Name
Delete receivers . . . . . . . .   *NO           *NO, *YES
Delete receiver delay time . . .   10            1-1440
Text 'description' . . . . . . .   *BLANK




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

Figure 24-11 ADDRMTJRN: add remote journal for an SAP system environment

6. Activate the remote journal, as shown in Figure 24-12.

```
                      Work with Remote Journal Information

Journal  . . . . . . :   QSQJRN         Library  . . . . . . :   R3RO1DATA

Journal type . . . . :   *LOCAL         Journal state  . . . :   *ACTIVE
Remote journal type  :                  Delivery mode  . . . :
Local journal  . . . :                  Source journal . . . :
  Library  . . . . . :                    Library  . . . . . :
  ASP group  . . . . :                    ASP group  . . . . :
  System . . . . . . :                    System . . . . . . :
Redirected receiver library . . . . . . . . . . . . . . . . :   *NONE
Number of remote journals . . . . . . . . . . . . . . . . . :   1

Type options, press Enter.
  13=Activate   14=Inactivate ...
                         --------------------Remote--------------------
     Relational                                 Journal     Delivery
Opt  Database            Journal     Library     State       Mode
13   ERPSF2              QSQJRN1     R3RO1RJRN1  *INACTIVE
                                                                     Bottom
===>
F3=Exit       F4=Prompt    F5=Refresh   F6=Work with remote journal list
F9=Retrieve   F12=Cancel   F23=More options
```

Figure 24-12 WRKJRNA: activate remote journal

7. Use the CHGRMTJRN command to change the delivery status or the attached remote journal receivers. Asynchronous mode recommended. See Figure 24-13 on page 447.

```
                        Change Remote Journal (CHGRMTJRN)

 Type choices, press Enter.

 Relational database  . . . . . . > ERPSF2        Name
 Source journal . . . . . . . . . > QSQJRN        Name
   Library  . . . . . . . . . . . >   R3RO1DATA   Name, *LIBL, *CURLIB
 Target journal . . . . . . . . . > QSQJRN1       Name, *SRCJRN
   Library  . . . . . . . . . . . >   R3RO1DATA1  Name
 Journal state  . . . . . . . . . > *ACTIVE       *SAME, *ACTIVE, *INACTIVE
 Delivery . . . . . . . . . . . .   *ASYNC        *ASYNC, *SYNC
 Starting journal receiver  . . .   *ATTACHED     Name, *ATTACHED, *SRCSYS
   Library  . . . . . . . . . . .                 Name, *LIBL, *CURLIB







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

Figure 24-13 CHGRMTJRN

> **Note:** Consider to implement remote journaling, if you are running a logical data replication software from a partner. Without a logical data replication software solution, you need to synchronize the SAP database and Integrated File System file objects manually.

For more information, refer to *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189.

24.3 SAP system implementation with a switchable IASP

This section describes how to implement an SAP system in a switchable IASP environment. Independent ASPs are discussed in 8.1.6, "Independent ASP" on page 53.

If you do not configure an Independent Auxiliary Storage Pool, you normally install an SAP system as shown in Figure 24-14 on page 448. In this example all data for an SAP system is placed in the system ASP (Sysbas). To optimize data integrity, consider installing the journal receivers in a separate User ASP. Since all objects and data for SAP applications are installed in the system ASP or an associated user ASP, the data of the SAP system cannot be switched over or assigned to another backup server without rebooting the environment.

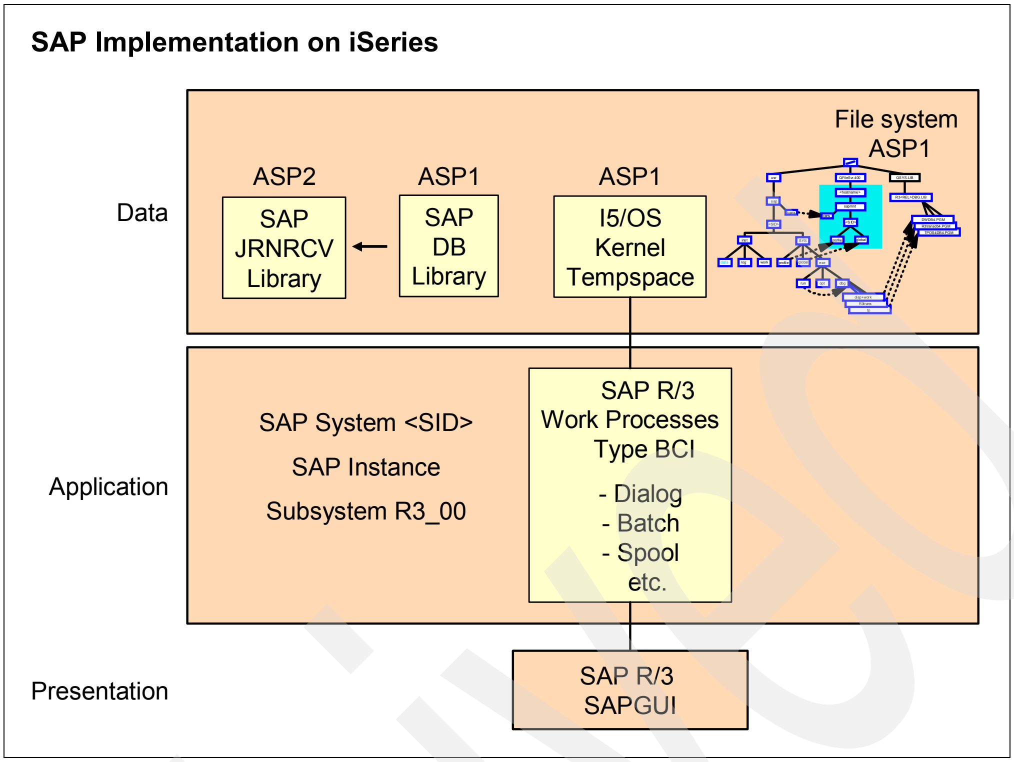

Figure 24-14 SAP system environment running in the system ASP

Configuring an SAP system (ABAP or J2EE) to run on a switchable IASP allows an SAP system to be switched from a primary host to a secondary one with minimal effort and without rebooting the secondary host. The switchover process is transparent to the end user and does not require any adjustments. After a given downtime for reconnection of the IASP and the database, the end user can log on again to the SAP system and resume working.

For downtime you have to consider the normal switchover time for an IASP from the primary server to a backup server, and in the event of an abnormal end of the SAP system, also consider the database recovery time which is needed to synchronize all open transactions.

As shown in Figure 24-15 on page 449, the SAP database, the SAP kernel and the /sapmnt directories are placed in the switchable IASP. The /sapmnt directories are linked to the static part of the SAP directory structures, which needed to be installed twice on the primary and the secondary system.

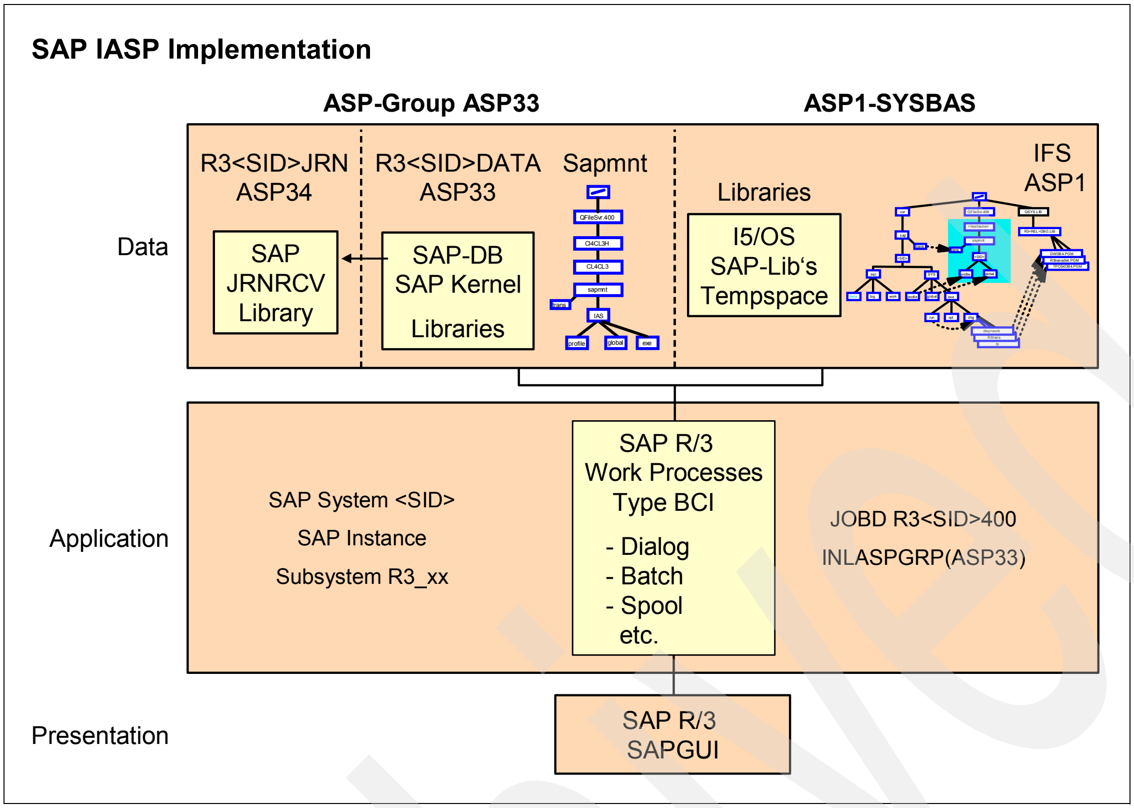

Figure 24-15 SAP implementation with an Independent Auxiliary Storage Pool

The associated journal receivers need to be located there or in a separate IASP which needs to be part of the IASP group.

Note: You also can install multiple SAP systems in one IASP.

The following description roughly outline how to configure an SAP system (ABAP or J2EE) in an IASP to become switchable from a primary production host to a secondary backup server with minimal effort and downtime. For a detailed description about how to configure an SAP system with a switchable IASP environment, refer to SAP note 568820.

The IASP configuration for an SAP system includes moving the SAP database library, Integrated File System directories and files, and SAP kernel library (6.40 and later only) to the IASP. The libraries R3<SID>400, R3WRK<nn>, R3400, and R3SYS remain on the system ASP, which means that they need to be available on both the primary and backup system. To enable switching from one host to another, configure the SAP system to run on a virtual host, which is a static IP address and host name that can be switched between the primary and the secondary host.

The guidelines presume that the SAP system is already installed and running in a system ASP environment, as shown in Figure 24-14 on page 448 and needs to be reconfigured into a switchable IASP environment.

To change the SAP system into a switchable IASP perform the following steps:

1. Check prerequisites, which include:
   a. OS/400 V5R2 or later
   b. OS/400 5722-SS1 Option 41, High Availability Switchable Resources

   c. SAP release 46D or SAP release 6.40 or later
   d. Switchable IASP and clustered environment configured for primary and secondary server
   e. SAP system running on primary server in system ASP
2. Preparation phase
   a. Install the latest operating system PTFs, as specified by the Information APAR
   b. Make sure the SAP system is at the latest patch level, especially the following components:
      i. DISP+WORK (DW)
      ii. LIB_DBSL
      iii. APYR3KRN
      iv. R3INLPGM
      v. TP
      vi. CMDMAINI
      vii. APYJ2EELIB (J2EE only)
   c. Make sure that the IASP is configured and available
   d. Check that a relational database entry is created for the IASP both on the primary and secondary system. Add an RDB entry if necessary:
      ADDRDBDIRE RDB(<iasp-name>) RMTLOCNAME(*LOOPBACK)
   e. Copy the kernel library to the system ASP of the secondary host
   f. For 6.40 and later, copy the kernel library under a new name to the IASP.
   g. Copy the J2EE library under a new name to the IASP
3. Setup phase
   a. Move the database, Integrated File System directories and files to the IASP and modify job descriptions, user profiles and data areas to utilize the IASP. Use the SAP command CFGSAPIASP.
   b. For 6.40 and later, apply the kernel that resides in the IASP of the SAP system (APYR3KRN command).
   c. Apply J2EE library that resides in IASP of the SAP system (APYJ2ELIB command, for WAS 6.40 only).
   d. Configure default, startup and instance profiles.
   e. For J2EE, configure connection and instance properties.
   f. For J2EE, check the JDBC driver symbolic links found at:
      /usr/sap/<SID>/SYS/jdbc/tbx.

      Ensure that these links still point to a valid driver. Modify if necessary.
   g. Test and verify your SAP system

For a more detailed and updated description refer to SAP note 568820. And download the IASP toolset (an i5/OS savefile) which contains tools to configure an SAP system running in an IASP.

This implementations allows you to run an SAP system in a switchable IASP environment. Based on this IASP implementation you can install a high availability solution to improve your business continuity by running the IASP for:

- Switchable device environment
- Cross-site mirroring environment
- System i copy services

You can find additional information in *8.1.6, "Independent ASP" on page 53* and *IBM @server IBM @server iSeries Independent ASPs: A Guide to Moving Applications to IASPs*, SG24-6802

The following sections shortly outline these environments. For more information refer to *Data Resilience Solutions for IBM i5/OS High Availability Clusters*, REDP-0888.

24.3.1 Switchable device environment

In a switchable device environment the IASP is connected via HSL cables between both the primary and the secondary server. This means that there is a limitation for the distance between both servers depending on the HSL cable length (6 meter for copper cable, 250 m for fibre cable). It also means that only a single copy of the data is maintained in the IASP which might become a single point of failure. Data resilience can be improved by running the IASP under RAID-1 with bus level protection, as shown in Figure 24-16.

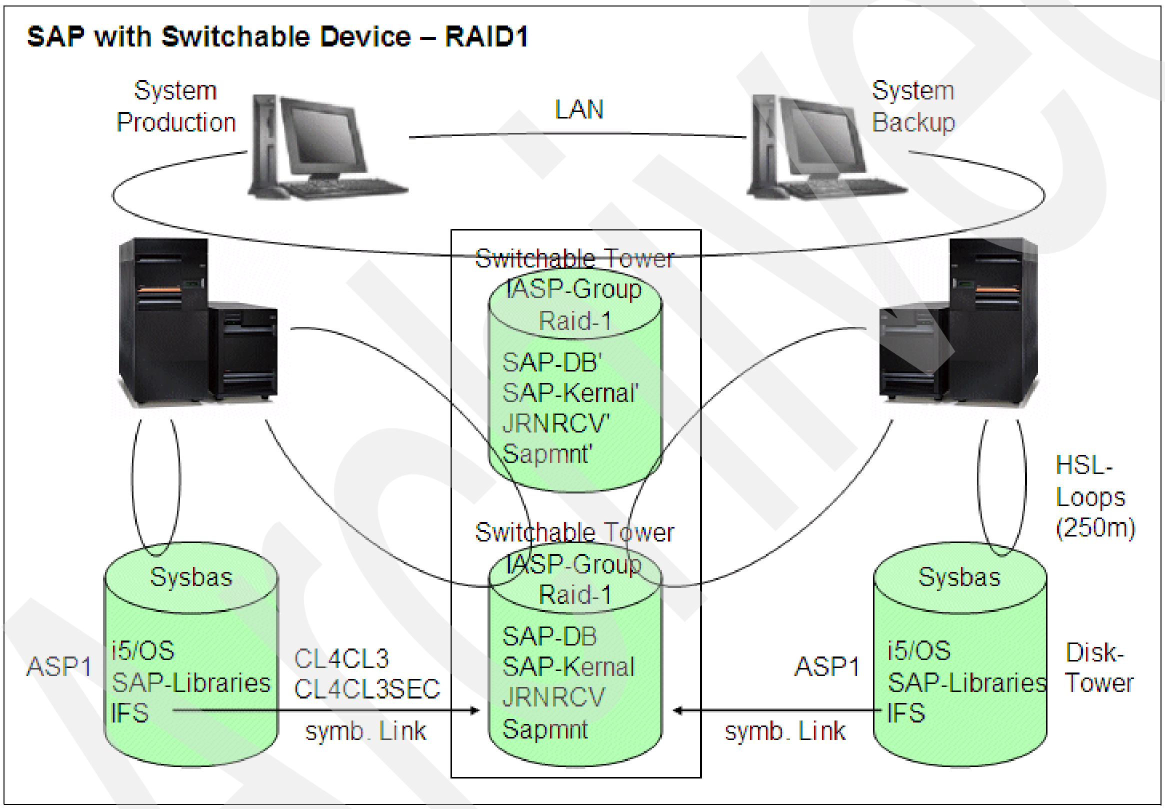

Figure 24-16 SAP system running in a switchable device environment

This technology does not involve replication and there are no additional copies of data, which causes only a minimal performance overhead. Another key benefit of this solution is zero-transmission latency. The major effort associated with this solution involves setting up the direct access storage device (DASD) configuration, to make the DASD switchable from one the primary to the secondary system in a transparent manner.

24.3.2 Cross site mirroring

Cross site mirroring technology (XSM) is based on the mirrored copy of an IASP which means that you have your IASP twice: one copy on the production system which is available for end user updates, and the other copy on the backup system which is not available at that time. XSM addresses the single point of failure issue of the basic switchable device structure. In addition, it allows a remote mirrored copy of IASP data. The benefits of this solution are essentially the same as the basic switchable device solution with the added advantage of providing disaster recovery to a second copy at increased distance.

Whenever you are performing a write or update into your IASP, the data is immediately sent to the backup machine and applied there. Be aware of the increased central processing unit overhead and I/O performance impacts in this environment.

Be aware that a full synchronization between the original and the mirrored IASP is required for any persistent transmission interruption, such as the loss of all communication paths between the source and the target system, or a planned detach of the mirrored copy to run a backup from there. During the synchronization you cannot switchover from your primary to the secondary system which means that during this time frame no backup system is available. Use redundant transmission links (multiple gigabit Ethernet connections recommended) to minimize the risk of loosing communication between both systems.

Figure 24-17 shows how an SAP system is running in a XSM environment.

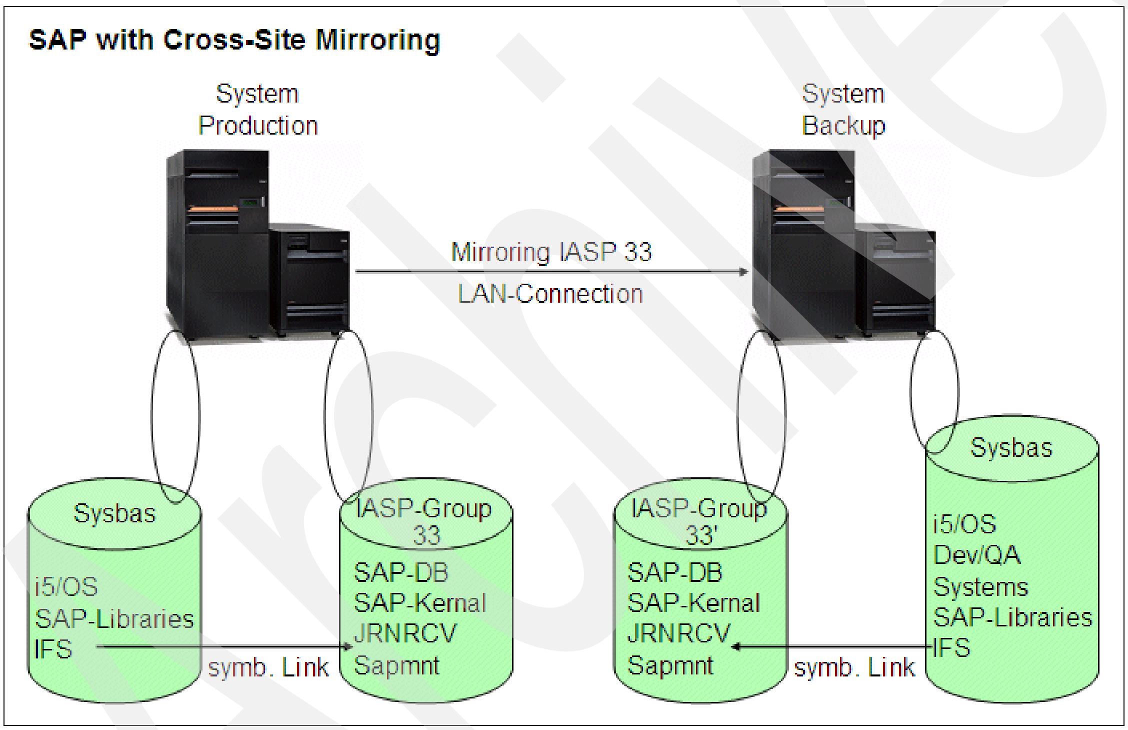

Figure 24-17 SAP applications running in cross-site mirroring environment

For more information about XSM performance behavior, refer to the IBM white paper “Geographic Mirroring in the SAP System Environment on @server iSeries”: at:
http://www.ibm.com/servers/eserver/iseries/perfmgmt/pdf/mirrsap.pdf

24.3.3 System i copy services

This solution involves the replication of data at the storage controller level to a second storage server using the Metro Mirror product (formerly known as PPRC) and being automated by System i copy services.

An IASP is the basic unit of storage for the storage server Metro Mirror function. Metro Mirror generates a second copy of the IASP on another storage server. The System i copy services offering provides a set of functions to combine the Metro Mirror, IASP and i5/OS cluster services for coordinated switchover and failover processing through a cluster resource group. Figure 24-18 shows how an SAP system can be implemented in such an environment.

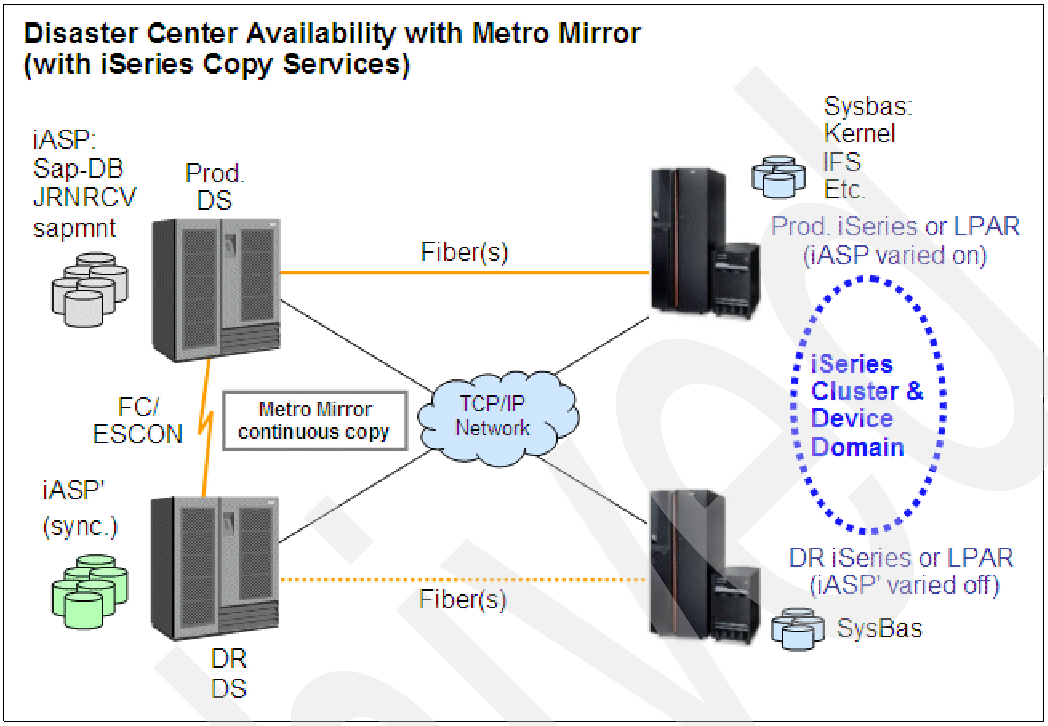

Figure 24-18 DS storage server with Metro Mirror and System i copy services

This solution provides the benefit of the remote copy function and coordinated switching operations which gives you good data resiliency capability if the replication is done synchronously. The copy services enable you automatically to attach the second copy to a backup server without an IPL or manual interventions.

24.4 Logical replication

High availability solutions based on logical database replication tools which means that a second copy of data is generated that is logical identical to the first. Replication is executed on object level (Integrated File System files for SAP applications) or record level (based on SAP database journal entries).

To build an efficient and reliable HA solution using logical replication, synchronous remote journaling as a transport mechanism is preferable. With remote journaling, IBM i5/OS continuously moves the newly arriving data in the journal receiver to the backup server journal receiver. At this point, a software solution is employed to ’replay’ these journal updates, placing them into the object on the backup server. For the database changes this is done on a commit level, which means that only all changes of a complete commit cycle are applied.

After this environment is established, there are two separate identical databases, one on the primary server and one on the backup server. With this solution in place, you can rapidly activate your production environment on the backup server without any delay because of an IPL or database recovery times (in the event of an unplanned outage). Figure 24-7 on page 442 illustrates the basic mechanism about how to run an SAP system in a logical replication environment.

24.4.1 Business Partner solutions

High availability middleware is the name given to the group of applications that provide the replication and management between System i servers. More recently they also provide the Cluster management middleware.

The high availability solutions for the system solutions are typically deployed via an High Availability Business Partner (HABP) solution package. This section contains information about high availability solutions for the System i SAP customers from business partners.

See the following HA Web site for a list of HA business partners:

http://www.ibm.com/servers/eserver/iseries/ha/

Part 4

Optimizing performance

You need a well balanced system in all components that affect performance to run applications like mySAP Business Suite solutions successfully. For System i installations, this means a balance of CPU utilization, memory paging, disk utilization, and load in the network. An optimized SAP system installation is targeted to need less resources.

A system administrative task is to diagnose problems and analyze poor performance. Involve the support organizations of SAP or IBM if you are unable to determine the problem yourself.

Part 4 discusses the tasks which an administrator can execute to analyze and to solve performance problems. The subjects include:

- System i work management
- SAP memory management
- System i and SAP monitoring
- System i and SAP tuning
- Database monitoring and tuning
- Helpful tools
- Other performance related aspects

Note: The chapters in Part 4 refer to i5/OS V5R3 operating system and SAP kernel Release 6.40 unless otherwise stated. Most statements also apply to i5/OS V5R4.

25

Performance concepts and monitoring

This chapter discusses System i performance and work management concepts. It presents an overview of some key factors that affect system, application, and network performance. It enables you to understand the performance and tuning concepts that are discussed later in this book.

Refer to 26.1, “Tuning System i models” on page 532 and Chapter 27, “Performance management tools and services” for information about analyzing performance and the tools and services available.

25.1 System i performance concepts

This section discusses performance concepts from the System i point of view, including client/server and interactive response time components, such as the impact of database growth. It examines the differences between the components of response time in a client/server System i environment such as an SAP system, and for a traditional interactive (such as 5250 terminal emulation) System i user.

25.1.1 Client/server response time

In a client/server environment, the response time perceived by the user is the sum of the response time of the following components involved in providing the service:

- Client system

  When a user at a client system, such as a personal computer (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, and network.

Typically, a server system functions in an environment with multiple requestors. The response time experienced by a requestor is affected by the function of the particular task, as well as 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 can become a significant contributor to overall client/server performance.

The response time is also increased by the number of times information moves between the client and server (communications flows), and the amount of data moved before a response is completed.

25.1.2 Components of System i interactive response time

There are many functions within a system that contribute to the response time experienced by the interactive user, including the processing time of the central processing unit (CPU) and disk input output (I/O) operations. Each transaction is also affected by the wait times associated with these services.

The interactive response time experienced by a System i user is the total of many components, as shown in Figure 25-1. Consider all components when analyzing performance problems.

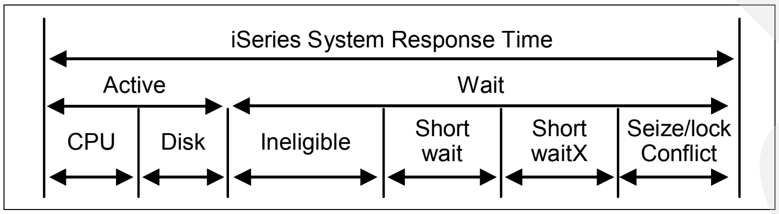

Figure 25-1 Components of System i response time

The following actions comprise the response time of an interactive transaction:

- A transmission time delay for the transaction to reach the CPU. This can be significant in some situations, such as on a remote workstation.
- The system's response time measurement begins once the transaction reaches the system. The following situations are included in the measurement of system response times:
  - The time a job waits for a system activity level.
  - Resource utilization begins once the activity level is entered. Resource utilization includes:
    - CPU processing time (including queuing)
    - Disk I/O time (including queuing)
  - There can be periods of inactivity when the transaction waits in a variety of states. The transaction states are reported in the performance tools report and include the following:
    - Short waits
    - Short waits – extended
  - The time a job waits in a long wait, if a job leaves the activity level (that is, object or record lock conflicts)

  Transaction states are discussed in *Performance Management/400 V4R4,* SC41-5347, and *Performance Tools for AS/400,* SC41-5340.
- There is a transmission delay in the response reaching the user.
- The time taken by the user workstation to process the information to present to the user.

25.1.3 Impact of database growth on System i performance

System i models use 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. This approach minimizes the number of searches involved in finding the required index. Also, the radix tree is partitioned at a higher level in the tree when a page segment of a binary radix tree becomes full. This minimizes the number of page segments that need to be searched. This structure is referred to as a *partitioned binary radix tree*.

On an 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. The total processing time is directly related to the file size. The response time is slow when many entries have to be processed, and likewise, the tables become larger. In other words, the greater the database becomes, the more important a suitable index is.

Another kind of index, the Encoded Vector Indexes (EVIs) are used in mySAP BW, for example, in the same manner as binary radix tree indexes. EVIs become greater with increasing tables.

25.2 System i work management concepts

Work management on System i models is the method used to manage System i resources to achieve optimum throughput. System i models have many objects that interact with each other and applications that process information. An understanding of the concepts of work management is an important prerequisite to maximize system performance.

This section discusses the System i work management concepts that affect performance.

25.2.1 Subsystems

A *subsystem* in i5/OS is an operating environment used to allocate main storage. A subsystem provides a degree of isolation for jobs with similar processing characteristics (such as batch or interactive). This minimizes the contention for system resources and increases system efficiency.

Use the Display Subsystem Description (**DSPSBSD**) command to view the components of the QBATCH subsystem that is supplied with the i5/OS operating system. The SAP subsystem descriptions can be found in the R3*SID*400 library, where *SID* is a three-character SAP system identifier. The SAP subsystem names are R3_*nn*, where *nn* is the instance number that is assigned during installation.

Some of the important components of a subsystem are as follows:

- Storage/memory pools

  Storage pool definitions provide information about:

  – Pool identification (within the subsystem)
  – Pool size
  – Maximum activity level

- Work entries

  Work entries identify the sources from where jobs can enter a subsystem. Specific types of work entries are used for different types of jobs. For example, an autostart job uses an autostart job entry.

- Autostart job

  These subsystem entries perform a one-time initialization job or a repetitive activity associated with a subsystem. The QSERVER subsystem uses a QPWFSERVER autostart job to initiate file and database servers.

- Routing entries

  In a subsystem, routing entries provides 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

    This is the run priority relative to the priority of other jobs running in the system. Run priority ranges from 1 (highest) to 99 (lowest). Jobs with the highest priority receive the best service from the CPU.

  - Time slice

    This shows the maximum amount of processor time, in milliseconds, given to each thread in the job before other threads in the job or other jobs are given the opportunity to run. The time slice establishes the amount of time needed by the thread to accomplish a meaningful amount of processing. At the end of the time slice, the thread may be made inactive so that other threads can become active in the storage pool.

  - Purge option

    This specifies whether the job is eligible to be moved out of main storage and put into auxiliary storage at the end of a time slice, or when there is a long wait (such as waiting for a user's response). This attribute is ignored when more than one thread is active within the job.

  The SAP execution classes have names in the format of R3_ *nn*, where *nn* is the instance number. The execution classes are located in the R3*SID*400 library, where *SID* is the SAP system identifier.

- Communications job

  A communications job, from an i5/OS work management standpoint, is a batch job started by a program start request from a remote system. In the case of servers, a client or PC application initiates the start request.

- Prestart job

  This is a batch job that starts 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.

In subsystems that are initiated by the SAP start (**STARTSAP**) command, communication jobs and prestart jobs are not used.

Other important work management objects are the job description (JOBD) and job queue (JOBQ).

- Job description

  A job description object contains a specific set of job-related attributes that can be used by one or more jobs. The attributes determine how each job is run on the system. The same job description can be used by multiple jobs. The values in the job description are usually used as the default values of the corresponding parameters in the Batch Job (**BCHJOB**) and Submit Job (**SBMJOB**) commands when their parameters are not specified.

- Job queue

  A job queue contains entries for jobs that are waiting to be processed by the system. Jobs can be placed on a job queue, for example, by using the **SBMJOB** command.

  After creating a new job queue, you must add an entry for it in the appropriate subsystem description using the Add Job Queue Entry (**ADDJOBQE**) command. A job queue entry identifies a job queue from which jobs are selected for running in the subsystem.

If you start the SAP system, the job queue entry transfers the SAPSTART job to the determined subsystem that SAP starts in the R3_*nn* subsystem. From there, the dispatcher job spawns the work processes. With this method, all attributes of the jobs are inherited from the original job.

Education courses are available for a greater understanding of work management and performance. These include:

- OS/400 Structure, Tailoring, and Basic Tuning (Course Code S6023)
- iSeries Performance Analysis and Capacity Planning (S6095)

Use the keywords “education” or “training’ to search your country-specific IBM internet Web site to find additional courses and those specific to your country at:

`http://www.ibm.com/`*`countrycode`*

For the United States, the Web site is:

`http://www.ibm.com/us/`

You can find more information about all aspects of work management in *OS/400 Work Management,* SC41-5306.

25.3 Memory management of an SAP system

This section introduces the memory management system used with mySAP applications. The concepts of an SAP memory management are outlined. The discussion includes which parameters are available, and what the optimal settings for your system are.

This section refers to memory management for SAP kernel releases 6.40 and higher. Changes occur as enhancements to the SAP kernel and i5/OS are developed. Considerations also depend on the type of SAP instance is involved (production or development), the total amount of physical memory and storage available, the number of SAP work processes, and the SAP modules and transactions being used.

Refer to the SAP note 808607 iSeries: Memory Management for the latest information relative to the 6.40 kernel.

Refer to the SAP online help documentation and the following SAP notes for more information about earlier releases, such as:

- SAP note 139326: Memory management in releases as of 4.6A, AS/400
- SAP note 103747: Performance: Parameter recommendations for Release. 4.0 and later
- SAP note 44695: AS/400: Memory management as of release 3.0C

SAP note 44695 is valid for releases 3.0C to 4.5B.

Refer to the following SAP notes for the 4.6D optimized kernel:

- SAP note 611946: iSeries: Optimized kernel for the POWER4™ technology
- SAP note 647616: iSeries: Improved performance on POWER4 hardware

25.3.1 Functions of the SAP memory management system

An application runs in an SAP work process where an Advanced Business Application Programming (ABAP) program normally executes. The process requires memory to do this. The memory is allocated to the process by the memory management system. The order of assignment from these memory areas arranges itself according to whether the user context runs in an SAP dialog work process or in another SAP work process, and otherwise depends on the underlaying operating system.

Figure 25-2 shows the location of the various memory areas.

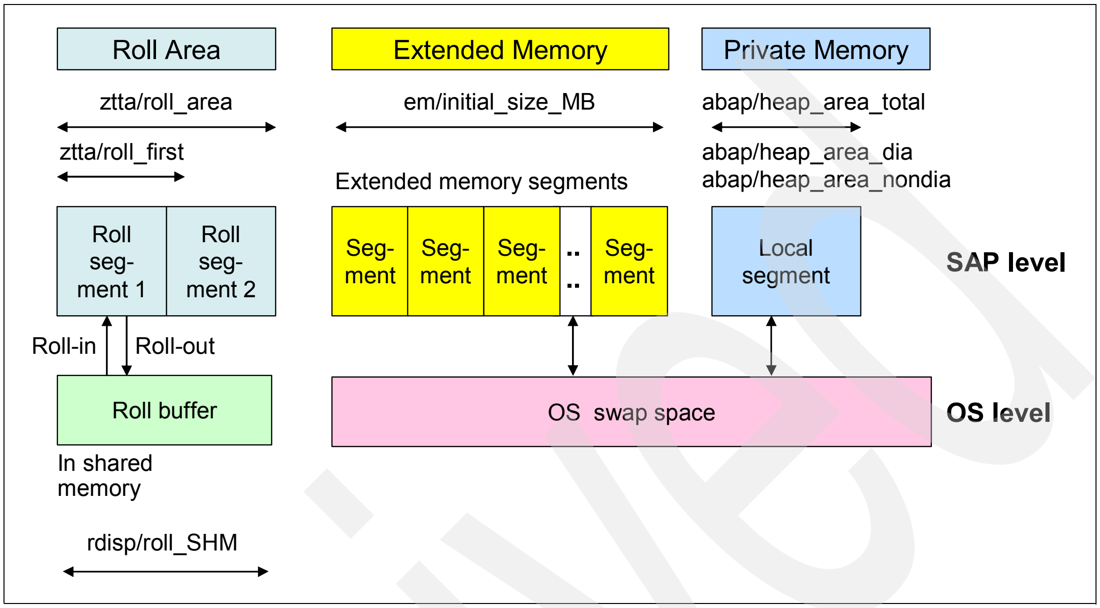

Figure 25-2 SAP memory management system

The area of a user context that is directly accessible is extended as needed if the user context has expanded. For dialog work processes, the data of the user context, including internal tables, is located in this extended area. Therefore, you can access all of the data that is in the user context. Only the extract type data and export to memory type data stay in the SAP paging area.

The SAP Roll Area is used for the initial memory assigned to a user context, and (if available) for additional memory if the extended memory is full.

Figure 25-2 displays the memory types that can be assigned to work processes on the SAP system and operating system level. These are the most important system profile parameters that control the availability of the memory types.

A roll action occurs between the roll buffer in the shared memory and the memory area whenever a dialog step is executed as follows:

- Roll-in

  Cross-user data is rolled in from the shared resources in the work process (and is processed there).

- Roll-out

  User-specific data is rolled out from the work process in the shared resources (after the dialog step has ended).

The user context is allocated according to ztta/roll_first in a dialog process. Then the shared memory area is accessed that belongs to this user context.

Figure 25-3 on page 464 displays the roll process performed by the dispatcher. The shared resource identifies the different SAP memory types.

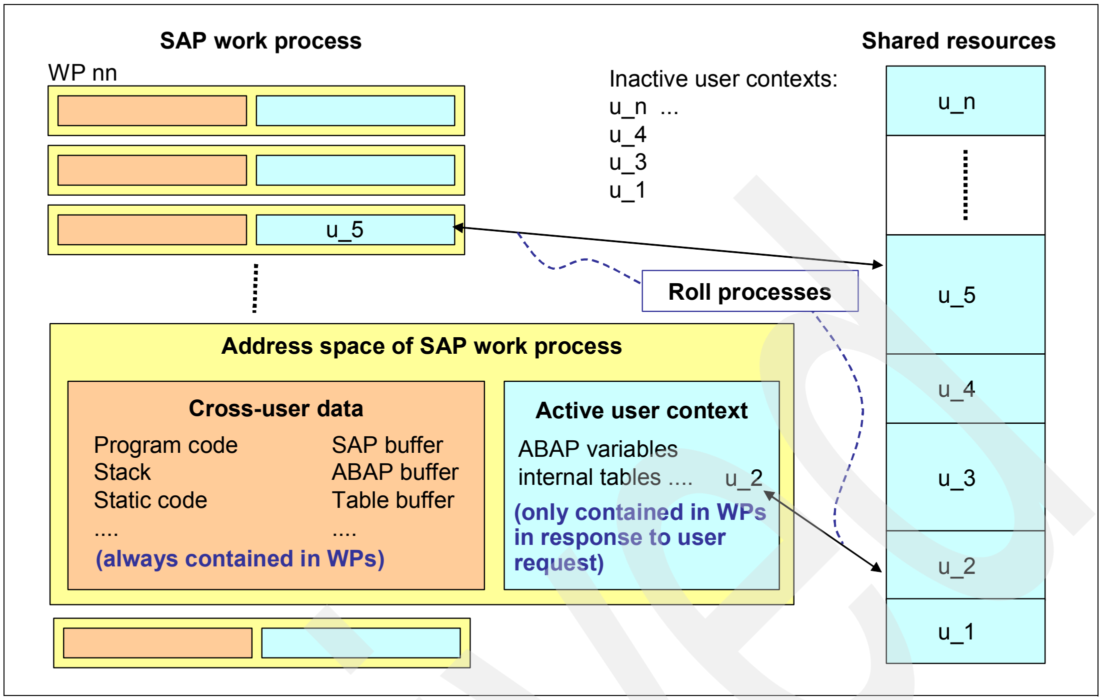

Figure 25-3 SAP work process and shared ressources

Refer to Chapter 16 of *Implementing SAP R/3 on OS/400*, SG24-4672 for an explanation of SAP memory implementations on releases prior to Version 6.40 and SAP applications on System i models.

AIX programs for PowerPC architecture Systems can be binary-compatible for System i models. Therefore kernel 6.40 is supported in the i5/OS PASE runtime environment.

System i memory management is similar to the UNIX model of memory management. However, the restrictions of AIX memory management are avoided with i5/OS PASE on System i models.

For further information about i5/OS PASE, refer to 8.6, “i5/OS Portable Application Solutions Environment” on page 84. Also refer to the following Web site:

`http://www-1.ibm.com/servers/enable/site/porting/iseries/pase`

Install OS/400 Portable Application Solutions Environment (5722-SS1 Option 33) to run the program code in i5/OS PASE. It is a no-charge feature of the operating system.

Refer to 8.6.1, “i5/OS PASE on System i models” on page 84 and the following Web site for more information about i5/OS PASE:

`http://www-1.ibm.com/servers/enable/site/porting/iseries/pase`

Every form of SAP memory (roll, extended, heap, and paging memory) works under i5/OS with i5/OS PASE at least in teraspace with the 6.40 kernel. An address range of one terabyte is available for an SAP system. The i5/OS-specific restrictions of the earlier releases are revoked.

Some of the key SAP values that affect memory and disk utilization, and are specified in the SAP instance profile are listed here. For example, *SID*_DVEBMGS*nn_hostname*, where *SID* is the SAP system identifier, *nn* is the instance number, and *hostname* is the Transmission Control Protocol/Internet Protocol (TCP/IP) host name.

- ztta/roll_area

  This value specifies the total roll memory available to each work process. The default value is platform-specific and is predetermined. Normally, you should not change the default value.

- ztta/roll_first

  The SAP 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 rolled out. The remaining part of roll memory is used after extended memory allocation. Normally, you should not change the default value.

- ztta/roll_extension

  This value limits the amount of SAP extended memory that one user context can allocate. The ztta/roll_extension parameter is specified in bytes. This value is valid for all types of work processes.

  You can control the quotas for dialog and non-dialog work processes separately. Use parameters ztta/roll_extension_dia and ztta/roll_extension_nondia to do this.

- em/initial_size_MB

  This value specifies the extended memory pool size for use by all dialog user contexts that the SAP system manages for the SAP extended memory. The value is specified in megabytes and should be between 512 and 32,000. User contexts use extended memory up to the limit specified in the roll_extension parameter after roll_first memory is allocated, until the extended memory is exhausted.

- rdisp/PG_SHM

  This value specifies the maximum amount of paging buffers in 8 KB blocks. All paging is done in the shared memory of the System i server. Set this value equal to the rdisp/PG_MAXFS value.

- rdisp/PG_MAXFS

  This value specifies the size of the SAP paging file in 8 KB blocks. All paging is done in the shared memory of the System i server. This value must be equal to the rdisp/PG_SHM value.

- abap/heap_area_dia

  This value restricts the amount of heap memory (private memory) that an SAP dialog work process can allocate.

- abap/heap_area_nondia

  This value restricts the amount of heap memory (private memory) that an SAP non-dialog work process can allocate.

- abap/heap_area_total

  This parameter determines the upper limit of the heap memory in bytes available to all work processes of an SAP application server.

  System i models run with the single-level storage concept, so there is sufficient memory available in total, without limitations caused by the main storage, as long as the direct access storage device (DASD) is not full.

- abap/heaplimit

  This value specifies the memory amount in bytes. A work process restarts after executing a dialog step when this amount is exceeded. The restart releases any private memory allocated to the work process.

- rdisp/wppriv_max_no

  This value specifies the maximum number of dialog work processes that can be in private mode. Private memory is assigned to the work process if a dialog work process has used up the roll area assigned to it and the extended memory. The work process goes into private mode. This parameter ensures that not too many dialog work processes are in the private mode simultaneously, as this lowers the system performance.

- rdisp/max_priv_time

  The system waits for the number of seconds specified before restarting the oldest work process in private mode once the limit of work processes in private mode is reached.

The following are some other parameters that are not used with i5/OS:

- ipc/shm_psize

  These are shared memory pool descriptor sizes that are used to define the shared memory segments. The number of shared memory segments per work process is limited in some platform implementations.

- abap/use_paging

  This value specifies the use of SAP's new memory management facilities.

- rdisp/ROLL_SHM

  This value specifies the size of the roll buffer in 8 KB blocks. The roll buffer functions as the cache for the roll file. The roll areas of all user contexts should find space in the roll buffer for a fast change of context to occur. Buffered data is in the shared memory. Accessing data here is faster than accessing a file.

- rdisp/ROLL_MAXFS

  This value specifies the size of the roll file in 8 KB blocks. The roll file functions as memory for the portion of the user contexts that finds space in the roll area.

i5/OS memory pools are allocated to subsystems. The memory requirement of an SAP instance is managed within a subsystem. The total memory available to an SAP instance is the memory that is allocated to the corresponding R3_*nn* subsystem, where *nn* is the instance number. Naturally, SAP jobs compete for resources with other jobs assigned to the same memory pools. The corresponding i5/OS subsystem description defines the amount of memory assigned to each SAP instance.

SAP implementation and System i architecture considerations include the following:

- Shared memory pools are not used in a System i implementation. The ipc/shm_psize parameters are therefore ignored.
- For paging memory, there is shared memory available on System i models, as long as the disk units are not full. This is why no distinction is made between paging-cache (rdisp/PG_SHM) and a paging file (rdisp/PG_MAXFS) in the SAP paging memory. The paging-cache is just as large as the paging file. The entire paging file is buffered and, therefore, does not exist.

  Ensure that rdisp/PG_SHM and rdisp/PG_MAXFS always have the same value. SAP paging memory is allocated when you start the SAP application server.

- SAP applications on System i models with the 6.40 kernel or later have the same implementation for the SAP roll memory as UNIX systems. However, the system does without the roll file for the same reasons as for the paging memory.

  A user context is generated with an external and internal mode when a user logs on to the SAP system. This means that the roll memory administration allocates an administration shared memory, as well as a roll area.

The roll area is preallocated up to the size of ztta/roll_first when you create an internal mode. The roll area is retrieved from the extended memory if you require a larger roll area, up to a maximum of the size of the ztta/roll_extension. Space is requested from the roll memory up to the upper limit (ztta/roll_area - ztta/roll_first) for any further requirements. This procedure should prevent temporary memory being unnecessarily consumed.

You can set the ztta/roll_area and ztta/roll_first profile parameters to any value as long as the following condition is met:

```
1 <= ztta/roll_first =< ztta/roll_area
```

Set the **roll_first** space to one byte. Specify the roll_area as a high value.

- Heap memory states process local private memory. This mode can no longer be rolled, since this heap memory is not addressable by another work process if heap memory is allocated to fulfill the memory request of a mode (user context). Therefore, memory is exclusively assigned to a mode. This work process status is called *PRIV mode* in the SAP context. The work process cannot edit any other user context while it is in the PRIV status.
- You should try to prevent dialog work processes from entering the PRIV mode on a System i model if you have a sufficiently large quota for the extended memory (ztta/roll_extension). The PRIV mode is irrelevant for non-dialog work processes, since rolling does not take place here.

  The extended memory is allocated in accordance with the value of em/initial_size_MB when you start the SAP instance. You cannot change the size of the extended memory during the runtime of an instance.

  Between the parameter em/initial_size_MB and em/blocksize_KB, there are some very complex dividing factors. The default of em/initial_size_MB is 4096 and the default of em/blocksize_KB is 4096 as well. Typical adapt your em/initial_size_MB parameter and do not change em/blocksize_KB at all.

  The blocks that are selected as free in the extended memory pool can be reused if the extended memory is released by the user context.

  All SAP shared memories are created during startup and then show up as temporary storage. Because of some other temporary objects, this temporary storage is increased more or less over the running time.

  Temporary memory is not released before an SAP instance is terminated.

  Set ztta/roll_extension high enough to avoid a dialog work process ever entering PRIV mode. The default of 2000000000 is typically an optimal size.

The SAP parameters are set to defaults when SAP is first installed. Review and adjust the parameters after implementing the system to ensure the best performance.

The values shown in the recommendation column in Table 25-1 apply to SAP kernel release 6.40 and are examples of a possible production SAP instance on the System i server.

Table 25-1 SAP parameter settings on System i models

Parameter name	Recommendation for System i	Explanation
ztta/roll_first	1 (This value is set automatically and cannot be changed.)	1 byte first roll area

Parameter name	Recommendation for System i	Explanation
ztta/roll_area	3000000 (This value is set automatically and cannot be changed.)	3 MB roll area per internal mode
ztta/roll_extension	2000000000	2 GB extended memory/context per user
rdisp/ROLL_SHM	32768 (default 16384)	number of 8 KB blocks of roll buffer
rdisp/ROLL_MAXFS	32768 (default)	number of 8 KB blocks of max roll file
em/initial_size_MB	default 4096 (512 to 32000)	4096 MB extended memory altogether
em/blocksize_KB	4096	4096 KB for each memory allocation request (Do not change this value)
rdisp/PG_SHM	32768 (same value as rdisp/PG_MAXFS)	number of 8 KB blocks of maximum SAP paging file
rdisp/PG_MAXFS	32768 (same value as rdisp/PG_SHM)	number of 8 KB blocks of paging buffer
abap/heaplimit	40000000	40 MB workprocess restart
abap/heap_area_dia	2000000000	2 GB heap per dialog WP
abap/heap_area_nondia	2000000000	2 GB heap per non-dialog WP
abap/heap_area_total	2000000000	2 GB heap per SAP system

As of kernel 6.40, there are various strategies to allocate memory for dialog and non-dialog work processes for the allocation sequence of memory on System i servers.

In all dialog work processes, the first part of the roll area (ztta/roll_first) is first exclusively used, followed by extended memory (ztta/roll_extension). Thereafter, the second part of the roll area (ztta/roll_area - ztta/roll_first) is used, which is finally followed by heap memory (abap/heap_area_dia). Figure 25-4 on page 469 illustrates this.

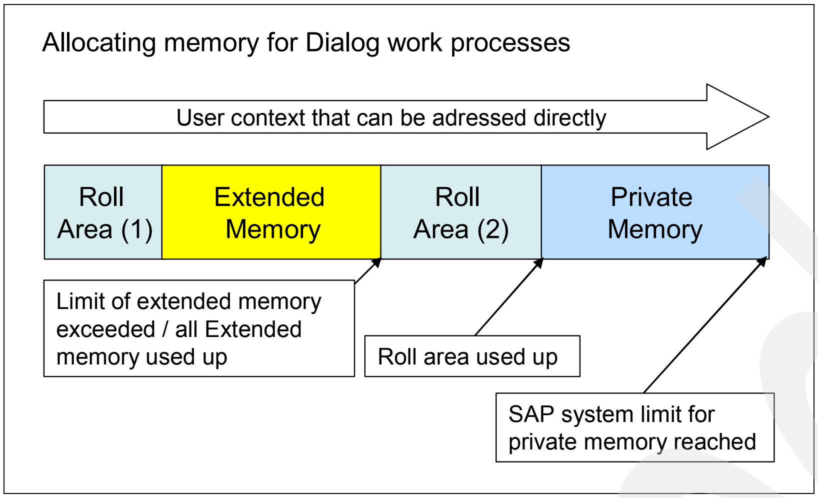

Figure 25-4 Allocation sequence for dialog work processes

In all other types of work processes, the first and the second part of the roll area (to ztta/roll_area) is used first, then the heap memory (abap/heap_area_nondia), followed by the extended memory (ztta/roll_extension) are used.

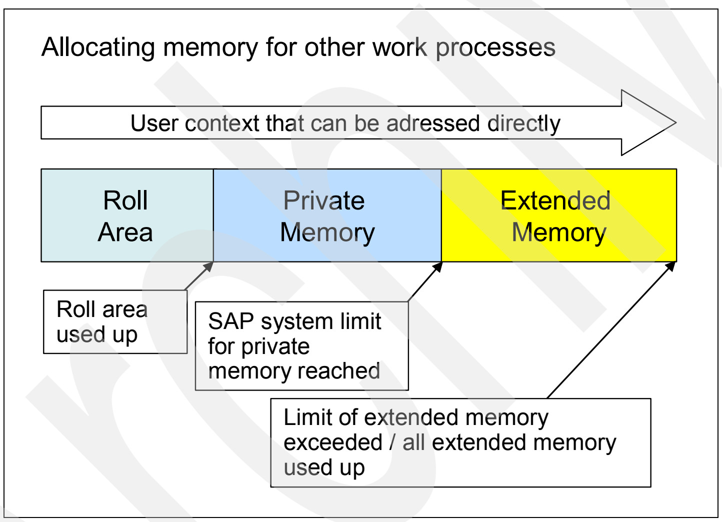

Figure 25-5 Allocation sequence for non-dialog work processes

25.4 Performance monitoring

This section introduces you to a performance management methodology of SAP applications on System i servers. It covers System i monitoring functions, as well as SAP monitoring functions.

Using operating system commands, as well as SAP tools and transactions is inevitable in order to estimate system performance. SAP extensively uses structured query language (SQL) queries to DB2 UDB for i5/OS as a System i client/server application. Therefore, this section also describes the functions to monitor database performance.

Routinely measuring and analyzing your system performance enables you to perform capacity planning and analyze specific performance problems. You can prevent bottlenecks in your system before real performance problems appear with tuning and capacity planning.

25.4.1 System i monitors

The i5/OS operating system includes a number of commands that can allow you to perform real-time monitoring of performance data from the 5250 interface. You can use these commands to answer specific questions about system performance and to help you tune your system. Similar real-time monitoring functions can be executed from the work management component of the iSeries Navigator.

Monitor your system in a preventive manner. When performance problems occur, they often affect specific components of the system first. Table 25-2 shows some of the methods available for analyzing performance on these system components.

Table 25-2 System i monitors

Component	Description	Command or tool
CPU	CPU utilization of system or logical partition (LPAR)	Work with Active Jobs (WRKACTJOB) Work with System Activity (WRKSYSACT) Work with System Status (WRKSYSSTS) iSeries Navigator - A work management function CPU utilization metrics within the iSeries Navigator system monitor Performance Tools reports IBM Insight Tool, SAP Early Watch Services
Main storage	Faulting rate Wait-to-ineligible transitions	Work with System Status (WRKSYSSTS) iSeries Navigator - a work management function - memory pools Performance Tools reports
Disk	Disk space utilization Disk arm busy utilization Balanced arms and space utilization	Work with Disk Status (WRKDSKSTS) iSeries Navigator - Disk arm utilization metrics Performance Tools - System Report and Component report

Component	Description	Command or tool
IOPs	IOP utilization Balanced IOP utilization	Performance Tools - System Report and Component Report iSeries Navigator - IOP utilization metrics
Communication	Line utilization / errors	Performance Tools - Component Report. iSeries Navigator - LAN utilization metrics

25.4.2 WRKSYSSTS

Run the Work with System Status (**WRKSYSSTS**) command for at least a few minutes to check the resource utilization. Press **F10** to restart the calculations and press **F5** to extend the elapsed time. Figure 25-6 illustrates this.

```
                             Work with System Status
                                                                    ITSO-SYS1
                                                           12/05/05  05:13:49
% CPU used . . . . . . . :       26.5    Auxiliary storage:
% DB capability  . . . . :        5.2      System ASP . . . . . . :     1336 G
Elapsed time . . . . . . :   00:04:51      % system ASP used  . . :    61.8614
Jobs in system . . . . . :       5072      Total  . . . . . . . . :     1336 G
% perm addresses . . . . :       .009      Current unprotect used :     9246 M
% temp addresses . . . . :       .022      Maximum unprotect  . . :    11365 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     512.00    198.90   +++++     .1     .1     .4     .6
   2    7647.54      5.44     300   21.4   13.1   65.6  200.6
                                                                       Bottom
Command
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F10=Restart   F12=Cancel
F19=Extended system status         F24=More keys
```

Figure 25-6 WRKSYSSTS (view 1)

The following is an explanation of the information shown in Figure 25-6:

- % CPU used

  This indicates the percentage of the CPU utilized during the elapsed time. This value can be high for a longer time, for example, if a large batch job is running. One job can consume the maximum (that is, 100% CPU divided by the number of processors). For instance, in the case of a four-way machine, this is 25%.

- % DB capability

  This indicates the percentage utilization of CPU during the elapsed time for database operations.

- % system ASP used

  This indicates the percentage of systems disk space used.

- Current unprotected used and Maximum unprotected

  This indicates the amount of disk space for temporary objects. This is often referred to as *Tempspace*.

- Pool Size (M)

  This indicates the current memory pools sizes in megabytes.

- DB Faults / Non-DB Faults

  This indicates page faults per second. A page fault occurs when an active program refers to a page that is marked as not in main storage.

Press **F11** to view the job transition information. Figure 25-7 illustrates this.

```
                              Work with System Status
                                                                     ITSO-SYS1
                                                             01/22/06 05:13:49
% CPU used . . . . . . . :       26.5    Auxiliary storage:
% DB capability  . . . . :        5.2      System ASP . . . . . . :      1336 G
Elapsed time . . . . . . :   00:04:51      % system ASP used  . . :    61.8614
Jobs in system . . . . . :       5072      Total  . . . . . . . . :     1336 G
% perm addresses . . . . :       .009      Current unprotect used :    9246 M
% temp addresses . . . . :       .022      Maximum unprotect  . . :   11365 M

Type changes (if allowed), press Enter.

System    Pool    Reserved    Max    Active->  Wait->  Active->
 Pool   Size (M)  Size (M)  Active    Wait     Inel     Inel
   1     512.00    198.20   +++++    12.5       .0       .0
   2    7647.54      5.44     300    3624       .0       .0

                                                                        Bottom
Command
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F10=Restart   F12=Cancel
F19=Extended system status         F24=More keys
```

Figure 25-7 WRKSYSSTS (view 2) - F11 transition data

- Wait → Inel

  This measures the average number of jobs transitioning from a wait state to the ineligible state per minute. When this occurs, it indicates that the Max Active jobs value in this memory pool is set too low.

25.4.3 WRKDSKSTS

The Work with Disk Status (**WRKDSKSTS**) command indicates the disk activity and disk utilization of your system. Figure 25-8 illustrates this.

```
                              Work with Disk Status
                                                                   ITSO-SYS1
                                                              02/27/06 05:26:59
Elapsed time:   00:06:27

               Size    %     I/O   Request   Read  Write   Read  Write    %
Unit  Type      (M)  Used    Rqs  Size (K)    Rqs   Rqs     (K)    (K)   Busy
   1  4326    26373  61.9    1.3     17.6    1.0     .3    15.8   23.2     1
   2  4326    30769  61.9    1.1     25.4    1.0     .1    26.8   13.0     1
   3  4326    30769  61.9     .8     31.3     .6     .1    29.9   37.1     0
   4  4326    30769  61.9     .6     30.1     .5     .1    29.8   31.5     0
   5  4326    35165  61.9    1.3     26.7    1.1     .1    23.0   49.3     1
   6  4326    30769  61.9    1.3     24.5     .9     .3    28.1   15.4     0
   7  4326    26373  61.9    1.1     23.6     .9     .1    23.5   24.4     0
   8  4326    30769  61.9    1.4     23.7    1.2     .1    24.5   18.5     1
   9  4326    30769  61.9    1.1     26.1     .9     .2    23.2   39.6     0
  10  4326    30769  61.9    1.1     19.9     .9     .1    22.6    3.4     0
  11  4326    26373  61.9     .9     27.9     .8     .0    30.7    4.0     0
  12  4326    30769  61.9    1.2     28.8    1.1     .1    26.4   46.9     0
  13  4326    26373  61.9     .8     32.2     .6     .1    28.6   52.4     0
                                                                    More...
Command
===>
F11=Display storage use   F16=Work with system status   F24=More keys
```

Figure 25-8 WRKDSKSTS

Following is an explanation of the information shown in Figure 25-8:

- **%** Used

  This indicates the percentage amount of disk space used of this disk unit. The usage of the disk should be balanced across the disk arms.

- % Busy

  This indicates the percentage disk arm utilization. This value is very important for disk performance.

- I/O Rqs

  This indicates the number of read and write requests per second. This value is determined by the maximum rate for a specific type of disk.

25.4.4 WRKACTJOB

Use the Work with Active Job (**WRKACTJOB**) command to review all active jobs in your system or LPAR. Each SAP instance has its own subsystem with a unique instance number that is created during the installation. For example, SAP jobs of the SAP instance 03 run in subsystem R3_03. Refer to Figure 25-9.

Specify the SBS parameter of the **WRKACTJOB** to limit the appearance of jobs for a specific subsystem. For example, **WRKACTJOB SBS(R3_03)** shows only the jobs of the SAP instance 03.

```
                                Work with Active Jobs
                                                                    ITSO-SYS1
                                                           03/20/06 07:29:50
 CPU %:     22.9     Elapsed time:    00:05:10     Active jobs:   198

 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 ...
                                                --------Elapsed---------
 Opt  Subsystem/Job  Type  Pool  Pty       CPU  Int    Rsp  AuxIO  CPU %
        IGS400       BCI      2  20      101.8                  0     .0
        IGS400       BCI      2  20      103.5                  0     .0
        MSG_SERVER   BCI      2  12       77.1                  0     .0
        RSLGCOLL     BCI      2  20        1.0                  0     .0
        RSLGSEND     BCI      2  20         .4                  0     .0
        SAPSTART     BCH      2  20         .3                  0     .0
        WATCHDOG     BCI      2  20        8.7                  0     .0
        WP00         BCI      2  20     7247.1                  1    5.3
        WP01         BCI      2  20     1148.4                  0    2.7
                                                                      More...
 Parameters or command
 ===>
 F3=Exit   F5=Refresh       F7=Find      F10=Restart statistics
 F11=Display thread data    F12=Cancel   F23=More options   F24=More keys
```

Figure 25-9 WRKACTJOB (view 1)

Following is an explanation of the information shown in Figure 25-9:

- Pool / Pty

  This indicates the pool and priority where the job runs.

- CPU

  This indicates the usage of CPU in seconds. The number displayed is the cumulated value since the job was started.

- CPU%

  This indicates the percentage utilization of CPU during the elapsed time

Press **F11** to change between different views. Refer to Figure 25-10:

```
                               Work with Active Jobs
                                                                       ITSO-SYS1
                                                              12/07/05  07:29:54
CPU %:     23.0    Elapsed time:   00:05:14     Active jobs:   196

Type options, press Enter.
  9=Exclude   10=Display call stack   11=Work with locks
  12=Work with threads   14=Work with mutexes ...

Opt  Subsystem/Job  User        Type  CPU %  Function        Status
       SAPSTART     ERP03       BCH      .0  PGM-sapstart     EVTW
       WATCHDOG     ERP03       BCI      .0  PGM-disp+work    SELW
       WP00         ERP03       BCI     5.3  PGM-disp+work    SEMW
       WP01         ERP03       BCI     2.7  PGM-disp+work    RUN
       WP02         ERP03       BCI      .3  PGM-disp+work    SEMW
       WP03         ERP03       BCI     1.0  PGM-disp+work    SEMW
       WP04         ERP03       BCI      .0  PGM-disp+work    LCKW
       WP05         ERP03       BCI      .2  PGM-disp+work    MTXW
       WP06         ERP03       BCI      .0  PGM-disp+work    SEMW
                                                                        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 25-10 WRKACTJOB (view 2)

The maximum CPU consumption of a single job is 100% divided by the numbers of processors. For example, this calculates to be 25% for a four-way system. If a single process consumes this value for a long time, a CPU bottleneck can be indicated. A solution could be to use more CPUs for higher throughput in general, or/and faster CPUs with higher clock speeds for shorter computing time that results in less CPU utilization.

`Status` indicates the current condition of this job, for example SEMW, if the SAP work process is idle. A status of `LCKW` or `MTXW` is normal. However, if processes stay for several minutes in this state, it can indicate an application (LCKW), kernel, or operating system (MTXW) problem. Use option **11** and **14** to explore the job states. Refer to Figure 25-11.

```
                              Work with Job Locks
                                                             System:   ITSO-SYS1
 Job:   WP02           User:   ERP03          Number:   021597

 Job status:   ACTIVE

 Type options, press Enter.
   5=Work with job member locks   8=Work with object locks

                              Object                              Member  ASP
 Opt  Object       Library     Type        Lock     Status   Locks   Device
      "CROSS+I"    R3ERPDATA   *FILE-LGL  *SHRRD     HELD     YES
                                          *SHRRD     HELD
                                          *SHRRD     HELD
                                          *SHRRD     HELD
                                          *SHRRD     HELD
      "DD03L+1"    R3ERPDATA   *FILE-LGL  *SHRRD     HELD     YES
      "DD03L+5"    R3ERPDATA   *FILE-LGL  *SHRRD     HELD     YES
                                          *SHRRD     HELD
                                          *SHRRD     HELD
                                                                        More...
 F3=Exit   F5=Refresh   F10=Display job record locks   F11=Display thread data
 F12=Cancel   F16=Job menu
```

Figure 25-11 Work with job locks

WRKSBSJOB

A similar command to explore and to check job details of a subsystem is the Work with Subsystem Jobs (**WRKSBSJOB**) command. All jobs dedicated to the SAP instance are seen, for example, R3_03 as shown in Figure 25-12.

```
                              Work with Subsystem Jobs
                                                                    ITSO-SYS1
                                                           12/05/05  05:48:05
Subsystem  . . . . . . . . . . :   R3_03

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  Job         User        Type     -----Status-----  Function
     DISP_WORK   ERP03       BATCHI   ACTIVE            PGM-disp+work
     GWRD        ERP03       BATCHI   ACTIVE            PGM-gwrd
     ICMAN       ERP03       BATCHI   ACTIVE            PGM-icman
     IGS400      ERP03       BATCHI   ACTIVE            PGM-igswd_mt
     IGS400      ERP03       BATCHI   ACTIVE            PGM-igsmux_mt
     IGS400      ERP03       BATCHI   ACTIVE            PGM-igspw_mt
     IGS400      ERP03       BATCHI   ACTIVE            PGM-igspw_mt
     MSG_SERVER  ERP03       BATCHI   ACTIVE            PGM-MSG_SERVER
                                                                      More...
Parameters or command
===>
F3=Exit      F4=Prompt    F5=Refresh    F9=Retrieve   F11=Display schedule data
F12=Cancel   F17=Top      F18=Bottom
```

Figure 25-12 WRKSBSJOB

25.4.5 WRKSYSACT

The Work with System Activity (**WRKSYSACT**) command is a part of the Performance Tools Licensed Program Product (5722-PT1). Refer to Figure 25-13:

```
                         Work with System Activity                 ITSO-SYS1
                                                           12/05/05  05:34:22
Automatic refresh in seconds  . . . . . . . . . . . . . . . . . . .     5
Elapsed time . . . . . . :   00:02:33    Average CPU util . . . . :       2.4
Number of CPUs . . . . . . :       2     Maximum CPU util . . . . . :     2.4
Overall DB CPU util  . . . :      .8     Minimum CPU util . . . . . :     2.3
                                         Current processing capacity:     2.00
Type options, press Enter.
  1=Monitor job   5=Work with job
                                                           Total  Total   DB
     Job or                                            CPU  Sync  Async  CPU
Opt  Task        User        Number  Thread     Pty   Util  I/O    I/O   Util
     WP02        ERP03       021597  00000002   20     .5     0      2    .3
     JLAUNCH     POP23       022675  000008DB   26     .1     2      2    .0
     SMXCAGER01                                 99     .0    21    799    .0
     WP00        ERP03       021595  00000002   20     .2     1      1    .1
     JLAUNCH     POP23       022675  000008DD   26     .0     5      2    .0
     JLAUNCH     POP23       022675  000008DC   26     .1     1      2    .1
     WP01        ERP03       021596  00000002   20     .2     0      1    .1
     JLAUNCH     POP23       022675  000008DE   26     .0     5      2    .0
                                                                     More...
F3=Exit   F10=Update list   F11=View 2   F12=Cancel   F19=Automatic refresh
F24=More keys
```

Figure 25-13 WRKSYSACT

This command allows you to automatically refresh the displayed data. This helps to find and monitor high-consuming jobs by viewing the following:

- CPU utilization
- Total number of Sync I/O
- Total number of Async I/O
- CPU utilization for database

Press **F16** to change the sort sequence. The sequence can be by:

- CPU
- I/O
- Net storage
- Allocated storage
- Deallocated storage
- Database CPU
- Total waiting time

Note: Run the monitoring commands during a typical level of system activity, and also if a performance bottleneck occurs. Refer to the colored bold items in the screen above for an example.

25.4.6 Start Data Collection Services

To monitor System i performance in intervals of time, bring up historical or archive data for examples of the performance of your System i server or in your SAP system landscape before changes are made. Use the i5/OS Data Collection Services. The performance collection tool is part of i5/OS. The Performance Tools product is not necessary to collect data.

There are different methods to start the data collection. Refer to Figure 25-14.

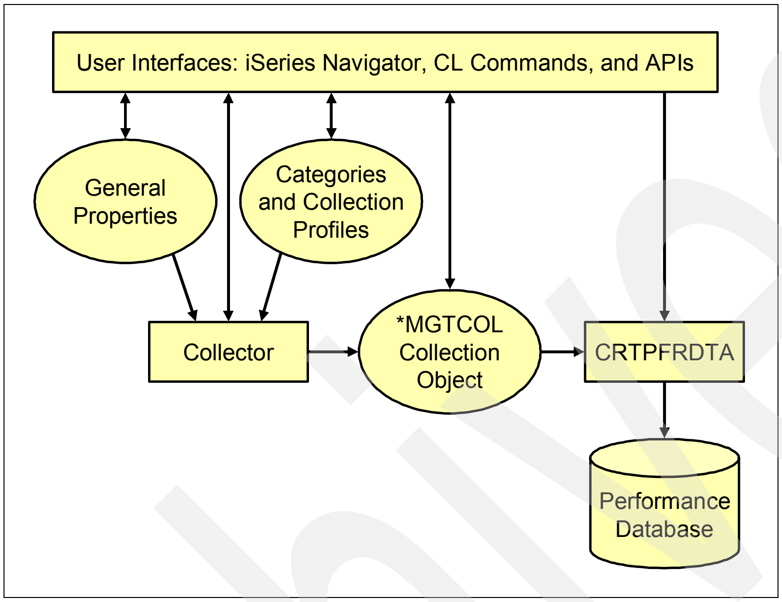

Figure 25-14 Methods to start Data Collection Services

The logic for all data collection methods is the same and is as follows:

1. Setup parameters for the Collection Services.
2. Start the data collection.
3. End the data collection.

To use the Data Collection Services with iSeries Navigator, refer to Figure 25-15

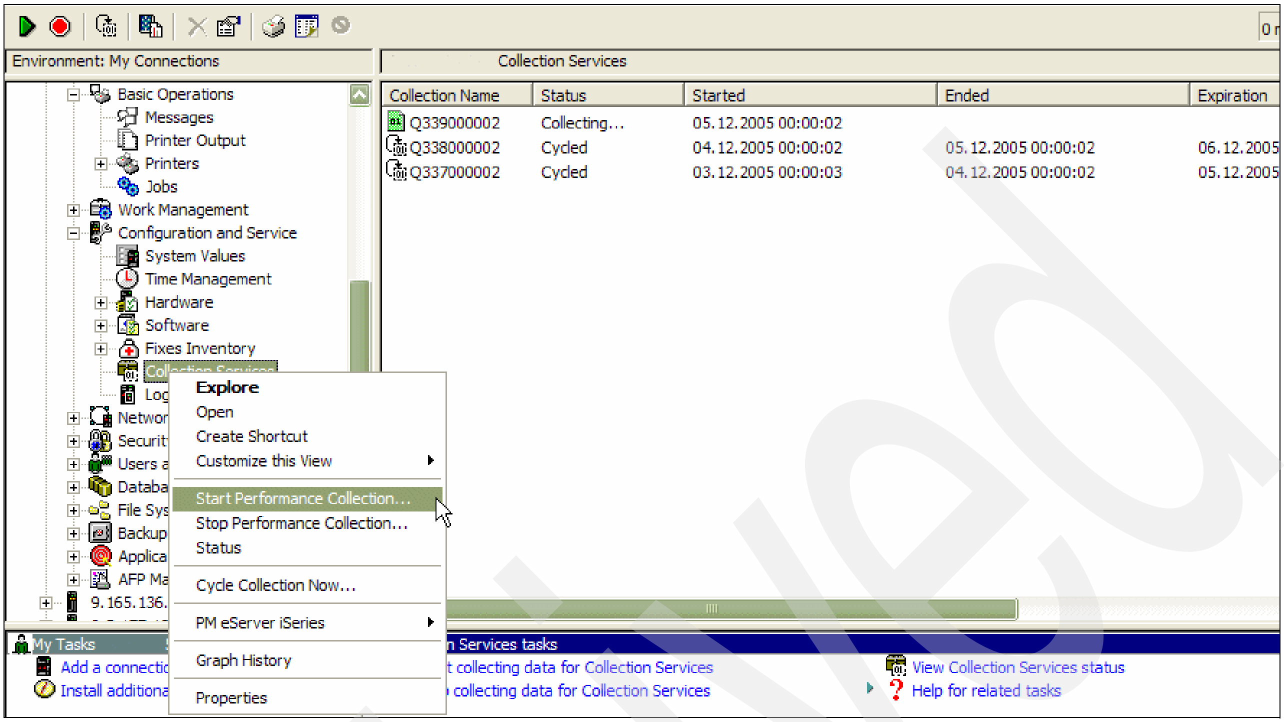

Figure 25-15 Manage Data Collection Services with iSeries Navigator

Use **Properties** to setup the Collection Parameters shown in Figure 25-15. Refer to Figure 25-16 on page 481.

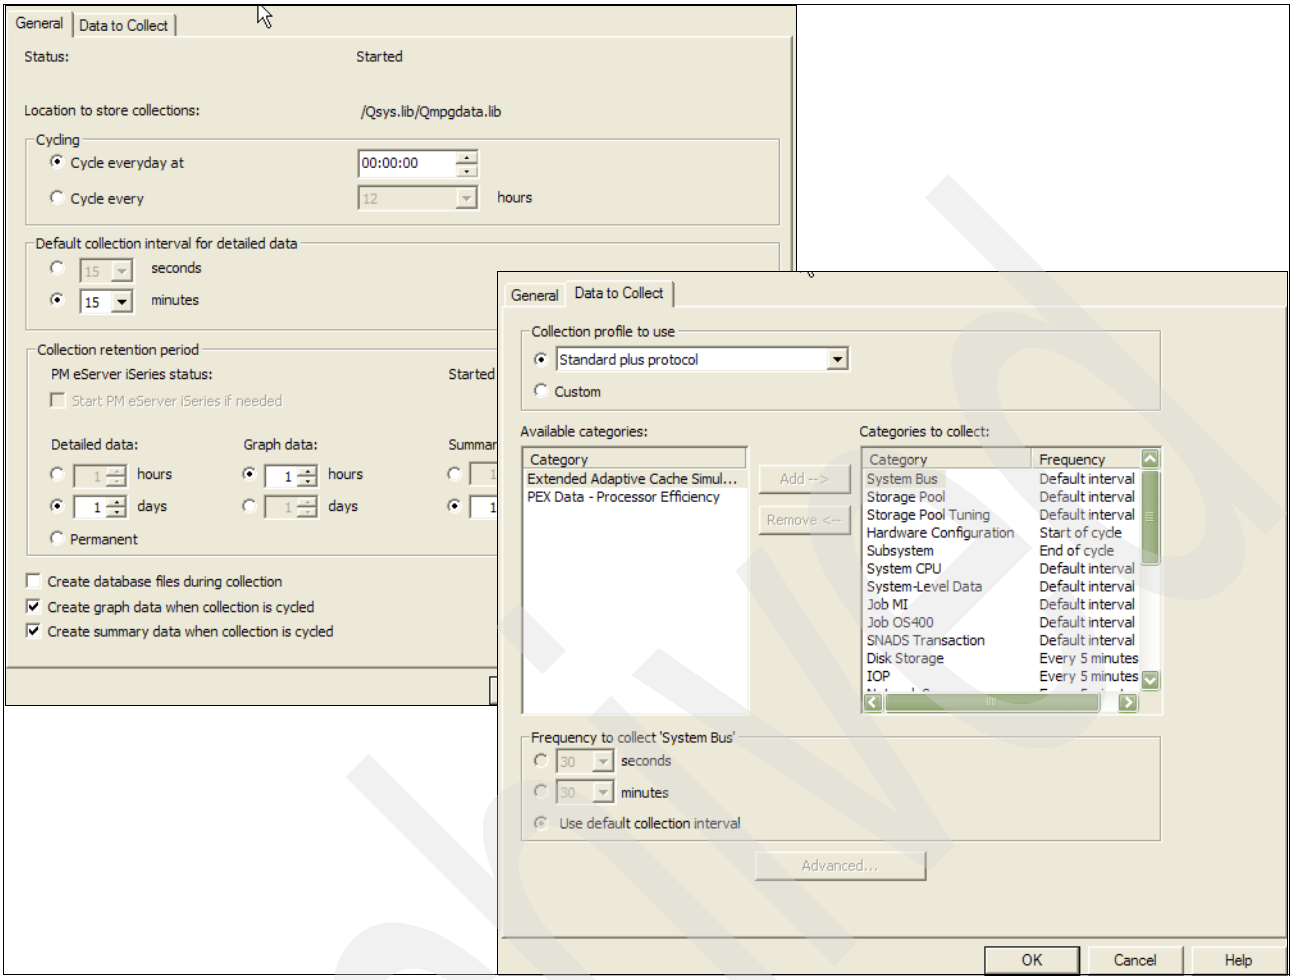

Figure 25-16 Setup parameters of Data Collection Services

You can also use the following i5/OS commands:

- STRPFRCOL
- ENDPFRCOL
- CFGPFRCOL to setup parameters. See Figure 25-17 on page 482.

These commands are not available before i5/OS V5R3.

```
                        Configure Perf Collection (CFGPFRCOL)

Type choices, press Enter.

Default interval . . . . . . . . INTERVAL      15.00
Collection library . . . . . . . LIB           QPFRDATA
Default collection profile . . . DFTCOLPRF     *STANDARDP
Cycle time . . . . . . . . . . . CYCTIME       085900
Cycle interval . . . . . . . . . CYCITV        24
Collection retention period:     RETPERIOD
  Number of units  . . . . . . .               00168
  Unit of time . . . . . . . . .               *HOURS
Create database files  . . . . . CRTDBF        *NO
Change PM/400 library  . . . . . CHGPMLIB      *NO
```

Figure 25-17 CHGPFRCOL

For a default collection profile, use ***STANDARDP**.

The other way to manage Data Collection Services is to use the following application programming interfaces (APIs):

- QYPSSTRC - Start Collector
- QYPSENDC - End Collector
- QYPSCSCA - Change System Collector Attributes

For detailed information, select **Performance Management APIs** → **Collection Services APIs** → **Collector APIs** at the following Web site:

http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp

Note: If you use Data Collection Services for your own data collection, suspend the PM/400 data collection during this time.

Performance Tools (5722-PT1) is required to evaluate and generate performance reports. Performance analysis is also a service available from IBM and some business partners.

For an introduction to using performance reports, refer to 27.4.2, “Performance Tools” on page 585.

25.4.7 iSeries Navigator

The iSeries Navigator offers you a Windows-like and a partial graphical view of your System i system performance, including the monitoring functions.

Figure 25-18 shows an example of System i monitors. **WRKACTJOB** is running in the background, while something like **WRKSYSSTS** is running in the foreground.

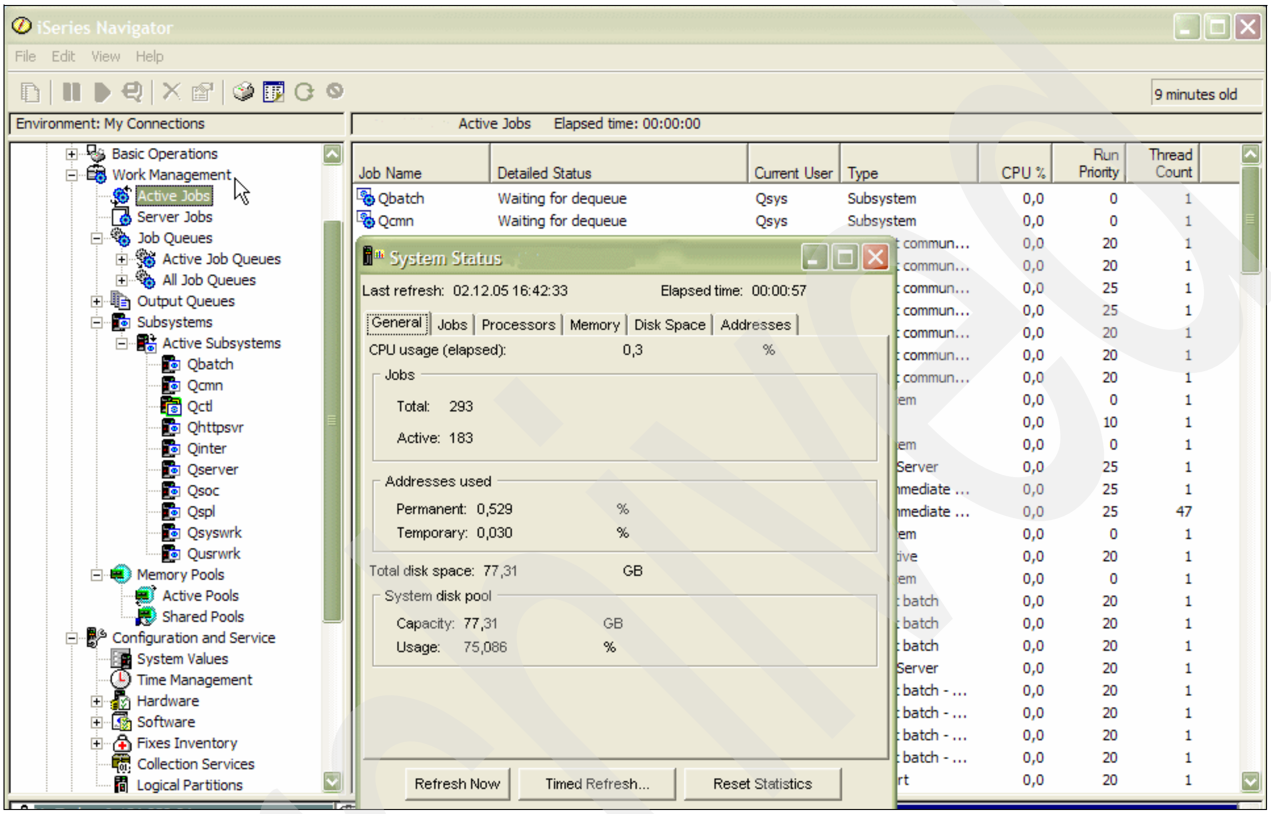

Figure 25-18 iSeries Navigator monitors

You can also define your own monitors by selecting **iSeries Navigator** → **Management Central** → **Monitors**. Right-click **System** and select **New Monitor**. Refer to Figure 25-19 on page 484.

For example, it could be helpful to monitor the values of:

- Machine pool faults
- User pool faults (Maximum)
- CPU utilization
- Disk Arm utilization (Maximum)
- Disk IOP utilization (Maximum)

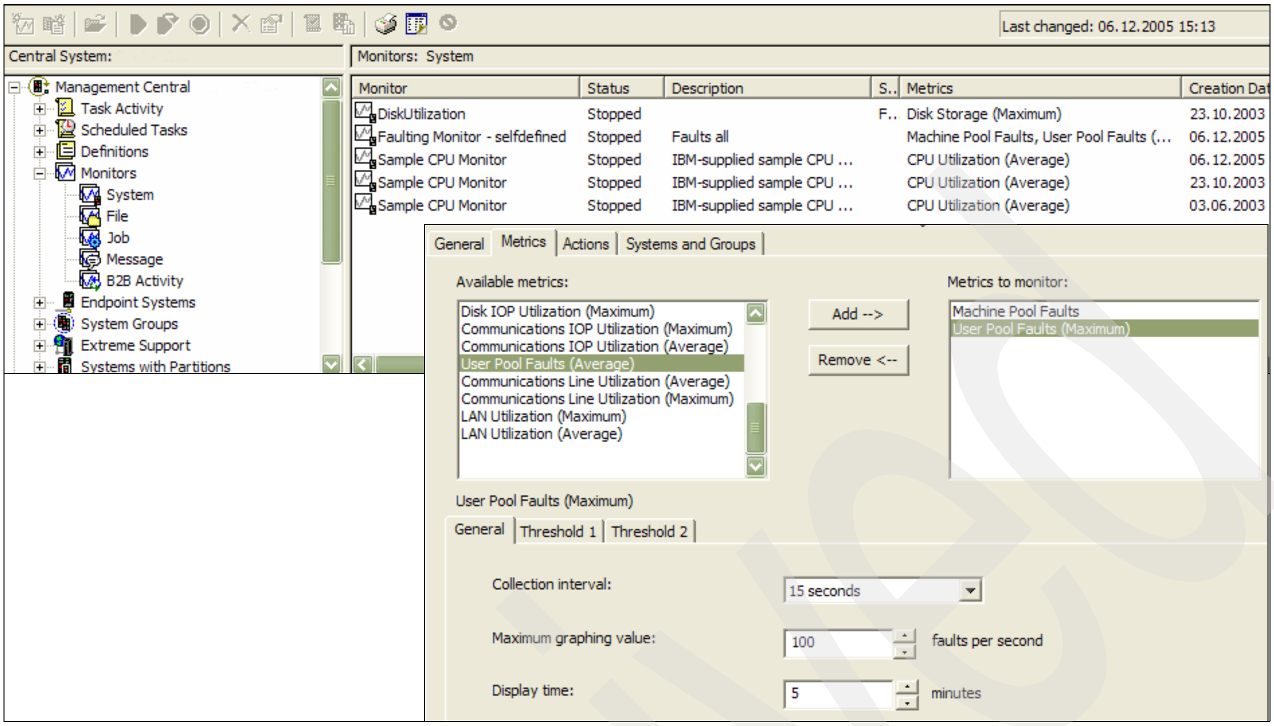

Figure 25-19 iSeries Navigator - define a monitor

For more information, search for **iSeries Navigator** at the following Web site:

http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp

25.4.8 SAP transaction ST06

If the saposcol job is running, some of the System i real-time monitoring values show the ST06 SAP transaction. Refer to Figure 25-20. Select **Snapshot analysis** for a current view (a one minute interval) or select **Previous hours** to view the last 24 hours in intervals of one hour.

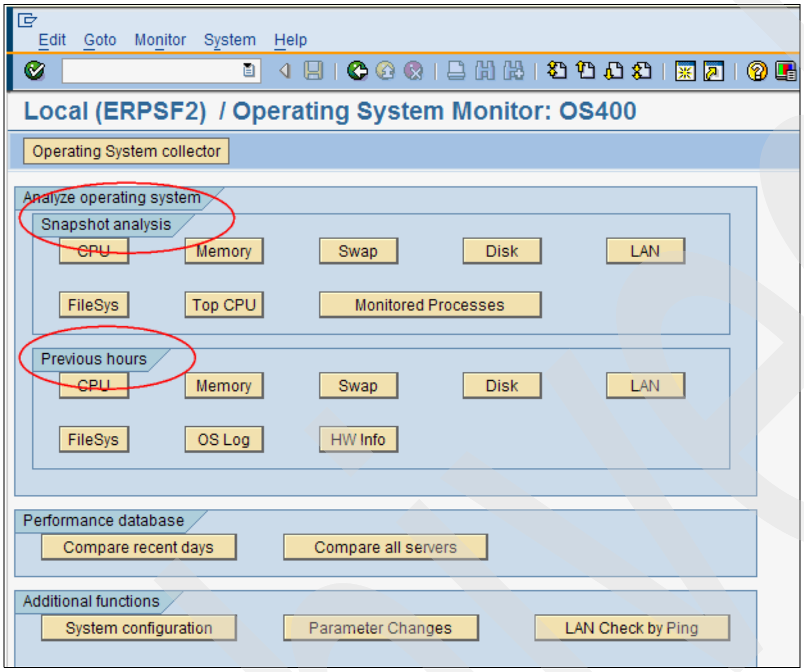

Figure 25-20 SAP transaction ST06

> **Note:** For recommendations, thresholds, and guidelines of System i components, refer to 26.1, "Tuning System i models" on page 532.

25.5 Monitoring SAP systems

Applications of mySAP Business Suite deliver many transactions, monitors, and analysis tools to observe the SAP system.

This section gives an overview about the choices that are of interest from the view of data processing staff with responsibility to a stable, free of disturbance execution of the SAP applications.

The transactions are usually related to each other, so that different approaches can lead to the identical destination. Determine their preferred use to supervise your SAP systems.

Monitor your system carefully.

25.5.1 CCMS

The Computing Center Management System (CCMS) is a standard part of the SAP system. The functionality and appearance differs somewhat from database to database, therefore, there are some tasks to perform outside the SAP system.

The previous monitoring and alert system in the CCMS is replaced by the monitoring architecture.

The CCMS analysis monitors perform the following functions:

- Check the system status and the operating modes.
- Detect and correct potential problems as quickly as possible.
- Diagnosis potential problems which could affect the SAP system, such as resource problems in a host or database system.
- Analyze and tune the SAP system and its environment (host and database system) to optimize the throughput of the SAP system.

These monitors are crucial in order to understand and evaluate the behavior of the SAP processing environment. The monitors provide you with the information required to fine-tune your SAP system in the case of poor performance values, and help ensure that your SAP installation runs efficiently.

Use the following monitors independently, or execute them as analysis methods in the alert monitor:

- Global Work Process Overview
- Workload Monitor
- Global Workload Monitor
- Operating System Monitor
- Operating System Collector
- SAP Buffer
- Database Monitor

The following CCMS transactions provide specific information for an installation running on i5/OS:

- Database monitoring and optimization (transaction DB4COCKPIT)
- Alert monitoring (transaction RZ20)
- Database check (transaction DB02)
- Database Performance Monitor (transaction ST04)
- Backup Overview (transaction DB12)

> **Note:** Use the operating system for database backup, operating system optimization and database reorganization.

Figure 25-21 illustrates how to navigate in CCMS.

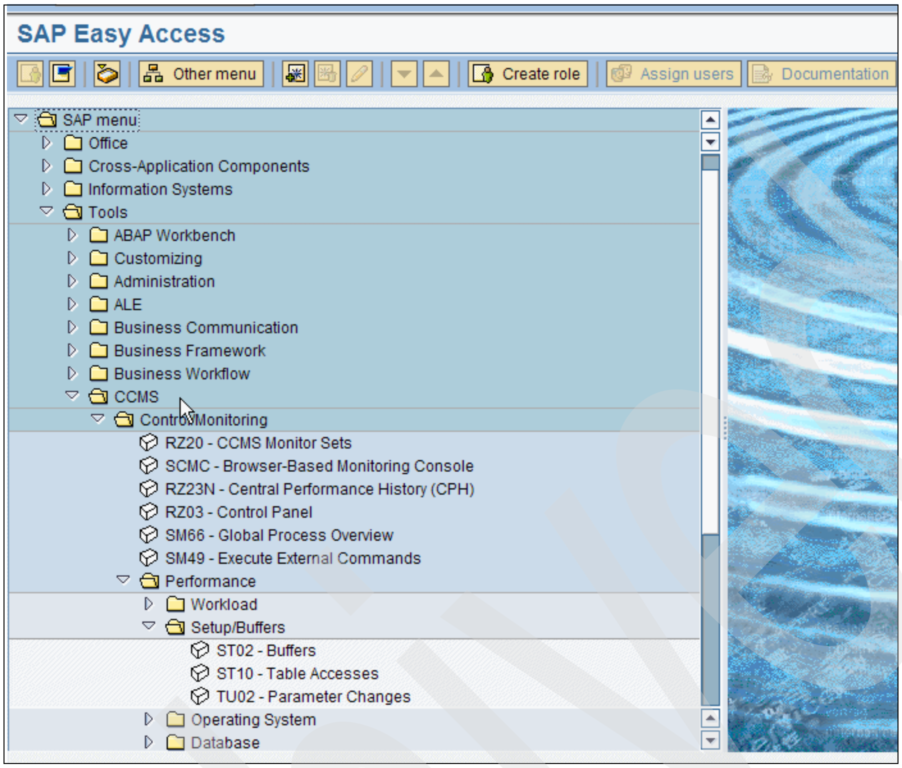

Figure 25-21 CCMS in the SAP easy access menu

1. Double-click **RZ20 CCMS Monitor Sets** and expand **SAP CCMS Monitor Templates**, as shown in Figure 25-22.

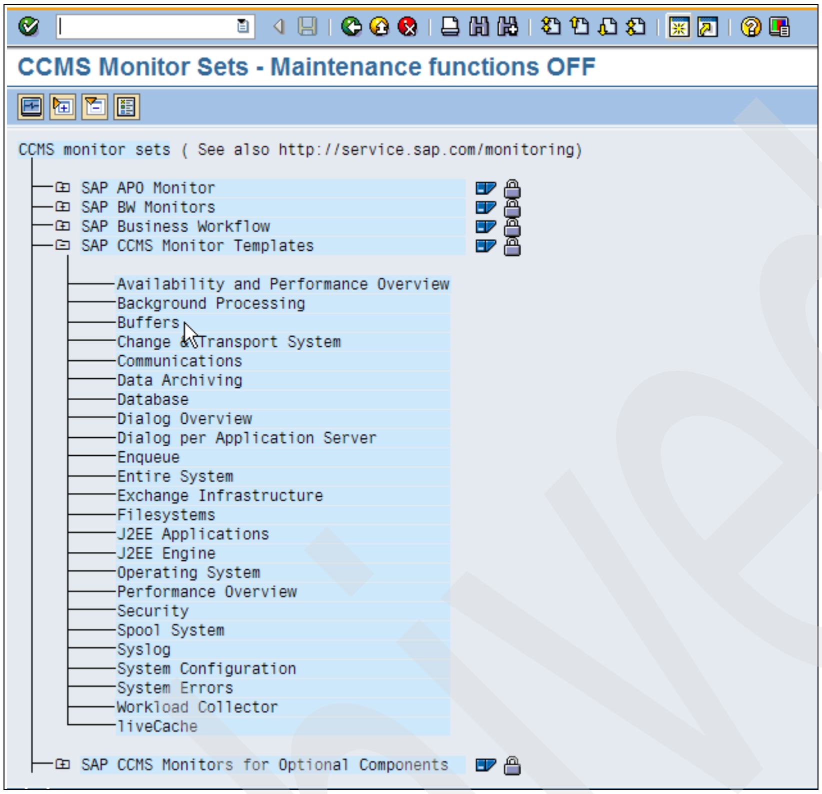

Figure 25-22 CCMS Monitor Sets (RZ20) - view CCMS Monitor Templates

2. Double-click **buffers** to view the current buffers status. Refer to Figure 25-23.

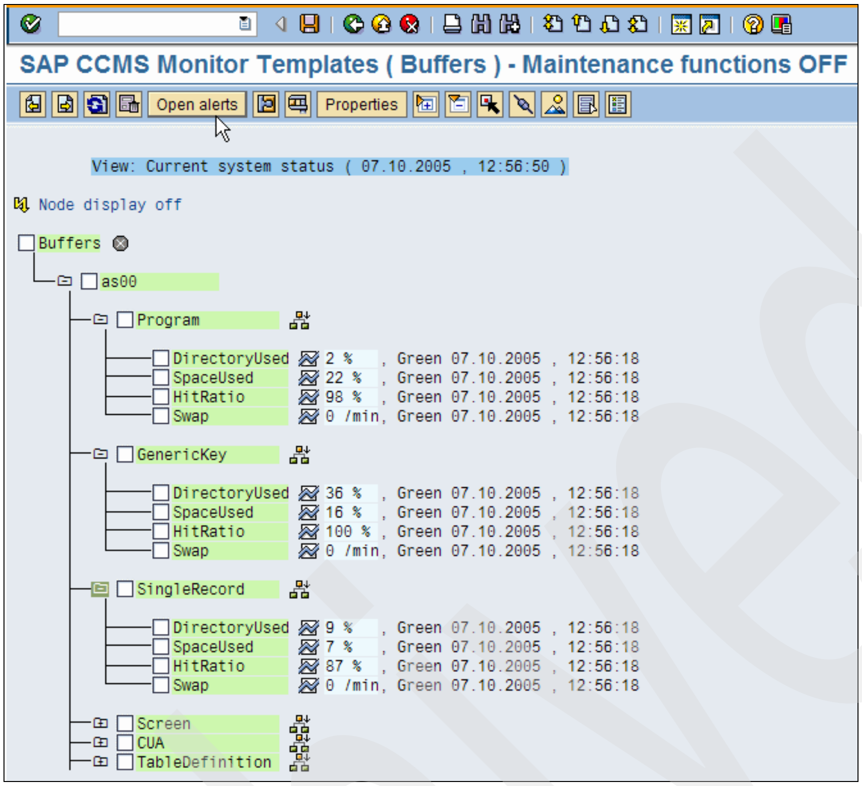

Figure 25-23 Buffers - view current system status

3. Click the **open alerts** button to change to view, as shown in Figure 25-24.

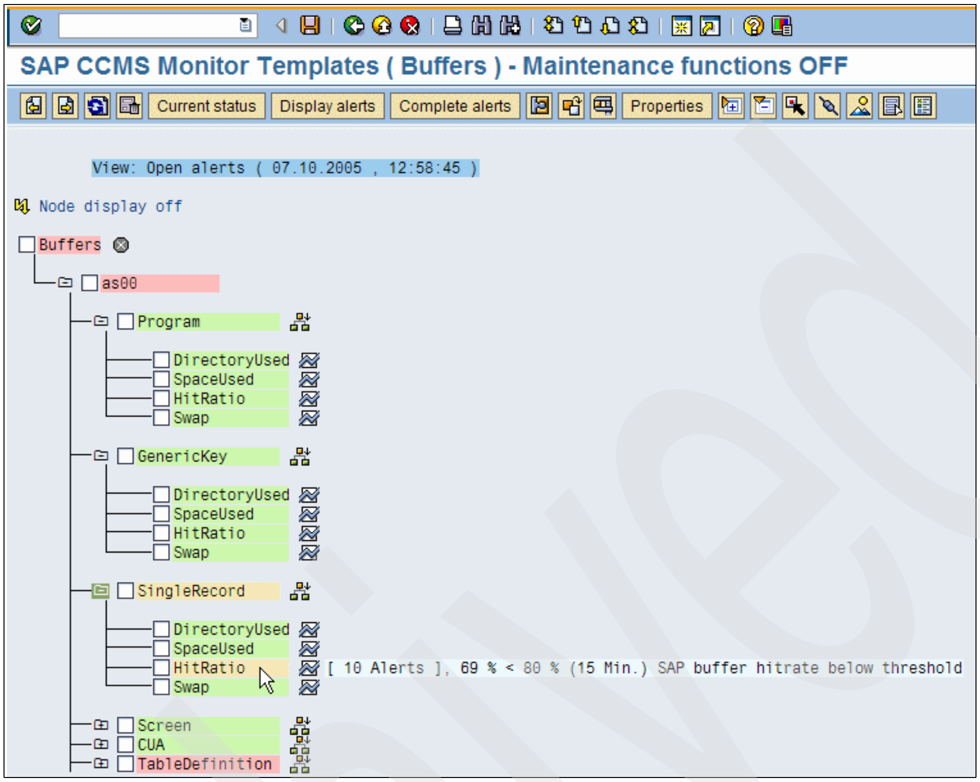

Figure 25-24 Buffers - view open alerts

Note the legend for colors and icons. Refer to Figure 25-25.

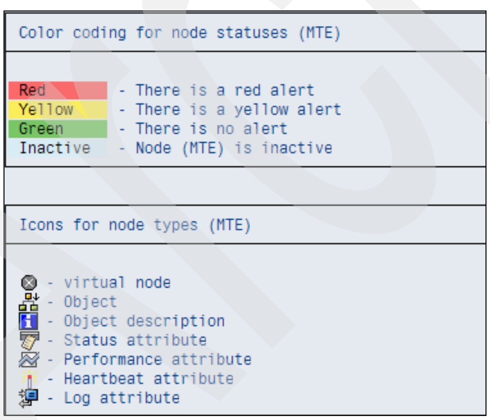

Figure 25-25 CCMS monitor legend

4. Double-click the yellow **HitRatio** entry to see a detailed alert display, as shown in Figure 25-26.

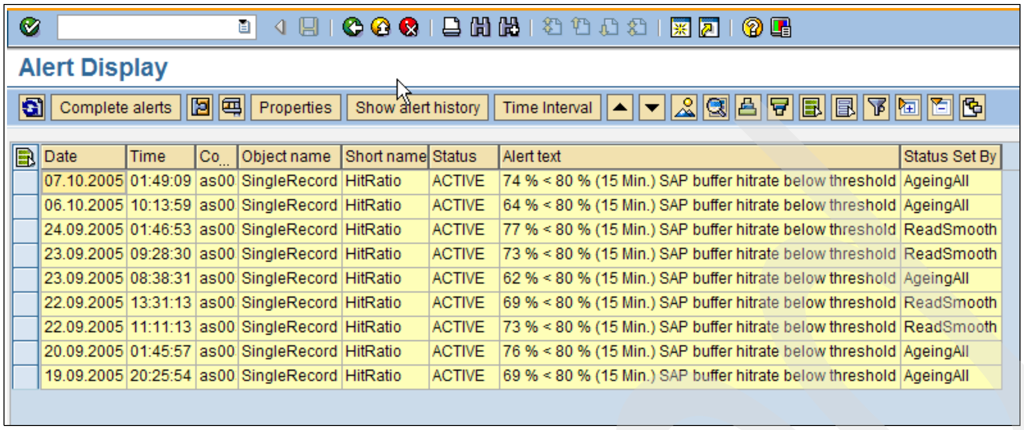

Figure 25-26 Buffers - single record - alert display

There were nine alerts in the past. Double-click any row to see a detailed view of an alert. Refer to Figure 25-27.

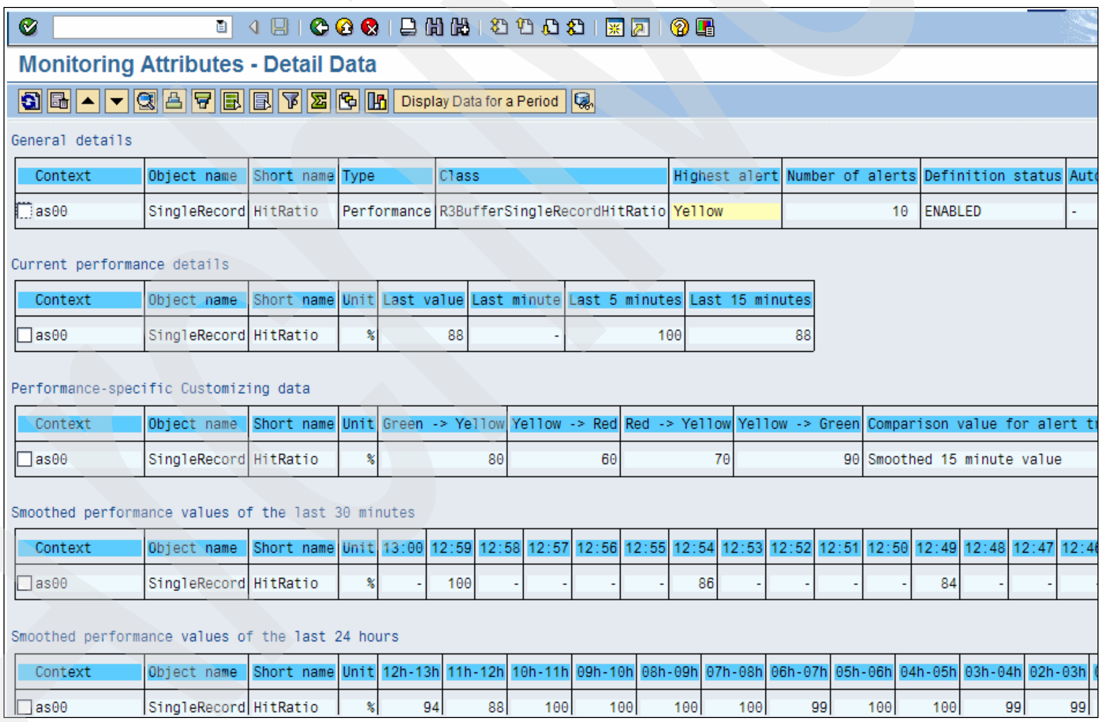

Figure 25-27 Buffers - single record - alert - detailed view

Present alerts are shown in Figure 25-23 on page 489. Figure 25-24 on page 490 shows alerts from the past. Double-click any row in the buffers current view, as shown in

Figure 25-23 on page 489, to see a current overview of buffers state, as shown in Figure 25-28. No alerts are currently shown.

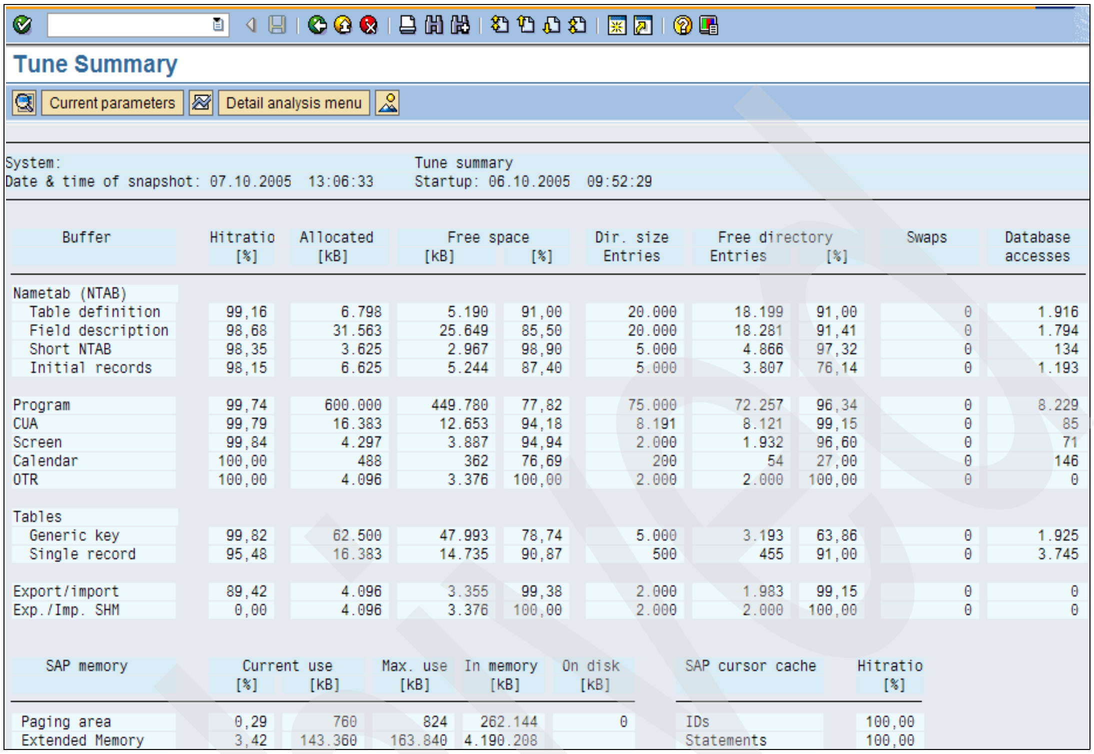

System:
Date & time of snapshot: 07.10.2005 13:06:33

Tune summary
Startup: 06.10.2005 09:52:29

Buffer	Hitratio [%]	Allocated [kB]	Free space [kB]	Free space [%]	Dir. size Entries	Free directory Entries	Free directory [%]	Swaps	Database accesses
Nametab (NTAB)									
Table definition	99,16	6.798	5.190	91,00	20.000	18.199	91,00	0	1.916
Field description	98,68	31.563	25.649	85,50	20.000	18.281	91,41	0	1.794
Short NTAB	98,35	3.625	2.967	98,90	5.000	4.866	97,32	0	134
Initial records	98,15	6.625	5.244	87,40	5.000	3.807	76,14	0	1.193
Program	99,74	600.000	449.780	77,82	75.000	72.257	96,34	0	8.229
CUA	99,79	16.383	12.653	94,18	8.191	8.121	99,15	0	85
Screen	99,84	4.297	3.887	94,94	2.000	1.932	96,60	0	71
Calendar	100,00	488	362	76,69	200	54	27,00	0	146
OTR	100,00	4.096	3.376	100,00	2.000	2.000	100,00	0	0
Tables									
Generic key	99,82	62.500	47.993	78,74	5.000	3.193	63,86	0	1.925
Single record	95,48	16.383	14.735	90,87	500	455	91,00	0	3.745
Export/import	89,42	4.096	3.355	99,38	2.000	1.983	99,15	0	0
Exp./Imp. SHM	0,00	4.096	3.376	100,00	2.000	2.000	100,00	0	0

SAP memory	Current use [%]	Current use [kB]	Max. use [kB]	In memory [kB]	On disk [kB]
Paging area	0,29	760	824	262.144	0
Extended Memory	3,42	143.360	163.840	4.190.208	

SAP cursor cache	Hitratio [%]
IDs	100,00
Statements	100,00

Figure 25-28 Buffers - current view

25.5.2 Solution Manager

The SAP Solution Manager is a dedicated SAP system and a central point in the SAP system landscape. The SAP Solution Manager is designed to support you throughout the entire life cycle of your solutions, from the Business Blueprint, through configuration, to production operation. It provides central access to tools, methods, and pre-configured content that you can use during the evaluation, implementation, and productive operation of your systems. Therefore, the SAP Solution Manager provides functions for project leaders, application, development and technical consultants, as well as operational and service staff.

The SAP Solution Manager provides functions to:

- Create and centrally manage the system landscape to use in your implementation and template projects.
- Monitor change requests.
- Distribute and examine the system customizing.
- Support operational tasks (support desk, services, solution monitoring.)

Adjust the SAP Solution Manager roles to suit your requirements. See 16.5, “SAP Solution Manager” on page 276 for more information about the SAP Solution Manager.

The monitoring and performance aspects of the SAP Solution Manager provide important functions for solution monitoring, central system administration, real-time system and business process monitoring, and analysis of your system landscape with service level reporting.

The system and business process monitoring (including interface monitoring) in the SAP Solution Manager each use the CCMS (transaction RZ20) architecture. This means that system alerts which occur in the local CCMS are passed between the SAP Solution Manager and the satellites via Remote Function Call (RFC) connections. You can also handle the alerts centrally without having to go to the local CCMS of the satellite systems.

You can see the alerts from several systems in a solution landscape in the SAP Solution Manager in a graphical overview, in contrast to the local CCMS of the SAP Solution Manager satellite system. This is the view of a central CCMS (CEN). You can also monitor non-SAP systems in a central CCMS (CEN) of a satellite system.

A view for Landscape Components of Solution Manager is shown in Figure 25-29.

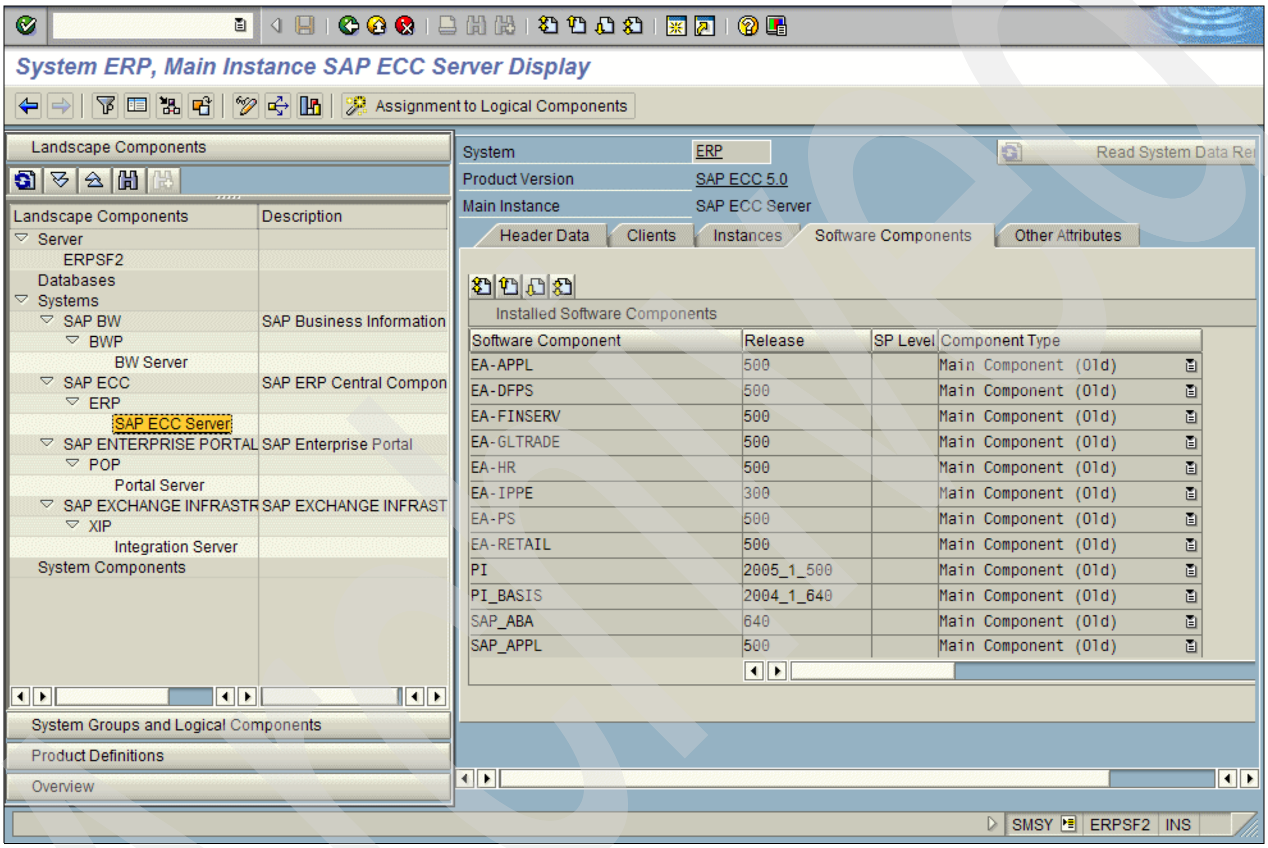

Figure 25-29 Solution Manager - Landscape Components

The system monitoring in the SAP Solution Manager is, as described above, like a system-wide central CCMS.

Figure 25-30 on page 494 illustrates a scenario of monitoring alerts of a J2EE Basis 6.40 (for example, EP6.0). The J2EE 6.40 can be connected to a central CCMS in a Central Monitoring

System (CEN) 6.40. You can therefore also monitor alerts of these components in SAP Solution Manager 3.20 with Basis 6.20.

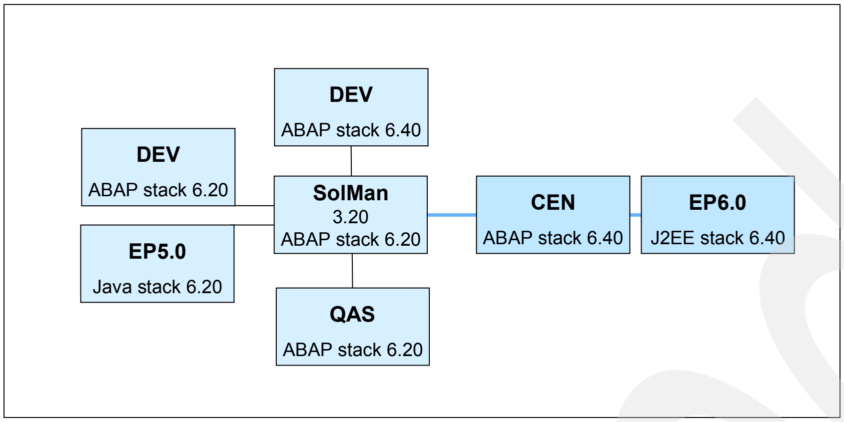

Figure 25-30 SAP Solution Manager as a system-wide central CCMS

Refer to 16.5, "SAP Solution Manager" on page 276 for further information. You can find a detailed description of SAP Solution Manager in the SAP online library at:

http://help.sap.com/saphelp_sm32/helpdata/en/index.htm

25.5.3 Alert Monitor (RZ20)

The monitoring architecture is a part of SAP NetWeaver. It centrally monitors any information technology (IT) environments.

Alerts form a central element of monitoring. They quickly and reliably report errors, such as values that exceed or fall below a particular threshold value, or an IT component that has been inactive for a defined period of time. These alerts are displayed in the Alert Monitor. This reduces the workload for the system administration, since they only need to watch the error messages instead of monitoring endless system data. Customize the Alert Monitor and extend the monitoring architecture to include SAP and non-SAP components.

The Alert Monitor is the central tool with which you can efficiently administer and monitor distributed SAP systems or client/server systems. The Alert Monitor displays problems quickly and reliably.

Whether you use a starting point for monitoring alerts, the CCMS alert component, or directly use the transaction code RZ20, the first view you see is as shown in Figure 25-22 on page 488. From there, navigate as described in 25.5.1, "CCMS" on page 485.

You must perform various configuration steps yourself if you want to use the Alert Monitor for central monitoring (that is, if you want to monitor the systems of your IT landscape from a central monitoring system).

To configure the monitoring architecture, refer to the following Web site:

http://help.sap.com/saphelp_erp2004/helpdata/en/02/411841a66af223e10000000a155106/frameset.htm

25.5.4 Transaction SSAA

The System Administration Assistant (SSAA) supports the system administrator in performing recurring tasks. SSAA provides many functions, including monitoring and tuning tasks. There is a highlighted list of daily, weekly and occasional activities.

The tasks include:

- Release and monitor transports in the Change and Transport System
- Monitor the system (SAP system and database)
- Monitor performance
- User administration
- Client administration
- Preupgrade and post-upgrade activities
- Post-installation activities
- Integration

Figure 25-31 illustrates the structure of the basic parts of the SSAA transaction.

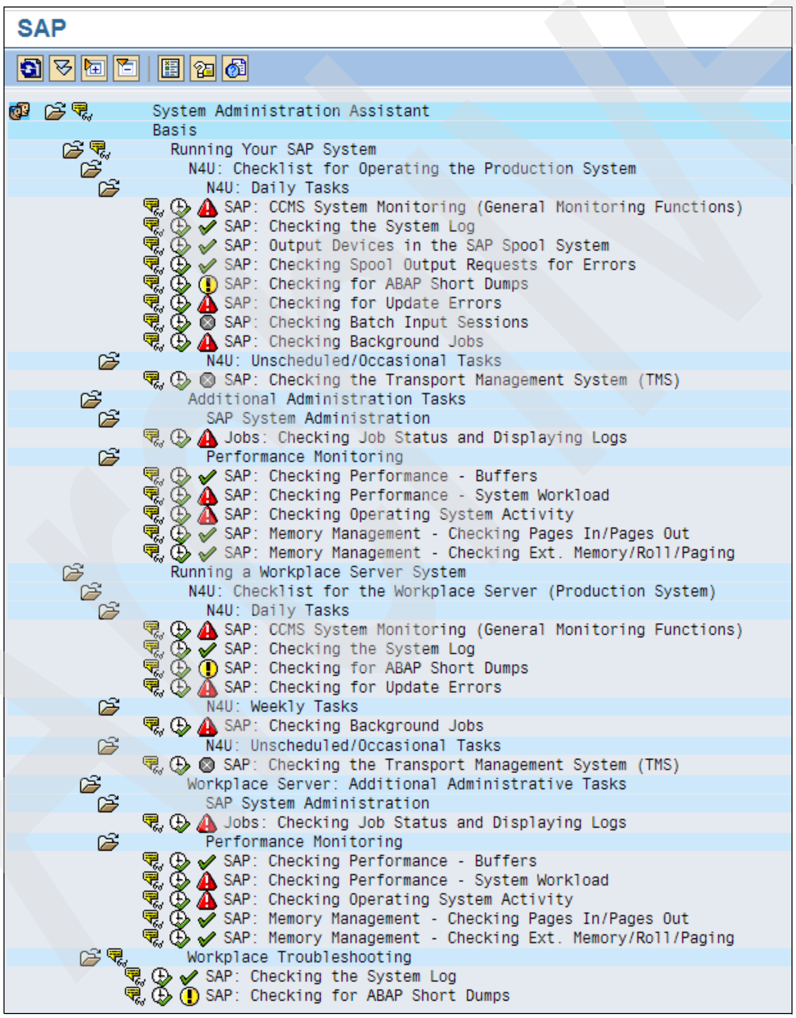

Figure 25-31 Transaction SSAA - alert view

25.5.5 Global (System Wide) Work Process Overview (SM66)

With the SM66 Global (System Wide) Work Process Overview transaction you can quickly investigate the potential cause of a system performance problem by checking the work process load. Use transaction **ST66** to monitor the work process load on all active instances across the system. For example, you can identify locks in the database (lock waits).

Using the Global Work Process Overview screen, you can see:

- The status of each application server
- The reason why it is not running
- Whether it has been restarted
- The CPU and request runtime
- The user who has logged on and the client that they logged on to
- The program that is running

Refer to Figure 25-32 to Figure 25-34 on page 497 for details.

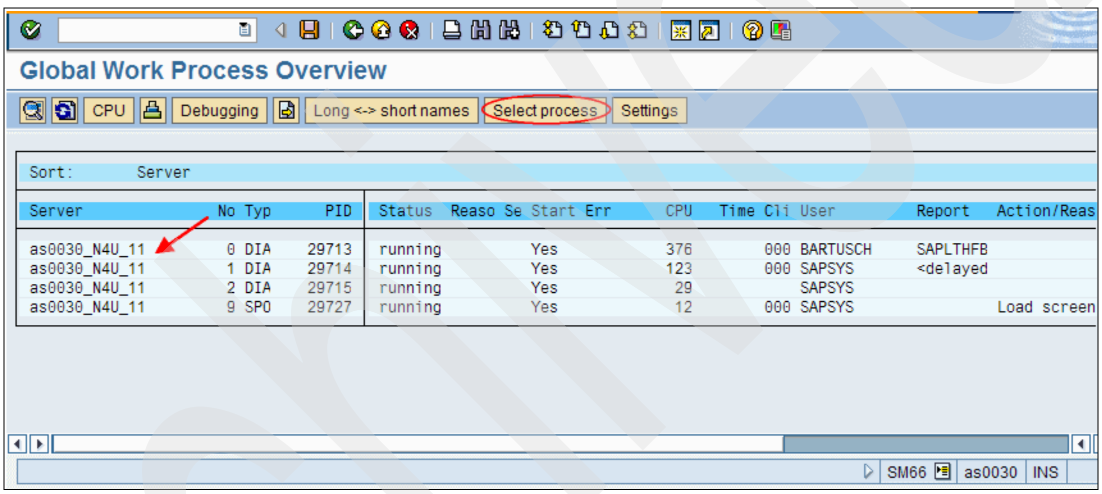

Figure 25-32 Global Work Process Overview (SM66)

Press the **Select process** button. You can filter the work processes of the view. For example, you can find all work processes for users beginning with the letter B with more than 5s CPU consumption which are in a wait state. Refer to Figure 25-33.

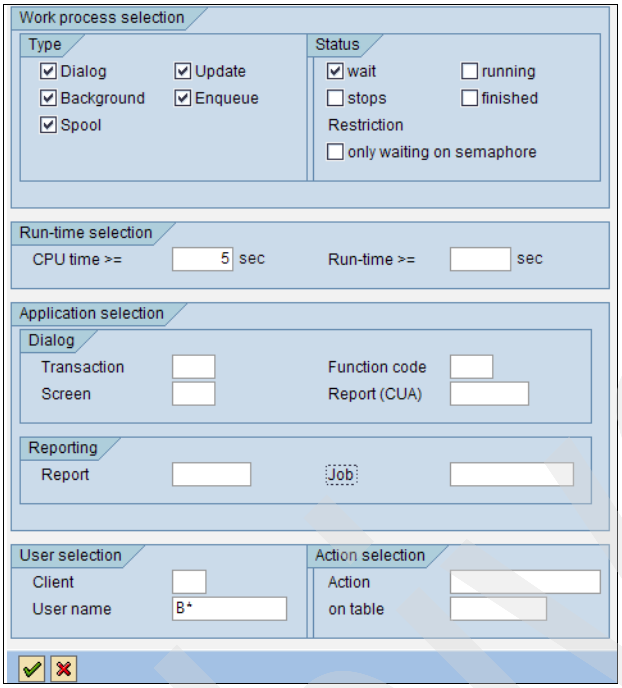

Figure 25-33 SM66 - Work process selection

Double-click a work process to view details. Refer to Figure 25-34.

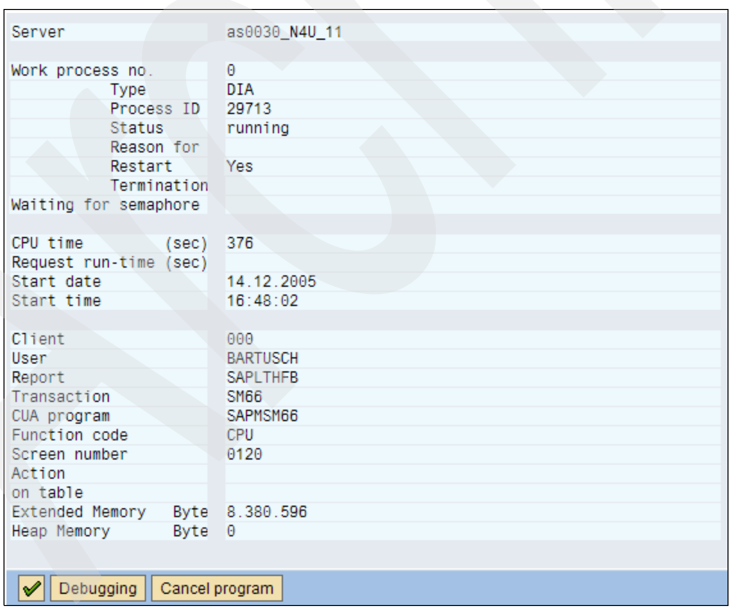

Figure 25-34 SM66 - Work process detail

For further information you can debug the program or cancel it from here.

25.5.6 Process overview (SM50)

With the SM50 transaction you see an overview of the work processes of an instance. Use the process analysis to determine details, such as:

- Type, process ID (PID) and status of work processes
- Status, reason
- User and client
- Accumulated CPU time of process, run time of task
- Report name
- DB table statistics

Figure 25-35 illustrates this.

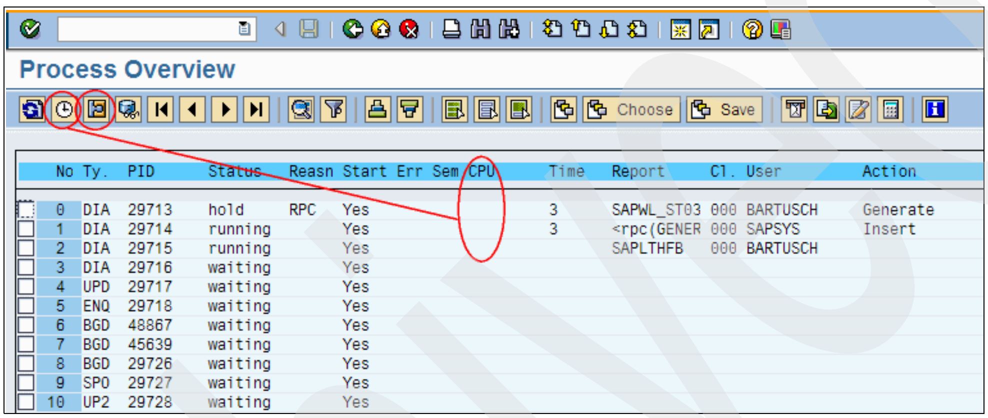

Figure 25-35 Process Overview (SM50)

Press the **clock** button to display the CPU time. Press the **details** button as shown in Figure 25-35 on page 498 to view details of a work process. Refer to Figure 25-36.

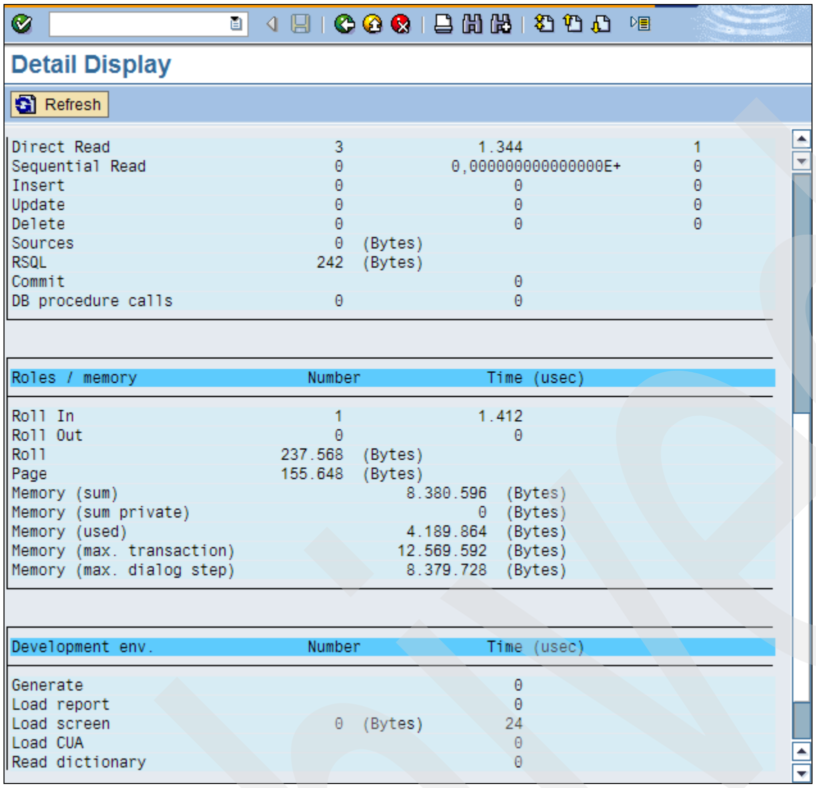

Figure 25-36 SM50 - Detail display

You can choose other tasks from the menu, for example, select **System** → **Utilities** → **Runtime analysis** as shown in Figure 25-37.

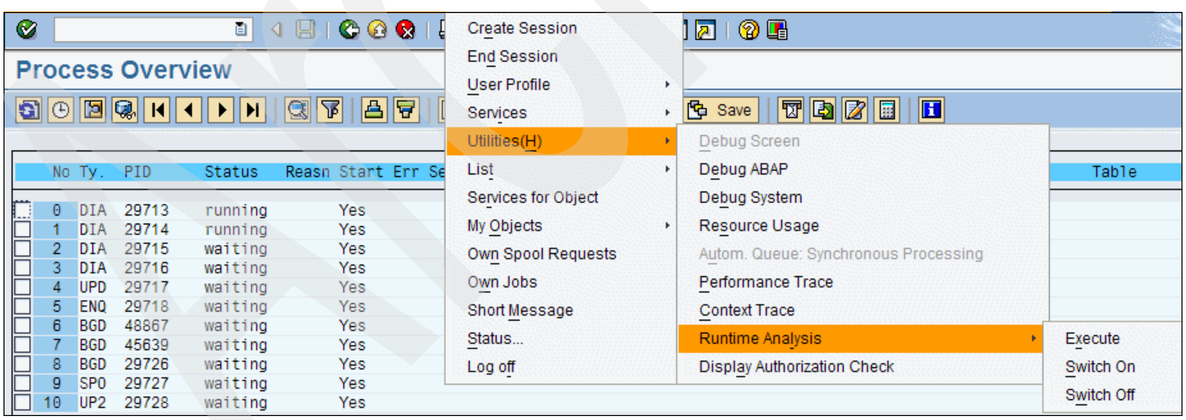

Figure 25-37 SM50 - Choose other tasks from Utilities

This takes you to transaction SE30 as shown in Figure 25-38 on page 500.

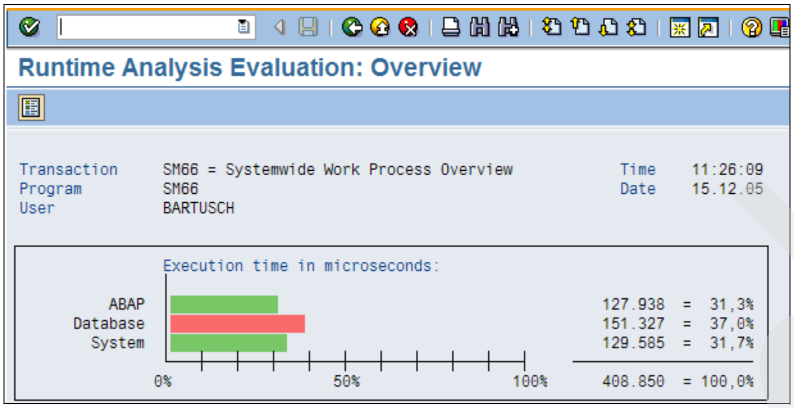

Figure 25-38 Example - Runtime analysis of a transaction

25.5.7 System Log - Analysis (SM21)

The System log collects system messages and classifies this messages by type. Check this log periodically to find problems which frequently occur. Transaction SM21 shows you this log. Initially, choose a time period and problem classes.

Select **System log** → **Choose** to switch between logs as shown in Figure 25-39.

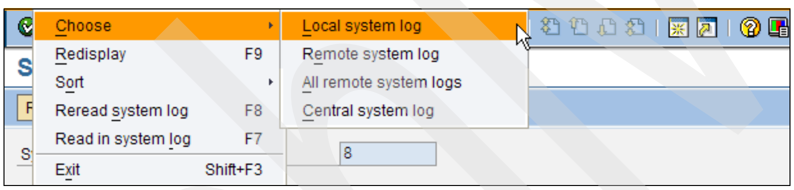

Figure 25-39 Choose a system log

A similar report is displayed as Figure 25-40 shows.

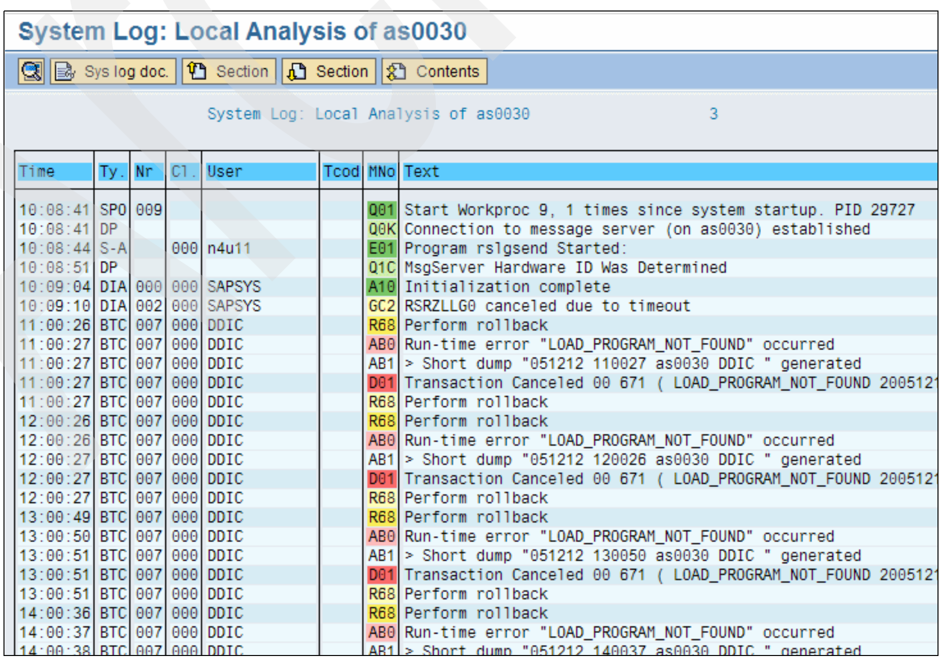

Figure 25-40 System Log - Analysis (SM21)

Double-click an entry to view the details as Figure 25-41 shows.

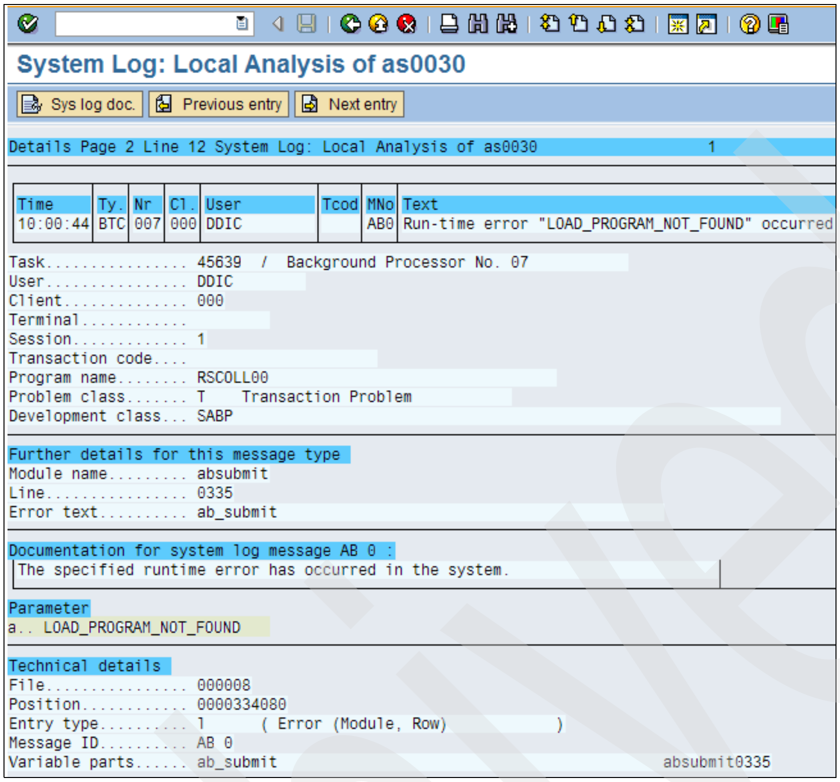

Figure 25-41 System log - detail

25.5.8 Logged on users (SM04)

Transaction ST04 allows you to view all the users logged on to your SAP system as Figure 25-42 shows.

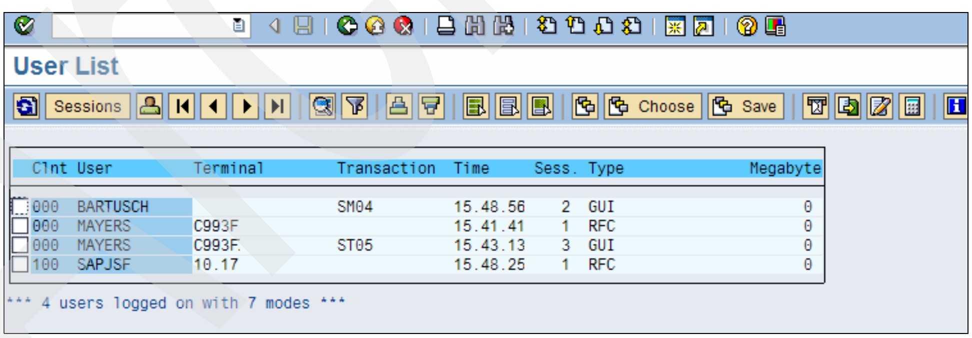

Figure 25-42 User list (ST04)

Select **Go to** → **Memory** to view the memory list as Figure 25-43 on page 502 shows.

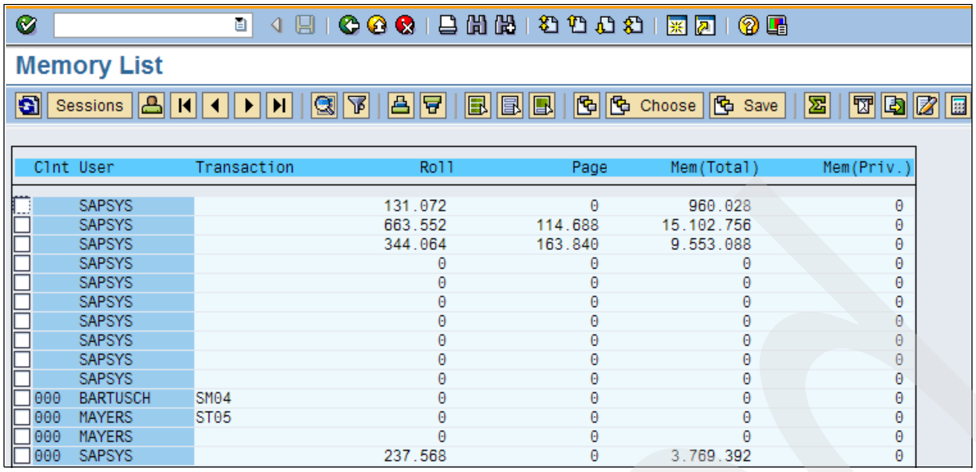

Figure 25-43 Memory list (ST04)

If any user allocated private memory, this means that the work process is reserved for processing the current user content until the content releases the work process again when the request ended. This impacts system performance.

25.5.9 Application Monitor(ST07)

The Application Monitor enables a view over system activities. The ST07 transaction includes different views for:

- SAP buffer
- Database accesses
- Database memory
- Response times
- Quantity structure
- History

Double-click an entry to see a more detailed view of several layers.

The default view shows you the user distribution as Figure 25-44 shows.

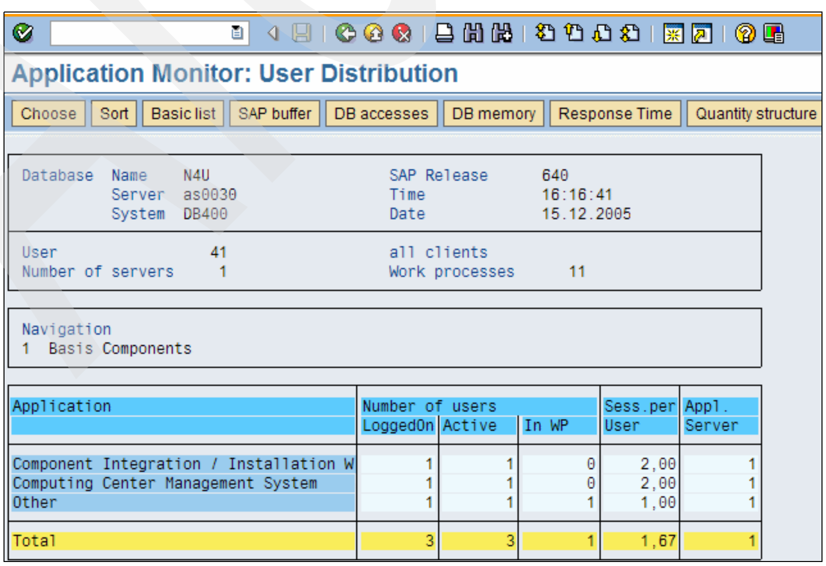

Figure 25-44 Application Monitor (ST07)

Select **Response time** in contrast to a general response time analysis to see transaction ST07, an application-based response time analysis as Figure 25-45 shows.

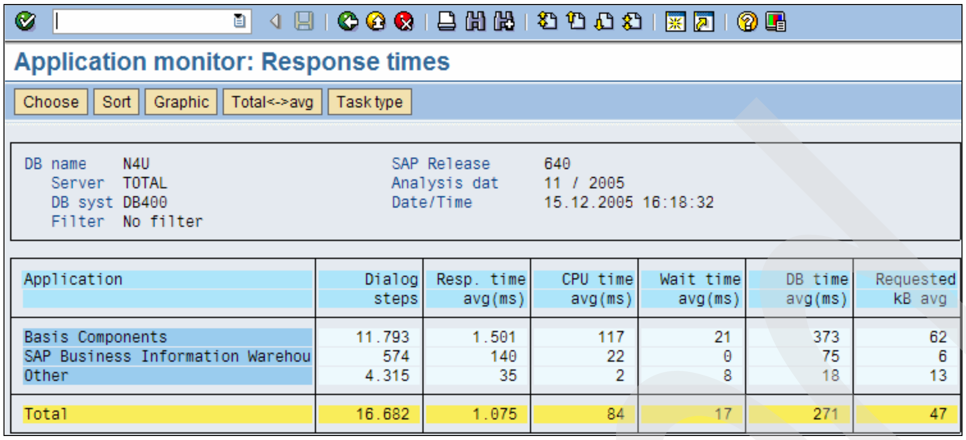

Figure 25-45 Application Monitor - Response times

Select **Database accesses** to determine the distribution of the database workload as Figure 25-46 shows.

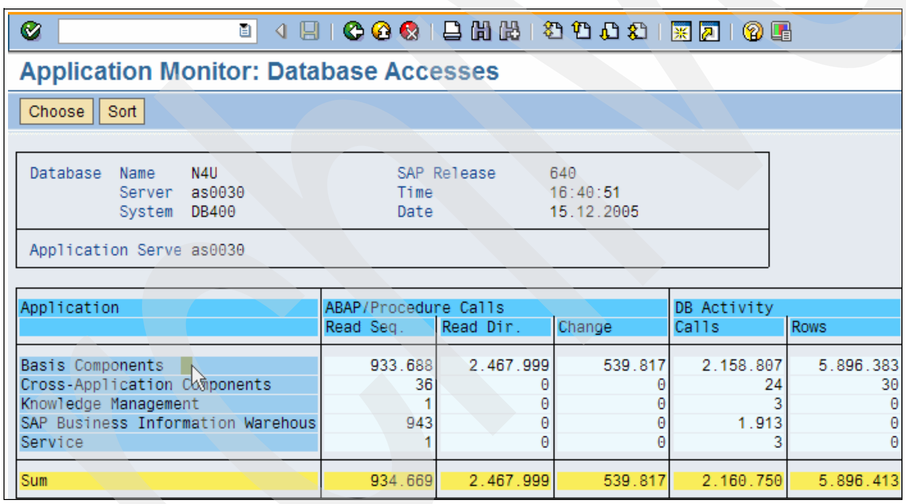

Figure 25-46 Application Monitor - Database accesses

Double-click an entry of the column `Application` to see a detailed view, similar to that shown in Figure 25-47 on page 504.

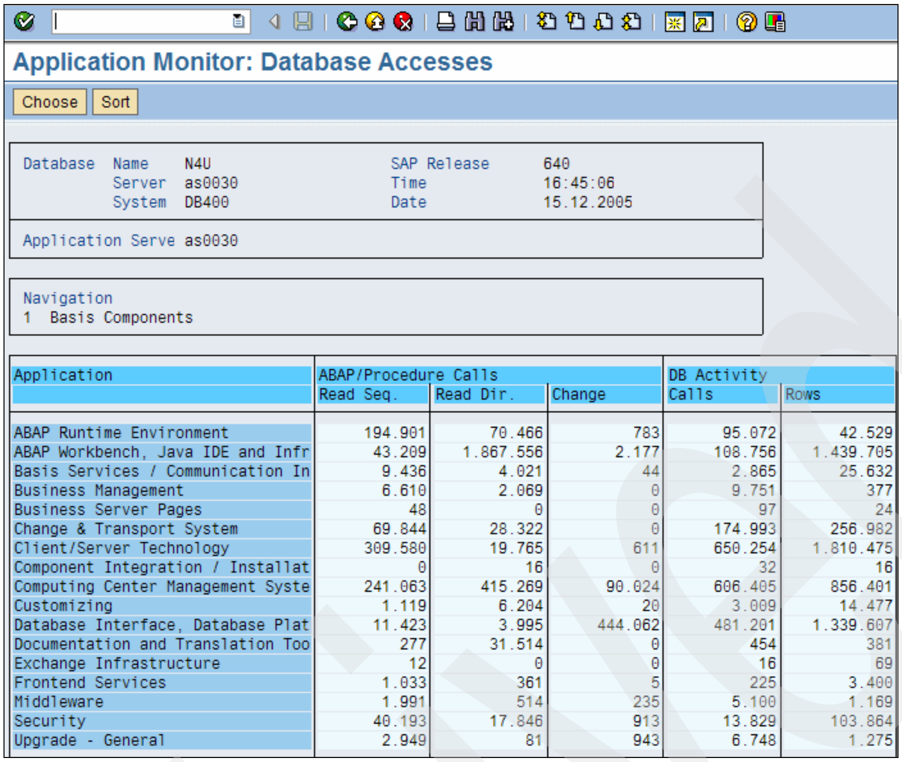

Application Monitor: Database Accesses

Choose | Sort

Database Name N4U — SAP Release 640
Server as0030 — Time 16:45:06
System DB400 — Date 15.12.2005

Application Serve as0030

Navigation
1 Basis Components

Application	ABAP/Procedure Calls			DB Activity	
	Read Seq.	Read Dir.	Change	Calls	Rows
ABAP Runtime Environment	194.901	70.466	783	95.072	42.529
ABAP Workbench, Java IDE and Infr	43.209	1.867.556	2.177	108.756	1.439.705
Basis Services / Communication In	9.436	4.021	44	2.865	25.632
Business Management	6.610	2.069	0	9.751	377
Business Server Pages	48	0	0	97	24
Change & Transport System	69.844	28.322	0	174.993	256.982
Client/Server Technology	309.580	19.765	611	650.254	1.810.475
Component Integration / Installat	0	16	0	32	16
Computing Center Management Syste	241.063	415.269	90.024	606.405	856.401
Customizing	1.119	6.204	20	3.009	14.477
Database Interface, Database Plat	11.423	3.995	444.062	481.201	1.339.607
Documentation and Translation Too	277	31.514	0	454	381
Exchange Infrastructure	12	0	0	16	69
Frontend Services	1.033	361	5	225	3.400
Middleware	1.991	514	235	5.100	1.169
Security	40.193	17.846	913	13.829	103.864
Upgrade - General	2.949	81	943	6.748	1.275

Figure 25-47 Application Monitor - Database accesses - detailed

If you choose **SAP Buffer** to find buffer intensive application parts, compare this view with the general buffer analysis with transaction ST02 in this chapter as Figure 25-48 shows.

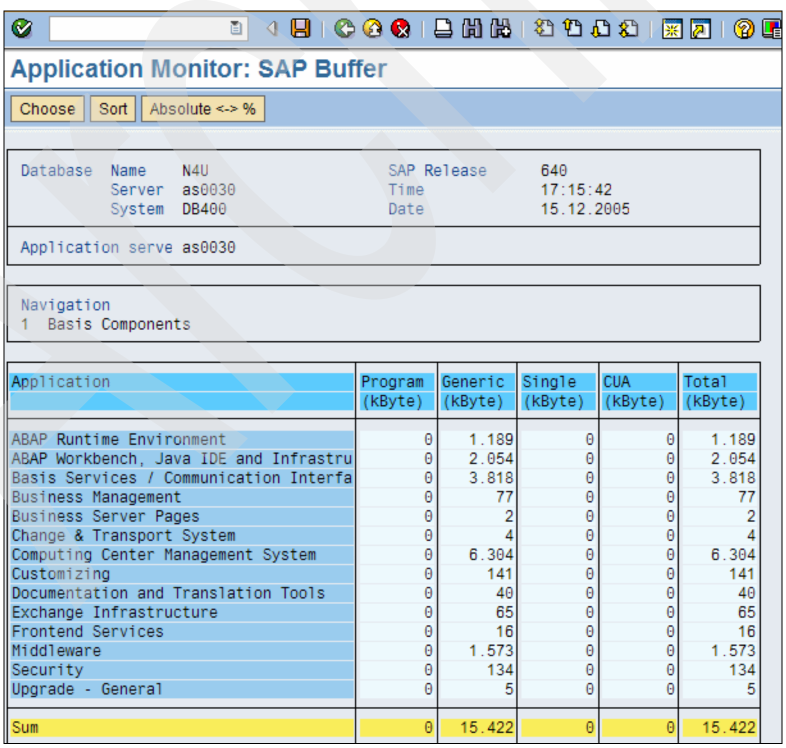

Application Monitor: SAP Buffer

Choose | Sort | Absolute <-> %

Database Name N4U — SAP Release 640
Server as0030 — Time 17:15:42
System DB400 — Date 15.12.2005

Application serve as0030

Navigation
1 Basis Components

Application	Program (kByte)	Generic (kByte)	Single (kByte)	CUA (kByte)	Total (kByte)
ABAP Runtime Environment	0	1.189	0	0	1.189
ABAP Workbench, Java IDE and Infrastru	0	2.054	0	0	2.054
Basis Services / Communication Interfa	0	3.818	0	0	3.818
Business Management	0	77	0	0	77
Business Server Pages	0	2	0	0	2
Change & Transport System	0	4	0	0	4
Computing Center Management System	0	6.304	0	0	6.304
Customizing	0	141	0	0	141
Documentation and Translation Tools	0	40	0	0	40
Exchange Infrastructure	0	65	0	0	65
Frontend Services	0	16	0	0	16
Middleware	0	1.573	0	0	1.573
Security	0	134	0	0	134
Upgrade - General	0	5	0	0	5
Sum	0	15.422	0	0	15.422

Figure 25-48 Application Monitor - SAP buffer

25.5.10 Workload analysis (ST03)

The ST03 SAP transaction helps you to review the workload on your SAP system. Over time the workload monitor ST03 has been enlarged, for example, with more monitors and functions, formatting options, and also simplifies the use and download of data. Refer to Figure 25-49 on page 506.

Note: In earlier releases, the ST03 transaction was the old-style workload monitor. In SAP release 4.6, the old-styled ST03 transaction and a new transaction ST03N are present. In newer SAP releases, there is only the new-styled workload monitor called transaction ST03 present.

Use the workload monitor to analyze statistical data from the SAP kernel. When analyzing the performance of a system, normally, you start by analyzing the workload overview. For example, you can display the total workload for all instances and compare it with another over specific periods of time. You can determine the source of possible performance problems using the large number of analysis views and determined data.

You can use the workload monitor to display the following:

- Number of configured instances for each SAP system
- Number of users working on the different instances
- Distribution of response times
- Distribution of workload by transaction steps, transactions, packages, subapplications, and applications
- Transactions with the highest response time and database time
- Memory usage for each transaction or each user per dialog step
- Workload through RFC, listed by transactions, function modules, and destinations
- Number and volume of spool requests
- Statistics about response time distribution, with or without the graphical user interface (GUI) time
- Table accesses
- Workload and transactions used listed by users, payroll number, and client
- Workload generated by requests from external systems

For all of this data, you can do the following:

- Display the data for a particular instance (not only the one to which you are logged on) or a total of all instances.
- Depending on your user mode, choose the time period for which you want to display the data between day, week and month (or determine the length of time yourself using the Last Minutes Load function).
- Display all or only certain task types for most analysis views

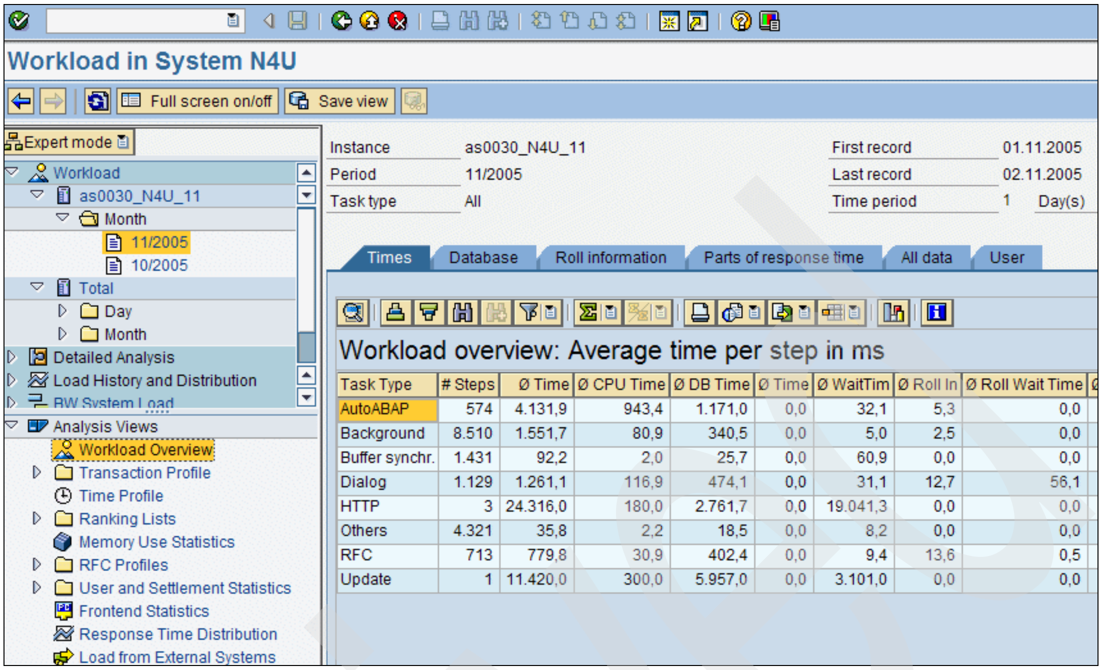

Figure 25-49 Workload Monitor (ST03)

Press the **modes** button (above left) to change user modes between Expert mode, Administrator and Service engineer as Figure 25-50 shows.

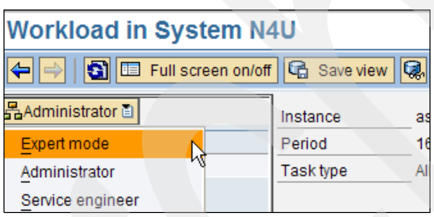

Figure 25-50 ST03 - choose a mode

You can set your preferred view as your initial view. Press the **Save view** button as Figure 25-51 shows. Press **F6** to save the current view as the initial view.

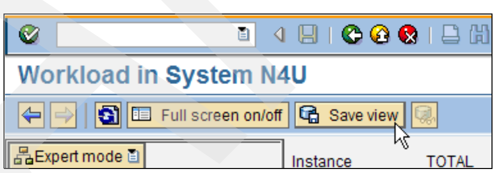

Figure 25-51 ST03 - save as initial view

In the upper left part of the navigation area, you can choose the workload for servers, consolidated workload for days, weeks or months, or a customized interval of the last number of hours, and setup the collection parameters, such as collecting file space, consolidation periods, how much history to collect and so on. Refer to Figure 25-52 on page 507.

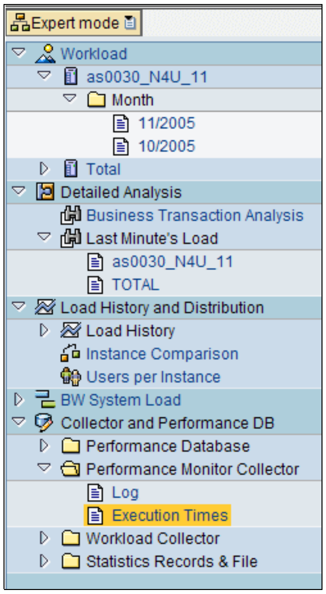

Figure 25-52 ST03 - upper navigation area

If you selected a workload, choose one of the analysis views in the upper navigation area. Refer to Figure 25-53. The selected workload for most analysis views can be changed for analysis.

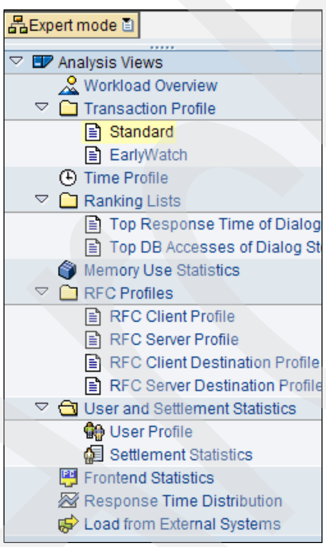

Figure 25-53 ST03 - lower Navigation area

The following are some examples of workload reports. Double-click a row to see a more detailed view. Refer to Figure 25-54 on page 508 to see a time profile analysis for dialog workload, with the detailed workload for an hour.

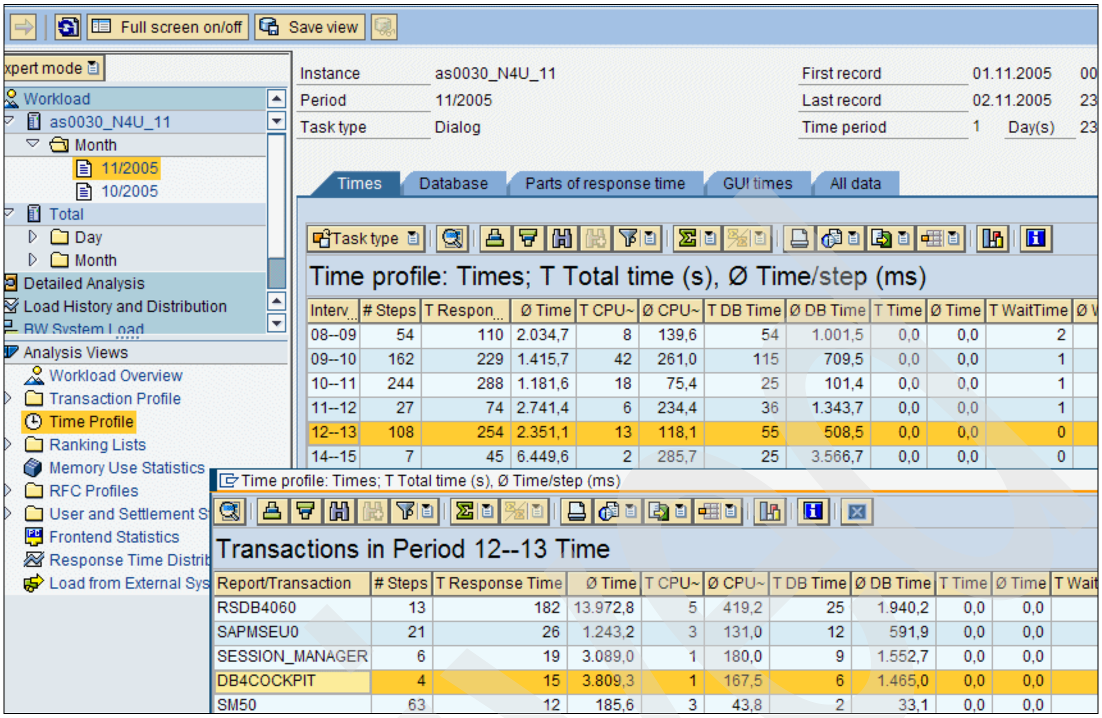

Figure 25-54 ST03 - dialog workload analysis by time

Refer to Figure 25-55 for a memory usage statistic of all work processes.

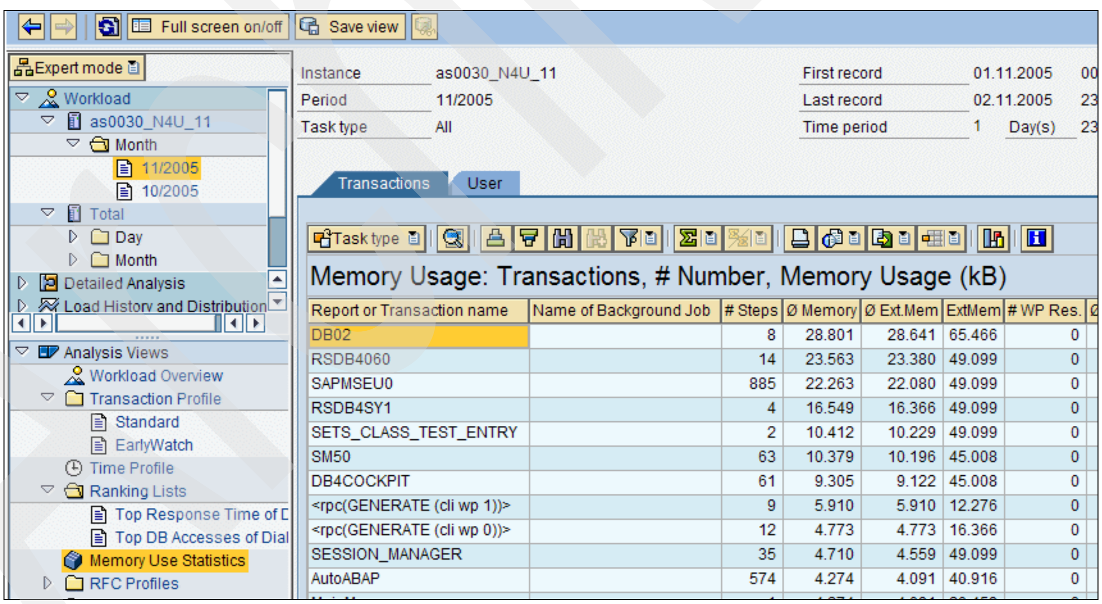

Figure 25-55 ST03 - memory usage statistic of all work processes

Refer to Figure 25-56 on page 509 for a percentage response time distribution statistic of all work processes.

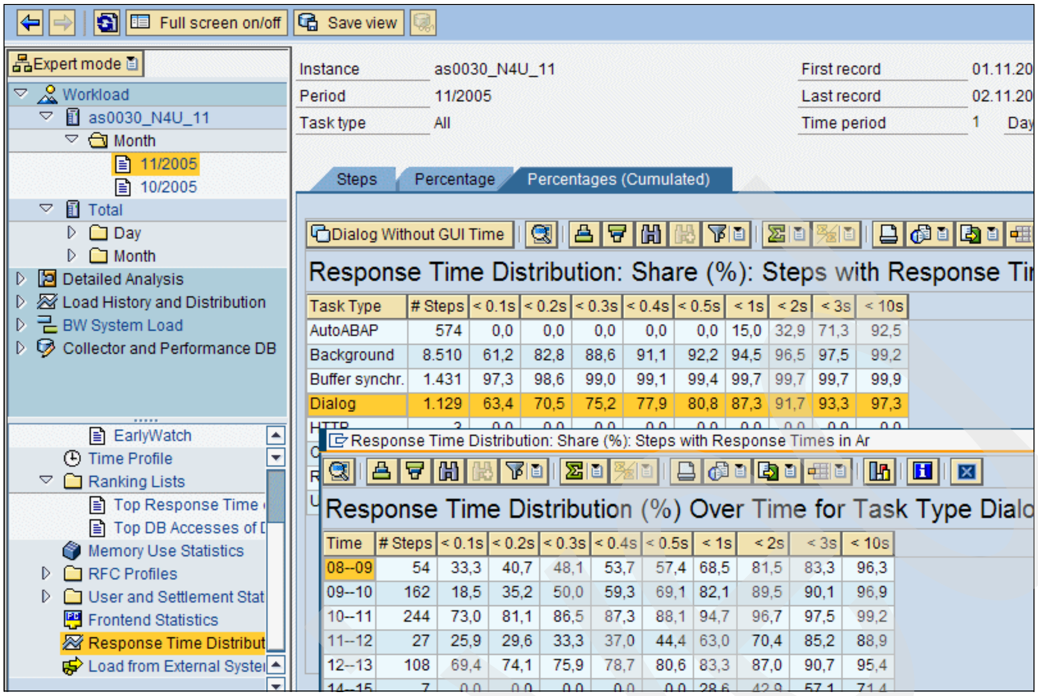

Figure 25-56 ST03 - Percentage Response Time distribution for all task types such as details

Refer to Figure 25-68 on page 519 for a list of the top response time dialog transactions, such as details of the first transaction.

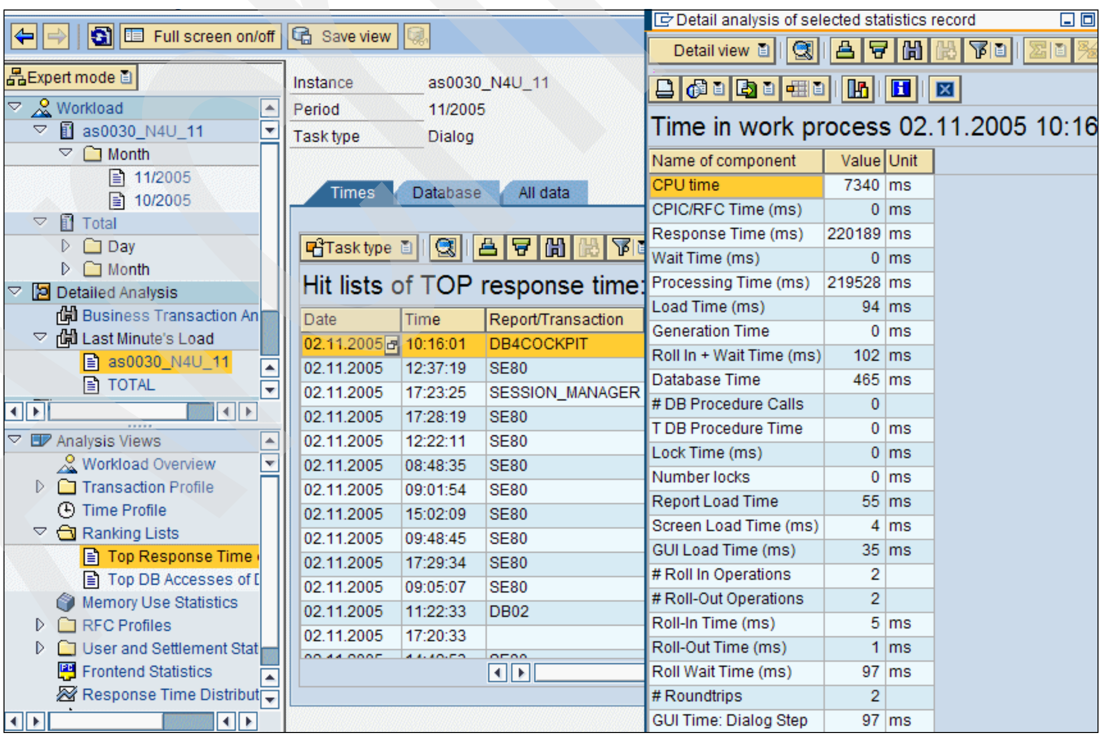

Figure 25-57 ST03 - top list of dialog transactions such as details for the first transaction

To create a last minutes work load analysis report, first set the filters and a minimum time period as shown in Figure 25-58 on page 510.

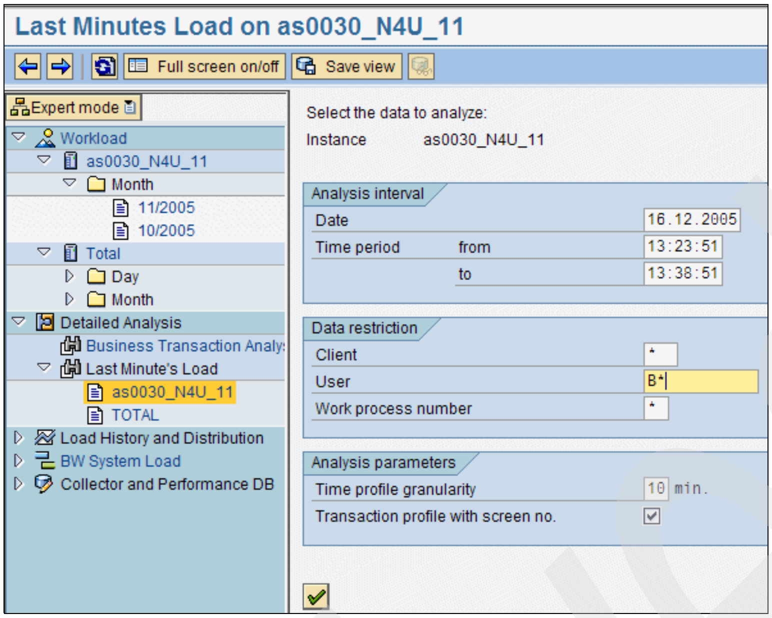

Figure 25-58 ST03 - filters to generate a last minutes workload analysis report

To use the sort function, first mark the column to be sorted, then press the **sort in descent order** button as Figure 25-59 shows.

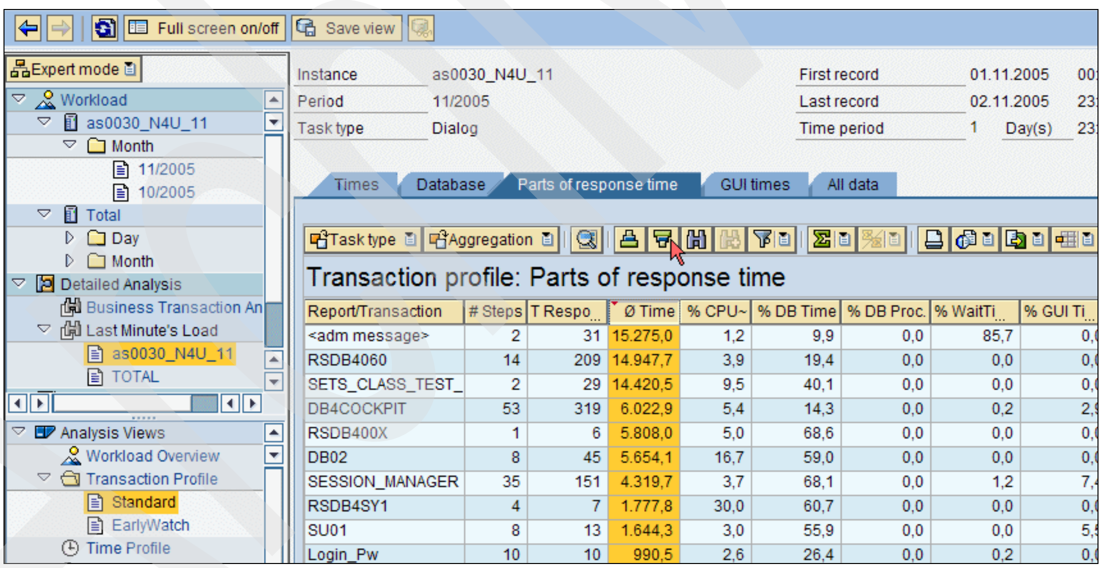

Figure 25-59 ST03 - format your report - sort in descent order

As Figure 25-60 on page 511 shows, you can export (download) your reports in different formats, such as Hypertext Markup Language (HTML), Extensible Markup Language (XML), unformatted or formatted text, spreadsheet formats and so on, to archive, send, publish, or download for further analysis, such as standard table calculation programs.

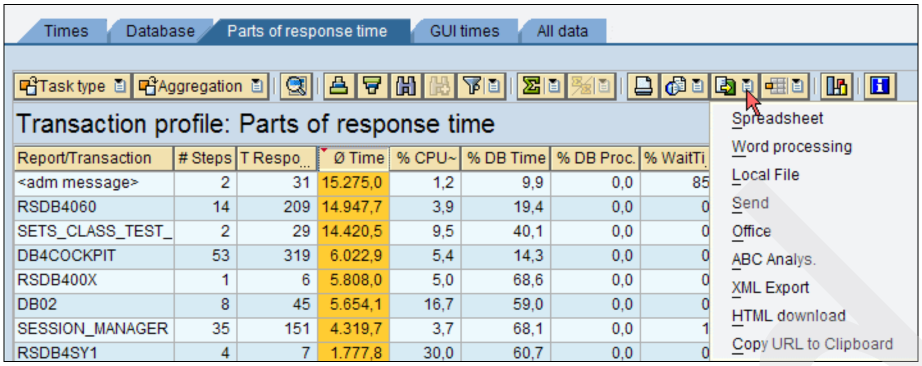

Figure 25-60 ST03 - export function

25.5.11 Tune summary and buffers (ST02)

The ST02 transaction provides information about SAP buffers as shown in Figure 25-61. You can monitor the values for the hit ratio and percentage of free directory entries or free space. If you find low hit ratios, or red fields for swaps, or both, you need to tune the buffers. Refer to “Buffers” on page 542 for recommendations.

> **Note:** For an application server on a 64-bit system, the value of the swaps should be zero as far as possible. Otherwise you should check and adapt the parameters for free space and the number of entries.

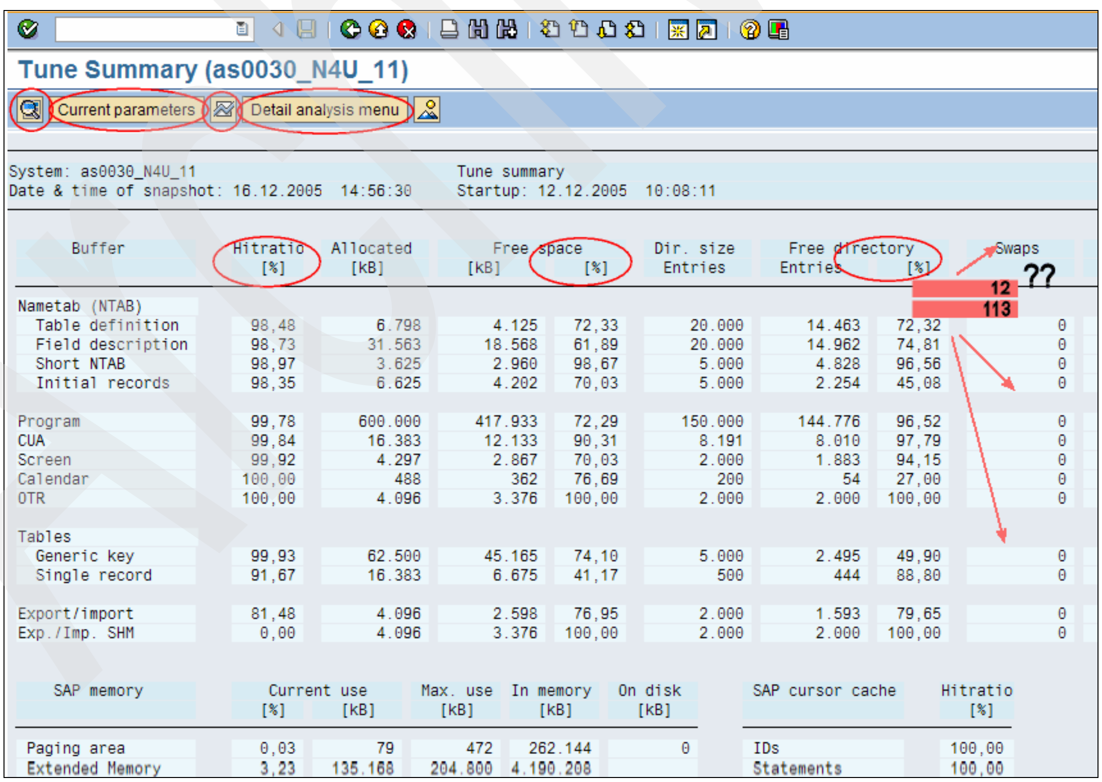

Figure 25-61 Tune summary and buffers (ST02)

You can also use the icons or buttons in the left upper area of the screen (from left to right) to view the following:

- Details
- Current parameters
- History
- The detail analysis menu

Mark a row with a parameter and select **choose** or press **F2** to view details. From this screen you can select **buffered objects** for further analysis as Figure 25-62 shows.

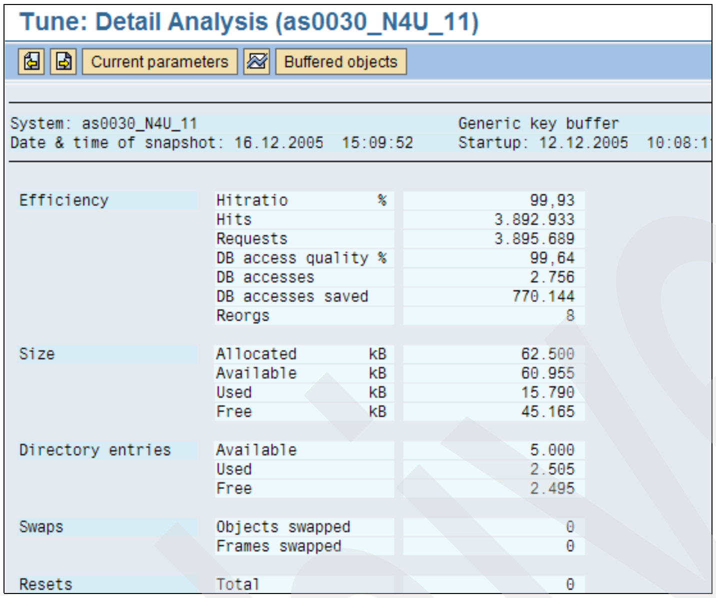

Figure 25-62 ST02 - detail analysis

Select **Current parameters** to view the current parameters as Figure 25-63 shows.

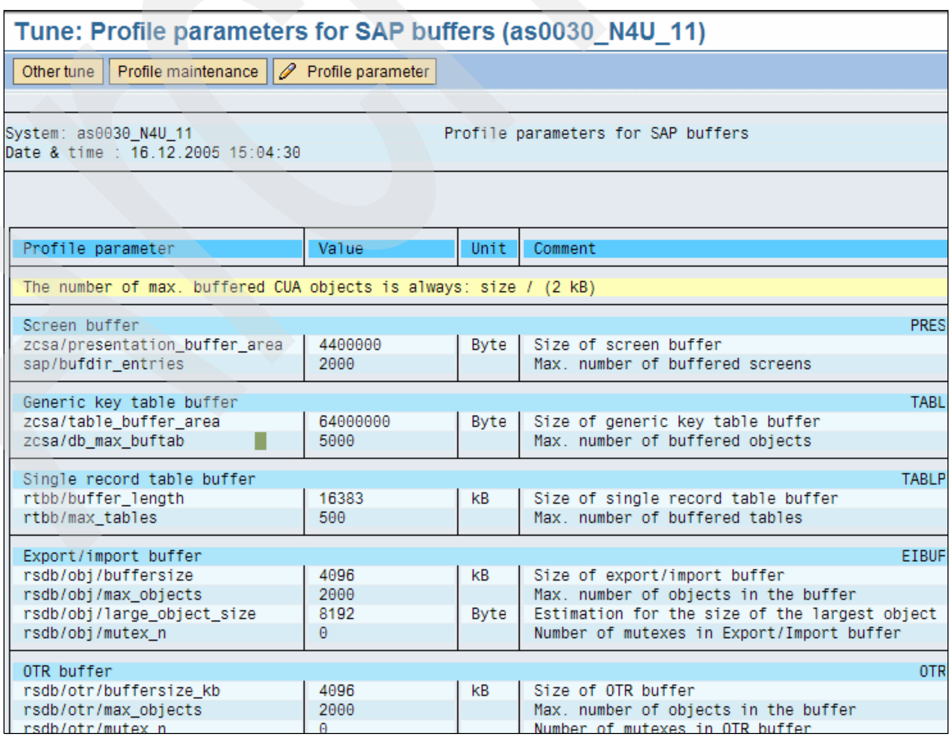

Figure 25-63 ST02 - Current parameters

As an example, assume you found ST02 for the generic key table buffer swaps (shown in red), a free space of 70%, and free directory entries of 15%. In this case, increase the parameter **zcsa/db_max_buftab**.

Select the icon to view the buffer history as Figure 25-64 shows.

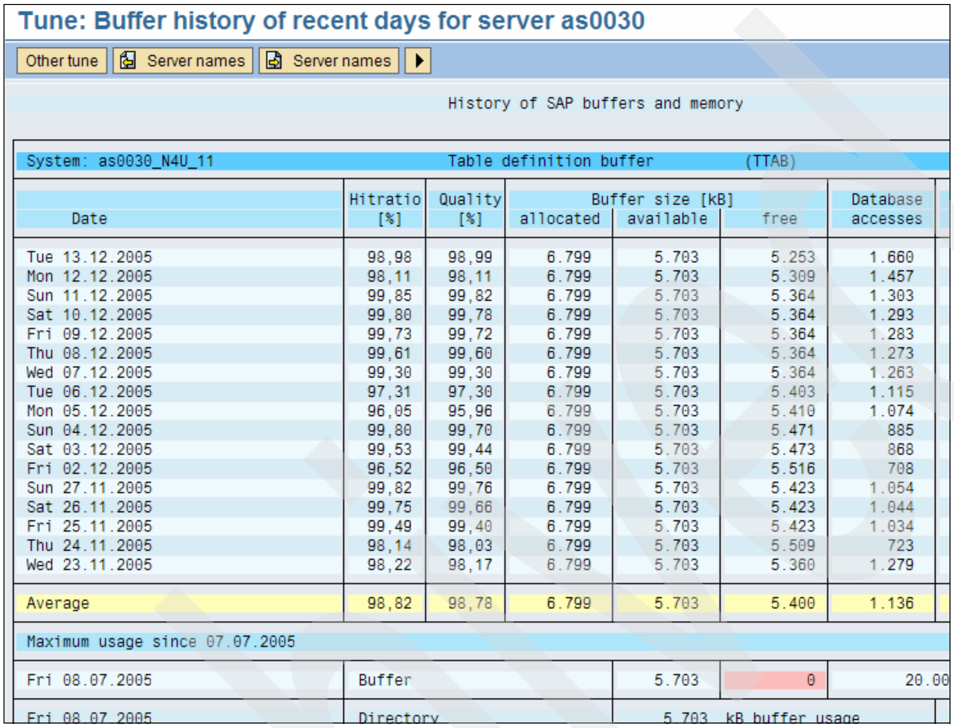

Tune: Buffer history of recent days for server as0030

Other tune | Server names | Server names

History of SAP buffers and memory

System: as0030_N4U_11 Table definition buffer (TTAB)

Date	Hitratio [%]	Quality [%]	Buffer size [kB] allocated	available	free	Database accesses
Tue 13.12.2005	98,98	98,99	6.799	5.703	5.253	1.660
Mon 12.12.2005	98,11	98,11	6.799	5.703	5.309	1.457
Sun 11.12.2005	99,85	99,82	6.799	5.703	5.364	1.303
Sat 10.12.2005	99,80	99,78	6.799	5.703	5.364	1.293
Fri 09.12.2005	99,73	99,72	6.799	5.703	5.364	1.283
Thu 08.12.2005	99,61	99,60	6.799	5.703	5.364	1.273
Wed 07.12.2005	99,30	99,30	6.799	5.703	5.364	1.263
Tue 06.12.2005	97,31	97,30	6.799	5.703	5.403	1.115
Mon 05.12.2005	96,05	95,96	6.799	5.703	5.410	1.074
Sun 04.12.2005	99,80	99,70	6.799	5.703	5.471	885
Sat 03.12.2005	99,53	99,44	6.799	5.703	5.473	868
Fri 02.12.2005	96,52	96,50	6.799	5.703	5.516	708
Sun 27.11.2005	99,82	99,76	6.799	5.703	5.423	1.054
Sat 26.11.2005	99,75	99,66	6.799	5.703	5.423	1.044
Fri 25.11.2005	99,49	99,40	6.799	5.703	5.423	1.034
Thu 24.11.2005	98,14	98,03	6.799	5.703	5.509	723
Wed 23.11.2005	98,22	98,17	6.799	5.703	5.360	1.279
Average	98,82	98,78	6.799	5.703	5.400	1.136

Maximum usage since 07.07.2005

Fri 08.07.2005	Buffer	5.703	0	20.00
Fri 08.07.2005	Directory	5.703 kB buffer usage		

Figure 25-64 ST02 - buffer history

Selecting the **Detail Analysis menu** offers you links to view details and some tasks for further analysis as Figure 25-65 on page 514 shows.

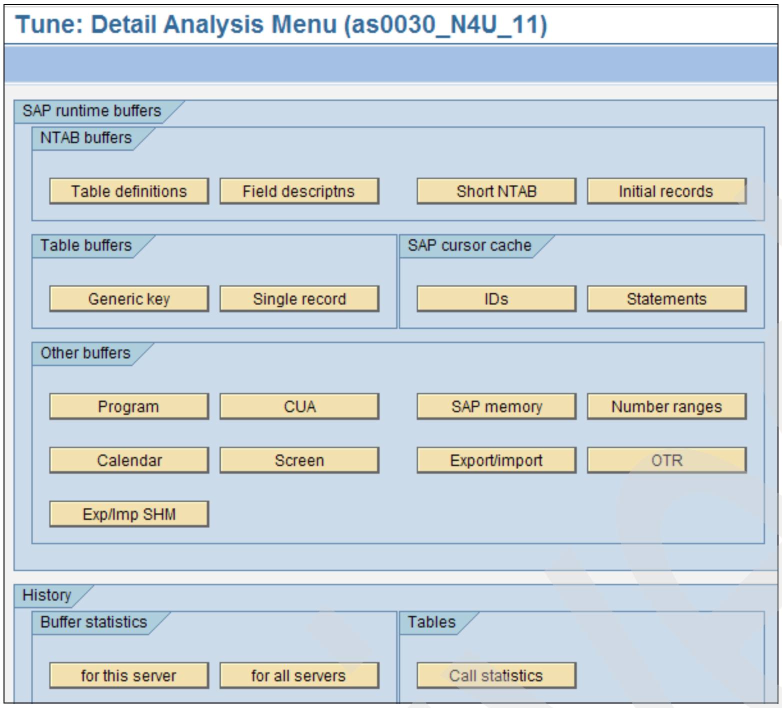

Figure 25-65 ST02 - Detail Analysis menu

25.5.12 Performance trace (ST05)

The transaction ST05 allows you to monitor the following:

- Database accesses of transactions or users (SQL Trace)
- Locking activities (Enqueue Trace)
- Remote Calls of reports (RFC Trace)
- Buffers (Buffer Trace)

ST05 records activities to a trace file and offers support for detailed trace analysis. It is integrated in the Advanced Business Application Programming (ABAP) workbench.

The SQL Trace tool allows you to see and analyze database accesses in detail such as the following:

- The SQL statements your monitored transaction or user carries out
- Which times the system uses for specific database accesses and changes
- How ABAP Open SQL statements are transformed into the SQL syntax for DB2 UDB for iSeries
- Where and how often your application makes database accesses
- Where your application positions COMMIT statements
- What database accesses or changes occur in the update section of your application
- The detailed duration of each step

The Enqueue trace records all locking and unlocking activities that occur for a user or a group of users. The trace file contains the following information:

- Lock statements executed
- Table names of the lock event

- Name of the program that set the lock
- Type of lock
- Lock owner
- Total time required for the lock
- Time required by the enqueue server to release the lock

With Enqueue trace, you can examine your applications, see whether lock conflicts occur, and therefore, if performance is affected.

With the RFC trace, you can monitor which remote calls of your application, or the SAP system are executed, and on which instance these calls are executed. In an RFC trace, all RFC calls, either for a specific user or for a user group are recorded, including the following:

- Which function modules have been called remotely by the program to be analyzed
- Whether the RFC was executed successfully
- The total time used to process the remote call
- The marking of the RFC communication (RFC client or RFC server)
- On which instance the remote call was executed
- Technical parameters of this instance
- The number of sent and received bytes at the RFC

The Buffer trace allows you to monitor all buffer-related activities. The trace file allows you a detailed buffer analysis. It contains information about the following:

- Buffering type
- Buffer access
- Buffer objects
- Buffer values

The trace list is directly linked to the ABAP workbench.

For detailed information about how to make an analysis with ST05, refer to 26.2, “Analyzing and optimizing single transactions” on page 554.

25.5.13 Monitoring Java

SAP recommends that you use CCMS to monitor a production application server. To be able to monitor Java, you need the following prerequisites:

- An installed and running J2EE engine
- An SAP Web AS ABAP 6.40 as the central monitoring system
- The J2EE engine connected to a Central Monitoring System (CEN) (SAPCCMSR agent with the option J2EE has to be registered at CEN)

For information about how to register the SAPCCMSR agent, select **Documentation** → **SAP ERP Central Component,** select your language and search for **Registering the SAPCCMSR agent with option –J2EE** in the SAP online documentation on the Web at:

http://help.sap.com

The Management of the SAP Web Application Server (Java) in SAP NetWeaver documentation can be found by selecting **Monitoring and Error Handling for SAP Web AS Java** on the Web at:

http://help.sap.com/saphelp_erp2004/helpdata/en/c8/cdfacc37efa84d914699ad31eb69b8/frameset.htm

When the J2EE engine starts, it creates the Java Management Extensions (JMX™) monitors and provides them with data during its runtime. The monitoring data of the JMX monitors is mapped in CCMS. To do this, the CCMS connector forwards the monitoring data from the J2EE engine to the shared memory segment of the SAPCCMSR agent with the J2EE option. The agent is part of the CCMS monitoring infrastructure. Run the agent on every host on which an SAP J2EE engine monitored by the CCMS is running. The agent is required even if the SAP J2EE engine is running on the same host as the monitoring SAP system.

You can display the J2EE monitoring data in the central monitoring system using the CCMS alert monitor. To do this, open the alert monitor (transaction RZ20), and select **SAP J2EE Monitor Templates**. The status data is displayed in the following monitor views:

- The monitor view engines display status data for the kernel, services, performance, and system.
- The monitor view applications displays application data.

Within the DB4COCKPIT transaction, you can trace Java SQL statements, although this is without the Explain from the ABAP engine site, if the ABAP and the Java database run on the same server as shown in Figure 25-84 on page 529.

Another way to get basic information of the J2EE engine is to use the http entry point. To access the system information application, log on to the J2EE engine server as an Administrator at the following page:

`http://<hostname>:5<inst.nbr>00`

From this page, choose the **System Information** view.

The System Information application is a good starting point to receive an overview of the most important information about an SAP Web AS system that is running.

25.6 Database monitoring

This section describes how to use database monitor transactions.

25.6.1 Database performance: State on Disk (DB02)

The DB02 transaction monitors the SAP database library, and file statistics for tables and indexes. Another way to run transaction DB02 is to call **DB4COCKPIT** and select **Space State on disk** in the navigation frame as shown in Figure 25-66 on page 517.

Use transaction DB02 to do the following:

- Analyze the current disk status (available and used disk space)
- Determine the current disk space used by tables, primary keys and indexes
- Perform a consistency check
- Perform detailed object analysis
- See the deleted rows analysis
- List damaged files
- Find missing indexes
- See statistics on space consumption

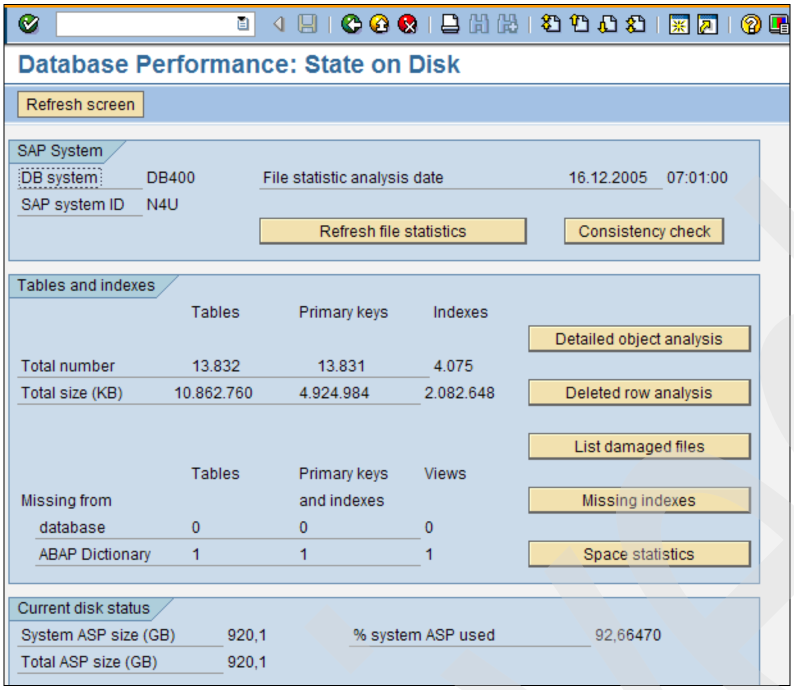

Figure 25-66 Database performance: State of Disk (ST02)

Refreshing the file statistics initiates a long running job to collect this data.

For an example of the Deleted rows analysis report, refer to Figure 25-67. This report can help you determine whether it makes sense to use the Reorganize Physical File Member (**RGZPFM**) command to reorganize some tables.

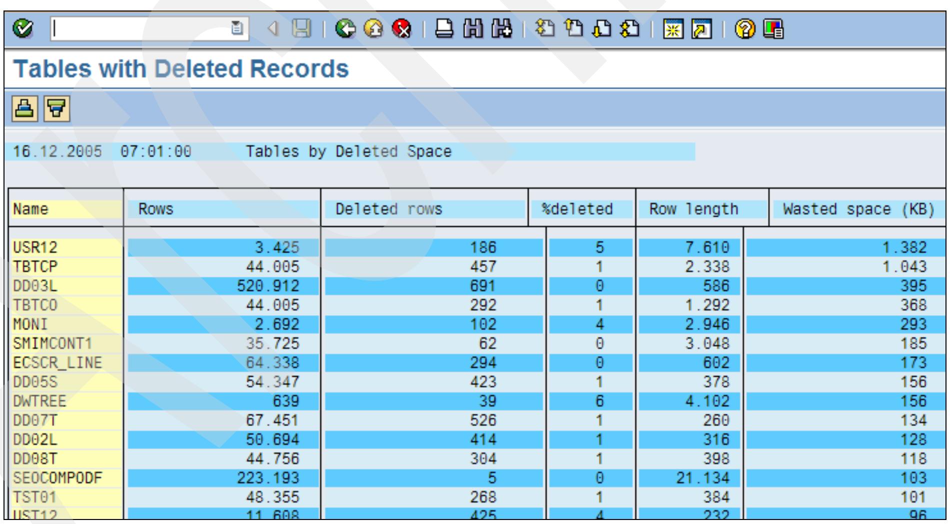

Tables with Deleted Records

16.12.2005 07:01:00 Tables by Deleted Space

Name	Rows	Deleted rows	%deleted	Row length	Wasted space (KB)
USR12	3.425	186	5	7.610	1.382
TBTCP	44.005	457	1	2.338	1.043
DD03L	520.912	691	0	586	395
TBTCO	44.005	292	1	1.292	368
MONI	2.692	102	4	2.946	293
SMIMCONT1	35.725	62	0	3.048	185
ECSCR_LINE	64.338	294	0	602	173
DD05S	54.347	423	1	378	156
DWTREE	639	39	6	4.102	156
DD07T	67.451	526	1	260	134
DD02L	50.694	414	1	316	128
DD08T	44.756	304	1	398	118
SEOCOMPODF	223.193	5	0	21.134	103
TST01	48.355	268	1	384	101
UST12	11.608	425	4	232	96

Figure 25-67 DB02 - Deleted rows analysis

25.6.2 Enabling database monitoring

The SAP database monitor for DB2 UDB for iSeries (DB2 UDB for i5/OS) is part of the Computing Center Management System (CCMS). The database monitor for System i models performs monitoring of the database and analysis of certain types of problems. Since all data comes from the database server, you can retrieve the same data from every application server of your SAP system.

If the parameter as4/dbmon/enable in the instance profile is set to `ON` (default), the memory-resident database monitor is started automatically for an SAP instance. Normally, you should not switch off the database monitor. The only situation when we recommend that you switch off the database monitor is if you have serious performance problems that have already been analyzed, and you are waiting for additional hardware to be delivered.

The database monitor is also activated for an SAP WebAS Java server and can be analyzed from the associated SAP WebAS ABAP server through transaction DB4COCKPIT if the Java database is on the same server (same LPAR) as the ABAP database.

Support for remote Java databases is to be added in SAP releases after 6.40.

A Java database is prepared by default, at the time of installation, to be monitored. To monitor a Java database, the parameter as4/dbmon/enable in the DEFAULT.PFL (not the instance profile) must be `ON` (default value is `ON`). In this case, the JDBC shadow processes uses pre-configured exit programs to initiate monitor functionality after the getConnection() to database. The Java native driver and the toolbox driver uses different exit programs.

For more details about of JDBC exit programs, refer to SAP note 654794.

To deactivate Java monitoring, *do not* touch the exit programs. However, only change **parameter as4/dbmon/enable** to `OFF`.

The SAP database monitor for i5/OS configurations gets its data from the memory-resident database monitor on your System i server.

The data is obtained as follows:

1. The SAPMSSY6 report runs every *n* seconds and submits report RSDB4DMP. N is the value of the profile parameter rdisp/autoabaptime.
2. The RSDB4DMP report initiates a regular database monitor dump every 300 minutes (by default).
3. The RSDB4DMP report schedules the RSDB4UPD report.
4. The RSDB4UPD report initiates an update of the performance tables.
5. This merge supplies all monitors, except the database monitor history, with the necessary information.
6. The RSDB4DMP report also initiates a history database monitor dump every 60 minutes (by default), and schedules the RSDB4DBH report.
7. The RSDB4DBH report calculates the workload statistics that are stored in the database table, MONI, and displays it in the database monitor history.

To configure the database monitor, call the **DB4COCKPIT** transaction and select **Performance** → **Database Monitor** → **Database Monitor Configuration**.

Search for **Database Monitor Configuration (iSeries)** in the SAP online documentation at:

http://help.sap.com

The data of transaction DB02 (Database Performance - State on disk) is refreshed by the program RSDB4090, which is called from the report RSDB_TDB.

The results are stored with a time stamp. You should execute the program on the database system at least once a day.

Considerable CPU resources are required on the database server to run RSDB_TDB, depending on the database size. This means that you should not schedule this program to run too frequently. Therefore, the determined data is not always up-to-date, but does reflect the status at the time of running the collector.

To change the execution times of RSDB_TDB, use **ST03** and select **Collector and Performance DB** → **Performance Monitor Collector** → **Execution Times**.

25.6.3 Database performance monitor (ST04)

The SAP ST04 transaction enables you to monitor and review the performance of your database. SAP kernel 6.40 for i5/OS implementations uses both ST04 and DB4COCKPIT transactions for the memory-resident database monitor.

Run transaction **ST04** to get the screen as shown in Figure 25-68.

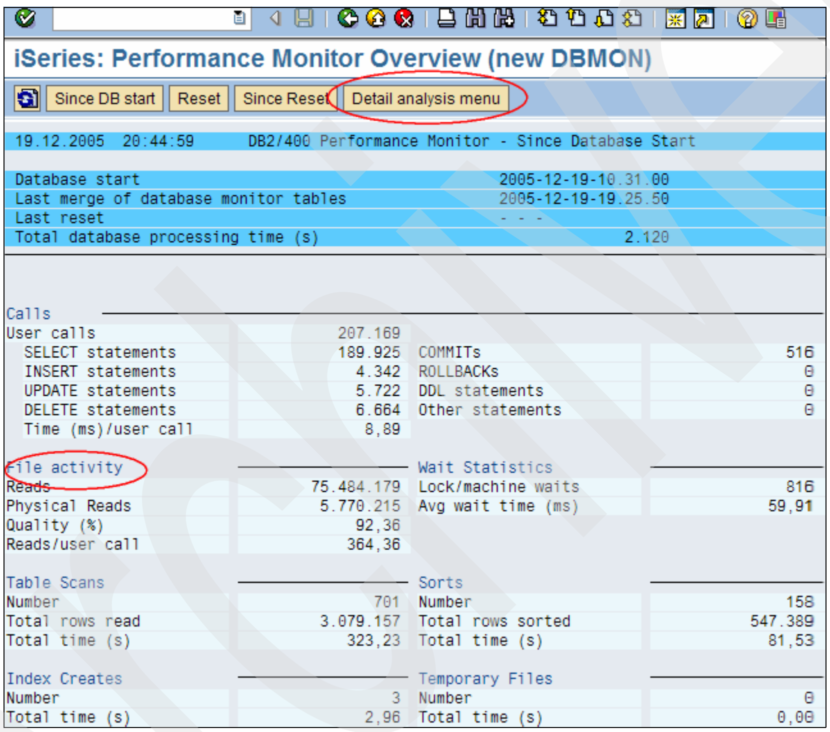

Figure 25-68 Database Performance Monitor (ST04)

Double-click an entry, for example, **File activity** as shown in Figure 25-69 on page 520.

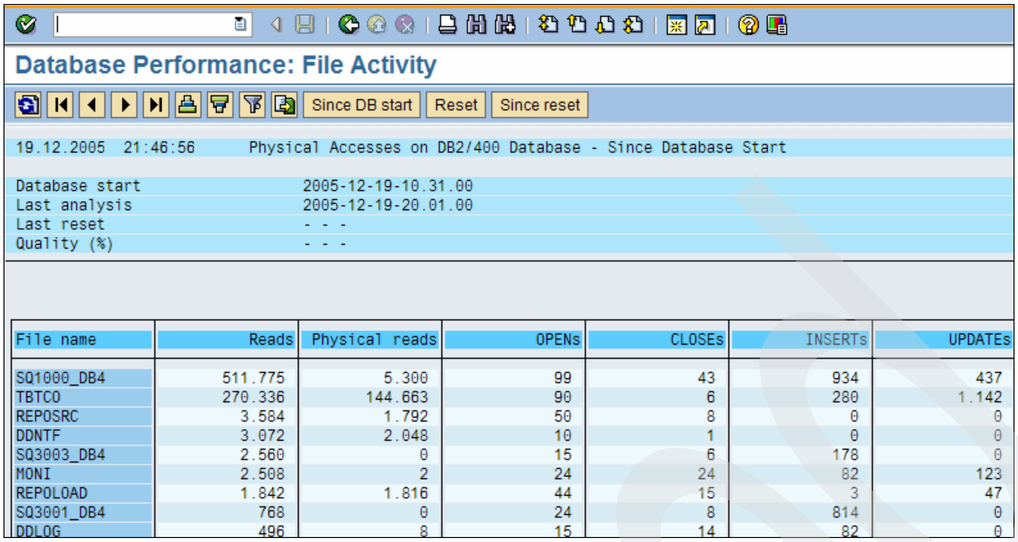

Figure 25-69 ST04 - file activity

Select **Detail analysis** as shown in Figure 25-68 on page 519 to view the selection screen as shown in Figure 25-70.

Figure 25-70 ST04 - Analyze database activity

The following options are available:

- Analyze overall activity
  - Database lock monitor
  - Wait situations

- File activity
- Resource consumption
  - Database monitor history
  - SQL requests
  - 50 slowest queries
- Analyze query activity
  - Table scans
  - Sorts
  - Temporary files
  - Index advised
  - Index usage
  - Index creates
- Additional functions
  - State on disk
  - System catalog views
  - Performance tables
  - SQL packages

With ST04, you can monitor the System Catalog Tables to provide information about all DB2 UDB for iSeries objects, such as table and view definitions, indexes, interrelationship between objects and so on as shown in Figure 25-71.

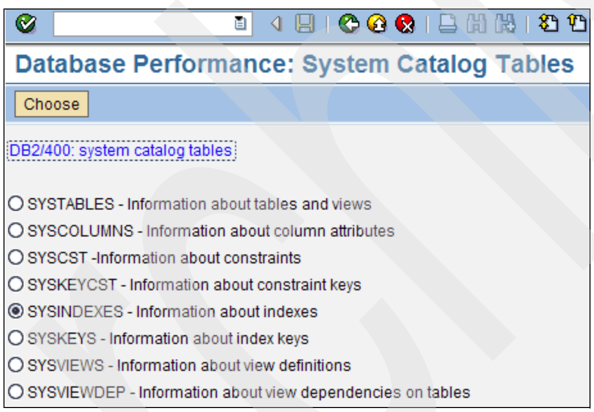

Figure 25-71 ST04 - System catalog view

If you choose **SYSINDEXES**, you get detailed information about all indexes in the SAP database library as shown in Figure 25-72 on page 522.

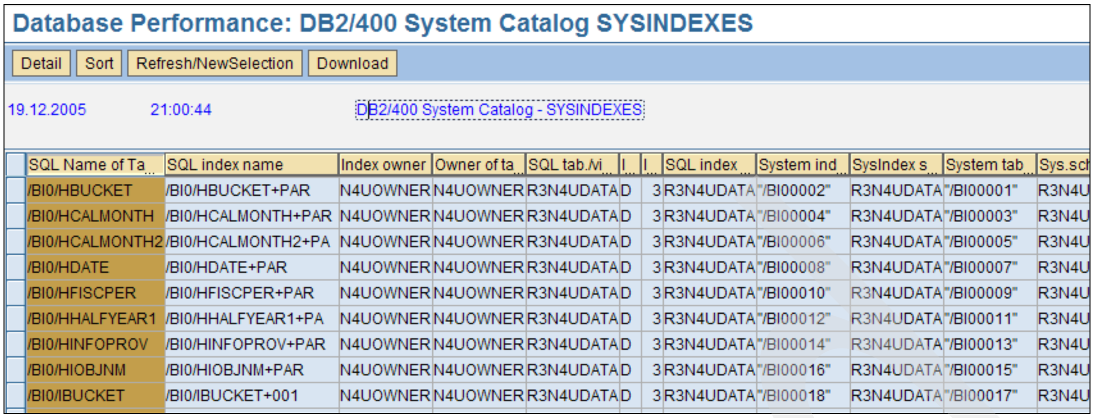

Database Performance: DB2/400 System Catalog SYSINDEXES

Detail | Sort | Refresh/NewSelection | Download

19.12.2005 21:00:44 DB2/400 System Catalog - SYSINDEXES

SQL Name of Ta...	SQL index name	Index owner	Owner of ta...	SQL tab./vi...	I...	I...	SQL index...	System ind...	SysIndex s...	System tab...	Sys.sch
/BI0/HBUCKET	/BI0/HBUCKET+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00002"	R3N4UDATA	"/BI00001"	R3N4U
/BI0/HCALMONTH	/BI0/HCALMONTH+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00004"	R3N4UDATA	"/BI00003"	R3N4U
/BI0/HCALMONTH2	/BI0/HCALMONTH2+PA	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00006"	R3N4UDATA	"/BI00005"	R3N4U
/BI0/HDATE	/BI0/HDATE+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00008"	R3N4UDATA	"/BI00007"	R3N4U
/BI0/HFISCPER	/BI0/HFISCPER+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00010"	R3N4UDATA	"/BI00009"	R3N4U
/BI0/HHALFYEAR1	/BI0/HHALFYEAR1+PA	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00012"	R3N4UDATA	"/BI00011"	R3N4U
/BI0/HINFOPROV	/BI0/HINFOPROV+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00014"	R3N4UDATA	"/BI00013"	R3N4U
/BI0/HIOBJNM	/BI0/HIOBJNM+PAR	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00016"	R3N4UDATA	"/BI00015"	R3N4U
/BI0/IBUCKET	/BI0/IBUCKET+001	N4UOWNER	N4UOWNER	R3N4UDATA	D	3	R3N4UDATA	"/BI00018"	R3N4UDATA	"/BI00017"	R3N4U

Figure 25-72 ST04 - System catalog SYSINDEXES

You can also monitor the database performance tables as shown in Figure 25-73.

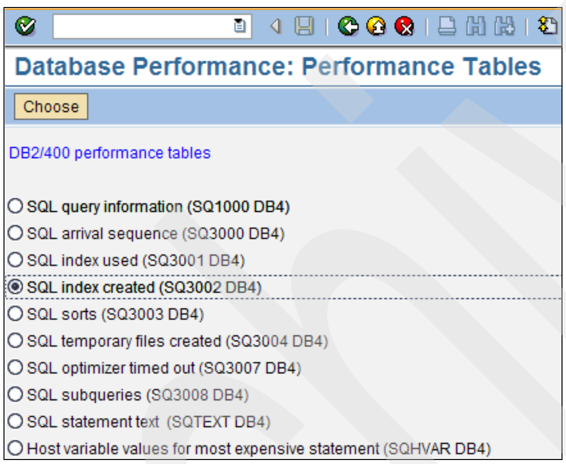

Figure 25-73 ST04 - Performance tables

25.6.4 Database monitor: Transaction DB4COCKPIT

For many tasks, we recommend you use the newer SAP database monitor transaction named DB4COCKPIT. DB4COCKPIT provides database information about the following:

- SQL statements
- Table scans
- Sorts
- Used indexes
- Created indexes
- Advised indexes
- Temporary files
- Database monitor activity
- Database monitor configuration
- SQL packages
- Lock monitor
- File statistics

This transaction includes the functionality of the easy-to-use ST04 transaction. Refer to Figure 25-74 on page 523 for the overview screen.

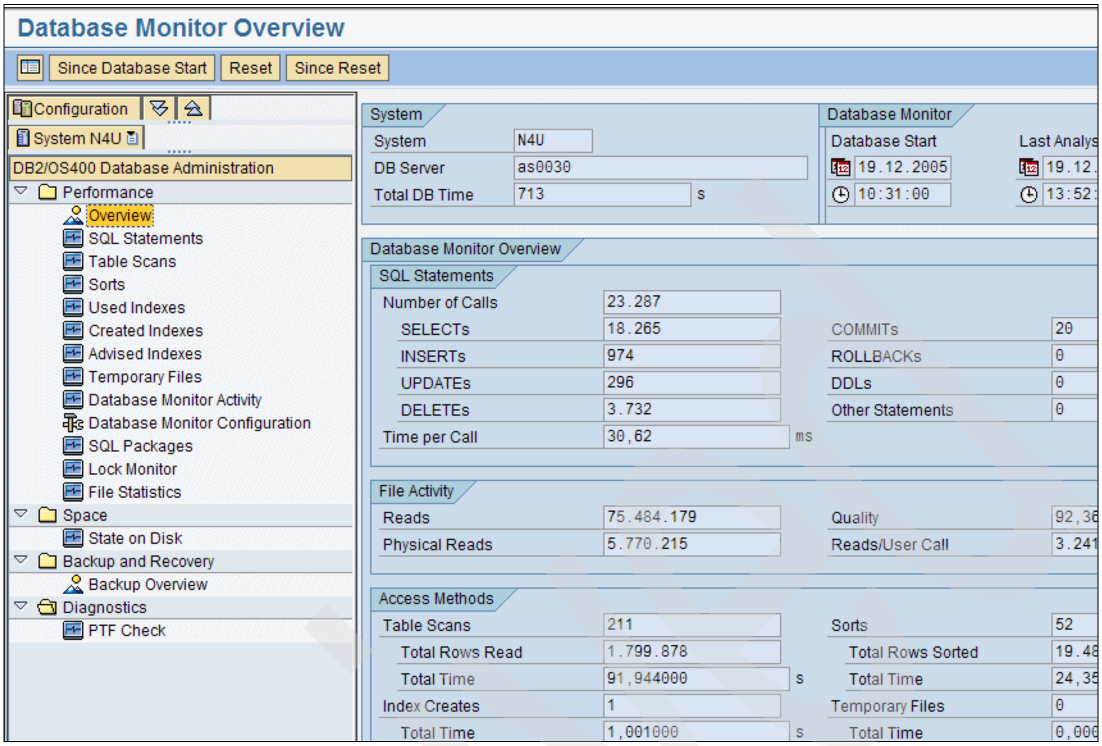

Figure 25-74 Database monitor (DB4COCKPIT)

The following examples show some important information displayed from DB4COCKPIT.

Select **SQL Statements** on the left navigation area, to see a screen similar to that shown in Figure 25-74 on page 523.

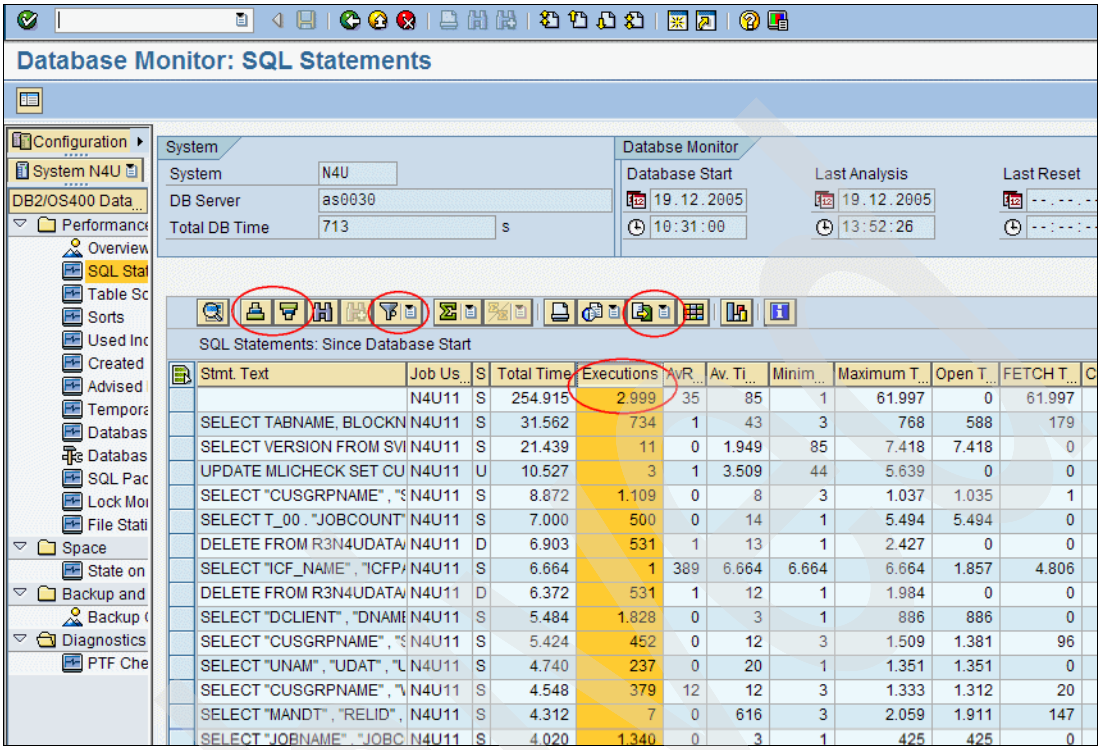

Stmt. Text	Job Us	S	Total Time	Executions	AvR	Av. Ti	Minim	Maximum T	Open T	FETCH T	C
	N4U11	S	254.915	2.999	35	85	1	61.997	0	61.997	
SELECT TABNAME, BLOCKN	N4U11	S	31.562	734	1	43	3	768	588	179	
SELECT VERSION FROM SVI	N4U11	S	21.439	11	0	1.949	85	7.418	7.418	0	
UPDATE MLICHECK SET CU	N4U11	U	10.527	3	1	3.509	44	5.639	0	0	
SELECT "CUSGRPNAME" , "S	N4U11	S	8.872	1.109	0	8	3	1.037	1.035	1	
SELECT T_00 . "JOBCOUNT"	N4U11	S	7.000	500	0	14	1	5.494	5.494	0	
DELETE FROM R3N4UDATA	N4U11	D	6.903	531	1	13	1	2.427	0	0	
SELECT "ICF_NAME" , "ICFP/	N4U11	S	6.664	1	389	6.664	6.664	6.664	1.857	4.806	
DELETE FROM R3N4UDATA	N4U11	D	6.372	531	1	12	1	1.984	0	0	
SELECT "DCLIENT" , "DNAME	N4U11	S	5.484	1.828	0	3	1	886	886	0	
SELECT "CUSGRPNAME" , "S	N4U11	S	5.424	452	0	12	3	1.509	1.381	96	
SELECT "UNAM" , "UDAT" , "U	N4U11	S	4.740	237	0	20	1	1.351	1.351	0	
SELECT "CUSGRPNAME" , "\	N4U11	S	4.548	379	12	12	3	1.333	1.312	20	
SELECT "MANDT" , "RELID" ,	N4U11	S	4.312	7	0	616	3	2.059	1.911	147	
SELECT "JOBNAME" , "JOBC	N4U11	S	4.020	1.340	0	3	1	425	425	0	

Figure 25-75 DB4COCKPIT - SQL statements

By default, this monitor displays the 300 slowest SQL statements recorded by the memory-resident database monitor since the last initial program load (IPL) of the database server.

The statements are initially ordered by total time in descending order. For more detailed information, double-click the SQL statement you are interested in. The SQL Explain panel appears as shown in Figure 25-76 on page 525.

You can set filters to limit the result set of displayed statements, to select a column to order the statements by another criteria, or use the export function to download this report. All options are marked red as shown in Figure 25-75 on page 524.

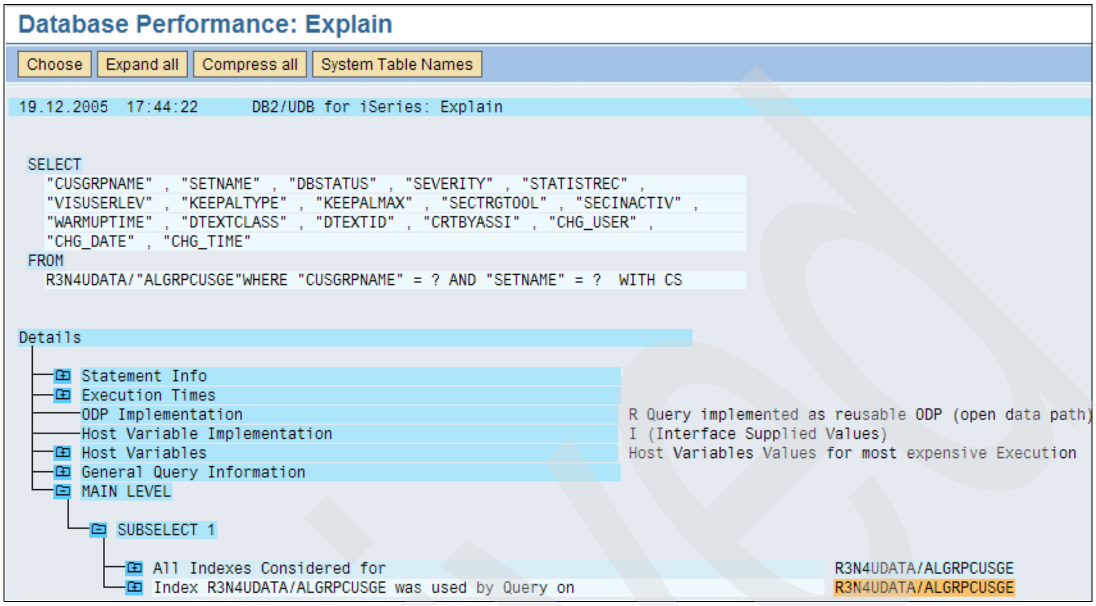

Figure 25-76 DB4COCKPIT - SQL statement detail

With this method, you can easily find slow-running SQL statements to optimize your database performance.

The Used Indexes monitor differentiates between temporary and permanent indexes as shown in Figure 25-77.

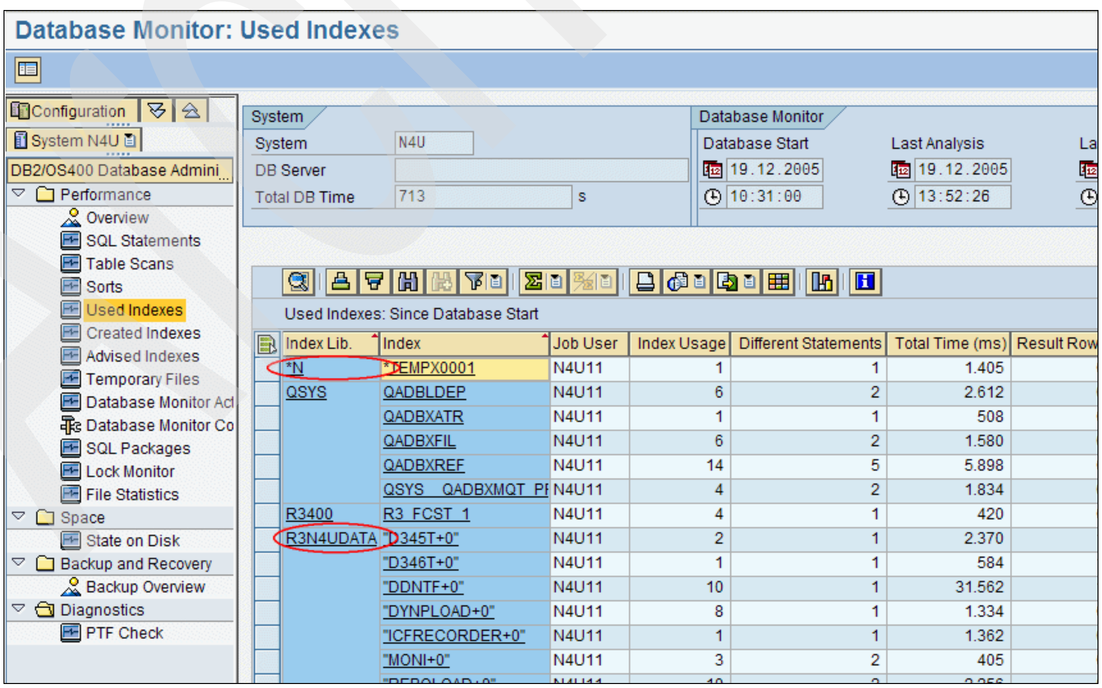

Figure 25-77 DB4COCKPIT - Used Indexes

The query optimizer can dynamically create temporary indexes to execute queries. Indexes can be created to execute queries that specify joins, ordering or grouping.

When the query optimizer creates a temporary index, this index can only be used for the query and work process (job) that initiated it. When the associated cursor is fully closed, the system deletes the index.

Indexes created by the query optimizer are sparse indexes. This means that all selections that can be applied at the time the index is created are built into the index. These indexes, once created, usually perform very well. However, before you create an additional permanent index, consider the cost to maintain this additional index.

To find the list of Advised Indexes, navigate to the related navigation entry in DB4COCKPIT as shown in Figure 25-78.

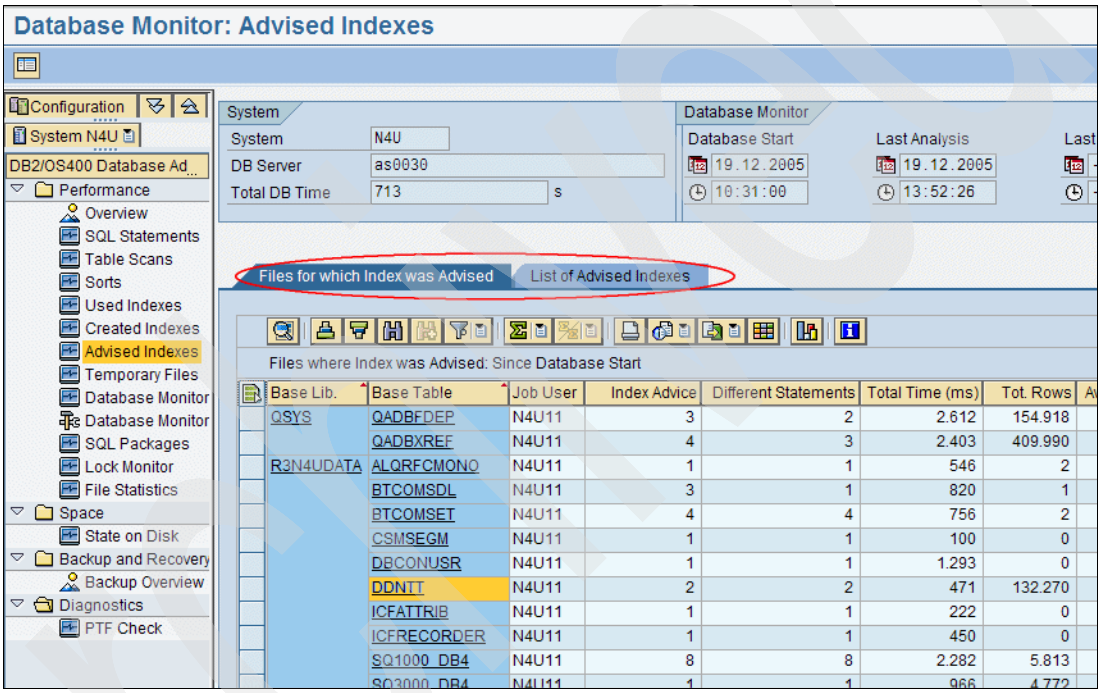

Figure 25-78 DB4COCKPIT - Advised Indexes

The initial view of this monitor shows you the files for which indexes were advised. Double-click an entry to see the details as shown in Figure 25-79. You can change the view to **List of Advised indexes**, as marked in red in Figure 25-78. Refer to Figure 25-80 on page 527.

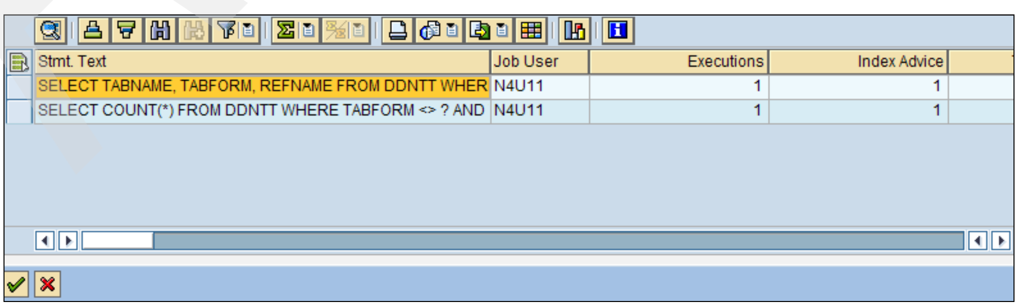

Figure 25-79 DB4COCKPIT - Advised Indexes - detail

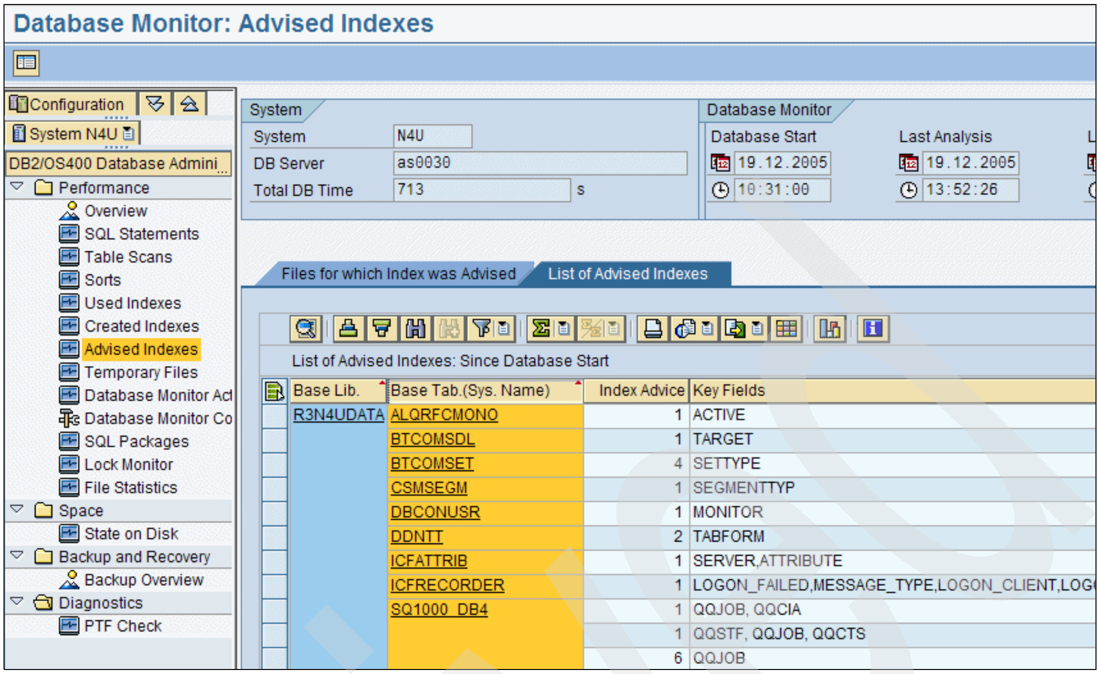

Figure 25-80 DB4COCKPIT - List of advised Indexes

These advised indexes may not necessarily be optimum, because the function takes into consideration only the Where clause. The Order by and Group by clauses are not taken into consideration on this occasion.

In i5/OS V5R4, there is a table QSYS2/SYSIXADV with index-advised information from SQE, independently from the database monitor. There are plans to make this information available as a DB4COCKPIT function later.

Note: Advised indexes may not necessarily be optimum.

The query optimizer might need to create temporary files for more complex queries.

This Temporary Files monitor provides an overview of temporary files that have been created in the database since the last IPL of the database server. The monitor displays the time required to build a temporary file (Total Time (ms)) and indicates why the query optimizer is required to create the temporary file (Reason Code). This information indicates how the query can be rewritten to avoid a temporary file being created in future.

The Database Monitor Activity monitor provides information about which jobs are monitored by which monitor type, and the times of regular dumps and merge times to update the database statistics as shown in Figure 25-81.

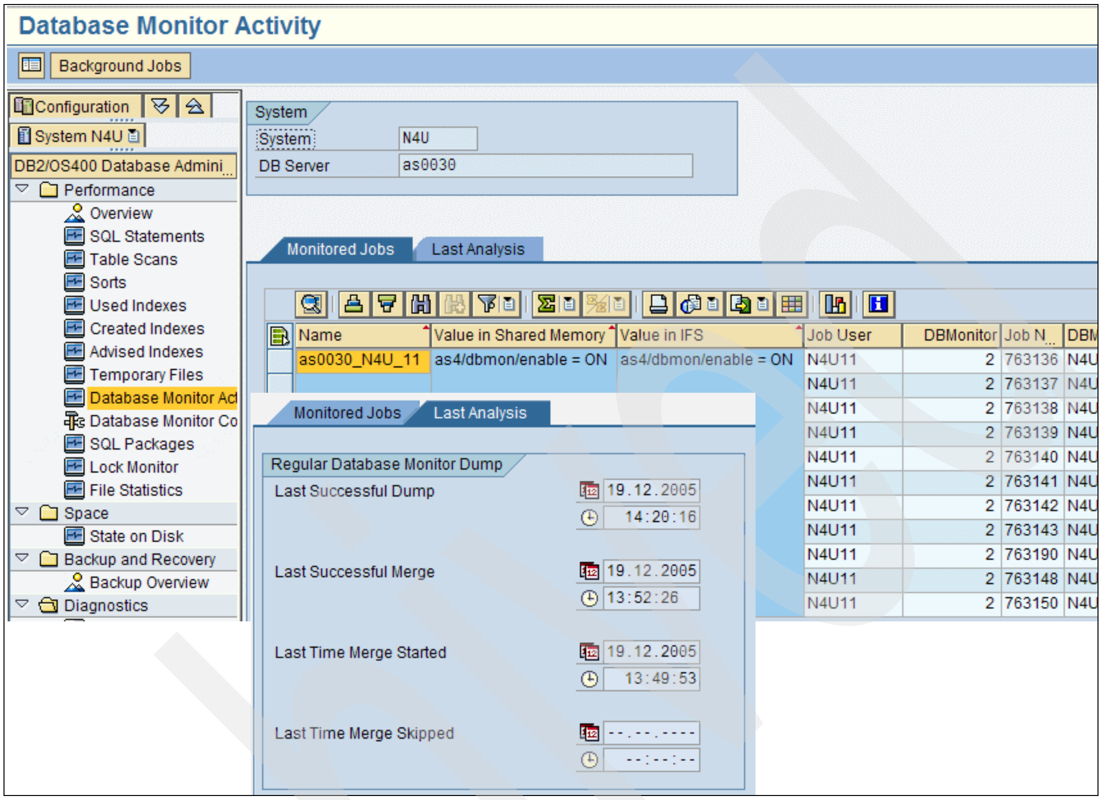

Figure 25-81   DB4COCKPIT - Database Monitor Activity - Monitored Jobs and Last Analysis

The value 2 in the DBMonitor column indicates that the job is monitored by the memory-resident database monitor.

Use **Database Monitor Configuration** to setup your own configuration as shown in Figure 25-82:

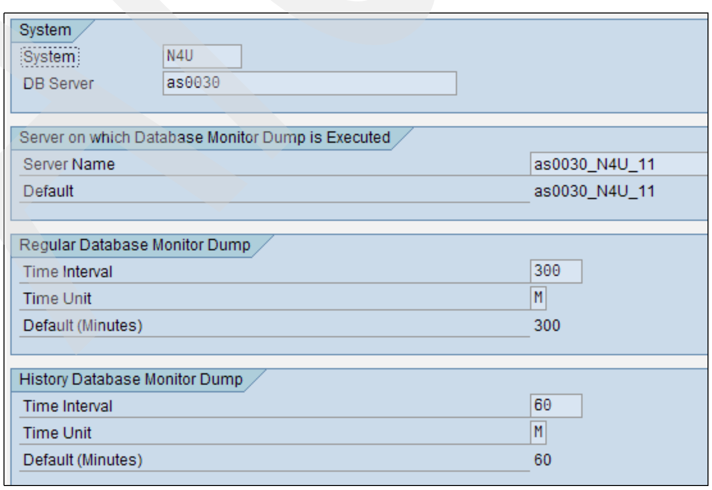

Figure 25-82   DB4COCKPIT - Database Monitor Configuration

The File Statistics monitor provides an analysis of activity for all physical files in the SAP database. This includes statistics on the number of OPENs, CLOSEs, INSERTs, UPDATEs, DELETEs, and logical and physical reads performed on the files.

To monitor your Java instance, select a related entry in the left navigation area of DB4COCKPIT as shown in Figure 25-83.

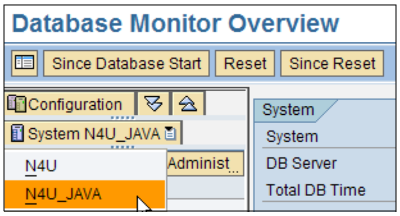

Figure 25-83 DB4COCKPIT - Select your Java instance

Figure 25-84 shows the initial overview screen.

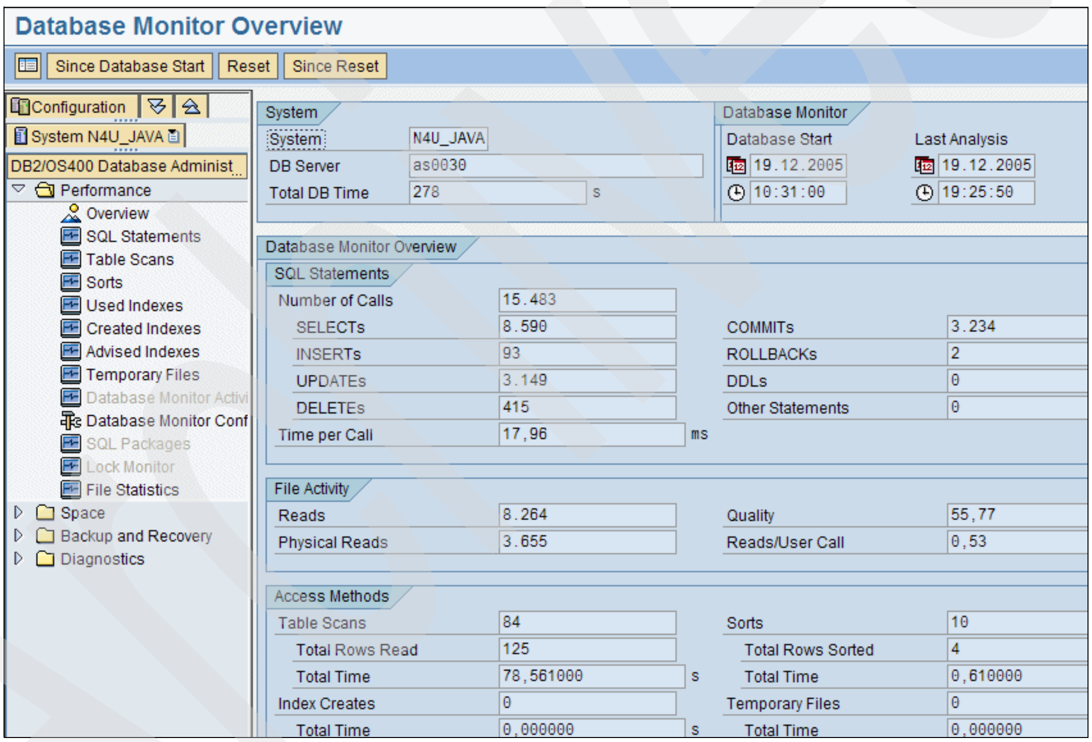

Figure 25-84 DB4COCKPIT - Database Monitor for the Java instance

Navigate in the same manner as the ABAP instance to see table scans, for example, as shown in Figure 25-85.

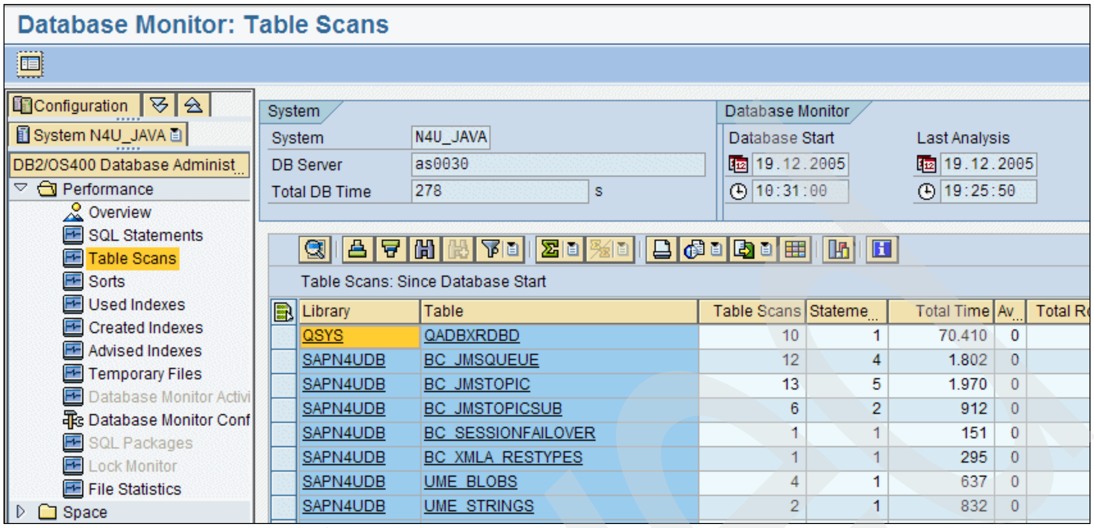

Figure 25-85 DB4COCKPIT - Table scans

There are less monitor options for the Java instance.

25.7 SAP Business Intelligence performance

SAP Business Intelligence, as an online analytical processing (OLAP) application, works very differently from online transaction processing (OLTP) applications, such as the SAP Central Component out of mySAP Enterprise Resource Planning.

For information about using OLAP and OLTP and other Business Information Warehouse information, refer to Chapter 28, “SAP Business Information Warehouse” .

26

Performance tuning

This chapter discusses analyzing and tuning your System i server and SAP implementations.

Refer to Chapter 25, “Performance concepts and monitoring” to understand the concepts of monitoring your System i server and SAP system, and Chapter 27, “Performance management tools and services” for information about the tools and services available.

26.1 Tuning System i models

To optimize system performance, tune the hardware together with all the system components, tune the application, optimize the database, and look systematically for possible bottlenecks in the system performance.

There are different ways to tune a complete system. If you are already aware of certain problems or bottlenecks, it is better to begin where the biggest improvements are required.

Thus, you can avoid purchasing new hardware by tuning your applications to avoid bottlenecks.

Figure 26-1 illustrates the general steps in tuning.

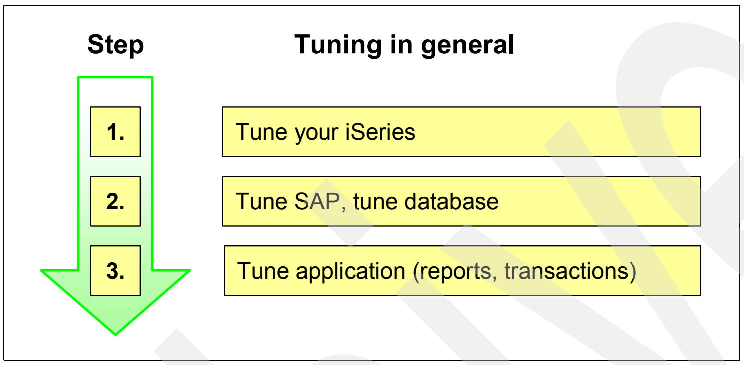

Figure 26-1 Tuning steps in general

To tune your System i server, you must first monitor your system. To determine what tool or command to use, refer to Table 25-2 on page 470. To understand the effects of resource utilization on system performance, refer to 25.4.1, “System i monitors” on page 470.

The figures obtained when you monitor your system are a function of the duration of the work. It makes a difference how 80% of central processing unit (CPU) utilization is ascertained, for example, as a snapshot of a few seconds, of one hour, or as an average of a whole day.

Most tuning steps influence each other. For example, if the main memory utilization grows because of an increase in paging rate, the amount of disk input/output (I/O) increases. Tuning an SAP program and inserting a helpful index in the database improves response time and decreases CPU utilization.

> **Note:** Evaluate data based on the duration of its appearance. Remember that tuning steps influence each other.

Furthermore, possible performance problems are often differentiated by duration. For example, you could observe some problems while monitoring a few snapshots in a short span of time (such as in a minute) or over a long time (the average in one or more hours).

Figure 26-2 illustrates how to tune your System i server.

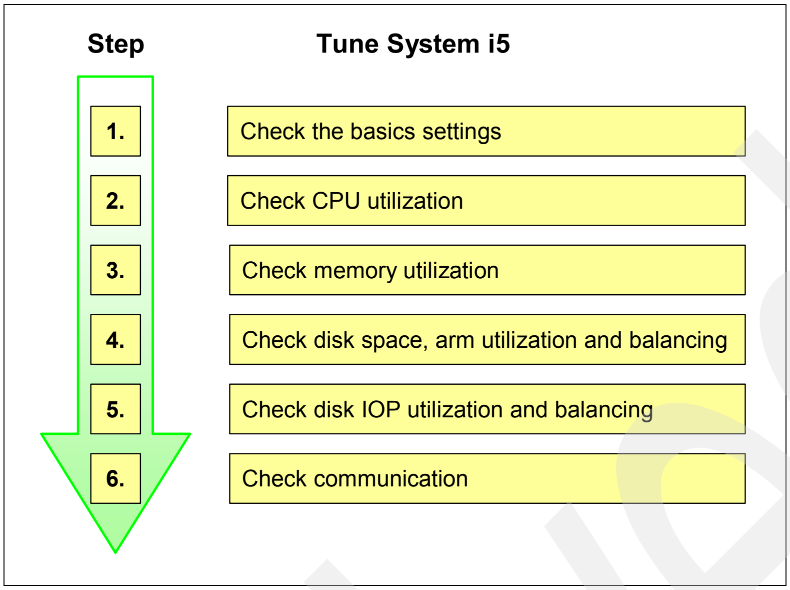

Figure 26-2 An approach to System i tuning

Basic settings

The following are the basic settings necessary to tune your System i server:

1. Set the QPFRADJ system value to `0` to prevent i5/OS from automatically changing the memory pool settings.
2. Set the QDYNPTYSCD system value to `1` to allow dynamic priority scheduling of jobs. A change to this system value takes effect at the next initial program load (IPL).
3. Set the SAP system memory pools to use expert cache. With expert cache turned on, the system adjusts storage pool paging characteristics and data handling. It reduces the number of I/O operations and disk arm contention.

   Use the Work with Shared Pool (**`WRKSHRPOOL`**) command to turn on expert cache. Set the Paging Option to `*CALC`. Turn off expert cache only when there are main memory constraints.
4. Set up the activity level of the *BASE pool. To do this, perform the following steps:
   a. Sum up all the work processes of all SAP systems which run in *BASE.
   b. Add all shadow work processes in the 3-tier SAP installations using Transmission Control Protocol/Internet Protocol (TCP/IP).

      Additionally, for Java installations, add all the threads of the Java 2 Platform, Enterprise Edition (J2EE) server jobs. Multiply this value by 1.2.
   c. Use the Change System Value (**`CHGSYSVAL`**) command to set the activity level of *BASE to the following calculated value:

      ```
      CHGSYSVAL QBASACTLVL n
      ```
5. Set the TCP/IP buffers to be a minimum of 1 MB. Use the **`CHGTCPA TCPRCVBUF (1048576) TCPSNDBUF (1048576)`** command. To activate the setting, restart TCP/IP by running the **`ENDTCP`** command followed by **`STRTCP`**.

CPU

If you determine the average CPU utilization for a typical business time, compare it to the rule of thumb sizing guideline of 65% utilization. This gives a reserve on many systems to manage the workload of some normal peak activities. High CPU utilization can be normal when there are many long-time running jobs.

Compare your monitored data with the general guidelines, as shown in Table 26-1 and Table 26-2. The detailed guidelines depend on the number of CPUs on your system.

The color designation seen in the columns in Table 26-1 and Table 26-2 can be interpreted as follows:

- Green: very good performance. Little potential to improve performance.
- Yellow: good to acceptable performance. Improvements have an effect.
- Red: unacceptable performance. A bottleneck with strong impact on performance.

Table 26-1 Guidelines for long-time CPU utilization

CPU utilization - long time			
Number of CPUs	good	good to still acceptable	Critical
1	< 80%	80 ... 90%	> 90%
2	< 85%	85 ... 93%	> 93%
3	< 87%	87 ... 94%	> 94%
4	< 89%	89 ... 95%	> 95%
> 4	< 93%	93 ... 97%	> 97%

Table 26-2 Guidelines for short-time CPU utilization

CPU utilization - short time			
Number of CPUs	good	good to still acceptable	Critical
1	< 85%	85 ... 94%	> 94%
2	< 88%	88 ... 96%	> 96%
3	< 90%	90 ... 97%	> 97%
4	< 92%	92 ... 98%	> 98%
> 4	< 95%	95 ... 99%	> 99%

Perform any of the following actions if you suspect a CPU bottleneck:

- Reallocate the partition size (if used).
- Use Capacity on Demand processors or prepaid temporary CPU days.
- Use lower priorities for jobs running in test and development systems. This does not decrease the CPU utilization, but does prioritize CPU utilization for more important tasks. Refer to 27.3.3, “Testing systems with lower priority” on page 581for more details.
- Use different priority levels for update, batch, spool, and gateway processes. Refer to 27.3.3, “Testing systems with lower priority” on page 581for more details.
- Optimize your applications.
- Upgrade the CPU.

Memory

There is no measured value for memory utilization. This makes the memory resource factor more complicated. When the faulting rate of the response time transcends a certain value, it is a sure sign that there is a bottleneck in the main memory.

To analyze the performance factors of memory, perform the following steps:

1. The initially setting of the *MACHINE pool should be 10% to 15% of the main storage size.
2. Make sure that the faulting rate in the *MACHINE pool is low (<4, and if possible, <1).

   If the measurement does not meet this recommendation, enlarge the *MACHINE pool. Alter the memory of other pools until the machine pool measurements show good parameters. Faulting is the most important tuning step to optimize memory utilization. Refer to Figure 25-6 on page 471 and Figure 27-20 on page 588.
3. Check the other pools.

   Use this rule of thumb: try to achieve a faulting rate of less than 100 per fully occupied processor, as represented by the formula in Example 26-1.

Example 26-1 100% processor formula

```
#CPU * %CPUutil = X;
FltRateAll < X ?
#CPU - number of CPU of the LPAR or system
%CPUutil - average CPU utilization in percent
FltRateAll - number of all Faults per second (DB + NonDB) in all pools
For example:
4 CPUs * 60% Utilization = 240

Measured FltRateAll = 190
190 < 240 --> Fault rate is good
```

This rule should meet most high performance requirements. However, this is not necessary or reasonable for all customers. Some applications can successfully run with hundreds of faults on large machines.

You can calculate the portion of response time that is caused by faulting. Performance tuning for the long-term is more critical as compared to performance tuning for short intervals.

A long-term method is based on the following measurements:

- TransInTime / DurationS = TransPerS
- FltRateAll / TransPerS = FltsPerTrans
- AvgDiskResp * FltsPerTrans = TransFltTime
- AvgDiaTransTime / TransFltTime = PerctFault
  - TransInTime: Transactions in a time interval (ST03)
  - DurationS: Duration of this interval in seconds
  - TransPerS: Transactions per second
  - FltRateAll: Number of all faults per second (DB + NonDB) (System report / ST06)
  - FltsPerTrans: Faults per transaction
  - AvgDiskResp: Average disk response time (System report / ST06)
  - TransFltTime: Transactions faulting time
  - AvgDiaTransTime: Average Dialog Response Time (ST03)
  - PerctFault

Refer to Example 26-2 for a working example.

Example 26-2 Response time calculation

```
36000 Transactions (in 1 hour) / 3600 s = 10 Transactions per s
225 (Faults all) / 10 Transactions per s = 22,5 Faults per Transaction
4,0 ms Avg. disk resp. * 22,5 Faults per Transaction = 90,0 ms Trans.faulting time
600 ms AvgDiaTransTime / 90,0 ms = 15 % --> it is acceptable
```

Further analysis is required for a more detailed analysis of performance.

A method to determine response time for a short time period or snapshot analysis is as follows:

- Use several sessions on the SAP system to determine what the following values are at the same time:
  - FltRateAll: Number of all faults per second (DB + NonDB) (WRKSYSSTS / ST06)
  - AvgDiskResp: Average disk response time (ST06 - calculate or estimate average)
  - NbrActWP: Number of active work processes (SM50, SM66)
- Calculate the faulting overhead per second using the following formula:

  ```
  FltRateAll / NbrActWP * AvgDiskResp = FltOverhead
  ```

  For example:

  ```
  320 faults per second / 8 act.WP * 5,0 ms = 200 ms (20 % of a second)
  ```

A fault rate of 20% is within the acceptable range for response time performance. Refer to Table 26-3 for the guidelines of response time relative to fault rates.

Table 26-3 Fault rate factor of response time guideline

Percentage of response time due to page faults		
good	good to still acceptable	Critical
< 10%	10 ... 25%	> 25%

Note: If you run SAP applications in an LPAR system, you must explore the work processes of all instances and systems. If you examine one of several memory pools, explore the work processes of the SAP applications which run in this pool.

Perform the following actions if you suspect a main storage bottleneck:

- Reallocate a partition's main storage setting if logical partition (LPAR) is used.

  To test the effects of increased main storage availability, consider suspending a less critical application (for example, a test system), and temporarily allocate its main storage to a partition constrained by main storage.
- Use Capacity on Demand main storage.
- Run several SAP systems in different memory pools and allocate an amount of main storage in order of importance of the SAP system. This method does not decrease the memory utilization, but can improve performance of more important tasks.
- Optimize your applications by reducing the amount of data transferred or build indexes to avoid table scans.
- Add main storage to the system.

Disk

To determine the impact of your disk configuration on performance, compare your measured results to guidelines, as outlined in the following process:

1. Differentiate between response time measurements taken over a short time (for example a snapshot) and a longer time (for example, a one hour interval). Compare the measured disk space utilization to the values shown in Table 26-4.

Table 26-4 Guidelines for disk space

Percentage Usage of disk space		
good	good to still acceptable	Critical
< 70%	70 ... 80%	> 80%

2. Compare your measured disk arm utilization to the values shown in Table 26-5 and Table 26-6.

Table 26-5 Guidelines for disk arm utilization - long time

Percentage disk arm utilization - long time / average		
good	good to acceptable	Critical
< 10%	10 ... 20%	> 20%

Table 26-6 Guidelines for disk arm utilization - short time / peak

Percentage disk arm utilization - short time / peak		
good	good to acceptable	Critical
< 30%	30 ... 40%	> 40%

3. Examine the effect of balancing disk space utilization and disk arm utilization.

   Consider the following actions if you suspect a disk bottleneck:

   – Housekeeping

     Remove the unnecessary objects. Archive less-used objects.

   – Use the Reorganize Physical File Member (**RGZPFM**) command to reorganize database files.

   – Add disk units if the percentage of disk space used is high.

   – Add disk units if the percentage of disk arm utilization is high.

   – Replace slower disk devices with faster disk drives.

   – If the disk space utilization within an auxiliary storage pool (ASP) is unbalanced, use the Start ASP Balance (**STRASPBAL TYPE(*CAPACITY)**) command to balance the capacity of an ASP.

   – If the disk arm utilization within an ASP is unbalanced, use the Start ASP Balance (**STRASPBAL TYPE(*USAGE)**) command to balance disk arm utilization.

     You can find instructions for using STRASPBAL by searching for `STRBASPBAL` at:

     http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp

     Also refer to *SAP note 517515: iSeries: When to use ASP balancing?*

Refer to Table 25-2 on page 470 for a list of tools used to monitor disk capacity, disk arm utilization, and disk load balancing.

For information about the types of disks and their performance, refer to the chapter on direct access storage device (DASD) performance in: *iSeries Performance Capabilities Reference - i5/OS Version 5, Release 3,* SC41-0607.

I/O processor

To determine the impact of your input/output processor (IOP) configuration on performance, compare the measured IOP utilization to the guidelines, as shown in the following process:

1. For disk IOP utilization, refer to Table 26-7.

Table 26-7 Guidelines for disk IOP utilization - long time

Percentage disk IOP utilization - long time		
good	good to still acceptable	Critical
< 10%	10 ... 20%	> 20%

Table 26-8 Guidelines for disk IOP utilization short time / peak

Percentage disk IOP utilization - short time / peak		
good	good to still acceptable	Critical
< 30%	30 ... 40%	> 40%

2. Examine the balancing of disk IOP utilization.

Consider the following actions if you suspect an IOP bottleneck:

- If the disk IOP utilization is unbalanced, check your configuration to determine if it is possible to balance the workload by moving disks to another IOP.
- Add IOPs if the percentage of utilization is too high.

Communication

In a complex network environment, as in any type of configuration, the slowest component is the bottleneck. For wide area network (WAN) configurations, the communications speed and latency times can be the largest contributor to overall response time. If a communications configuration setup is correct, however, a local network is usually not a bottleneck for an SAP system.

To determine the impact of your communications configuration on performance, compare your measured communication line utilization to the guidelines described in this section.

In order to achieve good performance in a multi-user interactive local area network (LAN) environment, manage the number of active users. Keep media utilization of LAN communication lines with multiple users below 50% for token ring and below 25% for ethernet environments, because media collisions cause thrashing. Operating at higher utilizations can cause poor response time due to an excessive queuing time for the communications line. For an interactive utilization of a LAN IOP, do not exceed 60% IOP utilization.

Use File Transfer Protocol (FTP) to transfer a large file to check the overall network data rate. The resulting FTP speed should be faster than 10% of the theoretical network speed. Perform this test in both directions (put and get command). The results should be similar.

In addition to a client/server connection in a complex SAP system landscape, a separate network, as well as other network connections, can be used for database and application servers, or for LAN connections used for Independent Auxiliary Storage Pool (IASP) with cross-site mirroring (XSM).

For an in-depth analysis of a large network, you require additional network monitoring software.

Consider the following actions if you suspect a network bottleneck:

- Replace components that are limited by technology.
- Increase the network capacity.

For further information about communications resource utilization, refer to:

- *SAP note 484548: iSeries: 3 tier system installation*
- *SAP note 500235: Network Diagnosis with NIPING*

To understand the effect of communication line utilization on system performance, refer to 25.4.1, "System i monitors" on page 470

26.1.1 Tuning an SAP system

You must monitor your SAP system in order to tune it. Refer to 25.5, "Monitoring SAP systems" on page 485 for a description of the commands used to monitor an SAP system.

Figure 26-3 shows the steps involved in tuning an SAP system.

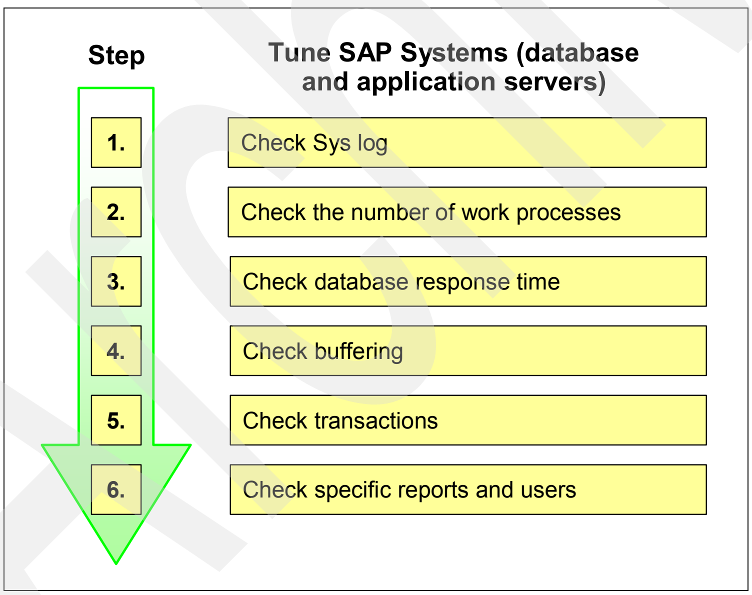

Figure 26-3 Steps to tune an SAP system

The steps represented in Figure 26-3 are described in this section.

Sys log

Check and analyze the System Log - Analysis (transaction SM21**)** for general problems, such as frequent errors, dumps, and failures. For example, frequent Structured Query Language (SQL) errors can indicate a database problem. Each time the newly developed source is transported with **`tp`** into the system that is running the **`tp`**, the entries of the related Advanced

Business Application Programming (ABAP) in the program buffer are invalidated. Use the /$SYNC command to clear the buffers.

See Figure 25-40 on page 500 for an illustration of the System Log - Analysis transaction.

Work processes

For good performance, minimize the queuing time of an SAP work process. Use the **SM50** transaction to monitor your SAP work processes. See 25.5, "Monitoring SAP systems" on page 485 to see the type and status of all work processes. Ensure there are a sufficient number of dialog, batch, and update processes available in the peak periods. Click the **clock** button to view the CPU consumption.

Jobs with low CPU consumption in a group indicate that there is a sufficient number of work processes. On the other hand, if the jobs of a group have more or less equal CPU consumption, queuing of some work processes occur.

The least-used work processes should be free most of the time for dialog work processes. This means CPU consumption should be near zero. In general, a high wait time (transaction ST03) indicates the need for more work processes.

Using large instances can help minimize the main memory consumption, but bottlenecks can occur on shared resources within one instance, for example, the dispatcher or gateway process.

For more information about limiting work processes, and the advantages and disadvantages of large instances, refer to *SAP note 9942: Maximum number of work processes*.

Response time components

Figure 26-4 shows the components comprising SAP system response time.

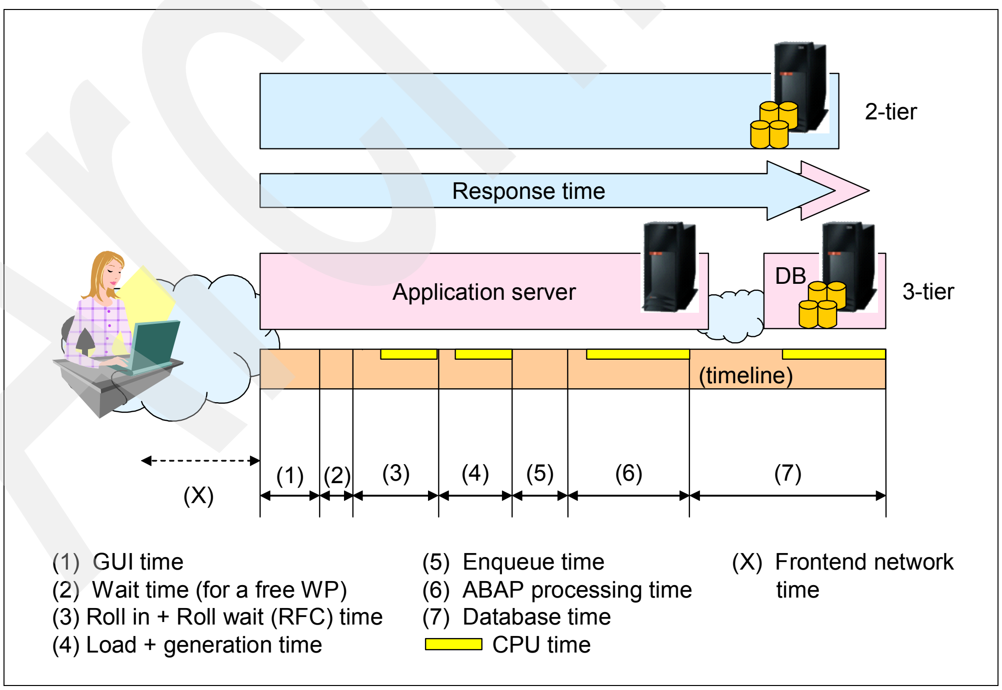

Figure 26-4 SAP system response time

Database response time

The measurement of the response time, either a transaction, a user, or an average of a time interval, contains different components which contribute to the whole response time. There is an additional part for front end network time which is difficult to measure. This time also passes before the user sees the result of his transaction.

In the SAP Enterprise Central Component (SAP ECC) system today (the successor product to R/3), response time is expected to be less than 0.6 seconds, and up to 1.0 second. In a well-tuned system, most of the response time consists of ABAP time and database time. The wait time should be near zero. Do not exceed 50% for the database factor. Database operations in 3-tier installations are somewhat higher, because network overhead is included.

Reports which show CPU time are based on the overall CPU used in the measured work process. This includes the CPU time for ABAP processing, and the load and generation, such as roll-in operations. In 2-tier installations where the database requests are processed in the work process directly, the CPU time includes the portion of the CPU used for database operations.

Some SAP applications, such as mySAP Business Information Warehouse (BW), can show higher database and response times, depending on the amount of data and the types of queries.

You can see significantly higher response times in development systems. This is caused in part by poor buffer quality, operations that are rarely used, no frequently used key transactions, and frequently recompiling programs.

Note: This behavior is not a particular factor of the SAP applications on System i models. It applies to every development system.

The steps following those shown in Figure 26-3 on page 539 include checking database response time, buffering, transactions, as well as specific reports or users as Figure 26-5 on page 542 shows.

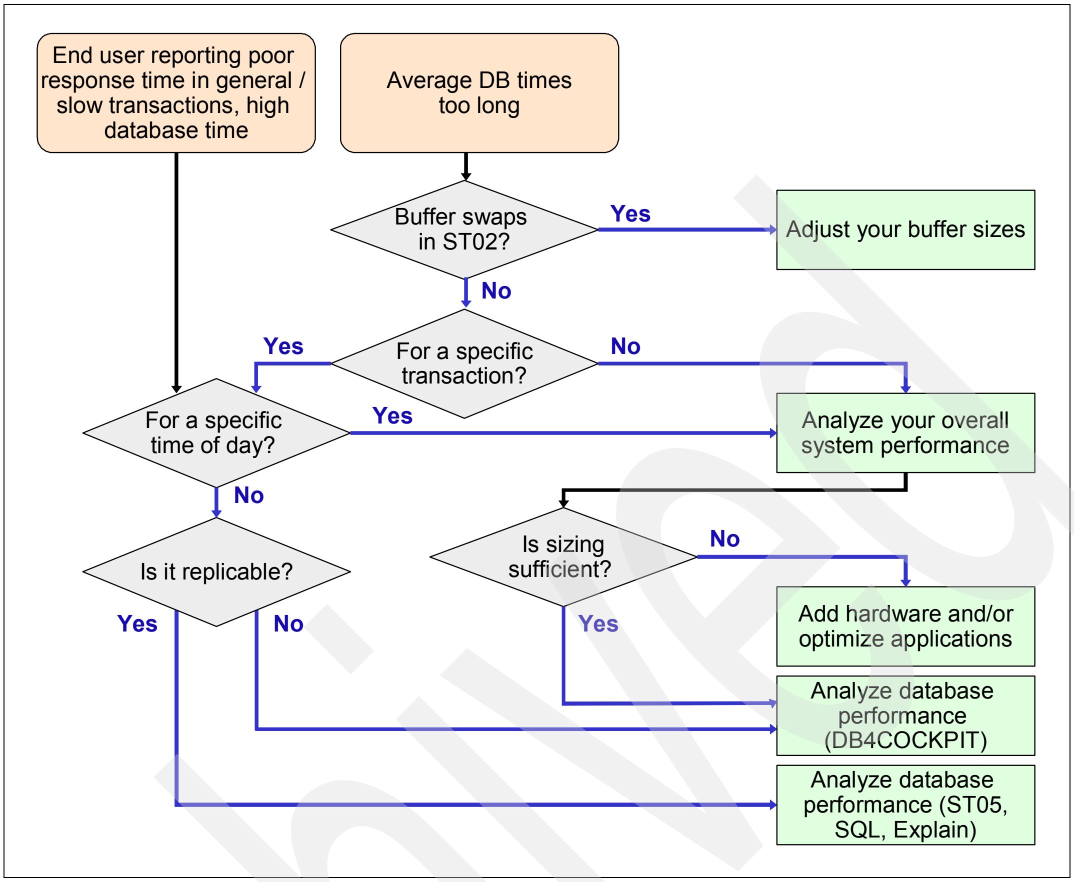

Figure 26-5 An approach to analyze database performance

Perform the following actions if you suspect a database bottleneck:

- For high average database times, check buffer sizes and adjust if necessary.
- Isolate a specific transaction, if possible.
- Determine if the problem occurs in a specific time, for example, in a peak workload situation.
- Determine if the sizing is sufficient for peak workload situations. Use i5/OS tools to monitor performance over a short time period.
- If the user problem or transaction can be repeated, use ST05 and SQL. Refer to step 5 and 6 in Figure 26-3 on page 539 to find long-running database statements.

Refer to 26.1.1, “Tuning an SAP system” on page 539 and 26.2, “Analyzing and optimizing single transactions” on page 554 for information about optimizing your database and applications.

Buffers

Internal buffers are used in the implementation of SAP applications. The buffers are used to store frequently used data that typically remains unchanged, such as the following:

- ABAP programs
- ABAP dictionary data
- Screens
- Company specific data

Keep as much data buffered as possible. This helps avoid repetitive database accesses and improves system performance.

Reduced performance occurs when buffers are too small. The required data cannot be stored in the buffers. Instead, objects have to be swapped out of the buffers. This causes costly response-time database accesses. On the other hand, if buffers are set too large, memory is inefficiently utilized.

Periodically check whether the buffer size is suited to your system requirements. Since buffers are important for the performance of the SAP system, adjust all buffers to their optimum value. Increasing the buffer size increases main storage consumption.

The optimum size for each buffer depends largely on the specific configuration of the server, that is, the applications, the number of users working in each module, and so on. Therefore, it is difficult to specify values suitable for all configurations. The criteria for the correct buffer size is the significance of buffer quality and avoidance of swapping memory. Typically, once a stable operating environment is reached, the buffer hit ratio should be 95% or more.

Buffer sizes are defined in the instance profiles. Following are the different groups of buffers:

1. Repository buffers (Nametab Buffers)

   The name table (nametab) contains the table and field definitions that are activated in the SAP system. The repository buffer consists of buffers for the following:

   – Table definitions (TTAB)
   – Filed descriptions (FTAB)
   – Initial record layouts (IRBD)
   – Short nametab (SNTAB)

   The qualities of the repository buffers can reach 99.9% in a system that has run for a few days. Investigate the situation further if the buffer quality is less than 95%.

2. Table buffers

   The partial table buffer (TABLP) stores single table entries (one record) with its field values. The generic table buffer (TABL) stores a range of table entries, that is, a range of records with their field values. The generic table buffer can also store all the entries (records) in a table. This is known as *resident buffering*, or *full buffering*.

   Use the **SE13** transaction to specify whether or not to buffer a table. Buffer tables that are changed infrequently.

   Changes to buffered information must be updated in the buffers of other application servers that share this information. These updates can adversely affect performance. The more servers that need updating, the more expensive the update process.

   If all operations are performed on a central server with a single instance, deactivate the buffer update messaging service using the **`rdisp/bufrefmode = sendoff,exeauto`** command.

   The quality of the single key buffer increases slowly from the time the system starts up. Therefore, bad quality (that is, less than 90%) should only be of concern if there is no free space left in the buffer.

   The quality of the generic key buffer should be greater than 96% and can go up to 99%.

3. Program buffers

   Program buffer (Program Execution Area - PXA) stores the compiled executable versions of ABAP programs (loads).

4. SAP graphical user interface (GUI) buffers

The presentation buffer (PRES) stores the generated screens (DYNPRO loads). The menu buffer (CUA®) stores objects from the SAP GUI, for example, menus and push button definitions.

5. Roll and paging buffers

   The roll and paging buffer stores among other user contexts and large lists.

6. Calender buffers

   The SAP calendar buffer (CALE) stores all defined factory and public holiday calendars.

7. Export/import buffer

   Import / export buffers (EIBUF) are used to store data that must be available to several work processes.

   As of SAP Basis Release 4.6B, the ABAP List Viewer (ALV) uses the export/import buffer to improve performance. The objects written by the ALV to the export/import buffer are not removed automatically again from the buffer, but are retained in the buffer until the system shuts down. An increased number of swaps can occur if several applications use the import/export buffer or if the buffer size selected is too small. The hit ratio should be higher than 80%.

If you find bad hit ratio and swaps in transaction ST02, enlarge the corresponding parameter. In most cases, there is a value for a number of entries and another value for the size of a buffer. Check the free space (%) and free directory entries (%) to determine which value is too small. Change only one value to be able to ascertain the effects of this change. Changes are in effect after an SAP system restart. Therefore, the steps of buffer optimizing last a longer time.

Note: Tune buffers for all SAP system instances on all servers.

For a detailed view of the transaction ST02, refer to 25.5.11, “Tune summary and buffers (ST02)” on page 511 in 25.5.11, “Tune summary and buffers (ST02)” on page 511.

For more information search for SAP system buffers on the Web at:

`http://help.sap.com`

For more information refer to the following:

- *SAP note 101113: AS/400: Analysis of performance problems*
- *SAP note 123418: AS/400: Performance 4.0B on AS/400*
- *SAP note 121625: AS/400: buffer size*
- *SAP note 373986: Overflow of the export/import buffer*

Transactions

To check the transactions (step 5 in Figure 26-3 on page 539), refer to 26.2, “Analyzing and optimizing single transactions” on page 554.

Specific users and reports

To check specific users and reports (step 6 in Figure 26-3 on page 539), refer to 26.1.1, “Tuning an SAP system” on page 539 and 26.2, “Analyzing and optimizing single transactions” on page 554.

26.1.2 Database tuning

There are multiple factors that influence database performance, such as the management of the physical data, the maintenance of database statistics, and many programming considerations. The storage management component of the i5/OS operating system manages the physical layout of the data on the disk and in the main storage when using DB2 UDB for i5/OS. A system administrative function is to help ensure there are adequate hardware resources for these functions. You do not need to worry about low-level programming considerations, because this is all handled in the SAP kernel. Performance problems encountered that can be solved in low-level programming are handled by SAP programmers. Install the **LIB_DBSL** patch.

> **Note:** Sometimes the SAP kernel patch level affects performance.

Create an appropriate index to improve performance in other areas, such as statistics collection, query optimization, or application development. Additional statistics can help the i5/OS SQL engine be more efficient by creating indexes. Use additional indexes for faster access methods to optimize queries and speed up application programs.

Use the database (DB) monitoring tasks that are described in 25.6, “Database monitoring” on page 516 to identify statements that can be improved.

> **Note:** Identify statements that can be improved with database monitoring.

Example 26-3 shows how to improve performance-expensive statements.

Example 26-3 Sample report ZFLIGHT_SQLTEST1

```
*&---------------------------------------------------------------------*
*& Report  ZFLIGHT_SQLTEST1
*&---------------------------------------------------------------------*

REPORT  ZFLIGHT_SQLTEST1.

DATA: BEGIN OF wa,
  carrid   TYPE sflight-carrid,
  fldate   TYPE sflight-fldate,
  seats    TYPE sflight-seatsmax,
  price    TYPE sflight-price,
END OF wa.

SELECT f~carrid f~fldate f~price sum( f~seatsmax ) as seats
  INTO CORRESPONDING FIELDS OF wa
  FROM ( spfli AS p
            INNER JOIN sflight AS f ON p~carrid = f~carrid AND
                                       p~connid = f~connid    )
 WHERE p~cityfrom = 'FRANKFURT'
AND p~cityto   = 'NEW YORK'
   AND f~fldate   BETWEEN '20050701' AND '20050831'
 GROUP BY f~carrid f~fldate f~price
 ORDER BY f~fldate seats descending.

WRITE: / wa-carrid, wa-fldate, wa-seats, wa-price.
ENDSELECT.
```

Figure 26-6 shows an SQL sample to demonstrate some elements of the SQL language that you need to understand in order to optimize your queries.

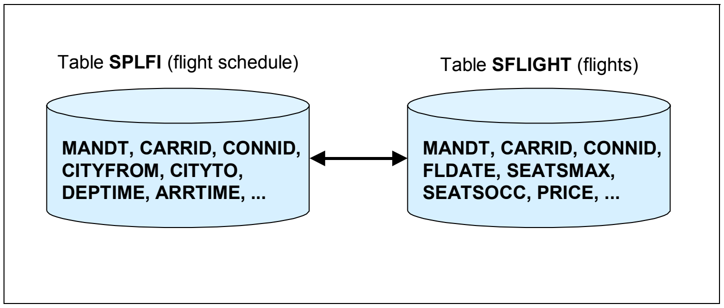

Figure 26-6 Database tables SPLFI and SFLIGHT

The following scenario illustrates this function.

Assume a travel agent wants to arrange a trip from Frankfurt to New York for a time between 01 July and 31 August. All of the participants can use the same airline, on the same day, but they do not need to use the same flight connection. You look for the airline and day with the most available seats.

The flight data is spread over multiple tables. Use table **SPFLI** to find the departure and destination city, and use table **SFLIGHT** to find the flight dates and the number of seats for each flight. Sum up the available seats if the same airline offers several flights on the same day.

Following are the key elements in the illustrated statement to optimize the execution:

- ► Join (combine the data from two or more tables)
- ► Record selection (the WHERE clause)
- ► Group (GROUP BY)
- ► Order (ORDER BY)

The output of the statement can look like what Figure 26-7 illustrates.

ZFLIGHT_SQLTEST 1

ZFLIGHT_SQLTEST 1

LH	02.07.2005	280	666,00
UA	25.07.2005	385	611,01
LH	25.07.2005	380	666,00
LH	28.07.2005	660	666,00
UA	28.07.2005	385	611,01
LH	30.07.2005	280	666,00
UA	22.08.2005	385	611,01
LH	22.08.2005	380	666,00
LH	27.08.2005	280	666,00
LH	28.08.2005	660	666,00
UA	28.08.2005	385	611,01

Figure 26-7 Report ZFLIGHT_SQLTEST1 results

The following are the few differences when comparing the SQL statement that is sent to the database with the ABAP syntax of the statement:

- ► Table names are qualified by a library (or schema) name (A).
- ► The join conditions are moved to the WHERE clause of the statement (B).
- ► Column MANDT is added to the WHERE clause and compared to the current client (C).

- Constants in the statement are replaced by variables, which are represented by `?` in the statement text (D).
- The SQL statement on the database is completely different from the ABAP statement when accessing Pool and Cluster tables.

Use the database syntax, not the ABAP syntax, for query optimization. The SQL trace (ST05) and the database monitor (DB4COCKPIT) always show the database syntax.

Note: Use the database syntax for query optimization.

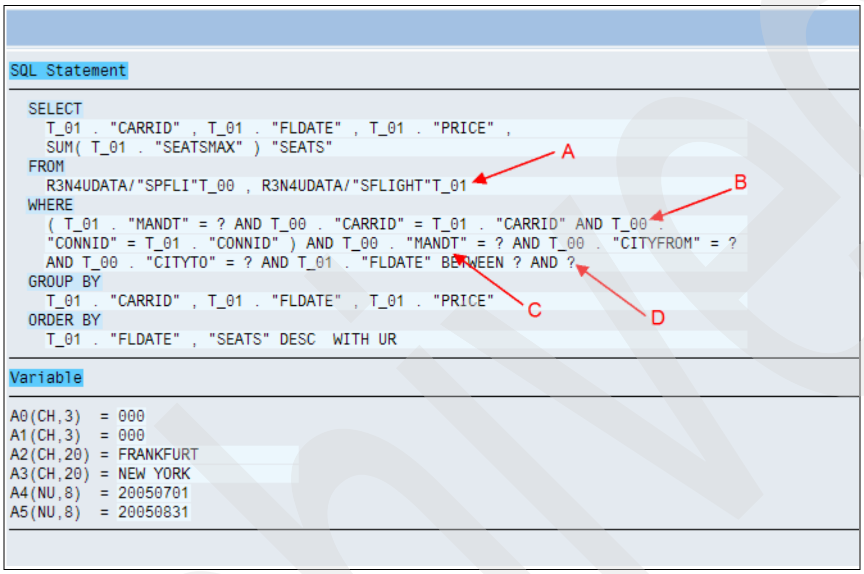

Figure 26-8 SQL Trace (ST05) of `ZFLIGHT_SQLTEST1`

The query optimizer for DB2 UDB on iSeries is a cost-based optimizer. It bases decisions on many environmental factors. The goal of the optimization is to come up with the fastest possible implementation. You can set an optimization goal of `first I/O` for traditional applications, so that the first screen of data, that is, the first 30 rows, is returned as quickly as possible, even though the total runtime can be slower. The SAP system use the optimization goal of all I/O because the SAP application usually requests only the data it needs to show on the display.

Sometimes indexes are not used as a result of the cost-based optimizer, even though they appear to be helpful. It can be faster to read the whole table in large blocks, rather than use an index and read the data row by row, if the table size is relatively small, or if a significant percentage of the data in the table is returned to the application.

IBM introduced a new query optimizer called *SQE*, or the *SLIC Query Engine*, in OS/400 V5R2. It replaces the optimizer known as the *Classic Query Engine* (CQE) step by step. Both CQE and SQE are cost-based optimizers. Both use indexes to obtain statistical information about the tables. SQE uses additional statistics that are not bound to indexes and are collected automatically.

SQE provides significant enhancements for certain types of SQL statements. However, not all SQL statements can go through SQE with i5/OS V5R3. Some go through CQE. In future releases, SQE is enhanced so that more statements can go through SQE.

The cost-based query optimizer tries to find the fastest access plan depending on the following:

- SQL statement
- Table size
- Available indexes
- Optimization goal (first I/O or all I/O)
- CPU speed and disk speed
- Amount of data returned
- Fair share of main memory

The query engine first selects records.

Query Optimization: Select records

Record selection limits the amount of returned data by allowing only certain values for selected columns, such as `FRANKFURT` for column `CITYFROM` in table SPFLI. The query optimizer can choose to perform a so-called table scan. The whole table is read into the memory in this case from the database. In the memory, the records that do not match the selection are discarded. The optimizer can also locate matching rows quickly through the index and only return rows that match the selection if an index exists for the column or columns in the WHERE clause. The optimizer can choose not to use an index, even though it exists, if the table is small or if a large percentage of the rows in the table are returned.

Figure 26-9 illustrates these concepts.

```
SQL Statement

SELECT
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE" ,
  SUM( T_01 . "SEATSMAX" ) "SEATS"
FROM
  R3N4UDATA/"SPFLI"T_00 , R3N4UDATA/"SFLIGHT"T_01
WHERE
  ( T_01 . "MANDT" = ? AND T_00 . "CARRID" = T_01 . "CARRID" AND T_00 .
  "CONNID" = T_01 . "CONNID" ) AND T_00 . "MANDT" = ? AND T_00 . "CITYFROM" = ?
  AND T_00 . "CITYTO" = ? AND T_01 . "FLDATE" BETWEEN ? AND ?
GROUP BY
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE"
ORDER BY
  T_01 . "FLDATE" , "SEATS" DESC  WITH UR
```

Figure 26-9 Query: record selection

> **Note:** For implementation, the query optimizer makes the decision between a table scan and index access.

Following is a review of the steps for optimization.

Query optimization: Join condition

You must join two or more tables if the required data results from more than one table. The most common join is an *inner join* where all rows in the first table can be combined with all rows in the second table (and subsequent tables) where the join fields match. Rows in one table that do not have corresponding rows in the other tables are not returned. An Open SQL statement in ABAP also supports the *left outer join*, where all the rows of the first table are returned, and the values from the second table are either filled with values from the second table, if a matching row was found, or set to null. In the case of an inner join, the query optimizer can change the order of the tables if it estimates a shorter execution time.

A *nested loop* is a concept where the database reads all the rows from the first table (reduced by the record selection if one exists), and tries to find matching rows in the second table. The nested loop can be implemented by using the index if an index exists over the join columns (CARRID and CONNID, as in our example illustrated in Figure 26-10). This is usually faster.

```
SQL Statement

SELECT
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE" ,
  SUM( T_01 . "SEATSMAX" ) "SEATS"
FROM
  R3N4UDATA/"SPFLI"T_00 , R3N4UDATA/"SFLIGHT"T_01
WHERE
  ( T_01 . "MANDT" = ? AND T_00 . "CARRID" = T_01 . "CARRID" AND T_00 .
  "CONNID" = T_01 . "CONNID" ) AND T_00 . "MANDT" = ? AND T_00 . "CITYFROM" = ?
  AND T_00 . "CITYTO" = ? AND T_01 . "FLDATE" BETWEEN ? AND ?
GROUP BY
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE"
ORDER BY
  T_01 . "FLDATE" , "SEATS" DESC  WITH UR
```

Figure 26-10 Query: join condition

The optimizer can either create a temporary index, a hash table, or a sorted list (in case of inequality) to locate matching rows in the second table if no index exists over the join columns. All three methods take some time, depending on the table size.

Note: Create indexes over the join columns to significantly improve join performance.

If the created index is not used, it could, nevertheless, help with additional statistical information to find the optimum join order.

Note: The optimizer can change the order of tables for an inner join. The implementation methods include a *nested loop via index*, *nested loop via hash table* or a *nested loop via sorted* list.

Query optimization: Group By

The *Group By* clause in an SQL statement, illustrated in Figure 26-11 on page 550, is used to aggregate values for all rows that have the same key. This includes the following:

- Summing up the values (SUM)
- Finding the minimum or maximum value (MIN or MAX)
- Calculating the average (AVG)
- Counting rows (COUNT)

The optimizer needs to group all rows with the same values in the columns that are specified in the Group By clause in order to process the data. An index over those columns can be used for the implementation if an index exists. The optimizer can choose to build a temporary index or a temporary hash table to implement the Group By if an index is not available. It is usually helpful to have an index as both methods take some time and require resources.

```
SQL Statement

SELECT
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE" ,
  SUM( T_01 . "SEATSMAX" ) "SEATS"
FROM
  R3N4UDATA/"SPFLI"T_00 , R3N4UDATA/"SFLIGHT"T_01
WHERE
  ( T_01 . "MANDT" = ? AND T_00 . "CARRID" = T_01 . "CARRID" AND T_00 .
  "CONNID" = T_01 . "CONNID" ) AND T_00 . "MANDT" = ? AND T_00 . "CITYFROM" = ?
  AND T_00 . "CITYTO" = ? AND T_01 . "FLDATE" BETWEEN ? AND ?
GROUP BY
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE"
ORDER BY
  T_01 . "FLDATE" , "SEATS" DESC  WITH UR
```

Figure 26-11 Query: Group By

Note: Possible implementation methods are indexes or temporary hash tables.

Query optimization: Order By

The *Order By* clause is used to return rows in a specific order. The order of rows is unpredictable without it. An index over the specified columns can be used to implement the Order By request if an index exists. The optimizer can decide to build a temporary index if an index does not exist. Another option is to first fetch all rows in random order, and then sort the result set. This is necessary if the Order By clause contains calculations, such as SUM or AVG, as illustrated in Figure 26-12.

```
SQL Statement

SELECT
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE" ,
  SUM( T_01 . "SEATSMAX" ) "SEATS"
FROM
  R3N4UDATA/"SPFLI"T_00 , R3N4UDATA/"SFLIGHT"T_01
WHERE
  ( T_01 . "MANDT" = ? AND T_00 . "CARRID" = T_01 . "CARRID" AND T_00 .
  "CONNID" = T_01 . "CONNID" ) AND T_00 . "MANDT" = ? AND T_00 . "CITYFROM" = ?
  AND T_00 . "CITYTO" = ? AND T_01 . "FLDATE" BETWEEN ? AND ?
GROUP BY
  T_01 . "CARRID" , T_01 . "FLDATE" , T_01 . "PRICE"
ORDER BY
  T_01 . "FLDATE" , "SEATS" DESC WITH UR
```

Figure 26-12 Query: Order By

Note: Possible implementation methods are indexes or a sort of the results.

Query optimization: Combine the table columns to an index

The optimizer can pick a maximum of one index per table to implement record selection, join conditions, group by, and order by when the query gets executed. Therefore, combine columns that are helpful for the different parts into one index.

You must make assumptions about which parts of the index have the biggest impact on performance because it is impossible to build one index that covers all the different elements. The efficiency of an index depends on the SQL statement, as well as the amount and distribution of the data in the tables.

The best way to determine which of the possible indexes works best is to create them all, collect an SQL trace over the statement, use **Explain** to check which one is picked, and drop the other ones again. Having too many indexes can cause a slowdown of INSERT, UPDATE, or DELETE operations because it adds index maintenance overhead.

Precede those columns by columns from the WHERE clause that return exactly one value (= comparison without OR) when using join conditions, Group By, or Order By columns in order to design an index. Therefore, the ideal index on table SPFLI for the given statement combines MANDT, CITYFROM, and CITYTO from the where clause with CARRID and CONNID from the join condition.

Multiple indexes can be helpful for SFLIGHT. The index with MANDT (for record selection), CARRID, CONNID (for the join), and FLDATE (for further record selection) can be useful when SPFLI is used in join position 1 and SFLIGHT in join position 2. The index with MANDT (for record selection), CARRID, FLDATE and PRICE (for Group By and Order By) is useful when SFLIGHT is used in join position 1. Which join order and index is actually chosen depends on the data in the tables. Therefore, try which one works best.

Table 26-9 Combine the columns of an SQL query

	Columns in table SPFLI	**Columns in Table SFLIGHT**	**Comparison**
Record selection	MANDT CITYFROM CITYTO	 MANDT FLDATE	= = = = BETWEEN
Join	CARRID CONNID	CARRID CONNID	= =
Group By		CARRID FLDATE PRICE	
Order By		FLDATE	

Combine the fields shown in Table 26-9 to get the best possible column for an index.

For the indexes for the tables SPFLI and SFLIGHT, the combination of entries in Table 26-10 are possible.

Table 26-10 Possible key columns for table SPLFI and table SFLIGHT

	table SPFLI	**table SFLIGHT**
Key columns for table	MANDT CITYFROM CITYTO CARRID CONNID	MANDT CARRID CONNID FLDATE
Key columns for table (alternative)		MANDT CARRID FLDATE PRICE

The best way to determine which of the possible indexes works best, is to create them all, collect an SQL trace (ST05) over the statement, use **Explain** to check which one is picked from the optimizer, and drop the unused index again. Having too many indexes can cause a slowdown of INSERT, UPDATE, or DELETE operations because it adds index maintenance costs.

Table SPFLI is in join position 1 in the sample report ZFLIGHT_SQLTEST1. The optimizer selected SPFLI0001 as the index.

The results of the SQL Explain operation appear as in Figure 26-13.

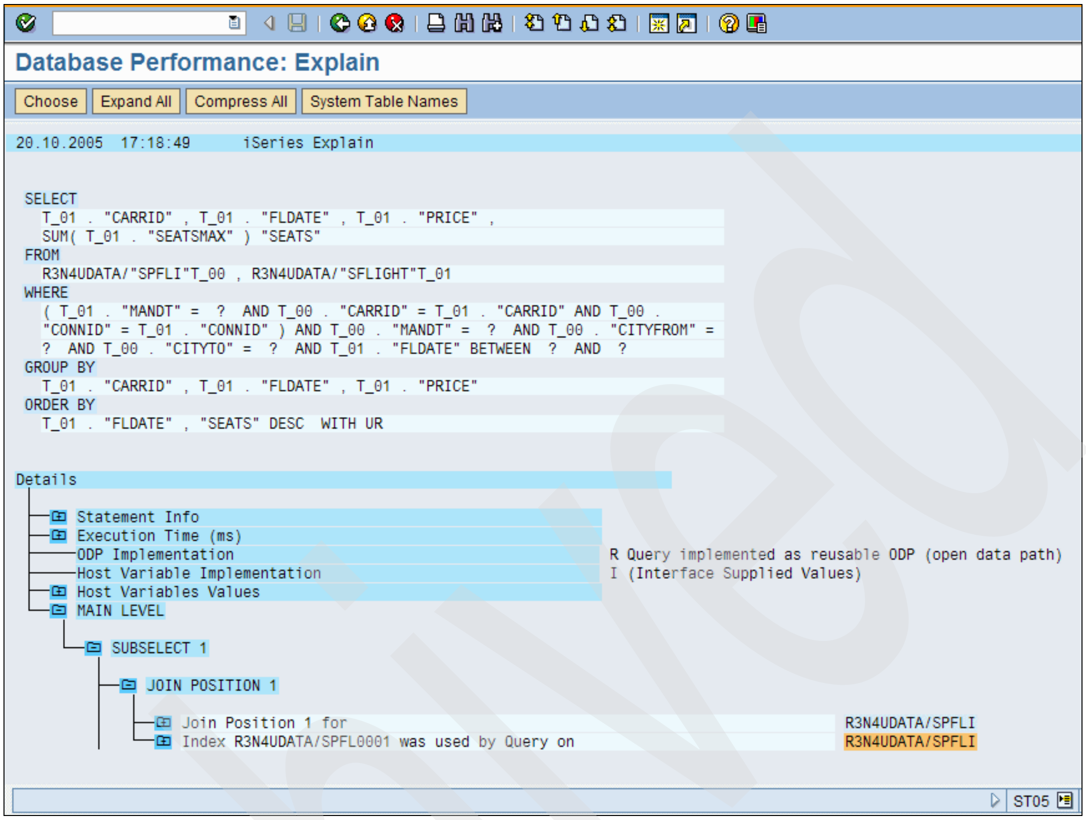

Figure 26-13 SQL explain results

Note: More indexes speed up queries, however, more overhead is required to maintain the additional indexes. Combine helpful indexes for multiple cases.

Query optimization: working with views

You can transform a statement with a view into a statement that only uses tables and the elements that are discussed above. A view in the SAP dictionary consists of three elements:

- The tables and their join conditions
- The view fields (usually a subset of the underlying tables)
- The selection conditions (optional)

The view fields are not relevant for the optimizer. However, you can convert the table and join conditions into a simple join, and add the selection conditions to the WHERE clause of your statement. This is what the query optimizer does as well. You can design the best indexes for your query as before after this conversion. Refer to Figure 26-14.

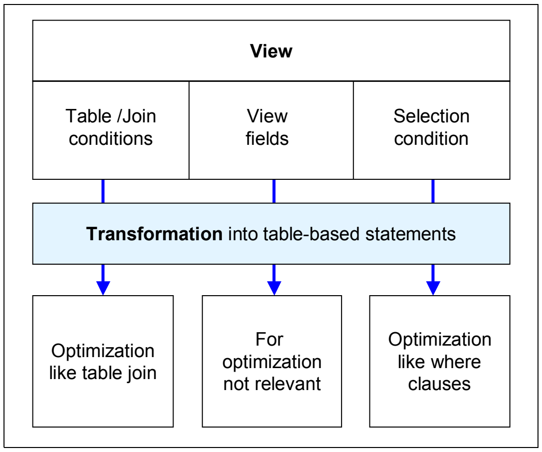

Figure 26-14 Use views to optimize a query

Note: Optimize a query after the transformation into table-based statements.

Query optimization: star join (snowflake) schema

You can see very complex joins that follow the *star* schema or *snowflake* schema especially in the environment of the Business Information Warehouse application (including APO/SCM and SEM). Typically, you have one large fact table that contains all the key figure, and many smaller dimension tables that contain additional values for the data. For example, you can see customer keys only in the fact table, and then you get the customer details from the customer dimension table.

Figure 26-15 illustrates the star join (snowflake) schema.

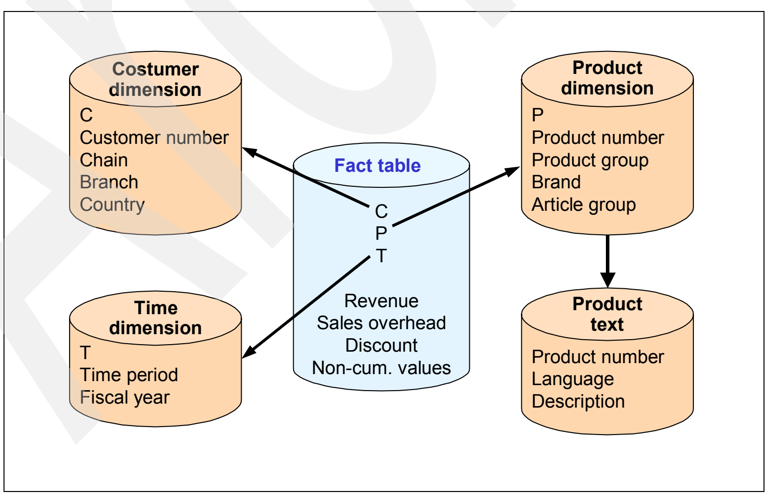

Figure 26-15 Query optimization: Star join (snowflake) schema

The SAP system uses special techniques, such as enhanced vector indexes (EVIs) on the fact table and special options in the query options file QAQQINI in order to help the query optimizer make the best decision for accessing the data in a star join or snowflake schema.

The setup is described in *SAP note 501572 iSeries: EVI stage 2 support.* Use the *Sanity check* tool as described in *SAP note 541508 iSeries: Checking the system parameters for BW* to verify the setup. Additional tuning beyond the contents of the SAP notes is not necessary.

Important: The settings can be dependent on the i5/OS release. Check the SAP notes for potential changes if you upgrade to a newer release.

You can use parallelism to improve performance if you have the *DB2 Symmetric Multiprocessing* feature installed (Option 26 of the operating system) and enough hardware resources.

Refer to *Preparing for and Tuning the V5R2 SQL Query Engine on DB2 Universal Database for iSeries*, SG24-6598 for more information about SQL Query Engine.

26.2 Analyzing and optimizing single transactions

This section shows you how to analyze and optimize single transactions.

Examine a single transaction to analyze slow-running applications. For example, information about slow-running transactions with ABAP programming can come from affected users, or can be found by analyzing statistics, for example, SAP transaction ST03.

Use two SAP-GUI sessions to analyze the transaction. Figure 26-16 shows the time line of the steps to monitor a single transaction.

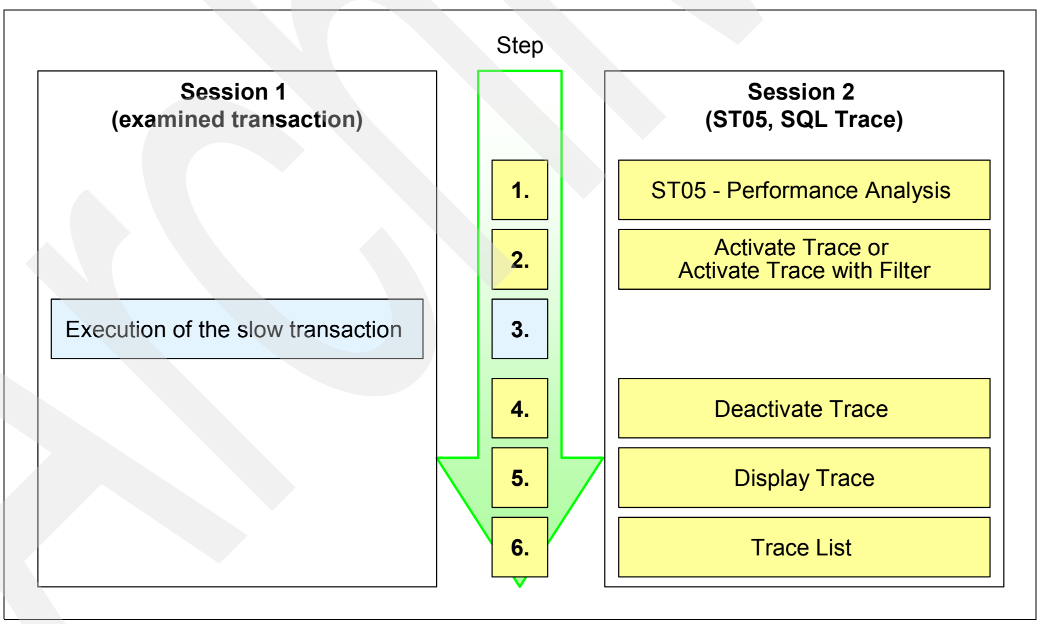

Figure 26-16 Time line: use the SQL trace for ABAP

The individual steps are described in 26.2.1, “Transaction ST05: performance analysis” on page 555 and 26.2.2, “SQL trace” on page 557.

26.2.1 Transaction ST05: performance analysis

The following steps show you how to use a trace to analyze a transaction:

1. Select **Session two** to perform Transaction ST05 as shown in Figure 26-17.

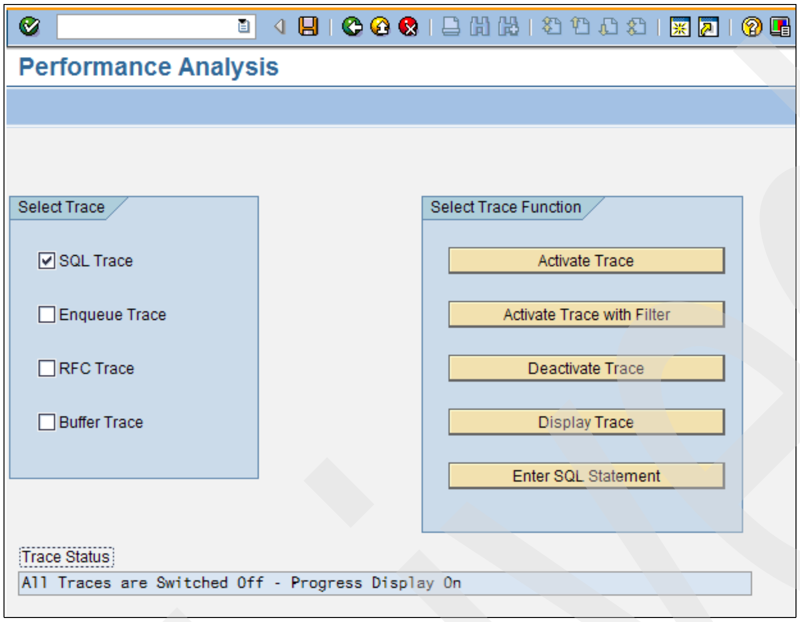

Figure 26-17 Session 2: Transaction ST05

2. Select **Activate Trace** to switch the trace on as Figure 26-18 shows.

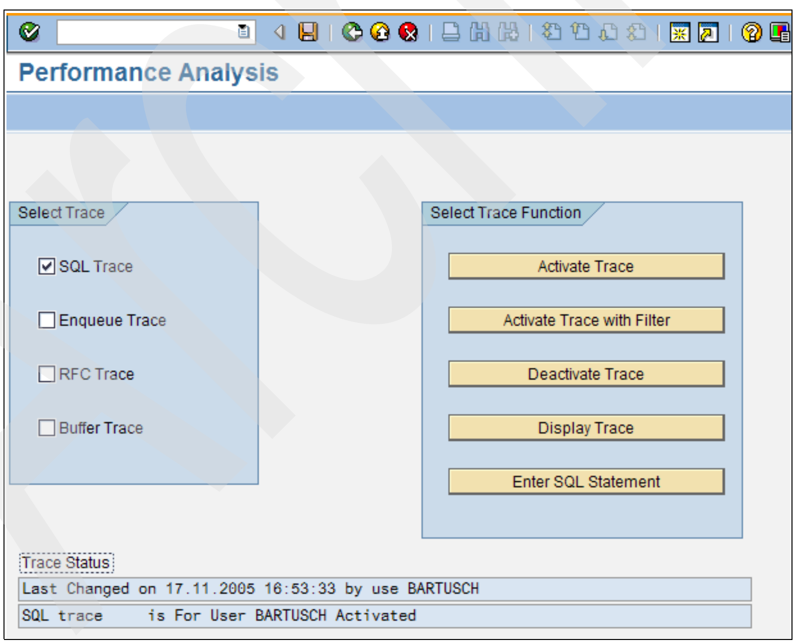

Figure 26-18 Session 2: Activate Trace

The Status message changes from `All traces are switched off` to `SQL Trace is for user` *`user name`* `activated`. Alternatively, use the **Activate Trace with Filter** button to trace another user, a definite table or transaction, or a combination of the possibilities, as shown in Figure 26-19.

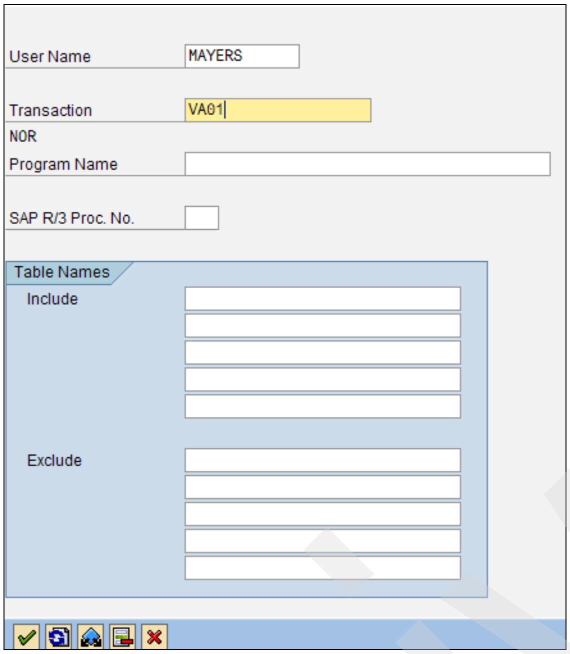

Figure 26-19 Activate Trace with filters: selection screen

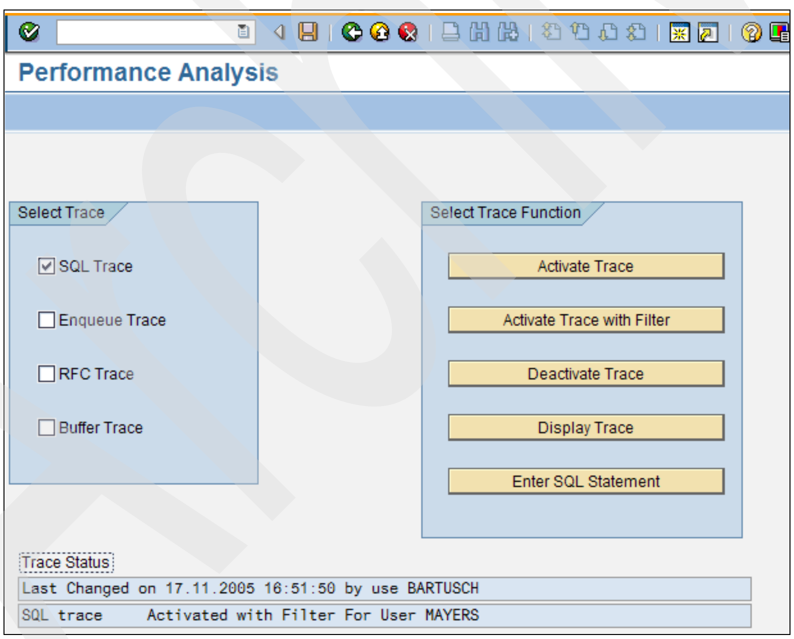

Figure 26-20 Session 2: Activate trace with Filter: status

The status message changes to `SQL Trace Activated with Filter for User` <*specified user*> as Figure 26-20 shows.

3. Change to **Session one** and run the **slow transaction**. Execute the transaction from the affected user. Activate the **Trace with filters - user not equal yourself** to trace another user's transaction. Note the system time and locate the right rows in the trace file.

4. Change to **Session two** and select **Deactivate the trace** as Figure 26-21 shows.

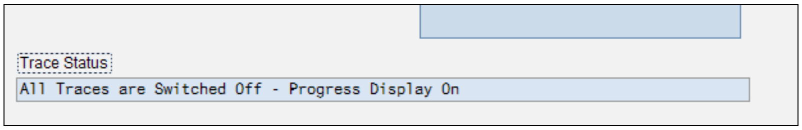

Figure 26-21 Session 2: Deactivate trace: status

The status changes to `All traces are switched off`.

26.2.2 SQL trace

Analyze the trace data after running the slow transaction, as follows:

1. Select **SQL Trace** for the left check boxes, and then select **Display Trace** in the ST05 screen to change to SQL Trace, as illustrated in Figure 26-22.

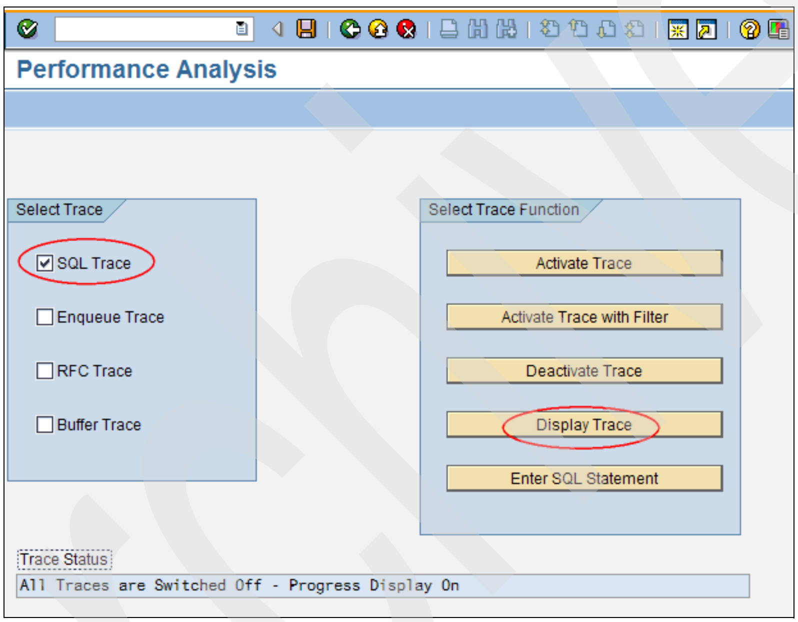

Figure 26-22 Session 2: selection for SQL Trace

2. A report similar to those shown in Figure 26-23 and Figure 26-25 on page 558 appears.

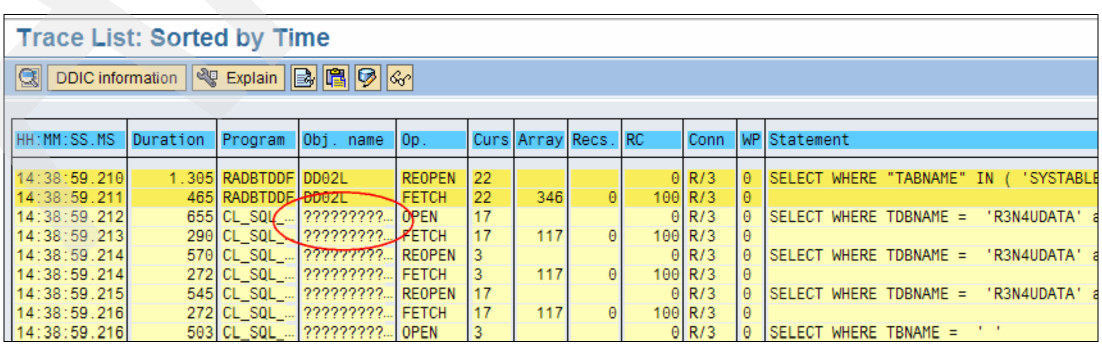

HH:MM:SS.MS	Duration	Program	Obj. name	Op.	Curs	Array	Recs.	RC	Conn	WP	Statement
14:38:59.210	1.305	RADBTDDF	DD02L	REOPEN	22			0	R/3	0	SELECT WHERE "TABNAME" IN ('SYSTABLE
14:38:59.211	465	RADBTDDF	DD02L	FETCH	22	346	0	100	R/3	0	
14:38:59.212	655	CL_SQL_...	?????????...	OPEN	17			0	R/3	0	SELECT WHERE TDBNAME = 'R3N4UDATA' a
14:38:59.213	290	CL_SQL_...	?????????...	FETCH	17	117	0	100	R/3	0	
14:38:59.214	570	CL_SQL_...	?????????...	REOPEN	3			0	R/3	0	SELECT WHERE TDBNAME = 'R3N4UDATA' a
14:38:59.214	272	CL_SQL_...	?????????...	FETCH	3	117	0	100	R/3	0	
14:38:59.215	545	CL_SQL_...	?????????...	REOPEN	17			0	R/3	0	SELECT WHERE TDBNAME = 'R3N4UDATA' a
14:38:59.216	272	CL_SQL_...	?????????...	FETCH	17	117	0	100	R/3	0	
14:38:59.216	503	CL_SQL_...	?????????...	OPEN	3			0	R/3	0	SELECT WHERE TBNAME = ' '

Figure 26-23 Session 2: SQL Trace: an ABAP Trace list (view 1)

The ?????????? entries in the column Obj. name indicate that the name is unknown in the SAP data directory. This is typical for the use of native SQL statements, for example, when accessing system catalog views.

Another possible entry is open data path (ODP) cleanup, as shown in Figure 26-24.

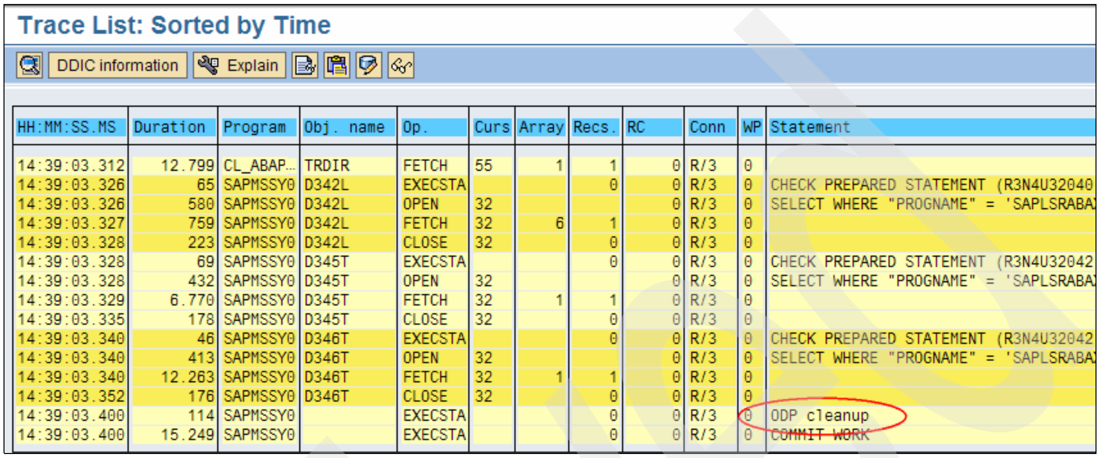

Trace List: Sorted by Time

DDIC information | Explain

HH:MM:SS.MS	Duration	Program	Obj. name	Op.	Curs	Array	Recs.	RC	Conn	WP	Statement
14:39:03.312	12.799	CL_ABAP...	TRDIR	FETCH	55	1	1	0	R/3	0	
14:39:03.326	65	SAPMSSY0	D342L	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT (R3N4U32040
14:39:03.326	580	SAPMSSY0	D342L	OPEN	32			0	R/3	0	SELECT WHERE "PROGNAME" = 'SAPLSRABA
14:39:03.327	759	SAPMSSY0	D342L	FETCH	32	6	1	0	R/3	0	
14:39:03.328	223	SAPMSSY0	D342L	CLOSE	32		0	0	R/3	0	
14:39:03.328	69	SAPMSSY0	D345T	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT (R3N4U32042
14:39:03.328	432	SAPMSSY0	D345T	OPEN	32			0	R/3	0	SELECT WHERE "PROGNAME" = 'SAPLSRABA
14:39:03.329	6.770	SAPMSSY0	D345T	FETCH	32	1	1	0	R/3	0	
14:39:03.335	178	SAPMSSY0	D345T	CLOSE	32		0	0	R/3	0	
14:39:03.340	46	SAPMSSY0	D346T	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT (R3N4U32042
14:39:03.340	413	SAPMSSY0	D346T	OPEN	32			0	R/3	0	SELECT WHERE "PROGNAME" = 'SAPLSRABA
14:39:03.340	12.263	SAPMSSY0	D346T	FETCH	32	1	1	0	R/3	0	
14:39:03.352	176	SAPMSSY0	D346T	CLOSE	32		0	0	R/3	0	
14:39:03.400	114	SAPMSSY0		EXECSTA			0	0	R/3	0	ODP cleanup
14:39:03.400	15.249	SAPMSSY0		EXECSTA			0	0	R/3	0	COMMIT WORK

Figure 26-24 Session 2: SQL trace: an ABAP trace list (view 2)

Usually the system cleans up open cursors at the commit time, as seen in Figure 26-24. You can also see an ODP cleanup entry between two database operations if many new cursors are opened within the same SAP transaction. Long ODP times in general indicate a lack of main memory. This can show up with an increase in page fault rates in the operating system.

The main task is to search and to explain the long-running statements, outlined in red in the Duration column, as seen in Figure 26-25.

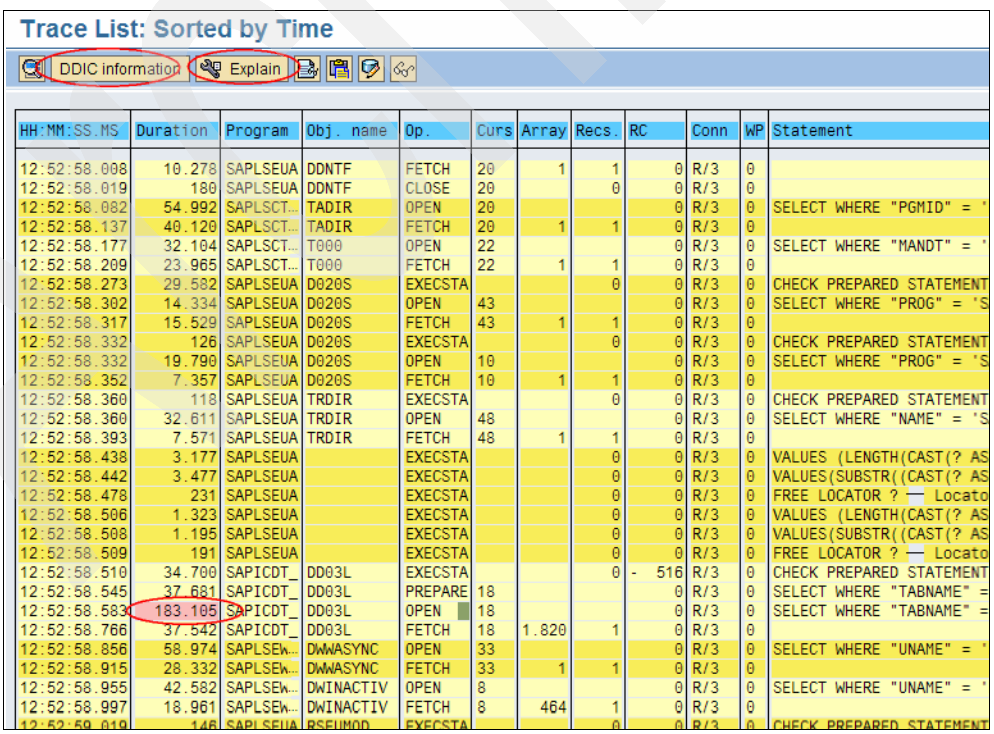

Trace List: Sorted by Time

DDIC information | Explain

HH:MM:SS.MS	Duration	Program	Obj. name	Op.	Curs	Array	Recs.	RC	Conn	WP	Statement
12:52:58.008	10.278	SAPLSEUA	DDNTF	FETCH	20	1	1	0	R/3	0	
12:52:58.019	180	SAPLSEUA	DDNTF	CLOSE	20		0	0	R/3	0	
12:52:58.082	54.992	SAPLSCT...	TADIR	OPEN	20			0	R/3	0	SELECT WHERE "PGMID" = '
12:52:58.137	40.120	SAPLSCT...	TADIR	FETCH	20	1	1	0	R/3	0	
12:52:58.177	32.104	SAPLSCT...	T000	OPEN	22			0	R/3	0	SELECT WHERE "MANDT" = '
12:52:58.209	23.965	SAPLSCT...	T000	FETCH	22	1	1	0	R/3	0	
12:52:58.273	29.582	SAPLSEUA	D020S	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT
12:52:58.302	14.334	SAPLSEUA	D020S	OPEN	43			0	R/3	0	SELECT WHERE "PROG" = 'S
12:52:58.317	15.529	SAPLSEUA	D020S	FETCH	43	1	1	0	R/3	0	
12:52:58.332	126	SAPLSEUA	D020S	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT
12:52:58.332	19.790	SAPLSEUA	D020S	OPEN	10			0	R/3	0	SELECT WHERE "PROG" = 'S
12:52:58.352	7.357	SAPLSEUA	D020S	FETCH	10	1	1	0	R/3	0	
12:52:58.360	118	SAPLSEUA	TRDIR	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT
12:52:58.360	32.611	SAPLSEUA	TRDIR	OPEN	48			0	R/3	0	SELECT WHERE "NAME" = 'S
12:52:58.393	7.571	SAPLSEUA	TRDIR	FETCH	48	1	1	0	R/3	0	
12:52:58.438	3.177	SAPLSEUA		EXECSTA			0	0	R/3	0	VALUES (LENGTH(CAST(? AS
12:52:58.442	3.477	SAPLSEUA		EXECSTA			0	0	R/3	0	VALUES(SUBSTR((CAST(? AS
12:52:58.478	231	SAPLSEUA		EXECSTA			0	0	R/3	0	FREE LOCATOR ? — Locato
12:52:58.506	1.323	SAPLSEUA		EXECSTA			0	0	R/3	0	VALUES (LENGTH(CAST(? AS
12:52:58.508	1.195	SAPLSEUA		EXECSTA			0	0	R/3	0	VALUES(SUBSTR((CAST(? AS
12:52:58.509	191	SAPLSEUA		EXECSTA			0	0	R/3	0	FREE LOCATOR ? — Locato
12:52:58.510	34.700	SAPICDT_	DD03L	EXECSTA			0	- 516	R/3	0	CHECK PREPARED STATEMENT
12:52:58.545	37.681	SAPICDT_	DD03L	PREPARE	18			0	R/3	0	SELECT WHERE "TABNAME" =
12:52:58.583	183.105	SAPICDT_	DD03L	OPEN	18			0	R/3	0	SELECT WHERE "TABNAME" =
12:52:58.766	37.542	SAPICDT_	DD03L	FETCH	18	1.820	1	0	R/3	0	
12:52:58.856	58.974	SAPLSEW...	DWWASYNC	OPEN	33			0	R/3	0	SELECT WHERE "UNAME" = '
12:52:58.915	28.332	SAPLSEW...	DWWASYNC	FETCH	33	1	1	0	R/3	0	
12:52:58.955	42.582	SAPLSEW...	DWINACTIV	OPEN	8			0	R/3	0	SELECT WHERE "UNAME" = '
12:52:58.997	18.961	SAPLSEW...	DWINACTIV	FETCH	8	464	1	0	R/3	0	
12:52:59.019	146	SAPLSEUA	RSEUMOD	EXECSTA			0	0	R/3	0	CHECK PREPARED STATEMENT

Figure 26-25 Session 2: SQL Trace: an ABAP Trace list (view 3)

A typical sequence for a SELECT statement consists of the following:

1. CHECK PREPARED STATEMENT (try to find the statement in a SQL package)
2. OPEN
3. One or more FETCHes (bring data)

Use the **Explain** button and a report similar to that shown in Figure 26-26 appears.

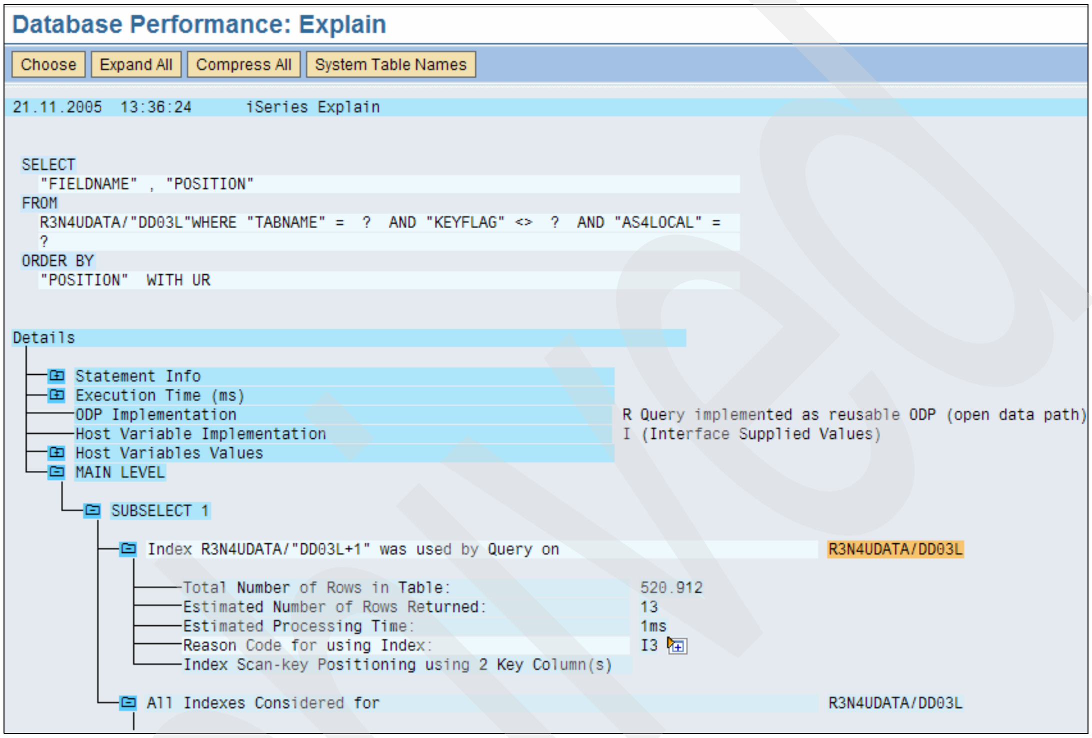

Figure 26-26 ABAP SQL Trace: Explain

The SQL statement that was sent to the database is shown on the top of the display. This statement can be different to the one in the ABAP coded statement.

The statement text is limited in characters in this report. Go back to the SQL trace list and double-click the statement if the statement is incomplete. This view is not limited. Expand the sections that are marked with +, to see more information. The Statement Info and Execution Time is not that useful. The statement is prepared again, but into an SQL package in the QTEMP library when you execute the **`SQL Explain`** from the SQL trace.

The Host Variables Values replace the `?` in the statement text. Expand the **Host Variables Values** section to see the values expressed as hexadecimal values. The values in the WHERE clause are generally not sent to the database as part of the statement text in order to make the statements reusable. They are used as a variable.

The MAIN LEVEL section shows the actually implementation of the statement. The type and amount of information depends on the statement and the kind of implementation. Use the **DDIC Information** button to see other helpful information, as shown in Figure 26-25 on page 558. Click **DDIC Information** and a report similar to that shown in Figure 26-27 on page 560 is displayed.

The index used in this example is DD03L+1. You can find this index under Indexes of DD03L in the row with Name equalling 1.

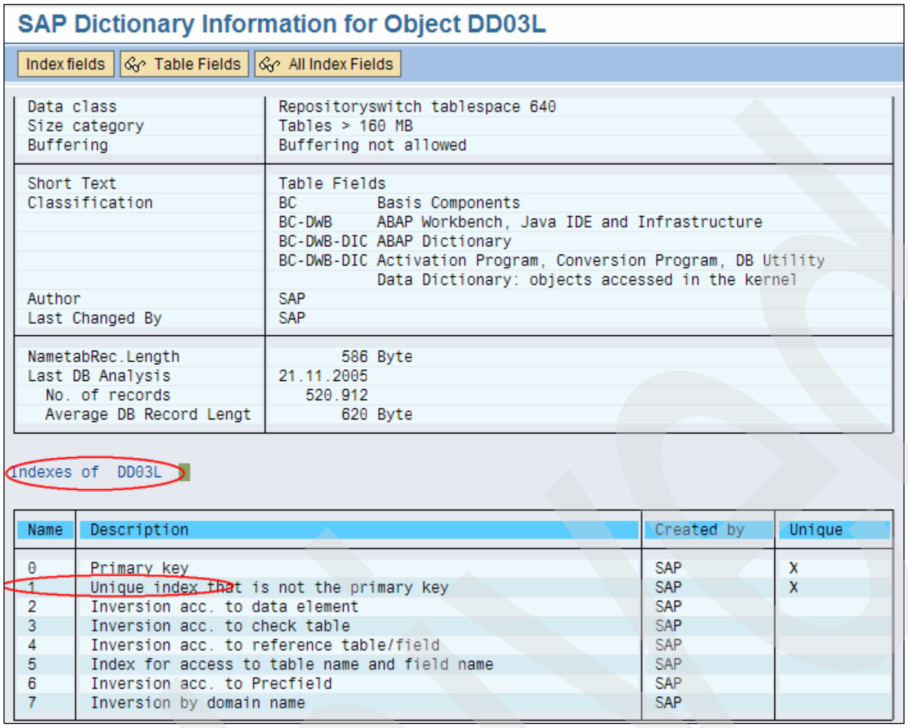

Figure 26-27 SQL Trace: DDIC Information

Click **used index DD03L+1** as shown in Figure 26-28.

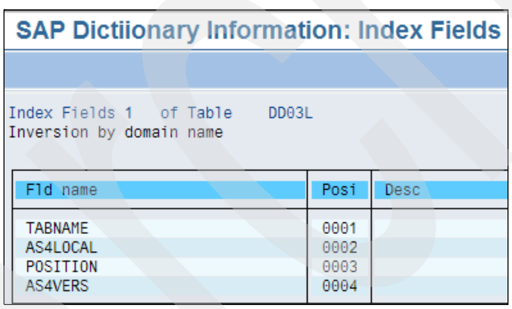

Figure 26-28 DDIC Information for Index DD03L+1

This is the easiest way to check the fields of the used index.

You can optimize the SQL statement with this information, as described in 26.1.1, "Tuning an SAP system" on page 539.

> **Note:** Use this transaction trace to find and optimize slow-running SQL statements.

The Explain function prepares the statement and opens the cursor (without fetching data) in order to collect information about the implementation of that statement when executing the SQL trace for SAP WebAS ABAP. You can use Explain functions only for read-only statements, not for INSERT, UPDATE, or DELETE operations because the statement is re-run.

You can implement the statement differently using the Explain function as compared to the original implementation. It is possible to map the SQL statements to the place in the ABAP source where they are executed. This helps to tune the application by modifying the SQL statement. The ABAP SQL trace can list and explain statements.

The database monitor is typically active whenever the SAP system is operational. You can control this through the as4/dbmon/enable profile parameter. The database monitor (transaction ST04) aggregates the runtimes of identical statements in order to keep the amount of collected data within a reasonable size. In contrast to the SQL trace (transaction ST05), it only keeps implementation details of the most expensive execution.

The implementation that was actually used is seen. However, the slower implementation still takes place for the most expensive statement, even if the indexes are added for a specific statement to improve runtime. *SAP note 727078* explains how you can clear the database monitor cache, so that the current implementation of a statement turns out to be the most expensive one.

The database monitor is also activated for an SAP WebAS Java server and can be analyzed from the associated SAP WebAS ABAP server through the DB4COCKPIT transaction if the Java database is on the same host (in the same LPAR) as the ABAP database. Support for remote Java databases is planned for SAP WebAS releases after 6.40.

You can use the SAP WebAS Java SQL trace to trace database access in Java, similar to the SQL trace in ABAP. Go to the following entry page:

`http://`*`servername: instance no.`*`00/SQLTrace`

Figure 26-29 illustrates this.

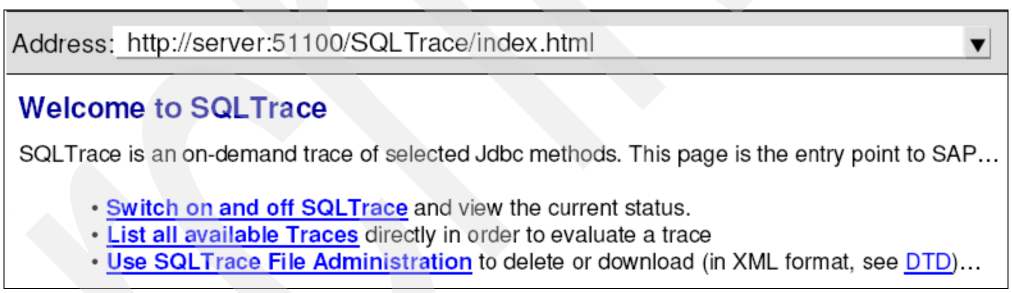

Figure 26-29 Java SQL trace

The SQL trace is written to a set of files that are stored in Extensible Markup Language (XML) format. You cannot limit the SQL trace to a specific user when collecting the data, other than in ABAP. You can limit the data by certain method names or execution times when backtrace information is added (that is the complete method call stack when calling Java Database Connectivity (JDBC) methods).

You can set filters (for user names, application, system or non-system thread) when you evaluate the trace data so that you have a better overview of the data. Delete SQL trace files that are no longer needed. The File Administration function helps you with that task. In addition, it allows you to download the files in XML format.

Figure 26-30 shows an example for the output from the Java SQL trace.

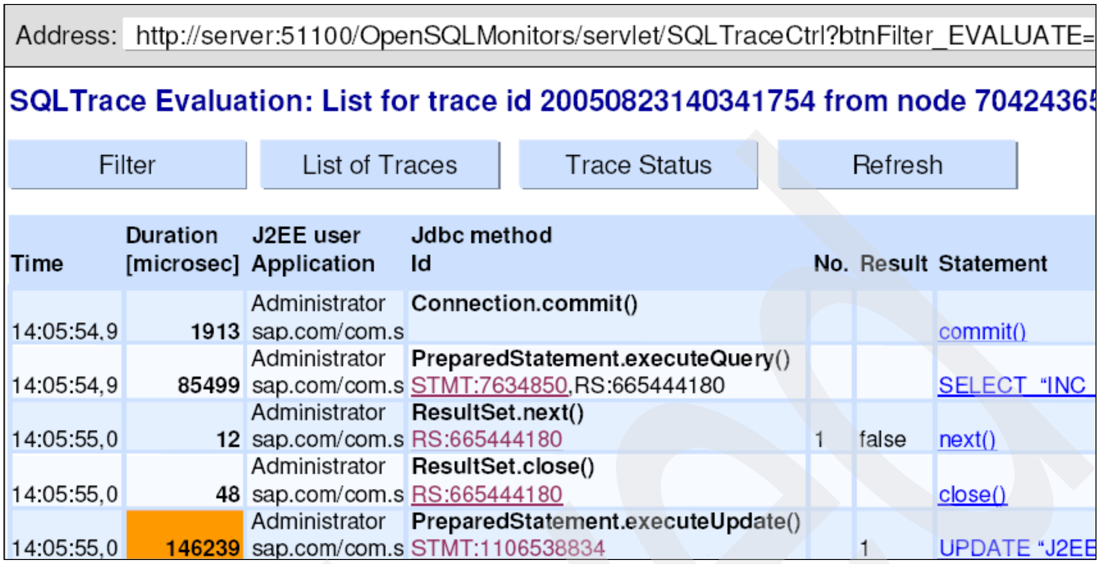
Address: http://server:51100/OpenSQLMonitors/servlet/SQLTraceCtrl?btnFilter_EVALUATE=

SQLTrace Evaluation: List for trace id 20050823140341754 from node 70424365

Filter | List of Traces | Trace Status | Refresh

Time	Duration [microsec]	J2EE user Application	Jdbc method Id	No.	Result	Statement
14:05:54,9	1913	Administrator sap.com/com.s	Connection.commit()			commit()
14:05:54,9	85499	Administrator sap.com/com.s	PreparedStatement.executeQuery() STMT:7634850,RS:665444180			SELECT "INC
14:05:55,0	12	Administrator sap.com/com.s	ResultSet.next() RS:665444180	1	false	next()
14:05:55,0	48	Administrator sap.com/com.s	ResultSet.close() RS:665444180			close()
14:05:55,0	146239	Administrator sap.com/com.s	PreparedStatement.executeUpdate() STMT:1106538834		1	UPDATE "J2EE

Figure 26-30 Java SQL trace: output

Long execution times are identified in color. You can see the JDBC functions with their related SQL statements as they are executed by the applications. Click on a statement text to get the statement details, including the stack of methods, including the JDBC method if you selected it, to collect backtrace information. Currently, you do not get any implementation background (the Explain) for Java.

To run Java SQL traces you must be a member of the Administrator group.

26.3 Java optimization

Performance of the application program itself is often a small contributor to the overall performance in traditional i5/OS operations. A large percentage of a program’s execution can be contributed to the system services used by the application (for example, database get record operations). There are two ways to improve application performance:

- Improving i5/OS
- Improving how the application uses the system services (especially DB2 UDB for iSeries services)

For Java, this can still be true. Key portions of Java (such as JDBC, encryption, security) can have a substantial portion of their support executing in i5/OS. For some applications, tuning Java’s use of these system services is sufficient performance tuning. However, for many complex applications, it is also true that Java, as part of its portability story, often has a higher percentage of the application's execution in Java programs and uses less of a given operating service's function. It is the performance of these Java middleware functions that are becoming important.

Java is maturing as a language. It is interesting to compare Java to traditional languages. Such comparisons are difficult, because Java computation occasionally reaches parity for some applications and is seldom twice as slow as traditional languages. However, as early as OS/400 V4R5, performance is not a significant barrier to Java deployment. Improvements to the just-in-time compilation (JIT) beginning with OS/400 V5R2 have improved performance by 15% or more.

Such comparisons are less relevant, however, because Java has made substantial progress over the last several releases. Java is important in its own right. Java becomes more important for applications such as mySAP Business Suite that simply require Java.

In general, Java uses fewer CPU cycles for the same commercial processing workload (CPW) rating. The default initial heap size of a Java virtual machine (JVM) has been changed from 2 MB to 16 MB in i5/OS V5R3. This provides significantly better performance for most JVMs that use the default initial heap size for everything except the smallest and shortest-running JVMs.

Specify the following statement on the Java command line to start a JVM that overrides the default heap size:

```
-XmsSize of JVM
```

This change is necessary as the JVMs become longer-lived, more complex, and with more objects. The larger initial heap size allows more objects to be created before the garbage collector (GC) runs. This reduces the number of GC cycles that run, which in turn reduces the amount of CPU time spent in the GC.

> **Note:** The default Java Development Kit (JDK) used with i5/OS V5R3 is 1.4. The default in OS/400 V5R2 is 1.3 JDK. Testing shows equivalent performance between the 1.4 and 1.3 versions of the JDK. Any performance updates are targeted at JDK 1.4, so it is beneficial to keep the operating system and database current in program temporary fix (PTF) level in order to reach optimum Java performance.

Refer to the following authorized program analysis reports (APARs) at the following Web site for the latest PTF packages:

- *SF99530 for i5/OS V5R3* and *SF99520 for OS/400 V5R2* for Cum PTF Package
- *SF99503 for i5/OS V5R3* and *SF99502 for OS/400 V5R2* for DB2 UDB Group PTF Package
- *SF99269 for i5/OS V5R3* and *SF99169 for OS/400 V5R2* for Java Group PTF Package
- *II13868 for i5/OS V5R3* and *II13337 for OS/400 V5R2* for SAP on System i5 Info APAR

You can find these APARs at the following Web site:

http://www-912.ibm.com/ImprovedSearch/searchoptions.jsp

Required memory and virtual machine (VM) settings are different from those set by the SAP J2EE engine on System i installation procedures. Adopt these settings to improve the runtime behavior.

Refer to SAP note 717376: Recommended VM settings for the iSeries (OS/400) VM for the latest information by selecting **SAP notes Search** at the following Web site:

http://www.service.sap.com

Refer to Chapter 28, “SAP Business Information Warehouse” to understand the factors affecting database performance.

27

Performance management tools and services

This chapter discusses helpful tools from SAP and IBM to monitor your mySAP Business Suite applications. Use these tools to obtain high-level overview reports to analyze the performance of your system.

Refer to Chapter 25, “Performance concepts and monitoring” to understand the concepts of monitoring your System i server and SAP system, and 26.1, “Tuning System i models” on page 532 for information about analyzing performance.

27.1 SAP EarlyWatch

SAP offers a special service with *EarlyWatch*. The SAP EarlyWatch Check analyzes the components of your SAP application, the operating system, and the database to determine how to optimize performance. Initiate the check when you encounter performance issues. The check can also be triggered by the SAP EarlyWatch Alert, which automatically monitors your SAP system.

If vital alerts are reported through SAP EarlyWatch Alert, SAP schedules an SAP EarlyWatch Check for a more detailed analysis of your system. Each productive system is entitled to a maximum of two SAP EarlyWatch Checks per year within your maintenance agreement with SAP.

Additional SAP EarlyWatch Checks are available for an additional fee. This offer is valid for SAP direct customers. Special conditions apply for customers of value-added remarketers (VARs).

A delivery date is set up when the SAP EarlyWatch Check is triggered manually or by the SAP EarlyWatch Alert. The delivery of the SAP EarlyWatch Check lasts one day and is performed by experienced service engineers over a remote connection. Programs run in the background to collect data two weeks prior to delivery of the service. This data is the basis for the analysis performed by the EarlyWatch Check. The collection of data and the delivery of this service do not affect production. Service engineers telephone the designated contact person at your company when the analysis is complete.

Analysis is the key of an EarlyWatch Check. The results provide you with information necessary to:

- Eliminate performance bottlenecks
- Make a more thorough and detailed performance analysis
- Improve response times
- Improve the acceptance of the system by end users
- Optimize use of your hardware investment

The concept of the SAP EarlyWatch Check is to ensure smooth operation of the SAP system by taking proactive action before severe technical problems occur. This EarlyWatch Check is based on numerous installations, giving you experience that saves time, reduces costs, and keeps your SAP system running smoothly.

The results of the analysis and the actions necessary to optimize your system's performance are listed in a final report consisting of the following topics:

- Global EarlyWatch summary
- Global analysis
- Application server analysis
- Database server analysis
- Summary and recommendations
- Appendix A: Checklists
- Appendix B: Glossary

Traffic light logic helps identify critical performance points of your SAP system.

Figure 27-1 shows the system status from the global summary report.

1 System Status

During this EarlyWatch Alert Session, we detected potential problems concerning your system.
We recommend that you take corrective action as soon as possible.
If you would like further information, create a customer message on component XX-SER-TCC, or call your SAP Local Support organization.

Note:
All recommendations provided in this report are based on our general experience only. We advise you to test our recommendations before using them in your production system. Also note that EarlyWatch Alert is an automatic service.

Service Contents

Overview: Alert Messages
The following table contains an overview of all alerts.

Rating	Performed Check
⚠	System Configuration
✔	System Performance
✔	Workload Distribution
✔	SAP System Operating
✔	Hardware Capacity
✔	DB2/400 Database Checks

Priority	Description	New Alert
High	At least one ABAP dump of type 'SAPSQL_ARRAY_INSERT_DUPREC' was found.	New

Figure 27-1 SAP EarlyWatch Alert: system status

Figure 27-2 shows the trend information about the top five transactions.

TOP 5 TRANSACTIONS

Date	Transaction 1	Transaction 2	Transaction 3	Transaction 4	Transaction 5
05.09.2004	1391	5440	2162	875	4818
12.09.2004	1424	4686	2962	884	4854
19.09.2004	1541	5727	3042	830	4317
26.09.2004	1610	3589	2899	836	4309
03.10.2004	598	3535	2876	801	4135
10.10.2004	937	1072	4885	859	4365
17.10.2004	980	2065	4432	827	4686
24.10.2004	997	4602	4408	829	4605
31.10.2004	871	2289	4679	800	4492
07.11.2004	786	2278	4757	833	4781
14.11.2004	353	6229	4541	858	4836
21.11.2004	1096	4442	4928	875	4840

No	Transaction Code	Load (%)
Transaction 1	SESSION_MANAGER	4.6
Transaction 2	KE30	4.4
Transaction 3	ZKV02	4.0
Transaction 4	VA01	3.8
Transaction 5	COHVOMPP	3.2

Figure 27-2 SAP EarlyWatch: trend of response time for the top five transactions

Information about Advanced Business Application Programming (ABAP) dumps is displayed in a report as shown in Figure 27-3.

6.3 Program Errors (ABAP Dumps)

In your system, 45 ABAP dumps have been detected since 18.09.2005. To view the ABAP dumps in your system, call transaction ST22 and choose 'Selection'. Then select a time frame.

ABAP DUMPS PER DAY

Date	Quantity of Dumps
15.09.2005	0
16.09.2005	0
17.09.2005	0
18.09.2005	12
19.09.2005	4
20.09.2005	3
21.09.2005	0
22.09.2005	9
23.09.2005	9
24.09.2005	8

DETECTED ABAP DUMPS

Error	Number	Server (e.g.)	User (e.g.)	Date (e.g.)	Time (e.g.)
DBIF_RSQL_SQL_ERROR	4	SLPAG06	ULLER	20050918	095613
DBIF_RSQL_INVALID_RSQL	1	SLPAG06	NEIDER01	20050920	071049
SYSTEM_CANCELED	3	SLPAG06	WARZULIG	20050922	082010
TABLE_INVALID_INDEX	4	SLPAG06	WALTER	20050922	122757
GETWA_NOT_ASSIGNED	3	SLPAG06	NEIDER01	20050923	093907
GEN_SAPSQL_TOO_MANY_ITABS	1	SLPAG06	BRECHT	20050923	131703
CONTROL_NOT_FOUND	1	SLPAG06	STOPPE	20050923	141326
SAPSQL_ARRAY_INSERT_DUPREC	2	SLPAG06	TENES	20050923	141632
MESSAGE_TYPE_X	4	SLPAG06	BAKER	20050924	133936
TIME_OUT	7	SLPAG06	NAHM	20050924	162501
RAISE_EXCEPTION	15	SLPAG06	D440FG	20050924	212828

It is important for you to monitor ABAP dumps using transaction ST22 on a regular basis. If ABAP dumps occur, you should determine the cause as soon as possible.
Based on our analysis, we expect no serious problems at the moment.

Figure 27-3 SAP EarlyWatch: history of ABAP dumps

You can find more information about SAP EarlyWatch and other SAP services on the Web at:

`http://service.sap.com/servicecatalog`

27.1.1 IBM Insight

The IBM Insight for SAP utility programs and its subsequent analysis process and report are designed to provide a high-level and convenient workload analysis for an in-production SAP system, for all servers. The analysis includes a count of active users, machine utilizations, user and module load distributions, a dialog count, information about batch and reporting usage, system and database information.

The IBM Insight for SAP utility programs is packaged as an all-in-one Windows 2000 Service Pack 3 (SP3) and Windows XP install image, ready to install on any customer's Windows or Intel system that is capable of communicating with the production SAP complex. Documentation is included with the software. IBM Insight for SAP has been validated to run against SAP production complexes comprised of Wintel (xSeries®), AIX (p5 and pSeries®), i5/OS or OS/400 (System i), OS/390 (zSeries®), and non-IBM installed UNIX or Wintel servers.

To begin the recording session to use IBM Insight, install the utility on your personal computer (PC), set up initial communication parameters, and ensure authorization and access. The software records performance data from the SAP complex using SAP's Remote Function Call (RFC) functionality until the session is ended.

The collected data is forwarded to IBM via e-mail for data reduction, analysis, and to produce customer reports. Only the data from the production system is analyzed. Non-production environments are excluded. Collect three days of data during a period of the month with reasonably high usage.

Use the following link to download the IBM Insight software:

ftp://ftp.software.ibm.com/software/sap/SetupInsight4.exe

IBM Insight Version 4 (available in November 2005) provides for collecting data from Unicode and non-Unicode based SAP systems.

The prerequisites for IBM Insight are:

- A PC dedicated for data collection use
  - An IBM-compatible local area network-attached (LAN-attached) PC with at least a 500 MHz Pentium® processor and 64 MB of RAM
  - Windows 2000 SP3 or Windows XP operating system
  - A minimum of several hundred MB of disk space.

    Multiply the total expected number of dialog steps the Insight Collector is to capture by 0.00025. For example, if the Insight Collector is expected to run for three days capturing 500,000 dialog steps per day, the formula is:

    3 x 500,000 x 0.00025 = 375 MB of disk space
  - Direct Transmission Control Protocol/Internet Protocol (TCP/IP) access to all production servers (except for a stand-alone database server)
  - TCP/IP addresses of all server short host names are resolved with either the hosts file entries, or the Domain Name System (DNS) with a consistent, fully-qualified domain. Do not specify a domain in the SAP logon parameters.
  - The PC must be powered on during the entire data collection process. Disable (suspend) the PC's power management during the collection process.
- A production SAP system
  - SAP Basis release 3.0D and higher
  - SAP Operating System collector (saposcol) setup to run on all application servers in the target system
  - SAP recording statistics enabled (profile parameter: stat/level = 1)
  - An SAP user ID is created with a type of CPI-C (System). A profile (that is, ZINSIGHT) is defined with the following authorizations:
    - S_ADMI_FCD(S_SACHBEARB)
    - S_C_FUNCT(S_C_FUNCT_AL)
    - S_DATASET(S_DATASET_AL)
    - S_LOG_COM(S_LOGCOM_ALL)
    - S_RFC(S_RFC_ALL)
    - S_RZL_ADM(S_RZL_SHOW)

Most of Insight's information is related to the SAP system where Insight is run. However, Insight's host CPU utilization has a wider perspective. The RFC mechanism provided by SAP collects only CPU statistics for an entire machine. SAP systems do not make a delineation for the potential source of the load. This means, for example, that if two SAP systems run on a single machine, both systems contribute to the same total reported CPU utilization. It is not possible for IBM to allocate specific resource consumption amounts against a specific SAP system. The same is true for significant loads coming from a non-SAP related or reported source.

Similarly, the Insight *Host CPU utilization* is for all SAP systems where two or more SAP systems are run or additional "other" workload is in the same logical partition (LPAR). However, the Host CPU Utilization does not include any work running in other partitions.

> **Important:** Host CPU utilization measures the total workload, including the "other" workload of this system if you run Insight on a system with significant "other" workload.

You can find additional information about IBM Insight on the Web at:

http://www.ibm.com/erp/sap/insight

Also refer to the README file provided with the Insight tool.

For instructions on installation, data collecting, and data reduction, refer to:

ftp://ftp.software.ibm.com/software/mktsupport/techdocs/insight4_readme.pdf

Also refer to the document with common questions and answers for Insight at:

ftp://ftp.software.ibm.com/software/mktsupport/techdocs/insight3_q_and_a.pdf

You are prompted for SAP system hardware information (machine model, number of CPUs, and memory) during the setup of Insight. Enter this information correctly. Incomplete or incorrect information can cause incorrect reporting of the system.

IBM produces one SAP system ID (SID) report per quarter. You can request additional reports based on your needs and the IBM workload. Send a request for your additional requirements to the following e-mail address before sending in new data:

eSizings@us.ibm.com

After the data collection and reduction, you can send an e-mail directly to IBM via Simple Mail Transfer Protocol (SMTP) with the Insight data attached for analysis. Send an e-mail to the following e-mail address for any technical questions or problems with IBM Insight:

eSizings@us.ibm.com

Mention "Insight Technical Question" in the subject line. Enter your company contact's e-mail address. Select the appropriate IBM destination. Press the **OK** button to send the e-mail. You do not require e-mail software on your PC to send the data. Refer to Figure 27-4 on page 571 and Figure 27-5 on page 571.

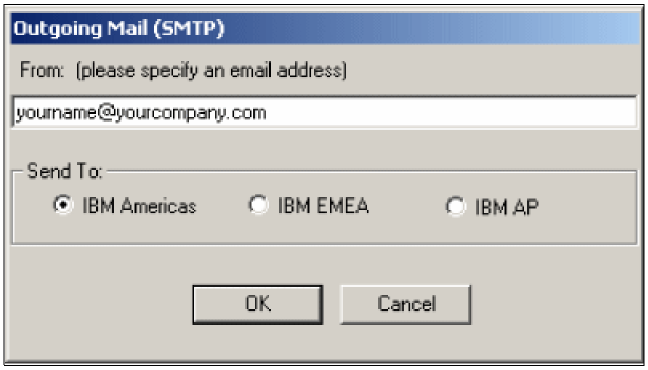

Figure 27-4 IBM Insight: send data mail dialog window

The locations are defined as follows:

- IBM Americas refers to North, Middle and South America.
- IBM EMEA refers to Europe, the Middle East, and Africa.
- IBM AP refers to Asia and Pacific.

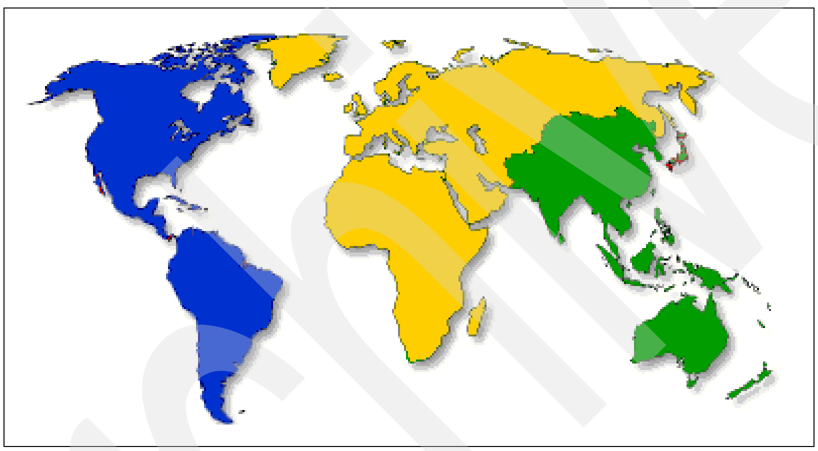

Figure 27-5 IBM Insight: select geography for the send data e-mail

When IBM receives the data, it is checked for data integrity. In general, about five working days are required. Insight is not a reactive process. Plan all activities well in advance.

The prepared Insight report includes the following sections:

- General
- Abbreviations
- System i information
- Component information
- Patch information
- Active users
- Dialog steps
- Dialog response time
- Dialog steps by task type
- Dialog steps by module
- Dialog steps by code type
- CPU time by task type
- CPU time by module
- CPU time by code type
- Database wait time

- Top transactions
- Top users
- Instance list
- Instance information
- Host list
- Host information
- Host CPU utilization summary
- Host memory utilization summary
- Host disk utilization summary
- Host statistics
- Top disk

Note: The memory statistics portion of the IBM Insight report is not applicable to a System i report due to the single-level storage architecture of System i models.

Figure 27-6 through Figure 27-10 on page 575 show examples from an Insight report for the following:

- Active users

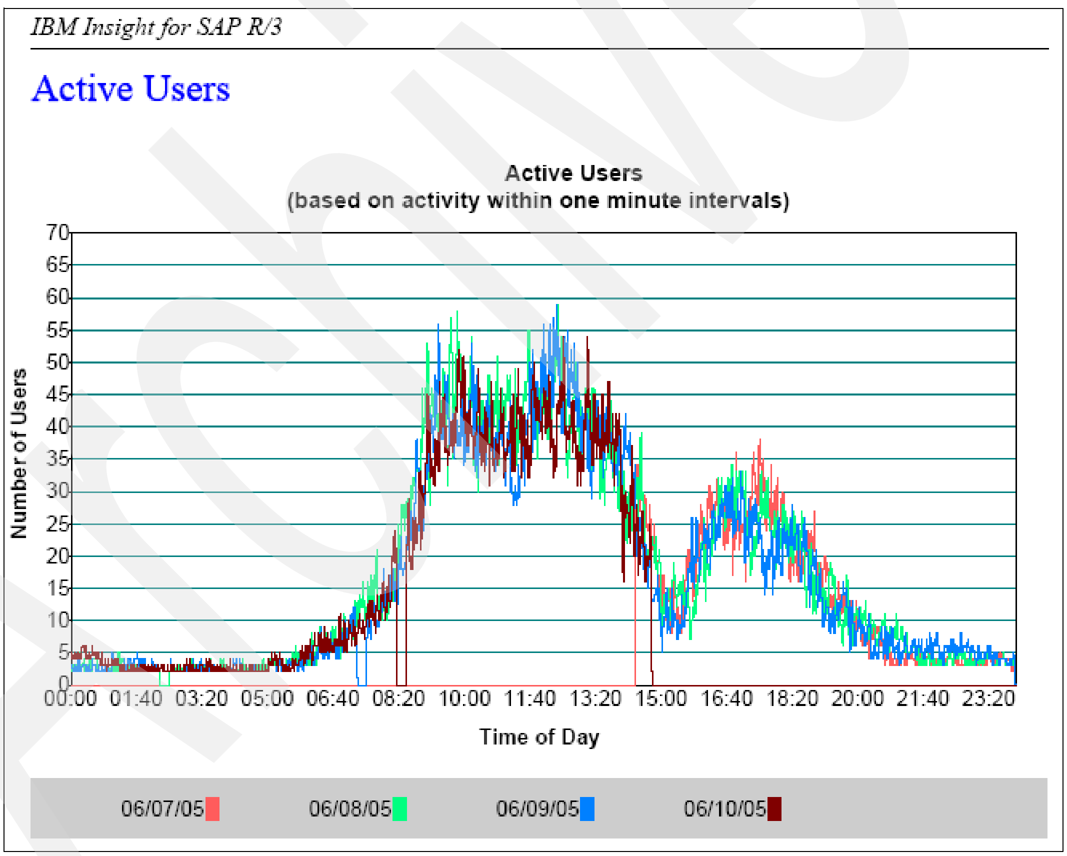

Figure 27-6 IBM Insight: active users

► Dialog steps

IBM Insight for SAP R/3

Dialog Steps

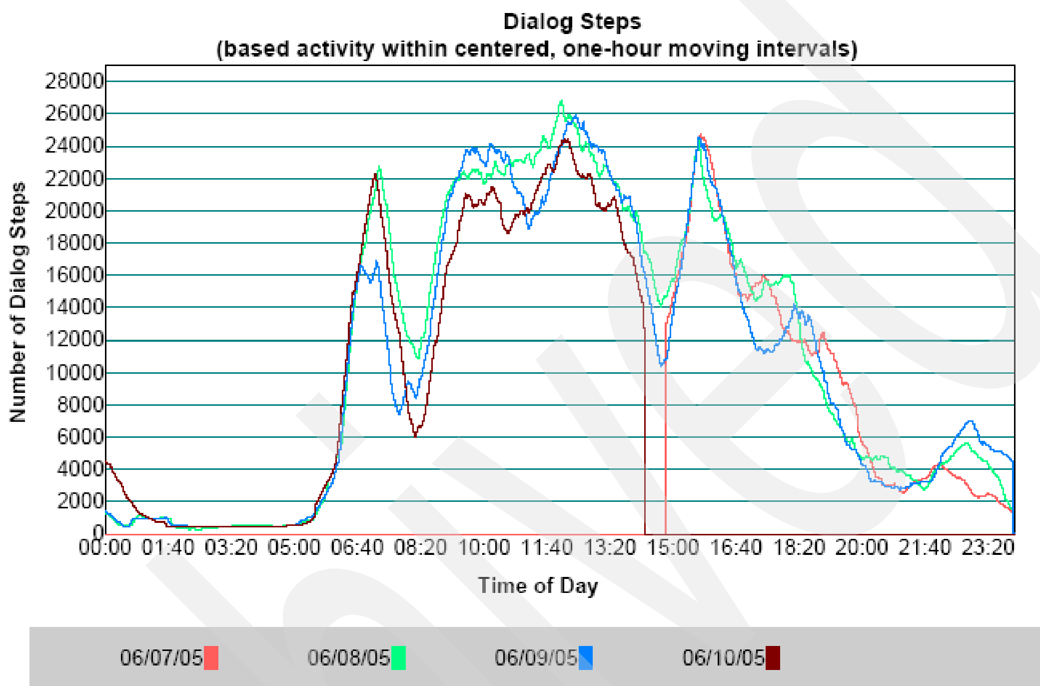

Statistics

26799	Highest number of dialog steps per hour at 12:03 on June 8, 2005.
10672	Average number of dialog steps per hour within the monitoring period.

Figure 27-7 IBM Insight: dialog steps

- Comparison of dialog steps versus CPU time by code type

  In the example, the average CPU consumption of a customer programming transaction is nearly twice that of an SAP standard transaction (13.1 / 7.2 = 1.8). Refer to Figure 27-8.

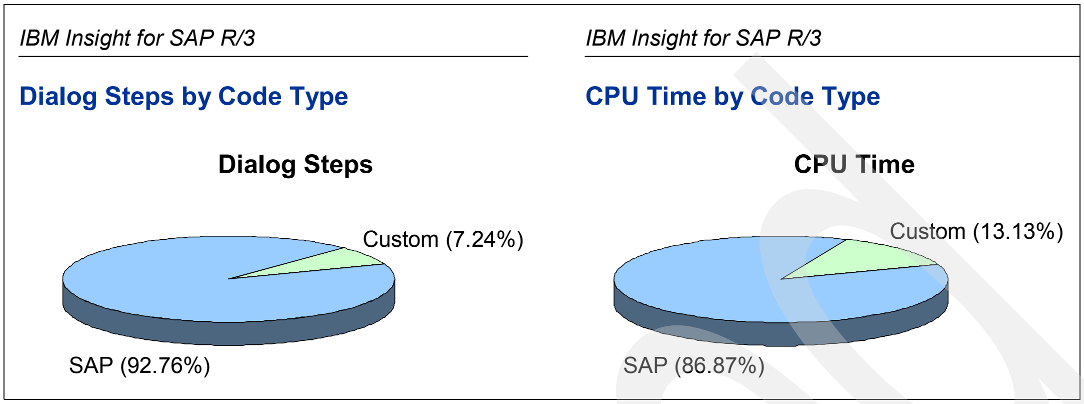

Figure 27-8 IBM Insight - Dialog steps versus CPU time by code type

- Comparison of dialog steps versus CPU time by module as shown in Figure 27-9.

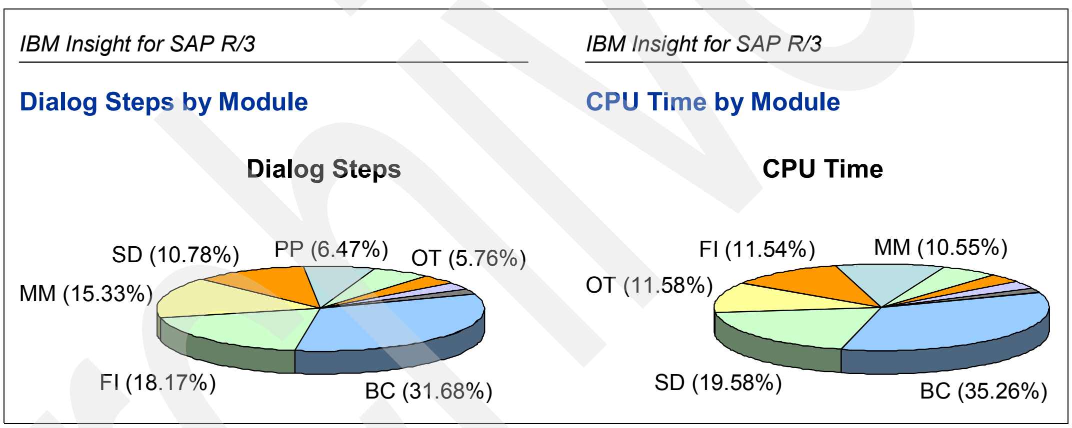

Figure 27-9 IBM Insight - Dialog Steps versus CPU Time by Module

- A top transaction by memory consumption report as Figure 27-10 shows.

IBM Insight for SAP R/3

Top Transactions by Memory Consumption

Memory (MB)	%	Steps	Avg (MB)	Module	TCode/Report
539603	10.1	2472	218.3	BC	SP01
536060	10.0	33640	15.9	FI	FBL5N
354505	6.6	20954	16.9	FI	FBL3N
315722	5.9	37661	8.4	SD	VA01
265256	5.0	16522	16.1	FI	FBL1N
124872	2.3	12045	10.4	SD	VA02
109690	2.1	18304	6.0	MM	ME57
101339	1.9	304	333.4	BC	SM30
98626	1.8	3923	25.1	CO	KKS5
97484	1.8	3370	28.9	CO	KSB1
96444	1.8	74022	1.3	BC	RSM13000

Figure 27-10 IBM Insight: top transaction by memory consumption

Important: Retain workload, and performance-related data and reports to better understand the trends in your SAP system landscape.

27.2 Performance aspects in 3-tier environments

The 3-tier systems experience higher database time than 2-tier systems. The network overhead between application servers and the database server is included in the database time shown in SAP system reports. A 3-tier system, nevertheless, reduces the utilization of the database server considerably. Therefore, more users can work at the same time when the application workload is removed from this server. A 2-tier installation is typically the fastest solution, unless the originating network overhead is more than compensated by shorter CPU times on a faster application server.

In 3-tier installations, application servers can communicate with the System i database server through the two following methods:

- OptiConnect
- TCP/IP

With OptiConnect, specified as `dbs/db4/opticonnect=1`, you require Option 23 of i5/OS (a chargeable feature) to communicate via high-speed link (HSL) OptiConnect, Virtual OptiConnect (between partitions of one server), or system product division (SPD) OptiConnect (a physical connection of older hardware). To communicate over TCP/IP, specified as `dbs/db4/opticonnect=0`, use Gigabit Ethernet adapters or Virtual Ethernet between partitions in one server.

Adapters between database and application servers that are slower than Gigabit Ethernet can work, but are not recommended and not supported because of performance reasons.

In 3-tier installations, each work process has at least one shadow process on the database siteFigure 27-11 shows.

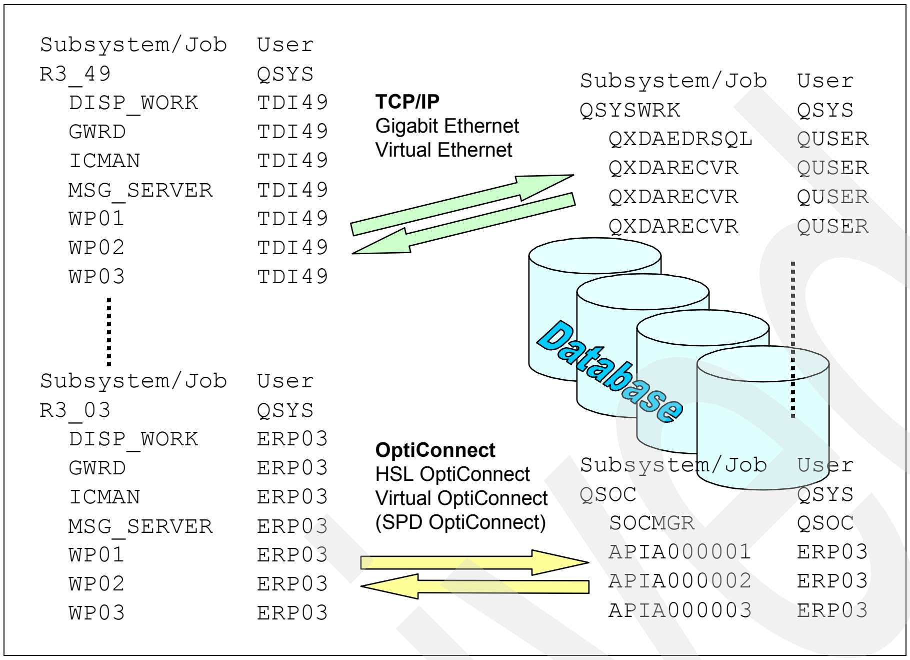

Figure 27-11 Shadow processes in a 3-tier installation

For example, in Business Information Warehouse, a work process can open more than one database connection. Controlling jobs like dispatcher and message server do not have a database connection.

In installations with TCP/IP connection, the shadow process jobs named QXDARECVR, run in the subsystem QSYSWRK with the QUSER user profile. If TCP/IP is used on a 3-tier system, install and permanently activate the **CHKXDA**, as described in SAP note 450351.

In installations with OptiConnect, the shadow process jobs named APIA*nnnnnn* run in the QSOC subsystem with the SAP instance user profile, for example ERP03.

Due to the higher database response times, the SAP system buffers are more important in a 3-tier than in a 2-tier system. For this reason, avoid daily restarts of the SAP system on 3-tier systems.

It is better to buffer fewer tables to improve performance, because of the need to synchronize the table buffers across all SAP instances of an SAP application.

Note: There can be additional administrative tasks on 3-tier systems. Additional skills are required to run these operating system tasks.

Another performance aspect in a 3-tier environment is related to the running of Batch Input jobs. You should try to run them on the same system that the sequential input file is located. This can make a significant difference in the runtimes, if the input Integrated File System file and the log for the job are physically located on the local server that the job is running on. This minimizes any communication delay in frequently accessing Integrated File System files. Usually the file system is located on the same server as the central instance.

For more information about System i 3-tier installations, refer to:

- *SAP note 484548: iSeries: 3 tier system installation*
- *SAP note 450351: iSeries: Perf-problems on 3-tier or Virtual Host systems*

27.3 Performance aspects in complex system landscapes

This section covers some aspects in complex system landscapes. Because of the efficient work management operations of i5/OS on System i models, several SAP instances or SAP systems run under the same operating system image.

27.3.1 SAP OS Collector

The SAP OS Collector is an i5/OS job named *SAPOSCOL*.

SAPOSCOL starts the first time the system is started. The collection starts once per LPAR, even if more than one SAP system is run on the system or in a partition. Subsequent executions of STARTSAP in other systems recognize that the program is already active and stop their SAPOSCOL within several minutes.

Special considerations apply if more systems run in one partition mixed with SAP kernels prior to and with Version 6.40.

SAPOSCOL uses the System i Single Level Store (SLS) memory model based on 16 MB memory segments for SAP Release 3.1I, 4.0B and 4.5B implementations to maintain compatibility with earlier i5/OS operating systems.

The memory model of the System i server has been enhanced due to various technological developments. SAPOSCOL now uses the teraspace model for SAPOSCOL of the Release 6.40 level SAP kernel. Teraspace storage overcomes the 16 MB limitation of segmented storage and provides up to one terabyte (TB), that is 1024 GB, of private storage to a single process.

SAP applications are designed in releases prior to the kernel 6.40 such that performance data that is recorded centrally by the SAPOSCOL job can be read and processed on all systems, regardless of the SAP release level, and in spite of the different memory and pointer models. An SLS memory segment also collects performance data on teraspace-capable systems.

As of SAP kernel release 6.40, SAP systems now use a memory model that can no longer process SLS segments. Therefore, data on systems with earlier releases collected in an SLS segment cannot be viewed on 6.40-based SAP systems and vice-versa.

The implementation is changed such that the performance data collection for all SAP kernel releases up to and including SAP kernel release 6.20 (for example, Releases 31I, 40B, 45B, 46D, and 6.20) still collect performance data in an SLS segment. Also, kernels that are based on SAP kernel release 6.40 (specifically Releases 6.40 and 6.20 with the 6.40 downward-compatible kernel) collect performance data in a teraspace segment to avoid these problems. Therefore, two SAPOSCOL jobs can be active on a System i server collecting performance data independently of each other when there is a combination of 6.40- and other release level kernels.

Some reported SAPOSCOL problems include the following:

- In Transaction ST06, under the Operating System Collector, the status message `SAPOSCOL not running` appears even though WRKACTJOB tells you that an SAPOSCOL job is active.
- The status message `Work processes terminate because they cannot connect to a shared memory segment` is displayed.
- Performance data is not displayed in Transaction 06.

Use the latest combination of kernel DW (DWDB4) and SAPOSCOL programs according to *SAP note 708136 - iSeries: No monitor data in ST06* to help resolve these problems.

Refer to the following SAPOSCOL notes for more information:

- *SAP note 753917 - Series: Version number of SAPOSCOL*
- *SAP note 637174 - SAPOSCOL cannot access Libraries of different SAP systems*.

27.3.2 Multiple system considerations

Because of the work management of System i servers, you can run several SAP systems under the same operating system image (in a partition or an entire System i server without LPAR). Ensure that you have sufficient hardware resources in the system.

System pools

By default on System i configurations, all SAP systems run in System Pool 2 (*Base), which has a defined size. Run the Work with Active Job (**WRKACTJOB**) command to see this as shown in Figure 27-12. Refer to Figure 27-13 on page 579 for the Work with System Status (**WRKSYSSTS**) command.

Jobs on an SAP system can compete for system resources (memory, CPU, and activity levels) with jobs of another SAP system or other applications. Refer to Figure 27-12 and Figure 27-13 on page 579 for an illustration.

```
                         Work with Active Jobs                    ITSO-SYS1
                                                         11/25/05  13:36:21
CPU %:    21.6     Elapsed time:   00:07:17     Active jobs:   521

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 ...
                                              --------Elapsed---------
Opt  Subsystem/Job  Type  Pool  Pty      CPU  Int    Rsp  AuxIO  CPU %
       WP04         BCI     2   12        .6                  0     .0
       WP05         BCI     2   20     222.6                 45     .0
       WP06         BCI     2   20       7.9                  0     .0
       WP07         BCI     2   20        .7                  0     .0
       WP08         BCI     2   20       4.6                  1     .0
       WP09         BCI     2   20        .6                  0     .0
     R3_03          SBS     2    0        .8                  0     .0
       DISP_WORK    BCI     2   12      16.4                  8     .0
       GWRD         BCI     2   12      25.6                  2     .0
```

Figure 27-12 WRKACTJOB

```
                              Work with System Status                    ITSO-SYS1
                                                                 11/25/05  13:39:27
% CPU used . . . . . . . :       27.7    Auxiliary storage:
% DB capability  . . . . :        1.6      System ASP . . . . . . :     987.9 G
Elapsed time . . . . . . :   00:00:38      % system ASP used  . . :     87.3989
Jobs in system . . . . . :       2331      Total  . . . . . . . . :     987.9 G
% perm addresses . . . . :       .013      Current unprotect used :     45176 M
% temp addresses . . . . :       .117      Maximum unprotect  . . :     47200 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      800.00    187.87   +++++     .0      .0     .7    1.0
   2     7391.06      3.91     162   44.0 112.3    17.3  113.9
```

Figure 27-13 WRKSYSSTS

As each SAP instance presents its own i5/OS subsystem, you can run these subsystems isolated from each other. Separate the subsystems in memory pools to help minimize a contention for memory. Use the change Subsystem Description (**CHGSBSD**) command to route the jobs of this SAP instance into another memory pool after stopping an SAP instance as Figure 27-14 shows.

```
                       Change Subsystem Description (CHGSBSD)

Type choices, press Enter.

Subsystem description  . . . . . > R3_03         Name
  Library  . . . . . . . . . . . >   R3ERP400    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
```

Figure 27-14 CHGSBSD

Run the Work with Shared Pool (**WRKSHRPOOL**) command to set or change the size of the shared pool.

> **Important:** Work without automatic performance adjustment. Set the i5/OS system value of QPFRADJ to **0**.
>
> Keep in mind that with running in different system pools, there is no complete separation of the applications. The applications should always share communication and I/O resources, disks and temporary work space (TEMP). Use partitions to isolate an SAP instance completely.

OLAP versus OLTP

From the technical point of view, running different processing types, Online Transaction Processing (OLTP) and Online Analytical Processing (OLAP), in one partition is not a problem. However, OLTP and OLAP run quite differently.

OLTP systems, like SAP ECC, are business process oriented, often running recurrent transactions with current data. Database access methods are read, write, update, or delete. Typical transactions are short read and write operations. Users expect split-second response times.

OLAP systems, like SAP Business Intelligence, use historical, multi-dimensional data. Few users perform a few transactions, need large amounts of data, access ad hoc and heuristic data, and read condensed data. With this manner of work, the users experience long response times of up to several minutes. OLAP transactions are more database-centric (DB-centric) as compared to OLTP transactions.

Because of these differences, OLTP and OLAP systems do not work well together. We recommend that you separate OLTP and OLAP systems into different partitions or separate servers.

CPU consumption

In many cases, for monitoring, sizing, planning tasks, and so on, you need information about CPU resource utilization. If you run multiple SAP instances in one LPAR or one server, it can be difficult to provide all the information you require.

On an operating system level, the CPU consumption is shown when you run the **WRKSYSSTS** command for the partition and respective system. **WRKACTJOB** lists values accumulated for single jobs.

The OS Collector which provides OS-related data into the SAP database, cannot distinguish subsystems nor determine where the workload comes from. On an SAP level, the transaction ST06 is based on the OS Collector. Both, the CPU values in the *snapshot analysis* view, and the values in the *previous hours* view, represent the CPU consumption of the whole LPAR server.

If you run multiple SAP instances mixed with other applications, you can see the total CPU consumption for all applications. The same is true for the other components measured with the OS Collector using transaction ST06 (disk, pool (faults), and so on).

Use the **Pool Report** from the Performance Tools licensed program product (5722-PT1) to see the CPU utilization of different SAP instances. Refer to 27.4.2, “Performance Tools” on page 585 for more information.

The Pool Report, as shown in Figure 27-15 on page 581, shows the CPU consumption of every subsystem over each interval of i5/OS of data collection.

```
                              Pool Interval Report            12/18/05 14:58:5
                                  Subsystem Activity
Page
Pool Report Selected Pools / 12 a.m. - 02 p.m.
Member  . . . : Q318091602 Model/Serial  . : 550/65-9FEBD
Main storage  . . : 8256.0 MB Started . . . : 12/14/05 09:16:03
Library . . : QMPGDATA   System name . . : ITSO-SYS1 Version/Release . :     5/3.0
Stopped . . . : 12/15/05 00:00:00
  Partition ID  : 001       Feature Code  . : 7463-8958
                                     --------- Physical I/O per Transaction -------
   Itv  Subsystem     CPU             ----- Synchronous ----- ---- Asynchronous -----
   End  Name       PL Util     Tns     DBR   DBW   NDBR  NDBW  DBR   DBW   NDBR
  ----- ---------- -- ----- -------- ----- ----- ----- ----- ----- ----- ----- -----
  12:00 QINTER      2    .0       31             11.0    .2
  12:00 QSERVER     2    .0        0
  12:00 QSYSWRK     2    .0        0
  12:00 R3_03       2  14.5        0
  12:00 R3_10       2    .0        0
  12:00 R3_11       2   2.6        0
  12:00 R3_25       2    .2        0
  12:00 R3_27       2   1.3        0
  12:00 R3_30       2    .2        0
  12:00 R3_42       2    .1        0
```

Figure 27-15   Pool Report

Use this data, or consolidate additional collections over more intervals by using a spreadsheet.

You also can use the **i5/OS Collection Services object** (***MGTCOL**). The Create Performance Data (**CRTPFRDTA**) command processes data from this collection object and stores the result in performance database files. The file QAPMSYSCPU contains System CPU usage data. You can use it and run your own evaluation query.

For more details, select **Systems Management** → **Performance** → **Manage iSeries Performance** in Online iSeries Information Center, Version 5 Release 3 on the Web at:

http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp?lang=en

Note: SAP transaction ST06 shows the CPU sum for all applications if you run multiple SAP instances or mix other applications. The same is true for all reports or tools which are based on the OS Collector performance data collected within the SAP application, for example, the SAP EarlyWatch Services and IBM Insight.

27.3.3 Testing systems with lower priority

Different priorities can be used when running multiple SAP systems on one system or in any partition. For example, an SAP system can run with a lower priority an SAP production system.

Use the Change Class (**CHGCLS**) command to edit the run priority parameter (RUNPTY) in the class (*CLS) object type to change the priority. The Class Object is found in the SAP System Library R3*sid*400. Figure 27-16 on page 582 illustrates this.

```
                              Change Class (CHGCLS)

 Type choices, press Enter.

 Class  . . . . . . . . . . . . . . > R3_04          Name
   Library  . . . . . . . . . . . >   R3ERP400       Name, *LIBL, *CURLIB
 Run priority . . . . . . . . . .   30             1-99, *SAME
 Time slice . . . . . . . . . . .   2000           Milliseconds, *SAME
 Eligible for purge . . . . . . .   *YES           *SAME, *YES, *NO
 Default wait time  . . . . . . .   120            Seconds, *SAME, *NOMAX
 Maximum CPU time . . . . . . . .   *NOMAX         Milliseconds, *SAME, *NOMAX
 Maximum temporary storage  . . .   *NOMAX         Kilobytes, *SAME, *NOMAX
 Maximum threads  . . . . . . . .   *NOMAX         1-32767, *SAME, *NOMAX
 Text 'description' . . . . . . .   *BLANK



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

Figure 27-16 i5/OS: Change Class command

The `run priority` parameter specifies the run priority of jobs that use the class. The run priority value ranges from 1 (highest priority) through 99 (lowest priority). All work processes of an SAP system run with a priority of 20 by default.

Refer to Figure 27-16 to see how to change the priority of the ERP System with an instance of 04 to a RUNPTY value of 30.

> **Note:** Changing the run priority of the i5/OS Class Object in the R3*sid*400 Library changes the run priority of all work processes of this SAP system, regardless of the SAP instance.

This method of changing priorities is independent from other methods to change the priority of work processes within an SAP system or SAP instance. It is possible to increase or reduce the run priority level in increments of 4 or 8 for the update, batch, spool, and gateway processes when considering the rules of interconnection.

Set the following parameters in the instance profile to determine the priority:

- rdisp/prio/upd
- rdisp/prio/btc
- rdisp/prio/spo
- rdisp/prio/gwrd

Both methods can be used at the same time. The effects are additive.

Refer to *SAP note 45335 - AS/400: Changing the priority of work processes* for more information about how to control the run priority within an SAP system, relative to an SAP instance.

> **Note:** Set the i5/OS system value QDYNPTYSCD to **1** to allow priorities to be changed.

27.4 Performance management services

This section covers additional services from IBM and SAP to manage your systems performance. It discusses what these services deliver and provides references for more information.

27.4.1 PM iSeries

The IBM Operational Support Services for Performance Management @server iSeries (PM iSeries) capability is an automated, self-maintaining tool for single, multiple or cross-LPAR systems. Similar services are available for AIX and Linux operating systems (PM pSeries and PM pSeries for Linux, respectively).

PM iSeries gathers performance data from a System i server and sends it to IBM on a daily or weekly basis. In return, you receive access to reports, tables, and graphs that show your system's growth and performance calculations. The interactive Web-based service automates the collection and summarization of your performance and capacity data. You can view graphical reports using a standard Web browser with this service to identify how well the systems in your enterprise perform.

iSeries Performance Management allows you to receive summary level cross-LPAR reports beginning with the availability of i5/OS V5R3, and the IBM Director Multiplatform Server option of the IBM Virtualization Engine. These reports show a system-wide view of individual partition utilization for all i5/OS, Linux and AIX 5L partitions.

PM iSeries also helps you with the following:

- Identifying performance bottlenecks before they affect your performance
- Identifying resource-intensive applications
- Maximizing your return on current and future hardware investments
- Planning and managing consistent service levels
- Forecasting data processing growth that is based on trends

PM iSeries offer two levels of service:

- A no-charge offering

  The no-charge offering is available if your System i configuration is covered under a warranty or an IBM Hardware Maintenance Agreement.

  The no-charge option includes the following functions:

  – PM iSeries Management summary graph
  – Cross partition, multiple operating system management
  – PM iSeries Letter of Notification
  – Access to Enterprise View function
  – PM iSeries integration with IBM Workload Estimator

- A full-service offering

  The following benefits are included in the iSeries Performance Management full-service offering. These services are in addition to the no-charge service features:

  – Over 40 detailed reports and graphs
  – Composite system wide view of all partitions
  – Customizable graphs
  – Managing your system with more flexibility
  – PM iSeries graph explanations
  – Explanations of growth and performance issues in non-technical terms
  – Time savings with automatic measuring of the interactive CPU utilization

Note: IBM recommends the daily submission of data. The advantage is being able to see the performance statistics the day following the transmission. Collect performance data 24 hours per day.

The IBM Redbook *Performance Management for IBM @server iSeries and pSeries: A Systems Management Guide*, SG24-7122 offers detailed documentation of the PM iSeries.

Figure 27-17 illustrates the PM iSeries homepage.

Figure 27-17 PM iSeries homepage

The main sections on the left side of the window are as follows:

- Guided Tour
- Existing Customers FAQs

- Upcoming Events
- PM iSeries News
- Activating PM @server System i models

You can find the homepage for iSeries Performance Management on the Web at:

http://www.ibm.com/eserver/iseries/pm

27.4.2 Performance Tools

IBM Performance Tools for iSeries (5722-PT1) is a program product that provides a set of reporting, analysis, and modeling functions to assist a System i administrator manage the performance of the system. It provides printed and online reports in graphic or tabular form.

The Performance Advisor function can assist you in analyzing system performance, and provides recommendations. Through its modeling facility, you can use Performance Tools for iSeries to help predict probable system performance before making changes.

The Performance Tools product offers you a menu interface that provides access to all of the performance tools. Use the Start Performance Report (**STRPFRT**) command or the **GO PERFORM** command on any command line to call the Performance Tools main menu interface.

Figure 27-18 shows the Performance Tools menu interface.

```
PERFORM                IBM Performance Tools for iSeries
                                                             System:   ITSO-SYS1
 Select one of the following:

      1. Select type of status
      2. Collect performance data
      3. Print performance report

      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
```

Figure 27-18 Performance Tools main menu

The data collection services are not an exclusive part of the Performance Tools for iSeries. Data collecting functions are provided within i5/OS.

The i5/OS data collection object (MGTCOL) can be transmitted to other System i models, converted and analyzed if the destination system has at least the same operating system release level.

The Analyze Performance Data (**ANZPFRDTA**) command, also known as the *Advisor*, analyzes performance data, makes suggestions, and offers recommendations for improving performance. To use the Tune System function, you must run the Advisor on the same system on which the data was collected.

Selecting option **3**, **Print Performance Report**, of the PERFORM menu gives you the choice of the following reports:

- System report (PRTSYSRPT)

  This report gives an overview of how the system is operating over a selected period of time.

- Component report (PRTCPTRPT)

  This report expands on the information given in the System Report. It shows the details of each interval to help you find which jobs are consuming high amounts of system resources, such as CPU and disk.

- Job interval report (PRTJOBRPT)

  This job-oriented report is organized by time interval. It helps you find, by interval, which jobs are consuming high amounts of system resources, such as CPU and disk.

- Pool report (PRTPOLRPT)

  The pool-oriented report is organized by time interval. It helps you find which subsystems or pools are consuming high amounts of system resources, such as CPU and disk.

- Resource report (PRTRSCRPT)

  The device resource usage report is organized by time interval. This report shows how system resources are used.

Select option **8** of the PERFORM menu to run the **WRKSYSACT** command. You can use the **WRKSYSACT** command to identify the jobs and tasks that might be causing performance problems in real-time, with an automatic refresh. Unlike the **WRKACTJOB** command of i5/OS, the **WRKSYSACT** command can display system tasks.

The **WRKSYSACT** command shows the CPU utilization and the total number of physical disk I/O operations of all jobs and system tasks based on the elapsed time. All jobs and system tasks are ordered by default by the amount of processing time they use during the interval. However, you can also change the component for sorting.

Performance Tools for iSeries also provides an optional graphical plug-in that can be used by an iSeries Navigator Windows client workstation. This plug-in interface can quickly access any of the sets of performance database files on the system, including up to the last interval of a currently active Collection Services collection. With this plug-in, you essentially get a graphical representation of the functions provided by the Display Performance Data (**DSPPFRDTA**) command on a 5250 interface. Most of the Performance Tools for iSeries reports can be printed. Run the System Report first to get an overview of your system.

Refer to the System Report in Figure 27-19 for examples of:

- CPU utilization
- System value settings

```
                                                     System Report                                  11,11,05 18:00:25
                                                      Workload                                             Page 0001
                                                     825-2 Okt
Member  . . . : Q298090314 Model/Serial  . : 825/65-CONNN       Main storage . . :   20,0 GB  Started   . . . . : 25,10,05 09:03:15
  Library . . : PFR_LIB825 System name . . : AS825AC2           Version/Release  :   5/ 2,0   Stopped   . . . . : 26,10,05 00:00:00
Partition ID  : 000        Feature Code  . : 7416-2473-7416     Int Threshold  . :   0,10 %
Virtual Processors:   5    Processor Units :   5,0
QPFRADJ . . . : 0          QDYNPTYSCD  . . : 1                  QDYNPTYADJ . . . : 1
Interactive Workload
     Job                Number            Average      Logical DB       -------- Printer ----------      Communications        MRT
     Type             Transactions        Response     I/O Count          Lines          Pages             I/O Count        Max Time
--------------    ----------------    --------  ------------------  -------------   ----------      ---------------     --------
Interactive                      2        0,30                   0              0            0                    0
PassThru                     1,302        0,21               1,302            439            9                    0
Total                        1,304                           1,302            439            9                    0
Average                                   0,21
-------------------------------------------------------------------------------------------------------------------------------------
Non-Interactive Workload
     Job           Number         Logical DB       --------- Printer ---------      Communications          CPU Per           Logical
     Type         Of Jobs         I/O Count           Lines           Pages            I/O Count            Logical I/O      I/O /Second
--------------  ---------   -----------------    -------------   -----------      --------------          -----------     -----------
Batch                 990          72,847,149          119,849         3,728                  45              0,0008           1,352,5
Spool                  29               2,989                0             0                   0              0,0118               0,0
iSeries Access          4                   0                0             0                   0              0,0000               0,0
COLLECTION              4               4,654                0             0                   0              0,0047               0,0
DIRSRV                  3                   0                0             0                   0              0,0000               0,0
OS400                   1                   0                0             0                   0              0,0000               0,0
SQL                     2              10,050                0             0                   0              0,0000               0,1
SMTP                    5                   4                0             0                   0              0,1937               0,0
TELNET                 20                   0                0             0                   0              0,0000               0,0
EDRSQL                 78             146,018                0             0                   0              0,0001               2,7
MGMTCENTRAL             2                   0                0             0                   0              0,0000               0,0
NETSERVER               1                   0                0             0                   0              0,0000               0,0
Total               1,139          73,010,864          119,849         3,728                  45
Average                                                                                                       0,0016           1,355,5
Average CPU Utilization . . . . . . . . . . . . . :       23,4
CPU 1 Utilization . . . . . . . . . . . . . . . . :        4,1
CPU 2 Utilization . . . . . . . . . . . . . . . . :        5,2
CPU 3 Utilization . . . . . . . . . . . . . . . . :        3,9
CPU 4 Utilization . . . . . . . . . . . . . . . . :        5,1
CPU 5 Utilization . . . . . . . . . . . . . . . . :        4,8
Total CPU Utilization (Interactive Feature) . . . :        0,2
Time exceeding Int CPU Threshold (in seconds) . . :          0
Total CPU Utilization (Database Capability) . . . :        7,3
```

Figure 27-19 System Report (view 1)

> **Note:** The term `Transaction` in the System Report, shown in Figure 27-19, and the term `Transaction` or dialog steps in connection with SAP, that is, in transaction ST03, are defined very differently. There is no coherence between the data of the two terms.

Refer to the System Report in Figure 27-20 on page 588 for examples of:

- Fault rate per second in Pool 01 (*MACHINE)
- Total Fault rate per second

```
                                              System Report                                        11,11,05 18:00:25
                                        Storage Pool Utilization                                          Page 0005
                                                825-2 Okt
Member  . . . : Q298090314 Model/Serial  . : 825/65-CONNN      Main storage . . :   20,0 GB  Started  . . . . : 25,10,05 09:03:15
  Library . . : PFR_LIB825 System name . . : AS825AC2          Version/Release  :   5/ 2,0   Stopped  . . . . : 26,10,05 00:00:00
Partition ID  : 000        Feature Code  . : 7416-2473-7416    Int Threshold  . :   0,10 %
Virtual Processors:   5    Processor Units :   5,0
                                                                ---------- Avg Per Second ---------   ---- Avg Per Minute -----
Pool  Expert       Size      Act   CPU        Number    Average   ------ DB ------  ---- Non-DB ----    Act-     Wait-    Act-
 ID   Cache        (KB)      Lvl   Util        Tns      Response   Fault    Pages    Fault    Pages     Wait     Inel     Inel
----  -------  ----------- -----  -----  -----------  --------  -------  -------  -------  -------  --------  -------  -------
 01      0      1,433,600      0    0,7            0      0,00      0,0      0,0      0,1      0,3       154        0        0
 02            10,424,640    200    8,1        1,304      0,21     16,1    812,9     32,6     81,2       941        0        0
 03               512,000     50    0,0            0      0,00      0,0      0,0      0,0      0,0         9        0        0
 04             8,192,000     70   14,3            0      0,00      9,4   1162,7     28,4    210,0       317        0        0
Total          20,562,240           23,2        1,304                25,5   1975,7     61,2    291,8     1,423        0        0
Average                                                   0,21
Pool ID                  --  Pool identifier
Expert Cache             --  Method used by the system to tune the storage pool
Size (KB)                --  Size of the pool in kilobytes at the time of the first sample interval
Act Lvl                  --  Activity level at the time of the first sample interval
CPU Util                 --  Percentage of available CPU time used.  This is the average of all processors
Number Tns               --  Number of transactions processed by jobs in this pool
Average Response         --  Average transaction response time
DB Fault                 --  Average number of data base faults per second
DB Pages                 --  Average number of data base pages per second
Non-DB Fault             --  Average number of non-data base faults per second
Non-DB Pages             --  Average number of non-data base pages per second
Act-Wait                 --  Average number of active to wait job state transitions per minute
Wait-Inel                --  Average number of wait to ineligible job state transitions per minute
Act-Inel                 --  Average number of active to ineligible job state transitions per minute
```

Figure 27-20 System Report (view 2)

Refer to the System report in Figure 27-21 on page 589 for the following:

- Average disk response time per I/O operation
- Average number of disk operations per second
- Average number of kilobytes transferred per disk operation
- Average disk operation utilization (percentage busy)
- Percentage of disk space capacity in use
- Disk arm resource name, identifier and disk type
- IOP name and percentage of utilization for each IOP

```
Unit                  --  Disk arm identifier
Unit Name             --  Disk arm resource name
Type                  --  Type of disk
Size (M)              --  Disk space capacity in millions of bytes
IOP Util              --  Percentage of utilization for each Input/Output Processor
IOP Name              --  Input/Output Processor resource name
Dsk CPU Util          --  Percentage of Disk Processor Utilization
ASP Rsc Name          --  ASP resource name to which the disk unit was allocated at collection time
ASP ID                --  Auxiliary Storage Pool ID
Percent Full          --  Percentage of disk space capacity in use
Percent Util          --  Average disk operation utilization (busy)
Op per Second         --  Average number of disk operations per second
K Per I/O             --  Average number of kilobytes (1024) transferred per disk operation
Average Service Time  --  Average disk service time per I/O operation
Average Wait Time     --  Average disk wait time per I/O operation
Average Response Time --  Average disk response time per I/O operation
                                                System Report                              11.11.05 18:00:25
                                               Disk Utilization                                    Page 0007
                                                  825-2 Okt
Member . . . : Q298090314 Model/Serial  . : 825/65-CONNN      Main storage . . :  20,0 GB  Started . . . . : 25.10.05 09:03:15
  Library . . : PFR_LIB825 System name . . : AS825AC2         Version/Release  :  5/ 2,0   Stopped . . . . : 26.10.05 00:00:00
Partition ID  : 000        Feature Code  . : 7416-2473-7416    Int Threshold . :  0,10 %
Virtual Processors:   5    Processor Units :   5,0
         Unit               Size   IOP  IOP         Dsk CPU   ASP Rsc   ASP  --Percent--  Op Per     K Per   - Average Time Per I/O --
  Unit   Name        Type   (M)    Util Name        Util       Name     ID   Full   Util  Second      I/O    Service   Wait   Response
  ----   ---------- ----  -------  ---- ---------- -------  ---------- ---  ----   ----  --------  --------- -------  ------  --------
  0029   DD029      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,42       33,2   ,0017   ,0015     ,0032
  0030A  DD022      4326  35,165   0,4 CMB04          1,0                2  12,8    0,1      5,86       12,5   ,0001   ,0000     ,0001
  0030B  DD057      4326  35,165   0,9 CMB05          1,0                2  12,8    0,1      5,91       13,0   ,0001   ,0002     ,0003
  0031   DD030      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,26       33,9   ,0017   ,0017     ,0034
  0032   DD045      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      6,92       34,2   ,0018   ,0018     ,0036
  0033   DD031      4326  35,165   0,4 CMB04          1,0                1  43,6    1,3      8,19       33,2   ,0015   ,0014     ,0029
  0034   DD056      4326  26,373   0,9 CMB05          1,0                1  43,6    1,0      5,96       32,1   ,0016   ,0014     ,0030
  0035   DD039      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,38       33,3   ,0017   ,0015     ,0032
  0036   DD040      4326  35,165   0,9 CMB05          1,0                1  43,6    1,4      8,13       32,5   ,0017   ,0014     ,0031
  0037   DD035      4326  26,373   0,9 CMB05          1,0                1  43,6    1,1      6,29       32,0   ,0017   ,0013     ,0030
  0038   DD044      4326  35,165   0,9 CMB05          1,0                1  43,6    1,4      8,29       31,7   ,0016   ,0015     ,0031
  0039   DD041      4326  30,769   0,9 CMB05          1,0                1  43,6    1,2      8,18       29,8   ,0014   ,0014     ,0028
  0040   DD058      4326  26,373   0,9 CMB05          1,0                1  43,6    1,0      6,47       28,1   ,0015   ,0012     ,0027
  0041   DD037      4326  26,373   0,9 CMB05          1,0                1  43,6    1,1      5,95       32,3   ,0018   ,0012     ,0030
  0042   DD038      4326  30,769   0,9 CMB05          1,0                1  43,6    1,2      6,91       33,6   ,0017   ,0013     ,0030
  0043   DD032      4326  26,373   0,9 CMB05          1,0                1  43,6    1,1      6,53       29,6   ,0016   ,0011     ,0027
  0044   DD046      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,34       31,4   ,0017   ,0017     ,0034
  0046   DD036      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,33       33,5   ,0017   ,0015     ,0032
  0047   DD043      4326  30,769   0,9 CMB05          1,0                1  43,6    1,2      7,15       33,1   ,0016   ,0016     ,0032
  0048   DD050      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,11       33,7   ,0018   ,0015     ,0033
  0049   DD059      4326  30,769   0,9 CMB05          1,0                1  43,6    1,3      7,24       33,6   ,0017   ,0016     ,0033
  0055A  DD023      4326  35,165   0,4 CMB04          1,0                3  62,3    0,4      8,08       26,3   ,0004   ,0000     ,0004
  0055B  DD047      4326  35,165   0,9 CMB05          1,0                3  62,3    0,4      8,20       26,5   ,0004   ,0002     ,0006
Total                   1,793,391
Average                                                                     43,2    1,1      7,20       31,1   ,0015   ,0014     ,0029
```

Figure 27-21 System Report (view 3)

Consider the time interval of the System Report. In this example, the time is analyzed between 09:03 a.m. and midnight.

Other results can be analyzed for the following time intervals:

- An entire day
- A typical business day (for example, between 08:00 a.m. and 05:00 p.m.)
- A high workload hour (for example, between 10:00 a.m. and 11:00 a.m.)
- A peak workload interval of 5 minutes

> **Note:** You can generate all reports for the whole measurement period or for one or some of the measurement intervals.

The shortest period for generating a report is a single measurement interval. The measurement intervals by default are 15 minutes. Use the Change Performance Collection command (**CHGPFRCOL**) to set up the interval between 30 minutes and 15 seconds.

> **Note:** Adapt your measurement intervals and report generating to suit your analysis.

For detailed information about all the functions of the Performance Tools for iSeries product, refer to the following Web sites:

http://www.ibm.com/servers/eserver/iseries/perfmgmt/pt400.html
http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp?lang=en

Also refer to the following publications:

- *IBM @server iSeries Performance Management Tools*, REDP-4026
- *IBM iDoctor iSeries Job Watcher: Advanced Performance Tool*, SG24-6474

27.4.3 iDoctor for iSeries

The growth in data processing capabilities on the System i family of servers along with the introduction of complex e-business applications has vastly increased the amount of data processed in only a few seconds. iDoctor for iSeries, software and associated services, is key to evaluate the health of your system. iDoctor automates the analysis of this vast amount of data in a timely manner, which is essential in order to resolve many of today's complex system and application problems.

The iDoctor for iSeries tools provides you with information at the application developer level. The iDoctor family of products and services includes the following:

- Job Watcher

  Job Watcher displays real-time tables and graphical data that represents in detail what a job is doing and why it is not running. Job Watcher provides several reports that provide detailed job statistics by interval. These statistics allow you to determine things like CPU utilization, direct access storage device (DASD) counts, waits, faults, call stack information, conflict information, and more.

- Performance Explorer (PEX) Analyzer

  PEX Analyzer evaluates the overall performance of your system and builds on what is available with Performance Tools. The Analyzer condenses volumes of trace data into reports that can be graphed or viewed to help isolate performance issues and reduce overall problem determination time.

  The Analyzer provides an easy-to-use graphical interface for analyzing CPU utilization, physical disk operations, logical disk input/output, data areas, and data queues. The Analyzer can also help you isolate the cause of application slowdowns.

- Consulting services

  If you want expert consultants to analyze your system using one of the in-depth software tools from the iDoctor for iSeries Suite (PEX Analyzer or Job Watcher), select the **Consulting Services** component.

- Java Watcher and Heap Analysis Tools for Java

  Java Watcher provides invaluable information to help in the debugging some of the most complex problems in the area of Java and WebSphere.

- Performance Trace Data Visualizer (PTDV)

  The Performance Trace Data Visualizer for iSeries (PTDV) is a tool for processing, analyzing, and viewing Performance Explorer collection trace data residing in performance explorer database files. PTDV is a free component of iDoctor for iSeries.

iDoctor for iSeries uses the two following performance data sets:

- The Collection Services Database, included in the i5/OS operating system
- The more detailed Performance Explorer database

Figure 27-22 on page 591 shows the performance data sets used by iDoctor.

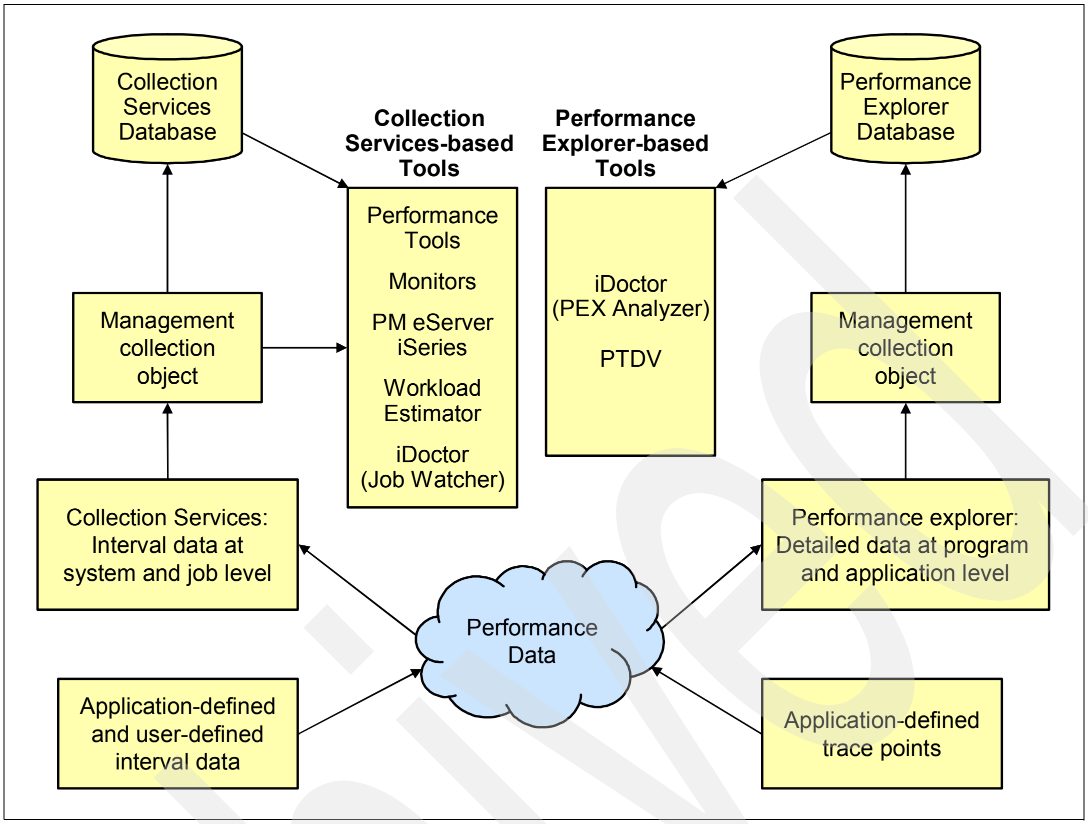

Figure 27-22 Performance data sets used by iDoctor

The Performance Explorer (PEX) controls the data collection for the Performance Explorer Database (QPEXDATA). Therefore, the PEX Explorer uses operating system commands, for example **STRPEX**, or application programming interfaces (APIs). Performance Explorer is a no-charge i5/OS detailed performance data collection tool.

The following interfaces are available to collect data for the Collection Services database:

- Operating system commands
- PM/400 menu
- iSeries Navigator
- APIs

Figure 27-23 shows the controlling interfaces for the Collection Services database.

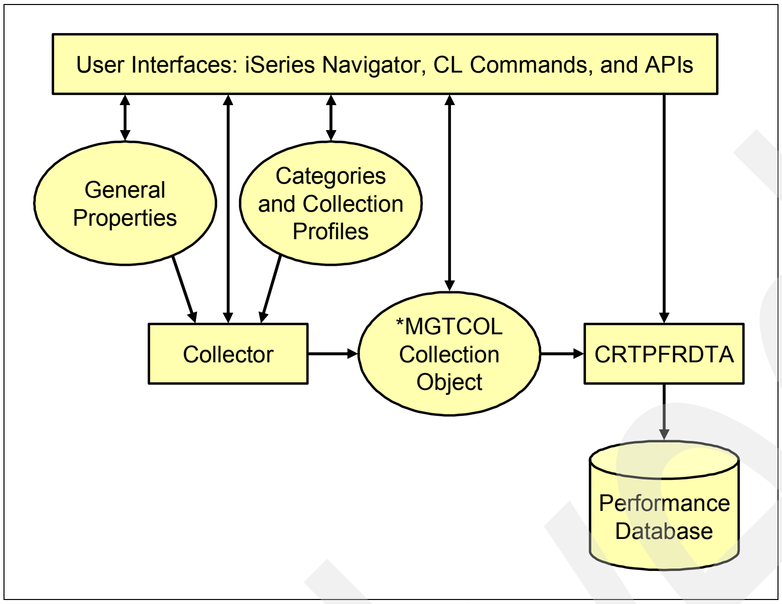

Figure 27-23 Controlling interfaces for Collection Services database

Job Watcher and PEX Analyzer work differently in analyzing the details about the workload of the system.

Job Watcher returns near-real-time information about a select set of jobs, threads, or licensed internal code (LIC) tasks. It is similar, as a sampling function, to the system commands **WRKACTJOB** and **WRKSYSACT**, in which each refresh computes delta information for the ending snapshot interval. In the Job Watcher, these refreshes can be set to occur automatically, even as frequently as every five seconds. Job Watcher harvests the data from the jobs, threads, or tasks being watched in a manner that does not affect other jobs on the system while it is collecting.

You can use PEX data that is collected by the PEX Explorer to help find application program-level causes of performance problems that cannot be identified by the Collection Services data and general performance monitors. You can define and start or end a PEX collection under one of three options:

- *STATUS
- *PROFILE
- *TRACE

The iDoctor for iSeries PEX Analyzer provides extended PEX Trace data analysis details, as compared to the Performance Tools for iSeries PEX reports. This output can include the PEX information provided by the Performance Tools for iSeries PEX printed data reports plus additional information about wait states, details about the objects referenced, and disk space consumption.

The Performance Tools for iSeries licensed program provides printed report options for PEX *STATUS and *PROFILE data that includes a low-level summary of disk operations, CPU utilization, file opens, and programs, including machine instruction (MI) programs invoked. You can use printed PEX Profile data to identify actual high-level languages statements causing significant CPU utilization, assuming that program compiler statement optimization levels have not been used on the programs.

To analyze SAP system performance problems, use iDoctor with guidance from IBM in critical cases. For detailed information about iDoctor for iSeries, refer to the following Web sites:

http://www-912.ibm.com/i_dir/idoctor.nsf/iDoctor.html
http://www-912.ibm.com/i_dir/idoctor.nsf/faqNewAtV5R3.html
http://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/$FILE/Job_Waits_White_Paper_V5R3.pdf

Also, refer to the *IBM iDoctor iSeries Job Watcher: Advanced Performance Tool*, SG24-6474.

27.4.4 Workload Estimator

The purpose of the IBM eServer Workload Estimator (WLE) is to provide general sizing recommendations for System i servers running one or more workloads. WLE allows you to combine different workloads from several systems to a new workload.

WLE is not particular to SAP applications. You can use the i5/OS Collection Services, Performance Tools for iSeries (5722-PT1) System Report and IBM Insight Tool for better SAP-related information.

Information about the IBM eServer Workload Estimator is available on the Web at:

http://www.ibm.com/eserver/iseries/support/estimator

27.4.5 SAP services

SAP provides many services about your SAP system landscape, as discussed in this section.

SAP EarlyWatch Alert Monitoring in SAP Solution Manager

SAP EarlyWatch Alert (EWA) is a free monitoring service for SAP customers that monitors all SAP systems in a system landscape. SAP recommends that customers perform centralized system monitoring of all connected satellite systems in the SAP Solution Manager.

Use the SAP Solution Manager to create service-level reports that are tailored to individual requirements (for example, to use customized transactions and threshold values to trigger the creation of the reports). One advantage of using the SAP Solution Manager is the ability to create real-time reports and simple archiving.

All EWA-enabled satellite systems connected to the SAP Solution Manager send their EWA data to SAP once a month. This automatically keeps the system and software component data for SAP systems that is stored in the SAP Support Portal up-to-date.

EWA data is immediately sent to SAP for further analysis if the EWA data for an SAP production system is marked with a red light after evaluation. You receive a call from SAP if the alert is critical.

You can find more information about EWA monitoring on the Web at:

http://service.sap.com/solutionmanager

SAP EarlyWatch Alert Monitoring by SAP

The SAP EarlyWatch Alert is a part of SAP Standard Support, included in SAP maintenance, and plays an important role in SAP's safeguarding program, which ensures technical robustness in your SAP application. You can continue to let SAP carry out SAP EarlyWatch Alert evaluation for you even if you do not use SAP Solution Manager to monitor your systems.

For more information refer to:

http://service.sap.com/ewa

SAP EarlyWatch Check

The *EarlyWatch Check* service, formerly known as the *EarlyWatch Service*, is divided into two services:

- SAP EarlyWatch Check for the analysis
- SAP Remote Performance Optimization for the optimization of an installation

SAP EarlyWatch Check is a service in the SAP Safeguarding portfolio. It analyzes the SAP system for resource and system bottlenecks, and potential error situations, as well as makes recommendations for optimization.

An SAP EarlyWatch Service Check session offers an SAP EarlyWatch report. The SAP EarlyWatch Check is carried out as a remote service.

For more information refer to 27.1, “SAP EarlyWatch” on page 566 and the following Web site:

http://service.sap.com/earlywatch

SAP Remote Performance Optimization

SAP Remote Performance Optimization is a service in the SAP Solution Management Optimization portfolio. It is designed to tackle bottlenecks which have a large impact on your system performance. This service focuses on Structured Query Language (SQL) statement optimization, customer program optimization, or application optimization. The SAP Remote Performance Optimization service is a remote service.

For more information, select **Support Services** → **SAP Solution Management Optimization** → **SAP Remote Performance Optimization** on the Web at:

http://service.sap.com/servicecat

SAP GoingLive Check

SAP GoingLive Check is a service in the SAP Safeguarding portfolio. It includes analysis, optimization, and verification of system and database sizing, system configuration, modifications, and technical aspects of the applications before going live. The project planning and the technical realization of key business processes for SAP implementations are also checked. The SAP GoingLive Check is carried out as a remote service.

For more information refer to:

http://service.sap.com/goinglivecheck

SAP GoingLive Functional Upgrade Check

SAP GoingLive Functional Upgrade Check is a service in the SAP Safeguard portfolio. It includes analysis, optimization, and verification of system and database sizing, system configuration, modifications, and technical aspects of the applications before entering production after the SAP system is upgraded to a higher release. The SAP GoingLive Functional Upgrade Check is carried out as a remote service.

For more information refer to:

http://service.sap.com/goinglive-fu

SAP Customer Services Network

The SAP Customer Services Network (formerly known as mySAP Services) creates value by providing consulting, trusted advice, proactive services, and collaborative learning through expertise in business system consulting, system operation services, education and support. Many of these services also cover performance aspects.

The SAP Customer Services Network includes the following parts:

- SAP Active Global Support
- SAP Consulting (Strategic Consulting Services, Solution Delivery Services, Operations Services, Life-Cycle Management Services)
- SAP Education
- SAP Custom Development
- SAP Hosting

For more information refer to the following Web sites:

http://service.sap.com/servicesmap
http://service.sap.com/servicecat

SAP newsletter

Select **Home** → **My Subscriptions** in the following Web site to maintain your user profile in the SAP Support Portal and to receive e-mail newsletters on subjects of interest:

http://service.sap.com/CCCNet

28

SAP Business Information Warehouse

mySAP Business Information Warehouse (BW) is the main component of SAP NetWeaver Business Intelligence (BI). mySAP BW provides support for strategic and operational decision making by integrating data from existing systems and returning valuable business information for analysis. It also serves as the backbone to many SAP applications, such as mySAP Customer Relationship Management (CRM), mySAP Strategic Enterprise Management (SEM) and mySAP Supply Chain Management (SCM).

This chapter discusses:

- Settings for enabling star join support in SAP BW systems
- Proper index strategy for InfoCubes and Operational Data Store (ODS) tables
- Identifying bottlenecks in BW query performance
- Tips for improving data load performance

General System i and SAP performance tips are described in Chapter 25, “Performance concepts and monitoring” on page 457. Database performance is discussed in 25.6, “Database monitoring” on page 516.

28.1 Query performance

Achieving good query performance requires a combination of factors. It starts with creating a solid index strategy, complementing the strategy with good statistics, and then modifying the strategy as users change their query behavior. This section discusses the table structure SAP BW uses for InfoCubes and Operational Data Stores (ODS), and the default index strategy for each table structure.

28.1.1 InfoCubes

Most SAP BW queries contain selection criteria over multiple dimensions of the fact table. This indicates that the query is a good candidate for using the star join optimizer support in DB2 UDB for i5/OS. However, at times there are queries that return large amounts of data or contain selection over few dimensions. These queries behave like traditional SAP queries. This indicates that a variety of encoded vector indexes (EVI) and radix indexes are needed to satisfy all of the user requests.

The default index strategy that is used when implementing the recommendations in "Settings for optimal BW query performance" on page 606 satisfy the star join queries, but additional indexes may be needed for other queries. This section discusses what makes up the extended star schema, the default indexes provided, and identifies steps for helping problem queries.

Extended star schema

SAP applications use an extended star schema when defining and creating InfoCubes. It contains two types of data used in analysis:

- Key figures: Such as sales revenue, fixed costs, sales quantity, or the number of employees
- Characteristics: For example, customer type, fiscal year, period, or region. Characteristics are used to create evaluation groups for analysis.

The underlying InfoObjects that make up an InfoCube are categorized in terms of these two types of data. That is, a given InfoObject represents either key figures or characteristics. A third type of InfoObject, *attributes*, contains metadata describing other InfoObjects.

An InfoCube is represented in the database as a set of relational tables, as shown in Figure 28-1 on page 599. These are arranged according to the star schema, a technique that organizes data in terms of data facts and business dimensions.

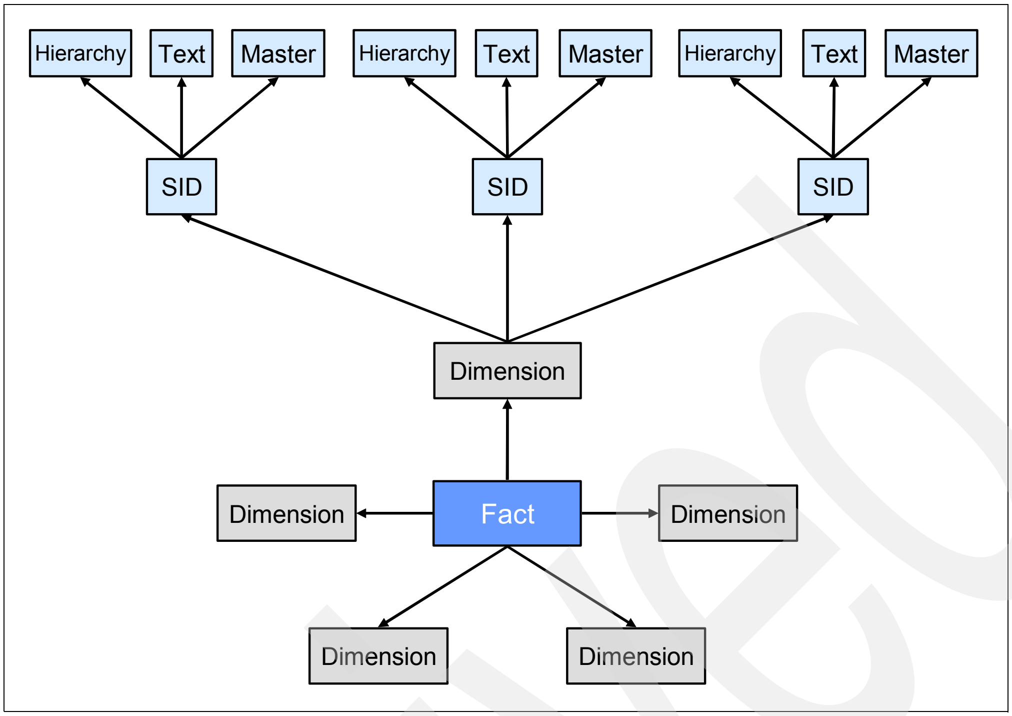

Figure 28-1 Extended star schema

The star schema places several dimension tables around a central fact table. The fact table stores key figures, while the surrounding dimension tables store the characteristics needed for evaluating and reporting on those key figures. Fact tables are the largest tables in star schemas and can contain billions of entries.

Dimension tables are independent of each other. The fact table links the dimension tables and the key figures. To link these tables, dimension identifiers are used. A dimension identifier uniquely identifies a particular combination of characteristic values in a dimension table, for example, a specific product and the corresponding product group. Characteristics that are correlated, such as product group, are usually put in the same dimension.

An InfoCube in SAP BW can have up to 16 dimensions. By default, every InfoCube has three standard dimensions: data package, time and unit.

SAP BW uses an extended star schema, which builds on the basic star schema by storing master data about attributes, hierarchies and text in separate tables that are shared between InfoCubes. This reduces data redundancies because master data is stored only once and then used by various InfoCubes.

Actual characteristic values, such as the name of a region, are shown in the dimension tables. In reality, characteristic values are replaced by surrogate identifiers (SIDs). These are numeric key values (4-byte integers) that are stored in the master table. Therefore, there are foreign key relationships between each characteristic in a dimension table and the corresponding attribute, hierarchy and text tables. SIDs are used to keep dimension tables as small as possible, since they can also grow very large. Figure 31-2 illustrates the use of surrogate identifiers in the fact, dimension and SID tables.

Fact Table: /BIC/FSAMPLE

KEY_SAMPLEP	KEY_SAMPLE_T	KEY_SAMPLEU	KEY_SAMPLE1	KEY_SAMPLE2	KEY_SAMPLE3	COST	QUANTITY
4	800	1	1743	20	8231	9283.04	283
4	932	1	2937	12	32013	2932.24	4726
4	820	1	231	39	3291	8249.90	74
4	173	1	103	20	239	2380.98	823
4	247	1	3092	93	1348	384.02	23

DIVISION	SID	CHCKFL	DATAFL	INCFL
	0	X	X	X
Division 2	2	X	X	X
Division 3	3	X	X	X
Division 4	4	X	X	X

SID Table: /BI0/SDIVISION

DIMID	SID_0DIVISION	SID_0SALESORG	SID_0DISTRCHAN
0	0	0	0
11	3	8	2
12	4	3	10
13	5	20	23

Dimension Table: /BIC/DSAMPLE2

Figure 28-2 Use of surrogate keys in SAP BW

SAP BW provides the option to define a very large dimension as a line item dimension. In this case, the corresponding dimension is not stored as a separate table, but rather is stored directly in the fact table. This eliminates the necessary join operation between the dimension and fact table during SAP BW query processing, which can provide improved query performance.

For an overview of the use of extended star schema and operational data store in an SAP BW environment, refer to Chapters 1 and 2 in *Building and Scaling SAP Business Information Warehouse on DB2 UDB ESE*, SG24-7094.

Star join support (LPG or EVI Stage 2)

Before discussing an index strategy for an InfoCube, we need to understand star join queries and how the optimizer handles these star joins. While the actual implementations for Lookahead Predicate Generation (LPG) and EVI Stage 2 are different, the basis for how they work are similar. For this discussion we describe the LPG method.

> **Note:** Additional information about star schema join support can be found in IBM whitepaper *Star Schema Join Support with DB2 UDB for iSeries* at:
>
> `http://www-03.ibm.com/servers/enable/site/education/wp/star/star.pdf`

Because of the design of the extended star schema, the selection criteria is pushed out to the outer boundaries. SAP BW attempts to push the criteria, but it usually falls into the dimension or SID tables. This leads to three scenarios:

- The selection criteria is very general and results in either returning a large number of rows from the fact table or summarizes a large number of rows from the fact table.
- The selection criteria is heavily dependent on one dimension. This selection reduces the number of rows read from the fact table.

- The selection is across multiple dimensions with no single dimension. This provides a drastic reduction in the number of rows read from the fact table. When the selection is combined together, the number of rows read from the fact table is drastically reduced.

The first two situations are handled like normal SAP-types of queries. A table scan is required (scenario 1) or one of the dimension tables can be placed first in the join order (scenario 2). The final scenario is considered a star join and is considered a candidate for LPG or EVI Stage 2 support.

Star join support relies on combining the bitmaps of EVIs over the key columns in fact tables to reduce the number of reads against the fact table. EVIs contain two parts: a symbol table and a vector. The symbol table contains statistical and descriptive information for each distinct key in the table. The vector is an array of codes listed in the same ordinal position as the rows in the table. Figure 28-3 shows an example of an EVI.

SQL: CREATE ENCODED VECTOR INDEX COLOR_EVI ON LINEITEM (COLOR)

Table: LINEITEM

Sales Order	Line Item	Product	Color	…
1	1	1029	Black	…
1	2	12	Blue	…
1	3	712	Green	…
…				
1242	11	34	Green	…

EVI: COLOR_EVI

Symbol Table

Value	Code	First Row	Last Row	Count
Black	1	1	3267	421
Blue	2	4	1415	642
Green	3	17	3345	294
…				
Yellow	11	34	2156	172

Vector

Row Number	Code
1	1
2	5
3	4
4	2
5	8
6	10
7	4
…	
3343	6
3344	6
3345	3

Figure 28-3 Encoded vector index example for column COLOR

The easiest way to understand how star joins are processed is to walk through an example: Let us say that we want to see all return data for a range of materials by division and month for a specific sales organization over a calendar year. Figure 28-4 on page 602 shows what the generated SAP BW query can look like.

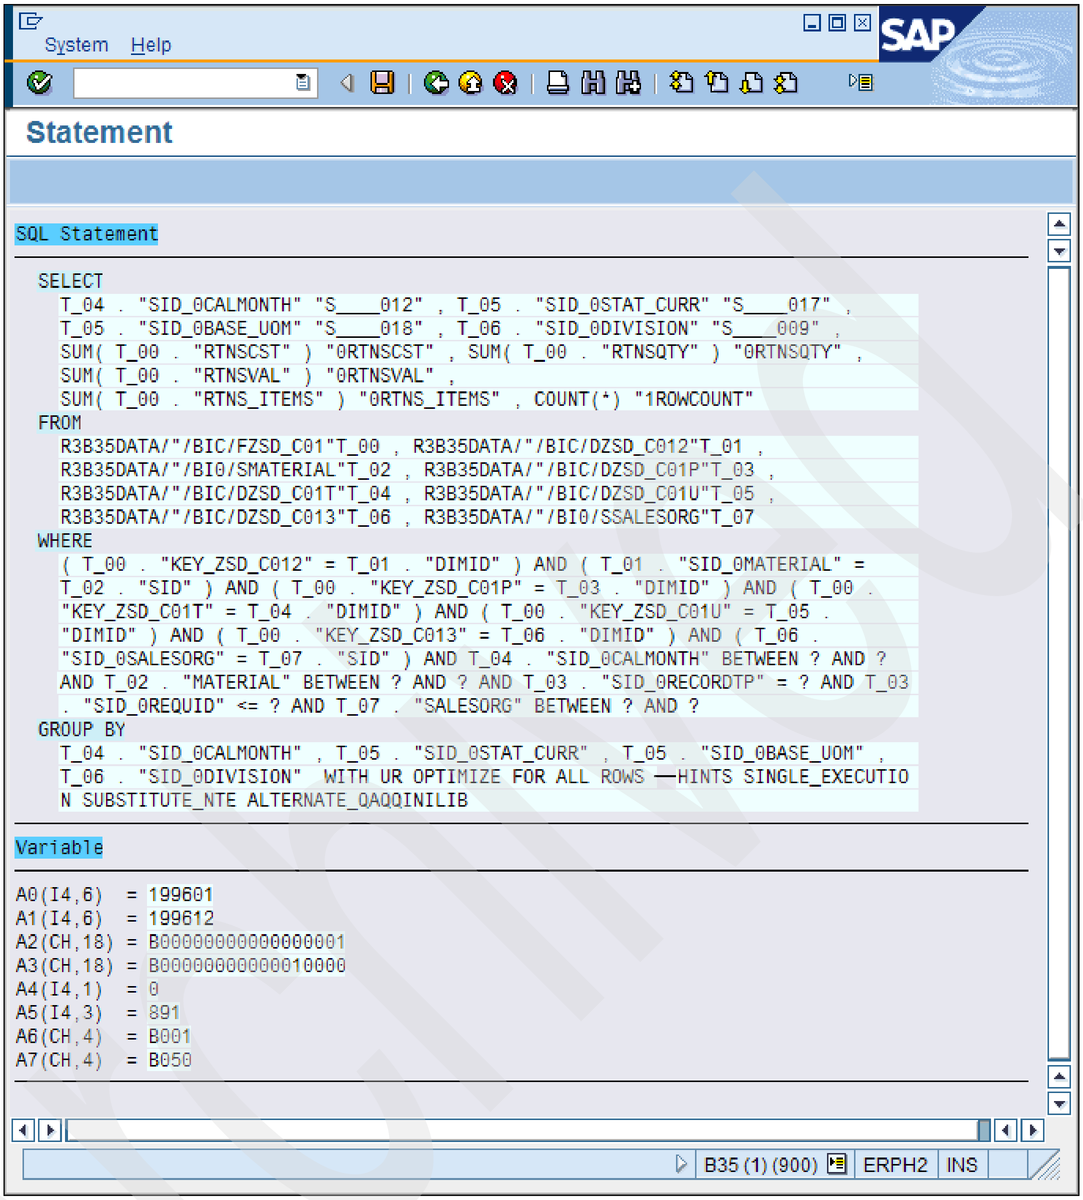

Figure 28-4 Star join example query

The optimizer goes through the following steps in building an LPG access plan:

1. Decide on the proper join order. The optimizer identifies the largest table as the fact table but may or may not place it in the first position. SQE has more access methods available, so it is not restricted in the join order, like CQE, for star join implementations. CQE requires the fact table to be in the first position for EVI stage 2 support.
2. Each dimension table is handled separately in order to create a hash table filled with the data returned by any selection criteria over that specific dimension. A list of distinct keys is created for identifying the join keys in the fact table. SQE does not restrict all dimensions to hash join. It can pick a different access method for dimensions not involved in the bitmap or relative record number (RRN) list processing. In CQE, all joins must be hash joins in order to use EVI stage 2 support.
3. The query is rewritten with the distinct key lists providing local selection over the fact table. This pushing of the key lists to the fact table is referred to as look-ahead predicate

generation. Figure 28-5 shows how the selection over column MATERIAL in the example query would be rewritten so the selection is now over the fact table (“/BIC/FSD_C01”) instead of the outer SID table (“/BIC/SMATERIAL”). This process of pushing down the selection to the fact table is repeated for all selection criteria that can be moved to the fact table.

Original Query:

```
SELECT …
FROM
  “/BIC/FSD_C01” F,
  “/BIC/DSD_C012” D2,
  “/BIC/SMATERIAL” S1,
…
WHERE
  F.KEY_SD_C012 = D2.DIMID AND
  D2.SID_0MATERIAL = S1.SID AND
  … AND
  S1.MATERIAL BETWEEN ‘B0001’
     AND ‘B0100’ AND
…
GROUP BY …
```

Rewritten Query:

```
SELECT …
FROM
  “/BIC/FSD_C01” F,
  “/BIC/DSD_C012” D2,
  “/BIC/SMATERIAL” S1,
…
WHERE
  F.KEY_SD_C012 = D2.DIMID AND
  D2.SID_0MATERIAL = S1.SID AND
  … AND
F. KEY_SD_C012 IN ( 1, 3, 10, …)
 …
GROUP BY …
```

Figure 28-5 Example of query rewrite process to pushdown predicates

4. The query is optimized again using the new local selection over the fact table. The query optimizer can now take advantage of the selection criteria that was pushed down to the fact table to identify which EVIs can be combined to reduce the number of records retrieved from the fact table.
5. A dynamic bitmap or RRN list is created for each EVI selected. Each bitmap or RRN list identifies all records in the fact table that satisfy the specific selection over the column the EVI is for. The keylists generated in step 2 provide the list of values that satisfy the selection criteria. These values can then be used by the EVI to generate either a dynamic bitmap or RRN list. Figure 28-6 on page 604 shows the process of creating the dynamic bitmaps.

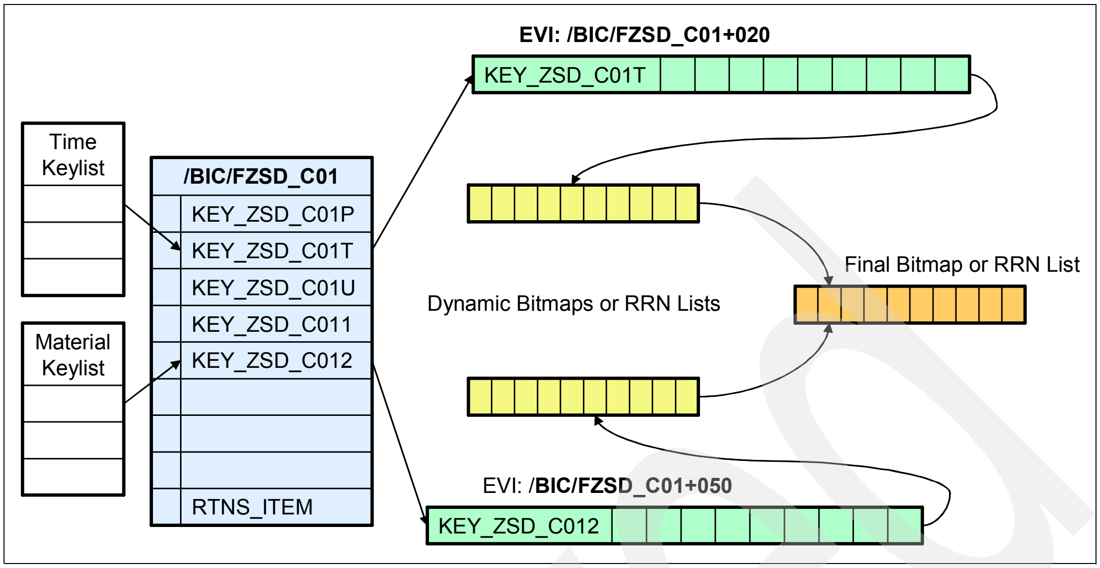

Figure 28-6 Creating dynamic bitmaps

The bitmap is a vector representing each record in the table and the value of each location of the vector is either a 1 indicating the record satisfied the selection criteria or a 0 indicating the record did not satisfy the selection criteria. The RRN list is a listing of all the relative record numbers that satisfy the selection criteria.

6. If more than one EVI is selected, then a final bitmap is created by merging the bitmaps or RRN lists from each of the selected EVIs.
7. The final bitmap or RRN list is used to process records from the fact table. This step is often called skip sequential processing in CQE or clustered table probe processing in SQE. Both of these methods use the final bitmap or RRN list to skip over rows that are not identified in the bitmap or RRN list as potential candidate rows.
8. Each selected row is then joined to the remaining tables.
9. After the join to all tables is complete, then any grouping and ordering is performed.

The use of LPG can be identified in Visual Explain with iSeries Navigator or in the Explain function of ST05. The Visual Explain diagram for the example query is displayed in Figure 28-7 on page 605 with the LPG nodes highlighted using function View → Highlight LPG. This access plan uses EVIs over the time (“/BIC0131”) and salesorg (“/BIC0135”) dimensions to build RRN lists that merged together for processing the fact table (/BIC/FZSD_C01).

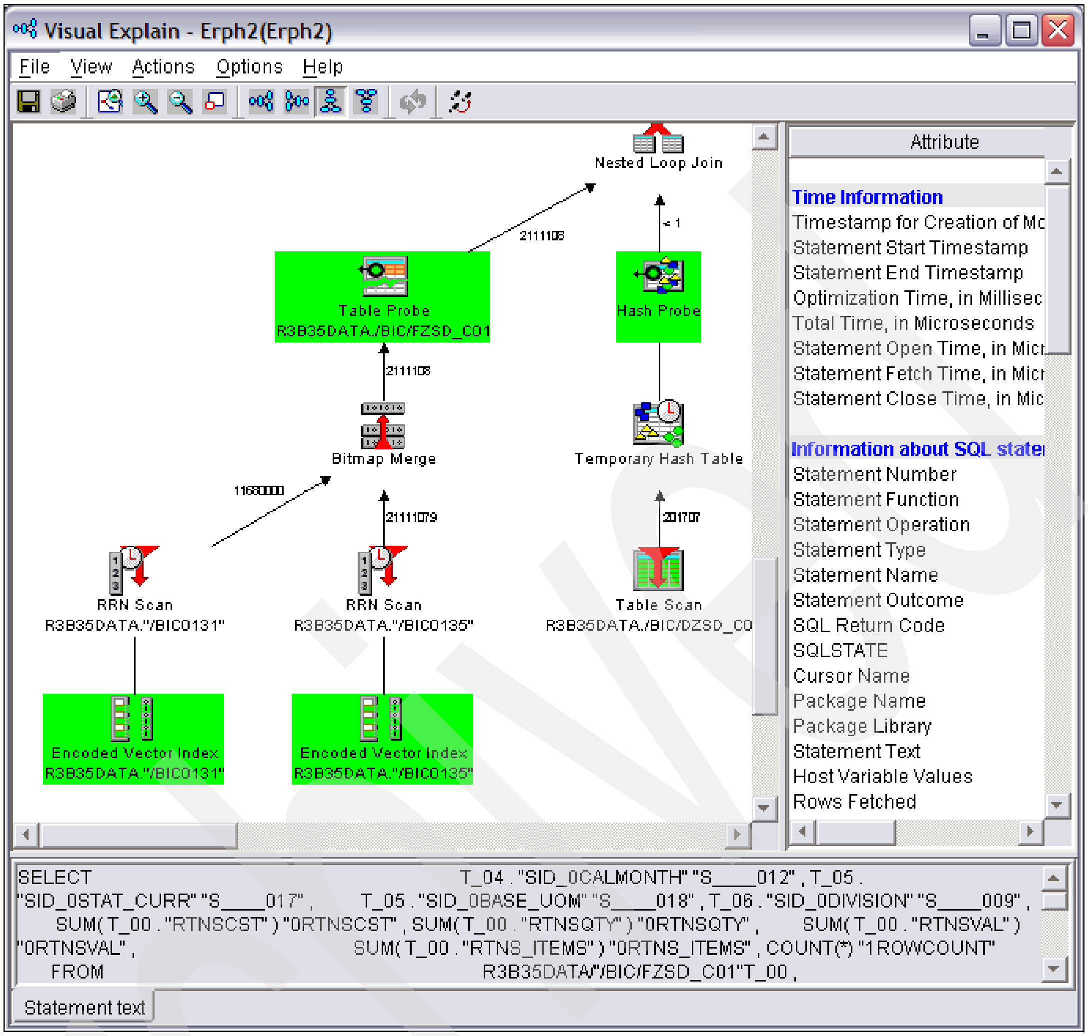

Figure 28-7 Visual Explain for statement using LPG

The same query with different host variables is shown in Figure 28-8 on page 606. The join order has changed, but the use of LPG is still present for processing the fact table. The circled area in Figure 28-8 on page 606 indicates the use of clustered table probe processing with the merging of two EVIs.

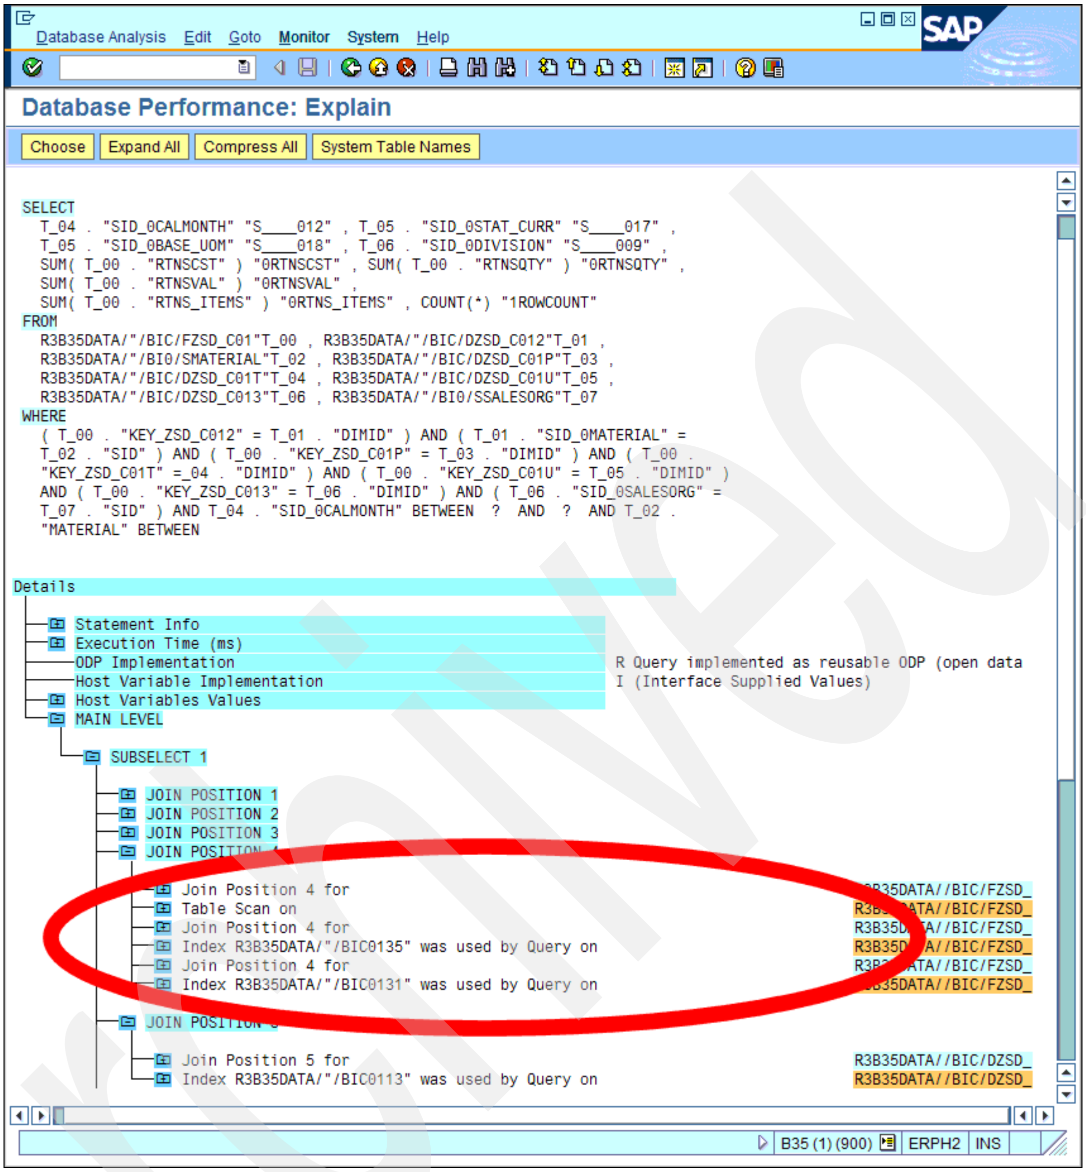

Figure 28-8 ST05 output for statement using LPG

Settings for optimal BW query performance

The SAP BW application and other SAP applications that are based on BW, such as SCM and SEM, run like most SAP applications, except for the extended star schema used in InfoCubes. The recommendations in Chapter 25, “Performance concepts and monitoring” on page 457 can be used for SAP BW. The difference is SAP BW needs additional database support for the recognition of star joins. This section describes why these enhancements are required.

> **Attention:** SAP note 501572 contains the steps required for implementing the star join support.

In OS/400 V4R5, Encoded Vector Indexes (EVI) stage 2 support was added to the optimizer for star join queries. This support requires specific options in the query options file (QAQQINI). The database optimizer was rewritten at OS/400 V5R2 and queries are either run with the new optimizer (SQE) or continued using the previous optimizer (CQE). Since star join support remains in CQE only, all SAP BW queries that are possible star joins continue down the CQE path.

In i5/OS V5R3, star join support was rewritten to use the Lookahead Predicate Generation (LPG) algorithm and included in the base SQE optimizer. The inclusion in the base costing of SQE makes the query options file unnecessary for most star join queries. There are a few instances, however, where star join queries continue to use CQE. Therefore, the use of the QAQQINI file is still required.

Table 28-1 shows the recommended options for enabling star join support in CQE.

Table 28-1 Recommended QAQQINI file options for star join support

Option	Value	Description
STAR_JOIN	*COST	Enables star join costing in CQE optimizer.
ALLOW_TEMPORARY_INDEXES	*ONLY_REQUIRED	Only uses a temporary index for joins, ordering or grouping if no other option is available. Used for forcing hash group by.
OPTIMIZATION_GOAL	*ALLIO	Optimizer hint indicating the goal of the query is to run to completion rather than returning the first page as soon as possible.

Note: The ALLOW_TEMPORARY_INDEXES is not required for star join support. It is used to help query performance. The default group by costing often overestimates the number of groups. This causes the slower index group by method to be selected. If the number of groups is high, then the hash group by method may not fit in memory and performance drastically degrades. Remove this option if performance degrades.

Note: For OS/400 V5R2 and prior releases, the QAQQINI option FORCE_JOIN_ORDER must be set to *PRIMARY. Refer to SAP note 501572 for details on older releases.

Additional SAP code changes are required to take advantage of the CQE star join support. The changes documented in a hint about SAP BW star join queries indicate an alternate QAQQINI file should be used.

Table 28-2 lists the required instance profile parameters for enabling star join support.

Table 28-2 Instance profile settings for star join support

Profile parameter	Recommended value	Description
as4/evistage2support	1	Enables star join support.
dbs/db4/alternate_qaqqinilib	library name	Location of QAQQINI file.
dbs/db4/use_hints	1	Enables use of hints.

The default BW index strategy is the use of single column radix indexes. The use of star join support requires Encoded Vector Indexes (EVI) over each of the dimension columns of the fact tables. Run the SAP_INFOCUBE_INDEXES_REPAIR report if you have existing

InfoCubes. If the existing InfoCubes contain a large number of rows, this report can run for a long time. Run it in the background.

Default index strategy

InfoCubes have two different index strategies depending on the settings described in “Settings for optimal BW query performance” on page 606. If the as4/evistage2support instance profile parameter is not defined or set to a value of 0, then the default index strategy relies on radix indexes. This strategy works best if there are not many multidimensional, or star join queries. If the as4/evistage2support instance profile parameter is turned on, or set to 1, then the index strategy relies on EVIs. Both strategies contain a single radix index over all dimension key columns and separate single column indexes over each of the dimension key columns.

28.1.2 Operational Data Store

An Operational Data Store (ODS) object stores consolidated transaction data from one or several InfoSources. In contrast to the multidimensional data models of InfoCubes, data in ODS objects is stored in flat and transparent database tables. ODS object data can be updated into InfoCubes or other ODS objects using a delta update, which loads only new or changed data into the destination object. Data in an ODS object can be analyzed with the SAP BW Business Explorer (BEx).

An ODS object contains a key (for example, order number) as well as data fields (key figures). As opposed to InfoCubes, fields in an ODS object can be overwritten. This is useful to process document structures. For example, if you change documents in an OLTP system, the changes not only affect numeric fields, such as order quantity, but also non-numeric fields, such as status and delivery date. To reflect these changes in the ODS object, the relevant fields in the ODS objects must be overwritten and set to the current value.

On the database level, every ODS object consists of three transparent tables:

- Active table: The data in this table is used for reporting.
- Activation queue: This table contains data that is new or modified since the last activation.
- Change log table: The change log is used for the delta update from the ODS object into other ODS objects or InfoCubes.

All end user queries are run over the active table. Because of the multifunctionality of these tables, the default index strategy for these tables is limited. The only index provided is a primary key index over all the key columns.

The default strategy does not provide much help for end user queries, unless the primary index can be used. If you are going to be using BEx to query the ODS table, then additional indexes are going to be needed. Since these queries usually resemble an OLTP type of query, the easiest method would be to create single column radix indexes over all the key columns of the ODS table. This allows for some joins with the SID tables that more than likely contain the selection criteria. Use either transaction ST04, see 25.6, “Database monitoring” on page 516 to information about using the database monitor, or end user complaints to identify poor performing queries that require additional tuning, such as multicolumn radix indexes or materialized query tables (MQT).

28.2 Troubleshooting poor performing queries

This section looks at what tools are available to debug the possible causes of slow BW queries and what can be done to improve the performance of the query.

28.2.1 Query Monitor

The query monitor tool, transaction RSRT, provides a debug environment for testing BEx reports over both InfoCubes and ODS tables.

To get information for a problem report, enter the query name and click the ‘Execute + Debug’ button.

A pop-up window opens with many options. Select the appropriate options to get the desired information, such as the runtime statistics (Display Statistics Data), test a new aggregate (Select Aggregate), disallow the use of aggregates (Do Not Use Aggregates), or to capture the SQL query text (Display SQL Query and Do not Use OPEN SQL).

After entering any necessary selection criteria, the BW query is executed and the results are displayed.

After reviewing the results, press either the Back or Exit buttons at the top of the screen. If the Display Statistics Data option is selected, then the next window contains the runtime statistics. Table 28-3 describes the key columns that can point to the source of the problem.

Table 28-3 Display statistics data column descriptions

Column heading	Description
QAGGRUSED	Name of the aggregate used.
QDBSEL	Number of records that had to be read on the database.
QDBTRANS	Number of records the database transferred to the server.
QTIMEOLAP	Time that the OLAP processor required.
QTIMEDB	Time the database + network required to select the transaction data. The time starts when the SQL command is sent to the database and ends when the requested data reaches the application server. Therefore, the time includes the transport in the network and the selection on the database.
QTIMECLIENT	Time spent formatting the data on the frontend. The time measurement starts when the OLAP processor has completed the data formatting and sent the data to the frontend and ends when the front end has inserted all of the data in the worksheet. The time therefore includes the transport in the network with the RFC connection and the output in the frontend.

Note: SAP note 130696 contains a description of all the columns.

The statistics data can point to the following problems:

- Network problems. This can show up with unusually high times in the QTIMEDB or QTIMECLIENT columns.
- Aggregate not used as expected or wrong aggregate selected. Verify the correct aggregate is in column QAGGRUSED. If different aggregate or no aggregate is selected, then review the query to see if it is defined as expected. If the query is correct, then verify that the preferred aggregate is correctly defined and that it actually runs better by forcing the aggregate selection (Select Aggregate option). Open a customer problem message with SAP if forcing the aggregate provides better performance.
- Aggregate is needed. If the number of records select (QDBSEL) is high but the number of records returned (QDBTRANS) is low, then this is a good candidate for an aggregate.

- Poor access plan. Time spent in database (QTIMEDB) is high and there are no apparent network issues. Verify the access plan using SQL Trace (ST05) or Visual Explain in iSeries Navigator.

28.2.2 Visual Explain

The Visual Explain (VE) feature within iSeries Navigator is an easy to use, graphical tool for viewing an access plan. A lot more information besides the access plan is available through the tool, such as a record of the estimates used, the amount of parallel processing used, and the index and statistics advisors.

To view a query using VE, open iSeries Navigator, expand the Database section and click Schemas. A list of database tasks should appear in the bottom right panel. Double-click the Run an SQL script task.

This opens a new window for entering SQL statements. Perform the following steps to simulate the SAP system environment:

1. Select Connection → JDBC Setup
2. Click the Package tab.
3. Select **Enable extended dynamic (package) support**.
4. Enter any name in the Package filed.
5. Enter a valid library name in the the Package library field.
6. Select **OK**.
7. If this is an SAP BW query and EVI Stage 2 support is being used, then continue with step 8. If EVI Stage 2 support was not enabled through the instance profile parameters, then continue with step 9.
8. Enter the following statement on the first line.

   ```
   cl: chgqrya qryoptlib(<alternate qaqqini library>)
   ```

   The library to specify is the same name used in the dbs/db4/alternate_qaqqinilib profile parameter.
9. Enter the following statement on the second line.

   ```
   cl: addlible r3<sid>data;
   ```

   Replace <sid> with the SAP system identifier (SID) of the BW system.
10. Select Run → All. The first statement should run successfully every time. The second statement can fail if the library already exists in the library list.
11. Enter the SQL statement. The query text can be retrieved from either RSRT, SQL Trace (ST05) or typed in manually.

> **Note:** SAP BW can use temporary tables within a query. These temporary tables are used when additional processing, such as hierarchies, is required before running the BW statement. These temporary tables only exist while the BW query is running, which means they cannot exist when attempting to run the statement in Visual Explain.
>
> A possible recreation of these tables can be attempted by locating the creation of the table or view in an ST05 trace. 28.5, “SAP BW naming convention for database tables” on page 617 lists the names used for temporary tables.

12. Select VisualExplain → Explain or VisualExplain → Run and Explain. VisualExplain → Explain provides the access plan and then exits the query, so the actual implementation

does not run to completion. VisualExplain → Run and Explain runs the query to completion and returns the results in a separate window. See Figure 28-9.

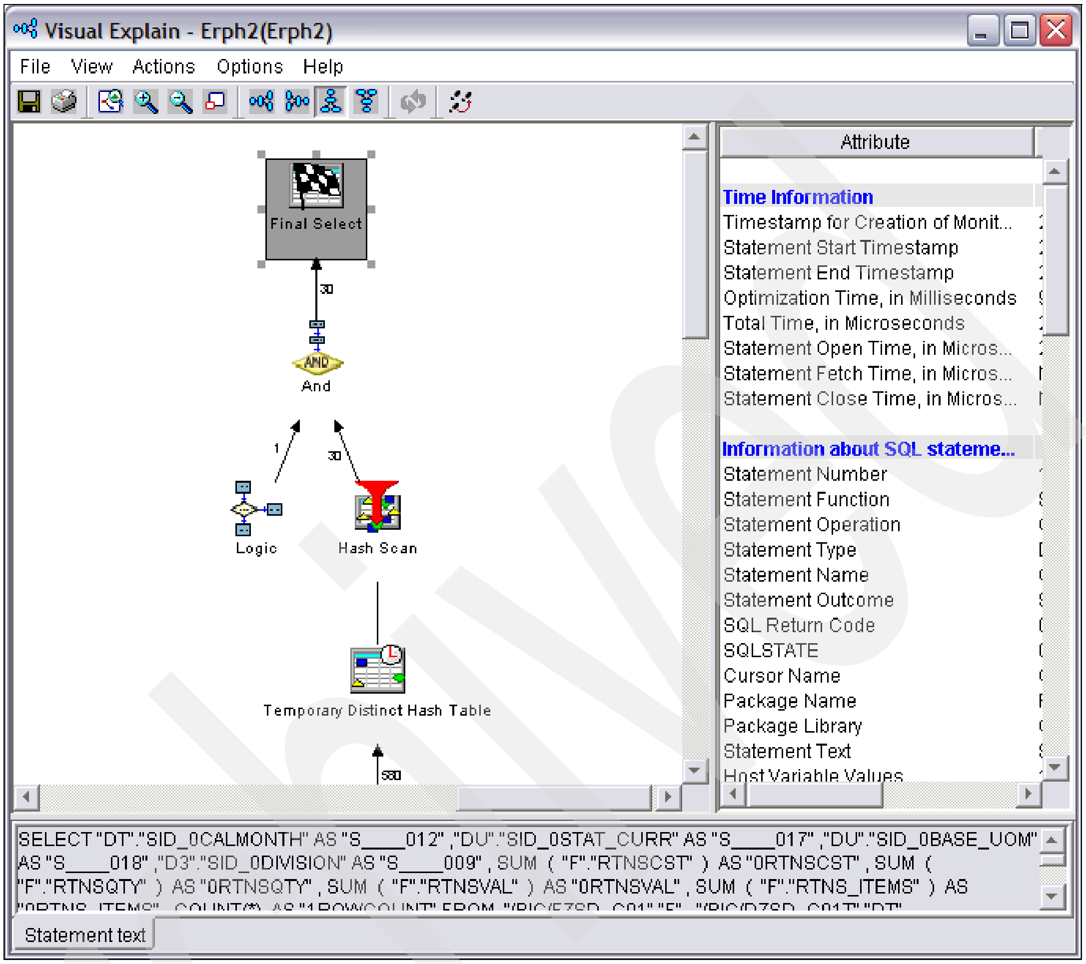

Figure 28-9 Visual Explain example

From this window you can see the query implementation in the top left pane, information about the selected node of the query implementation in the top right pane and any optimizer messages in the bottom pane. To see if LPG was used in the query select View → Highlight LPG. The nodes involved in LPG is now highlighted. See Figure 28-7 on page 605 for an example. To see the estimated record counts, estimated processing time or the degree of parallelism used, select View → Arrow Labels and the appropriate option.

The VE output can indicate several problems:

1. Incorrect join order. View the estimated record counts to see if these are reasonable. If they are incorrect, then additional indexes or statistics may be needed. Select Actions → Advisor to see if any additional indexes or statistics were requested by the optimizer.

> **Note:** Indexes can be created through the advisor pop-up window. However, it is recommended to create all new indexes within the SAP BW application.

2. Poor grouping strategy. If the estimated number of groups is quite different than the real number of rows returned, then additional indexes or statistics may be needed. If all the grouping columns are from the same table, then create a binary radix index over those

columns. If they belong to different tables, then create statistics over each individual column of the group by columns. Also, statistics over the join columns help detect any skew of the data that can lead to a bad grouping estimate. Statistics for a table can be viewed in VE by selecting the node associated with the table and selecting Action → Statistic Data.

3. Aggregate or Materialized Query Table (MQT) needed. If a large number of rows are estimated to be read from the fact table and very few rows are returned in the final result, then an aggregate or MQT can improve the performance.

   An aggregate is a structure that can be defined for an InfoCube in SAP BW that summarizes the data using a subset of the characteristics of the InfoCube. The structure used for an aggregate table is an extended star schema containing the summarized data. An MQT is a structure that can be defined at the database level outside of SAP and uses any number of the characteristics defined in the InfoCube or ODS to summarize the data.

   The structure of an MQT is a single database table containing the results of the defined summary table. Both aggregates and MQTs need to be maintained as part of the normal data load process. Neither aggregates nor MQTs need to be known by the end user to achieve better performance, as either the OLAP processor in SAP BW decides to use an aggregate in place of the base InfoCube tables in the generated query or the database optimizer decides to rewrite the query to use an MQT in place of all or part of the query submitted by SAP BW.

   Note: SAP note 849720 contains additional information concerning the usage of MQTs.

28.3 Data load performance

This section provides tips for improving data load performance.

28.3.1 Initial InfoCube data loads

When inserting records into InfoCubes, each dimension identifier (DIMID) is verified against the dimension tables. If the DIMID does not exist, then a new DIMID is created and added to the dimension table. Since the dimension tables are empty when the query is started, the reading of these tables are performed by a table scan. As these dimension tables grow, the table scan starts to become more expensive. This section discusses the two methods for avoiding the penalty of the table scan.

Note: SAP note 425593 discusses initial data load performance.

Method 1: Small initial load

The first method is to start with a small initial load of up to 20,000 records. This small load runs rather quickly and provides a good start to the building of the dimension tables. After the load is complete, delete the SQL packages as described in SAP note 63037. This allows the optimizer to choose a new access method for reading the dimension tables. Since the dimension tables are not empty, the optimizer more than likely chooses an index access method. Reading the dimension tables with an index allows for a constant data load rate rather than the gradual degradation that the table scan provides.

Method 2: Initialize dimension tables

The second method involves initializing the dimension tables to a predetermined size. The space allocated appears as deleted records. These records do not show up as records in the database when queried, but the optimizer takes them into account and can decide on index access rather than table scan when reading the table. The dimension tables are defined with the reuse deleted records option, so that the allocated space is used first when adding new records.

Existing dimension tables can be initialized by running the SAP_TABLE_INZPFM_DB4 report. This report provides three options:

1. All Dimension Tables. This option increases the size of all the dimension tables of existing InfoCubes. This is not recommended since this can allocate space to tables that do not increase and leads to creating wasted space.
2. InfoCube Dimension Table. This is the default option and increases the dimension table of the InfoCube specified in the InfoCube Name field.
3. Any Table. This option increases the size of the table specified in the Table Name field by the amount of records specified in the Number of Records field. This option can be used to initialize any table, since the table name specified does not need to be a dimension table

New dimension tables can be initialized automatically at creation time by setting the RSADMIN parameter DB4_AUTO_INZPFM. Use the following steps to maintain the parameter:

1. Start the SAP_RSADMIN_MAINTAIN report.
2. Enter DB4_AUTO_INZPFM in the OBJECT field.
3. Enter an 'X' in the VALUE field to automatically initialize the dimension table or ' ' (blank) to deactivate the automatic initialization of the dimension tables.
4. Select the appropriate action (INSERT, UPDATE or DELETE).
5. Select Program → Execute (F8).

Note: If you change the value of the DB4_AUTO_INZPFM parameter and the behavior does not change, then the RSADMIN table is buffered. Transaction /$tab can be used to reset the buffers.

The small initial load (Method 1) is the recommended method because it is more likely to avoid the table scan degradation. Changes in the optimizer behavior can lead to the failure of index access being selected in method 2.

28.3.2 Number range buffering

When loading data in an InfoCube, it is important to increase the number range buffers for all dimensions and InfoObjects that have a high number of records.

Increasing number range buffers for dimensions

The following steps identify the number range for each dimension and how to increase the buffer size.

1. Use function module RSD_CUBE_GET to find the object name of the dimension.
2. Fill in the I_INFOCUBE field with the name of the InfoCube, the I_OBJVERS field with the character 'A' and I_BYPASS_BUFFER with the character 'X'.
3. Select Function modules → Execute (F8).
4. Double-click the table E_T_DIME.
5. Column NOBJECT contains the name of the number range for each dimension.
6. Run transaction SNRO to change the number range buffer.

7. Enter the value from step 5 in the Object field and select Number range object → Change (F2).
8. To turn on number range buffering, select Edit → Set-up buffering → Main memory.
9. Enter a value in the No. of numbers in the buffer field. This value should be set to the expected number of new values in the initial and delta loads.
10. Select Number range object → Save.

> **Attention:** Do not buffer the number range for the package dimension.

Increasing number range buffers for InfoObjects

The following steps identify the number range for each dimension and how to increase the buffer size.

1. Use function module RSD_IOBJ_GET to find the object name of the dimension.
2. Fill in the I_IOBJNM field with the name of the InfoObject, the I_OBJVERS field with the letter ‘A’ and I_BYPASS_BUFFER with the letter ‘X’.
3. Select Function modules → Execute (F8).
4. Double-click the table E_S_VIOBJ.
5. Column NUMBRANR contains the name of the number range for each dimension. Append BIM to the front of the value in NUMBRANR to get the number range object.
6. Run transaction SNRO to change the number range buffer.
7. Enter the value from step 5,example BIM0000312, in the Object field and select Number range object → Change (F2).
8. To turn on number range buffering, select Edit → Set-up buffering → Main memory.
9. Enter a value in the No. of numbers in the buffer field. This value should be set to the expected number of new values in the initial and delta loads.
10. Select Number range object → Save.

> **Important:** Do not buffer the number range for InfoObject 0REQUEST.

28.3.3 Index maintenance

When starting a data load, it is important to understand what indexes exist over either the InfoCube or ODS tables. All the defined indexes are defined with immediate access path maintenance. This means that every change to a row in the table forces maintenance to be performed on all of the indexes defined for that table. This means that index maintenance during a data load can carry a significant performance cost.

There are no hard rules for when to keep indexes on during a data load and when to drop an index before the data load and rebuild the index after it completes. Since every cube is different, this needs to be done on a trial and error basis for each set of tables. What is known is that EVIs carry less of a penalty than binary radix indexes. Test shows EVIs can be left on during any data load, initial or delta. Binary radix indexes are where most of the tests need to be done, because the more binary radix indexes that exist, the higher the chance that dropping and rebuilding the index is the best method. Experience shows that all binary radix indexes should be dropped during initial data loads and for delta data loads if the number of new rows is greater than 10% of the total number of rows in the table.

Note: SAP note 935299 discuss naming conventions to use on user created indexes that allow for the automatic dropping and rebuilding of indexes during data loads.

Building indexes in parallel using SMP

If you have multiple CPUs available for use by SAP BW and have installed DB2 Symmetric Multiprocessing (SMP) on the database server, then the time to create indexes can be reduced dramatically. Tests show that the time to build an index with SMP is inversely proportional to the number of CPUs available, provided the necessary resources (CPU, disk and memory) are available. For example, the time to build an index on a four CPU system with SMP is done in one fourth of the time on a non-SMP system.

Parameter DB4_PARALLEL_INDEX_BUILD in RSADMIN is used to determine what level of SMP to use for index creation. If this parameter is maintained, then the work process follows the normal system settings until an index build is requested. It then changes the parallel degree for that work process to the value in DB4_PARALLEL_INDEX_BUILD for the index request. Once the index is created, the work process is changed back to the normal system settings.

Note: SAP note 601110 discusses building indexes with SMP.

To maintain the DB4_PARALLEL_INDEX_BUILD parameter:

1. Start report SAP_RSADMIN_MAINTAIN
2. Set the Object field to DB4_PARALLEL_INDEX_BUILD
3. Set the Value field to one of the following values:
   - *OPTIMIZE - the database determines how much resource can be consumed
   - *NBRTASKS N - the database is allowed up to N tasks
4. Select INSERT, UPDATE or DELETE
5. Select Execute (F8)

If table RSADMIN is buffered, a reset of the buffers using transaction /$tab can be required before the change goes into effect.

28.3.4 ODS activation settings

The activation of ODS tables can be a time consuming process. The following are settings that can reduce the activation times.

ODS object settings

Two settings in the definition of the ODS object can speed up the activation process. The BEx Reporting flag is used to indicate whether users can query this ODS object directly. If this flag is set, then the values in the key fields are validated against the SID tables of the corresponding InfoObjects during the activation. This validation step can consume up to half the total activation time. Leave this flag off if the ODS object is only used as input to other ODS objects or InfoCubes.

The second flag is the Unique Data Records flag. If this flag is checked, then the ODS object assumes that all input records satisfy the primary key constraint. This assumption saves time by allowing the activation to directly insert records into the active table and skip the verification step that ensures the record does not already exist in the active table. When this flag is turned on and a record from the activation queue has the same values for the key fields as an existing record in the active table, then the ODS activation fails and all previous records inserted during that activation need to be rolled back, or removed from the active table. Be careful when using this flag as roll backs can be very time consuming.

Global settings

Two settings in the global ODS settings can help improve the ODS activation time. To get to the customizing screen, run RSA1 → Settings → Global Settings → ODS Object.

The first field contains the number of parallel processes to use during the activation. If there are additional CPUs and dialog work processes available during the activation, then set the number of parallel processes to the smaller value of the number of available CPUs, the number of available dialog work processes, or the value of 6. The validation of surrogate identifiers (SIDs) step for when the BEx Reporting flag is turned on and the active table insert/update step can take advantage of the parallel processes.

The second field contains the minimum number of data records per data package. Optimal performance is achieved by keeping the number of data packages under 1000. The number of data packages can be calculated with the following formula:

(number of records to be activated * 2) / minimum number of data records

The default value for the minimum number of records is 10,000.

28.4 SAP notes for SAP BW

Table 28-4 lists the important SAP notes for SAP BW performance and configuration.

Table 28-4 Primary SAP notes related to SAP BW on System i configurations

SAP note number	Title	Description
307077	iSeries: Performance Optimization for BW Systems	Summarizes the current settings that are specific to System i configurations.
501572	iSeries: EVI stage 2 support	Describes the requirements and steps for implementing star join support.

Table 28-5 lists other SAP notes that contain either general performance tips for SAP BW or answers to specific SAP BW performance problems that have been encountered previously.

Table 28-5 Other SAP notes related to SAP BW on System i configurations

Note number	Title	Description
130243	General tips for uploading transaction data to BW	Ten tips for efficient mass data uploads.
130696	Performance trace in BW	Contains the definition of the columns in the BW Statistics table (RSDDSTAT), such as QTIMEOLAP and QTIMEDB.
425593	iSeries: Initial data load takes a long time	Methods for improving performance of initial data loads.
485420	iSeries: Database hints for Open SQL/Native SQL	Describes the hints that can be used in queries.

Note number	Title	Description
541508	iSeries: Checking the system parameters for BW	Describes how to run a sanity check for various System i model specific SAP BW setting.
557870	FAQ: BW Query Performance	Provides answers to general SAP BW performance questions.
562224	iSeries: Performance Master Data Loading	Tips for loading master data.
567745	Composite note BW 3.x performance: DB-specific settings	Lists all SAP notes related to DB-specific settings on all platforms.
567746	Composite note BW 3.x performance: Query & Web Applications	Lists all SAP notes related to query and Web application performance on all platforms.
567747	Composite note BW 3.x performance: Extraction and loading	Lists all SAP notes related to extracting and loading data on all platforms.
568139	iSeries: Performance Aggregate Filling and Complex Queries	Describes possible performance problem while loading aggregates or running queries that filter output for special attributes.
601109	iSeries: Performance Improvement During Compression	Tips for improving the InfoCube compression process.
601110	iSeries: Setting up indexes in parallel	Steps required for taking advantage of SMP for index builds.
795106	DB2/400 Enhanced Nested Table Expression® Support	Possible performance enhancement for queries that use views, such as SAP BW queries involving hierarchies.
820325	iSeries: Using QAQQINI with SAP	Describes the use and recommendations for QAQQINI file.
849720	iSeries: Using Materialized Query Tables (MQT)	Describes MQTs and how MQTs can be used in SAP BW.
892513	Consulting: Performance: Loading data, number of package, required size	Recommendations for the maximum number of data packages per data load request and the maximum number of records per data package.
935299	iSeries: Non Standard Indexes on InfoCube Fact Tables	Describes how to automatically drop and rebuild user created indexes during data loads.

28.5 SAP BW naming convention for database tables

Table 28-6 on page 618 lists the various naming conventions used for SAP BW tables.

Table 28-6 SAP BW table naming convention

Starting characters of SAP table	Description
/BI0/...	SAP name space used for objects delivered with SAP BW Business Content.
/BIC/...	Customer name space used for objects created by the customer.
/BI*/F<InfoCube>	InfoCube F-fact table
/BI*/E<InfoCube>	InfoCube E-fact table
/BI*/V<InfoCube>F	InfoCube fact view over both fact tables
/BI*/B...	Persistent Storage Area (PSA) tables
/BI*/A<ODS>00	ODS active table
/BI*/A<ODS>10	ODS activation queue table
/BI*/A<ODS>20	ODS change log table
/BI0/01...	temporary tables for query execution
/BI0/02...	buffered hierarchy tables for query execution
/BI0/03...	temporary views for query execution
/BI0/04...	temporary stored procedures for query execution
/BI0/05...	temporary triggers for query execution
/BI0/06...	temporary SID-selection tables for query execution
/BI0/0T...	temporary tables for query execution
/BI0/0D...	temporary tables for query execution
/BI0/0V...	temporary views for query execution

Note: Some of these BW objects or names are only used in specific BW releases.

28.6 Additional DB2 information

The following IBM white papers provide additional database information:

- *The creation and use of materialized query tables within IBM DB2 UDB for iSeries* at:

  `http://www-03.ibm.com/servers/enable/site/education/wp/438a/438a.pdf`

- *Index and statistics strategies for DB2 UDB for iSeries* at:

  `http://www-03.ibm.com/servers/enable/site/education/abstracts/indxng_abs.html`

- *Star Schema Join Support within DB2 UDB for iSeries* at:

  `http://www-03.ibm.com/servers/enable/site/education/abstracts/16fa_abs.html`

Part 5

Data exchange between applications

Part 5 describes techniques to enable mySAP Business Suite systems that are running on the System i server to interact with System i non-SAP applications (for example, RPG programs). The objective is to:

- Exchange data
- Call or trigger programs and system functions:
  - The SAP system calls programs outside SAP applications or i5/OS commands.
  - i5/OS commands start SAP reports or jobs in the SAP system.

29

Data exchange scenarios and examples

The examples in this chapter are based on scenarios using programs written in Advanced Business Application Programming (ABAP) and RPG/400®, COBOL/400®, or Java. These examples demonstrate:

- Using an ABAP program to access non-SAP data
- Java Servlets accessing non-SAP data
- Accessing SAP data using an RPG/400 program
- i5/OS jobs starting SAP applications
- ABAP applications calling i5/OS commands
- The Java connector
- Remote Function Call (RFC) program communications between ABAP programs and Integrated Language Environment (ILE) RPG/400, ILE COBOL/400 or Java.

Note: All program examples shown in this book are created and tested with i5/OS V5R3 and SAP ECC with Web Application Server 6.40.

All program source codes (ABAP, RPG/400, COBOL/400, Java, and CL) provided in this chapter are for demonstration purposes only. They are not intended for use in any production system. Use this program code as a basis to create the required programs according to your specific requirements. Adjust program names and library names to fit your particular environment.

For programming examples and further explanations (based on OS/400 V4R5 and SAP release 4.6C) for the following topics, refer to *Implementing SAP R/3 on OS/400*, SG24-4672:

- Data porting
- Using the Common Programming Interface for Communications (CPI-C) connection
- MQSeries® link for R/3

- SAP R/3 and Domino connection
- Access Builder for SAP R/3
- For SAP Extended Binary Coded Decimal Interchange Code (EBCDIC) environments:
  - Using an RPG/400 program to access SAP R/3 data
  - Using an ABAP program to access non-SAP R/3 data

29.1 Considerations

SAP applications on System i servers operate in their own environment. All SAP objects are stored in specific libraries on the system, and SAP jobs run in specific subsystems, using their own work management objects (for example, subsystems and job descriptions). SAP applications can run concurrently with other applications on the same System i server or on different servers connected together.

Figure 29-1 on page 623 shows an integrated system, where SAP and SAP note business applications are running under the same i5/OS environment. The SAP system (with a system ID of *ERP*, an instance of *03*, and the SAP database library *R3ERPDATA*) and a non-SAP application being placed in a library named *RFC*. They run concurrently on the same System i server.

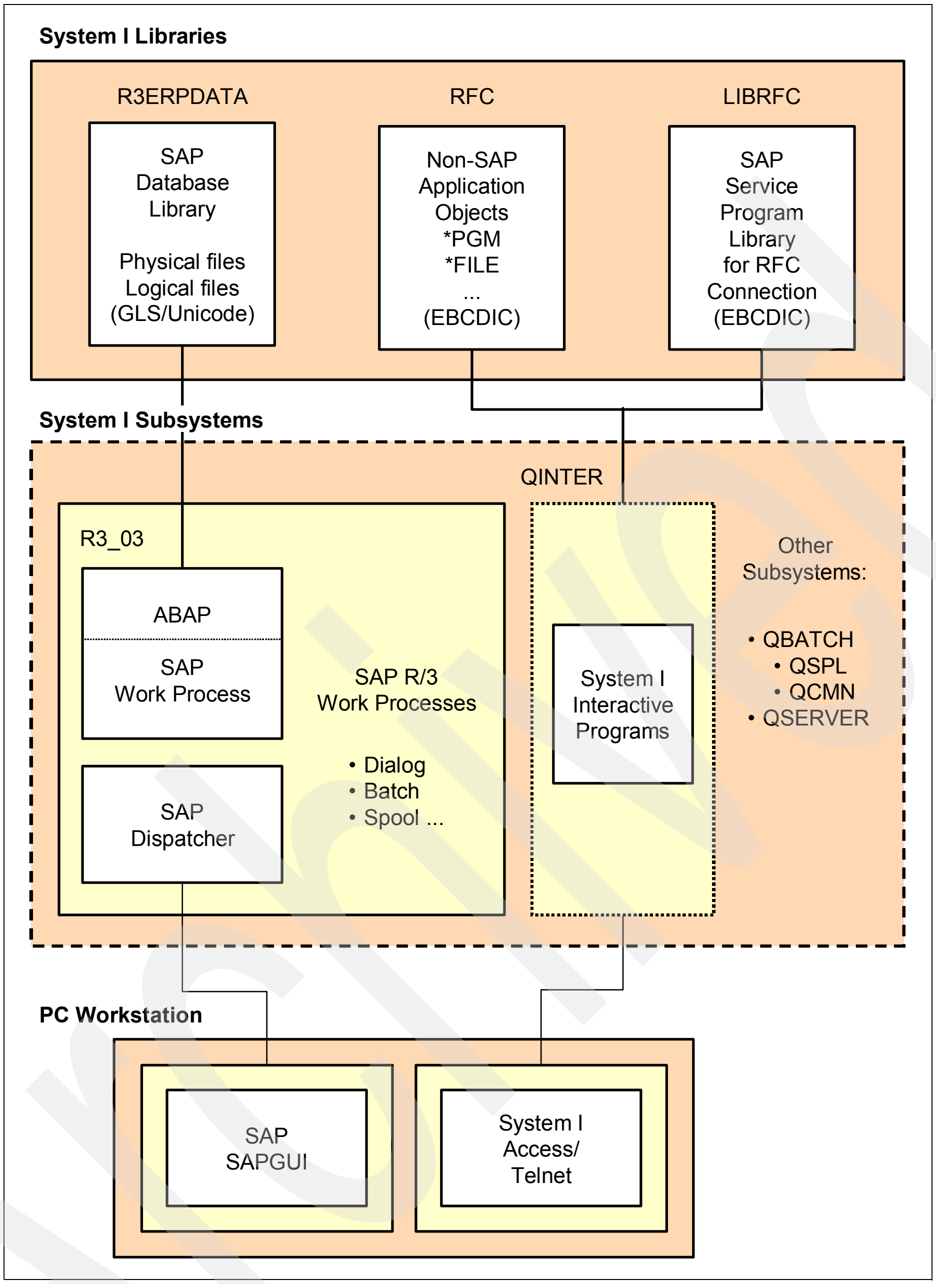

Figure 29-1 SAP system and other applications

All SAP work processes run in the R3_03 subsystem. For RFC connections to SAP systems, the non-SAP application uses the functions that are provided by the SAP service-program library *LIBRFC*. It runs in the default System i subsystem environment (QINTER and QBATCH).

Note: The SAP system and the non-SAP applications use different code pages: EBCDIC for non-SAP applications and a Global Language System (GLS) or Unicode implementation for SAP applications. Refer to Chapter 6, "Encoding data used in SAP systems" on page 35 for more information about an SAP GLS implementation.

Use parallel work station sessions to access the system, as follows:

- SAP graphical user interface (GUI) connects to SAP instance 03.
- 5250 emulation (for example, iSeries Access or Telnet) connects to the QINTER subsystem.

When using a 5250 emulation, it is important that the CCSID setting of your emulation session and the CCSID setting of the interactive job on System i match. With IBM Personal Communications or the 5250 emulation of iSeries Access for Windows, you can configure the Host Code Page through the Session Parameters in the Customize Communication screen. The default depends on the language version of the emulation program.

For System i implementations, your interactive job must use the same value for the job CCSID to get proper character conversion. The job CCSID can be controlled through the CCSID user profile parameter, the QCCSID system value, or with the CCSID parameter on the CHGJOB command. The shipped defaults are 65535 for the QCCSID system value and CCSID(*SYSVAL) for the CRTUSRPRF command, so that no character conversion takes place and character data can be displayed incorrectly.

Consider the following points when planning to interface non-SAP applications with SAP applications or exchange data between these systems:

- SAP applications are developed in ABAP or Java. These applications can only run in an SAP system environment.

  SAP applications do not run directly in an i5/OS work management environment. They are dispatched to an instance work process (batch or dialog). From the i5/OS call level interface, You can only start the execution of ABAP programs by using the STRREPORT command that is provided in the SAP kernel library. The ABAP program is then executed in an SAP work process. It is also possible to execute i5/OS commands from within ABAP programs by using an operating system call.

  SAP applications provide their own interface and tools for communication outside of the SAP system environment. These include:

  – Operating system command call
  – Event handler
  – RFC interface
  – Java connection

- SAP applications define and maintain their 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 System i database library using the American National Standards Institute (ANSI) Structured Query Language (SQL) industry standard database interface of DB2 UDB for iSeries.

  You can access SAP dictionary transparent tables from outside an SAP system if you are familiar with the SAP table structures. However, we do not recommend this because table structures can change in different SAP releases. Also, the ABAP and Java language provide functions to access tables and stream files outside of SAP applications. These functions include:

  – The EXEC-SQL interface to execute SQL host commands
  – Commands for reading from and writing to files and stream files

  > **Note:** To ensure data consistency, do not modify any SAP dictionary tables from outside the SAP system.

29.2 Example programs

The following examples show programs written in RPG/400 and ABAP. These programs read a data description specifications (DDS)-described DB2 UDB iSeries physical file and an ABAP dictionary table and display the contents of these files.

29.2.1 RPG/400 example

The RPG program named T8189RP1 (shown in Figure 29-2) reads the externally described physical file named T8189DB and prints the file contents.

```
 FMT FX .....FFilenameIPEAF........L..I........Device+......KExit++Entry+A....U
        *************** Beginning of data *************************************
             FT8189DB IF  E                    DISK
             FQSYSPRT O   F     132            PRINTER
             C           *IN20     DOUEQ'1'
             C                     READ Z8189DB                  20
             C           *IN20     IFEQ '0'
             C                     EXCPTDETAIL
             C                     ENDIF
             C                     ENDDO
             C                     MOVE '1'       *INLR
             OQSYSPRT E 11             DETAIL
             O                         ITEMM
             O                         ITEMD
        ****************** End of data ****************************************
```

Figure 29-2 RPG program T8189RP1

The file discretion for the DDS-described file T8189DB is shown in Figure 29-3 on page 626. Both objects are stored in the user library named RFC.

```
                            Display File Field Description
  Input parameters
    File  . . . . . . . . . . . . . . . . . . . :  T8189DB
      Library . . . . . . . . . . . . . . . . . :  *LIBL
  File Information
    File  . . . . . . . . . . . . . . . . . . . :  T8189DB
      Library . . . . . . . . . . . . . . . . . :  RFC
    File location . . . . . . . . . . . . . . . :  *LCL
    Externally described  . . . . . . . . . . . :  Yes
    Number of record formats  . . . . . . . . . :      1
    Type of file  . . . . . . . . . . . . . . . :  Physical
    File creation date  . . . . . . . . . . . . :  11/09/05
    Text 'description'. . . . . . . . . . . . . :  Non-SAP File
  Record Format Information
    Record format . . . . . . . . . . . . . . . :  Z8189DB
    Format level identifier . . . . . . . . . . :  2DD1DEE277EBD
Number of fields  . . . . . . . . . . . . . . :    2
    Record length . . . . . . . . . . . . . . . :     30
  Field Level Information
             Data       Field  Buffer    Buffer        Field    Column
    Field      Type      Length  Length  Position        Usage    Heading
    ITEMM      CHAR           5       5         1        Both     ITEMM
      Coded Character Set Identifier  . . . . . :     37
    ITEMD      CHAR          25      25         6        Both     ITEMD
      Coded Character Set Identifier  . . . . . :     37
```

Figure 29-3 DSPFFD: File T8189DB

File T8189DB is created with the i5/OS Create Physical File command (CRTPF) using a Coded Character Set-Identifier (CCSID) value of 37, which is a single-byte EBCDIC code page. The example program shown in Figure 29-2 on page 625 reads this T8189DB file and prints a list.

29.2.2 ABAP example

In this example, the Z8189AB1 ABAP program reads the Z8189DB SAP table sequentially and prints the file contents. See Example 29-1.

Example 29-1 ABAP program Z8189AB1

```
REPORT  Z8189AB1.
TABLES: Z8189DB.
SELECT * FROM Z8189DB.
WRITE: / Z8189DB-ZITEMM, Z8189DB-ZITEMD.
ENDSELECT.
```

Table Z8189DB is defined as a transparent table in the ABAP dictionary. A transparent table is stored in the assigned System i 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 System i library without using an ABAP program. Figure 29-4 on page 627 shows the ABAP dictionary table named Z8189DB.

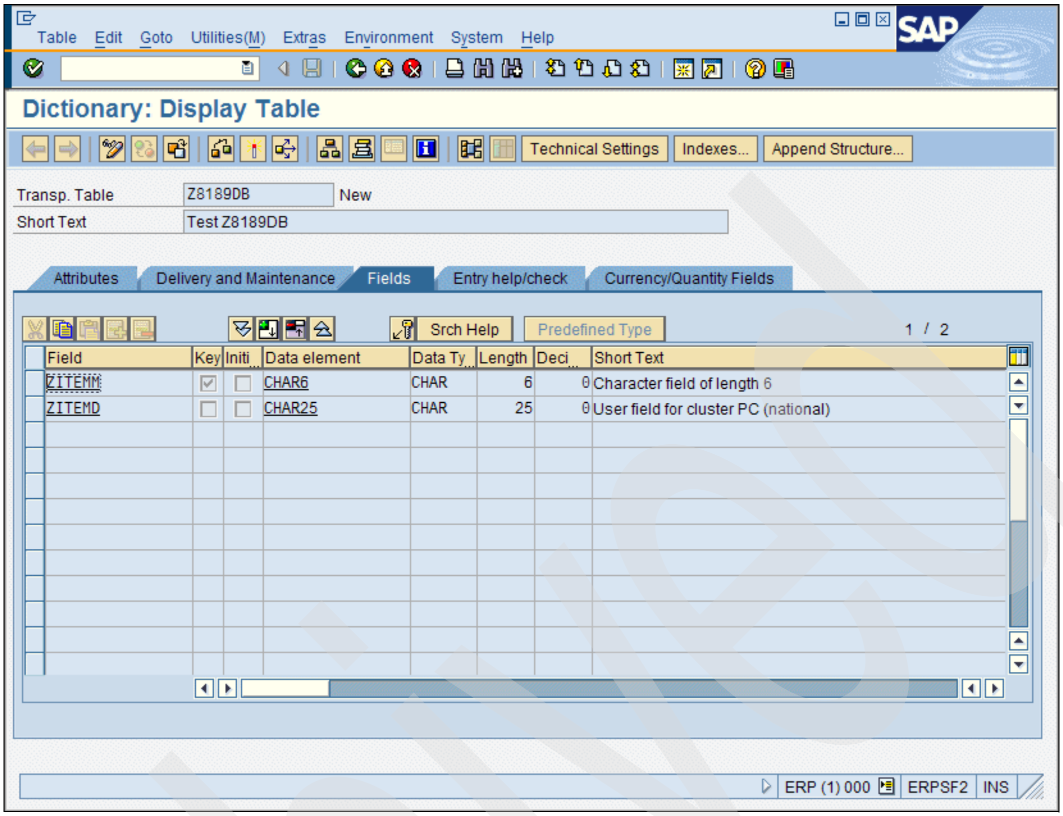

Figure 29-4 ABAP dictionary table Z8189DB

ABAP program Z8189AB1 reads the dictionary table Z8189DB sequentially using the SAP-SQL SELECT statement. SAP-Open-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 in the SAP database library R3ERPDATA. The resulting output of ABAP program Z8189AB1 is shown in Figure 29-5 on page 628.

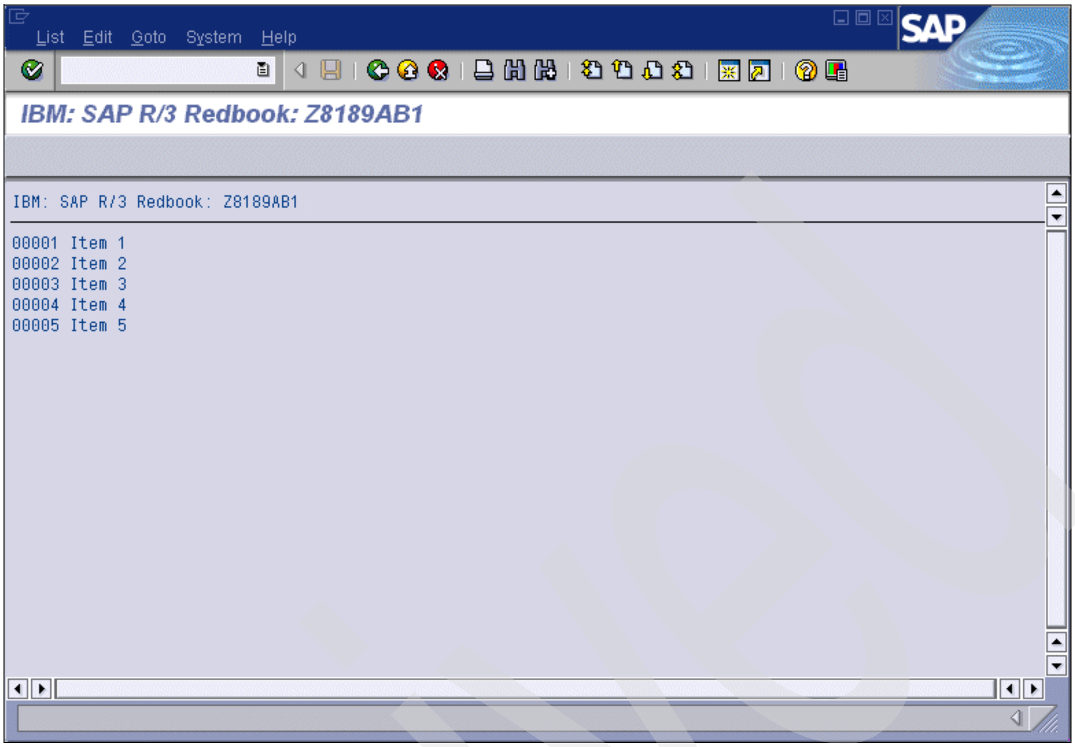

Figure 29-5 Output of ABAP program Z8189AB1

29.3 Accessing SAP systems data using an RPG/400 program

This section describes how you can directly access a table of the SAP data dictionary from outside of the SAP system.

Note: Refer to *Implementing SAP R/3 on OS/400*, SG24-4672 for examples of how to access SAP EBCDIC database tables. Refer to *SAP note 554792* for more information about how to access SAP GLS and Unicode database tables from outside the SAP system.

The example described in this section works with a customer-created database table in the SAP system.

Note: To preserve data consistency, do not update SAP database tables from outside the SAP system. Be aware that the structure of SAP database tables can change, for example, through customizing, support packages, or upgrade to a newer SAP release.

The ABAP dictionary table named Z8189DB is defined as a transparent table in the SAP data dictionary and stored as a physical file in the R3ERPDATA i5/OS library. This is the assigned database library for the SAP system ERP used in this example. See Figure 29-6 on page 629.

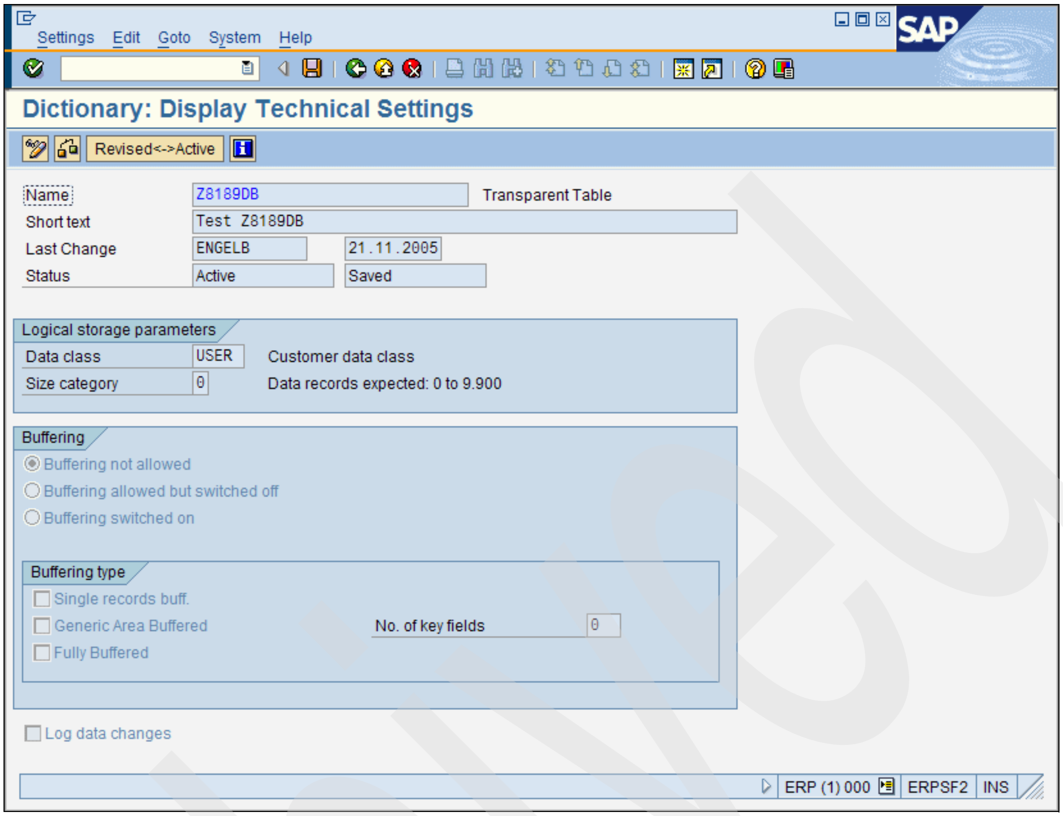

Figure 29-6 Transparent SAP database table Z8189DB

Figure 29-4 on page 627 and Figure 29-6 show the Z8189DB SAP dictionary table as being described in the SAP system environment. Figure 29-7 on page 630 shows the file layout of the associated i5/OS database file named Z8189DB.

```
                            Display File Field Description
 Input parameters
   File  . . . . . . . . . . . . . . . . . . . . :  Z8189DB
     Library . . . . . . . . . . . . . . . . . . :  *LIBL
 File Information
   File  . . . . . . . . . . . . . . . . . . . . :  Z8189DB
     Library . . . . . . . . . . . . . . . . . . :  R3ERPDATA
   File location . . . . . . . . . . . . . . . . :  *LCL
   Externally described  . . . . . . . . . . . . :  Yes
   Number of record formats  . . . . . . . . . . :      1
   Type of file  . . . . . . . . . . . . . . . . :  Physical
   SQL file type . . . . . . . . . . . . . . . . :  TABLE
   File creation date  . . . . . . . . . . . . . :  11/21/05
 Record Format Information
   Record format . . . . . . . . . . . . . . . . :  Z8189DB
   Format level identifier . . . . . . . . . . . :  2C0D5EB3D4F4A
Number of fields  . . . . . . . . . . . . . . :      2
   Record length . . . . . . . . . . . . . . . . :     62
 Field Level Information
              Data       Field  Buffer    Buffer        Field    Column
   Field      Type      Length  Length  Position        Usage    Heading
   ZITEMM     GRAPHIC        6      12         1        Both     ZITEMM
     Default value . . . . . . . . . . . . . . . :
         UX'0020'
     Coded Character Set Identifier  . . . . . . :  13488
     UCS2 or Unicode conversion  . . . . . . . . :  *CONVERT
   ZITEMD     GRAPHIC       25      50        13        Both     ZITEMD
     Default value . . . . . . . . . . . . . . . :
         UX'0020'
     Coded Character Set Identifier  . . . . . . :  13488
     UCS2 or Unicode conversion  . . . . . . . . :  *CONVERT
```

Figure 29-7 DSPFD for file Z8189DB

File T8189DB (Figure 29-3 on page 626) and file Z8189DB (Figure 29-7) appear similar and have the same field structure. However, they are not compatible.

Figure 29-7 shows the file description of the Z8189DB physical file in the associated R3ERPDATA database library, which is created by the SAP database interface. All SAP database tables are created with the Unicode CCSID of 13488. All character fields are represented in a double-byte structure.

In this example, the files are named ZITEMM and ZITEMD. This does not allow you to access this data from a single byte EBCDIC environment without any conversion. Therefore, interfaces without any automatic code page conversion (for example, the i5/OS commands Display Physical file Member (DSPPFM) and Update Data (UPDDTA)) or direct access from an RPG or COBOL program do not run without an interface that is able to convert the data.

An SQL-based environment such as STRSQL, SQLUTIL, or RUNQRY runs for SAP Unicode or GLS Latin-1 tables because the SQL interface automatically performs the code page conversion.

Note: Automatic code page conversion is done only if the job does not run under CCSID 65535. The shipped defaults for the QCCSID system value is 65535, which means that no code page conversion takes place. To enable automatic code page conversion, change this system value or change the job default values with the CHGJOB command into appropriate CCSIDs (for example 500, 37, or 273).

Jobs running under SAP user profiles are automatically routed into CCSID 500.

The Work Query (WRKQRY) command is also supported for these tables. A QRY-2361 warning message is issued if you attach such a table, as shown in Figure 29-8.

```
                           Specify File Selections

 Type choices, press Enter.  Press F9 to specify an additional
   file selection.

   File . . . . . . . . .   Z8189DB        Name, F4 for list
     Library  . . . . . .     R3ERPDATA    Name, *LIBL, F4 for list
   Member . . . . . . . .   *FIRST         Name, *FIRST, F4 for list
   Format . . . . . . . .   Z8189DB        Name, *FIRST, F4 for list













 F3=Exit          F4=Prompt          F5=Report           F9=Add file
 F12=Cancel       F13=Layout         F24=More keys
 File Z8189DB in R3ERPDATA may have DBCS data or text.
```

Figure 29-8 WRKQRY: Work with Unicode tables

You can ignore this warning because the SAP table does not contain DBCS data. It works the same as in an EBCDIC environment.

The next objective is to use this T8189RP1 program to access the Z8189DB SAP database file instead of the original T8189DB EBCDIC file. An interface for the RPG program is necessary to access the Unicode file. SAP applications provide two interfaces (logical files) to access SAP GLS or SAP Unicode tables from outside the SAP system:

- The CRTSAPLF command:
  - Converts SAP Unicode to i5/OS native EBCDIC code
  - Converts SAP GLS Latin-1 to i5/OS native EBCDIC Latin-1
  - Cannot be used for non-Latin-1 or MDMP in an GLS environment
  - Can be used for read and write access to the SAP database table
  - Can be used for key positioning in RPG or Cobol

- The CRTSAPVIEW command:
  – Can be used to convert SAP GLS Latin-X to i5/OS native EBCDIC Latin-X (X= code page different from Latin-1)
  – Can only be used for read access to the SAP database table
  – Can only be used for sequential access in RPG or Cobol

Figure 29-9 illustrates both the interfaces necessary to access SAP database tables from outside of an SAP system.

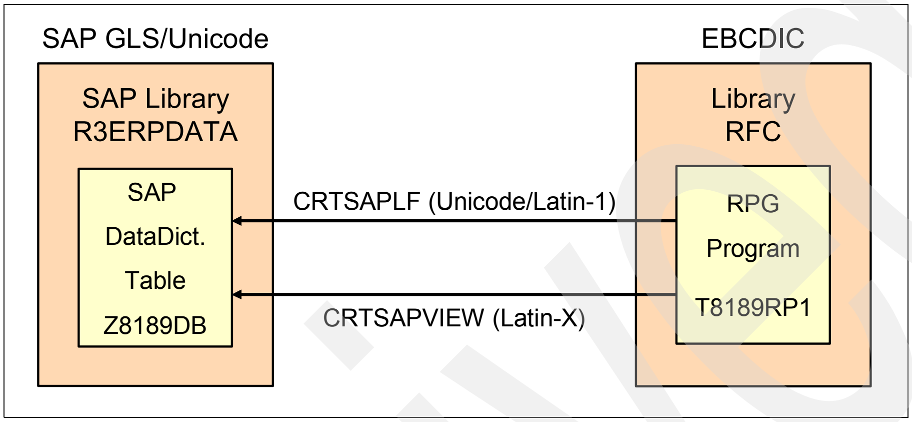

Figure 29-9 Accessing SAP dictionary tables from outside the SAP system

Note: Only characters that can be mapped between both sides for data exchange between SAP Unicode and i5/OS EBCDIC can be exchanged.

There is no support to access SAP dictionary tables from outside an SAP system for the following scenarios:

- SAP GLS MDMP to i5/OS EBCDIC (multiple code page support for SAP GLS)
- SAP GLS Latin-X to i5/OS EBXDIC Latin-Y (different code pages)

Use other SAP-supported interfaces for these scenarios, for example, individually coded data exchange between SAP systems and i5/OS programs based on RFC connection.

29.3.1 CRTSAPLF: Accessing SAP Unicode or GLS Latin-1 code page data

Use the Create SAP Logical File (CRTSAPLF) command to access SAP database tables with Unicode or ASCII Latin-1 code page from a native i5/OS environment. The CRTSAPLF command is available in a zip attachment in *SAP note 554792*. This zip file contains a savefile with a library name of CRTSAPLF. Download it from the SAP note, decompress the zip file and transfer it via File Transfer Protocol (FTP) binary onto your System i server. See Figure 29-10 on page 633 for an illustration.

```
                    Create LF on SAP Unicode DB (CRTSAPLF)

Type choices, press Enter.

SAP System ID  . . . . . . . . . > ERP          Name
Table Name . . . . . . . . . . . > Z8189DB      Name
LF Name in R3SIDLF . . . . . . .   *TABLE       Name, *TABLE
ASP for library R3SIDLF  . . . .   1            Number
CCSID of the fields  . . . . . .   00500        Number
Option key field 1 . . . . . . . > ZITEMM       Name
Option key field 2 . . . . . . .                Name
Option key field 3 . . . . . . .                Name
Option key field 4 . . . . . . .                Name
Option key field 5 . . . . . . .                Name
Option key field 6 . . . . . . .                Name
Option key field 7 . . . . . . .                Name
Option key field 8 . . . . . . .                Name
Option key field 9 . . . . . . .                Name
Option key field10 . . . . . . .                Name

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

Figure 29-10 CRTSAPLF: Create Logical File for SAP Latin-1 or Unicode data

CRTSAPLF creates a logical file named Z8189DB in the R3ERPLF library. It points to the Z8189DB table in the R3ERPDATA SAP library, as shown in Figure 29-11 on page 634.

```
   File  . . . . . . . . . . . . . . . . . . . : FILE       Z8189DB
     Library . . . . . . . . . . . . . . . . . :            R3ERPLF
   Type of information . . . . . . . . . . . . : TYPE       *ALL
   File attributes . . . . . . . . . . . . . . : FILEATR    *ALL
   System  . . . . . . . . . . . . . . . . . . : SYSTEM     *LCL
 File Description Header
   File  . . . . . . . . . . . . . . . . . . . : FILE       Z8189DB
   Library . . . . . . . . . . . . . . . . . . :            R3ERPLF
   Type of file  . . . . . . . . . . . . . . . :            Logical
   File type . . . . . . . . . . . . . . . . . : FILETYPE   *DATA
   Auxiliary storage pool ID . . . . . . . . . :            00001
 Data Base File Attributes
   Externally described file . . . . . . . . . :            Yes
   File level identifier . . . . . . . . . . . :            1051122092942
Creation date . . . . . . . . . . . . . . . . :          11/22/05
   Text 'description'  . . . . . . . . . . . . : TEXT       Z8189DB
Access path journaled . . . . . . . . . . . . :          No
  Access path . . . . . . . . . . . . . . . . . :          Keyed
Number of key fields  . . . . . . . . . . . . :            1
   Record format . . . . . . . . . . . . . . . :            Z8189DB
     Key field . . . . . . . . . . . . . . . . :            ZITEMM
       Sequence  . . . . . . . . . . . . . . . :            Ascending
       Sign specified  . . . . . . . . . . . . :            UNSIGNED
       Zone/digit specified  . . . . . . . . . :            *NONE
       Alternative collating sequence  . . . . :            No
     Files accessed by logical file            PFILE
        File           Library           LF Format
        Z8189DB        R3ERPDATA         Z8189DB
   Sort Sequence . . . . . . . . . . . . . . . : SRTSEQ     *HEX
   Language identifier . . . . . . . . . . . . : LANGID     ENU
 Member Description
   Member  . . . . . . . . . . . . . . . . . . : MBR        CRTSAPLF
     Member level identifier . . . . . . . . . :            1051122092942
     Member creation date  . . . . . . . . . . :            11/22/05
```

Figure 29-11 DSPFD for logical file T8189DB in library R3ERPLF

You can predefine a key field to use an index (named ZITEMM in this example) and the CCSID for the fields that allow a mapping from the Unicode or GLS Latin-1 code page. In this example, CCSID 500 is specified, as shown in Figure 29-10 on page 633. Therefore, this logical file converts all fields of the Z8189DB table from code page 13488 (Unicode) into code page 500 (EBCDIC), as shown in Figure 29-12 on page 635.

```
                              Display File Field Description
  Input parameters
    File  . . . . . . . . . . . . . . . . . . . :  Z8189DB
      Library . . . . . . . . . . . . . . . . . :  R3ERPLF
  File Information
    File  . . . . . . . . . . . . . . . . . . . :  Z8189DB
      Library . . . . . . . . . . . . . . . . . :  R3ERPLF
    File location . . . . . . . . . . . . . . . :  *LCL
    Externally described  . . . . . . . . . . . :  Yes
    Number of record formats  . . . . . . . . . :      1
    Type of file  . . . . . . . . . . . . . . . :  Logical
    File creation date  . . . . . . . . . . . . :  11/22/05
    Text 'description'. . . . . . . . . . . . . :  Z8189DB
  Record Format Information
    Record format . . . . . . . . . . . . . . . :  Z8189DB
    Format level identifier . . . . . . . . . . :  29BD4C6351025
Number of fields  . . . . . . . . . . . . . :      2
    Record length . . . . . . . . . . . . . . . :     31
  Field Level Information
             Data       Field  Buffer    Buffer        Field    Column
    Field      Type      Length  Length  Position        Usage    Heading
    ZITEMM     CHAR           6       6         1        Both     ZITEMM
      Coded Character Set Identifier  . . . . . :    500
    ZITEMD     CHAR          25      25         7        Both     ZITEMD
      Coded Character Set Identifier  . . . . . :    500
```

Figure 29-12 DSPFFD for logical file Z8189DB

This logical file provides an interface for the specified code page of 500. The view of this file is compatible with the T8189DB file, see Figure 29-3 on page 626. With this view it is possible to print the contents of the file Z8189DB using the RPG program, which is defined in Figure 29-2 on page 625.

Figure 29-13 summarizes how RPG programs can access SAP database tables by using the CRTSAPLF command.

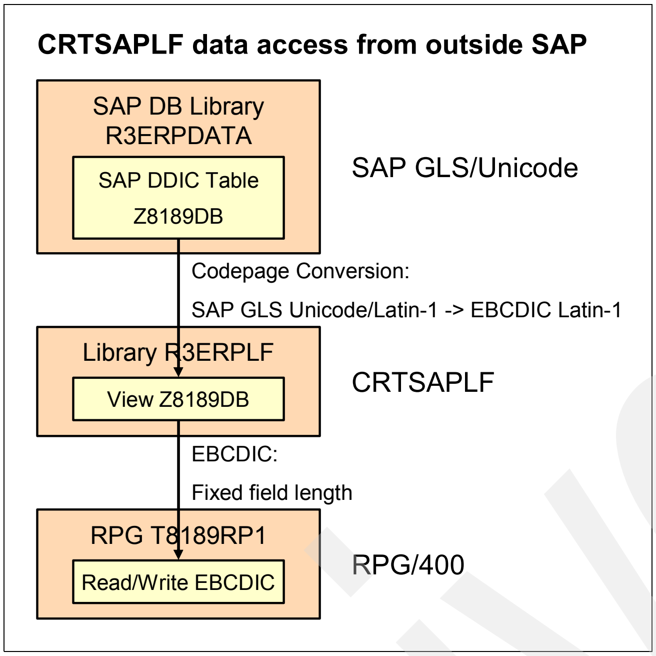

Figure 29-13 CRTSAPLF: Data access from RPG/400

Modifying the parameters of the CRTSAPLF command

Update the source file that is automatically generated by the CRTSAPLF command if you need some special parameters defined in the DDS source, which you cannot select via the CRTSAPLF command. Figure 29-14 on page 637 shows the source file for this example.

```
 Columns . . . :    1  71           Browse                   CRTSAPLF/QDDSLFWRK
 SEU==>                                                                CRTSAPLF
 FMT LF .....A..........T.Name++++++.Len++TDpB......Functions++++++++++++++++++
        *************** Beginning of data *************************************
0001.00                R Z8189DB                   PFILE(Z8189DB   )
0002.00 SOURCE* GENERATED WITH CRTSAPLF IDEATED
0003.00       *          MODELL0123   111P 3
0004.00       *          MODELL0123   111A         CCSID(00500)
0005.00                  ZITEMM     00006A         CCSID(00500)
0006.00                  ZITEMD     00025A         CCSID(00500)
0007.00                K ZITEMM
        ****************** End of data ****************************************









 F3=Exit   F5=Refresh   F9=Retrieve   F10=Cursor   F11=Toggle   F12=Cancel
 F16=Repeat find        F24=More keys
                                          (C) COPYRIGHT IBM CORP. 1981, 2003.
```

Figure 29-14 CRTSAPLF: Source file for Z8189DB

For example, edit the source with 'DESCEND' after the key field to order the ZITEMM key field in descending sequence. Then, recreate the file with the Create Logical File command (CRTLF) and change the name of the logical file from CRTSAPLF to the proper name. Use this method to create any DDS-LF.

> **Note:** A query typically reverts back to Classic Query Engine (CQE) from Standard Query Engine (SQE) whenever the database optimizer processes table objects that have any of the following defined:
>
> - Logical files with SELECT/OMIT DDS keyword specified
> - Non-standard indexes or derived keys, for example logical files specifying the DDS keywords RENAME or Alternative Collating Sequence (ACS) on any field referenced in the file
>
> This change in the Optimizer version can influence the performance behavior of the query.

29.3.2 CRTSAPVIEW: Accessing SAP GLS data with non Latin-1 code pages

Use the Create SAP View (CRTSAPVIEW) command to access SAP GLS database tables with the non Latin-1 code page from a native i5/OS environment. This command is delivered as a part of the SAP kernel library. It allows you to create a view in the library R3*SID*VIEW for each table of the SAP database to be accessed from outside the SAP application. The view converts all character data from the SAP GLS database table into the CCSID of the job that ran your CRTSAPVIEW command. This command supports all EBCDIC CCSIDs, including the double-byte codes.

These views do not allow an update of character fields. The character fields are pure input fields without any update capability. The views read the whole table sequentially because the views are unable to use any indexes. This can be slow if only a few records of a very large table are selected. In this case, it may be better to use this view in order to replicate the complete table into a format of your choice and then access the copy.

The CRTSAPVIEW command takes the session CCSID and assumes that you are running the same code page both in your SAP GLS system and your EBCDIC environment from where the table is to be accessed in the future.

> **Note:** The CRTSAPVIEW command is used to convert data inside the same code page only, for example, to convert from GLS Latin-2 (SAP code page 1401) to EBCDIC Latin-2 (IBM code page 870).

Views that are created by the CRTSAPVIEW command provide the same columns as the underlying tables. All character type data (GRAPHIC, VARGRAPHIC, and DBCLOB) is converted into EBCDIC characters by using user-defined functions (UDFs). The actual implementation of these functions is provided by service programs named O4UDFLIB or O4PORTLIB in the kernel library.

The mapping between the UDFs and the service programs is provided through the catalog view QSYS2/SYSFUNCS. When accessing data through these views, the UDFs convert the data from the GLS code page in the SAP system to the EBCDIC code page of the job where the CRTSAPVIEW command is run under.

Use the CRTSAPVIEW command to create view Z8189DB in the library R3ERPVIEW, as illustrated in Figure 29-15. This view can be used to access table Z8189DB in the SAP database library R3ERPDATA.

```
                    Create SAP View for ASCII (CRTSAPVIEW)

Type choices, press Enter.

SAP System ID  . . . . . . . . .   ERP           Character value
Table Name . . . . . . . . . . .   Z8189DB       Character value










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

Figure 29-15 CRTSAPVIEW for table Z8189DB

> **Note:** Run CRTSAPVIEW again when the name of the kernel library changes to recreate the association in QSYS2/SYSFUNCS.

The output in this example is a logical file in the R3ERPVIEW library, as illustrated in Figure 29-16 on page 639.

```
                              Display File Field Description
  Input parameters
    File  . . . . . . . . . . . . . . . . . . . :  Z8189DB
      Library . . . . . . . . . . . . . . . . . :  R3ERPVIEW
  File Information
    File  . . . . . . . . . . . . . . . . . . . :  Z8189DB
      Library . . . . . . . . . . . . . . . . . :  R3ERPVIEW
    File location . . . . . . . . . . . . . . . :  *LCL
    Externally described  . . . . . . . . . . . :  Yes
    Number of record formats  . . . . . . . . . :      1
    Type of file  . . . . . . . . . . . . . . . :  Logical
    SQL file type . . . . . . . . . . . . . . . :  VIEW
    File creation date  . . . . . . . . . . . . :  11/29/05
  Record Format Information
    Record format . . . . . . . . . . . . . . . :  Z8189DB
    Format level identifier . . . . . . . . . . :  32BD568BDB846
Number of fields  . . . . . . . . . . . . . :      2
   Record length . . . . . . . . . . . . . . . :     66
   Format text . . . . . . . . . . . . . . . . :  FORMAT0001
 Field Level Information
             Data        Field  Buffer    Buffer        Field    Column
   Field     Type       Length  Length  Position        Usage    Heading
   ZITEMM    CHAR           12      14         1        Input    ZITEMM
     Variable length field
     Allows the null value
     Coded Character Set Identifier  . . . . . :     37
     Derived field text  . . . . . . . . . . . :
       R3_UDF_GR6 ( ZITEMM )
   ZITEMD    CHAR           50      52        15        Input    ZITEMD
     Variable length field
     Allows the null value
     Coded Character Set Identifier  . . . . . :     37
Derived field text  . . . . . . . . . . . :
       R3_UDF_GR25 ( ZITEMD )
```

Figure 29-16 CRTSAPVIEW for logical file Z8189DB

This logical file points to the SAP database file Z8189DB in library R3ERPDATA and represents an interface for the specified EBCDIC code page of 37, see Figure 29-7 on page 630. However, the view of this file is not compatible with the original T8189DB file, see Figure 29-3 on page 626. The view created by the CRTSAPVIEW command takes the EBCDIC field length, doubles it, and makes it the length of a VARCHAR column, as shown in Figure 29-16. This allows the storing of double byte EBCDIC data as well.

Clients that have RPG programs with fixed field lengths and want to keep their fixed length RPG programs require an additional view on top on this view (Z8189DB in library R3ERPVIEW), which converts to the original field definitions from a VARCHAR into a CHAR format. In this example, this is the view named Z8189DB1 in library RFC. This view provides a record with single byte EBCDIC fields and fixed field length. In this example, these are:

- ZITEMM: CHAR(6)
- ZITEMD: CHAR(25)

Figure 29-17 on page 640 shows how to create this view.

```
                              Enter SQL Statements

 Type SQL statement, press Enter.
> CREATE VIEW RFC/Z8189DB1 (ZITEMM, ZITEMD) AS SELECT CAST(ZITEMM as
     CHAR(6)), CAST(ZITEMD as CHAR(25)) FROM r3erpview/Z8189DB
     View Z8189DB1 created in RFC.
 ===>











                                                                      Bottom
 F3=Exit   F4=Prompt   F6=Insert line   F9=Retrieve   F10=Copy line
 F12=Cancel            F13=Services     F24=More keys
```

Figure 29-17 Z8189DB1: View for RPG with fixed field length

Example 29-2 shows a simple RPG program, which lists the content of the SAP database file Z8189DB in library R3ERPData by using the Z8189DB1 view.

Example 29-2 Program listing: RPG program T8189RP2

```
FZ8189DB1IF  F      31            DISK
FQSYSPRT O   F     132            PRINTER
IZ8189DB1KF  01
I                                       1   6 ZITEMM
I                                       7  31 ZITEMD
C           *IN20     DOUEQ'1'
C                     READ Z8189DB1                20
C           *IN20     IFEQ '0'
C                     EXCPTDETAIL
C                     ENDIF
C                     ENDDO
C                     MOVE '1'       *INLR
OQSYSPRT E 11             DETAIL
O                         ZITEMM
O                         ZITEMD
```

Figure 29-18 on page 641 summarizes how RPG programs with fixed field length descriptions can access SAP database tables by using the CRTSAPVIEW command.

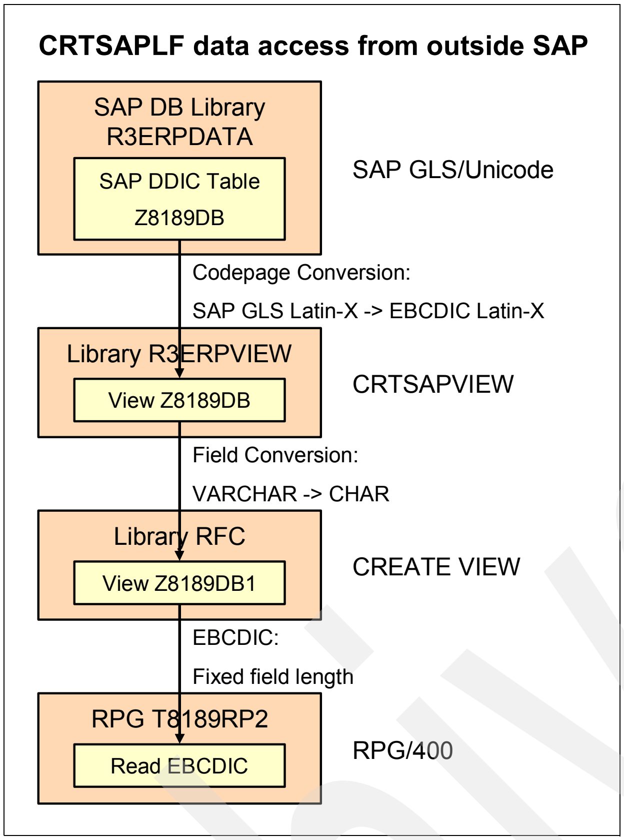

Figure 29-18 CRTSAPVIEW: Data access from RPG/400

> **Note:** The RPG program can also be defined according to the requirements of the CRTSAPVIEW command. In this case, directly specify VARCHAR fields in the RPG program. The Z8189DB1 view to convert VARCHAR into a CHAR field type is not necessary.

29.4 Accessing non-SAP data with ABAP programs

You can directly access the following objects from an ABAP program:

- Non-SAP database tables on the same server or on a remote server
- Sequential files (stream files and program-described physical files)

Use the following functions to access these objects:

- EXEC-SQL: All SQL commands are passed directly to the DB2/400 database manager to be run when you use the EXEC-SQL interface. This interface enables you to embed all functions supported by SQL/400® and work with all DB2/400 database objects to which you have access. Code page conversion between EBCDIC and the SAP Unicode or SAP GLS Latin-1 environment is automatically done by the SQL interface. This allows direct access to non-SAP database files in the same i5/OS environment where the SAP system is running.
- Database multiconnect with EXEC SQL: This function is based on EXEC SQL and allows an ABAP program to connect to non-SAP databases on the same or on a remote System i server. These databases are not subject to SAP data dictionary definitions, therefore, they can be configured and used within the System i environment in any way.

Note: You can access a DB2 UDB for iSeries database from any application server that supports DB2 UDB for iSeries, such as i5/OS, Windows 32-bit, Windows 64-bit, and Linux on Power.

You can access other remote databases, such as DB2 UDB for Linux/Unix/Windows, MaxDB for MySQL, or Oracle from application servers running Windows 32-bit (IA32), Windows 64-bit (IA64 or x64), or Linux on Power.

You can access the MS SQL Server database from application servers running Windows 32-bit (IA32) or Windows 64-bit (IA64 or x64).

- OPEN DATASET: Use text mode to access Integrated File System stream files or program-described physical files. The ENCODING parameter provides code page conversion between the SAP code page environment (UNICODE or GLS ASCII) and the i5/OS EBCDIC environment. Figure 29-19 illustrates the logic of accessing non-SAP data with an ABAP program.

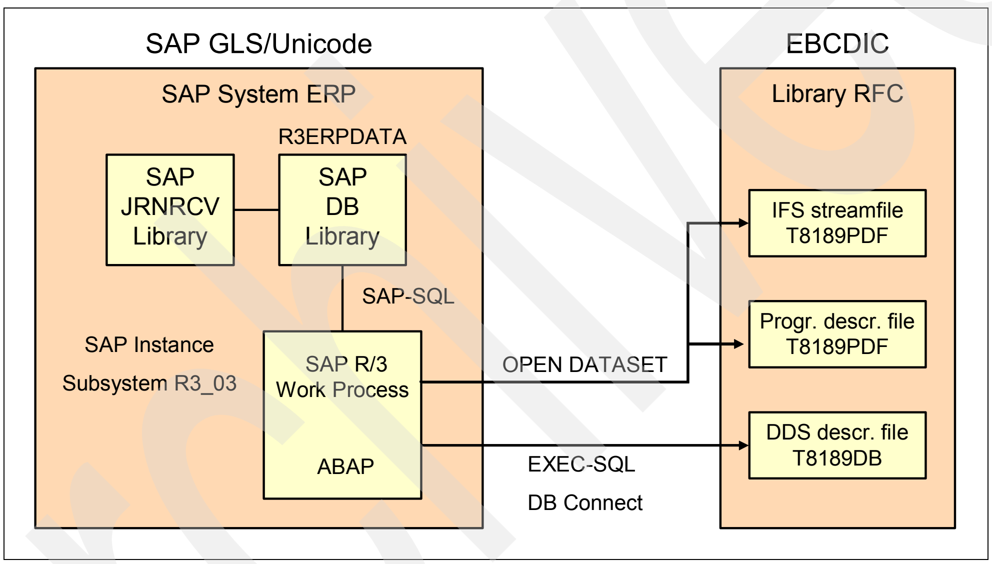

Figure 29-19 ABAP data access

The following sections describe example ABAP programs that are written to access non-SAP database data:

- EXEC-SQL (29.4.1, “Accessing non-SAP data with EXEC-SQL” on page 642)
- Database multiconnect (29.4.2, “Accessing non-SAP data with database multiconnect” on page 643)
- OPEN DATASET (29.4.3, “Accessing non-SAP data with open dataset” on page 648)

29.4.1 Accessing non-SAP data with EXEC-SQL

On the same host you can directly access non-SAP database tables from ABAP programs through the native SQL interface. The native SQL interface allows you to execute SQL statements directly in the syntax of the underlying database without using the SAP data dictionary. In the ABAP program, native SQL statements must be enclosed between the statements EXEC SQL and ENDEXEC. You can execute all valid SQL statements that are supported by DB2 UDB with i5/OS as long as the SAP instance user profile (*SID inst#*) is authorized to use the affected objects. Code page conversions between the accessed objects (usually EBCDIC) and the SAP code page of the application server are automatically *done* by the native SQL interface.

> **Note:** Code page conversion can only be done for SAP Unicode independent of the code page of the non-SAP data to be read or for ASCII Latin-1 with EBCDIC Latin-1 data to be read.

Example 29-3 shows the ABAP program named Z8189SQL using an EXEC-SQL statement to select data from outside the SAP database (the physical file named T8189DB in library RFC). The OUTPUT-ITEM subroutine is performed for each record retrieved from the T8189DB file.

Example 29-3 ABAP program using EXEC-SQL

```
REPORT  Z8189SQL.
DATA: BEGIN OF REC,
      RITEMM(5) TYPE C,
      RITEMD(25) TYPE C,
      END OF REC.
EXEC SQL PERFORMING OUTPUT-ITEM.
  SELECT ITEMM, ITEMD
        INTO :REC
        FROM RFC/T8189DB
ENDEXEC.
FORM OUTPUT-ITEM.
  WRITE: / REC-RITEMM, REC-RITEMD.
ENDFORM.
```

29.4.2 Accessing non-SAP data with database multiconnect

Use the database multiconnect function in order to access database objects on a remote database. You can maintain connection information for a remote System i server in the DBCON table and use the native SQL statement CONNECT TO *connection* to establish a connection to that remote database. Subsequent native SQL statements use this connection until it is reset to the local database with the native SQL statement SET CONNECTION DEFAULT. For more information about database multiconnect, refer to *SAP note 146624*.

ABAP example for database multiconnect

This section shows an ABAP example, which uses the database multiconnect function to access a database table outside of the SAP database. Example 29-4 shows the ABAP program named Z8189AB5 using the ‘MYCON’ database connection to select data from outside the SAP database. In this example, this is the physical file T8189DB in library RFC.

Example 29-4 Program listing: ABAP program using database multiconnect

```
REPORT  Z8189AB5.
DATA: BEGIN OF REC,
      RITEMM(5) TYPE C,
      RITEMD(25) TYPE C,
      END OF REC,
      CON_NAME(30) VALUE 'MYCON'.
EXEC SQL.
  CONNECT TO :CON_NAME
ENDEXEC.
EXEC SQL PERFORMING OUTPUT-ITEM.
  SELECT ITEMM, ITEMD
      INTO :REC
```

```
      FROM RFC/T8189DB
ENDEXEC.
EXEC SQL.
  SET CONNECTION DEFAULT
ENDEXEC.
FORM OUTPUT-ITEM.
  WRITE:/ REC-RITEMM, REC-RITEMD.
ENDFORM.
```

It uses the database connection ’MYCON’, as shown in Figure 29-20. The OUTPUT-ITEM subroutine is performed for each record retrieved from the T8189DB file.

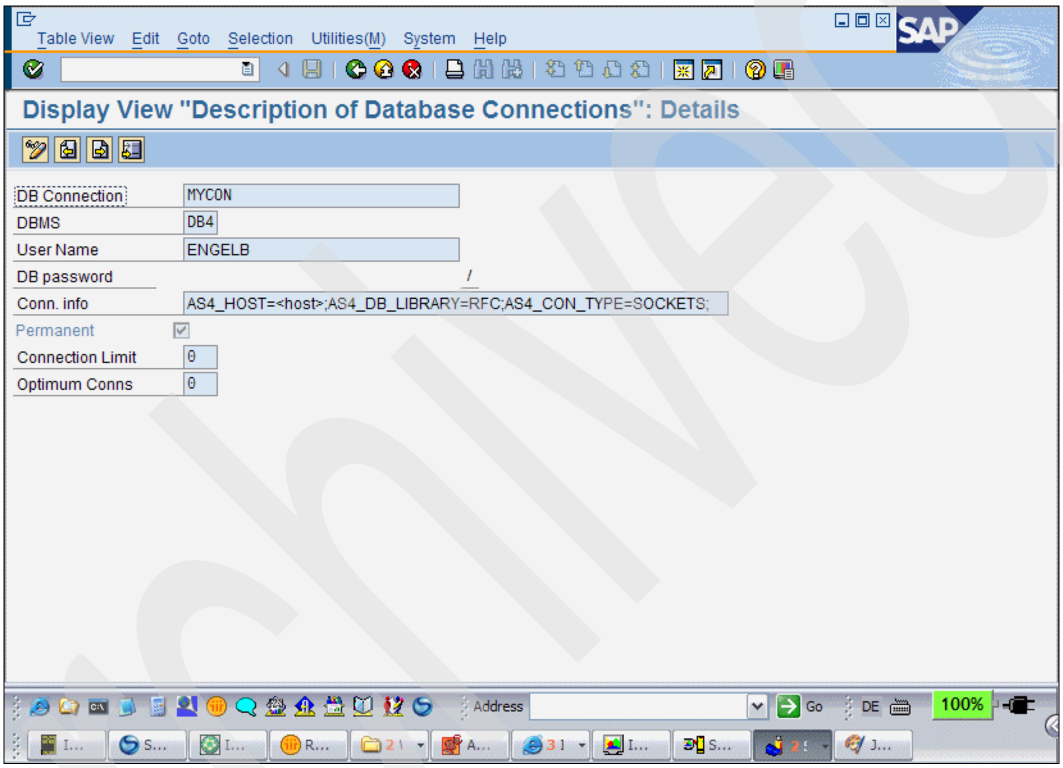

Figure 29-20 SM30: Maintain table DBCON

Java example for database multiconnect

In the following example, table T8189DB in library RFC is read using the Toolbox JDBC driver. Figure 29-21 on page 645 shows that Toolbox JDBC driver is registered as ’DB4_Toolbox_JDBC_Driver’ in JDBC Connector Service.

Furthermore, a JDBC 1.x data source AS0012_RFC is already configured. It is also known as CUSTOM/MY_DATASOURCE. In the Additional tab, it is specified that this should be a Native SQL connection, as shown in Figure 29-22 on page 646.

Use Java multi-connect for direct access to non-SAP database tables on the same host or on a remote host based on the so-called Java Database Connectivity (JDBC) interface.

Use the JDBC Connector Service in Visual Admin Tool to configure connections to non-SAP databases. This is a two-step process:

1. If you cannot use the system JDBC driver (SYSTEM_DRIVER) that SAP uses to access the Java database SAP*sid*DB, for example, because the external data resides in a

database of a different type, you must first configure an appropriate JDBC driver for that database type. In principle, you can use any JDBC driver available, but if possible, we recommend that you use one of the drivers particularly certified by SAP.

For a list of supported databases and JDBC drivers refer to *SAP note 907733*. For access to a DB2 UDB for iSeries, you can, for example, configure the iSeries Toolbox JDBC driver. You can use the iSeries Native JDBC driver as well, but only for local access and when already being configured as the system JDBC driver.

2. Define a data source. This is a connection definition similar to a DBCON entry from which your Java 2 Platform, Enterprise Edition (J2EE) application can then derive certain types of database connections. It contains user/password information, database host, schema name, JDBC driver to be used, information about the connection type as well as connection parameters. It is possible to define data sources according to different JDBC specifications:

   – From a JDBC URL (JDBC 1.x): In this case, SAP J2EE Engine takes care of connection pooling. Distributed transactions are not allowed.

   – From a native JDBC Data source object (JDBC 2.0): Use an XADatasource in case you require your connections to participate in distributed transactions. Otherwise, use a ConnectionPoolDatasource. In this case, the JDBC driver itself takes care of the connection pooling.

Figure 29-21 and Figure 29-22 on page 646 show how the connection needs to be defined:

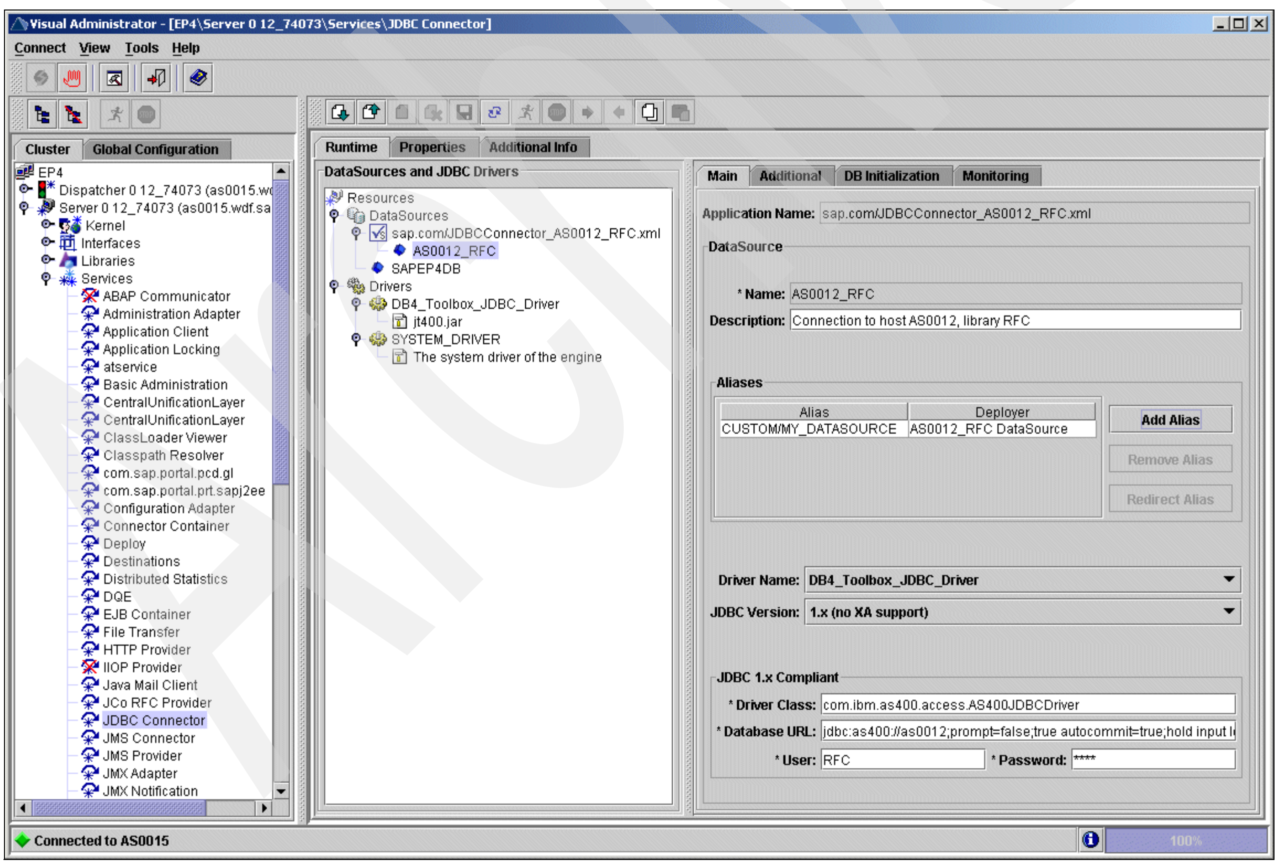

Figure 29-21 Visual administrator: JDBC connector, register connection

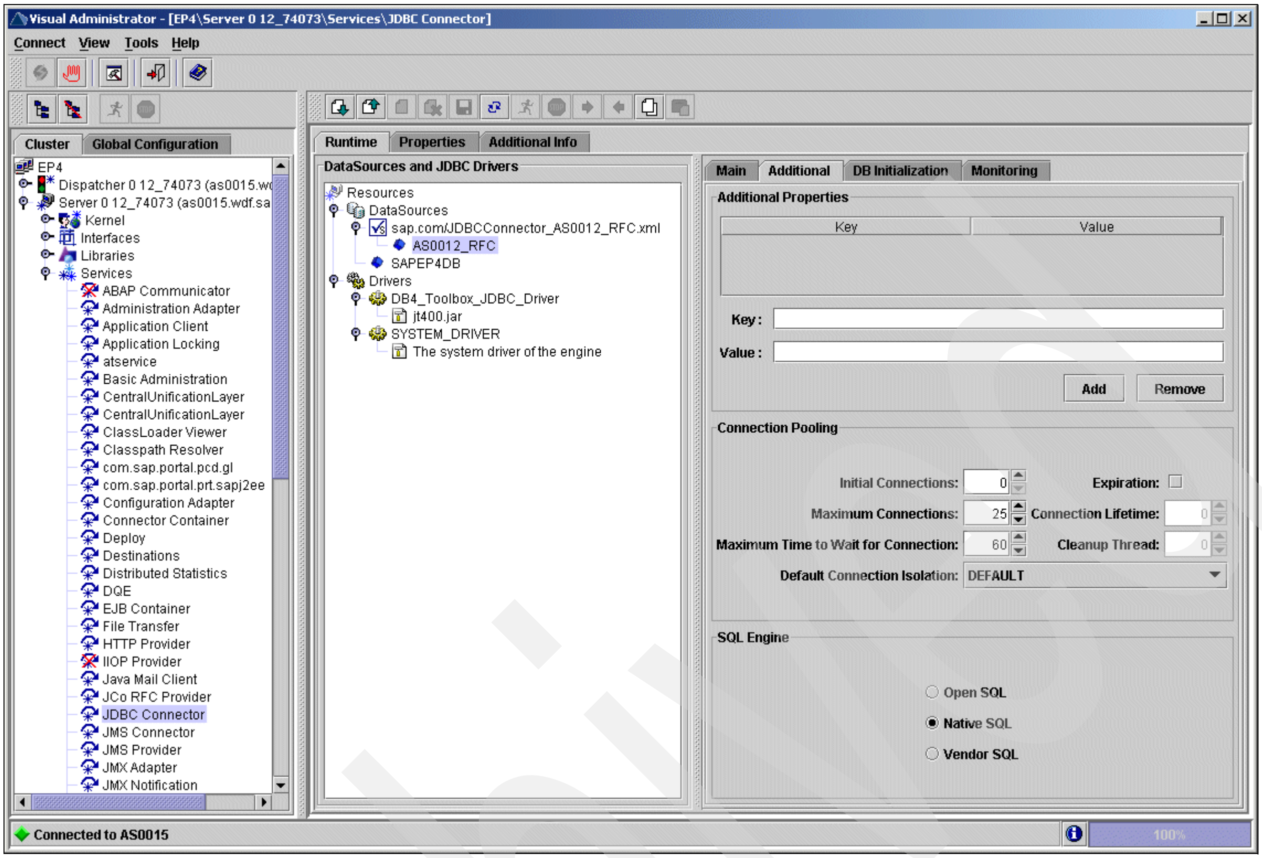

Figure 29-22 Visual administrator: JDBC connector, define SQL connection

You can choose between Vendor SQL, Native SQL, and Open SQL for the level of support given by the Open SQL Framework:

- Vendor SQL: This means that connections are exactly as returned by the JDBC driver. The Open SQL Framework for Java does not interfere at all. If your JDBC driver does not belong to the set of SAP-supported JDBC drivers as mentioned in *SAP note 907733*, this connection type is your only choice.
- Native SQL: This connection type is enriched by SAP with corrections for known driver bugs, as well as performance enhancements such as a transparent PreparedStatement pool, and tracing capabilities, namely the SQL trace for Java. However, the SQL dialect used by this connection is still the dialect of the underlying database system.
- Open SQL: This gives full support by the Open SQL Framework. That is, in addition to Native SQL, it provides you data access independent of the underlying database system, and further performance enhancements such as the table buffers. Open SQL is not an option for multi-connect to any schema different from the Java database SAP*sid*DB.

> **Note:** If possible, use Native SQL for multi-connect to non-SAP data. Use Vendor SQL only if you have no other choice.

Example 29-5 on page 647 shows a servlet that retrieves this data source through Java Naming and Directory Interface™ (JNDI) lookup specifying jdbc/CUSTOM/MY_DATASOURCE. From this, it derives a JDBC connection in order to read from the table. (In fact, the exact JNDI lookup name depends on the J2EE application kind.)

Note that in this example the table name need not be fully qualified with the library name. This is because the connection uses SQL naming and the database user name equals the library name. This sets the default schema to the database user name. Alternatively, you can specify the schema name on the URL to set the default schema.

Example 29-5 Program listing: Java program using database multiconnect

```
package myPackage;

import javax.naming.Context;
import javax.naming.InitialContext;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;
import javax.servlet.ServletException;
import javax.servlet.ServletOutputStream;
import javax.sql.DataSource;

import java.io.IOException;
import java.sql.Connection;
import java.sql.PreparedStatement;
import java.sql.ResultSet;

public class TestServlet extends HttpServlet {
    protected void service(HttpServletRequest req, HttpServletResponse res)
            throws ServletException, IOException {
        Context jndiContext = null;
        ResultSet rs = null;
        PreparedStatement pstmt = null;
        Connection con = null;
        try {

            // Get the datasource from alias
            jndiContext = new InitialContext();
            DataSource ds = (DataSource) jndiContext.lookup("jdbc/CUSTOM/MY_DATASOURCE");

            // Get connection to from datasource
            con = ds.getConnection();

            // Prepare and execute query
            pstmt = con.prepareStatement("select * from T8189DB");
            rs = pstmt.executeQuery();

            // Return table contents to browser
            res.setContentType("text/plain");
            ServletOutputStream out = res.getOutputStream();
            out.println("Contents of table T8189DB");
            out.println("-------------------------");
            while (rs.next()) {
                out.println(rs.getString(1) + ", " + rs.getString(2));
            }

        } catch (Exception e) {
            throw new ServletException(e);
        } finally {
            try {
                // Make sure all resources get closed or returned to their pools
                rs.close();
                pstmt.close();
                con.close();
```

```
            } catch (Exception ex) {}
        }
    }
}
```

Note: In some respects, Java gives more flexibility compared to ABAP multi-connect. For example:

- Each database system supported by SAP provides at least one type 4 JDBC driver. Since the client libraries of type 4 drivers are completely written in Java, they can be used from any SAP J2EE Engine instance independent from the underlying operating system.
- Database vendors other than the ones supported by SAP also often provide type 4 JDBC drivers. They can be used in Vendor SQL connections.

29.4.3 Accessing non-SAP data with open dataset

ABAP supports the reading and writing of System i stream files (object type *STMF) and System i program described files (object type *file, without DDS-description). These System i files are stored outside of the SAP system in the code page of the operating system environment. A code page conversion is required when passing data between the GLS ASCII or Unicode environment of the SAP system and the EBCDIC-based environment of the i5/OS operating system. Therefore, text mode is required in the ABAP program to access these data files.

Writing System i stream files with an ABAP program

The ABAP program shown in Example 29-6 uses an EXEC-SQL interface to read data from the non-SAP database table named T8189DB.

Example 29-6 ABAP program to write System i stream files

```
REPORT  Z8189AB2.
DATA: FILE(20) Value '/t8189/t8189seq'.
DATA: BEGIN OF REC,
      RITEMM(5) TYPE C, RITEMD(25) TYPE C,
      END OF REC.
  OPEN DATASET FILE FOR OUTPUT in TEXT MODE ENCODING UTF-8.
EXEC SQL PERFORMING OUTPUT-ITEM.
      SELECT ITEMM, ITEMD
            INTO :REC
            FROM RFC/T8189DB
ENDEXEC.
FORM OUTPUT-ITEM.
      TRANSFER REC to FILE.
ENDFORM.
```

The table is DDS-described and stored in the library named RFC. It writes data to a stream file *named t8189seq located in the Integrated File System directory of the i5/OS operating system (directory /t8189). The stream file is opened in text mode. The data is automatically converted into the correct code page.

Reading System i stream files with an ABAP program

The ABAP program shown in Example 29-7 on page 649 reads the data sequentially from the existing stream file named t8189seq, which is located in the System i /t8189 Integrated File

System directory, and prints the file content. The stream file is opened in text mode. Data conversion is done between both code page environments.

Example 29-7 ABAP program to read System i stream files

```
REPORT  Z8189AB3.
DATA: FILE(20) VALUE '/t8189/t8189seq'.
DATA: BEGIN OF REC,
      RITEMM(5) TYPE C, RITEMD(25) TYPE C,
      END OF REC.
     OPEN DATASET FILE FOR INPUT IN TEXT MODE ENCODING UTF-8.
DO.
READ DATASET FILE INTO REC.
      IF SY-SUBRC NE 0. EXIT. ENDIF.
        WRITE: / REC-RITEMM, REC-RITEMD.
ENDDO.
CLOSE DATASET FILE.
```

Accessing System i program described files with an ABAP program

The ABAP program shown in Example 29-8 reads data from a System i program described file and prints the contents. The file is opened in text mode. Data conversion is done between both code page environments.

Example 29-8 ABAP program to access program described files

```
REPORT  Z8189AB4.
DATA: FILE(60) VALUE '/qsys.lib/rfc.lib/t8189pdf.file/t8189pdf.mbr'.
DATA: BEGIN OF REC,
      RITEMM(5) TYPE C, RITEMD(25) TYPE C,
      END OF REC.
    OPEN DATASET FILE FOR INPUT IN TEXT MODE ENCODING UTF-8.
DO.
READ DATASET FILE INTO REC.
   IF SY-SUBRC NE 0. EXIT. ENDIF.
       WRITE: / REC-RITEMM, REC-RITEMD.
ENDDO.
CLOSE DATASET FILE.
```

Program Z8189AB4 is similar to the program Z8189AB3 shown in Example 29-7. The only difference is that the file used in this example is a System i program described physical file. The physical file named T8189DB is accessed using the Integrated File System directory structure.

> **Note:** This access type is not valid for DD- described physical files. Use EXEC-SQL to access DDS-described physical files. See Example 29-3 on page 643 for an illustration.

29.5 Process i5/OS commands from SAP applications

The SAP system environment provides transactions, from where you can work with i5/OS system commands. For example:

- SAP transaction SM69: To define i5/OS commands in the SAP system
- SAP transaction SM49: To execute i5/OS commands directly or in the background
- SAP transaction SM36: To define SAP background jobs. In this case, you can define one or several job steps to execute an external (i5/OS) command.

For more information about working with i5/OS commands from inside SAP systems, refer to *SAP note 677345*.

Note: All commands are executed by the assigned work process. The work process job must have the appropriate right to execute this command.

29.5.1 Defining i5/OS commands

Additional i5/OS commands can be created using the SAP transaction named SM69, as shown in Figure 29-23.

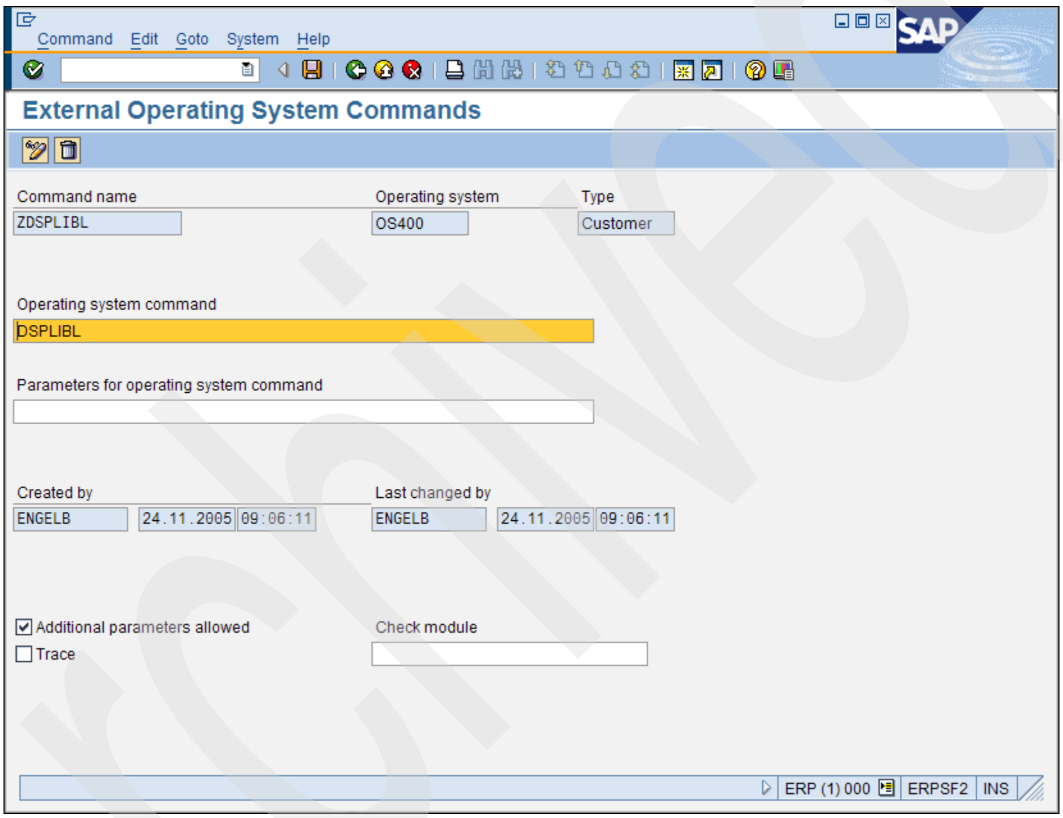

Figure 29-23 SM69: SAP transaction creating the i5/OS command

See 29.5.3, “Defining i5/OS job steps for SAP system background jobs” on page 652 for more detail about how to use the background job command interface.

29.5.2 Executing i5/OS commands

Use the SAP transaction named SM49 to execute a predefined i5/OS command. Figure 29-24 on page 651 shows the input parameters for this transaction. Note that transaction SM49 allows you to execute the i5/OS command on another connected host system.

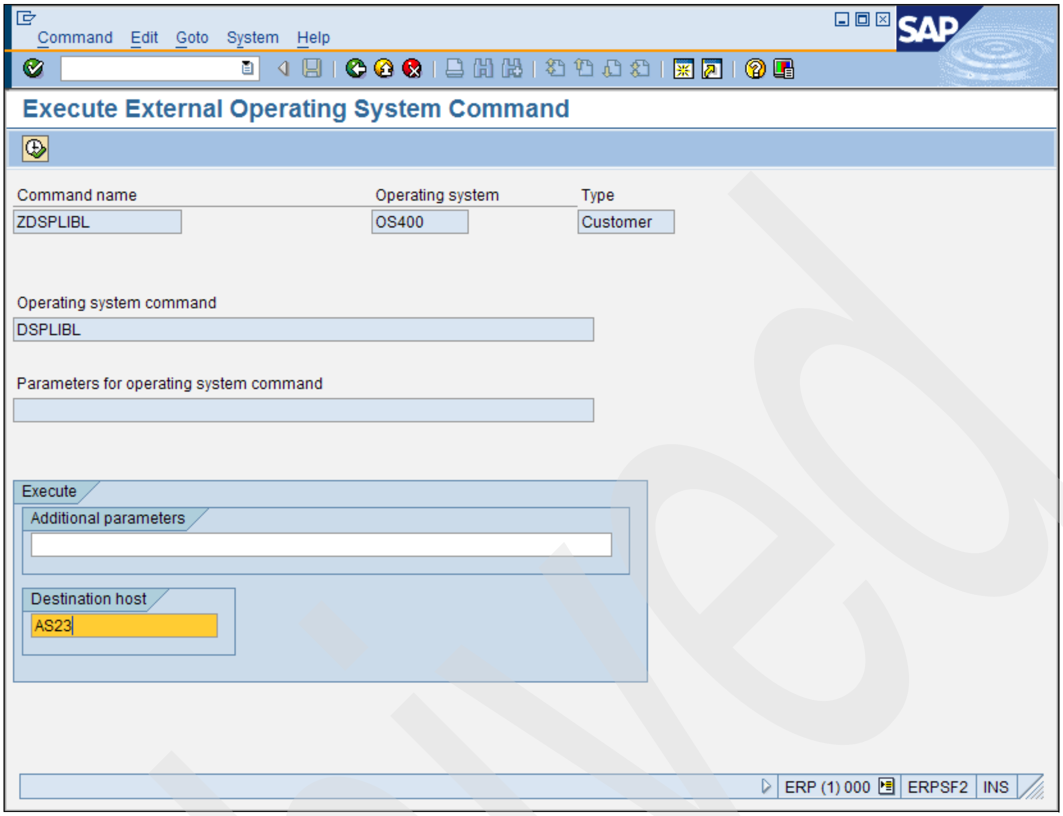

Figure 29-24 SM49: SAP transaction executing the i5/OS command

The command results are shown in Figure 29-25.

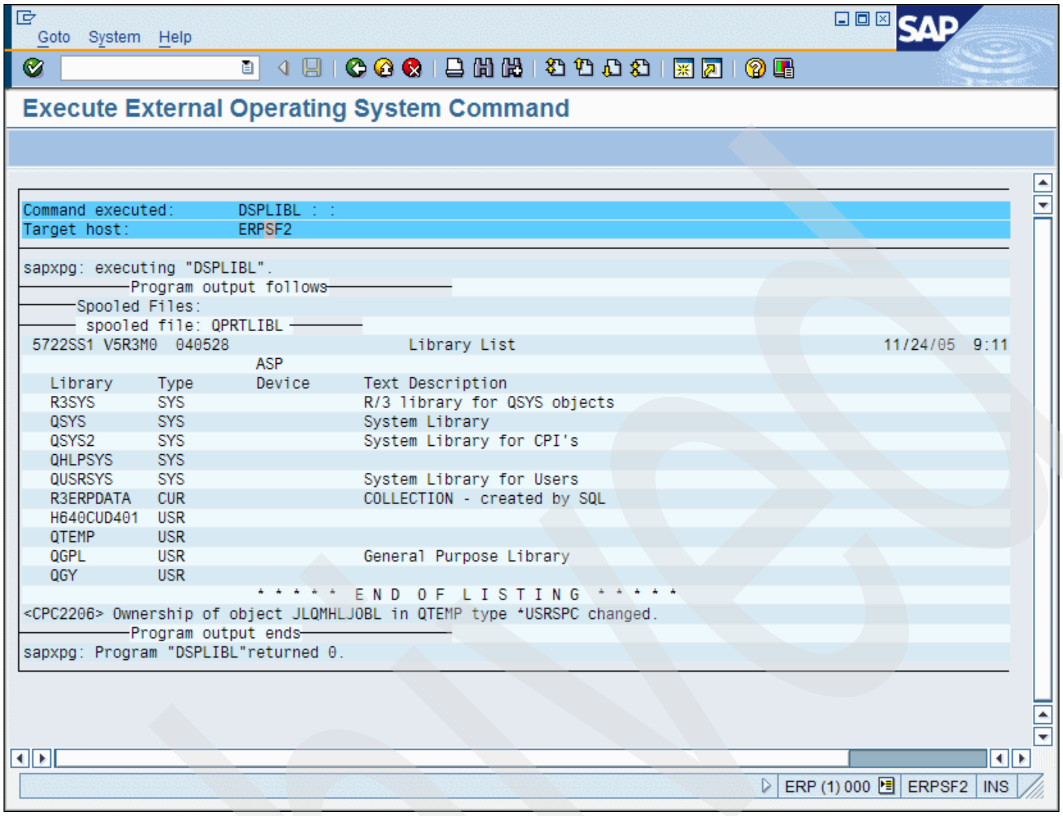

Figure 29-25 SAP program output results

SAP function modules for external commands

SAP provides function modules (based on predefined i5/OS commands) to check and execute 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

Use these function modules in ABAP programs to execute predefined i5/OS commands. Refer to the SAP online documentation for detailed information about these function modules.

29.5.3 Defining i5/OS job steps for SAP system background jobs

Use the SM36 SAP transaction to define background jobs in the SAP system. The example described in this section shows how to execute System i commands from an SAP background job.

Figure 29-26 on page 653 shows how to define the job name and assign the job priority. Note that you can select to run the background job on a specific server of the SAP system. Next, define the job steps to be executed by the background job.

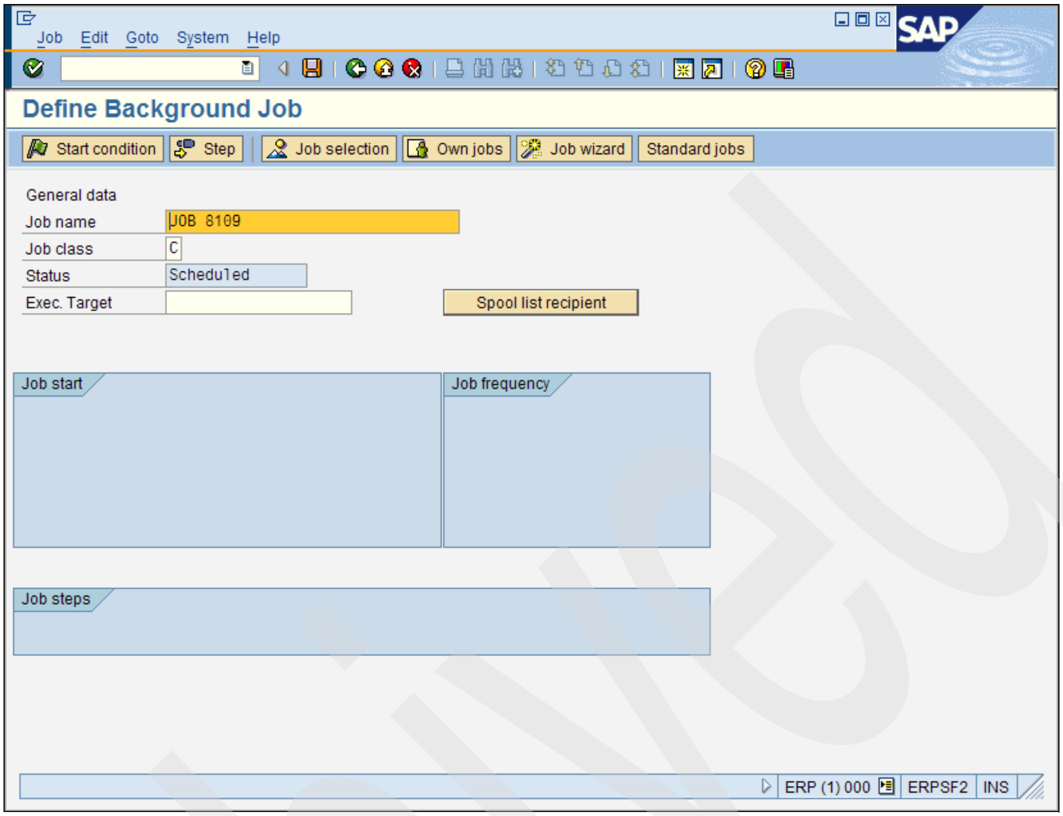

Figure 29-26 SM36: Defining an SAP background job

Figure 29-27 on page 654 shows how to specify a predefined i5/OS external command as a job step. Use the SAP transaction SM69 to create external commands. Refer to 29.5.1, "Defining i5/OS commands" on page 650 for more information about how to predefine and execute i5/OS commands in an SAP system environment.

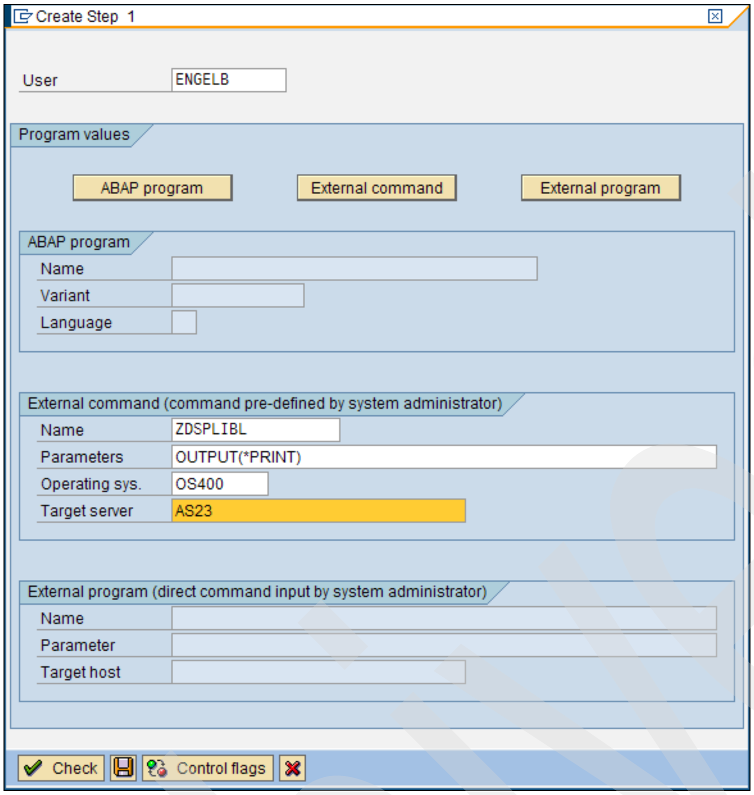

Figure 29-27 SAP program defining external command job steps

Figure 29-28 on page 655 shows how to specify an external System i program as a job step. This can be a System i server program (as in our example) or a user program.

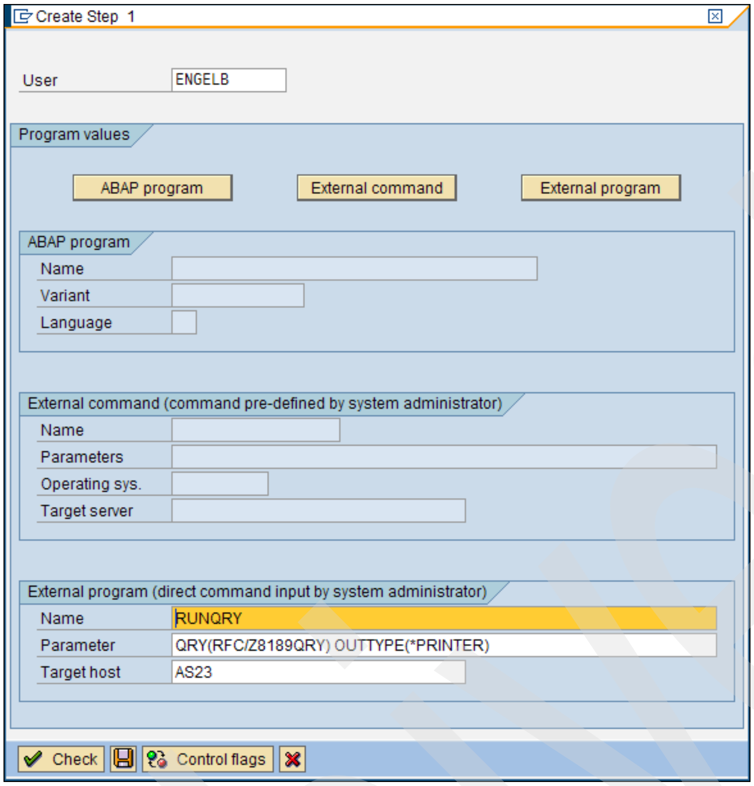

Figure 29-28 SAP program defining external program job steps

The job steps are run in the sequence specified. Schedule the background job to run at the required date and time.

29.6 System i jobs to start SAP applications

Use one of the following methods to start SAP jobs on a System i server from outside the SAP system:

- Use the SAP Event (SAPEVT) CL command to send an event to the SAP background job scheduler.
- Use the Start Report (STRREPORT) CL command to start an ABAP program.

Both the SAPEVT and STRREPORT commands are SAP supplied CL commands and reside in the SAP kernel library. The following sections further describe how to use these two methods.

29.6.1 The SAPEVT command

Use the SAP event interface SAPEVT command to enable System i jobs to trigger SAP background jobs. Use SAP transaction SM36 to specify the job start criteria when you define an SAP background job. The criteria options include:

- Immediately
- At a specific date and time
- After a specific job has finished
- At the start of a specific operation mode
- Depending on a specific event being raised

An event is a signal that can be sent to the SAP background job scheduler to trigger a process. Send this signal from either inside the SAP system (for example, SM64) or from an external i5/OS job. Use the SAP transaction SM62 to maintain events.

The following steps outline how to use SAPEVT to trigger an SAP job:

1. Use transaction SM36 to define the SAP background job.
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 reschedule the job upon completion (if required). This option automatically schedules and releases a new job (with the same name) upon completion and waits for the event to be raised again.
5. Save and release the background job. The job cannot be triggered if it is not in a released state.

This background job starts running when the event is raised. Figure 29-29 shows how to specify the event settings in the background job definition.

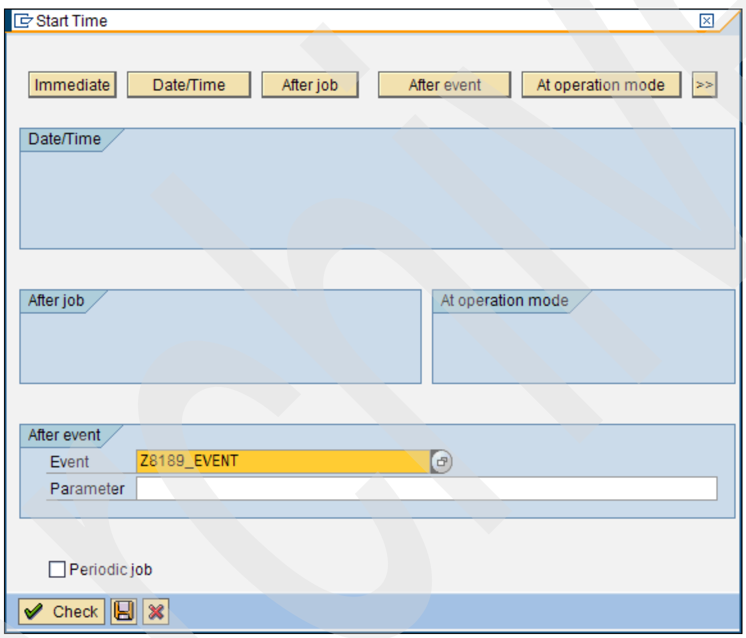

Figure 29-29 SAP job starting after an event

Use the following CL command to trigger the above event from the System i server:

```
R3640UOPT/SAPEVT EVTID(Z8189_EVENT)
PROFILE('/usr/sap/ERP/sys/profile/ERP_DVEBMGS03_AS23')
```

Incorporate this command into a CL program to trigger the event.

29.6.2 The STRREPORT command

The Start Report command (STRREPORT) starts an ABAP program using the i5/OS CL command interface. This command can be called interactively or from an i5/OS job stream.

The following i5/OS CL command executes the ABAP program named Z8189AB3 in the specified SAP system:

```
R3640UOPT/STRREPORT REPORT(Z8189AB3) JOB(Z8189AB3) SID(ERP) INSTANCE(03)
CLIENT(100) USER(ENGELB) PASSWORD() LANGUAGE(E)
```

29.7 Interactive program communication

This section describes how SAP ABAP programs and i5/OS ILE programs can evoke another and exchange data. It also shows how ABAP can communicate with Java programs. All examples are based on the RFC protocol.

29.7.1 Using an RFC connection

RFC is a consistent set of functions for program-to-program communication in a heterogeneous environment. Use it for the following types of communications:

- From ABAP to ABAP and vice versa
- From ABAP to external non-ABAP programs and vice versa
- From external non-ABAP programs to ABAP and vice versa

External programs must be written in ILE compiler languages (ILE RPG/400, ILE COBOL/400, or ILE C/400®).

RFC connections can:

- Establish and control communications
- Exchange data
- Convert data (from ASCII to EBCDIC and vice versa)
- Deallocate communications

This is similar to the CPI-C connection, which is described in *Implementing SAP R/3 on OS/400*, SG24-4672.

The functions listed previously are automatically performed by the CALL FUNCTION statement in the ABAP program. See ABAP program Z8189ABB in Example 29-9 on page 665 for an example of how to use this statement for i5/OS ILE programs. However, you must specify this information in detail, by using the RFC procedures for i5/OS ILE compilers that are delivered within the service program LIBRFC. Service program LIBRFC must be bound with the i5/OS ILE program to access these functions.

> **Note:** Use the LIBRFC for EBCDIC, because your ILE programs execute in an EBCDIC environment for System i configurations. Download the 4.6D EBCDIC LIBRFC from the kernel patch directory. The 4.6D version should be sufficient for higher releases as well.

RFC interaction between ABAP and external program

SAP applications have RFC interfaces for ABAP programs outside of an SAP implementation for an RFC-based interaction. This section shows two examples based on ILE RPG/400 and ILE COBOL/400:

- ABAP: Use statement CALL FUNCTION
- ILE program: Bind the LIBRFC Service program to the program

Table 29-1 on page 658 lists the procedures that must be called from within an ILE program to perform an RFC-based interaction.

Table 29-1 SAP RFC communication call functions

Procedure call	Description
RfcEnvironment	Sets up the RFC environment for this job
RfcAccept	Establishes a conversation
RfcInstall	Points to an RFC-definition, server program
RfcDispatch	Invokes a server program
RfcGetData	Receives data from ABAP program
RfcSendData	Sends data from ABAP program
RfcClose	Closes conversation

> **Note:** For more information about ILE RPG/400 and ILE COBOL/400, refer to:
>
> - *ILE RPG/400 Programmer's Guide*, SC09-2507
> - *ILE RPG/400 Reference*, SC09-2508

For ILE RPG/400 and ILE COBOL/400 examples, refer to "ILE RPG/400 example evoked by ABAP" on page 666 and "ILE COBOL program example evoked by ABAP" on page 670.

29.7.2 ABAP calling ILE or Java: flowchart of RFC connection

Figure 29-30 shows the interaction between the ABAP program (Example 29-9 on page 665) calling an ILE RPG/400 example (Example 29-10 on page 666 and Example 29-11 on page 668).

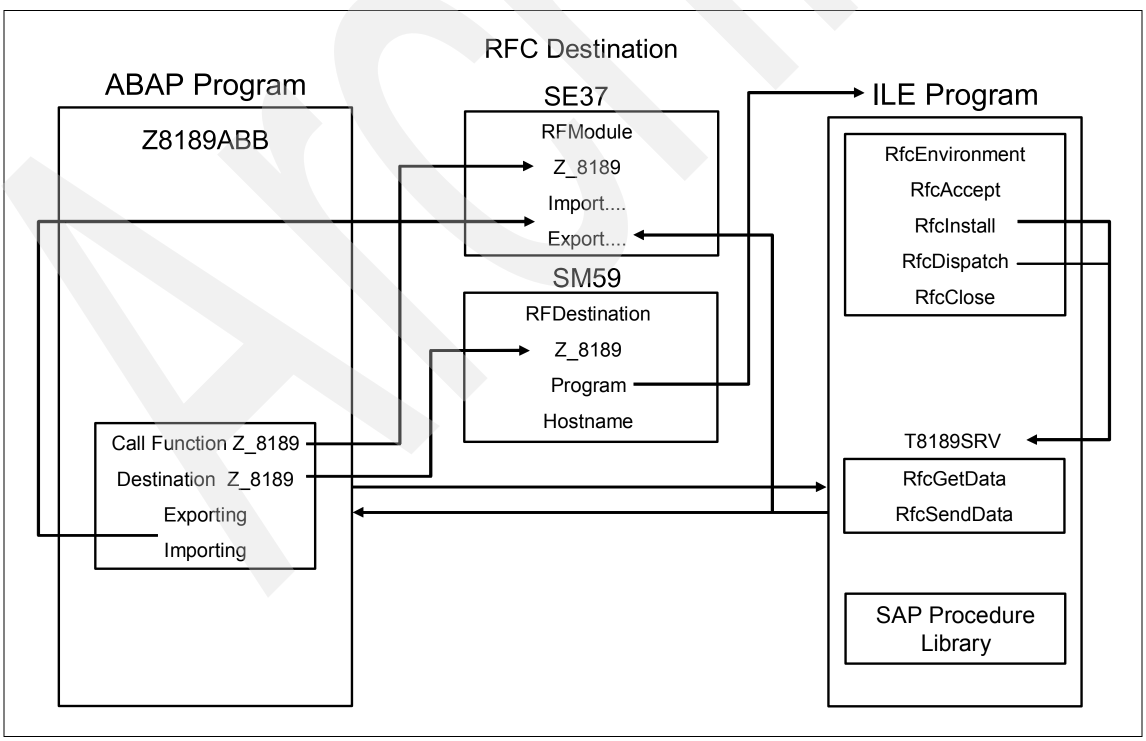

Figure 29-30 Flowchart: ABAP program triggering an ILE program

The process shown in Figure 29-30 on page 658 is explained in the following steps:

1. ABAP program Z8189ABB calls the remote function module Z_819 and establishes a session to program T8189RFC.
2. The function module is configured to send data (export) and to receive data immediately afterwards (import).
3. RFC destination Z8189 invokes the ILE program.
4. The ILE main procedure is started and performs the following actions:
   a. RfcEnvironment: Provides the environment for the RFC session.
   b. RfcAccept: A session is established and the RFC-handle is passed back to the program.
   c. RfcInstall: A connection to the remote function module and server program is established.
   d. RfcDispatch: The T8189SRV server program (a 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.

Defining RFC destination

Use SAP transaction SE37 to define a remote function group (if required), and the remote function module Z_8189, which is used by the ABAP program Z8189ABB. Figure 29-31 on page 660, Figure 29-32 on page 661, and Figure 29-33 on page 662 illustrate how to define the remote function module.

Ensure that the Function Module Processing type is set to Remote enabled module, see Figure 29-31.

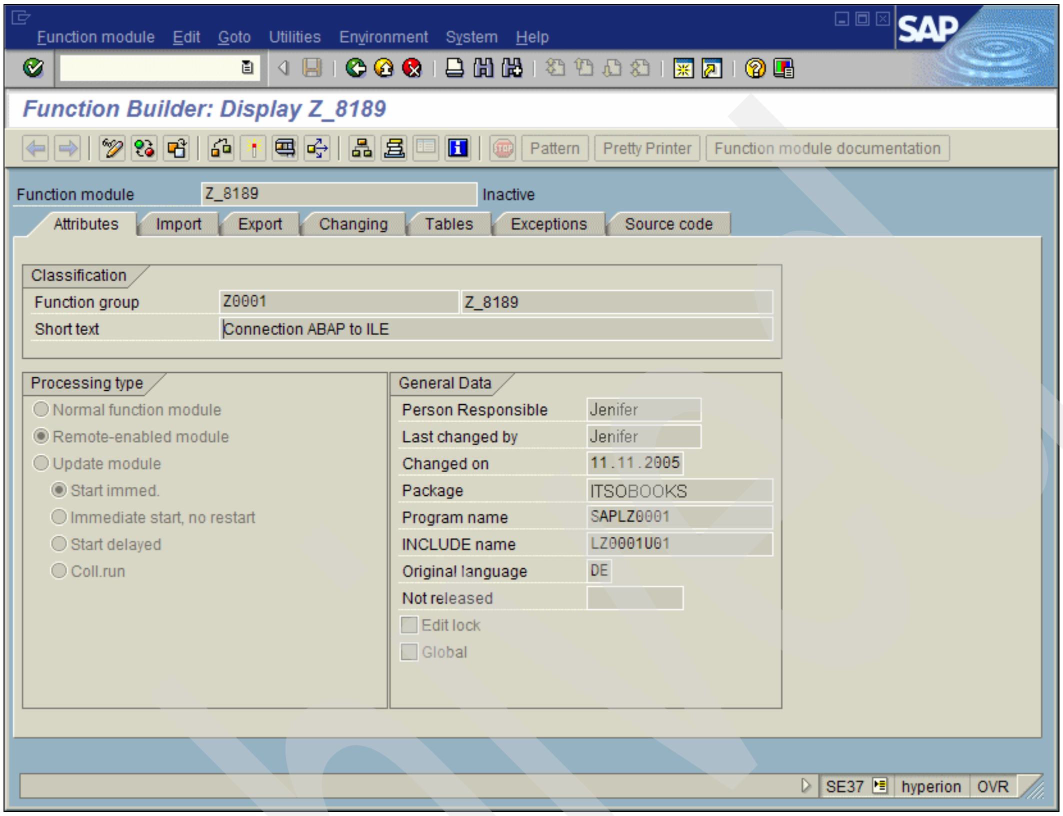

Figure 29-31 Defining the function module attributes (Part 1 of 3)

Figure 29-32 and Figure 29-33 on page 662 show the import and export settings required to communicate with the ILE program.

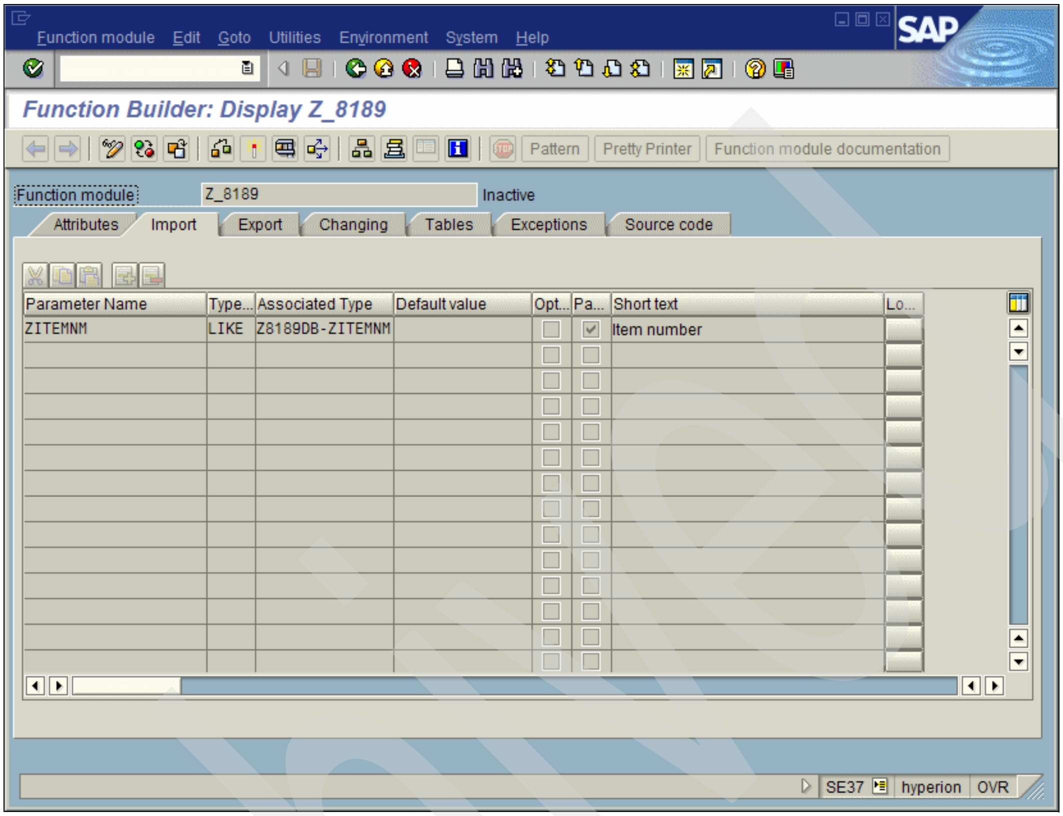

Figure 29-32 Defining the function module import parameters (Part 2 of 3)

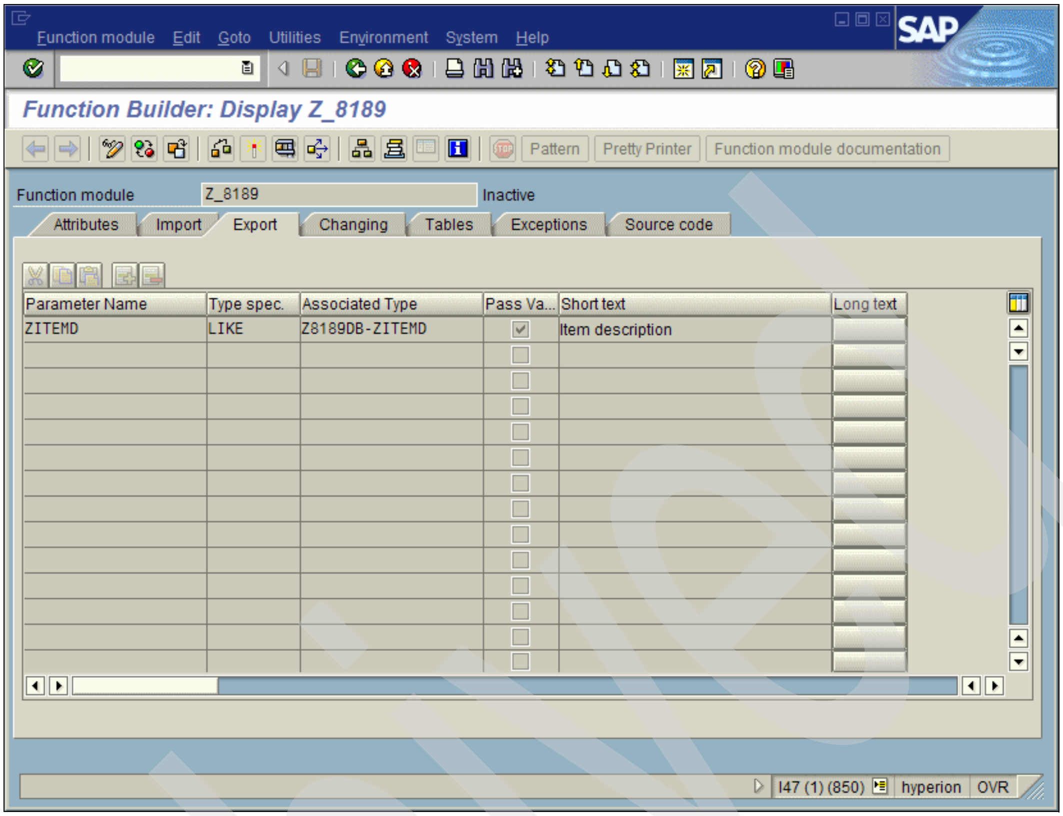

Figure 29-33 Defining the function module export parameter (Part 3 of 3)

Use the SAP transaction SM59 to create the RFC destination, as shown in Figure 29-34. The RFC connection type is ’Registered’. This means that the destination is registered under Program-ID T8189RFC after startup of the SAP system so that you can establish a communication.

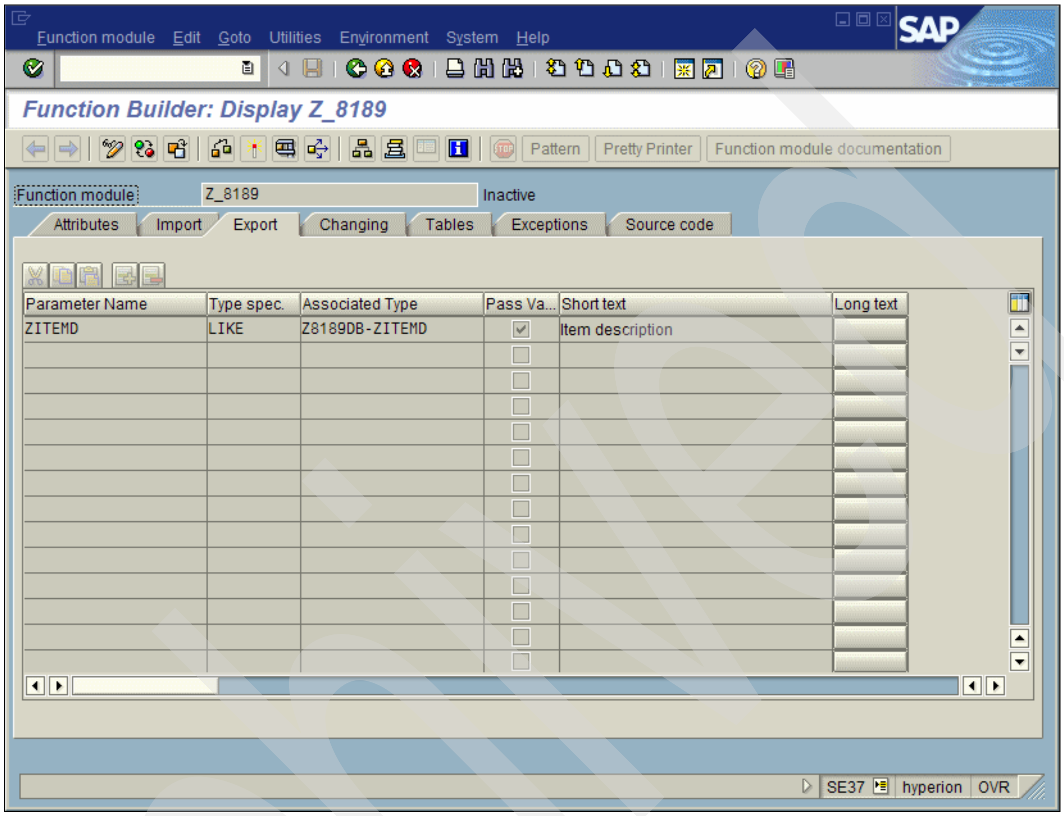

Figure 29-34 Creating an RFC destination

Specify an Activation Type for this transaction to determine how the RFC connection is activated. There are two options for this parameter:

- Start: Directly starts a program on the target host. Directly specify the library and name of the ILE program that is to be evoked from the ABAP program. For example, `qsys.lib/rfc.lib/T8189RFC.pgm`.

  Install a gateway instance on the target host if your ABAP program needs to activate an ILE program on a different System i server. For remote servers and for performance reasons, it is easier to use registration mode.

- Registered Server Program: Register the ILE program before the ABAP program establishes a connection. Register the ILE program with a program-ID, which is a logical name for the connection. In our example, this is T8189RFC.

  “CL program to register communication for ILE programs” on page 664 explains how to register a communication. You can also perform the registration from a remote server. This enables a communication between an ABAP program and an ILE program running on different servers. See Figure 29-34.

Our example (see “CL program to register communication for ILE programs” on page 664) shows an RFC connection based on an Activation Type of Registered Server Program.

Note: Use Registered Server Program to activate an ILE program on a different server. You do not need to install an SAP gateway instance on the target server.

CL program to register communication for ILE programs

Register the communication under the program-ID T8189RFC before establishing a communication between the ABAP program and the ILE program, see Figure 29-34 on page 663. Use a Control Language (CL) program to register a communication, as illustrated in Figure 29-35.

```
 Columns . . . :    1  71           Edit                            RFC/QCLSRC
 SEU==>                                                                T8189CL
 FMT **  ...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7
        *************** Beginning of data *************************************
0001.00             PGM
0001.01 /**********************************************/
0001.02 /* REGISTER TFC CONNECTION TO ABAP            */
0001.03 /*                                            */
0001.04 /*    -a<registername in SM59>                */
0001.05 /*    -g<hostname of SAP system>              */
0001.06 /*    -xsapgw<instance number of SAP system> */
0001.07 /*                                            */
0001.08 /**********************************************/
0004.00             CALL      PGM(RFC/T8189RFC) PARM('-aT8189RFC' +
0004.01                         '-gmyhost' '-xsapgw00')
0005.00             ENDPGM
        ****************** End of data ****************************************




 F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F10=Cursor   F11=Toggle
 F16=Repeat find       F17=Repeat change          F24=More keys
```

Figure 29-35 Registering connection to ABAP program

Pass the following parameters with the CL program:

- The program-ID, as defined in Figure 29-34 on page 663
- The host name of the SAP system
- The instance number of the SAP system

Transaction SMGW monitors the registered transactions, as shown in Figure 29-36.

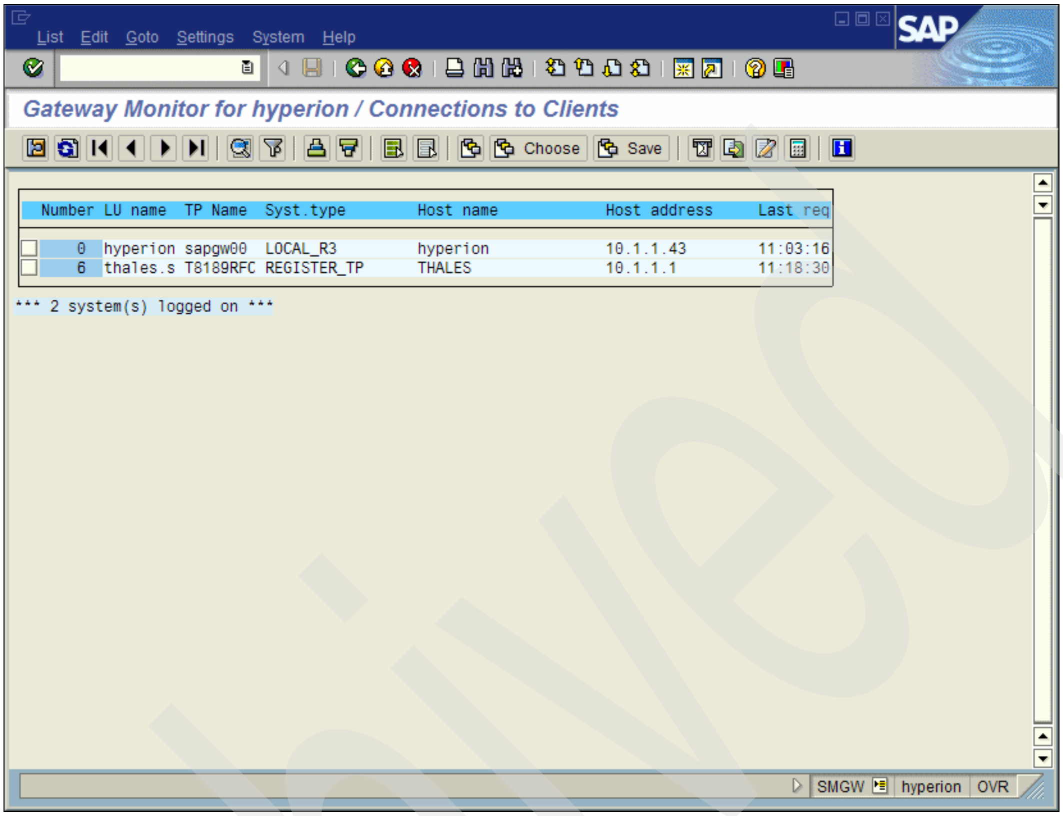

Figure 29-36 SMGW: Registered RFC destinations T8189RFC

Note: Register the communication after startup of your SAP system. Add this CL program in your autostart routine if you use an autostart program to start the SAP system after an initial program load (IPL) of the system.

ABAP program to evoke ILE or Java programs

The ABAP program named Z8189ABB calls the remote function module, which:

- Establishes a connection to the RFC destination T8189 (refer to Figure 29-34 on page 663)
- Sends and receives data (ITEMNM, ITEMD)
- Passes back the return code
- Closes the connection
- Processes the results
- Performs error recovery

The ABAP program Z8189ABB is shown in Example 29-9.

Example 29-9 Program listing: ABAP program Z8189ABB

```
REPORT Z8189ABB.
PARAMETERS: ITEMNM(5) DEFAULT '    1'.
DATA: ITEMD(25) TYPE C.
CALL FUNCTION 'Z_8189'
```

```
    DESTINATION 'T8189'
    EXPORTING
      ZITEMNM = ITEMNM
    IMPORTING
      ZITEMD = ITEMD
    EXCEPTIONS
      OTHERS = 1.
IF SY-SUBRC = 0.
  WRITE: / ITEMNM, ITEMD.
  EXIT.
ENDIF.
CASE SY-SUBRC.
  WHEN OTHERS. WRITE: / SY-SUBRC, '/Error'.
ENDCASE.
```

ILE RPG/400 example evoked by ABAP

This section shows an ILE RPG program example, which is invoked from the ABAP program Z8189ABB. See Example 29-9 on page 665. Refer to the remarks inside the ILE program source code in Example 29-10 about how to communicate with the ABAP program.

> **Note:** This ILE RPG/400 program is compatible with the ILE COBOL program example, as discussed in "ILE COBOL program example evoked by ABAP" on page 670 and the Java program example as described in "Java program example evoked by ABAP" on page 676.

Example 29-10 and Example 29-11 on page 668 list ILE program T8189RFC and the associated subroutine T8189SRV.

Example 29-10 Program listing: ILE RPG program T8189RFC

```
 **********************************************************
 * Pointer Structure for RfcAccept (argv = Dummy-Pointer) *
 *
 * This Pgm has to called always with 3 parameters as
 * follows:
* CALL PGM(T8189RFC) PARM('-a<RFC-Name>' '-g<Gateway-Host>' '-x<Gateway-Port>')
 *
 * e.g.:
 * CALL PGM(T8189RFC) PARM('-at8189rfc' '-gas0030' '-xsapgw45')
 **********************************************************
D NULL          C                CONST(x'00')
D* NULL: String Terminator
DPT             DS
D argv                         *
D PT1                          *
D PT2                          *
D PT3                          *
D PTNULL                       *   INZ(*NULL)
D**********************************************************
D* RfcEnvironment Pointer Definition
D**********************************************************
DRfc_Env           DS
DErrorHandler                   16*   PROCPTR INZ(*NULL)
DAllocate                       16*   PROCPTR INZ(*NULL)
D*********************************************************
D* RfcInstall Pointer Definitions                        *
D*********************************************************
Dins1              S              *
Dins2              S              *   PROCPTR
```

```
Dins3            S               *
D*********************************************************
D* RFC - Handle Definition                               *
D*********************************************************
DhRfc            S             8B 0
D*********************************************************
D* Returncode Definitions                                *
D*********************************************************
DRfcRc           S             8B 0
DRcen            S             8B 0
D*********************************************************
D* Defining Procedure Prototypes                         *
D* Symbolic Procedure Names are pointing to them         *
D*********************************************************
DRfcEnvironment   PR           8B 0 ExtProc('RfcEnvironment')
D Parm1                             like(Rfc_Env)
DRfcAccept       PR            8B 0 ExtProc('RfcAccept')
D argv                           *
DRfcDispatch     PR            8B 0 ExtProc('RfcDispatch')
D hRfc                         8B 0 VALUE
DRfcInst         PR            8B 0 ExtProc('RfcInstallFunction')
D ins1                           *  VALUE
D ins2                           *  PROCPTR VALUE
D ins3                           *  VALUE
DRfcClose        PR                 ExtProc('RfcClose')
D hRfc                         8B 0 VALUE
D*********************************************************
D* RFC Communication Environment parameters
D* RfcAccept (Parameter argv) is pointing to them
D*********************************************************
C     *ENTRY        PLIST
C                   PARM                    F1               32
C                   PARM                    F2               32
C                   PARM                    F3               32
C*
C* Add a Blank after the last "useful" character
C     F1            CAT       NULL:0        F1
C     F2            CAT       NULL:0        F2
C     F3            CAT       NULL:0        F3
C*********************************************************
C* Initializing Fields and Pointers
C*********************************************************
C                   MOVEL     'Z_8189'      RFCC              6
C                   MOVEL     'Z8189'       RFCC1             5
C                   EVAL      argv = %ADDR(F1)
C                   EVAL      PT1 = %ADDR(F1)
C                   EVAL      PT2 = %ADDR(F2)
C                   EVAL      PT3 = %ADDR(F3)
C*
C                   EVAL      ins1 = %ADDR(RFCC)
C                   EVAL      ins2 = %PADDR('T8189SRV')
C                   EVAL      ins3 = %ADDR(RFCC1)
C*********************************************************
C* Evoking RFC - Procedures (Service-Program LIBRFC)
C* - RfcEnvironment: Starting Environment
C* - RfcAccept : Getting RFC-Handle
C* - RfcInstallFunction: Pointing to Procedure T8189SRV
C* - RfcDispatch : Evoking Server Program T8189SRV
C* - RfcClose : Close Conversation
C*********************************************************
```

```
C* Rfc_Env - Points to *NULL
C                   CALLP     RfcEnvironment(Rfc_Env)
C* argv - Points to Entry Parameter Fields
C* hRfc - RFC - Handle
C                   EVAL      hRfc = RfcAccept(argv)
C* ins1 - Points to Remote Function Call Name Z_8189
C* ins2 - Points to RPG-Module T8189SRV (Get and Send Data)
C* ins3 - Points to RFC Destination Name Z8189
C* RfcRc - Returncode
C                   EVAL      RfcRc = RfcInst(ins1 : ins2 : ins3)
C* RfcRc - Returncode
C* hRfc - RFC - Handle
C**
C** When you want to use the "Start on"-Mode in SM59, this program
C** should only run ones and shouldn't "Re-RfcDispatch" afterwards.
C** When you want to use the "Register"-Mode in SM59, this program
C** has to restart via "Re-RfcDispatch" similar to the "Do"-Loop,
C** that is currently active.
C                   DoU       RfcRC <> 0
C                   EVAL      RfcRC = RfcDispatch(hRfc)
C                   EndDo
C* hRfc - RFC - Handle
C                   CALLP     RfcClose(hRfc)
C***********************************************************
C                   MOVE      '1'           *INLR
C                   RETURN
C***********************************************************
```

Example 29-11 lists the T8189SRV subprocedure.

Example 29-11 Program listing: RPG subprocedure T8189SRV

```
H***********************************************************
H* Defines T8189SRV as a sub-procedure
H***********************************************************
HNOMAIN
FT8189DB   IF   E           K DISK
D***********************************************************
D* Prototye Definition for Entry Parmstructure
D* Returncode: 8 Bytes Binary
D* Parameter : 8 Bytes Passed by Value
D***********************************************************
DT8189SRV         PR             8B 0
D hRfc                           8B 0 VALUE
D***********************************************************
D* Procedure Begin Definition
D***********************************************************
PT8189SRV         B                   EXPORT
D***********************************************************
D* Procedure Interface Definition
D***********************************************************
DT8189SRV         PI             8B 0
D hRfc                           8B 0 VALUE
D***********************************************************
D* Field Structure RfcGetData
D* Point : Points to Field Name (RFC Definition)
D* NLEN : Field Name Length
D* DTYPE : Data Type (ð=Char)
D* DLEN : Received Data Length
D* PTKEY : Points to Received Data
```

```
D* PTNULL : NULL-Pointer (Table Reference)
D**********************************************************
DPTGET            DS
D POINT                           *
D NLEN                           8B 0
D DTYPE                          8B 0
D DLEN                           8B 0
D PTKEY                           *
D PTNULL                          *   INZ(*NULL)
D**********************************************************
D* Field Structure RfcSendData
D* (Refer to Strucutre of RfcGetData)
D**********************************************************
DPTSEND           DS
D POINTS                          *
D NLENS                          8B 0
D DTYPES                         8B 0
D DLENS                          8B 0
D PTKEYS                          *
D PTNULLS                         *   INZ(*NULL)
D**********************************************************
DRcge             S              8B 0
D**********************************************************
D* Prototype Definitions
D**********************************************************
DRfcGet           PR             8B 0 ExtProc('RfcGetData')
D hRfc                           8B 0 VALUE
D POINT                           *
D PTNULL                          *
DRfcSend          PR             8B 0 ExtProc('RfcSendData')
D hRfc                           8B 0 VALUE
D POINTS                          *
D PTNULLS                         *
C**********************************************************
C* Initializing Pointers RfcGetData and RfcSendData
C**********************************************************
C                   EVAL      POINT = %ADDR(FNAME)
C                   EVAL      POINTS = %ADDR(FNAME1)
C                   EVAL      PTKEY = %ADDR(ITEMKEY)
C                   EVAL      PTKEYS = %ADDR(ITEMD)
C**********************************************************
C* Initializing Fields RfcGetData
C**********************************************************
C                   MOVEL     'ZITEMNM'     FNAME             7
C                   Z-ADD     7             NLEN
C                   Z-ADD     5             DLEN
C                   Z-ADD     0             DTYPE
C**********************************************************
C* Initializing Fields RfcSendData
C**********************************************************
C                   MOVEL     'ZITEMD'      FNAME1            7
C                   Z-ADD     6             NLENS
C                   Z-ADD     25            DLENS
C                   Z-ADD     0             DTYPES
C                   Z-ADD     0             Rcge
C                   MOVEL     ' '           ITEMKEY           5
C**********************************************************
C* Evoking RFC - Procedure (Service-Program LIBRFC)
C* - RfcGetData : Reads Data from ABAP Z8189ABB
C**********************************************************
```

```
C                   EVAL      Rcge = RfcGet(hRfc : POINT : PTNULL)
C**********************************************************
C                   MOVEL     *BLANKS       ITEMD
C     ITEMKEY       CHAIN     T8189DB                            99
C*
C* When not found: Send a message
C                   If        *in99 = *on
C                   Eval      ItemD = '*** Record not found ***'
C                   Endif
C**********************************************************
C* Evoking RFC - Procedure (Service-Program LIBRFC)
C* - RfcSendData : Sends Data back to ABAP Z8189ABB
C**********************************************************
C                   EVAL      Rcge = RfcSend(hRfc : POINTS : PTNULLS)
C**********************************************************
C                   RETURN    Rcge
PT8189SRV         E
```

Creating ILE RPG/400 program T8189RFC

Perform the following steps to create an executable program:

1. Use the i5/OS Create Program Module (CRTRPGMOD) command to create each of the RPG modules T8189RFC and T8189SRV:
   - `CRTRPGMOD MODULE(RFC/T8189RFC) SRCFILE(RFC/QRPGLESRC)`
   - `CRTRPGMOD MODULE(RFC/T8189SRV) SRCFILE(RFC/QRPGLESRC)`
2. Use the I5/OS Create Program (CRTPGM) command to create the program. Specify all the RPG modules created in the first step. Remember to bind the SAP service program named LIBRFC to this program:

   `CRTPGM PGM(RFC/T8189RFC) MODULE(T8189RFC T8189SRV) BNDSRVPGM(LIBRFC)`

ILE COBOL program example evoked by ABAP

The ILE COBOL program named T8189RFC and its embedded ILE module named T8189SRV are invoked from the ABAP program Z8189ABB. See Example 29-9 on page 665. Refer to the remarks inside the ILE program source code illustrated in Example 29-12 and Example 29-13 on page 673 for instructions about how to communicate with the ABAP program.

Note: This ILE COBOL program is compatible with the ILE RPG program example, as shown in “ILE RPG/400 example evoked by ABAP” on page 666 and the Java program, as described in “Java program example evoked by ABAP” on page 676.

Example 29-12 Program listing: ILE COBOL program T8189RFC

```
        IDENTIFICATION DIVISION.
        PROGRAM-ID.     T8189RFC.
       *
       * **********************************************************
       **                                                        **
       * Procedures of *SRVPGM LIBRFC
       * - RfcEnvironment  : Starting Environment
       * - RfcAccept       : Getting RFC-Handle
       *                     (hRfc needs to be DIFFERENT from 0)
       * - RfcInst         : (RfcInstallFunction)
       *                   : Provide Connection Data for Procedure T8189SRV
       * - RfcDispatch     : Call LibRfc and wait for Request from SAP
```

```
      *                    On Request, call T8189SRV internally via pointer
      * - RfcClose       : Close Conversation
      * ********************************************************
       DATE-COMPILED.
       ENVIRONMENT DIVISION.
       CONFIGURATION SECTION.
       SOURCE-COMPUTER.  IBM-AS400.
       OBJECT-COMPUTER.  IBM-AS400.
       SPECIAL-NAMES.
       INPUT-OUTPUT SECTION.
       FILE-CONTROL.
       DATA DIVISION.
       FILE SECTION.
       WORKING-STORAGE SECTION.
       77  WK-HEX-ZERO                PIC X VALUE X"00".
      * -------------------------------------------------------------*
      * Pointer structure for RfcAccept (argv = Dummy-Pointer)       *
      * -------------------------------------------------------------*
       01  PT.
        05 argv                   USAGE POINTER VALUE NULL.
        05 PT1                    USAGE POINTER VALUE NULL.
        05 PT2                    USAGE POINTER VALUE NULL.
        05 PT3                    USAGE POINTER VALUE NULL.
        05 PTNULL                 USAGE POINTER VALUE NULL.
      * -------------------------------------------------------------*
      * RfcEnvironment pointer definition                            *
      * -------------------------------------------------------------*
       77  env                   USAGE PROCEDURE-POINTER VALUE NULL.
      * -------------------------------------------------------------*
      * RfcInstall pointer definition                                *
      * -------------------------------------------------------------*
       77  ins1                  USAGE POINTER VALUE NULL.
       77  ins2                  USAGE PROCEDURE-POINTER VALUE NULL.
       77  ins3                  USAGE POINTER VALUE NULL.
      * -------------------------------------------------------------*
      * RFC - Handle definition                                      *
      * -------------------------------------------------------------*
       77  hRfc                  PIC  9(8) USAGE BINARY.
      * -------------------------------------------------------------*
      * Returncode definition                                        *
      * -------------------------------------------------------------*
       77  RfcRc                 PIC  9(8) USAGE BINARY.
       77  Rcen                  PIC  9(8) USAGE BINARY.
      * -------------------------------------------------------------*
      * Defining procedure proto types                               *
      * symbolic procdure names are pointing to them                 *
      * -------------------------------------------------------------*
       77  RfcEnvironment        USAGE PROCEDURE-POINTER VALUE NULL.
       77  RfcAccept             USAGE PROCEDURE-POINTER VALUE NULL.
       77  RfcDispatch           USAGE PROCEDURE-POINTER VALUE NULL.
       77  RfcInst               USAGE PROCEDURE-POINTER VALUE NULL.
       77  RfcClose              USAGE PROCEDURE-POINTER VALUE NULL.
      * -------------------------------------------------------------*
       77  FuncModule            PIC  X(32).
       77  RfcDest               PIC  X(32).
      * -------------------------------------------------------------*
       LINKAGE SECTION.
      * -------------------------------------------------------------*
       77  Parm1                 PIC  X(32).
       77  Parm2                 PIC  X(32).
```

```
       77  Parm3                 PIC  X(32).
      * -------------------------------------------------------------*
       PROCEDURE DIVISION  USING  Parm1 Parm2 Parm3.
      * -------------------------------------------------------------*
       050-SET-PROCEDURE-POINTERS.
           SET RfcEnvironment TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "RfcEnvironment".
           SET RfcAccept      TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "RfcAccept".
           SET RfcDispatch    TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "RfcDispatch".
           SET RfcInst        TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "RfcInstallFunction".
           SET RfcClose       TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "RfcClose".

       100-INIT-PARA.

* Build List of call arguments for C
ROB        INSPECT Parm1 REPLACING FIRST " " BY WK-HEX-ZERO
ROB        INSPECT Parm2 REPLACING FIRST " " BY WK-HEX-ZERO
ROB        INSPECT Parm3 REPLACING FIRST " " BY WK-HEX-ZERO
           SET argv       TO   ADDRESS OF Parm1.
           SET PT1        TO   ADDRESS OF Parm1.
           SET PT2        TO   ADDRESS OF Parm2.
           SET PT3        TO   ADDRESS OF Parm3.
* How can we add "hex-00" to a string ???
* LibRfc expects everything as "zero-terminated-strings" ...
   MOVE SPACE TO FuncModule
   STRING "Z_8189", WK-HEX-ZERO
   DELIMITED BY SIZE INTO FuncModule
   MOVE SPACE TO RfcDest
   STRING "Z8189", WK-HEX-ZERO
   DELIMITED BY SIZE INTO RfcDest
ROB *    MOVE "Z_8189" TO   FuncModule.
ROB *    MOVE "Z8189"  TO   RfcDest.

           SET ins1      TO   ADDRESS OF FuncModule.
           SET ins2      TO   ENTRY LINKAGE TYPE IS
                              PROCEDURE "T8189SRV".
           SET ins3      TO   ADDRESS OF RfcDest.

       200-RFC-CALL-PARA.
      *
      * - RfcEnvironment  : Starting Environment
           CALL RfcEnvironment
                USING BY REFERENCE env.
      *
      * - RfcAccept       : Getting RFC-Handle
      *                     (hRfc needs to be DIFFERENT from 0)
           CALL RfcAccept
                USING BY REFERENCE argv
                GIVING hRfc.
      *
           DISPLAY hRFC.
      *
      * - RfcInst         : (RfcInstallFunction)
      *                   : Provide Connection Data for Procedure T8189SRV
           CALL RfcInst
                USING BY VALUE ins1
```

```
                   BY VALUE ins2
                   BY VALUE ins3
             RETURNING RfcRc.
      *
      * - RfcDispatch     : Call LibRfc and wait for Request from SAP
      *                     On Request, call T8189SRV internally via pointer
           PERFORM UNTIL RfcRc NOT EQUAL 0
              CALL RfcDispatch
                   USING  BY VALUE hRFC
                   RETURNING RfcRc
           END-PERFORM.
      *
      * - RfcClose        : Close Conversation
           CALL RfcClose
                USING BY VALUE hRfc.
      *
       999-END-PARA.
           EXIT.
```

Example 29-13 lists the COBOL T8189SRV subprocedure.

Example 29-13 Program listing: COBOL subprocedure T8189SRV

```
       IDENTIFICATION DIVISION.
       PROGRAM-ID.     T8189SRV.
      *
      * ***********************************************************
       DATE-COMPILED.
       ENVIRONMENT DIVISION.
       CONFIGURATION SECTION.
       SOURCE-COMPUTER.  IBM-AS400.
       OBJECT-COMPUTER.  IBM-AS400.
       SPECIAL-NAMES.
       INPUT-OUTPUT SECTION.
       FILE-CONTROL.
           SELECT INPUT-T8189DB ASSIGN TO DATABASE-T8189DB
           ORGANIZATION IS INDEXED
           RECORD KEY IS ITEMNM WITH DUPLICATES
           ACCESS IS DYNAMIC
           FILE STATUS IS INPUT-FILE-STATUS.
       DATA DIVISION.
       FILE SECTION.
       FD INPUT-T8189DB LABEL RECORDS STANDARD.
       01 Z8189DB.
          COPY DDS-Z8189DB OF T8189DB.
       WORKING-STORAGE SECTION.
       77 ITEMKEY                PIC X(5)      VALUE SPACES.
       77 FNAME                  PIC X(7)      VALUE SPACES.
       77 FNAME1                 PIC X(7)      VALUE SPACES.
       77 INPUT-FILE-STATUS      PIC X(2)      VALUE SPACES.
       77 WS-ITEMD               PIC X(25)     VALUE SPACES.
      * -----------------------------------------------------------------*
      * Field structure RfcGetData                                       *
      * -----------------------------------------------------------------*
       01  PTGET.
        05 POINT                 USAGE POINTER VALUE NULL.
        05 NLEN                  PIC 9(8) USAGE BINARY VALUE ZERO.
        05 DTYPE                 PIC 9(8) USAGE BINARY VALUE ZERO.
        05 DLEN                  PIC 9(8) USAGE BINARY VALUE ZERO.
        05 FILLER                PIC X(4).
```

```
      05 PTKEY                 USAGE POINTER VALUE NULL.
    * Sign, that last field is reached here
      05 PTNULL                USAGE POINTER VALUE NULL.
    * -------------------------------------------------------------*
    * Field structure RfcSendData                                  *
    * -------------------------------------------------------------*
     01 PTSEND.
      05 POINTS                USAGE POINTER VALUE NULL.
      05 NLENS                PIC 9(8) USAGE BINARY VALUE ZERO.
      05 DTYPES               PIC 9(8) USAGE BINARY VALUE ZERO.
      05 DLENS                PIC 9(8) USAGE BINARY VALUE ZERO.
      05 FILLER               PIC X(4).
      05 PTKEYS                USAGE POINTER VALUE NULL.
    * Sign, that last field is reached here
      05 PTNULLS               USAGE POINTER VALUE NULL.
    * -------------------------------------------------------------*
    * Procedure pointers                                           *
    * -------------------------------------------------------------*
      77 RfcGetData            USAGE PROCEDURE-POINTER VALUE NULL.
      77 RfcSendData           USAGE PROCEDURE-POINTER VALUE NULL.
    * -------------------------------------------------------------*
     77  Rcge                  PIC 9(8) USAGE BINARY.
    * ------------------------------------------------------------*
     LINKAGE SECTION.
     77 hRfc                   PIC 9(8) USAGE BINARY.
    * ------------------------------------------------------------*
     PROCEDURE DIVISION USING BY VALUE hRfc.
    * ------------------------------------------------------------*
    *
     050-SET-PROCEDURE-POINTERS.
         SET RfcGetData  TO ENTRY LINKAGE TYPE IS
                         PROCEDURE "RfcGetData".
         SET RfcSendData TO ENTRY LINKAGE TYPE IS
                         PROCEDURE "RfcSendData".
    *
     100-INIT-PARA.
    * Name of the variable in the function module Exporting in ABAP
         MOVE "ZITEMNM"  TO   FNAME.
    * Length in bytes of the field name
         MOVE 7          TO   NLEN.
    * Length in Bytes of the data to receive
         MOVE 5          TO   DLEN.
    * Type of data to receive (0=CHAR)
         MOVE 0          TO   DTYPE.

    * Name of the variable in the function module Importing in ABAP
         MOVE  "ZITEMD"  TO   FNAME1.
    * Length in bytes of the field name
         MOVE  6         TO   NLENS.
    * Length in Bytes of the data to send
         MOVE  25        TO   DLENS.
    * Type of data to receive (0=CHAR)
         MOVE  0         TO   DTYPES.

    * Sets the pointer to the fieldnames
         SET POINT  TO ADDRESS OF FNAME.
         SET POINTS TO ADDRESS OF FNAME1
    * Sets the pointer to the real data fields
         SET PTKEY  TO ADDRESS OF ITEMKEY.
    *    SET PTKEYS TO ADDRESS OF ITEMD.
```

```
           SET PTKEYS TO ADDRESS OF WS-ITEMD.

      *  Initialize variables
           MOVE  0         TO   Rcge.
           MOVE SPACES     TO   ITEMKEY.
      *
       150-OPEN-FILES-PARA.
           OPEN INPUT INPUT-T8189DB.
           IF INPUT-FILE-STATUS NOT = "00"
              CLOSE INPUT-T8189DB
              PERFORM 999-END-PARA.
      *
       200-RFC-CALL-PARA.
      *
      * Call RFC - Procedure (Service-Program LIBRFC)
      * - RfcGetData : Reads Data from ABAP Z8189AB7
      *
           CALL RfcGetData
                USING BY VALUE hRfc
                BY REFERENCE POINT
                BY REFERENCE PTNULL
                RETURNING Rcge.
      *
      * ************************************************************
      * ** Your Program Code - START                              **
      * ************************************************************
           MOVE SPACES TO ITEMD.
           MOVE ITEMKEY TO ITEMNM.
           READ INPUT-T8189DB INVALID KEY
                MOVE "INVALID ITEM CODE" TO ITEMD.
           MOVE ITEMD  TO WS-ITEMD.
      * ************************************************************
      * ** Your Program Code - END                                **
      * ************************************************************
      *
      * Call RFC - Procedure (Service-Program LIBRFC)
      * - RfcSendData : Sends Results back to ABAP Z8189AB7
      *
           CALL RfcSendData
                USING BY VALUE hRfc
                BY REFERENCE POINTS
                BY REFERENCE PTNULLS
                RETURNING Rcge.
      *
       300-CLOSE-PARA.
           CLOSE INPUT-T8189DB.
      *
       999-END-PARA.
           EXIT.
```

Creating ILE COBOL/400 program T8189RFC

Perform the following steps to create an executable program:

1. Use the i5/OS Create Cobol Module (CRTCBLMOD) command to create each of the COBOL modules T8189RFC and T8189SRV:
   - `CRTCBLMOD MODULE(RFC/T8189RFC) SRCFILE(RFC/QRPGLESRC)`
   - `CRTCBLMOD MODULE(RFC/T8189SRV) SRCFILE(RFC/QRPGLESRC)`

2. Use the i5/OS Create Program (CRTPGM) command to create the program. Specify all the COBOL modules created in the first step. Remember to bind the SAP service program named LIBRFC to this program:

```
CRTPGM PGM(RFC/T8189RFC) MODULE(T8189RFC T8189SRV) BNDSRVPGM(LIBRFC)
```

Java program example evoked by ABAP

This section shows an Java program example, which is invoked from the ABAP program Z8189ABB (Example 29-9 on page 665). Refer to the remarks inside the program source code in Example 29-14 for information about how to communicate with the ABAP program.

> **Note:** This Java program is compatible with the ILE COBOL program example, as discussed in “ILE COBOL program example evoked by ABAP” on page 670 and the RPG program example as described in “ILE RPG/400 example evoked by ABAP” on page 666).

Example 29-14 Program listing: Java program ABAP2JAVA

```
import com.sap.mw.jco.*;
import java.util.*;
import java.sql.*;
import java.io.*;

/* Java sample program to react on incoming ABAP RFC requests         */

/* Enhanced version: Accessing an OS/400 library BRTEST via Java      */

/* This sample program creates a Java Connector Server that           *
 * listens on incoming requests from an SAP system.                   *
 *                                                                    *
 * Java then processes the request and returns a value to ABAP.       *
 *                                                                    *
 * Be sure to have RFC destination and remote function module         *
 * in place.                                                          *
 *                                                                    *
 * ABAP passes 1 parameter to the JCo server:                         *
 *    itemnm   Item number                                            *
 *                                                                    *
 * Java returns 1 parameter to the ABAP function module:              *
 *    itemd    Item description                                       */

/**
 *
 */
public class ABAP2JAVADB extends JCO.Server {
  private static String user = "GUEST";
   private static String pwd = "password";
   private static String lib = "BRTEST";
   private static String tab = "STMT";

  /* Native Driver information */
   private static String nativeConnectStr = "jdbc:db2:*LOCAL";
  private static String nativeDriverStr = "com.ibm.db2.jdbc.app.DB2Driver";

  private static Connection con;

  public ABAP2JAVADB(Properties properties, JCO.Repository mRepository) {
    super(properties, mRepository);
  }
```

```
/* Processing of incoming requests */
public void handleRequest(JCO.Function function) {
 /* Get input parameter */
  String zitem  = function.getImportParameterList().getString("ZITEMNM");

  System.out.println("handleRequest(" + function.getName() + ")");

  try {
  /* Respond to the right remote function module */
    if (function.getName().equals("Z_8190")) {
    /* Process data */
    String zitemd = returnDescription(con, zitem, lib, tab);

    /* Set output parameter */
    function.getExportParameterList().setValue(zitemd, "ZITEMD");
    }
  }
  catch(JCO.AbapException ex) {
    throw ex;
  }
  catch(Throwable t) {
    throw new JCO.AbapException("SYSTEM_FAILURE", t.getMessage());
  }
}

public void startServer() {
  try {
    this.setTrace(true);

    /* Do all necessary RFC handling, like registering */
    this.start();
  }
  catch (Exception ex) {
    System.out.println("Could not start server " + this.getProgID() + ":\n" + ex);
  }
}

/* Take care for the Database connection */
private static Connection getConnection() throws Exception {
  Class.forName(nativeDriverStr);
  con = DriverManager.getConnection(nativeConnectStr, user, pwd);
  return con;
}

/* Recreate an arbitrary test table */
 private static void createTable (Connection con, String lib, String tab, String
statement) throws Exception {
     String crtstmt = null;
     String primkeystmt = null;
     String dropstr = "";
     Statement stmt = null;
     stmt = con.createStatement();

     try {
          dropstr = "DROP TABLE " + lib + "." + tab;
          System.out.println(dropstr);
          stmt.executeUpdate(dropstr);
     }
     catch (Exception e) {}
```

```
        try {
             crtstmt = "CREATE TABLE " + lib + "." + tab + statement;
             System.out.println(crtstmt);
             stmt.executeUpdate(crtstmt);

            primkeystmt = "ALTER TABLE " + lib + "." + tab + " ADD PRIMARY KEY (ZITEMNM)";
             System.out.println(primkeystmt);
             stmt.executeUpdate(primkeystmt);
        }
        catch (Exception e) {
           System.out.println("Error creating table. " + e.toString());
        }
     }

    /* Insert some arbitrary test data (The three muses) */
    public static void prepare(Connection con, String lib, String tab) throws SQLException {
        String ins = "INSERT INTO " + lib + "." + tab + " VALUES (?,?)";
        PreparedStatement stmt = con.prepareStatement(ins);
        stmt.setString(1, "00001");
        stmt.setString(2, "Melete");
        stmt.executeUpdate();

        stmt.setString(1, "00002");
        stmt.setString(2, "Mneme");
        stmt.executeUpdate();

       stmt.setString(1, "00003");
        stmt.setString(2, "Aoede");
        stmt.executeUpdate();

       stmt.close();
    }

    /* Retrieve requested data from database */
    public static String returnDescription(Connection con, String key, String lib, String
  tab) throws SQLException {
     String sel = "SELECT * FROM " + lib + "." + tab + " WHERE ZITEMNM = ?";
      PreparedStatement pstmt = con.prepareStatement(sel);
      String result;

      pstmt.setString(1, key);
      ResultSet rs = pstmt.executeQuery();

      rs.next();
      result = rs.getString(2);

      rs.close();
      pstmt.close();
      return result;
    }

    public static void main (String args[]) throws Exception {
      JCO.Repository mRepository;
      String POOL = "MY_CLIENT_POOL";

      /* Provide logon data */
      String client = "001";
      String sapuser = "ddic";
      String sappasswd = "sappaswrd";
      String lang = "";
```

```
    String ashost = "host";
    String sysnr = "78";

    try {
      /* Definition of RFC connection */
      JCO.addClientPool(POOL, 5, client, sapuser, sappasswd, lang, ashost, sysnr);

      /* Definition of metadata for RFC */
      mRepository = new JCO.Repository("JAVA2ABAP", POOL);
    }
    catch (Exception ex) {
      ex.printStackTrace();
      throw ex;
    }

   String gwhost = "gwhost";
   String gwserv = "sapgw78";
   String program_id = "T8189RFC";

   Properties properties = new Properties();
    properties.setProperty("jco.server.gwserv", gwserv);
    properties.setProperty("jco.server.gwhost", gwhost);
    properties.setProperty("jco.server.progid", program_id);

    /* Only necessary, if communicating SAP System is Unicode */
    properties.setProperty("jco.server.unicode", "1");

    Connection con = getConnection();

     /* GRAPHIC(x) CCSID 13488: Unicode */
     String halfcrtstmt = "(ZITEMNM GRAPHIC(5) CCSID 13488, " +
                          "ZITEMD GRAPHIC(26) CCSID 13488)";

    /* Create arbitrary test table */
    createTable(con, lib, tab, halfcrtstmt);

    /* Insert arbitrary data */
    prepare(con, lib, tab);

    ABAP2JAVADB app = new ABAP2JAVADB(properties, mRepository);
    app.startServer();
  }
}
```

Starting and registering Java program ABAP2JAVA

Perform the following steps to start the Java program and to register the RFC connection:

- `qsh`
- `cd /JCo install path`
- `./runExample2.sh`

Example 29-15 Script to set the environment and evoke ABAP2JAVA: runExample2

```
export LIBPATH=/JCo install path
export CLASSPATH=/JCo install path
export QIBM_JAVA_PASE_STARTUP=/usr/lib/start64
/QIBM/ProdData/Java400/jdk14/bin/java -classpath /JCo install path/sapjco.jar ABAP2JAVA
```

29.7.3 ILE RPG or Java program calling ABAP: Flowchart of RFC connection

This section shows an example of an ILE RPG/400 and Java programs triggering an ABAP program, which starts the report RDDIMPDP in the SAP system. In our example, this is the function module SUBST_START_BATCHJOB. See Figure 29-37 for an illustration of the RFC flow chart. Both the RPG/400 and Java programs perform the following steps.

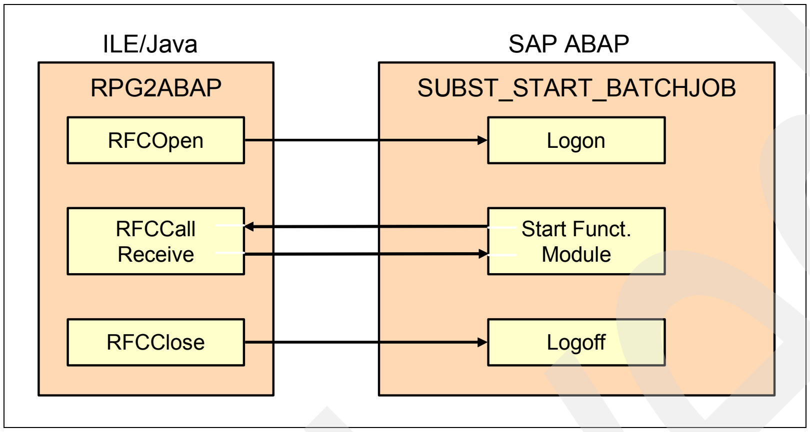

Figure 29-37 Flowchart: RPG program RPG2ABAP

The example program named RPG2ABAP (shown in Example 29-16) and JAVA2ABAP (shown in Example 29-17 on page 686) perform the following steps:

1. Log into the SAP system
2. Call the SAP functional module, which starts reports in the SAP system. In this example, it is the SAP standard functional module SUBST_START_BATCHJOB.
3. Pass the following parameters:
   - Job name (named RPG2ABAP for RPG and JAVA2ABAP for Java in this example)
   - Name of the report (named RDDIMPDP in this example)
4. Receive the following parameters:
   - Internal SAP job ID where the report is running in
   - Return code of the report RDDIMPDP

ILE RPG program example

This section shows the ILE RPG/400 program RPG2ABAP, which starts an ABAP program in the SAP system, see Example 29-16.

Example 29-16 Program listing: RPG program RPG2ABAP

```
**********************************************************************
** ILE RPG sample program to call an ABAP function module via RFC
**********************************************************************
** This sample program invokes the ABAP function module
** SUBST_START_BATCHJOB (see SE37)
** to start an R/3 background job named RPG2ABAP (see SM37).
**
** The R/3 background job executes the ABAP program RDDIMPDP.
** The ABAP function module SUBST_START_BATCHJOB returns immediately.
**
** RPG passes 2 parameters to the ABAP function module:
**   JOBNAME    Name of the R/3 background job
**   REPNAME    Name of the ABAP report to be executed
**
```

```
** RPG receives 1 parameter from the ABAP function module:
**   JOBCOUNT   Id of the submitted R/3 background job
**
*********************************************************************
** Bind the SAP supplied service program LIBRFC to this module:
*********************************************************************

*********************************************************************
** Some constants
*********************************************************************
D NULL            C                    CONST(x'00')
terminator

D TYPC            C                    CONST(0)                        saprfc.h
D TYPDATE         C                    CONST(1)
D TYPP            C                    CONST(2)
D TYPTIME         C                    CONST(3)
D TYPX            C                    CONST(4)
D TYPTABH         C                    CONST(5)
D TYPNUM          C                    CONST(6)
D TYPFLOA         C                    CONST(7)
D TYPINT          C                    CONST(8)
D TYPINT2         C                    CONST(9)
D TYPINT1         C                    CONST(10)
D TYP1            C                    CONST(12)
D TYP2            C                    CONST(13)

***************************************************************************
** Data structure similar to C structure RFC_ENV in saprfc.h
***************************************************************************
DRFC_ENV          DS                   ALIGN
D ERRORHANDLER                  16*    PROCPTR INZ(*NULL)
D ALLOCATE                      16*    PROCPTR INZ(*NULL)

***************************************************************************
** Procedure prototype for C function exported from LIBRFC. See saprfc.h
***************************************************************************
DRfcEnvironment   PR                   EXTPROC('RfcEnvironment')
D PARM1                                LIKE(RFC_ENV)

***************************************************************************
** Data structure similar to C structure RFC_CONNOPT_R3ONLY in saprfc.h
***************************************************************************
DRFC_ERROR_INFO   DS                   ALIGN
D KEY                           33A
D STATUS                       128A
D MESSAGE                      256A
D INTSTAT                      128A

***************************************************************************
** Procedure prototype for C function exported from LIBRFC. See saprfc.h
***************************************************************************
DRfcLastError     PR            10I 0 EXTPROC('RfcLastError')
D PARM1                               LIKE(RFC_ERROR_INFO)

***************************************************************************
** Data structure similar to C structure RFC_CONNOPT_R3ONLY in saprfc.h
***************************************************************************
D RFCHOSTNAME     S            257A
D RFCGATEWAYHOST  S            257A
```

```
D RFCGATEWAYSRV   S              9A

DRFC_CONNOPT      DS                  ALIGN
D HOSTNAME                       *    INZ(%ADDR(RFCHOSTNAME))
D SYSNR                        10I 0  INZ(0)
D GATEWAYHOST                    *    INZ(%ADDR(RFCGATEWAYHOST))
D GATEWAYSRV                     *    INZ(%ADDR(RFCGATEWAYSRV))

*****************************************************************************
** Data structure similar to C structure RFC_OPTIONS in saprfc.h
*****************************************************************************
D RFCSID          S              4A
D RFCCLIENT       S              4A
D RFCUSER         S              9A
D RFCPASSWORD     S              9A
D RFCLANG         S              3A

DRFC_OPTIONS      DS                  ALIGN
D DESTINATION                    *    INZ(%ADDR(RFCSID))
D MODE                         10I 0  INZ(0)
D CONNOPT                        *    INZ(%ADDR(RFC_CONNOPT))
D CLIENT                         *    INZ(%ADDR(RFCCLIENT))
D USER                           *    INZ(%ADDR(RFCUSER))
FMT D  DName+++++++++++ETDsFrom+++To/L+++IDc.Keywords+++++++++++++++++++++++++++++Comments+
D PASSWORD                       *    INZ(%ADDR(RFCPASSWORD))
D LANGUAGE                       *    INZ(%ADDR(RFCLANG))
D TRACE                        10I 0  INZ(0)

*****************************************************************************
** Procedure prototype for C function exported from LIBRFC. See saprfc.h
*****************************************************************************
DRfcOpen          PR           10I 0  EXTPROC('RfcOpen')
D PARM1                               LIKE(RFC_OPTIONS)

*****************************************************************************
** Return variable of RfcOpen. See saprfc.h
*****************************************************************************
D RFC_HANDLE      S            10U 0  INZ(0)

*****************************************************************************
** Procedure prototype for C function exported from LIBRFC. See saprfc.h
*****************************************************************************
DRfcClose         PR                  EXTPROC('RfcClose')
D PARM1                               LIKE(RFC_HANDLE) VALUE

*****************************************************************************
** Variables for names and values passed to the R/3 function module
*****************************************************************************
D VARNAME1        S             10A
D VARNAME2        S             10A
D VARNAME3        S             10A
D VARNAME4        S             10A
D JOBNAME         S             32A
D REPNAME         S             40A
D JOBCOUNT        S              8A
D SySubRC         S             10I 0

*****************************************************************************
** Data structure similar to C structure RFC_PARAMETER in saprfc.h
** This array is used to pass parameters from RPG to ABAP
```

```
 ****************************************************************************
DRFC_EXPORTPARM   DS                 ALIGN OCCURS(15)
D ENAME                        *
D ENLEN                      10U 0
FMT D  DName+++++++++++ETDsFrom+++To/L+++IDc.Keywords+++++++++++++++++++++++++++++Comments+
D ETYPE                       10U 0
D ELENG                       10U 0
D EADDR                         *

 ****************************************************************************
** Data structure similar to C structure RFC_PARAMETER in saprfc.h
** This array is used to return parameters from ABAP to RPG
 ****************************************************************************
DRFC_IMPORTPARM   DS                  ALIGN OCCURS(15)
D INAME                         *
D INLEN                       10U 0
D ITYPE                       10U 0
D ILENG                       10U 0
D IADDR                         *

 ****************************************************************************
** Data structure similar to C structure RFC_TABLE in saprfc.h
** This is not used use in this sample program.
 ****************************************************************************
DRFC_TABLE        DS                  ALIGN OCCURS(15)
D TNAME                         *   INZ(*NULL)
D TNLEN                       10U 0 INZ(0)
D TTYPE                       10U 0 INZ(0)
D TLENG                       10U 0 INZ(0)
D TITHANDLE                     *   INZ(*NULL)
D TITMODE                       *   INZ(*NULL)
** 'NewITab' is wrong, even when it is sometimes documentated like this
**TNEWITAB                    10I 0 INZ(0)

D RFC_FUNCTION    S           31A
D RFC_EXCEPTION   S             *   INZ(*NULL)
D RC              S           10I 0 INZ(0)
D SYSNRA          S            2A

 ****************************************************************************
** Procedure prototype for C function exported from LIBRFC. See saprfc.h
 ****************************************************************************
DRfcCallReceive   PR          10I 0 EXTPROC('RfcCallReceive')
D PARM1                             LIKE(RFC_HANDLE) VALUE
D PARM2                             LIKE(RFC_FUNCTION)
D PARM3                            LIKE(RFC_EXPORTPARM)
D PARM4                            LIKE(RFC_IMPORTPARM)
D PARM5                            LIKE(RFC_TABLE)
D PARM6                            LIKE(RFC_EXCEPTION)

 ****************************************************************************
** Call C function to initialize the RFC environment
 ****************************************************************************
C                   CALLP     RfcEnvironment(RFC_ENV)

 ****************************************************************************
** Provide logon data.
 ****************************************************************************
C     'I47'         CAT       NULL          RFCSID
C     '000'         CAT       NULL          RFCCLIENT                      Client
```

```
C     'DDIC'        CAT       NULL          RFCUSER                      User
C     'SAPADM01'    CAT       NULL          RFCPASSWORD                  Password
C     'D'           CAT       NULL          RFCLANG                      Language

C     'hyperion'    CAT       NULL          RFCHOSTNAME                  Hostname
C                   EVAL      SYSNR = 00
C*
C                   Movel     SYSNR         SYSNRA
C                   Eval      RFCGATEWAYHOST = RFCHOSTNAME
C                   Eval      RFCGATEWAYSRV  = 'sapgw' + SYSNRA + NULL

***************************************************************************
** Call C function to open the RFC connection to R/3
***************************************************************************
C                   EVAL      RFC_Handle = 999
C                   EVAL      RFC_HANDLE = RfcOpen(RFC_OPTIONS)
C     RFC_HANDLE    IFEQ      0
 * ERROR HANDLING
C                   RETURN
C                   ENDIF

***************************************************************************
** Provide function module parameters.
***************************************************************************
C                   EVAL      RFC_FUNCTION = 'SUBST_START_BATCHJOB' + NULL

** Parameters that are exported by RPG and imported by ABAP
C     1             Occur     RFC_ExportParm
C                   EVAL      VARNAME1 = 'JOBNAME'
C                   EVAL      ENAME    = %ADDR(VARNAME1)
C                   EVAL      ENLEN    = 7
C                   EVAL      ETYPE    = TYPC
C                   EVAL      JOBNAME  = 'RPG2ABAP'
C                   EVAL      ELENG    = 32
C                   EVAL      EADDR    = %ADDR(JOBNAME)

C     2             Occur     RFC_ExportParm
C                   EVAL      VARNAME2 = 'REPNAME'
C                   EVAL      ENAME    = %ADDR(VARNAME2)
C                   EVAL      ENLEN    = 7
C                   EVAL      ETYPE    = TYPC
C                   EVAL      REPNAME  = 'RDDIMPDP'
C                   EVAL      ELENG    = 40
C                   EVAL      EADDR    = %ADDR(REPNAME)

C     3             Occur     RFC_ExportParm
C                   EVAL      ENAME  = *NULL

** Parameters that are exported by ABAP and imported by RPG
C     1             Occur     RFC_ImportParm
C                   EVAL      VARNAME3 = 'JOBCOUNT'
C                   EVAL      INAME    = %ADDR(VARNAME3)
C                   EVAL      INLEN    = 8
C                   EVAL      ITYPE    = TYPC
C                   EVAL      JOBCOUNT = 'abcdedfg'
C                   EVAL      ILENG    = 8
C                   EVAL      IADDR    = %ADDR(JOBCOUNT )

C     2             Occur     RFC_ImportParm
C                   EVAL      VARNAME4 = 'RC_START'
```

```
C                   EVAL      INAME    = %ADDR(VARNAME4)
C                   EVAL      INLEN    = 8
C                   EVAL      ITYPE    = TYPC
C                   EVAL      SySubRC  = 987
C                   EVAL      ILENG    = 8
C                   EVAL      IADDR    = %ADDR(SySubRC  )

C     3             Occur     RFC_ImportParm
C                   EVAL      INAME  = *NULL

** Parameters that are exported by RPG to ABAP by table
** (Here are  N O  parameters exported)
C     1             Occur     RFC_Table
C                   EVAL      TNAME  = *NULL


****************************************************************************
** Call C function to call the ABAP function module
****************************************************************************
 * The arrays have to be at position 1, in orderto give the pointer to the next
 * function correct
C     1             Occur     RFC_ExportParm
C     1             Occur     RFC_ImportParm
C     1             Occur     RFC_Table

C                   EVAL      RC = RfcCallReceive(RFC_HANDLE     :
C                                                 RFC_FUNCTION   :
C                                                 RFC_EXPORTPARM :
C                                                 RFC_IMPORTPARM :
C                                                 RFC_TABLE      :
C                                                 RFC_EXCEPTION)

****************************************************************************
** Error handling
****************************************************************************
C     RC            IFNE      0
C                   EVAL      RC = RfcLastError(RFC_ERROR_INFO)
C     RC            IFEQ      0
C                   DSPLY                   KEY
 * Variable MESSAGE contains error text
C                   ENDIF
C                   ENDIF

****************************************************************************
** Close the connection

************************************************************************************
C                 CALLP     RfcClose(RFC_HANDLE)

C                   RETURN
```

Creating ILE RPG/400 program T8189RFC

Perform the following steps to create an executable program:

1. Use the i5/OS Create RPG Module (CRTRPGMOD) command to create the RPG module named RPG2ABAP:

   ```
   CRTRPGMOD MODULE(RFC/RPG2ABAP) SRCFILE(RFC/QRPGLESRC)
   ```

2. Use the i5/OS Create Program (CRTPGM) command to create the program. Specify the RPG module created in step 1 on page 685. Bind the SAP service program named LIBRFC to this program:

```
CRTPGM PGM(RFC/RPG2ABAP) MODULE(RPG2ABAP) BNDSRVPGM(R3640AOPT/LIBRFC)
```

Java program example

This section lists the Java program JAVA2ABAP, which starts an ABAP program in the SAP system, see Example 29-17.

Example 29-17 Program listing: JAVA2ABAP

```
import com.sap.mw.jco.*;

/* Java sample program to call an ABAP function via RFC              */

/* This sample program invokes the ABAP function module              *
 * SUBST_START_BATCHJOB (see SE37)                                   *
 * to start an SAP background job named JAVA2ABAP (see SM37).        *
 *                                                                   *
 * The SAP background job executes the ABAP program RDDIMPDP.        *
 * The ABAP function module SUBST_START_BATCHJOB returns immediately. *
 *                                                                   *
 * Java passes 2 parameters to the ABAP function module:             *
 *    jobname   Name of the SAP background job                       *
 *    repname   Name of the ABAP report to be executed               *
 *                                                                   *
 * Java receives 2 parameters from the ABAP function module:         *
 *    jobcount  ID of the submitted SAP background job               *
 *    rc_start  Return code of function module                       */

/**
 *
 */
public class JAVA2ABAP extends Object {

  JCO.Repository mRepository;
  static String POOL = "MY_CLIENT_POOL";

  public JAVA2ABAP() throws Exception
  {
    /* Provide logon data */
    String client = "001";
    String user = "sapuser";
    String passwd = "sappaswrd";
    String lang = "";
    String ashost = "host";
    String sysnr = "78";

    try {
      /* Definition of RFC connection */
      JCO.addClientPool(POOL, 5, client, user, passwd, lang, ashost, sysnr);

      /* Definition of metadata for RFC */
      mRepository = new JCO.Repository("JAVA2ABAP", POOL);
    }
    catch (Exception ex) {
      ex.printStackTrace();
      throw ex;
    }
```

```
}

private void callSUBST_START_BATCHJOB() throws Exception
{
  /* Provide function module parameters */
  JCO.Function function = null;
  String rfcFunction  = "SUBST_START_BATCHJOB";
  String jobname = "JAVA2ABAP";
  String repname = "RDDIMPDP";

  String jobcount; // CHAR 8
  int rc_start; // INT 4

  JCO.Client connection = JCO.getClient(POOL);
  try {
    function = this.createFunction(rfcFunction);
    if (function == null) {
      throw new Exception(rfcFunction + " not found in SAP.");
    }

    function.getImportParameterList().setValue(jobname, "JOBNAME");
    function.getImportParameterList().setValue(repname, "REPNAME");

    /* Call ABAP function module */
    connection.execute(function);

    /* Print result */
    jobcount = function.getExportParameterList().getString("JOBCOUNT");
    System.out.println("Job name: " + jobcount);
    rc_start = function.getExportParameterList().getInt("RC_START");
    System.out.println("Return code: " + rc_start);
  }
  catch (Exception ex)
  {
    ex.printStackTrace();
    throw ex;
  }
  finally
  {
      /* Release client */
      JCO.releaseClient(connection);
  }
}

private void cleanUp()
{
  JCO.removeClientPool(POOL);
}

public JCO.Function createFunction(String name) throws Exception {
  try {
    IFunctionTemplate ft = mRepository.getFunctionTemplate(name.toUpperCase());
    if (ft == null)
      return null;
    return ft.getFunction();
  }
  catch (Exception ex) {
    throw new Exception("Problem retrieving JCO.Function object.");
  }
}
```

```
  public static void main (String args[]) throws Exception {
    JAVA2ABAP app = new JAVA2ABAP();
    app.callSUBST_START_BATCHJOB();
    app.cleanUp();
  }
}
```

Starting Java program JAVA2ABAP

Perform the following steps to start the Java program:

```
qsh
cd /JCo install path
./runExample.sh
```

Example 29-18 Script to set the environment and evoke JAVA2ABAP: runExample

```
export LIBPATH=/JCo install path
export CLASSPATH=/JCo install path
export QIBM_JAVA_PASE_STARTUP=/usr/lib/start64
/QIBM/ProdData/Java400/jdk14/bin/java -classpath /JCo install path/sapjco.jar JAVA2ABAP
```

29.7.4 Monitoring communication status

Use SAP transaction SM37 to monitor the status of the RFC jobs, as shown in Figure 29-38.

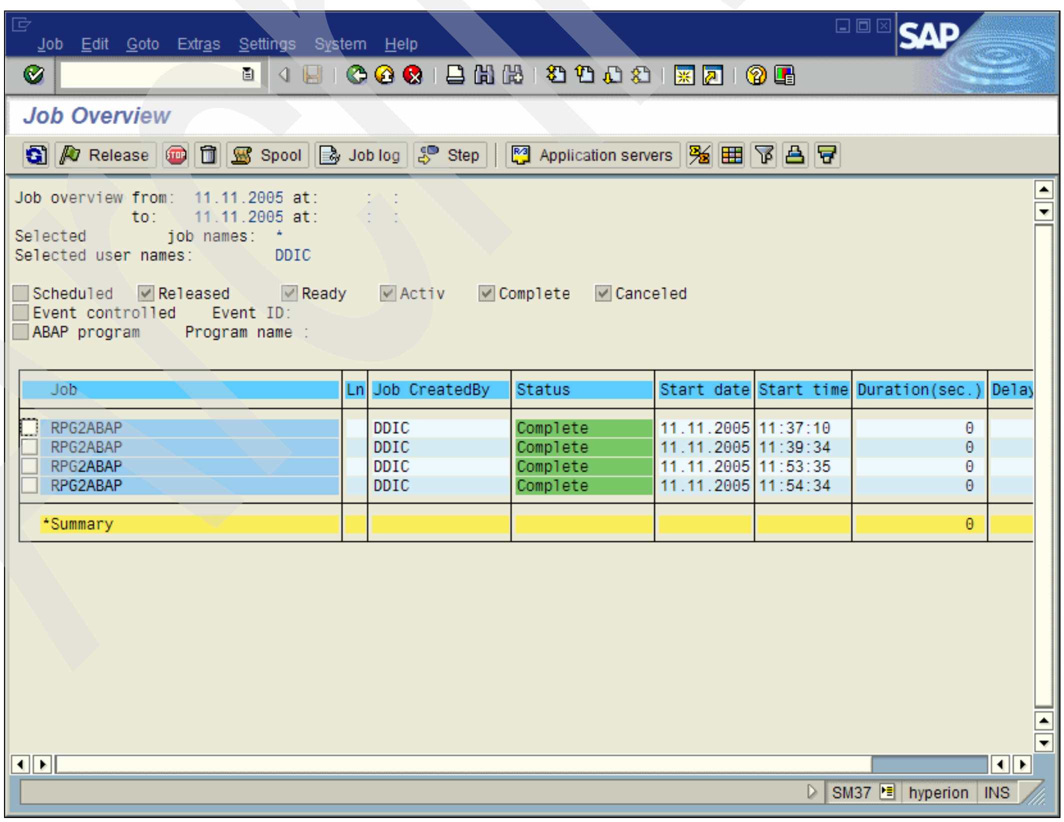

Figure 29-38 SM37: Monitoring status of SAP jobs

30

Data porting

This chapter describes data porting from non-SAP applications into an SAP system environment. It discusses the steps necessary to perform data porting and provides examples of using the data porting tools.

See Chapter 29, “Data exchange scenarios and examples” on page 621 for techniques to enable SAP systems running on the System i server to interact with System i non-SAP applications.

30.1 Concept of data porting to SAP applications

When you implement new application software such as SAP applications, one of the major tasks is to migrate the data from existing applications into the new environment. This activity is typically performed only once. Normally, it is not necessary that existing applications run in addition to the SAP application. Therefore, data can be exchanged between the two environments frequently. This activity is called *data porting*.

The implementation of an SAP system normally begins with the installation of the new software on a development or test system. The 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 customer's requirement. Sometimes, an additional specific function must be developed. The test system with its parameters, additional functions, programs, and data files are transferred to the new production environment after the test is completed.

SAP applications provide a facility to transfer the existing application data into the SAP database using a function called *batch input*. This function reads the data from a sequential text database file with fixed format and stores it into the SAP database using normal SAP 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 required by the batch input program.

This sequential file is used as a transfer medium between the existing application and the SAP database. The client is responsible for producing this sequential file by using a data transfer program that reads the existing data file, checks and converts it, and exports it into the sequential file. A user-written data transfer program or data porting tools can be used for this purpose.

The data porting concept is shown in Figure 30-1.

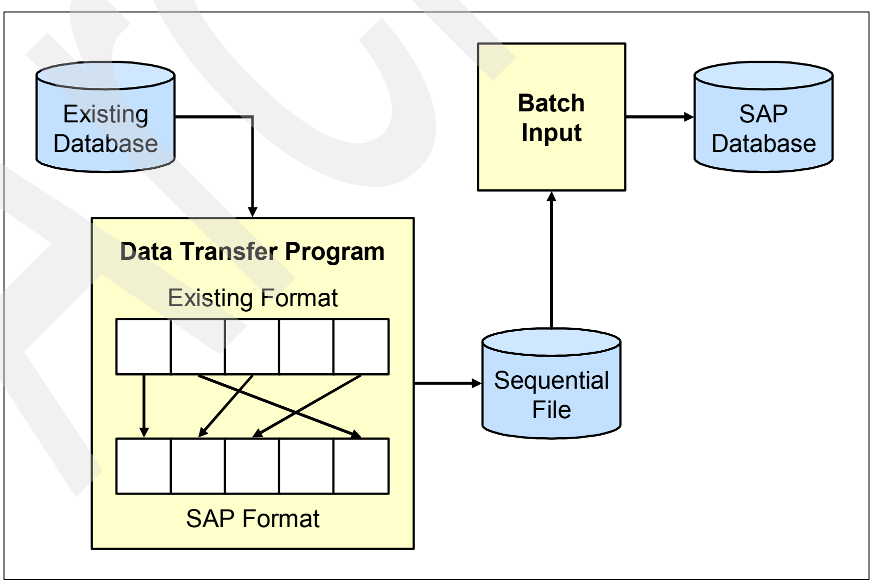

Figure 30-1 Concept of data porting to an SAP system

30.2 Programming aspects for data porting to an SAP system

Data transfer programs and batch input programs are required to perform data porting to an SAP system environment. This section discusses these programs.

30.2.1 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 required by the batch input program.

The data transfer program performs the following tasks:

- Checks to see whether 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 system 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 applications provide a facility to generate a data structure for SAP tables in the Common Business Oriented Language (COBOL), PL/I, or C languages. You can incorporate this data structure into your data transfer program. Select **Environment** and select the **Generate table description** in the Advanced Business Application Programming (ABAP) dictionary to generate the data structure.
2. Analyze the SAP system 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:
   a. The data field must be in character format.
   b. No data field is longer than those in the SAP system format.
   c. 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 or Report Program Generator (RPG).

30.2.2 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. The system runs the transaction automatically when you start batch input. Therefore, it is 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 method so that it ensures data integrity.

The batch input program is written in ABAP and 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 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 within the SAP-specific ABAP programming language. Then, you can either explicitly start and monitor the session or have the session run in the background. To do this, in the SAP system environment choose **System** → **Services** → **Batch Input** → **Sessions** or call transaction SM35.

  Use the BDC_OPEN_GROUP function module to create a new session. After you have created a session, you can insert batch input data into it using BDC_INSERT. Use the BDC_CLOSE_GRUP function module to close a session after you have inserted all of your batch input data into it. After a session is closed it can be processed.

- The program uses the CALL_TRANSACTION_USING statement to run the SAP transaction.

  With CALL_TRANSACTION_USING, the system processes the data faster than with batch input sessions. Unlike batch input sessions, CALL_TRANSACTION_USING does not automatically support interactive correction or logging functions.

  Your program prepares the data and then calls the corresponding transaction that is then processed immediately.

- The program uses the CALL_DIALOG statement. Your program prepares data for a sequence of dialog screens, and calls a dialog module for immediate processing.

  We do *not* recommend this method because it is outdated and more complex than the other methods.

Each of the preceding methods uses the SAP system 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 Batch Input is an SAP system technique that allows automating the input in transactions. It relies on a Batch Data Command (BDC) scenario. The following statement is used to declare the SAP system 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 30-2 on page 693 shows the batch input technique.

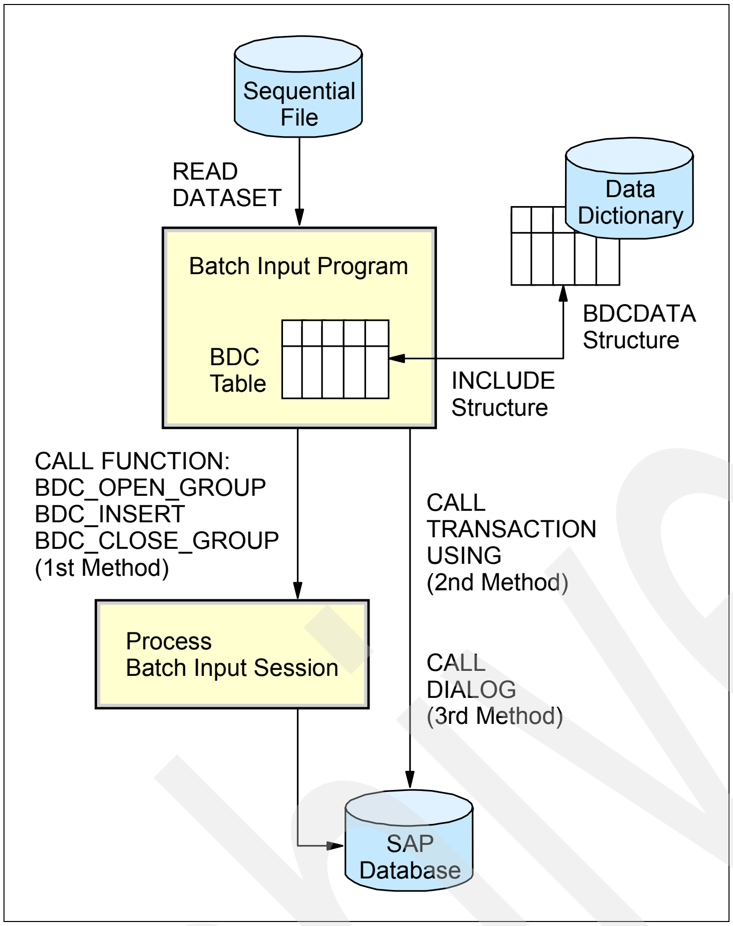

Figure 30-2 Batch input technique

SAP applications provide a ready-to-run standard batch input program for most of the SAP applications. You can also create your own batch input program, if necessary. To do this, analyze the SAP transaction as shown in the following steps:

1. Go through the application function just as you do in a normal transaction.
2. Select **System** → **Status** and note the program names and display (dynpro) number.
3. Press **F1** (Help) on the field, check box, or button, and choose **Technical Info**. Note any field names, check boxes, or buttons that require input.
4. Note the display (dynpro) sequence and function codes.
5. Create a BDC table structure.

   The structure of BDCDATA is described in Example 30-1:

Example 30-1 BDCDATA structure

```
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. 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 a 3-tier environment, run the Batch Input jobs on the same System i server where the sequential input file is located. The sequential input file comes from the System i Integrated File System. Because the input data has to be updated in the SAP database, the fastest situation for the data load is when you run both the Batch Input jobs and the related sequential input files on the SAP database server.

With this configuration there is a significant difference in the run times if the input Integrated File System file and the log for the job are physically located on the local server that the job is running on. This setup helps avoid delays during frequent access to these Integrated File System files.

30.3 Data porting services

IBM can assist you in connecting to data porting services through the iSeries Porting Team in PartnerWorld® for Developers. Refer to the Web site at:

http://www.ibm.com/servers/enable/site/porting/

Information about non-IBM porting services and tools from System i Business Partners is available on the Web at:

http://www.ibm.com/software/

More information for porting services from the SAP partner portal is available at:

http://www.sap.com/partners/index.epx

30.4 Data porting tools

Specific one-time-use programs can be developed for data porting purposes when you start implementing an SAP system in a core business environment. You can 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. In many cases, using such a tool can result in more efficient and faster data porting compared to developing many one-time-use programs.

There are some data porting tools available on the market today. This section focuses on the official SAP data porting tool Legacy System Migration Workbench (LSMW).

30.4.1 Before working with LSMW

Before you use the LSMW, we recommend that you perform the following tasks:

- Ensure that the SAP system customization is complete.

> **Note:** We know that data porting is often started for project reasons before the SAP customizing is complete. So the definition and the activities for the data porting are iterative processes if the SAP customizing changed between the start and the final data load.

- Analyze the data that exists in the existing system (a company's core business applications) and which data (from a business point of view) will be required in the future.
- Identify the transactions that you want to use to transfer the data into the SAP system.
- Process the respective transaction in the SAP system manually using test data from the core business applications and note which fields must be filled. There may be mandatory fields that do not correspond to a data field in the existing system. In this case, assigning a fixed value or constant or setting up a field as an optional field can be useful for the data transfer.

  See Chapter 29, "Data exchange scenarios and examples" on page 621 for a description of iterative data loads onto an SAP system.
- Map the fields on paper by assigning the source fields to the SAP application fields (target fields). You can do this by using the object overview after you define the object.
- Determine the conversion rule according to which the contents of the source fields are to be converted into SAP application fields.
- Define the way in which the data is to be extracted from the non-SAP system. Note that LSMW does not extract data itself.
- Check the format of the core business data. Decide which standard import technique you want to use, or whether you need to define an extra *object class* by means of recording.
- In case only a part of your existing system is to be replaced with the SAP system, you have to determine the functions that are to be covered by an SAP application and those that are to be covered by the core business applications. If necessary, create a concept of the data structure and data flow to define
  - Where which data must be available in which timeliness
  - Which interfaces must maintained or new created

30.4.2 Legacy System Migration Workbench

The LSMW is an SAP-based tool that supports the one-time or periodic transfer of data from non-SAP or core business systems to an SAP system. It uses standard SAP interfaces.

The LSMW 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 LSMW. Reusable conversion rules assure consistent data conversion for different data objects.

Figure 30-3 on page 696 shows the steps in the data migration.

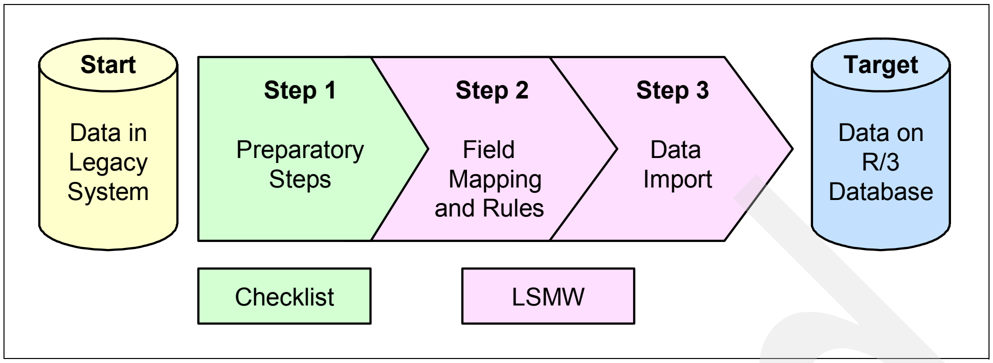

Figure 30-3 Data migration in three steps

The LSMW performs the following tasks:

- Reads the existing 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, (Business Application Programming Interface (BAPI), and SAP Intermediate Document (IDoc)).

  IDoc (for intermediate document) is a standard data structure for electronic data interchange (**EDI**) between application programs written for the popular SAP business system or between an SAP application and an external program. IDocs serve as the vehicle for data transfer in SAP's Application Link Enabling (**ALE**) system. IDocs are used for asynchronous transactions: each IDoc generated exists as a self-contained text file that can then be transmitted to the requesting workstation without connecting to the central database. Another SAP mechanism, the Business Application Programming Interface (**BAPI**) is used for synchronous transactions.

See Figure 30-4 for an illustration.

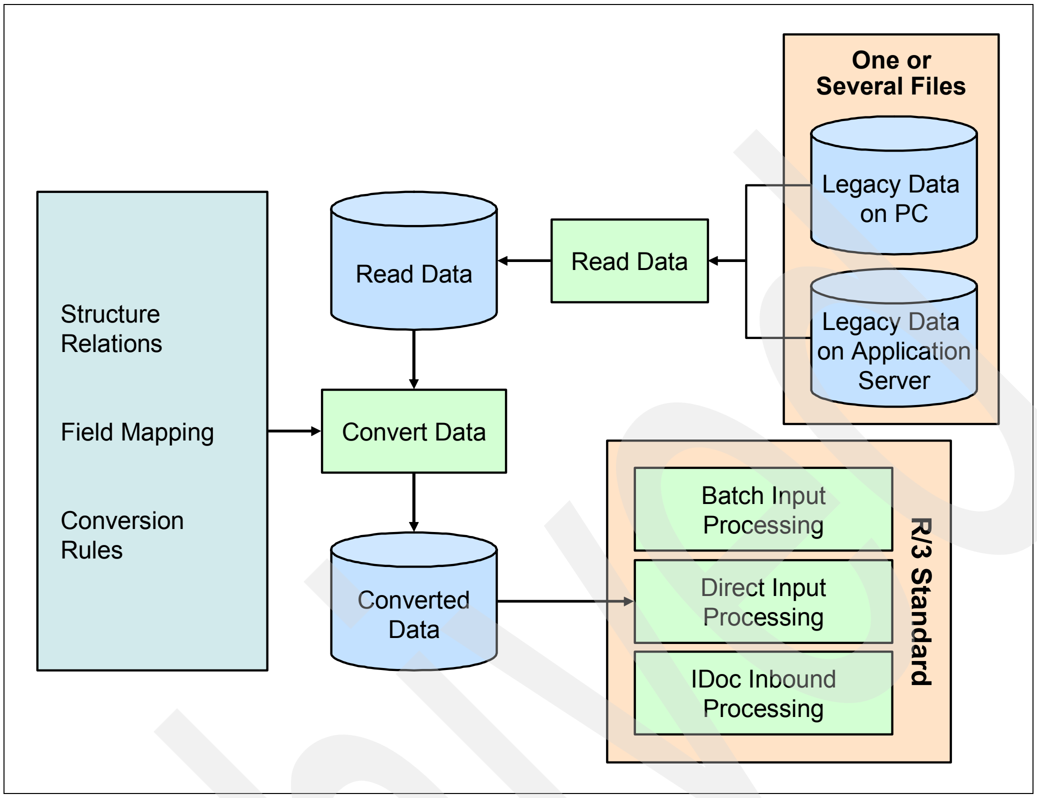

Figure 30-4 How LSMW works

The benefits of LSMW include:

- It is free of charge for all SAP customers and business partners.
- Provides maximum data consistency.
- Standard support is provided via SAP Service Marketplace (component BC-SRV-DX-LSM).
- Leads you smoothly through all the steps of data migration.
- Can be downloaded easily and quickly from SAP Service Marketplace.
- Meets your requirements because it is a highly flexible tool.
- Is independent from SAP releases, platforms, and the kind of data to be migrated.
- Can be used without deep knowledge of ABAP.
- Comes with different control functions.
- Allows reusability of data mapping and conversion rules.
- Can be used in conjunction with DX-Workbench.
- Can significantly accelerate data migration.

LSMW releases

Available releases for LSMW include:

- Version 1.7.2 of the LSMW (LSMW 1.7.2) available

  LSMW 1.7.2 requires an SAP system with SAP R/3 4.0 or SAP R/3 4.5.

- Version 1.8.0 of the LSMW (LSMW 1.8.0) available

  LSMW 1.8.0 requires an SAP system with SAP R/3 4.6.

- Version 3.0 of the LSMW (LSMW 3.0) available for Web Application Server 6.10

  LSMW 3.0 requires an SAP Web Application Server 6.10. It is functionality equivalent to LSMW Version 1.7.

- Version 4.0 of the LSMW (LSMW 4.0) integrated in Web Application Server 6.20 and later

  LSMW 4.0 is an integrated part of SAP Web Application Server with 6.20. Do not install any LSMW Add-On version in an SAP system with Web Application Server 6.20 or later.

Installing LSMW

As of LSMW Version 4.0, it is part of the Web Application Server and you cannot install it separately. You only have to install LSMW until Version 3.0. Take care to have the right LSMW version for your SAP application release (see the previous section).

No objects are imported that are also part of the standard version delivered for your SAP system by installing LSMW. Installing LSMW, therefore, does not affect the functions of the SAP system in any way.

Prerequisites

The prerequisites for the installation of LSMW 1.8.0 are:

- The SAP R/3 system has one of these maintenance levels: 4.6A, 4.6B, 4.6C, or 4.6D.
- The SAP correction and transport system is set up correctly and tested.
- The following background jobs are released:
  - RDDIMPDP in client 000
  - RDDIMPDP_CLIENT_*client* in the remaining clients

  Otherwise, start report RDDNEWPP as user DDIC in all clients.

See **System** → **Services** → **Jobs** → **Job overview**. Enter '*' in the Start after event field, where * means the SAP system including SAP R/3, SAP APO, and SAP CRM.

LSMW has its own namespace, so there should be no naming conflicts with other applications.

Installing the software

To install the software, follow these steps:

1. Copy the archive LSMW180.CAR to an arbitrary directory *sourcedir*
2. Switch to directory /usr/sap/trans:

   ```
   cd /usr/sap/trans
   ```

3. Unpack the archive named LSMW180.CAR:

   ```
   CAR -xvf sourcedir/LSMW180.CAR
   ```

4. Switch to directory /usr/sap/trans/bin:

   ```
   cd bin
   ```

5. Import the transport request named U46K001009:

   ```
   tp addtobuffer U46K001009 SID [pf=TMS PROFILE FILE]
   tp import U46K001009 SID [pf=TMS PROFILE FILE] U9
   ```

To distribute the security profiles, follow these steps:

1. As user *sid*adm or *SID*OFR, create the LSMW.CMD file with the following contents:

   ```
   Import
   client cascade = yes
   file = '/usr/sap/trans/data/R001009.U46'
   ```

```
Including 'R3TRTABU'
including 'R3TRTDAT'
```

2. As user *sid*adm or *sid*ofr, execute the following command:

   ```
   R3trans -u 1 LSMW.CMD
   ```

3. Check the file trans.log in the current directory. The maximum admissible return code is 4 (see end of trans.log).

4. Reset the buffer using the following command in the command field:

   ```
   /$SYNC
   ```

For the latest LSMW version updates and for additional information, see:

http://service.sap.com/LSMW

Getting started with LSMW

This section provides an introduction to LSMW and shows simple LMSW windows, with steps and functions.

1. Use transaction LSMW to start working with the LSMW. Figure 30-5 shows the welcome window.

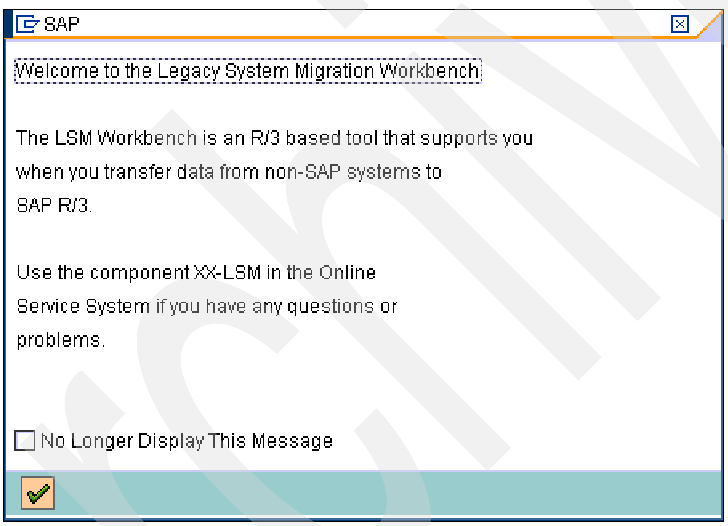

Figure 30-5 Transaction LSMW (Welcome Legacy System Migration Workbench)

2. Select the **Enter** key and you see the LSMW start window. In this window define the following:

   – LSMW project
   – LSMW subproject
   – LSMW object

   Figure 30-6 on page 700 shows the LSMW start window.

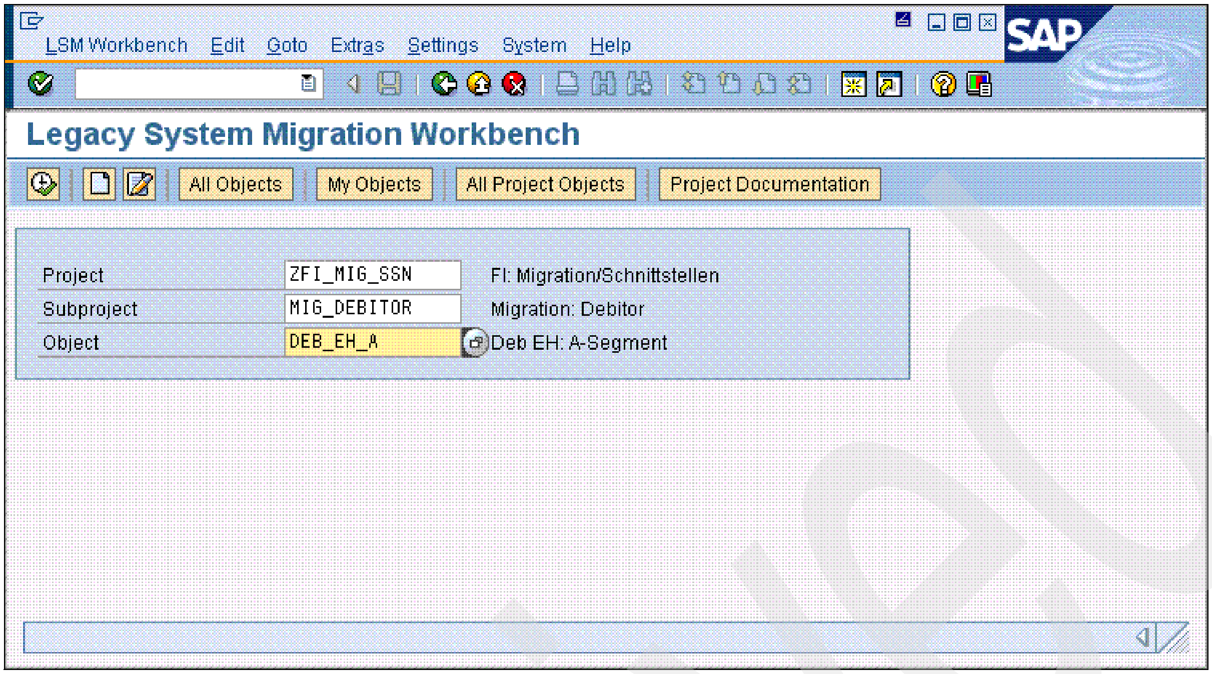

Figure 30-6 Defining the project, subproject, and object

3. Click **Edit** → **Create new entry** on the initial window to create a new project, corresponding subprojects, and objects. Select one of the following buttons:
   - **All objects** provides a list of all projects and their objects 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 a 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 30-7 on page 701 shows an example of a project with several subprojects and objects. Click the **All objects of the project** button to display this representation.

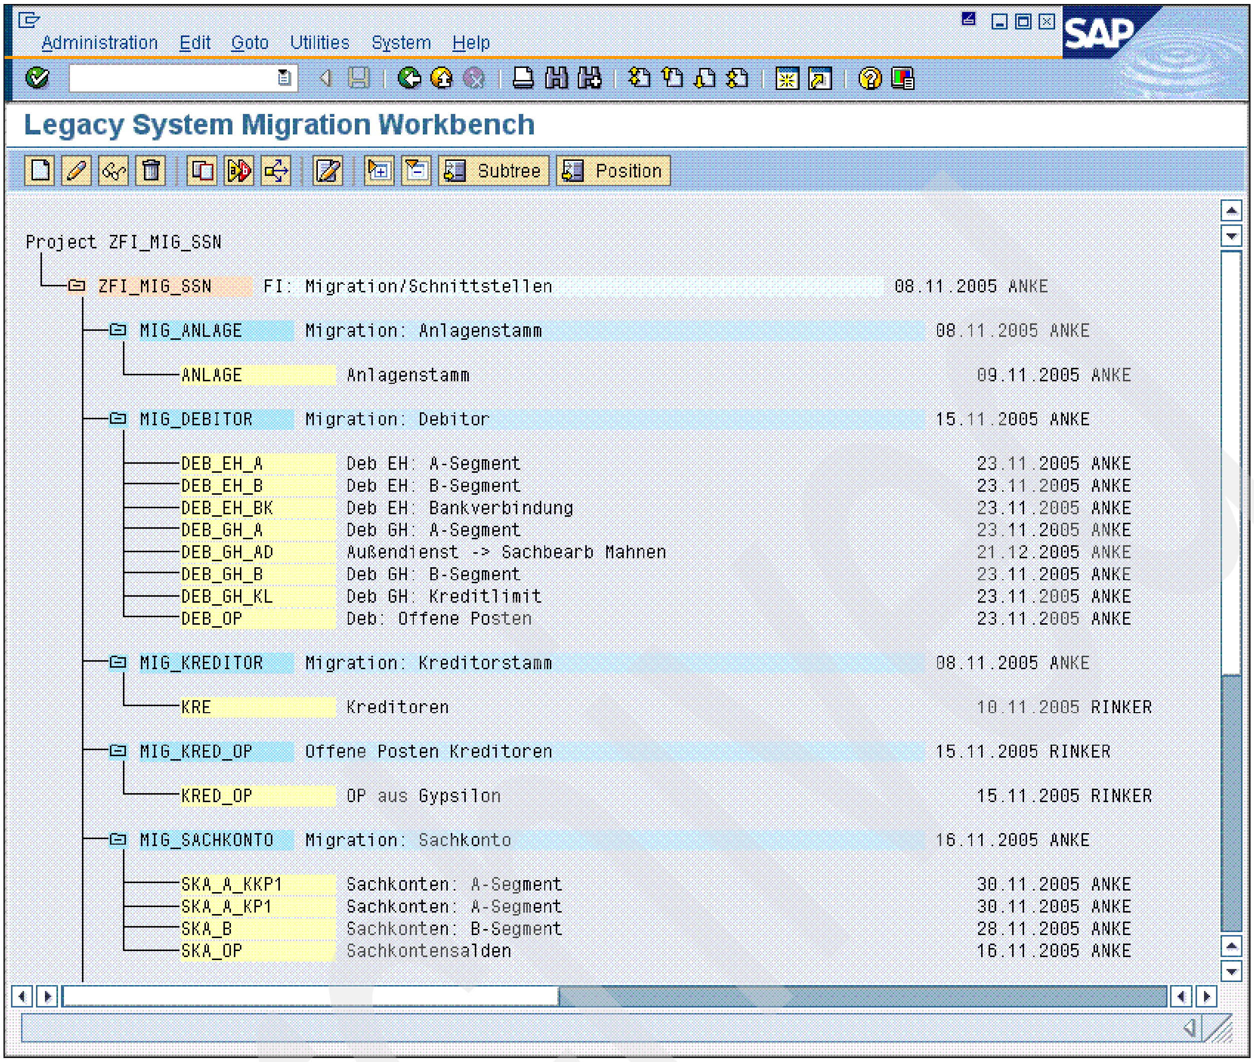

Figure 30-7 Project structure

4. Select **Enter** or **Continue** to go to the interactive process guide after you select an object. This takes you through the individual steps of data migration. The interactive process guide is shown in Figure 30-8 on page 702.

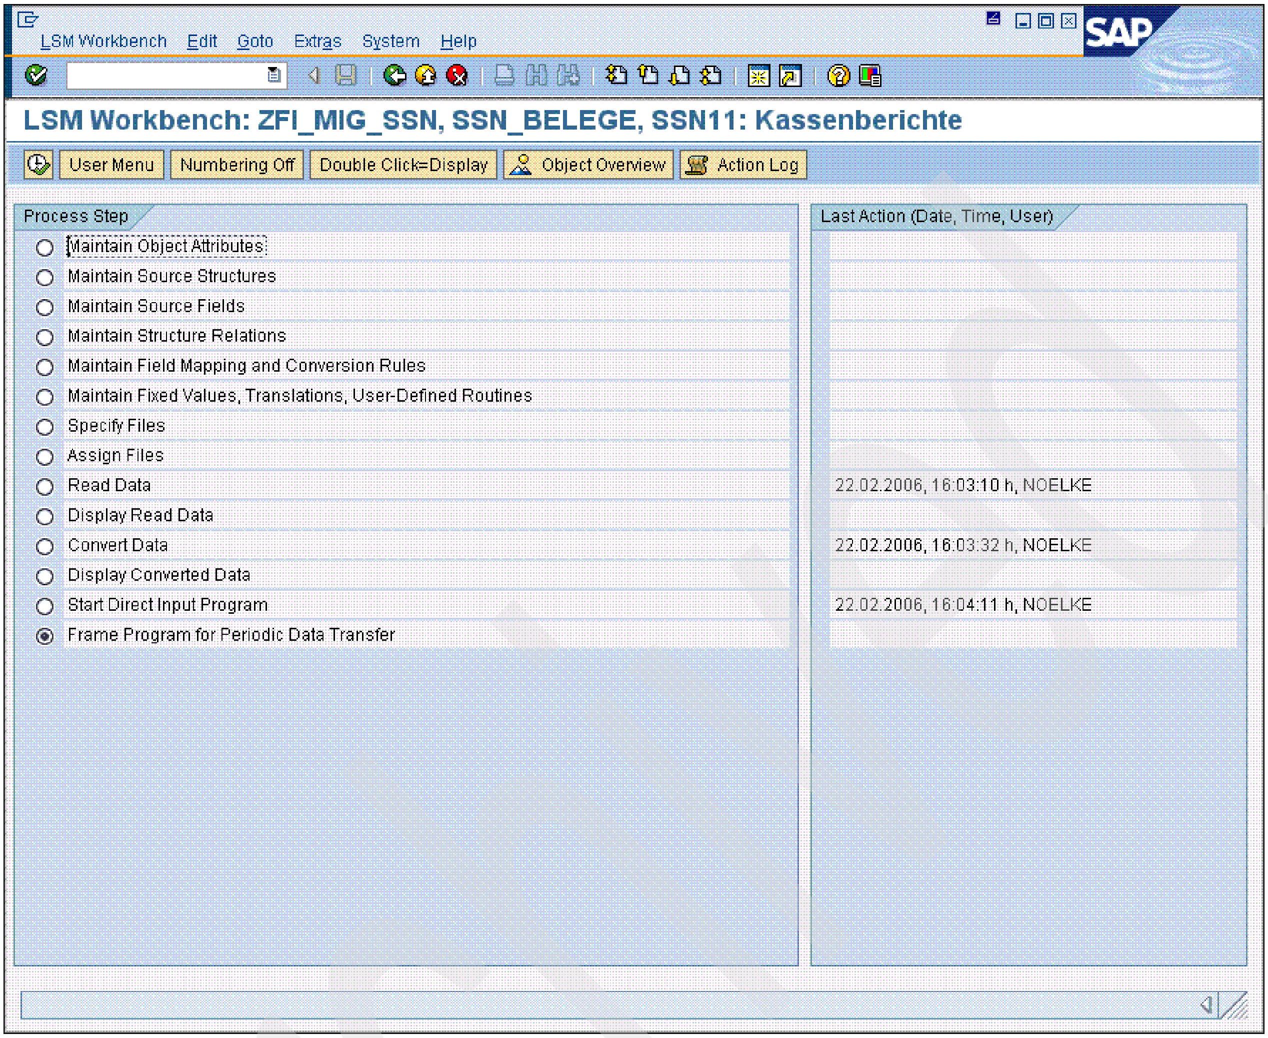

Figure 30-8 LSMW project steps

In the window shown in Figure 30-9, the object type and import technique are selected.

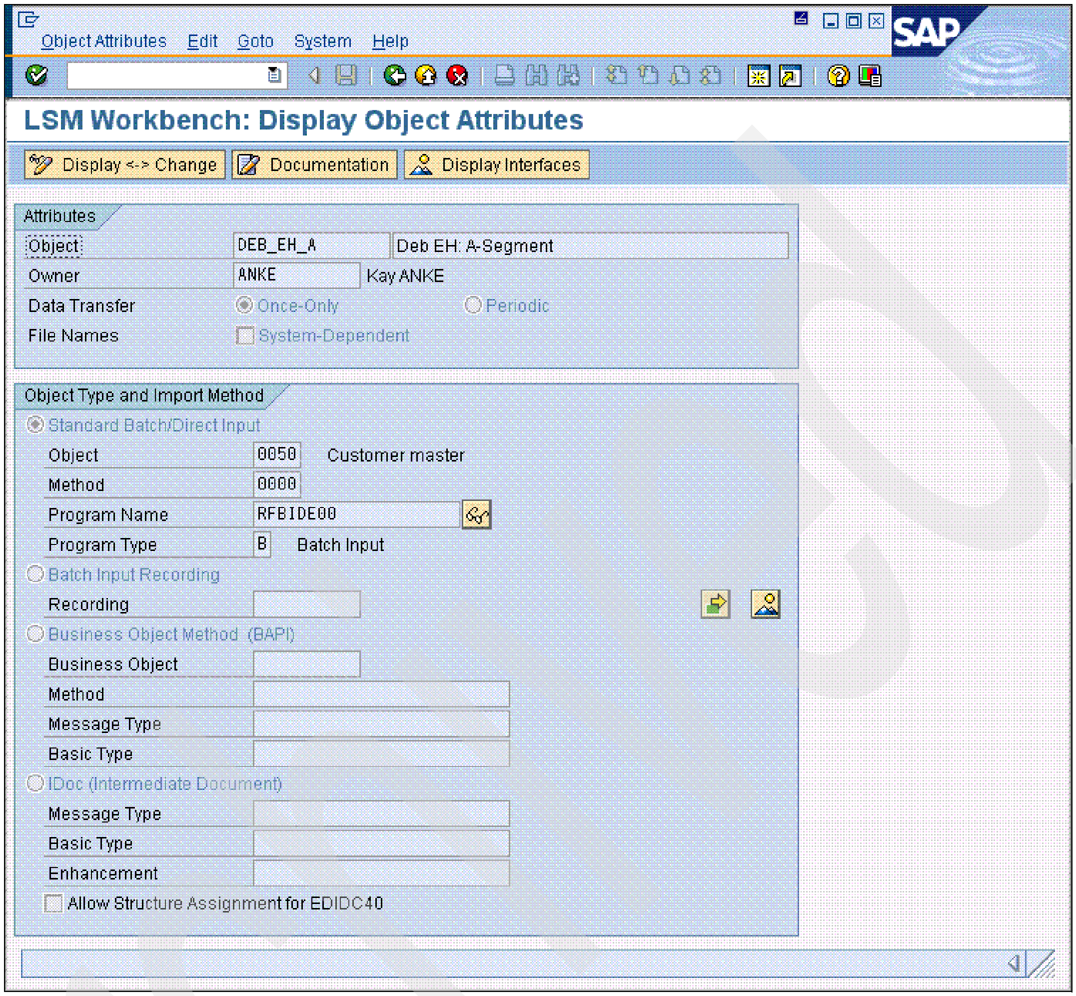

Figure 30-9 Object attributes

You can perform the following steps in this display:

a. Name your object. Enter data in the Owner field and add the project to the list of all projects you created. Display it afterwards in the initial window under My objects.

b. 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.

c. Flag whether the file names are system dependent. If selected, file names can be entered at a later time per system ID.

d. Select the object type and import technique. 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 30-10, the data structures are defined for the existing system, so that they can be mapped in the SAP system. The structures of the object are defined with a name, description, and the hierarchical relationships.

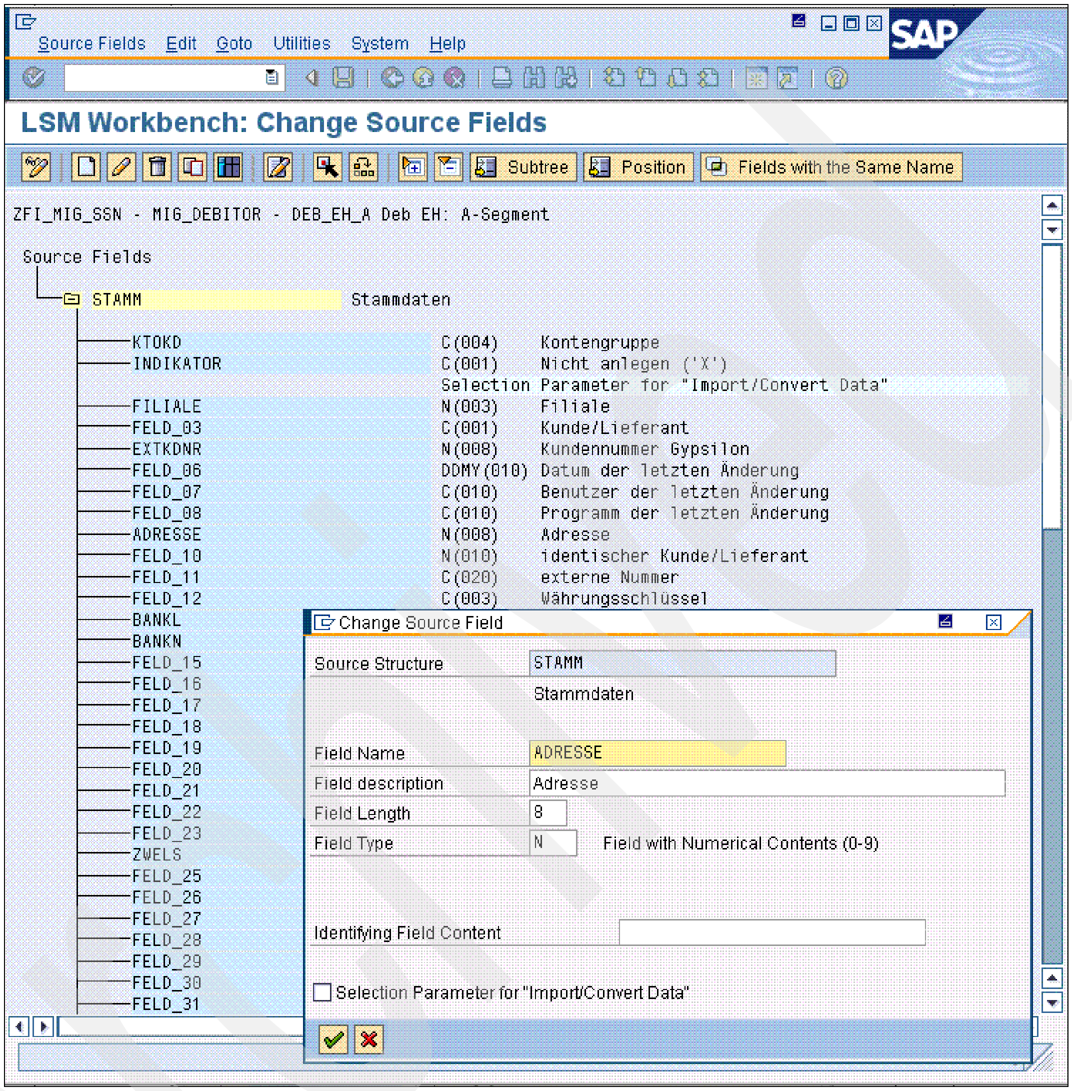

Figure 30-10 Entering the structure for the existing system in an SAP application (source fields)

5. In the pop-up window, click **Change**. You can now define, change, relink, or remove structures. All these functions are available via push buttons. When you define more than one structure, a pop-up display appears that queries the relationship between the structures, that is, same level/subordinated.

   Figure 30-11 shows the rules definitions for the source fields to target fields. It also defines how the field contents are to be converted.

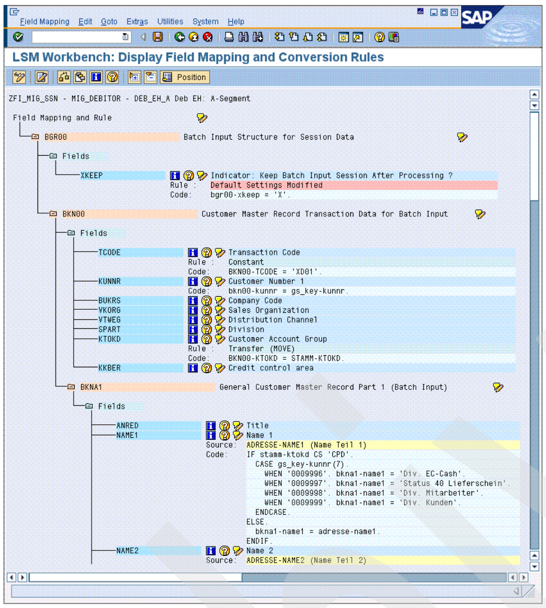

Figure 30-11 Defining conversion rules and field mapping

The display in Figure 30-12 shows the generated conversion code in ABAP.

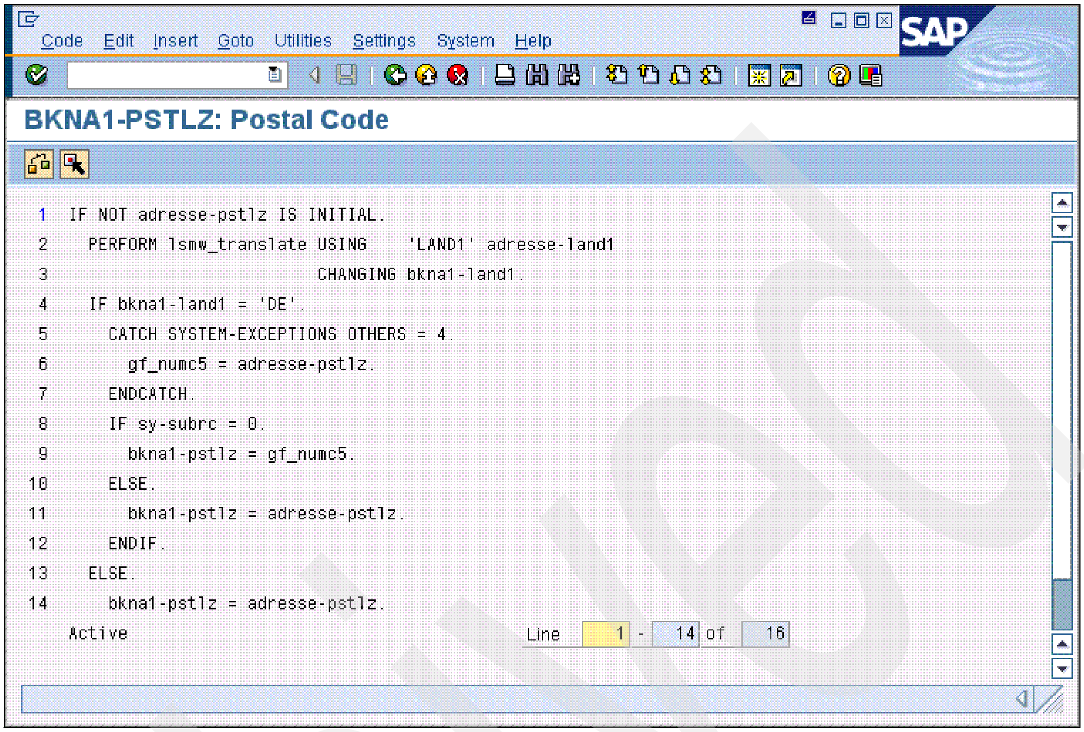

Figure 30-12 ABAP code for the generated conversion program

Data from PC applications or data already converted to a PC-based file can also be converted to SAP system compatible files using LSMW. The display in Figure 30-13 shows the input interface, as from the PC or from the Integrated File System.

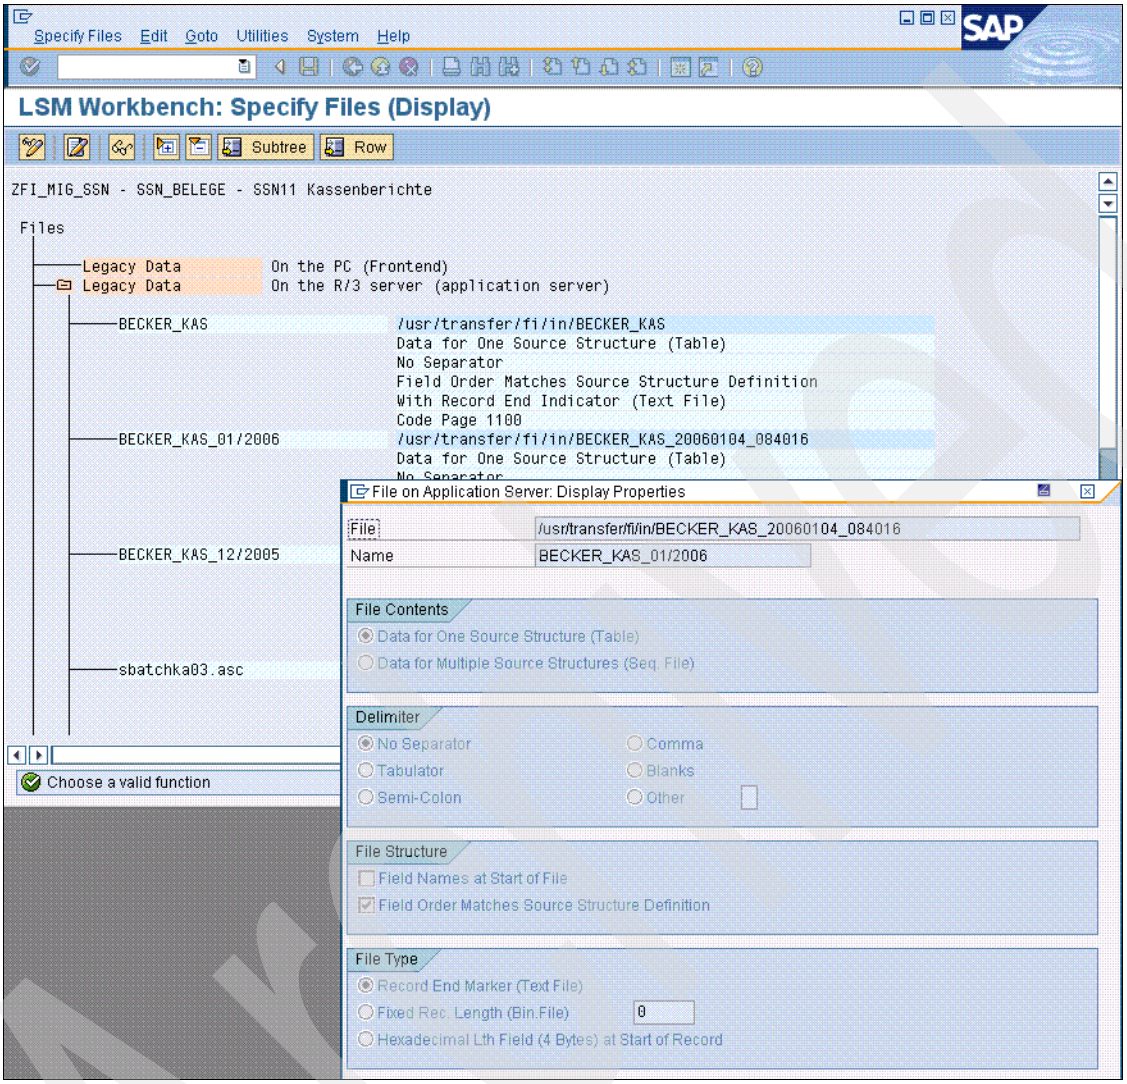

Figure 30-13 Interface for the input file (from the PC or from the Integrated File System)

For additional help and detailed examples of how to use LSMW, refer to:

http://service.sap.com/LSMW

Part 6

Appendices

Part 6 contains appendixes and complementary information to the other parts of this book. Appendix A outlines support options that are available for SAP clients for application and system support.

A

Support for SAP applications

This appendix outlines the support available for SAP with System i models.

A.1 Marketing and technical support

This appendix can help with marketing situations and opportunities. If more technical assistance with SAP applications on System i configurations is required, the iSeries Technology Center (iTC) in Rochester, Minnesota, can assist you.

A.2 Defect support

There are two options to report or work on problems or defects for SAP applications:

- The SAP support channel. This is appropriate if you suspect that the problem is related to SAP's software. Use SAPNet - Frontend (formerly known as OSS) as the entry point to the to report problems or defects to the SAP support channel. (Refer to A.5.2, "SAPNet - Frontend" on page 714.)
- The local IBM Support Center. This avenue is appropriate when you suspect that the problem is linked to the IBM hardware or i5/OS software.

IBM and SAP follow the normal process of defect management and escalation. Both sides can involve their respective partner and work together if necessary. The System i account team can contact the iTC in special situations to seek assistance and information about problems reported by their customers.

A.3 IBM SAP International Competence Center (ISICC) support

The IBM SAP International Competence Center (ISICC) is jointly staffed by SAP and IBM to provide support for SAP applications on IBM platforms. Among its various tasks and services are to:

- Provide marketing support material (presentations, brochures, customer references)
- Provide technical support (benchmarks, sizing recommendations)
- Provide second-level regional support (Infoservice)
- Be the second-level customer briefing center
- Communicate news to the field
- Initiate and lobby for the development of solution packages

The ISICC is located at:

IBM SAP International
Competence Center (ISICC)
Altrottstraße 31
D-69190 Walldorf
Germany
e-mail: ISICC@de.ibm.com

There are two organizations providing international support: IBM and SAP. Base the decision of where to submit a request on the following guidelines:

- Use the ISICC Infoservice for questions related to:
  - IBM hardware platforms in the SAP system environment (PC servers, p5 servers, and System i servers)
  - IBM software and add-ons in the SAP system environment
  - IBM pre-sales support (references, availability, and so on)
- SAP customers (with an SAP customer number) use the SAP regional support when they have questions related to all problems and defects in an SAP installation. SAP handles all

requests centrally. SAP routes problems related to a specific system platform to the respective partner company.

A.3.1 ISICC InfoService

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 SAP and IBM clients.

The ISICC InfoService serves as a point of entry for all IBM SAP-related ERP pre-sales questions directed to the ISICC. Their primary focus is pre-sales situations relative to IBM products within SAPs influence. They also work to improve the flow of information between both companies. As a managed question and answer service, the ISICC InfoService ensures that questions are assigned to the most appropriate experts and answered as soon as possible.

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
- Client briefings
- Access to “SAP Service Marketplace”

The ISICC InfoService can be contacted by:

- e-mail: ISICC@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. Send the request as early as possible so that the Infoservice can provide a quality answer in a timely fashion. Requests received from customers are forwarded to the appropriate country organizations.

A.3.2 SAP regional support

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
- SAP system support Fax: +49 (0) 180/534 34 30
- SAP EarlyWatch support: +49 (0) 180/534 34 35
- SAP system support 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

A.4 Sizing support

You can use the Quicksizer available on `http://service.sap.com/sizing` or ask the local SAP office or the local IBM team for assistance. If they need help to propose a specific hardware configuration IBM Techline helps them. IBM Techline is the first line of contact to assist business partners or IBM colleagues to do a proper sizing.

A.5 Information access

This section highlights the options available to you for having your questions answered.

A.5.1 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 system 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 a valid SAPNet user ID and password. If your company does not have SAPNet access, you can apply through SAPNet Support at +49 (0) 180 534- 34 33.

A license for an SAP system at the client site is required to receive an SAPNet user ID and password. The NET provides an easy to use front end for making comments and responses.

A.5.2 SAPNet - Frontend

SAP provides an online service called SAPNet - Frontend. Clients are required to access SAPNet - Frontend to report and monitor problems. SAPNet - Frontend 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 a X.25, ISDN, or frame relay interface. Notify SAP about the relevant network data.

Refer to SAP's Remote Connection to the Online Services document for details. Information can be found in the SAP system online help text.

A.5.3 System i Informational APARs

A System i Informational Authorized Program Analysis Report (APAR) is created for each SAP supported release. This Information APAR lists the current recommended PTFs and cumulative package level that SAP customers should have on their system. The Informational APARs are:

- V4R4 II11832
- V4R5 II12399
- V5R1 II12833
- V5R2 II13337
- V5R3 II13868 for modification level 0 and II14125 for modification level 5
- V5R4 II14126

A.5.4 iSeries Technology Solutions Center

The iSeriesTechnology Solutions Center's ERP Team provides the following services for SAP customers on the System i platform.

A.5.4.1 Service offerings

The service offerings include:

- System i System Performance Evaluation Offering

  Validate System i performance parameters prior to or after placing a System i SAP implementation into production.

- System i SAP System Capacity Planning Offering

  IBM can assist in assessing System i hardware requirements. This recommendation can help the client meet their information processing objectives based on current utilization and workloads on the system. It also serves as an assessment of any additional requirement in the future which results from upgrading the SAP application software from the current version or increasing the scope of a client's overall operations.

Refer to the iTC's informational Web page for more information about these offerings:

http://iws.as400.ibm.com/Service/bms/support.htm

This page also contains information about tools, performance tips, and fixes.

A.5.5 SAP Support Services

SAP provides the following services free of charge within the maintenance contract:

- GoingLive Check

  The GoingLive Check is done just prior to the system entering production. Prerequisites of going live are checked with the hardware and parameter settings. Otherwise, recommendations are given to improve the performance.

- EarlyWatch

  EarlyWatch service supports SAP system operation by providing remote diagnosis for installations worldwide. It carries out detailed analysis of SAP applications and configurations in addition to database, operating system, and network components. Specific client problems are analyzed, and appropriate solutions are developed.

- GoingLive Functional Upgrade Check

  This check is recommended when the customer decides to upgrade the production SAP system. It consists of two parts. In the first part, the old release is checked (prior to the upgrade). Secondly, a check is done to improve and adjust the settings of the system after the upgrade.

- EarlyWatch Alert

  EarlyWatch Alert gathers information about the customer's system and transfers the data to SAP. SAP analyzes this data automatically, and the generated report is sent to the client. The client 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. The additional information improves SAP's ability to give advice without looking at the client's system, because a lot of the data is already available at SAP.

- Migration Service

  The Migration Service is needed when a client wants to migrate the SAP system to another database. The client then receives the necessary software to export the old database and then import onto the new system after the conversion. Before the migration, checks are done on the old production system. After the migration, the data is imported in order to optimize the performance.

To arrange TCC services, contact SAP at:

Tel: +49 62277 43766
Fax: +49 62277 44214

Also, refer to SAP note 35360 for details.

A.6 Education

We briefly summarize which courses are available that offer content specific to System i configurations with SAP applications.

A.6.1 Basis Administration for SAP on the System i platform

SAP offers a class for basis administration for SAP on iSeries ADM525 SAP NetWeaver on DB2 UDB for iSeries: Database Administration. This course is intended for:

- SAP system administrators
- Database administrators
- Technology consultants

It is a five day course and is available worldwide.

A.6.2 Code Page Conversion

Clients running earlier versions of SAP applications need to migrate to newer versions. This can imply a codepage conversion. The Code Page Conversion class is offered to educate the consultants and basis administrators to perform this conversion.

Abbreviations and acronyms

ABAP	Advanced Business Application Programming
ACS	Alternative Collating Sequence
ALV	ABAP List Viewer
ALV	ABAP List Viewer
APAR	authorized program analysis report
API	application programming interface
APO	Advanced Planning and Optimizing
ASCII	American Standard Code for Information Interchange
ASP	auxiliary storage pool
AWT	Abstract Window Toolkit
BAPI	Business Application Programming Interface
BC	Basis Component
BDC	Batch Data Command
BEx	BW Business Explorer
BI	Business Intelligence
BPP	Business Process Platform
BRMS	Backup Recovery and Media Services
BW	Business Information Warehouse
CA	certificate authority
CCMS	Computing Center Management System
CCSID	coded character set identifier
CEN	central CCMS
CEN	Central Monitoring System
CIC	Customer Interaction Center
CL	control language
CLI	command-line interface
CLI	call level interface
COBOL	Common Business Oriented Language
CPC	code page conversion
cpi	characters per inch
CPW	commercial processing workload
CQE	Classic Query Engine
CRG	cluster resource group
CRM	customer relationship management
CRS	Cluster Resource Services
CTS	Change and Transport System
CUA	Central User Administration
CUoD	Capacity Upgrade on Demand
CUS	customizing system
DASD	direct access storage device
DBA	database administrator
DBCS	double-byte character set
DDS	data description specifications
DIMID	dimension identifier
dLPAR	Dynamic LPAR
DML	Data Manipulation Language
DS	Data Server
DST	dedicated service tools
DW	Disp+Work
ECC	Enterprise Core Component
EJB	Enterprise JavaBeans
EP	Enterprise Portal
ERP	enterprise resource planning
ESA	Enterprise Systems Architecture
ESS	Enterprise Storage Server
EVI	enhanced vector index
EWA	EarlyWatch Alert
GB	gigabyte
GC	garbage collector
GID	group ID
GLS	Global Language Support
GUI	graphical user interface
HA	high availability
HABP	High Availability Business Partner
HIPER	High Impact PERvasive
HMC	Hardware Management Console
HSL	high-speed loop
HSL	high-speed link
HSM	hierarchical storage management
I/O	input/output
i5/OS PASE	i5/OS Portable Application Solution Environment
IASP	Independent Auxiliary Storage Pool
ICM	Internet Communication Manager
IDE	integrated development environment
IDoc	SAP Intermediate Document(

IFS	Integrated File System
ILE	Integrated Language Environment
IMG	Implementation Guide
IMIG	Incremental migration
IOP	input/output processor
IOT	Integrated Operating Team
IPDS	Intelligent Printer Data Stream
IPL	initial program load
IRBD	Initial record layouts
ISICC	International SAP IBM Competence Center
IT	information technology
iTC	iSeries Technology Center
IXA	Integrated xSeries Adapter
IXS	Integrated xSeries Server
JDBC	Java Database Connectivity
JDK	Java Development Kit
JFC	Java Foundation Classes
JRE™	Java runtime environment
KB	kilobyte
KPI	key performance indicator
KPRO	Knowledge Provider
KW	SAP Knowledge Warehouse
LDAP	Lightweight Directory Access Protocol
LIC	Licensed Internal Code
LOB	large object
LOB	Long Object Management
LPAR	logical partition
LPD	line printer daemon
LPG	Lookahead Predicate Generation
LPR	line printer requester
LRC	longitudinal redundancy check
LSMW	Legacy System Migration Workbench
LSMW	LSM Workbench
MCOD	Multiple Components on One Database
md	Make Directory
MI	machine instruction
MI	machine interface
MI	Mobile Infrastructure
MMC	Microsoft Management Console
MQT	materialized query table
NFS	Network File System
NIS	Network Information Service
NVS	nonvolatile storage
NVS	non-volatile shared
ODS	Operational Data Store
OLAP	online analytical processing
OLTP	online transaction processing
OMS	output management system
OS	operating system
OSS	Online Service System
OTF	output text format
PAM	Product Availability Matrix
PCL	printer control language
PDA	personal digital assistant
PDM	Programming Development Manager
PEX	Performance Explorer
PID	process ID
PIN	personal identification number
PJL	Printer Job Language
PLM	Product Lifecycle Management
PM	Performance Management
PMR	Problem Management Record
PPRC	Peer-to-Peer Remote Copy
PSE	Problem Solving Environment
PSP	preventative service planning
PTDV	Performance Trace Data Visualizer
PTDV	Performance Trace Data Visualizer
PTF	program temporary fix
QAS	Quality Assurance System
QTST	Quality Assurance Client
RAID	Redundant Array of Independent Disks
RAM	random access memory
RDB	relational database
RETAIN®	Remote Technical Assistance and Information Network
RFC	Remote Function Call
rpm	revolutions per minute
RRN	relative record number
SAN	storage area network
SAP	Systems, Applications, Products in Data Processing
SAPEVT	SAP Event
SAPS	SAP Application Benchmark Performance Standard

SAPS SAP Application Benchmark Performance Standard
SBCS Single Byte Character Set
SCM supply chain management
SCS SAP Central Services Instance
SCS services instance
SCS SNA character string
SCSI Small Computer System Interface
SDM Software Deployment Manager
SEM Strategic Enterprise Management
SID SAP system name
SID surrogate identifier
SID System ID
SLA service level agreement
SLD System Landscape Directory
SLS single level store
SM Solution Manager
SMAPP system-managed access-path protection
SMP symmetric multiprocessor
SMS Short Message Service
SMT simultaneous multithreading
SNC Secure Network Communications
SOA service-oriented architecture
SP Support Package
SPCN system power control network
SQE SLIC Query Engine
SQE SQL Query Engine
SQL Structured Query Language
SRM Supplier Relationship Management
SSAA System Administration Assistant
SSCR SAP Software Change Registration
SST system service tools
TCO total cost of ownership
telnet Terminal session
TemSe temporary sequential database
TIMI technology-independent machine interface
TMS Transport Management System
TO Transport Organizer
TREX Search and Classification Engine
UDB Universal Database
UDF user-defined function
UDF user-defined functions
UDFS user-defined file system
UID User-ID
UPS uninterruptible power supply
VAR value-added reseller
VE Visual Explain
WAS SAP Web Application Server
WLE Workload Estimator
XI Exchange Infrastructure
XSM cross-site mirroring

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 723. Note that some of the documents referenced here may be available in softcopy only.

- *IBM AS/400 Printing V*, SG24-2160
- *Advanced Functions and Administration on DB2 Universal Database for iSeries*, SG24-4249
- *Implementing SAP R/3 on OS/400*, SG24-4672
- *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189
- *Clustering and IASPs for Higher Availability on the IBM @server iSeries Server*, SG24-5194
- *The IBM TotalStorage Solutions Handbook*, SG24-5250
- *iSeries in Storage Area Networks A Guide to Implementing FC Disk and Tape with iSeries*, SG24-6220
- *IBM iDoctor iSeries Job Watcher: Advanced Performance Tool*, SG24-6474
- *Preparing for and Tuning the V5R2 SQL Query Engine on DB2 Universal Database for iSeries*, SG24-6598
- *Sizing I5/OS Work on @server i5 Partitions*, SG24-6656
- *IBM @server iSeries Security Guide for IBM i5/OS Version 5 Release 3*, SG24-6668
- *IBM @server iSeries Independent ASPs: A Guide to Moving Applications to IASPs*, SG24-6802
- *Microsoft Windows Server 2003 Integration with iSeries*, SG24-6959
- *Building and Scaling SAP Business Information Warehouse on DB2 UDB ESE, SG24-7094*
- *iSeries and IBM TotalStorage: A Guide to Implementing External Disk on @server i5*, SG24-7120
- *Performance Management for IBM @server iSeries and pSeries: A Systems Management Guide*, SG24-7122
- *IBM System i5 Handbook IBM i5/OS Version 5 Release 4 January 2006*, SG24-7486
- *Data Resilience Solutions for IBM i5/OS High Availability Clusters*, REDP-0888
- *Independent ASP Performance Study on the IBM @server iSeries Server*, REDP-3771
- *Best Practices for Microsoft Windows and Linux Integration in iSeries Systems*, REDP-4000
- *IBM @server iSeries Performance Management Tools*, REDP-4026
- *Improve Whole System Backups with the New Save-While-Active Function*, REDP-7200

Other publications

These publications are also relevant as further information sources:

- *IBM @server i5 and iSeries System Handbook: IBM i5/OS Version 5 Release 3 October 2004*, GA19-5486
- *iSeries: Printing with SAP R/3 and AFP,* S544-5412
- *IBM WebSphere Development Studio: ILE RPG Programmer's Guide,* SC09-2507
- *IBM WebSphere Development Studio: ILE RPG Reference*, SC09-2508
- *V5R4 iSeries Security Reference*, SC41-5302
- *Backup and Recovery V5R4,* SC41-5304
- *OS/400 Work Management V4R4*, SC41-5306
- *Performance Tools for iSeries*, SC41-5340
- *Backup Recovery and Media Services for iSeries*, SC41-5345
- *Performance Management/400 V4R4,* SC41-5347
- *APPC Programming V4R1*, SC41-5443
- *Integrated File System Introduction*, SC41-5711
- *V5R3 Printer Device Programming*, SC41-5713
- *AS/400 CL Reference V4R5*, SC41-5722
- *System API Reference*, SC41-5801
- *BC-CCM SAP Printing Guide*
- *BC-ABA ABAP Programming*
- *BC-SRV-SCR SAPscript*
- *SAP R/3 Basis Administration Training 4.6 - BC370*
- *ADM525 - SAP NetWeaver on DB2 UDB for iSeries: Database Administration*
- *Sizing i5/OS Partitions on @server i5* (a presentation)
- *Network Configuration Guide (from dieter's install overview file 10 November)*

Online resources

These Web sites and URLs are also relevant as further information sources:

- SAP Service Marketplace

  http://www.service.sap.com
- SAP Marketplace for Output Management Systems

  http://service.sap.com/output
- iSeries software

  http://www-03.ibm.com/servers/eserver/iseries/software/
- System Copy Guide

  http://service.sap.com/instguides
- iSeries Porting Team in PartnerWorld for Developers

  http://www.iseries.ibm.com/developer/porting/

- IBM white paper: *Geographic Mirroring in the SAP System Environment on @server iSeries*

  http://www-03.ibm.com/servers/eserver/iseries/perfmgmt/pdf/mirrsap.pdf
- HA Business partners

  http://www-01.ibm.com/servers/eserver/iseries/ha/
- Non-IBM porting services and tools from iSeries Business Partners

  http://www.ibm.com/software/
- XSM performance behavior

  http://www-03.ibm.com/servers/eserver/iseries/perfmgmt/pdf/mirrsap.pdf
- LSMW latest version updates and additional information

  http://service.sap.com/LSMW
- iSeries Information Center

  http://publib.boulder.ibm.com/infocenter/iseries/v5r3/index.jsp
- *BC-CCM SAP Printing Guide* in the SAP Online Help

  http://help.sap.com/saphelp_nw04/helpdata/en/d9/4a8eb751ea11d189570000e829fbbd/content.htm
- SAP R/3 printing

  http://sapnet.sap.com/output
- CCSID and globalization

  http://www-03.ibm.com/servers/eserver/iseries/software/globalization/\
- Creation and use of materialized query tables within IBM DB2 UDB for iSeries
  http://www-03.ibm.com/servers/enable/site/education/wp/438a/438a.pdf
- Index and statistics strategies for DB2 UDB for iSeries
  http://www-03.ibm.com/servers/enable/site/education/abstracts/indxng_abs.html.
- Star Schema Join Support within DB2 UDB for iSeries
  http://www-03.ibm.com/servers/enable/site/education/abstracts/16fa_abs.html.

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

5250 interface 628
64-bit 38

A

ABAP 628, 630, 647
ABAP dictionary table 632
ABAP dumps 263
ABAP List Viewer 548
abnormal end 378
access path
    save 394
access path protection 381
activation queue 612
active table 612
ADDENVVAR 135
adding CPU resources 325
adding languages 328
adding memory 325
ADDRMTJRN 450
ADM525 720
administration task 255
adopted authority 124
Advanced Function Printing 225
AFP 225
agent 283
aggregate 613
AL02 290
AL08 261, 287
Alert Monitor 498
alerts 282
All-in-One 11
ALV 548
American Standard Code for Information Interchange 331
API
    QjoRetrieveJournalInformation 411
application
    EarlyWatch 719
    Legacy System Migration Workbench 698
    mySAP Business Suite 9, 11
    SAP Solution Manager 12
    SAPNet - R/3 Frontend 718
    SSAPNet - Frontend 716
    Transport Management System 196
application integration 14
Application Monitor 506
application servers 646
application software 392
Apply Journaled Changes command 417
APYJRNCHG 426
APYJRNCHGX 396, 417, 425–426
APYR3FIX 170
archiving 323
AS4FIXFILE 112
AS525 341
ASCII 331
ASP 55, 441
ASP overflow 117, 409, 416
asynchronous mode 448
attributes 602
auxiliary storage pool 55, 382, 441
availability
    continuous 377
availability solution 379

B

background job 272
background work process 273
BACKUP 387
backup
    offline 396, 428
    online 396, 429
    online with Rapid Checkpoint Save-While-Active 396
backup performance 393
Backup Recovery Media Services 383
backup system 439
barcode printing 250
Basis Administration for SAP on System i 720
Basis Component 141
Batch Data Command (BDC) 696
batch input 694
batch input program 695
batch input session 696
Batch jobs 370
BC 141
BC-CCM-PRN SAP Printing Guide in the SAP Online Help 246
BEx Reporting flag 619
BI 534
binary patch 158
boot device 95
boot from SAN 99
BRMS 383
BRMS recovery plan 434
buffer trace 519
Buffers 546
buffers
    tune 515
Business Intelligence 534
business operations 438
Business Partner Solution 458
Business Process Monitoring 281
BW 36

C

capacity
- CPU 325
- disk load 325
- memory 325
- storage 322

Capacity Upgrade on Demand (CUoD) 326
capped mode 35
catch-up mode 447
CCMS 276
CCMS agent 283
CCSID 628, 630, 634–635, 638
central instance 438
CFGPFRCOL 485
CFGSAPIASP 454
Change and Transport System 185
change log table 612
change request 185, 203
Character Set-Identifier 630
CHGENVVAR 135
CHGR3SHLOC 190
CHGRMTJRN 451
CHKXDA 83, 353
CHKXDASM 354
CL 54
class 465
Classic Query Engine 611
client
- customizing 186
- Customizing Client 186
- default clients 191
- Early Watch Client 191
- production 186
- Production Client 186
- Quality Assurance Client 186
- sandbox client 187
- test 186
- Training client 187

client / server response time 462
client attributes 202
client delete 323
client dependent 201
client independent 201
cluster 440
cluster management middleware 458
Cluster Resource Services 440
clustering 440
CMDRST 387
CMDSAV 387
Code Page Conversion 373
code page conversion 647, 720
code pages 627
codepage
- GLS 330

Codepage Conversion (CPC) 40, 329
command
- ADDRMTJRN 450
- APYJRNCHG 426
- APYJRNCHGX 396, 417, 425–426
- AS4FIXFILE 112
- BDC 696
- CFGPFRCOL 485
- CFGSAPIASP. 454
- CHGR3SHLOC 190
- CHGRMTJRN 451
- CRTCBLMOD 679
- CRTDEVPRT 227, 229
- CRTJRN 112
- CRTJRNRCV 112
- CRTOUTQ 228, 234
- CRTPF 630
- CRTPGM 674, 680, 690
- CRTRPGMOD 674, 689
- CRTSAPLCL 328
- CRTSAPLF 635–636, 640
- CRTSAPVIEW 636, 641–642
- DLTR3PKG 357, 426
- DSPPFM 634
- DSPSBSD 464
- DSPTMPSTG 323
- ENDPFRCOL 485
- GO BACKUP 387
- GO CMDRST 387
- GO CMDSAV 387
- GO RESTORE 387
- GO SAVE 387
- i5/OS printer 227
- LODR3KRN 153
- MKDIR 195
- PRTDSKINF 323
- R3trans 185
- RCVJRNE 445
- restore commands 387
- RGZPFM 323,
- RST 387, 417
- RSTAUT 387, 417
- RSTCFG 387, 417
- RSTDLO 387, 417
- RSTLIB 387, 417
- RSTOBJ 387
- RSTUSRPRF 387, 417
- RTVDSKINF 323
- RZGPFM 323, 521
- SAP Event 659
- SAPEVT 659
- SAPROUTER 307
- SAV 387
- SAVCFG 387
- SAVCHGOBJ 387
- SAVDLO 386–387
- SAVDLTRCV 411
- SAVDLTRCVE 412
- SAVE 386
- Save Changed Objects 387
- Save Configuration 387
- Save Document Library Objects 386–387
- Save Integrated File System 387
- Save Library 387
- Save Objects 387
- Save Security Data 387

Save Storage 387
Save System 387
SAVLIB 387
SAVOBJ 387
SAVR3SYS 83
SAVSECDTA 387
SAVSTG 387
SAVSYS 387
STARTSAP 344
STOPSAP 400
STRASPBAL 322–323
STRPFRCOL 485
STRPFRT 589
STRPRTWTR 228
STRRCYBRM 434
STRREPORT 628, 659–660
STRRMTWTR 228
tp 185
UPDDTA 634
Work with Active Jobs (WRKACTJOB) 350
Work with Job by PID (WRKPID) 348
WRKACTJOB 477
WRKACTJOB (Work with Active Jobs) 350
WRKDSKSTS 322, 477
WRKJRN 424
WRKJRNA 410
WRKLNK 54
WRKLNK (Work with Link) 72
WRKLNKSAP (Work with Object Links) 72
WRKMEDIBRM 432
WRKOUTQ 228, 237
WRKPID (Work with Job by PID) 348
WRKPRB 336
WRKQRY 635
WRKSBSJOB 481
WRKSYSACT 482, 590
WRKSYSSTS 322, 475
WRKWTR 228
Command Language 54
command MD 195
commitment boundary 422
commitment control 382
Component report 590
Computing Center Management System 276
configuration file 69
CONNECT TO 647
considerations
journal receiver 409
continuous availability 377
continuous operations 377
CPI-C 661
CPU utilization 584
CPU virtualization 326
CQE 551, 611
Cross Site Mirroring 438
cross site mirroring 455
cross-site mirroring 454
CRS 440
CRTCBLMOD 679
CRTDEVPRT 227, 229
CRTJRN 112
CRTJRNRCV 112
CRTOUTQ 228, 234
CRTPF 630
CRTPGM 674, 680, 690
CRTRPGMOD 674, 689
CRTSAPLCL 328
CRTSAPLF 635–636, 640
CRTSAPVIEW 636, 641–642
CTS 185
cumulative PTF 79
Cumulative PTF package 336
CUST 186
Customer Services Network 599
customer subtree 277
customized ABAP code 371
customizing 201
customizing client 186

D

data
TemSe 221, 224
Data Collection Services 483
data porting 694
data source 649
data transfer program 695
database 393
consistency 416
SAP 438
database error 357
database journal entries 457
database lock monitor 349
database monitor 526
database multiconnect 646–647
database multiconnect with EXEC SQL 645
database performance 520
database performance monitor 523
DB01 349
DB02 290, 520
DB2 UDB 4
DB2 UDB for iSeries 628
DB2 UDB for Linux/Unix/Windows 646
DB4COCKPIT 526
DBCON 647–648
DBSL interface 353
DDS 629, 641
dedicated system 387
default profile 177
defect support 716
Deleted Row Analysis 323
delta update 612
developer trace 139, 352
Development and QAS in one SAP system 26
device parity protection 97, 380
dialog work process 263
dimension identifier 603
dimension identifier (DIMID) 616
directory
/QFileSvr.400 194, 360
/sapmnt 452

/sapmnt/trans 193
/usr/lib/nls/loc 88
/usr/sap/trans 193
QFileSvr.400 194
disaster recovery 370, 377, 438, 442
disk failure 379
disk load 325
dLPAR 34
DLTR3PKG 357, 426
driver
JTOpen 71
native 71
toolbox 71
DSPCURSQL 83
DSPPFM 634
DSPSBSD 464
DSPTMPSTG 323
dual systems 381
dw.SAR kernel patch 166
dynamic LPAR 34
dynamic SQL 64

E

Early Watch Alert 285
EarlyWatch 570, 719
EarlyWatch Alert 719
EarlyWatch Check 570
EBCDIC 373, 630, 634–636, 638, 643, 646
EBCDIC Latin-1 635
EDRSQL server 139
education 341
Basis Administration for SAP on System i 720
code page conversion 720
SAP course UABC92, i5/OS Codepage Conversion 329
SAP NetWeaver on DB2 UDB for iSeries Database Administration 720
SAP on iSeries 720
Encoded Vector Index 464, 610, 611
ENCODING 646
ENDPFRCOL 485
enqueue 263
enqueue trace 518
enqueue work process 263
Enterprise Resource Planning 36
Environment variables 135
event handler 628
EVI 464, 610–611
Example environment descriptions 26
EXEC-SQL 628, 645–646, 652
explain function 608
Extended Binary Coded Decimal Interchange Code 331
extended dynamic SQL 64
extended star schema 602–603

F

fact table 603
failure
disk 379
power 379
program 379
fault rate 539, 591

file
.CAR archive 167
/usr/sap///log/SLOG 352
/usr/sap///work/sapstart.log 352
configuration 69
developer trace 139
file server 66
library file system 66
log 351
open systems file 66
optical file 66
Personal Security Environment 69
program described file 653
QFileSvr 66
QOpenSys 66
QOPT 66
QSYS.LIB 66
root file system 66
sequential 645
stream file 652
system log 351
trace 351
File Server file system 66
file system 68
Integrated File System 54
NFS 66
QFileSvr 68
QSYS 68
UDFS 66
FIXR3OWNS 174
FOR ALL ENTRIES 371
format 245, 245
frontend printer 244

G

Global Language Solution 40, 329
Global Language System 627
global transport directory 190
GLS 40, 329, 627, 635–636
GLS ASCII 646
GLS codepage 39, 330
GLS Latin-1 634–635
GO CMDRST 387
GO RESTORE 387
GoingLive Check 719
GoingLive Functional Upgrade Check 599, 719
group PTF 336
growth management 322
GUI 628

H

hard link 66
Heap Analysis 594
heterogeneous environment 213
heterogeneous system copy 138, 328–329

high availability 370, 377, 438, 457
hiper group PTF 337
homogeneous environment 212
homogeneous system copy 138, 328
host print transform 225
host spool access method 240
hot package 201
HotNews 338

I

i/5 Portable Application Solution Environment 87
I/O performance 456
i5/OS commands 653–654
i5/OS environment 628
i5/OS PASE 54, 87
i5/OS Portable Application Solutions Environment 54
i5/OS printer command 227
i5/OS release upgrade 326
IASP 55, 441, 453
IBM Backup Recovery and Media Services (BRMS) 28
IBM Insight
    download 573
    prerequisites 573
IBM Insight for SAP 324, 572
IBM Insight Tool 48
IBM SAP International Competence Center (ISICC) 716
IBM TotalStorage solutions 98
iDoctor for iSeries 594
ILE 662
ILE RPG 684
Independent Auxiliary Storage Pool 55, 409, 441, 453
indexes
    advised 526
    created 526
    used 526
industry-specific solutions 11
InfoCube 602
InfoObjects 602
information integration 13
Informational APAR 337, 367, 718
Insight report 575
installation guide 138
instance profile 177, 468
instance profile tuning 326
integrated database 5
Integrated File System 54, 438
    problems 360
Intelligent Printer Data Stream 225
Interactive program communication 661
interactive response time 462
internal tables 371
IPDS 225
IPL 371
iSeries Navigator 487
iSeries Technology Center 716
ISICC 716
ISICC InfoService 717
ISICC infoservice (Walldorf) 717
iTC 716

J

Java 14, 628, 648, 662, 684
    monitoring 519
    native driver 91
Java connection 628
Java database 648
Java Watcher 594
JDBC 649
JDBC Connector Service 648
JDBC driver 648–649
job
    attributes 346
    job description 465
    QDBFSTCCOL 54
    QDBSRVXR 54
    QDBSRVXR2 54
    SCPF 54
    status 345
    structure 344
    system 54
job description 465
Job interval report 590
job lock 420
job logs 346
Job Watcher 594
journal 382
    shared 103
journal entries 381
journal management 381
journal receiver 370, 382, 408
    considerations 409
    save 409
journal switch 62
journaling 382
JTOpen driver 71

K

kernel 142
kernel library 642
kernel name 144
    R3 143
    R3OPT 143
    R3OPT 143
    standard 144
kernel patch 338
kernel patches 175
key performance indicator 284
kpi 284

L

landscape
    one-system landscape 189
    three-system landscape 189
    two-system landscape 189
Latin-1 647
Latin-2 642
legal change patch 201
library 59
    sapcrypto 304

Library File System 66
library name 145
  J2EE 145
  R3 143
  R3OPT 143
  R3OPT 143
LIBRFC 661
LIBRFC library 627
Linux on Power 646
livecache 36
load source 95
load source unit 94
locating R/3 developer trace files 352
lock 263
lock monitor 526
lock table 263
locks 262, 420
LODR3KRN 153
LODSAPKRN 154
log file 351
logical data replication 451
logical file 643
logical partition 34, 370
logical partitioning 439
logical processor 34
logical replication 457
logical replication environment 439
Lookahead Predicate Generation 604, 606, 611,
LPAR 34, 36
LPG 604, 606, 611
LSM Workbench 699
LSMW 698

M

main storage 53
materialized query table 612
MaxDB for MySQL 646
MD 195
MDMP 38, 635–636
Media Definition Object 393
memory management
  SAP 466
memory pool 53, 464
  pool
    memory pool 464
Metro Mirror copy services 442
MGTCOL 585, 589
micro partitions 35
Migration Service 719
mirrored protection 97, 380, 442
mirrored Sysbas 438
mirrored system environment 442
mirroring
  remote 97
MKDIR 195
mode
  asynchronous 448
  catch-up 447
  synchronous 448
monitor
  Application Monitor 506
  database 526
  database performance 523
  lock 526
monitoring 692
  Java 519
monitoring task 255
more users 324
MQT 612
MS SQL Server 646
MTXW status 362
multi-data multi-processing 38
multiple code page support 636
mySAP Business Suite 9, 11

N

native driver 71
native SQL 650
native SQL connection 648
Network File System 66
NFS 66
niping 301
no-data character 695
non-Latin-1 635
non-SAP data 645
non-SAP database table 645

O

ODS 612
offline backup 396, 428
OFR 121, 129
OLAP 584
OLAP systems 36
OLTP 584
OLTP systems 36
one-system landscape 189
Online Analytical Processing 584
online backup 393, 396, 429
online backup with Rapid Checkpoint Save-While-Active 396
Online Service System (OSS) 718
Online Transaction Processing 584
OPEN DATASET 646
open SQL 650
Open Systems File System 66
operating system call 628
operating system images 34
Operational Data Store 612
OPR 122, 130
Optical File System 66
optimization
  Java 566
optimized kernel 142
optimizer 641
Oracle 646
OS Collector 581
OSS 716
OSS (Online Service System) 718
outage 377

overflow 409
  ASP 416

P

page format 245
parallel backup and restore 393
parameter
  *SYNCLIB 406, 429
  commitment boundary 422
  of SAVDLTRCV 411
  precheck option 394
  SAVACTWAIT 402
parameter read sequence 270
parameter replace sequence 270
parameter setting 471
partitioned binary radix tree 463
PCL 229
Peer to Peer Remote Copy 383
pending changes 406
people integration 13
performance 372
  3-tier 579
  application program 566
  backup 393
  communications 542
  complex system landscapes 581
  database 520, 549
  I/O 456
  IOP 542
  Java optimization 566
  LPAR 540
  multiple system 582
  system 536
Performance Explorer Analyzer 594
Performance Explorer Database 595
Performance Report
  Print 590
performance report
  Resource 590
Performance Tools 325, 589
performance trace 518
Performance Trace Data Visualizer 594
Personal Security Environment file 69
PEX 594
PJL driver 229
PM iSeries 587
PMR 340, 372
pool
  memory 464
  system 582
Pool Report 584, 590
power failure 379
PPRC 383
PRCED 64
precheck option parameter 394
Preventative Service Planning 337
primary key constraint 61
primary system 440
Printer Command Language 229
printer data stream
  AFP 225
  IPDS 225
  SCS 225
printer definition 238
priority
  lower 585
problem analysis 355
problem determination 343
Problem Management Record (PMR) 340
process integration 13
process overview 502
production client 186
profile
  default 177
  instance 177
  start 177
program
  batch input 695
  CHKXDASM 354
  data transfer 695
  LIBRFC 661
  QSTRUP 310
  QXDARECVR 353
  R3INLPGM 136
  RSP01041 252
  safeguard 597
  SAP2AFP 227
  SAPLPD 243
  SAPlpd 243
  SAPOSCOL 364
  SAProuter 298
  STAT 290
  Z8189AB1 630
program described file 653
program exceptions 358
program failure 379
Programming Temporary Fix 336
PRTDSKINF 323
PSP 337
PSP bucket 337
PTDV 594
PTF 336
  Cumulative PTF package 336
  group 336
  hiper group 337

Q

QDBFSTCCOL 54, 64
QDBSRVX2 54
QDBSRVXR 54
QFileSvr.400 194
QFileSvr.400 file system 68
QjoRetrieveJournalInformation API 411
QPEXDATA 595
QSTGLOWACN 322
QSTGLOWLMT 322
QSTRUP 310
QSTRUPPGM system value 310
QSYS.LIB file system 68
QTST 186

query
- Group By 553
- index 554
- Join condition 552
- Order By 554
- Select records 552
- views 556

query optimization 551
query optimizer 551
Quicksizer 44
QXDAEDRSQL 315
QXDARECVR 353

R

R3 143
R3INLPGM 136
R3OPT 143
R3OPT 143
R3RMTDB 315
R3trans 185
Rapid Checkpoint Save-While-Active 396
Rapid Checkpoint Save-while-Active 430
RCLSTG 371
RCVJRNE 445
RDB 448
Receive Journal Entry command 445
recovery
- BRMS recovery plan 434
- reload 438

Redbooks Web site 727
- Contact us xxxii

regional support 716
register communication 668
relational database entry 448–449
release upgrade 327
reload 438
Remote Function Call Software Development Kit 146
remote journal environment 446
remote journaling 438, 445, 451
remote mirroring 97
Remote Performance Optimization 598
reorganization of SQL tables 63
report
- Component 590
- Insight 575
- Job interval 590
- Pool 584, 590
- print 590
- System 590

resilient data 440
Resource report 590
response time 44
- client / server 462
- components 544
- database 545
- interactive 462

RESTORE 387
restore commands 387
restore menu options 384
restricted state 387
Reuse Deleted Records 63
REUSEDLTRCD 63
RFC 628, 661–662
RFC connection 627, 661, 683
RFC destination 663
RFC interaction 661
RFC quota 183
RFC trace 519
RFC-SDK 146
RGZPFM 323, 371, 521
ROLLBACK 371
root file system 66
routing entries 464
RPG program 629
RSP01041 252
RSRT 613
RST 387, 417
RSTAUT 387, 417
RSTCFG 387, 417
RSTDLO 387, 417
RSTLIB 387, 417
RSTOBJ 387
RSTUSRPRF 387, 417
RTVDSKINF 323
RUNPTY 585
RZ10 271
RZ11 271
RZ12 182–183
RZ20 251, 290, 497–498

S

SA38 271
safeguard program 597
SAN-based boot support 95
SAP
- regional support 717
- work process 544

SAP applications 36
SAP certification 152
SAP data dictionary 632
SAP database 438, 452, 647
SAP database interface 634
SAP dictionary 628
SAP dictionary table 633
SAP Early Watch 325
SAP identification number 247
SAP kernel 452
SAP locks 262
SAP memory management 466
SAP message 340, 372
SAP NetWeaver architecture 12
SAP note
- 011842 181
- 101113 548
- 051651 179
- 103747 180, 183, 466
- 107899 222
- 110157 180
- 114426 243–244
- 119754 180

12103 364
121625 180, 548
121808 180
123418 548
126210 182
130243 620
130696 613, 620
130978 252
138498 74
139326 466
146624 180, 647
146642 65
14754 139
148208 174
150533 244
156554 153
156557 148–149, 151
167229 182
17054 250
173579 119, 135
17852 364
182207 183
19466 167
20176 221
202593 83, 139, 394
212876 167
25344 243
26317 139
27532 206
307077 183, 620
330793 167
33576 180, 183
3467 183
35360 720
371029 181, 183
371786 364
373986 548
376827 179
384971 181
39031 243
395365 178
410783 148–149, 151, 153
416994 5
425593 183, 616, 620
428855 139
445872 180
44695 180, 466
450351 84, 139, 353–354, 580–581
450870 146
45335 89, 586
46063 119, 135
484548 543, 581
485420 178, 620
48803 243
49201 53
49365 159–160, 171–172, 356
494086 178
49701 160, 172
497608 83
500235 543
500950 330
501572 178, 183, 558, 610–611, 620
517515 541
522878 284
52505 150
531337 181–182
541508 183, 558, 621
554792 632, 636
557870 621
562224 621
567745 621
567746 621
567747 621
568139 183, 621
568820 179, 453–454
597312 183
601109 183, 621
601110 183, 619, 621
60252 139
611946 150, 466
63037 616
637174 364, 582
641336 183
647616 466
65224 227
654794 522
654800 91
654801 115
663811 150
66467 167
67213 195–196
677345 654
682281 315–316
68440 151, 327
68487 69
68732 81–82, 85
693802 220, 226
700778 11
705888 178
708136 364, 582
727078 565
717376 355, 359, 567
732453 156
74141 183
743113 152, 327
744055 105, 112
751131 160, 172
751132 143, 156
753917 582
760162 158
787127 326, 364
787302 150
79478 181
795106 621
800791 331
805447 327
807912 153, 328
808607 466
809693 91
818091 213

820325 65, 178–179, 621
82079 411
825473 361, 402
826449 91
828354 112
83292 81, 152
834218 119, 135
84081 323, 378
84548 178
84972 616
849720 621
850218 364
86161 206
86329 299–300
863803 88
870871 384
874028 179
874480 364
875679 178
889292 364
892513 621
8928 225, 243
898203 178
904977 156
907733 649–650
924753 65
92589 139
93316 315
935299 619, 621
94998 384
97497 180
9942 544
SAP OS Collector 581
SAP OS/DB Migration Check 329
SAP printer definition 238
SAP program
QXDAEDRSQL 315
R3RMTDB 315
Z8189AB2 652
Z8189AB3 653
Z8189AB4 653
Z8189AB5 647, 651
Z8189SQL 647
SAP R/3 10
SAP release upgrade 327
SAP Service Marketplace 718
SAP Software Change Registration 204
SAP Solution Manager 12
SAP system environment 628
SAP System ID 188
SAP system libraries 388
SAP tool
APYR3FIX 168
SAP transaction 656
AL02 290
AL08 261, 287
DB01 349
DB02 290, 323, 520
DB4COCKPIT 526
RSRT 613
RZ10 90, 271
RZ11 90, 271
RZ12 182–183
RZ20 251, 290, 497–498
SA38 271
SCC1 187
SCC3 191
SCC4 186, 191
SCCL 191
SE09 207, 212
SE10 207
SE13 547
SE38 271, 359
SE73 250
SM02 270
SM04 261, 270, 287
SM12 262, 288
SM13 261, 270, 287
SM21 264, 274, 289, 351, 356, 358, 504, 543
SM35 270, 288, 696
SM36 653, 659–660
SM37 269, 288, 363
SM49 653, 655
SM50 84, 253, 261, 269, 287, 346, 348–349, 352, 502
SM50 363
SM51 221, 260, 356
SM59 667
SM62 660
SM66 261, 500
SM69 653–654, 658
SP01 242, 251, 253, 289
SP12 289
SPAD 219, 238–239, 243–245, 248, 250
SSAA 499
ST02 180, 290, 359, 515
ST03 289, 509, 558
ST04 505, 523
ST05 352, 518, 559, 608, 614
ST06 290, 489, 585
ST07 506
ST11 352, 355–356
ST22 263, 288, 352, 358
STMS 196, 198, 209
SAP2AFP 227
SAPcomm support 240
sapconf 69
sapcrypto library 304
SAPEVT 659
SAPLPD 243
SAPlpd 243
SAPNet - Frontend 716
SAPNet - R/3 Frontend 718
SAP-Open-SQL 631
SAPOSCOL 364, 373, 581
sappfpar 183
SAPROUTER 307
SAProuter 298
SAPS 44
SAV 387

SAVACTWAIT 402
SAVCFG 387
SAVCHGOBJ 387
SAVDLO 386–387
SAVDLTRCV 411
SAVDLTRCVE 412
SAVE 387
save
  Rapid Checkpoint Save-while-Active 430
  Save-While-Active 402
  user data 399
save access paths 394
SAVE command 386
save journal receiver 409
save menu options 384
Save-While-Active 402
SAVLIB 387
SAVOBJ 387
SAVR3SYS 83
SAVSECDTA 387
SAVSTG 387
SAVSYS 387
SCC1 187
SCC3 191
SCC4 186, 191
SCCL 191
scheduled outages 378
schema 59
SCPF 54
SCS 225
SE09 207, 212
SE10 207
SE38 271, 359
SE73 250
SEM 36
sequential file 645
server
  EDRSQL 139
server consolidation 35
Service Level Agreement 284
Service Level Management 281, 284
services
  business partner 458
  Metro Mirror copy services 442
  System i Copy Service 438
  System i copy service 454
  System i copy services 456
SET CONNECTION DEFAULT 647
shared journal 103
shared processor pool 326
short dump 352
Short Message Service 255
SID 188, 603
simultanious multithreading 34
single level storage 52
sizing 370
sizing of SAP systems 43
SLA 284
SLIC 417
SLM 284
SM02 270
SM04 261, 270, 287
SM12 262, 288
SM13 261, 270, 287
SM21 264, 274, 289, 351, 356, 358, 504
SM22 358
SM35 270, 288, 696
SM36 653, 659–660
SM37 269, 288, 363
SM49 653, 655–656
SM50 84, 253, 261, 269, 287, 346, 348–349, 352, 502
SM51 221, 260, 356
SM59 667
SM62 660
SM66 261, 500
SM69 653–654, 658
SMS 255
SMT 34
SNA Character Stream 225
Software Change Management 282
Software Development Kit 165
Solution Manager 280, 496
SP Stack 160, 338
SP stack 338
SP01 251, 253, 289
SP12 289
SPAD 219, 238–239, 243–245, 248, 250
SPCN 380
SPOI 242
spool request 224
spool work process 221
SQE 551, 611, 641
SQL 628, 647
SQL collection 59
SQL long name 61
SQL Query Engine 611
SQL statements 526
SQL Table Statistics 64
SQL trace 352, 518
    transaction 614
SSAA 277, 499
SSCR 204
ST02 180, 290, 359, 515
ST03 289, 509
ST04 505, 523
ST05 352, 518, 608, 614
ST06 290, 489, 585
ST07 506
ST11 352, 355–356
ST22 263, 288, 352
Standard Print Control table 247
star schema 602
start profile 177
Start Report 659
STARTSAP 344
STAT 290
static LPAR 34
static SQL 64
STMS 196, 198, 209
stop save and delete journal receivers 412

STOPSAP 400
storage capacity 322
storage management 5
STRASPBAL 322–323, 371
stream file 158, 652
StreamServe 227
STROSCOL 373
STRPFRCOL 485
STRPFRT 589
STRPRTWTR 228
STRRCYBRM 434
STRREPORT 628, 659–660
STRRMTWTR 228
subsystem 54, 464
Supply Chain Management 36
Support Package 338
Support Package Stack 160
support package stack 338
surrogate identifier 603
surrogate identifiers 603
switchable device 454
switchable device environment 438, 455
swtichable IASP 451
symbolic link 66, 393
symbolic name 247
synchronous mode 448
Sys log 543
SYSINDEXES 525
system 392
System Administration Assistant 277, 499
system ASP 452
system copy 138
System Copy Guide 329
system failure 378
System i 3
System i Copy Service 438
System i copy services 454, 456
System i Informational APAR 718
System i value proposition 6
system job 54
System Licensed Internal Code 417
system log 351, 504
system loss 378
System Monitoring 281
system pools 582
System Power Control Network 380
System report 590
system value
    QDBFSTCCOL 64
    QSTGLOWACN 322
    QSTGLOWLMT 322
    QSTRUPPGM 310

T

technology independent machine interface 53
Technology Solutions Center (TSC) 718
temporary sequential database 220
temporary storage 322
TemSe 220
TemSe data 221, 224
test client 186
three-system landscape 189
thresholds
    CPU utilization 538
    disk arm utilization 541
    disk space 541
    IOP utilization 542
    LAN utilization 542
    memory utilization 539
throughput 44
TIMI 53
Tivoli Storage Manager 383
TMS 69, 185
TO 185
tool
    CHKXDA 83, 353
    DSPCURSQL 83
    IBM Insight Tool 48
        for SAP 324
        Performance Tools 325
    LODSAPKRN Installation Tool 154
    SAP Early Watch 325
    SAP tool
        Instance Profile Check Tool 183
    SQL trace 352
    Toolbox driver tool 71, 91
    Visual Admin Tool 648
    Visual Administration 355
TopNote 338
Total Cost of Ownership (TCO) 373
tp 185
trace
    buffer 519
    developer 352
    enqueue 518
    performance 518
    RFC 519
    SQL 518
    XDA 353
trace file 351
trace level 352
transaction SSAA 277
transparent table 630, 632
transparent tables 628
transport domain 192
transport domain controller 192
transport group 192
Transport Management System 69, 185, 196, 198
Transport Organizer 185, 207
transport request 185
transporting 185
TSC (Technology Solutions Center) 718
TSM 383
tune
    database 549
    SAP system 543
    System i 536
tuning
    instance profile 326
two-system landscape 189

U
UCS-2 39, 330
UCS2 39, 330
UDFS 66
unconditional mode 210–211
UNICODE 646
Unicode 38–39, 330, 627, 634–636, 647
Uninterruptible Power Supply 380
Unique Data Records flag 619
update work process 262
UPDDTA 634
upgrade guide 140
UPS 380
user-defined file system 66
utilization 34

V
V1 update 262
V2 update 262
VARCHAR 645
vendor SQL 650
virtual processor 34
virtualization 34, 36
Visual Administration tool 355, 648
Visual Explain 614

W
Walldorf 717
Web Application Server 344
Windows 32-bit 646
Windows 64-bit 646
WLE 47
work management 628
work process 346
work processes 627
workload analysis 509
Workload Estimator (WLE) 47
workload growth 324
WRKACTJOB 350, 477
WRKDSKSTS 322, 477
WRKENVVAR 135
WRKJRN 424
WRKJRNA 410
WRKLNK 54, 72
WRKMEDIBRM 432
WRKOUTQ 228, 237
WRKPID (Work with Job by PID) 348
WRKPRB 336
WRKQRY 635
WRKSBSJOB 481
WRKSYSACT 482, 590
WRKSYSSTS 322, 475
WRKWTR 228

X
XDA trace 353
XSM 438, 455

Z
Z8189AB1 630
Z8189AB2 652
Z8189AB3 653
Z8189AB4 653
Z8189AB5 647, 651
Z8189SQL 647

Implementing SAP Applications on the IBM System i Platform with IBM i5/OS

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Implementing SAP Applications on the IBM System i Platform with IBM i5/OS

Comprehensive yet simple SAP business applications with i5/OS

Successful business solutions in an on demand complex world

SAP applications and System i integration

This IBM Redbook features a collection of knowledge gained from IBM and SAP system experts who work with customers that use SAP applications on the IBM System i platform with i5/OS. It is written to assist SAP basis consultants and other IT professionals with the implementation of a total business solution, consisting of System i servers, i5/OS Version 5 Release 4 (V5R4), DB2 UDB for i5/OS database, and SAP products.

The ideas in this book apply to the entire spectrum of business solutions within the mySAP Business Suite. These solutions share SAP NetWeaver, the common technical application platform that includes the SAP Web Application Server and SAP Business Information Warehouse.

The first part of this book features an overview of SAP applications and landscape options for i5/OS followed by a discussion of SAP and i5/OS implementation and infrastructure. Then, we provide information about operational topics such as administration, printing, and how to avoid problems. After that, there is a detailed discussion about availability and backup and recovery and a section dedicated to performance topics. The book ends with an examination of the different ways that you can exchange data between your SAP applications and other applications.

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

SG24-7166-00 ISBN 0738495905