IBM i
7.6

Systems management
Recovering your system

Note

Before using this information and the product it supports, read the information in “Notices” on page 513.

This edition applies to IBM® i 7.6 (product number 5770-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC41-5304-13.

This document might contain references to Licensed Internal Code. Licensed Internal Code is Machine Code and is licensed to you under the terms of the IBM license Agreement for Machine Code.

© Copyright International Business Machines Corporation 1997, 2025.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures xiii

Tables xv

Recovering your system xix

What's new for IBM i 7.6 xxi

Part 1. Saving information on your system 1

Chapter 1. Saving the system 3
Overview of the GO SAVE command 3
GO SAVE command menu options 4
GO SAVE: Option 21 (saving the entire system) 6
Performing a complete save by using the GO SAVE checklist 6
Features that affect your backup 12
Are you using user-defined file systems 12
Are you using software encryption for tapes 12
Do you use virtual storage 12
Do you use independent disk pools 13
Are independent disk pools configured to switch between systems in a cluster 13
Do you use IBM MQ for IBM i on this system 13
Do you use network servers 13
Do you use the Hardware Management Console 13
Printing system information 14
Save considerations when checking out objects 16
ObjectConnect function 16

Part 2. Recovering information on your system 17

Chapter 2. Restoring the system: Introduction 19
The relationship between save and restore commands 21
What happens when you restore objects 22
Sequence for restoring related objects 25
Deferring the restore of dependent objects 26
Putting your system in a restricted state 27
RCLSTG command 28
Reclaiming storage 32
Reclaiming object links 33
Controlling restoration of security-sensitive objects 34
QALWOBJRST system value 34
Setting the QALWOBJRST system value to allow complete recovery 35
Setting the QALWOBJRST system value to restrict restore operations 36
QVFYOBJRST system value 36
Setting the QVFYOBJRST system value to allow complete recovery 37
Setting the QVFYOBJRST system value to restrict restore operations 37
Minimizing object scans after restore operations 37
Setting the QPFRADJ system value for a different processor or memory 38
Locked objects while restoring 38

Verifying whether objects are restored successfully..... 39
Recovery from an unsuccessful restore operation.....41
Recovering from an error while restoring libraries..... 41
Recovering from an error while restoring document library objects.....42
Recovering documents and folders..... 42
Recovering from unsuccessful system signon.....43
Performing a normal initial program load.....43
Parallel restore operations.....44
Restore operation output file information..... 44
Field descriptions for restore output file..... 47

Chapter 3. Selecting the appropriate recovery strategy.....53
Some common recovery terminology.....53
Recovery procedure for a power failure..... 55
Recovery procedure for a system failure..... 55
Recovery procedure for a program failure or human error..... 55
Recovering the vital product data information if the partition does not IPL in A-mode or B-mode.. 56
Choosing the recovery procedure for a disk failure or disk errors..... 56
Checklist 1: Actions for load-source disk unit failure..... 59
Checklist 2: Actions for load-source disk unit failure..... 60
Checklist 3: Actions for load-source disk unit failure..... 61
Checklist 4: Actions for load-source disk unit failure..... 62
Checklist 5: Actions for load-source disk unit failure..... 66
Actions for non-load source disk unit failure or disk units in basic user ASP disk failure–Checklist 6..... 70
Checklist 7: Actions for non-load source disk unit failure.....71
Checklist 8: Actions for non-load source disk unit failure.....72
Checklist 9: Actions for non-load source disk unit failure.....73
Checklist 10: Actions for non-load source disk unit failure.....77
Checklist 11: Actions for a failure in a basic auxiliary storage pool disk unit..... 81
Checklist 12: Actions for a failure in a basic auxiliary storage pool disk unit..... 82
Checklist 13: Actions for a failure in a basic auxiliary storage pool disk unit..... 84
Checklist 14: Actions for non-load source disk unit failure.....87
Checklist 15: Actions for non-load source disk unit failure.....87
Checklist 16: Actions for non-load source disk unit failure.....88
Checklist 17: Actions for independent auxiliary storage pool disk failure.....89
Checklist 18: Actions for a failure in an independent auxiliary storage pool disk unit.....90
Checklist 19: Actions for a failure in an independent auxiliary storage pool disk unit.....91
Checklist 20: Recovering your entire system after a complete system loss.....93
Checklist 21: Recovering your entire system after a complete system loss including independent auxiliary storage pools.....97
Checklist 22: Restoring a logical partition to another logical partition.....101
Checklist 23: Actions for a failed cache card.....104
Choosing the procedure to recover user information..... 105
Checklist 24: Recovering user information using commands.....107
Checklist 25: Using option 21 from the restore menu..... 110
Checklist 26: Using options 22 and 23 from the restore menu..... 113
Checklist 27: Recovering user information using tapes from Operational Assistant backup.... 116

Chapter 4. Recovering the Licensed Internal Code.....121
Preparing for loading the Licensed Internal Code.....122
Task 1: Getting ready to load the Licensed Internal Code..... 122
Task 2: Powering down the system.....123
Task 3: Preparing to start the system from an alternate IPL device..... 124
Task 4: Loading the Licensed Internal Code from media..... 124
Loading the Licensed Internal Code.....128
Recovering the vital product data information if the partition does not IPL in A-mode or B-mode131

Setting up your disk configuration after installing the Licensed Internal Code and initializing the system....132
Using Navigator for i to set up your disk configuration after you install Licensed Internal Code....133
Recovering your disk configuration....136
Starting your system after restoring the Licensed Internal Code....138

Chapter 5. Restoring the operating system....141
Choosing the right procedure for restoring the operating system....142
Loading the operating system that uses a manual initial program load (IPL)....142
Tasks to restoring the operating system....143
Task 1: Starting to restore the operating system....143
Task 2: Selecting the installation options....147
Task 3: Selecting initial program load options....151
Task 4: Setting major system options....153
Task 5: Defining or changing the system at initial program load (IPL)....153
Task 6: Completing the initial program load....155
Recovering from SRC A900 2000....157
Creating a configuration for other tape units....157

Chapter 6. Starting the system after it ends abnormally....159
What happens when your system stops....159
Using the Disk Configuration Error Report display....160
Using the Main Storage Dump Occurred display....160
Restarting your system....161
Task 1: Performing an attended initial program load....161
Task 2: Editing the rebuilding of access paths....163
Task 3: Editing check pending constraints....164
Task 4: Recovering from damaged objects and unreadable sectors....165
Recovering a damaged journal....167
Recovering a damaged journal receiver....168
Recovering a journaled object that is damaged or not synchronized....169
Recovering damaged objects in the integrated file system....169
Recovering other types of damaged objects....170

Chapter 7. Recovering information in a user ASP....171
Describing the contents of your user ASP....171
Choosing the procedure to recover user ASPs....172
Recovering a basic user ASP after recovering the system ASP....172
Task 1: Reclaiming storage....173
Task 2: Restoring user profiles....174
Task 3: Restoring the configuration....175
Task 4: Recovering journals and journal receivers in the QRCL library....175
Task 5: Restoring libraries to the system auxiliary storage pool....176
Task 6: Restoring document library objects to the system auxiliary storage pool....177
Task 7: Restoring user-defined file systems to the user auxiliary storage pool....178
Recovery steps for unmounted user-defined file systems in a user ASP....178
Recovery steps for a mounted UDFS if data is not restored....178
Recovery steps for a mounted UDFS if data is restored....178
Task 8: Reclaiming document library objects....179
Task 9: Recovering save files from the QRCL library....179
Task 10: Associating journal receivers with journals....179
Task 11: Restoring object ownership....181
Recovering an overflowed user ASP....182
Resetting an overflowed user ASP without an initial program load IPL....182
Resetting an overflowed user ASP during an IPL....184
Deleting overflowed objects during recovery....186
Recovering a damaged basic ASP....186
Task 1: Restoring user profiles....186

Task 2: Determining the contents of the lost auxiliary storage pool....187
Task 3: Determining tasks to restore objects....188
Task 4: Restoring libraries to a basic auxiliary storage pool....188
Task 5: Restoring journals to a basic auxiliary storage pool....188
Task 6: Restoring documents to a basic auxiliary storage pool....190
Task 7: Restoring user-defined file systems to a basic auxiliary storage pool....190
Recovery steps for unmounted user-defined file systems in a basic ASP....190
Recovery steps for a mounted UDFS if data is not restored....190
Recovery steps for a mounted UDFS if data is restored....191
Task 8: Restoring journal receivers to a basic auxiliary storage pool....191
Task 9: Restoring save files to a basic auxiliary storage pool....191
Recovering an independent ASP....192
Task 1: Restoring user profiles....192
Task 2: Determining tasks to restore objects to an independent auxiliary storage pool....192
Task 3: Restoring libraries to the independent auxiliary storage pool....193
Task 4: Restoring user-defined file systems to the independent auxiliary storage pool....194
Recovery steps for unmounted user-defined file systems in an independent ASP....194
Recovery steps for a mounted UDFS if data is not restored....194
Recovery steps for a mounted UDFS if data is restored....195
Task 5: Restoring authority for an independent auxiliary storage pool....195
Removing a failed disk unit from the system ASP....196
Task 1: Accessing dedicated service tools....196
Task 2: Deleting the auxiliary storage pool data....197
Task 3: Removing the disk unit from the auxiliary storage pool configuration....198

Chapter 8. The restore menu....201
What the restore menu options do....201
Using restore menu options 21, 22, and 23....203

Chapter 9. Restoring specific types of information....209
Recovering system information....209
Recovering TCP configuration information....210
Sequence for restoring security information....210
Restoring user profiles....211
What happens when you restore user profiles....213
What you should know about restoring user profiles....214
How the system establishes ownership for restored objects....216
How the system establishes the authorization list for a restored object....216
How the system establishes the primary group for restored objects....216
Restoring object authorities....217
Overview of restoring authorities....217
Restoring private authorities in a full-system recovery....218
Restoring private authorities for selected objects....218
Restoring authority on a system in a nonrestricted state....219
What you should know before running the RSTAUT command....220
Job log considerations....221
Restoring authority on a system in a restricted state....223
Examples: How the system restores authority....223
Restoring authorities for independent ASP data....225
Restoring configuration objects....227
Correcting problems with the SRM information....228
Recovering devices that will not vary on....229
Recovering when you change the console type....230
Recovering the System/36 environment configuration....230
Restoring logical partitions....231
Restoring libraries....231
Restoring a library from a previous release....232
Restoring multiple libraries....232

Considerations and restrictions of the Restore Library command....233
Restoring all libraries from a single save operation....233
Restoring all libraries from multiple save operations....234
Restoring objects....235
Restoring journaled objects....236
What happens when you restore journaled objects to a different library or directory....237
Restoring UDFSs....238
Restoring an unmounted UDFS....238
Restrictions while you restore an unmounted UDFS....238
Restoring an individual object from an unmounted UDFS....238
Restoring a mounted UDFS....239
Restore actions for mounted UDFS....239
Restoring database files....240
Comparing file attributes during a restore operation....241
How the system matches file members during a restore operation....243
Restoring members to a file....244
Restrictions on the FILEMBR parameter....244
Restoring logical files....244
How the system restores access paths....246
Examples: Restoring a file network....247
Preventing the system from rebuilding a large access path....248
How the system restores files with shared formats....248
How the system restores files with referential constraints....249
Example: Referential constraint network....250
How the system restores files with triggers....250
Restoring SQL materialized query tables....251
Deleting a physical file....252
Restoring physical files that are temporal....252
Restoring journals and journal receivers....252
Restoring journals....253
Deleting a journal....253
Restoring journal receivers....255
Resolving name conflicts when restoring journal receivers....255
Correcting the journal receiver directory....255
Deleting a journal receiver....256
How the system restores programs....256
Restoring programs to a different release....258
Restoring save file data....258
Saving and restoring spooled files....259
Restoring licensed programs....260
Restoring documents and folders....260
RSTDLO command options....261
Using multiple concurrent DLO commands....261
Output from the RSTDLO Command....261
Considerations and restrictions....262
Moving documents....262
Searching tape files....262
Selecting files from DVD-RAM optical media....262
Searching index database errors....262
Authority required to restore DLOs....262
How the system restores new DLOs....263
How the system restores existing DLOs....263
Size limitations when restoring DLOs....263
Restoring folders....263
Renaming documents when restoring....264
How the system restores descriptive information for DLOs....264
How the system restores authority and ownership for DLOs....264
When to run the RNMDIRE command....265

When to Run the RNMDLO command ... 265
Restoring objects in directories ... 265

Chapter 10. Restrictions when using the RST command ... 269

Chapter 11. Restoring PTFs ... 273

Chapter 12. Restoring system information ... 275

Chapter 13. Restoring changed objects and apply journaled changes ... 277
Task 1: Restoring changed objects ... 277
Restoring changed objects by library ... 278
Restoring changed objects individually ... 278
Task 2: Restoring changed objects in directories ... 279
Task 3: Determining whether you need to apply journaled changes ... 280
Task 4: Determining what journal receivers to use ... 280
Task 5: Applying journaled changes for user journals ... 282
Task 6: Applying journaled changes for the QAOSDIAJRN journal ... 284
Task 7: Replaying journaled changes in a disaster recovery ... 285
Task 8: Restoring changed documents and folders ... 285

Chapter 14. Mirrored protection recovery actions ... 287
System actions for permanent errors ... 287
Suspending mirrored units ... 288
Resuming mirrored units ... 289
Replacing a mirrored unit ... 289
Using spare nonconfigured units for replacement ... 291
Mirrored protection recovery actions performed by the service representative ... 292
Actions when concurrent maintenance is possible ... 292
Actions when concurrent maintenance is not possible ... 292
Other recovery considerations for mirrored protection ... 293
Mirrored protection disk-error handling ... 293
Missing disk units ... 294
Saving a unit ... 295
Restoring a unit ... 295
Active mirrored load source failure ... 295
System cannot find active mirrored load source for IPL ... 295
Active mirrored load source being used for IPL fails ... 296
Active mirrored load source fails late in the IPL or at run time ... 297
Cannot read system configuration data from active mirrored load source ... 297
Unknown unit 1 status ... 297
Recovering the state of the unknown load source ... 298
Displaying incorrect Licensed Internal Code installation ... 298
Recovering remote load source mirroring ... 299
Recovering after the remote load source has failed ... 299
Recovering after the local load source has failed on the system ... 299
Dedicated recovery of local load source – Local system is still functional ... 300
Dedicated recovery from remote disk units – After a local site disaster ... 300
Using the recover mirrored load source function ... 300

Chapter 15. Restoring your system using Operational Assistant tapes ... 303
Restoring your libraries ... 304
Restoring libraries that you saved by using a backup list ... 305
Restoring changed objects that you saved by using Operational Assistant ... 305

Chapter 16. Restoring encrypted backups and encrypted ASPs ... 307
Recovering from an encrypted backup using software encryption ... 307
Recovering from an encrypted backup using an encrypted tape ... 308

Backing up the Encryption Key Manager....308
Restoring encrypted auxiliary storage pools....309

Part 3. Release-to-release support....311

Chapter 17. Release-to-release support....313
Current release to previous release support....313
Creating the object for the previous release....314
Saving the object for the previous release....315
Testing the object on the current release....322
Restoring and using the object on the previous release....322
Restrictions for current release to previous release support....322
Previous release to current release support....322
Considerations when moving system customization information....323
Restoring previous release user data to a new system....323
Prerequisites for the recovery....326
Restoring previous release user data to a new system....326
Saving spooled files....338
Restrictions when going from previous release to current release....339

Chapter 18. System synchronization: Planning and procedures....341
Synchronization methods: Overview....342
Moving changed objects....343
Saving changed objects....343
Restoring changed objects....344
Problems when restoring changed objects....346
Problems restoring journal receivers....346
Problems restoring database file members....347
Problems with object authority or ownership....348
Moving entire libraries....348
Considerations for moving entire libraries....349
Moving individual objects....349
Applying journaled changes....350
Refreshing your new system....352
Additional synchronization tips....353

Part 4. Tips for merging two or more IBM i operating systems....355

Chapter 19. Tips for merging two systems into a single system....357
Guidelines for restoring information from the development system....357

Part 5. Alternate Installation device....359

Chapter 20. Using an alternate installation device....361
Alternate installation device: Overview....361
Setting up an alternate installation device....361
Disabling an alternate installation device....364
Verifying and selecting an alternate installation device during a manual installation....364

Part 6. Disk configuration and protection....367

Chapter 21. Configuring disks and disk protection....369
Choosing the right procedure for configuring disks....369
Checklist 1: Configuring disks on a new system....370
Checklist 2: Adding disk units without device parity protection....371
Checklist 3: Adding disk units to an existing input/output adapter....372
Checklist 4: Adding a new input/output adapter....373

Checklist 5: Moving disk units between non mirrored auxiliary storage pools..... 375
Checklist 6: Moving disk units between mirrored auxiliary storage pools.....376
Checklist 7: Deleting an auxiliary storage pool.....377
Checklist 8: Removing disk units without device parity protection..... 378
Checklist 9: Removing disk units that have device parity protection from an auxiliary storage pools without mirrored protection.....379
Checklist 10: Removing disk units that have device parity protection from an auxiliary storage pool with mirrored protection.....380
Using SST and DST.....381
Starting DST.....382
Stopping DST.....383
Starting SST.....384
Stopping SST..... 384
Displaying your disk configuration..... 385
Displaying your disk configuration (hardware view)..... 385
Displaying disk hardware configuration (command method)..... 385
Displaying disk hardware configuration (service tools method)..... 386
Displaying your disk configuration (software view)..... 386
Disk configuration and status displays.....388

Chapter 22. Working with auxiliary storage pools.....391
Initializing and formatting disk units..... 391
Adding disk units to an existing ASP.....392
Creating an unencrypted ASP..... 394
Creating an encrypted ASP..... 397
Changing the storage threshold for an ASP.....399
Changing the storage threshold for the system ASP.....400
Moving a disk unit to a different ASP..... 402
Removing a disk unit from an ASP..... 403
Deleting an ASP.....405
Calculating space requirements for an auxiliary storage pool.....406
Displaying the objects in a user ASP..... 406
Balancing an ASP..... 407
Capacity balancing.....407
Usage balancing.....407
HSM balancing..... 407
Media preference balancing..... 408
Transferring objects between ASPs.....408
Transferring a library to a different ASP.....408
Transferring a folder to a different auxiliary storage pool..... 409
Transferring journals and objects to a different ASP..... 409
Creating objects in a library user ASP..... 410
Placing a document in a basic ASP–Example..... 411
Example: Placing an object in a user auxiliary storage pool.....411
Example: Creating a user-defined file system in user auxiliary storage pool..... 412
Placing journal receivers in a user ASP..... 412
Placing journal receivers in a user ASP..... 412
Moving journal receivers from an overflowed basic user auxiliary storage pool..... 413
Resetting a journal with a status of overflowed.....413
Working with nonlibrary user ASPs..... 415
Creating objects in a nonlibrary user ASP.....415
Transferring an object to a nonlibrary ASP..... 416
Transferring a journal to a nonlibrary user ASP..... 416
Placing journal receivers in a nonlibrary basic ASP..... 417

Chapter 23. Working with device parity protection..... 419
Starting device parity protection..... 419
Starting device parity protection for an input/output adapter..... 419

Stopping device parity protection....421
Stopping device parity protection on an input/output adapter....421
Removing disk units that have device parity protection for System i while powered on....422
Removing disk units that have device parity protection for System i while powered off.... 424
Removing disk units that have device parity protection for a IBM i.... 426
Including a disk unit in device parity protection.... 427
Excluding a disk unit from device parity protection....429
Displaying device parity status.... 430

Chapter 24. Working with mirrored protection....431
Mirrored protection: Configuration rules....431
Starting mirrored protection....431
What the system does when you start mirrored protection....433
Mirrored protection configuration errors....433
Stopping mirrored protection.... 434

Chapter 25. Working with disk compression.... 437
Introduction to disk compression....437
Restrictions and considerations....437
Disk compression and capacity....438
Disk unit full considerations....439
How the system responds to disk unit full 448 reference code A6xx 0277 440
SRC code A6xx 0277....441
User action 1....441
User action 2....441
User action 3....442
User action 4....442
Examples of A6xx 0277....443
Starting disk compression.... 443
Stopping disk compression....445
Procedural sequences for configuring disks and protection.... 447
Checklist: Adding a new I/O compression-capable storage controller.... 447
Checklist: Adding disk units to an existing compression-capable storage controller....448
Checklist: Moving disk units from the system ASP to a user ASP.... 449
Recovering from error codes....450
Recovering from SRC 6xxx 7051....450
Recovering from SRC 6xxx 7052....451

Chapter 26. Managing auxiliary storage pools....453
Working with ASP trace and ASP balance.... 453
Capacity balance....454
HSM balance.... 455
Usage balance....455
Auxiliary storage pool trace....456
Determining adequate disk storage....456

Part 7. Backup and recovery tools and techniques.... 457

Chapter 27. Examples: Techniques and programming for backup and recovery....459
Examples: Backup and recovery....459
Using the RTVJRNE command in a program....459
CL program to handle escape conditions....459
Writing output to save media using the Receive Journal Entry command.... 460
Considerations when writing to tape....461
Journal entries written to an ICF file....462

Appendix A. Licensed Internal Code installation error displays.... 465

Appendix B. Example: Disaster recovery plan..... 473
Section 1. Example: Major goals of a disaster recovery plan..... 473
Section 2. Example: Personnel..... 473
Section 3. Example: Application profile..... 474
Section 4. Example: Inventory profile..... 474
Section 5. Information services backup procedures..... 475
Section 6. Disaster recovery procedures..... 476
Disaster action checklist..... 476
Recovery startup procedures for use after actual disaster..... 477
Section 7. Recovery plan for mobile site..... 477
Mobile site setup plan..... 478
Communication disaster plan..... 478
Electrical service..... 478
Section 8. Recovery plan for hot site..... 478
Alternate-site system configuration..... 479
Section 9. Restoring the entire system..... 479
Section 10. Rebuilding process..... 480
Section 11. Testing the disaster recovery plan..... 480
Section 12. Disaster site rebuilding..... 482
Vendors..... 483
Floor plan..... 483
Section 13. Record of plan changes..... 483

Appendix C. Recovering your system..... 485

Appendix D. Recovering your system to a different system..... 499

Notices..... 513
Programming interface information..... 514
Trademarks..... 514
Terms and conditions..... 515
Bibliography..... 517

Index..... 519

Figures

1. Save commands and menu options....4

2. Restore procedures....20

3. Save procedures and restore procedures for file systems....21

4. User ASP configuration before failure....171

5. Basic user ASP configuration after restoring operating system....173

6. User ASP configuration after reclaiming storage....174

7. User ASP configuration after recovering isolated journal receiver....176

8. Restore menu–first display....201

9. Sample job log for RSTAUT on a system in a restricted state....221

10. Expanded text for message CPF3736....221

11. Expanded text for message CPF3845....221

12. Sample job log for RSTAUT on a system in a nonrestricted state....222

13. Expanded text for message CPF3845....222

14. Save data from an independent ASP and restore it to an independent ASP with an identical name.. 225

15. Save data from *SYSBAS and restore it to an independent ASP....226

16. Save data from an independent ASP and restore it to an independent ASP with a different name.... 226

17. Save data from an independent ASP and restore it to *SYSBAS....227

18. Example: Restoring a journaled object to a different library....237

19. Example of a database file with two members....240

20. Restoring a copy of a file....241

21. Restoring database files with different creation dates....242

22. Restoring database files with different creation dates....242

23. Restoring access paths....247

24. Restoring a referential constraint network..... 250

25. An object with hard links–Example..... 266

26. An object with a symbolic link–Example..... 267

27. Sample recovery time line..... 277

28. Receiver directory–Saving attached receivers..... 281

29. Receiver directory–Saving detached receivers..... 281

30. How the system is saved with Operational Assistant backup..... 303

31. Recovery steps for restoring previous release user data to a new system..... 325

32. Overview of synchronization process..... 342

33. Display hardware resource listing..... 386

34. Program for retrieving journal entries..... 459

35. Program for writing the RCVJRNE command output to save media..... 461

Tables

1. Spooled files created by the system....15
2. Relationship between save and restore commands....21
3. Restoring objects with ALWOBJDIF....23
4. Possible values for the OPTION parameter of the ENDSBS command....28
5. Resolving problems detected by the RCLSTG procedure....30
6. Possible values for the QALWOBJRST system value....35
7. Lock type needed for restore operation....38
8. QASRRSTO (restore output file) information....44
9. Choosing the correct recovery procedure for disk media failure....57
10. Recovery checklist for disk failure–Checklist 1....59
11. Recovery checklist for disk failure–Checklist 2....60
12. Recovery checklist for disk failure–Checklist 3....62
13. Recovering from a disk failure–Checklist 4....63
14. Recovering from a disk failure–Checklist 5....67
15. Recovery checklist for disk failure–Checklist 6....70
16. Recovery checklist for disk failure–Checklist 7....71
17. Recovery checklist for disk failure–Checklist 8....72
18. Recovering from a disk failure–Checklist 9....74
19. Recovering from a disk failure–Checklist 10....78
20. Recovery checklist for disk failure–Checklist 11....82
21. Recovery checklist for disk failure–Checklist 12....82
22. Recovery checklist for disk failure–Checklist 13....85
23. Recovery checklist for disk failure–Checklist 14....87

24. Recovery Checklist for disk failure–Checklist 15....88
25. Recovery checklist for disk failure–Checklist 16....88
26. Recovering from a disk failure–Checklist 17....90
27. Recovering from a disk failure - Checklist 18....90
28. Recovering from a disk failure–Checklist 19....91
29. Recovery checklist for complete system loss Checklist 20....94
30. Recovering from a complete system loss Checklist 21....97
31. Recovering from a complete system loss -Checklist 22....102
32. Recovery checklist for failed cache card–Checklist 23....104
33. Choosing the correct recovery procedure for user information....106
34. Recovering user information by using commands....107
35. Recovering user information with option 21....111
36. Recovering user information by using options 22 and 23....113
37. Recovering user information from Operational Assistant backup tapes....116
38. Options from the Install the Licensed Internal Code (LIC) Menu....121
39. SRC codes when loading the Licensed Internal Code....125
40. Configuring disk while installing the operating system....145
41. Recovery for damaged objects by object type....166
42. Object types that require special procedures for deleting....183
43. Tasks for restoring basic ASP objects....188
44. Example of restore order for independent ASPs saved with GO SAVE: Option 21 or 23....192
45. Tasks for restoring independent ASP objects....193
46. Commands for changing system information....209
47. How user profiles are restored....211
48. Copying user profiles without affecting the Digital Certificate Manager information....212

49. Results of restoring user profiles....213
50. Restoring an object linked to an authorization list....216
51. Saving and restoring private authorities using the PVTAUT parameter....219
52. How configuration objects are restored....227
53. Methods for restoring all libraries–single save operation....234
54. Methods for restoring all libraries–multiple save operations....234
55. Restoring a file network....247
56. Restoring files that have trigger programs....251
57. Save and restore spooled files....259
58. Restoring objects that have hard links....267
59. Using the RST Command for QSYS.LIB Objects....270
60. New name options on the RST command–Examples....271
61. Restore procedures for changed objects....278
62. Values for TGTRLS parameter....313
63. Language support for the target release parameter....314
64. Previous-release support by object types....316
65. Comparison of synchronization methods....343
66. Choosing the right disk procedure....369
67. Configuring disks on a new system–Tasks....370
68. Adding disk units without device parity protection–Tasks....372
69. Adding disk units to an existing IOA–Tasks....373
70. Adding a new IOA–Tasks....374
71. Moving disk units between ASPs–Tasks....375
72. Moving disk units between ASPs with mirrored protection–Tasks....376
73. Deleting an user ASP–Tasks....377

74. Removing disk units that do not have device parity protection–Tasks..... 378

75. Removing disk units from an IOA and a non-mirrored ASP–Tasks..... 379

76. Removing disk units from an IOA and a mirrored ASP–Tasks..... 380

77. Word formats for SRC codes.....441

78. Adding a new I/O storage controller and disk units.....447

79. Adding compressed disk units to an existing storage controller.....448

80. Moving disk units from the system ASP to a user ASP..... 449

81. Checklist for testing the disaster recovery plan.....480

Recovering your system

This topic provides general information about backup and recovery options for the IBM i environment. It describes the options available on the system, compares and contrasts them, and tells where to find more information about them.

This topic covers the following information:

- Procedures on how to save your system using the GO SAVE commands.
- Restoring to a different release of IBM i.
- Selecting the right recovery strategy.
- Procedures for restoring information to your system.
- Device parity protection procedures.
- Mirrored protection procedures.

Related information

Backing up your system

© Copyright IBM Corp. 1997, 2025

What's new for IBM i 7.6

Read about new or changed information for the Recovering your system topic collection.

ObjectConnect documentation moved

Documentation for ObjectConnect functions were moved to ObjectConnect.

How to see what's new or changed

The documentation uses the following symbols to indicate where technical changes were made:

- The » image to mark where new or changed information begins.
- The « image to mark where new or changed information ends.

In PDF files, you might see revision bars (|) in the left margin of new and changed information.

To find other information about what's new or changed in this release, see the Memo to users.

© Copyright IBM Corp. 1997, 2025

Part 1. Saving information on your system

© Copyright IBM Corp. 1997, 2025

Chapter 1. Saving the system

If this is your first experience with your IBM i model, use the instructions in this topic collection to save all of the information about your system. Do this with the GO SAVE menu options.

You can browse the Backing up your system topic in IBM i documentation, or print a copy of the information about how to back up your entire system.

Overview of the GO SAVE command

Use the GO SAVE command to save your entire system or parts of your system that change regularly.

Using the GO SAVE command is a simple way to make sure that you have a good backup of your entire system. The GO SAVE command presents you with Save menus to back up your system, no matter what backup strategy you decide to use. It is a good idea to use menu option 21 of the GO SAVE command right after you install your system.

Menu option 21 of the GO SAVE command is the basis for all save strategies. This option saves all the data on your system. After you use menu option 21, you can use other menu options to save parts of the system, or to use a manual save process.

Another save method uses Backup, Recovery, and Media Services (BRMS), which automates your save processes. BRMS provides a comprehensive solution for your backup and recovery needs.

Important: Be sure to permanently apply all Licensed Internal Code PTFs (fixes) before you use the SAVSYS command, or the **GO SAVE** menu option 21 or 22.

The following figure illustrates the commands and menu options that you can use to save the parts of the system and the entire system.

© Copyright IBM Corp. 1997, 2025

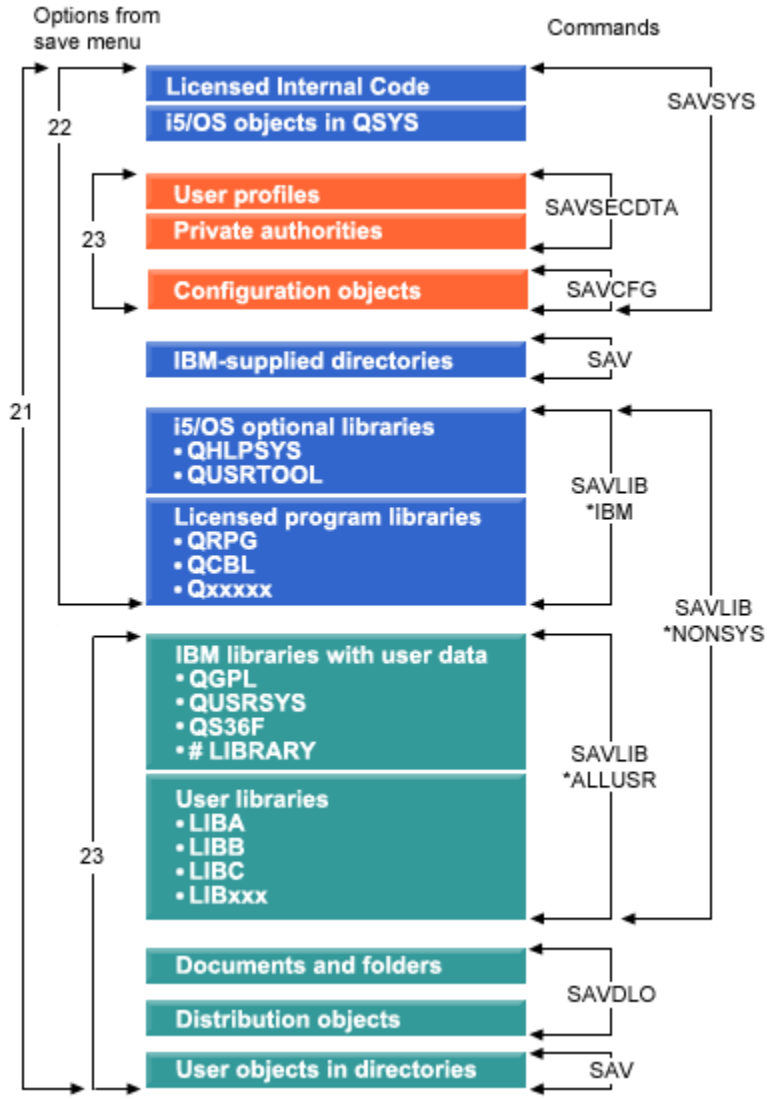

Figure 1. Save commands and menu options

GO SAVE command menu options

This information describes the GO SAVE command and the most common menu options that you can use.

Access the GO SAVE command menu by typing GO SAVE from any command line. From the Save menu, you see option 21, option 22, and option 23 along with many more save options.

```
SAVE                                Save
                                                          System:   SYSTEMA
Select one of the following:

  Save Data
     1. Files
     2. Libraries
     3. Documents and folders
     4. Programs
     5. Other objects
     6. Changed objects only
     7. Licensed programs
     8. Security data

    10. Configuration
    11. Objects in directories

                                                                        More...
Selection or command
===>

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

```
SAVE                                Save
                                                          System:   SYSTEMA
Select one of the following:

  Save System and User Data
    20. Define save system and user data defaults
 ++ 21. Entire system
 ++ 22. System data only
 +  23. All user data

  Save Document Library Objects
    30. All documents, folders, and mail
    31. New and changed documents, new folders, all mail
    32. Documents and folders
    33. Mail only


                                                                        More...
Selection or command
===>

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

A single plus sign (+) indicates that the option places your system into a restricted state if you run it without prompting for commands, which means that nothing else can be running on your system when the menu option is selected. If you do prompt for commands, you can skip the ENDSBS command that places the system in restricted state.

A double plus sign (++) indicates that your system must be in a restricted state to run this option. You cannot skip the ENDSBS command that runs as part of the option.

The options that place your system into a restricted state (options 21, 22, 23, and 40) also end TCP prior to ending the subsystems. Commands issued as part of these options are:

- ENDTCPSVR
- ENDHOSTSVR
- DLYJOB
- ENDTCP
- DLYJOB
- ENDSBS SBS(*ALL)

The two Delay job (DLYJOB) commands allow time for TCP processing to end. You can choose the amount of time that each DLYJOB command waits. The default value is 300 seconds for a total delay of 10 minutes before the End subsystem (ENDSBS) command is run.

GO SAVE: Option 21 (saving the entire system)

Option 21 saves everything on your system and you can run the save while you are not there.

Option 21 saves all of your data for other licensed programs when you select to vary off your network server. You can also back up data that is stored on a logical partition. If you have Linux® or AIX® installed on a guest logical partition, you can back up that partition when you vary off your network server.

Option 21 puts your system into a restricted state. When the save begins, no users can access your system and the backup is the only thing that is running on your system. It is best to run this option overnight for a small system or during the weekend for larger system. If you schedule an unattended save, make sure that your system is in a secure location. After you schedule the save, you will not be able to use the workstation where the backup is initiated until the save is complete.

Note: Before you run an option 21 save, if you are saving information on independent ASPs, make sure the independent ASPs are varied on.

If the save/restore master key is set when you do a full system save, the master key also gets saved.

Tip: If you are using the Hardware Management Console (HMC), you can run system saves by using Option 21 while you are at a remote location. For example, you can open a shared console on your HMC from your office, travel to another location and link to the shared session from there.

Option Number	Description	Commands
21	Entire server (QMNSAVE)	ENDTCPSVR ENDHOSTSVR DLYJOB ENDTCP DLYJOB ENDSBS SBS(*ALL) OPTION(*IMMED) CHGMSGQ MSGQ(QSYSOPR) DLVRY(*BREAK or *NOTIFY) SAVSYS SAVLIB LIB(*NONSYS) ACCPTH(*YES) SAVDLO DLO(*ALL) FLR(*ANY) SAV DEV('/QSYS.LIB/*media-device-name*.DEVD') + OBJ(('/*') ('/QSYS.LIB' *OMIT) + ('/QDLS' *OMIT))[1] UPDHST(*YES) STRSBS SBSD(*controlling-subsystem*)

[1]The command omits the QSYS.LIB file system because the SAVSYS command and the SAVLIB LIB(*NONSYS) command both save it. The command omits the QDLS file system because the SAVDLO command saves it.

"Performing a complete save by using the GO SAVE checklist" on page 6 provides you with step-by-step instructions on how to save your entire system with menu option 21 of the GO SAVE command.

Performing a complete save by using the GO SAVE checklist

Use the GO SAVE checklist to do a complete save operation.

Use the following checklist for menu options 21, 22, and 23 of the GO SAVE command. When appropriate, select the option that you require. If you choose to, you can print system information during the procedure. The Printing system information topic contains detailed instructions on how to print system information if you do not want the **Save** menu option command to print your system information automatically.

Some of the steps in this checklist might not apply to your system configuration. See Identify optional features that affect your backup for help with determining whether you use the features in your environment. If you are still unsure how your system is configured, contact your system administrator.

Attention: If you are using the Hardware Management Console (HMC), you must back up the HMC in addition to using the **GO SAVE** menu Option 21 to obtain a complete save of your system. See Backing up and recovering data.

If you need to restore the entire POWER® server, you must back up the Virtual I/O Server, which also contains Integrated Virtualization Manager, in addition to using the **GO SAVE** menu Option 21 command. Virtual I/O Server must be restored on the blade server before you can restore any IBM i client logical partitions. For saving and restoring Virtual I/O Server, see the *PowerVM Editions Operations Guide*.

1. If you are using software encryption for backup tapes and saving system data (GO SAVE option 21 or 22), set the save/restore master key before you run the save operation. The save/restore master key is a special purpose master key that is used to encrypt all the other master keys when they are saved by using a SAVSYS or GO SAVE operation. For the instructions, see Loading and setting the save/restore master key.
2. Sign on with a user profile that has *SAVSYS and *JOBCTL special authorities, and also has sufficient authority to list different types of system resources. (The QSECOFR user profile contains all of these authorities.) The sign-on makes sure that you have the authority that you need to place the system in the necessary state and to save everything.
3. If you want to exclude virtual images from a full-system save, use one of the following strategies. Virtual images can significantly increase the time that it takes to complete an Option 21 save operation, even if the image catalog entries do not contain data:
   - Use the following example for the Change Attribute (CHGATR) command to mark the image catalog directory as allow save *NO.

     ```
     CHGATR OBJ('/MYINFO') ATR(*ALWSAV) VALUE(*NO)
     ```

   - Use the Load Image Catalog (LODIMGCLG) command to make the image catalog ready. Image catalogs with a ready status are omitted from the save operation.
   - In an attended save, you can specify to omit the image catalog directories on the Save Object (SAV) command.
4. If you have independent auxiliary storage pools (ASPs), make them available if you want them to be included in an Option 21 or 23 save operation.

   Note: If your system includes geographically mirrored independent ASPs, eliminate them from the backup by taking them offline. Save independent ASPs that are geographically mirrored separate from this GO SAVE operation.

   If the geographically mirrored ASPs remain available during the GO SAVE operation, geographic mirroring is suspended when the system becomes restricted. When you resume mirroring after the save, the geographically mirrored ASPs must be synchronized. Synchronization can be a lengthy process.
5. If you are operating in a clustered environment and want to save independent ASPs without causing a failover or you want to save the cluster environment for a node, end the device cluster resource group and end clustering before you end subsystems.

   Use the End Cluster Resource Group (ENDCRG) command and the End Cluster Node (ENDCLUNOD) command.
6. If you plan to run the save procedure immediately, make sure that no jobs are running on the system, type the Work with Active Jobs (WRKACTJOB) command.

   If you plan to schedule the save procedure to run later, send a message to all users to inform them when the system is not available.
7. To run an attended save of your system, go to step 11.
8. To run an unattended save operation, continue with the following steps. An unattended save operation prevents your save operation from stopping because of unanswered messages:
   a. Display the reply list sequence numbers to find what numbers are available for use:

      ```
      WRKRPYLE
      ```

b. If MSGID(CPA3708) is not already in your reply list, add it. For *xxxx*, substitute an unused sequence number from 1 through 9999:

```
ADDRPYLE SEQNBR(xxxx) +
         MSGID(CPA3708) +
         RPY('G')
```

c. If you are using virtual media for your save media, specify automatic load in the reply list, MSGID(OPT149F), to avoid receiving a message that interrupts the unattended save operation. If necessary, virtual media uses the autoload feature to create more images with the same capacity as the last image you loaded, provided the disk storage is available.

d. Change your job to use the reply list and to notify you of any break messages that are sent:

```
CHGJOB INQMSGRPY(*SYSRPYL) BRKMSG(*NOTIFY)
```

Note: You can also set up a default so that whenever you select menu options 21, 22, or 23, the system uses the reply list. To set up the default, select menu option 20 from the **Save** menu. Specify `Yes` on the `Use system reply list` option.

9. Type `GO SAVE` at a command prompt to display the **Save** menu.
10. Select the option (21, 22, or 23) from the **Save** menu and press the Enter key.

    A prompt display describes the function of the menu option that you selected.
11. After reading the **Specify Command Defaults** prompt display, press the Enter key to continue.

```
                          Specify Command Defaults
Type choices, press Enter.
  Devices  . . . . . .   TAP01        Names
                         ----------
                         ----------
                         ----------
  Prompt for commands  . . . . .   Y            Y=Yes, N=No
  Check for active files . . . .   Y            Y=Yes, N=No
  Message queue delivery . . . .   *BREAK       *BREAK, *NOTIFY
  Start time . . . . . .   *CURRENT     *CURRENT, time
  Start date . . . . . .   *CURRENT     *CURRENT, date
  Vary off network servers . . .   *ALL         *ALL, *NONE
                                                                I
                                                                    More...
F3=Exit   F12=Cancel
```

```
                          Specify Command Defaults
Type choices, press Enter.
  Unmount file systems . . . . .   Y            Y=Yes, N=No
  Print system information . . .   N            Y=Yes, N=No
  Use system reply list  . . . .   N            Y=Yes, N=No
  Spooled file data  . . . . . .   *NONE        *NONE, *ALL
  End TCP/IP wait time . . . . .   300          *NONE, 1-9999 seconds
  Print joblog . . . . . .   Y            Y=Yes, N=No
  Start controlling subsystem      Y            Y=Yes, N=No
```

12. Type your choices for the **Devices** prompt. You can specify as many as four tape device names. If you specify more than one device, the system automatically switches to the next tape device when the

current tape is full. You can select only one DVD-RAM optical media device, tape media library device, or virtual tape device.

The first device for options 21 and 22 must be your alternate IPL device. If you are creating media to install on another system, the device must be compatible with the alternate IPL device for that system. Using compatible devices helps make sure that the system can read the SAVSYS media when you need to restore your Licensed Internal Code and the operating system.

13. Type your choice for the **Prompt for commands** prompt. Specify N (No) when you want to run an unattended save operation. Specify Y (Yes) when you want to change the defaults on the SAV*xx* commands.
14. Type your choice for the **Check for active files** prompt. Specify Y (Yes) when you want the system to warn you when active files exist on the save media. The warning that you receive gives the following choices:
    - Cancel the save operation.
    - Insert new media in the device and try the command again.
    - Initialize the current media and try the command again.

    Note: If you use DVD-RAM optical media for your save, the system sends inquiry messages to the QSYSOPR message queue when it encounters identical active files. The system sends the inquiry message for each identical active file that it finds.

    Specify N (No) when you want the system to write over any active files on the save media without warning you.
15. Type your choice for the **Message queue delivery** prompt. Specify *NOTIFY when you want to do an unattended save operation. *NOTIFY prevents messages from stopping the save operation. If you specify *NOTIFY and a severity 99 message is sent to the QSYSOPR message queue by another job, it does not interrupt the save job.
    For example, messages that request a new volume to be loaded interrupt the save operation because they are associated with the job. You cannot continue until you reply to these messages.

    Specify *BREAK when you want the save to be interrupted for severity 99 messages that require a reply.
16. Type your choices for the **Start time** and **Start date** prompts. To start the save operation immediately, specify *CURRENT for both prompts. You can schedule the start of the save operation up to 7 days from the current job date and time.

    Note:

    a. The system uses the Delay Job (DLYJOB) command to schedule the save operation. Your workstation is not available from the time that you request that the menu option until the save operation is completed.

    b. **Make sure that your workstation is in a secure location**. Your workstation remains signed on, waiting for the job to start. If the system request function is used to cancel the job, your workstation displays the **Save** menu. The workstation remains signed on with your user profile and your authority.

    c. Make sure that the value for the QINACTITV system value is *NONE. If the value for QINACTITV is other than *NONE, the workstation is varied off in the amount of time specified. If you changed the value to *NONE, write down the old value.

    d. Do the following when you specify a delayed start and you want your save operation to run unattended:
        - Set up the system reply list.
        - Specified *NONE on QINACTITV system value.
        - Specified *NOTIFY on message queue delivery.
        - Specify *NOTIFY for any break messages.
        - Responded N to the **Prompt for commands** prompt.

- Responded N to **Check for active files**.

17. Type your choice for the **Vary off network servers** prompt. If you use network servers, you can vary off the network server descriptions before you begin the save procedure. Examples of network servers include an IBM i hosting an IBM i, Linux, or AIX guest partition.

    Select one of the following options to specify which hosted network servers need to be varied off before the save operation and varied on after the save operation:

    ***NONE**
    This option does not vary off network servers. No data is saved for the network servers because the system is in a restricted state. Saving the individual objects requires an active connection to the system.

    ***ALL**
    Varies off all network servers before the save operation starts. The save operation takes less time but the network server data is not saved in a format that allows restoration of individual objects. All the data from the network servers must be restored.

18. Type your choice for the **Unmount file system** prompt. Specify Y (Yes) when you want to allow all dynamically mounted file systems to be unmounted. This step enables user-defined file systems (UDFSs) and their associated objects to be saved.

    Note: After the save operation is completed, the system attempts to remount the file systems.

    Specify N (No) when you do not want all dynamically mounted file systems to be unmounted. The file system attributes for the UDFS are saved for all UDFSs that are saved when mounted. To rebuild these mounted file UDFSs during a restore, you must specify the RBDMFS(*UDFS) parameter on the Rational® Service Tester for SOA Quality command.

19. Type your choice for the **Print system information** prompt. Specify Y (Yes) when you want to print the system information. The system information might be useful for disaster recovery. The "Printing system information" topic explains how to print your system information manually without using the automatic **GO SAVE** menu option.
20. Type your choice for the **Use system reply list** prompt. Specify Y (Yes) when you want to use the system reply list when the system sends an inquiry message.
21. Type your choice for the **Spooled file data** prompt. Specify whether spooled file data for output queues is saved by using one of the following values:

    ***NONE**
    No spooled file data is saved.

    ***ALL**
    For each output queue that is saved, all available spooled file data on the output queue is saved.

    Note: Saving spooled files might require more save media and might take more time.

22. Type your choice for the **End TCP/IP wait time** prompt. Specify the number of seconds to delay the job to allow TCP/IP processing to end before all subsystems end. Specify *NONE to end all subsystems without ending TCP/IP processing first.
23. Type your choice for the **Print job log** prompt. Specify Y (Yes) when you want to print the job log when the menu operation ends.
24. Type your choice for the **Start controlling subsystem** prompt. Specify Y (Yes) when you want to start the controlling subsystem when the menu operation ends.
25. Press the Entered key. If you chose a later start date or time, your display shows a message that indicates the start date and time of the save operation. You cannot use the display until the save operation is completed. The input-inhibited indicator might appear. The steps for setting up the save operation are complete.

    If you did not choose a later start date or time, continue with step 27. If the value for QSYSOPR message queue delivery is *BREAK with a severity level of 60 or fewer, you must respond to the End Subsystem (ENDSBS) messages. Respond to these messages even if you plan to run an unattended save operation by specifying a start date and time of *CURRENT.

26. If you responded Y to the system prompt, Prompt for commands, the End Subsystem display appears after TCP ends. Type any changes and press the Enter key. While the system is ending subsystems, you see the following messages. Respond to the messages when the QSYSOPR message queue is set to *BREAK with a severity level of 60 or fewer. Each message appears at least twice. Press the Enter key to respond to each message.
    a. CPF0994 ENDSBS SBS(*ALL) command being processed
    b. CPF0968 System ended to restricted condition

    If you responded N to the **Prompt for commands** prompt, skip to step 29.
27. When the system is ready to run each major step in the save operation, you are shown the prompt display for that step. The time between prompt displays might be long.

    Note: If independent ASPs are available, you might see more prompts for options 21 and 23, as described in the Saving independent ASP's topic.

    These prompt displays appear for option 21 (Entire system):

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

    These prompt displays appear for option 22 (System data only):

    ```
    ENDSBS SBS(*ALL) OPTION(*IMMED)
    SAVSYS
    SAVLIB LIB(*IBM) ACCPTH(*YES)
    SAV DEV('/QSYS.LIB/media-device-name.DEVD') +
        OBJ(('/QIBM/ProdData') +
            ('/QOpenSys/QIBM/ProdData')) +
        UPDHST(*YES)
    STRSBS SBSD(controlling-subsystem)
    ```

    These prompt displays appear for option 23 (All user data):

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

    Type your changes at each prompt display and press the Enter key.
28. When the system sends a message that indicates to load the next volume, load the next media and respond to the message.

    For example, if the message is the following, load the next volume and then enter R to try again (C cancels the operation):

    ```
    Device was not ready or next volume was
    not loaded (C R)
    ```

    Attention: If an unrecoverable media error occurs during the SAVLIB procedure, see Recovering from a media error during a SAVLIB operation.
29. Change the QINACTITV system value back to its original value. You wrote down this value in step 18 c.

30. When the save operation is completed, print the job log. It contains information about the save operation. Use it to verify that the operation saved all objects. Run one of the following commands:

    ```
    DSPJOBLOG * *PRINT
    ```

    Or

    ```
    SIGNOFF *LIST
    ```

    The save operation is complete. Make sure that you mark all of your media and store it in a safe, accessible place.
31. If you ended clustering before the save operation was run, restart clustering on the save node from a node where clustering is already active.

    For more information, see the online help for Cluster Resource Services, or see IBM i Clusters technology.
32. Restart the device cluster resource group to enable resiliency.

Features that affect your backup

Features that affect your backup might include user-defined files, virtual storage, encryption keys, independent disk pools, and network servers.

Are you using user-defined file systems

A user-defined file system (UDFS) is a file system that a user creates and manages. To determine whether you have any UDFS on your system, use one of the following methods.

Identify UDFSs by using IBM Navigator for i:

Using **Navigator for i**, expand ***your_system*** > **File Systems** > **Integrated file system** > **Root** > **dev** > **QASP*xx*** or select the name of an independent disk pool. If UDFS objects exist, they appear in the right panel.

Identify UDFSs by using the character-based interface:

Follow these steps to see whether you have user-defined file systems on the system.

1. At a command line, specify `wrklnk '/dev'`.
2. On the Work with Object Links display, select option 5 (Display) to display the contents of the dev directory.
3. Locate object link names that begin with QASP*xx* or the name of an independent disk pool, and select **Option 5** to display the UDFS within the auxiliary storage pool (ASP).

Are you using software encryption for tapes

If you are using software encryption for backup tapes and saving system data (GO SAVE option 21 or 22), set the save and restore master key before you run the save operation. For the instructions, see the Loading and setting the save/restore master key topic.

Do you use virtual storage

Virtual media simulates tape, CD, or DVD images that are stored directly on your system disk units. Do the following to determine whether you store virtual images in image catalogs:

1. At a command line, specify `WRKIMGCLG`.

   Note: The Work with Image catalogs (WRKIMGCLG) window displays the name of the image catalog, the status, and the virtual type.

Do you use independent disk pools

An independent disk pool is a collection of disk units that can be brought online or taken offline independent of the rest of the storage on a system. If you have the necessary authority, you can check whether independent disk pools are configured on your system. Using **Navigator for i**, expand ***your_system*** > **Configuration and Service** > **Hardware** > **Disk Units** > **Disk pools** folder. Independent disk pools are numbered 33 – 255.

Are independent disk pools configured to switch between systems in a cluster

A IBM i cluster is a collection or group of one or more systems or logical partitions that work together as a single system. If you have authority, you can check to see whether your independent disk pool is switchable between systems in a cluster.

1. Using **Navigator for i**, expand ***your_system*** > **Configuration and Service** > **Hardware** > **Disk Units** > **Disk pools** folder.
2. Independent disk pools are numbered somewhere between 33 and 255. Right-click the independent disk pool and select **Properties**.
3. On the **Disk Pool Properties** page, the General tab displays the field **Switchable: Yes** if you have configured your independent disk pool to switch between systems.

Do you use IBM MQ for IBM i on this system

The IBM MQ for IBM i (5724-H72) licensed program provides application programming services that enable indirect program-to-program communications that use message queues. This licensed program enables programs to communicate with each other across all supported platforms.

Use one of the following methods to check whether IBM MQ is installed:

Check if IBM MQ is installed by using Navigator for i:

Using **Navigator for i**, expand ***your_system*** > **Configuration and Service** > **Software** > **Installed Products**.

Check if IBM MQ is install by using the character-based interface:

1. At a command line, specify GO LICPGM.
2. Specify option 10 (Display installed licensed programs) to display installed licensed programs.

   If IBM MQ is installed, IBM MQ for IBM i appears in the Description column for licensed program 5724-H72.

   If IBM MQ is installed, the Work with Queue Managers (WRKMQM) command indicates whether any queue managers are configured.

Do you use network servers

Several solutions are available that enable other operating systems to run on your IBM i product. Examples include Linux or AIX running in a logical partition.

Do you use the Hardware Management Console

Your system might be equipped with a Hardware Management Console (HMC). An HMC is needed if you use Capacity on Demand or logical partitions.

Printing system information

Printing the system information provides valuable information about your system that is useful during a system recovery. It is especially useful if you cannot use your SAVSYS media to recover and must use your distribution media.

Printing this information requires *ALLOBJ, *IOSYSCFG, and *JOBCTL authority and produces many spooled file listings. You might not need to print this information every time you run a backup. However, print it whenever important information about your system changes.

1. Print your current disk configuration. This output is essential if you plan to do a model upgrade and you are using mirrored protection. This information is also vital if you need to recover an independent ASP. Do the following to print the disk configuration:
   a. Sign on with a user profile that has *SERVICE special authority.
   b. Type STRSST on a command line and press the Enter key.
   c. Specify the service tools user ID and service tools password. These values are case-sensitive.
   d. Select option 3 **Work with disk units** on the System Service Tools (SST) display.
   e. Select option 1 **Display disk configuration** on the Work with Disk Units display.
   f. Select option 3 **Display disk configuration protection** on the Display Disk Configuration display.
   g. Print the displays. Press the PRINT key for each display.
   h. Press F3 until you see the Exit System Service Tools display.
   i. On the Exit System Service Tools display, press the Enter key.
2. If you are operating in a clustered environment, print the cluster configuration information. Use the following commands to print cluster information:
   a. Display Cluster Information - DSPCLUINF DETAIL(*FULL) OUTPUT(*PRINT)
   b. Display Cluster Resource Group - DSPCRGINF CLUSTER(*cluster-name*) CRG(*LIST) OUTPUT(*PRINT)
3. If independent ASPs are configured, record the relationship between the independent ASP name and number. You can find this information in IBM Navigator for i. In the **Disk Units** folder, select **Disk Pools**.
4. Sign on with a user profile that has *ALLOBJ special authority, such as the security officer. The system lists information only if you have the correct authority. If you sign on as a user with less than *ALLOBJ authority, some of the listings in these steps might not be complete. You need to be enrolled in the system directory before you can print a list of all the folders on the system.
5. If you use the history log or if you need to keep it, follow these steps -
   a. Display the system log QHST. Type -

      ```
      DSPLOG LOG(QHST) OUTPUT(*PRINT)
      ```

   b. Display all copies of the system log:

      ```
      WRKF FILE(QSYS/QHST*)
      ```

      Look at the list to verify that you saved all copies of the log that you might need later.

      Note: The history (QHST) log contains information such as date created, and the last change date and time. To get more information about the history (QHST) log, select option 8 (Display file description) on the Work with Files display.
   c. To prevent confusion about the date of the log, select the `Delete` option on the Work with Files display. Delete all but the current copies of the system log. This step improves the performance of the SAVSYS command.
6. Print the system information by using one of the following methods:

a. Using the **GO SAVE** command, on the Specify Command Defaults display, select Y at the **Print system information** prompt.

b. Use the PRTSYSINF command.

The following table describes the spooled files that the system creates. The PRTSYSINF command does not create empty spooled files. If some objects or types of information do not exist on your system, you might not have all the files that are listed in the following table.

Table 1. Spooled files created by the system

Spooled File Name	User Data	Description of Contents
QPEZBCKUP	DSPBCKUPL	List of all user libraries
QPEZBCKUP	DSPBCKUPL	List of all folders
QSYSPRT	DSPSYSVAL	Current settings for all system values
QDSPNET	DSPNETA	Current settings for all network attributes
QSYSPRT	DSPCFGL	Configuration lists
QSYSPRT	DSPEDTD	User-defined edit descriptions (a separate spooled file for each)
QSYSPRT	DSPPTF	Details of all fixes that are installed on your system
QPRTRPYL	WRKRPYLE	All reply list entries
QSYSPRT	DSPRCYAP	Settings for access path recovery times
QSYSPRT	DSPSRVA	Settings for service attributes
QSYSPRT	DSPNWSSTG	Network server storage spaces information
QSYSPRT	DSPPWRSCD	Power® on/off schedule
QSYSPRT	DSPHDWRSC	Hardware configuration reports (a separate spooled file for each resource type, such as *CMN or *LWS)
QSYSPRT	WRKOPTCFG	Optical device descriptions (if your system has an optical device and optical support is started when you run the command)
QSYSPRT	DSPRJECFG	Remote job entry configurations
QPDSTSRV	DSPDSTSRV	SNADS configuration
QPRTSBSD	DSPSBSD	Subsystem descriptions (a separate spooled file for each subsystem description on your system)
QSYSPRT	DSPSFWRSC	Installed licensed programs (Software Resources List)
QPRTOBJD	DSPOBJD	A list of all the journals on your system
QPDSPJNA	WRKJRNA	The journal attributes for each journal that is not in the QUSRSYS library (a separate file for each journal). Typically, journals in the QUSRSYS library are IBM-supplied journals. If you have your own journals in the QUSRSYS library, you need to manually print information about those journals.
QSYSPRT	CHGCLNUP	Settings for automatic cleanup
QPUSRPRF	DSPUSRPRF	Current values for the QSECOFR user profile
QPRTJOBD	DSPJOBD	Current values for the QDFTJOBD job description

Table 1. Spooled files created by the system (continued)		
Spooled File Name	**User Data**	**Description of Contents**
QPJOBLOG	PRTSYSINF	The job log for this job[1]

1
On your system, this spooled file might be in the QEZJOBLOG output queue.

7. Print a list of directories in the "root" (/) directory.

```
DSPLNK OBJ('/*') OUTPUT(*PRINT)
```

8. Print any IBM-supplied objects that you modified, such as the QSYSPRT print file.
9. If you maintain a CL program that contains your configuration information, use the Retrieve Configuration Source (RTVCFGSRC) command to verify that the CL program is current.

```
RTVCFGSRC CFGD(*ALL) CFGTYPE(*ALL) +
          SRCFILE(QGPL/QCLSRC) +
          SRCMBR(SYSCFG)
```

10. Print these spooled files. Keep this information with your backup log or your save system media for future reference. If you choose not to print the lists, use the Copy Spooled File (CPYSPLF) command to copy them to database files. See Save spooled files for more information. Make sure that the database files are in a library that is saved when you run the **Save** menu option.

Save considerations when checking out objects

The topic provides information about the save considerations if you have objects that are checked out.

If you use IBM Navigator for i or the Check Out Objects (CHKOUT) command to check out objects to a specific user, check these objects back in before you run the save. When an object is checked out to a user, other users can read the object, but they cannot update the object. If an object is checked out during a save that specifies update history UPDHST(*YES), message CPFA09E is sent for the object. CPFA09E indicates that the update history function is unable to change the attributes for the object.

ObjectConnect function

ObjectConnect is a set of control language (CL) commands that can be used to copy objects between systems easily and efficiently.

The ObjectConnect commands use save commands on the source system and restore commands on the target system to copy objects from the source to the target. The save data is sent from the source to the target in buffers so an intermediate copy of the save data is not stored to disk on either system.

For more information about ObjectConnect functions, see the ObjectConnect topic.

Related tasks

Configuring the ObjectConnect server

Managing the ObjectConnect server

Related information

ObjectConnect Overview

Using the ObjectConnect commands

Part 2. Recovering information on your system

© Copyright IBM Corp. 1997, 2025

Chapter 2. Restoring the system: Introduction

This topic collection provides information about the menu options and commands that are available for restoring information and what restore commands can be used in different file systems.

Figure 2 on page 20 shows the menu options and commands that are available for restoring information. It also shows the normal sequence for restoring information, working from top to bottom. Figure 3 on page 21 shows what restoring commands can be used for information in the different file systems.

Look for comprehensive information about how to save your system. Compare these figures to see the relationship between how things are saved and how they are restored. Use them to gain a general understanding of what you need to restore and how you might do it. Use the information in Chapter 3, "Selecting the appropriate recovery strategy," on page 53 topic to plan the correct recovery strategy for your situation.

© Copyright IBM Corp. 1997, 2025

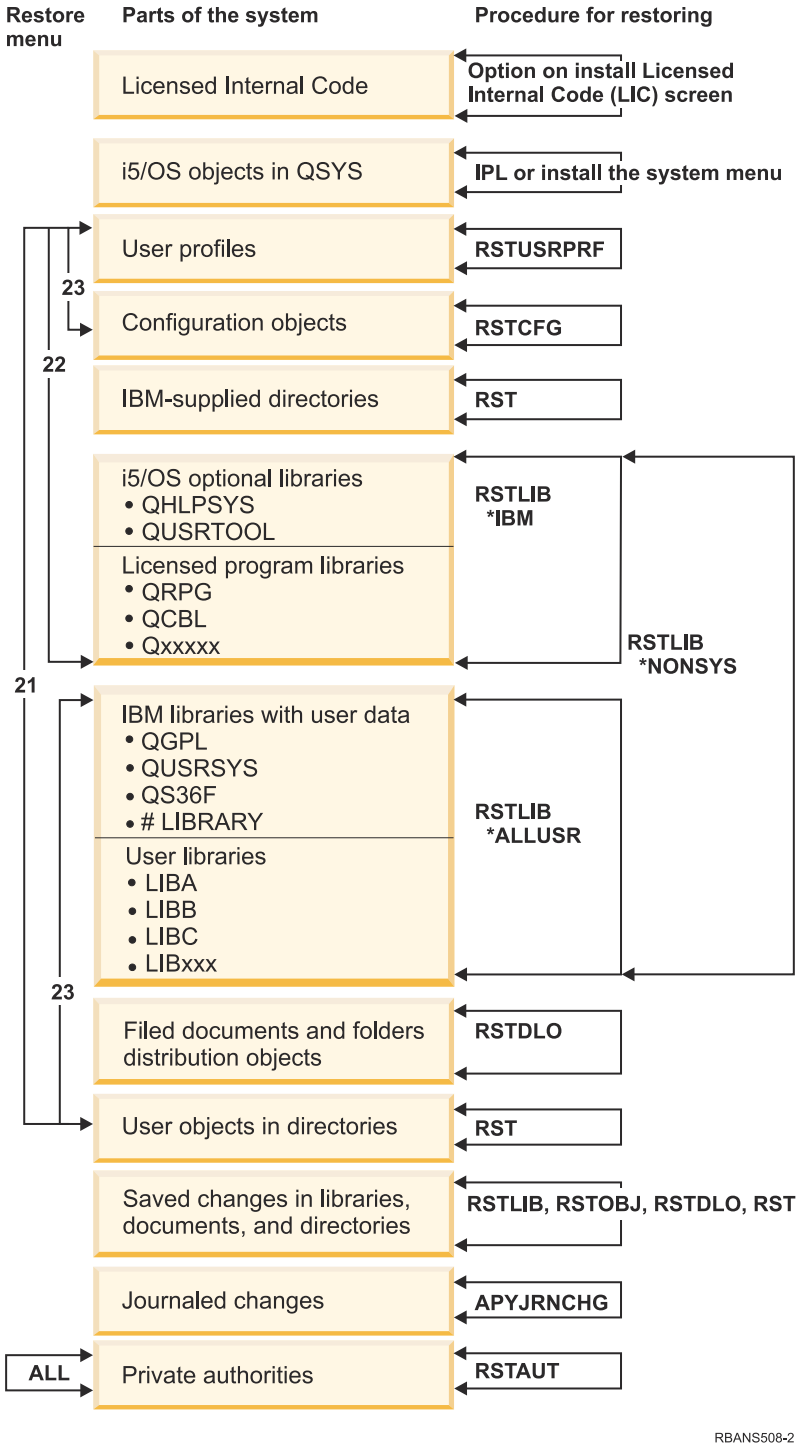

Figure 2. Restore procedures

Note: Restore Object (RSTOBJ) command can also be used where Restore Library (RSTLIB) command is shown to restore objects.

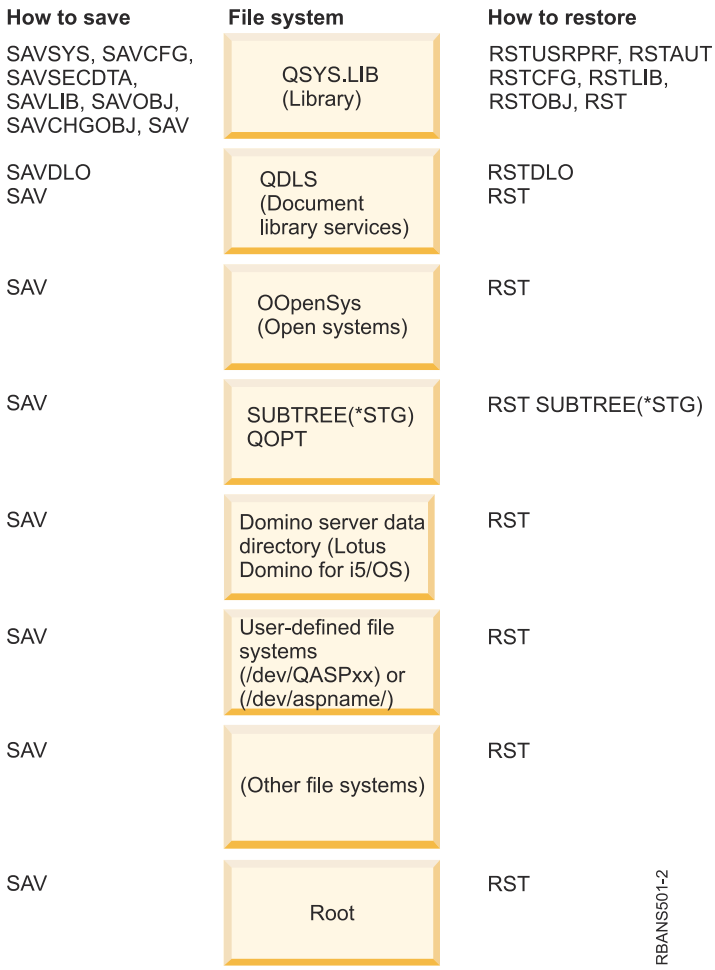

Figure 3. Save procedures and restore procedures for file systems

The relationship between save and restore commands

You can learn which store commands can be used, based on how the objects were saved.

Table 2. Relationship between save and restore commands

Save command used	Possible restore command
SAVOBJ	RSTOBJ
	RST
SAV	RST
SAVLIB LIB(*NONSYS)	RSTLIB SAVLIB(*NONSYS)
	RSTLIB SAVLIB(*IBM)
	RSTLIB SAVLIB(*ALLUSR)
	RSTLIB SAVLIB(*library-name*)
	RST
SAVLIB LIB(*ALLUSR)	RSTLIB SAVLIB(*ALLUSR)
	RSTLIB SAVLIB(*library-name*)
	RST

Table 2. Relationship between save and restore commands (continued)

Save command used	Possible restore command
SAVLIB LIB(*IBM)	RSTLIB SAVLIB(*IBM)
	RSTLIB SAVLIB(*library-name*)
	RST
SAVLIB LIB(*library-name*)	RSTLIB SAVLIB(*library-name*)
	RST
SAVSECDTA	RSTUSRPRF
	RSTAUT[1]
SAVCFG	RSTCFG
SAVSYS	Restore Licensed Internal Code.
	Restore operating system.
	RSTUSRPRF
	RSTCFG
	RSTAUT[1]
SAVDLO	RSTDLO
	RST
SAVSYSINF	RSTSYSINF

1

The Restore User Profiles (RSTUSRPRF) command restores authority information to temporary tables. The Restore Authority (RSTAUT) command regrants private authorities using tables that are built as a part of the RSTUSRPRF command.

Related concepts

Recovering the Licensed Internal Code
Licensed Internal Code is the layer of IBM i architecture just above the hardware.

Restoring the operating system
Use this information to recover the operating system.

Related tasks

Task 5: Restoring libraries to the system auxiliary storage pool
Before you restore libraries to the system auxiliary storage pool (ASP), decide which libraries to restore.

Related information

Backing up your system

What happens when you restore objects

When you restore an object, the system takes different actions depending on several conditions.

An object on this system is like a container. The object has information about the container itself, such as the owner of the object and the last time it was saved. This is the information you see when you display the object description (DSPOBJD command). The object also has contents, such as the records in a database file or the instructions in a program.

When you restore an object, the system takes different actions depending on the following conditions:

- Whether the object to be restored already exists.
- The allow object differences (ALWOBJDIF) parameter on the restore command.

- Whether the object was saved on a different system (serial number of the processor) or logical partition.

With a few exceptions that relate to security, the contents of the object are always restored. If the object exists, the system compares the object description information about the system copy and the media copy and then makes decisions. For most information, the media version of the information is restored. For security-relevant information, such as the public authority and the object owner, the system version is left unchanged. In a few cases, such as the size of the object and the date it was restored, the system determines a value when the object is restored.

The allow object differences (ALWOBJDIF) parameter on the restore commands is primarily for security protection and integrity protection. For example, if system security is important to you, you might want to take special action if someone attempts to restore an object whose owner has been changed. Or, if the member information about a database file does not match, you might have problems with the integrity of your data. You can use the ALWOBJDIF parameter to prevent this.

The default value for the ALWOBJDIF parameter is *NONE. This means that if important differences exist between the media version and the system version of an object, you want the system to notify you and not restore any differences. Normally, you should use the default value. However, when you are restoring your information to a different system or logical partition, such as during a disaster recovery, you should specify ALWOBJDIF(*COMPATIBLE) which allows differences that are compatible with existing database files.

You can specify a combination of up to four values on the ALWOBJDIF parameter to allow specific types of differences for the restore operation: *FILELVL, *AUTL, *OWNER, and *PGP. The *FILELVL value attempts to restore physical file data when the file level ID or the member level ID of the physical file on the system is different from that of the physical file on the save media. The *AUTL value allows differences in authorization lists. The *OWNER value allows differences in object ownership. The *PGP value allows differences in the primary group.

Table 3 on page 23 shows examples of the effect of the ALWOBJDIF parameter.

Table 3. Restoring objects with ALWOBJDIF. Effect of ALWOBJDIF parameter when the value on the media and on the system are different.

Object characteristic that differs	**ALWOBJDIF parameter value and result after restore operation**		
	***NONE**	***ALL**	***COMPATIBLE**
Object owner	Object is not restored	Existing value[1]	Existing value[1]
Object primary group	Object is not restored	Existing value[3]	Existing value[3]
Object auditing	Existing value	Existing value	Existing value
Authorization list, restore over existing object:			
Object on media is secured by an authorization list and object on system is not secured by an authorization list	Object is not restored	Object is restored and is secured by authorization list of object on system[2]	Object is restored and is secured by authorization list of object on system[2]
Object on media is not secured by an authorization list and object on system is secured by an authorization list	Object is restored and is secured by authorization list of object on system	Object is restored and is secured by authorization list of object on system[2]	Object is restored and is secured by authorization list of object on system[2]

Table 3. Restoring objects with ALWOBJDIF. Effect of ALWOBJDIF parameter when the value on the media and on the system are different. *(continued)*

Object characteristic that differs	ALWOBJDIF parameter value and result after restore operation		
	***NONE**	***ALL**	***COMPATIBLE**
Object on media is secured by an authorization list and object on system is secured by a different authorization list	Object is not restored	Object is restored and is secured by authorization list of object on system; message is sent to user[2]	Object is restored and is secured by authorization list of object on system; message is sent to user[2]
Authorization list, new object restored:			
Object is restored onto different system or logical partition than it was saved from	Object is restored and is not secured by an authorization list	Object is restored and is secured by the same authorization list that secured the object when it was saved, if the authorization list exists[2]	Object is restored and is secured by the same authorization list that secured the object when it was saved, if the authorization list exists[2]
Database files:			
Creation date for file	File is not restored	File is renamed on the system; copy is restored from media with media creation date; message is sent to user.	Logical file is not restored. System attempts to restore physical file data[4]
Creation date for member	Member is not restored	Member is renamed on the system; copy is restored from media with media creation date; message is sent to user.	Logical member is not restored. System attempts to restore physical member data[4]
Physical file data:			
Level identifier for file	Physical file data is not restored	File is renamed on the system; copy is restored from media with media creation date; message is sent to user.	System attempts to restore physical file data[4]
Level identifier for member	Physical file data is not restored	Member is renamed on the system; copy is restored from media with media creation date; message is sent to user.	System attempts to restore physical member data[4]

1 Also applies to restore commands with ALWOBJDIF(*OWNER)

2 Also applies to restore commands with ALWOBJDIF(*AUTL)

3 Also applies to restore commands with ALWOBJDIF(*PGP)

4 Also applies to RSTLIB and RSTOBJ commands with ALWOBJDIF(*FILELVL)

Related concepts

How the system establishes ownership for restored objects
Most objects on the system have an owner.

How the system establishes the authorization list for a restored object
This topic shows what happens when you restore an object that already exists if the object is linked to an authorization list. These rules do not apply when you are restoring a document or a folder.

How the system restores programs
Restoring programs to your system represents a security exposure. A restored program might be altered to do functions that you do not intend, or the program might adopt the authority of a powerful user profile.

Related tasks

Comparing file attributes during a restore operation
When you restore a database file or member that exists on the system, the system expects the creation dates for the system copy and the media copy to be the same. If they are not the same, the system cannot ensure that the contents of the saved copy match the format of the copy on the system.

Related information

Restoring changed objects

Sequence for restoring related objects

Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

If possible, restore objects in the following sequence:

- Journals before journaled objects. If you restore a journaled object when the journal is not on the system, you must start journaling again after the journal is restored. Use the Start Journal Physical File (STRJRNPF) command, the Start Journal Access Path (STRJRNAP) command, the Start Journal Object (STRJRNOBJ) command, the Start Journal (STRJRN) command, or the Start Journal Library (STRJRNLIB) command.
- Journals before journal receivers. If you restore a journal receiver when the journal is not on the system, you must associate the journal receivers with the journal after it is restored. Use the Work with Journal (WRKJRN) command.
- Physical files before logical files or MQTs. You cannot restore a logical file or MQT if the based-on physical files are not on the system.

Related concepts

Restoring journaled objects
This section describes options for restoring journaled objects (for example, libraries, data areas, data queues, database files, and integrated file system objects).

How the system restores access paths
The description for a database file contains a description of its access path, if it has one. When you save a database file, you can save the access path with the file. This depends on the type of file, the type of access path, and how you performed the save operation.

Deferring the restore of dependent objects
You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Restoring logical files
When you restore a logical file, the system uses the description for the logical file to establish its relationship with the based-on files. For example, database views or indexes are logical files.

Restoring objects in directories

You can use the Restore Object (RST) command to restore objects that you saved with the Save Object (SAV) command.

Related tasks

Restoring journals and journal receivers
Typically, you can restore journals or journal receivers only to the same library from which they were saved. The one exception is journal receivers that are part of a remote journal network. These receivers can often be restored to the remote receiver library as well.

Task 5: Restoring libraries to the system auxiliary storage pool
Before you restore libraries to the system auxiliary storage pool (ASP), decide which libraries to restore.

Deferring the restore of dependent objects

You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Database logical files, indexes, and SQL *materialized query tables* (MQTs) depend on other files for their information. Objects that are journaled depend on journal objects to record changes. When objects are in the same library as the objects they depend on, the Save Library (SAVLIB) and Restore Library (RSTLIB) commands save and restore the objects in the correct order. The dependent objects are saved and restored after the objects that they depend on.

When objects depend on objects in other libraries, or if you do not save and restore entire libraries, you might attempt to restore dependent objects before the objects that they depend on. Logical files, indexes, and MQTs cannot be restored when the files that they depend on do not exist. Objects cannot start journaling when the journals do not exist. However, you can defer the processing that is needed to restore these objects and start journaling by specifying the Defer ID (DFRID) parameter on the Restore Library (RSTLIB) or Restore Object (RSTOBJ) command. When you specify a defer ID, the system stores the information that is needed to complete the restore and start journaling. The information is stored in files in library QRECOVERY or QRCYxxxxx, where xxxxx is the number of an independent auxiliary storage pool. The information is available later when the related files and journals are available.

When the related files and journals are available, use the Restore Deferred Objects (RSTDFROBJ) command with the same defer ID to complete the restore and start journaling for the dependent objects. If you restore the related files with the same defer ID, the system automatically attempts to complete the restore of the deferred files; you do not need to use the RSTDFROBJ command in this case.

When you use the Restore Library (RSTLIB) command to restore all libraries (*NONSYS), all user libraries (*ALLUSR), or IBM-supplied libraries (*IBM), QRSTLIB is the default DFRID parameter value. These RSTLIB operations automatically perform the RSTDFROBJ operation after the libraries are restored. Therefore, these RSTLIB operations automatically handle dependent objects that are restored in the wrong sequence. For other restore operations, you must specify a defer ID to manage dependent objects.

If you use IBM Backup, Recovery and Media Services for i (BRMS), Q1ARSTID is the default defer ID value for all recovery operations.

Some types of deferred objects are created during the initial restore operation. Other types of deferred objects are not created. A deferred logical file does not exist on the system. A deferred MQT exists on the system, but you cannot perform functions on the MQT that require access to the based-on files. A deferred object that is journaled exists on the system, but journaling is not started.

When your data recovery is complete and you no longer need the information that is stored for deferred objects, use the Remove Defer ID (RMVDFRID) command to remove the information. For example, use the RMVDFRID command if you manually create the deferred objects or start journaling for the deferred objects. Information that is no longer needed might cause future restore operations to send messages that you do not expect.

Note: You can run multiple restore operations concurrently with the same defer ID. Wait until all the restore operations are complete before you run the RSTDFROBJ and the RMVDFRID commands. Running the RSTDFROBJ command while restore operations are still running might cause extra processing.

Running the RMVDFRID command while restore operations are still running might remove information too soon.

Deferred restore examples

This command restores all the saved nonsystem libraries to the system from tape. The printout lists all the objects that were deferred, and a final completion message states how many libraries were restored and how many were not restored. In this example, the system uses the special defer ID value of QRSTLIB so that dependent files that are restored before their dependencies in another library are restored, are deferred. The system attempts to automatically restore the deferred files when the based-on files are finally restored.

```
RSTLIB SAVLIB(*NONSYS) DEV(TAP01) OUTPUT(*PRINT) DFRID(*DFT)
```

This example shows libraries that were saved in alphabetical order being restored. Library LFLIB contains files that have based-on files in library PFLIB. The files in LFLIB cannot be restored until after the files in PFLIB are restored. A defer ID is specified, so the files in LFLIB are deferred. Use the RSTDFROBJ command to complete the restore operation of these files. If you had not specified a defer ID on the command, you would have had to rewind or reload the tape to restore the files in LFLIB.

```
(Restore libraries before LFLIB.)
RSTLIB SAVLIB(LFLIB) DEV(TAPE01) ENDOPT(*LEAVE) DFRID(ABC)
(Restore libraries between LFLIB and PFLIB.)
RSTLIB SAVLIB(PFLIB) DEV(TAPE01) ENDOPT(*LEAVE) DFRID(ABC)
(Restore libraries after PFLIB).
RSTDFROBJ DFRID(ABC)
```

The RMVDFRID command removes any remaining information about objects that were deferred during the restore operation.

```
RMVDFRID DFRID(ABC)
```

Related concepts

Verifying whether objects are restored successfully
You can use the job log or an output file to determine which objects are restored successfully.

Restoring logical files
When you restore a logical file, the system uses the description for the logical file to establish its relationship with the based-on files. For example, database views or indexes are logical files.

Related tasks

Restoring SQL materialized query tables
SQL materialized query tables (MQTs) are physical files, like logical files, that have dependencies on other files. You can defer the restore of logical files and MQTs whose based-on files or members are missing.

Restoring multiple libraries
You can use the Restore Library (RSTLIB) command to restore multiple libraries.

Related information

Restore Library (RSTLIB) command
Restore Object (RSTOBJ) command
Restore Deferred Objects (RSTDFROBJ) command
Remove Defer ID (RMVDFRID) command

Putting your system in a restricted state

Many recovery procedures require that your system have no other activity on it. When no subsystems except the controlling subsystem are active on your system, it is in a *restricted state*.

Use the End Subsystem (ENDSBS) command to put your system in a restricted state. You can specify how you want the subsystems to end.

Table 4. Possible values for the OPTION parameter of the ENDSBS command

Value	Description
*CNTRLD	Allow active jobs to end themselves (if they are checking to see if the job is being ended). If you specify *CNTRLD, you can use the delay parameter to set a time for the system to wait before ending subsystems immediately.
*IMMED	End the subsystem immediately. Use this option if there are no users on the system and no batch jobs running.

Note: Even if you have no activity on the system, jobs might be running under a few system-provided subsystems, such as the QSYSWRK (subsystem monitor) subsystem and the QCALSRV (calendar server) subsystem. You can end all subsystems immediately without first ending these jobs. You will receive messages that these subsystems ended abnormally.

To put your system in a restricted state, follow these steps:

1. Before putting your system in a restricted state, ensure that all users are signed off and all jobs are ended.
2. To receive notification that the subsystems have ended, type the following command and press the Enter key:

```
CHGMSGQ MSGQ(QSYSOPR) DLVRY(*BREAK)
        SEV(60)
```

3. To end all subsystems, type the following command:

```
ENDSBS SBS(*ALL) OPTION(*CNTRLD)
       DELAY(600)
```

   Note: For the delay parameter, specify a number of seconds that allows your system time to bring most jobs to a normal end. On a large busy system, you might need a longer delay.

A message is sent that indicates that the procedure for ending subsystems is in progress. A final message is sent when the system is in a restricted state.

Related tasks

Task 2: Restoring user profiles
Use these steps to restore user profiles.

Restoring user profiles
You can restore a single user profile, a list of user profiles, or all user profiles. You restore a user profile to move a user from one system to another system and to recover a damaged user profile.

Restoring all libraries from a single save operation
Use these steps to restore all libraries that were saved with a single command or menu option.

Reclaim Storage (RCLSTG) command

You can use the RCLSTG command to recover the addressability of lost or damaged objects. This allows you to identify and then restore those objects that were damaged.

If an authorization list is found damaged during reclaim storage, the objects secured by the damaged authorization list are associated with the system authorization list QRCLAUTL.

The RCLSTG command has four parameters. These parameters allow you to perform reclaim functions in the following ways:

SELECT
Specifies all or a subset of the reclaim functions that are performed.

OMIT
Specifies the subset of the reclaim functions that are omitted.

ASPDEV

- Reclaim the system auxiliary storage pool (ASP) and all basic ASPs. The system ASP has an ASP number of 1. Basic ASPs have ASP numbers of 2 through 32.
- Reclaim a specific independent ASP. Independent ASPs have a device name and a number greater than 32.

ESTIMATE
Estimates the amount of time that a RCLSTG command will take to run. The estimate is calculated by using statistics collected during previous RCLSTG operations. The statistics are contained in data area QRCLSTG in library QUSRSYS.

Note: The RCLSTG procedure requires auxiliary storage. If you are already using a very high percentage of auxiliary storage, the RCLSTG procedure might not complete successfully.

What happens when you reclaim storage: The purpose of the RCLSTG command is to ensure the following states:

- Objects that reside permanently in auxiliary storage can be accessed.
- All auxiliary storage either is used properly or is available for use.

The system checks every object that resides permanently in auxiliary storage for loss or damage.

- If an object does not address a library or directory, it is placed in an IBM-supplied library or directory based on the object type. The system might not be able to retrieve description information for the object, such as:
  - Program temporary fix (PTF) status.
  - Save and restore information.
  - Object attributes and text description.
- For objects that normally reside in libraries (the QSYS.LIB file system), the system performs the following tasks:
  - If a lost object with the same name and object type is already in the Recovery (QRCL) library, the system gives the object that it has just encountered a new name. The name has the format QRCL*nnnnn*, where *nnnnn* is a unique number. The original object name is placed in the text description for the object in the QRCL library.

    Note: You cannot rename journals and journal receivers. If the system encounters two journals (or journal receivers) with the same name and they both need to be placed in the QRCL library, the system renames one of them. You cannot rename that journal or journal receiver back to its original name. You must restore a previous version with the correct name or re-create the journal or journal receiver. For this reason, you should use a naming convention for journals and journal receivers that is unique for the entire system, not just for a library.
  - If data exists for a lost physical file, the system attempts to rebuild the file and place it in the QRCL library. To use the physical file, create it again in the correct library with the correct attributes. Then copy the data from the rebuilt file in the QRCL library to the new physical file. The data in the file might not be complete.
  - Independent ASPs have their own unique QRCL library, QRCL*nnnnn* where *nnnnn* is the number of the primary ASP. The text description for the object in the QRCL library indicates that it has been rebuilt.
  - A user domain object can be placed in the QRCL library only if the QALWUSRDMN system value includes QRCL or specifies *ALL. Otherwise, a lost user domain object is deleted. Most objects are system domain objects. User domain objects have type *USRSPC, *USRIDX, or *USRQ.
  - If an object does not have an owner, it is assigned to an IBM-supplied user profile based on the object type. Most objects are assigned to the QDFTOWN user profile.
  - If descriptions for objects in a library cannot be accessed, the library is rebuilt.

  - If an object is secured by a damaged authorization list or authority holder, the system makes QRCLAUTL the authorization list for the object. You can use the Display Authorization List Objects (DSPAUTLOBJ) command to determine which objects are secured by the QRCLAUTL authorization list.
- If a lost object was in the "root" (/) file system, the object is placed in the /QReclaim directory.
- If a lost object was in the QOpenSys file system, the object is placed in the /QOpenSys/QReclaim directory.
- If an object in a directory is damaged to the extent that it is unusable, the system deletes it. The RCLSTG command does not attempt to rebuild damaged objects.
- If a lost object was in a user-defined file system (UDFS), it is placed in the QReclaim directory located in the "root" (/) directory of the UDFS.
- If a lost object that was in a directory cannot be placed in the proper QReclaim directory based on its original location, then it is placed into the "root" (/) directory of a special file system within the ASP in which the object resides. This special file system is created by RCLSTG when needed. The file system is named '/dev/QASPxx/QReclaimFS.udfs' where '*xx*' is the number for system and basic ASPs. The file system is named '/dev/*iasp-name*/QReclaimFS.udfs' where *iasp-name* is the name of the independent ASP.
- For objects in the "root" (/) directory, QOpenSys, or UDFS, the system takes actions for duplicate names or for unidentified object owners similar to the actions taken for objects in the QSYS.LIB file system.

What to do after running the rclstg procedure: Table 5 on page 30 describes both where to look for problems that the RCLSTG procedure detects and how to correct the problems:

Table 5. Resolving problems detected by the RCLSTG procedure

Where to look for problems	**How to fix the problem**
Type `DSPMSG QSYSOPR` to display the QSYSOPR message queue. Look for messages about damaged objects.	Type `DSPLOG QHST` to display the history log. Look for messages about either damaged objects or rebuilt files. 1. Delete unusable objects by using the appropriate DLT*xxx* command. Restore them by using the Restore Object (RSTOBJ) or Restore Object (RST) command. 2. Copy data from rebuilt files to new files by using the Copy File (CPYF) command. **Note:** You might see a message indicating that objects were deleted by the reclaim storage procedure. These are internal system objects that are no longer needed.
Type `DSPLIB QRCL` to display the QRCL library. **Note:** If the reclaim storage procedure did not place any objects in the QRCL library, you might receive a message that the library is not found. Ignore the message and continue with the next step.	Move objects from the QRCL library to the correct library by using the Move Object (MOVOBJ) command. **Note:** • If IBM-supplied objects are in the QRCL library, contact software support for assistance. • If you are unsure what to do with objects in the QRCL library, use the Save Library (SAVLIB) command to save the library to save media. Mark the media volume in case you need the objects later.

Table 5. Resolving problems detected by the RCLSTG procedure (continued)	
Where to look for problems	**How to fix the problem**
Display the /QReclaim directory by using the Display Link (DSPLNK) command. **Note:** If the reclaim storage procedure did not place any objects in the /QReclaim directory, you might receive a message that the object is not found. Ignore the message and continue with the next step.	Move objects from the /QReclaim directory to the correct directory by using the Move (MOV) command.
Display the /QOpenSys/QReclaim directory by using the Display Link (DSPLNK) command. **Note:** If the reclaim storage procedure did not place any objects in the /QOpenSys/QReclaim directory, you might receive a message that the object is not found. Ignore the message and continue with the next step.	Move objects from the /QOpenSys/QReclaim directory to the correct directory by using the MOV command.
Type `DSPMSG QSYSOPR` to display the QSYSOPR message queue. Look for CPFA0D7 messages. For each CPFA0D7 message that contains a directory name beginning with '/dev/QASPxx/' (where 'xx' is the number of a system or basic ASP) or '/dev/*iasp-name*' (where *iasp-name* is the name of an independent ASP), perform the action specified in the "How to fix the problem" column.	Use the Add Mounted File System (ADDMFS) command to mount the UDFS specified in the CPFA0D7 message over a directory of your choice. Then use the Display Link (DSPLNK) command to view the contents of this UDFS. You might see objects with names beginning with 'QRCL' or you might see a directory named 'QReclaim'. If you see the 'QReclaim' directory, look inside it to find the object names beginning with 'QRCL'. These objects were previously lost but have been relocated by the RCLSTG command. Use the Move (MOV) command to move these objects back to their original location. The original object names can be specified in the CPFA0D7 message. If the original names are not available, use the "Display Attributes" option in DSPLNK to view an object's attributes in order to attempt to identify it.
Type `WRKOBJOWN QDFTOWN` to display objects that are owned by the QDFTOWN user profile.	Use option 9 (Change owner) from the Work with Objects by owner display to transfer ownership to the correct user profile.
Type `DSPAUTLOBJ QRCLAUTL` to display objects that are secured by the QRCLAUTL authorization list. **Note:** If the reclaim storage procedure did not assign any objects to the QRCLAUTL authorization list, you might receive a message that the authorization list is not found. Ignore the message.	If necessary, assign the object to the correct authorization list by using the Edit Object Authority (EDTOBJAUT) command.

Related concepts

Reclaiming object links

The Reclaim Object Links (RCLLNK) command reclaims objects in the "root" (/), QOpenSys, and mounted user-defined file systems (UDFSs).

Related tasks

Recovering damaged objects in the integrated file system

To recover damaged objects in the integrated file system, run the Reclaim Storage (RCLSTG) command. To recover damaged objects without putting your system in a restricted state, you can use the Reclaim Object Links (RCLLNK) command. Perform these recovery steps if the object was not journaled.

Task 1: Reclaiming storage

Use these steps to reclaim storage.

Related information

Restoring authorization lists

Reclaiming spooled file storage

Reclaim operation of the "root" (/), QOpenSys, and user-defined file systems

Reclaiming storage

Follow these steps to reclaim storage. You can reclaim storage of part of the system or the entire system.

1. Sign on the system with a user profile that is authorized to the Reclaim Storage (RCLSTG) command. Either sign on at the console or use the Transfer Job (TFRJOB) command to transfer your job to the controlling subsystem.
2. Type `DSPSYSVAL QALWUSRDMN`. If the value does not include the QRCL (Reclaim Storage) library or does not specify *ALL, use the CHGSYSVAL command to add QRCL to the list of libraries for this system value. Write the current value here: _____
3. Type `DSPSYSVAL QCTLSBSD` to display the name of your controlling subsystem. Write the current value here: _____
4. If you are reclaiming storage for system and basic auxiliary storage pools (ASPs), verify that your system is in a restricted state.

   To obtain a restricted state, follow the procedure in "Putting your system in a restricted state" on page 27 topic.

   Note: If you use independent ASPs, you must make the independent ASPs unavailable first. You can use the Vary Configuration (VRYCFG) command or the IBM Navigator for i to make the independent disk pool unavailable. If independent ASPs remain available, SYSBAS reclaim performance is degraded because primary ASPs contain SYSBAS cross-reference files that are rebuilt during a SYSBAS reclaim operation.
5. Start the reclaim storage process by typing one of the following commands:

RCLSTG	Reclaim storage of the entire system.
RCLSTG SELECT(*DBXREF)	Reclaim storage of the database cross-reference table.
RCLSTG SELECT (*DIR)	Reclaim storage of the directory recovery portion.
RCLSTG OMIT(*DBXREF)	Reclaim storage of the entire system except the database cross-reference table.
RCLSTG ASPDEV(*SYSBAS)	Reclaim the system ASP and all basic ASPs.
RCLSTG OMIT(*DIR)	Reclaim storage of the entire system except the directory recovery portion.

   Note: See "Reclaiming object links" on page 33 for an alternative command to reclaim your integrated file system.
6. If you have independent ASPs, make them available now by using the VRYCFG command.
7. Reclaim independent ASPs one at a time by using one of the following commands:

RCLSTG ASPDEV(auxiliary-storage-pool-device-name)	Reclaim a user-defined file system (UDFS) ASP. Enter the name of the device description for the independent ASP. Make sure the independent ASP is available before you begin the reclaim storage process.
RCLSTG ASPDEV (auxiliary-storage-pool-group-name)	Reclaim the primary ASP and all the secondary ASPs within the ASP group. The name of any ASP within the ASP group is accepted. Primary and secondary ASPs have ASP numbers greater than 32. Make sure the independent ASP group is available before you begin the reclaim storage process.

Note: If you prefer, you can reclaim storage for independent disk pools in parallel after you reclaim storage for SYSBAS and start the controlling subsystem. To reclaim storage for independent disk pools, submit a reclaim storage job for each UDFS ASP and each primary ASP.

8. Use the Change System Value (CHGSYSVAL) command to set the QALWUSRDMN system value back to its original setting. (You wrote the setting in step 2.)
9. When the reclaim storage procedure finishes, start the controlling subsystem by typing the following command -

```
STRSBS SBSD(controlling-subsystem)
```

(The name of the controlling subsystem was recorded in step 3.)

Reclaiming object links

The Reclaim Object Links (RCLLNK) command reclaims objects in the "root" (/), QOpenSys, and mounted user-defined file systems (UDFSs).

The RCLLNK command has three parameters. These parameters allow you to perform reclaim functions in the following ways:

OBJ
Specifies the object to be reclaimed.

SUBTREE
Specifies whether to reclaim the objects within the subtree if the object specified by the OBJ parameter is a directory.

DMGOBJOPT
Specifies how damaged objects are handled during the reclaim.

Use these examples to better understand situations in which the RCLLNK command can be used to reclaim objects in the "root" (/), QOpenSys, and UDFSs.

Correcting problems for an object in a directory subtree

In this situation, the known problems are isolated to one object. The object is damaged and unusable, and a backup version of the object cannot be restored from media. To correct the problem quickly without disrupting normal file system operations, you can use this command:

RCLLNK OBJ('/MyDir/MyBadObject')
SUBTREE(*NONE)

where '/MyDir/MyBadObject' is the damaged and unusable object.

Correcting problems that exist in a directory subtree

In this situation, the known problems are isolated to a group of objects within a directory subtree. The problems need to be corrected quickly without disrupting normal file system operations. To reclaim the objects within the directory subtree, you can use this command:

RCLLNK OBJ('/MyApplicationInstallDirectory') SUBTREE(*ALL)

where '/MyApplicationInstallDirectory' is the "root" (/) directory of the subtree containing the problems.

Finding all damaged objects in the "root" (/), QOpenSys, and mounted user-defined file systems

In this situation, a disk failure has caused damage to a number of objects. The damaged objects need to be identified before determining how to properly recover them. The solution required needs to identify the damaged objects, but not take action against them and must not disrupt normal file system operations. To identify the damaged objects, you can use this command:

RCLLNK OBJ('/') SUBTREE(*ALL) DMGOBJOPT(*KEEP *KEEP)

It will also correct problems other than damaged objects as it identifies damaged objects.

Deleting all damaged objects in the "root" (/), QOpenSys, and mounted user-defined file systems

In this situation, a disk failure has caused a number of objects to become damaged. The damaged objects need to be deleted so backup copies can be restored from media. To delete the damaged objects, you can use this command:

RCLLNK OBJ('/') SUBTREE(*ALL) DMGOBJOPT(*DELETE *DELETE)

The damaged objects will be deleted without disruption to normal file system operations. In addition, problems other than damage are corrected as the damaged objects are being deleted.

Related concepts

Reclaim Storage (RCLSTG) command
You can use the RCLSTG command to recover the addressability of lost or damaged objects. This allows you to identify and then restore those objects that were damaged.

Related tasks

Recovering damaged objects in the integrated file system
To recover damaged objects in the integrated file system, run the Reclaim Storage (RCLSTG) command. To recover damaged objects without putting your system in a restricted state, you can use the Reclaim Object Links (RCLLNK) command. Perform these recovery steps if the object was not journaled.

Related information

Reclaim Object Links (RCLLNK) command

Controlling restoration of security-sensitive objects

You can use two different system values to control the restoration of security-sensitive objects.

Here are the two system values:

- Allow object restore operation (QALWOBJRST).
- Verify object on restore (QVFYOBJRST).

Related concepts

How the system restores programs
Restoring programs to your system represents a security exposure. A restored program might be altered to do functions that you do not intend, or the program might adopt the authority of a powerful user profile.

Related information

System values
Digital signatures
Scanning support

QALWOBJRST system value

You can use the QALWOBJRST system value to prevent anyone from restoring a system-state object or an object that adopts authority.

The QALWOBJRST system value determines whether objects that are security-sensitive can be restored to your system. It allows you to determine how the system restores objects with digital signatures. This system value affects programs, service programs, modules, and Structured Query Language (SQL) packages.

When your system is shipped, the QALWOBJRST system value is *ALL. This value is necessary to install your system successfully.

Note: It is important to set the QALWOBJRST value to *ALL before performing some system activities. Here are several examples:

- Installing a new release of the IBM i licensed program.
- Installing new licensed programs.
- Recovering your system.

These activities might fail if the QALWOBJRST value is not *ALL.

If you are applying program temporary fixes (PTFs), set the QALWOBJRST value to *ALWPTF.

To ensure system security, return the QALWOBJRST value to your normal setting after completing the system activity. Make sure that the entire restore operation has completed before changing the QALWOBJRST system value, or some objects might not restore successfully.

You can specify multiple values for the QALWOBJRST system value, unless you specify *ALL or *NONE.

Table 6. Possible values for the QALWOBJRST system value

Value	Description
*ALL	Any object can be restored to your system by a user with the proper authority.
*NONE	Security-sensitive objects, such as system state programs or programs that adopt authority, might not be restored to the system.
*ALWSYSSTT	System state objects can be restored to the system.
*ALWPGMADP	Objects that adopt authority can be restored to the system.
*ALWPTF	Security-sensitive objects can be restored only when the restore is part of a PTF operation.
*ALWSETGID	Allows the system to restore files with the S_ISGID attribute enabled
*ALWSETUID	Allows the system to restore files with the S_ISUID attribute enabled
*ALWVLDERR	Allows the system to restore objects with validation (CRC) errors.

Setting the QALWOBJRST system value to allow complete recovery

You can learn the procedure of how to set the QALWOBJRST system value to allow complete recovery.

1. Type WRKSYSVAL QALWOBJRST and press the Enter key.
2. You are shown the Work with System Values display.
   Type 5 (Display) in the Opt column next to QALWOBJRST and press the Enter key.
3. You are shown the Display System Value display.
   Write down the current setting for use after you complete your recovery. If the value is *ALL, you do not need to change it for your recovery. Skip to step 6.
4. Press F12 to return to the Work with System Values display.
   Type 2 (Change) in the Opt column next to QALWOBJRST and press the Enter key.
5. You are shown the Change System Value display.
   Type *ALL for the value and press the Enter key.
6. Press F12 to cancel the Work with System Values display.

Setting the QALWOBJRST system value to restrict restore operations

You can learn the procedure of how to set the QALWOBJRST system value to restrict restore operations.

1. Type `WRKSYSVAL QALWOBJRST` and press the Enter key.
2. You are shown the Work with System Values display.

   Type 2 (Change) in the Opt column next to QALWOBJRST and press the Enter key.
3. You are shown the Change System Value display.

   Type the value you wrote down in step3 in How to set the QALWOBJRST system value to allow complete recovery topic. Press the Enter key.
4. Press F12 to cancel the Work with System Values display.

QVFYOBJRST system value

You can add digital signatures to objects so that users can verify the object's integrity and origin.

The objects affected by the QVFYOBJRST system value are as follows:

- *CMD
- *PGM
- *SRVPGM
- *SQLPKG
- *MODULE
- *STMF objects with attached Java™ programs

You need to install Digital Certificate Manager (DCM) (IBM i option 34) before you can use the QVFYOBJRST system value to verify objects not signed by IBM. If DCM is not installed, the system treats user-state objects being restored as unsigned objects even if they contain digital signatures. You do not need to restart your system for changes to this value to take effect.

The QVFYOBJRST system value allows you to control signature verification of objects during a restore operation. The QVFYOBJRST system value has the following five options (option 3 is the default value):

1. **Do not verify signatures on restore.**

   This is the only option that restores system-state or inherit-state objects without valid IBM-generated signatures.

   This option should not be used unless you have a large number of signed objects to restore that will fail their signature verification for some acceptable reason. Allowing a system-state or inherit-state object without a valid signature to restore represents an integrity risk to your system. If you choose to restore such an object onto your system by selecting this option, be sure to change it back to its previous value after the object has been restored.
2. **Verify: Restore unsigned objects; Restore signed object, even if signatures are not valid.**

   Restores unsigned user-state objects. Restores signed user-state objects, even if signatures are not valid. Does not restore system-state or inherit-state objects without valid IBM-generated signatures. This option should be used only if there are specific objects with signatures that are not valid which you want to restore. In general, it is dangerous to restore objects with signatures that are not valid on your system.
3. **Verify: Restore unsigned objects; Restore signed objects only if signatures are valid.**

   Restores unsigned user-state objects. Restores signed user-state objects only if signatures are valid. Does not restore system-state or inherit-state objects without valid IBM-generated signatures.
4. **Verify: Do not restore unsigned objects; Restore signed objects, even if signatures are not valid.**

   Does not restore unsigned user-state objects. Restores signed user-state objects, even if signatures are not valid. Does not restore system-state or inherit-state objects without valid IBM-generated signatures.

5. **Verify: Do not restore unsigned objects; Restore signed objects only if signatures are valid.**

   Does not restore unsigned user-state objects. Restores signed user-state objects only if signatures are valid. Does not restore system-state or inherit-state objects without valid IBM-generated signatures. This option is the most restrictive option and should be used when the only objects you want to allow to be restored are those which have been signed by trusted sources.

Attention: It is important to set the QVFYOBJRST value to 1 before performing some system activities, such as:

- Installing a new release of the IBM i licensed program.
- Recovering your system.

These activities might fail if the QVFYOBJRST value is not 1.

Setting the QVFYOBJRST system value to allow complete recovery

You can learn the procedure of how to set the QVFYOBJRST system value to allow complete recovery.

1. Type `WRKSYSVAL  QVFYOBJRST` and press the Enter key.
2. You are shown the Work with System Values display. Type 5 (Display) in the Opt column next to QVFYOBJRST and press the Enter key.
3. You are shown the Display System Value display.

   Write down the current setting for use after you complete your recovery. If the value is 1, you do not need to change it for your recovery. Skip to step 6.
4. Press F12 to return to the Work with System Values display.

   Type 2 (Change) in the Opt column next to QVFYOBJRST and press the Enter key.
5. You are shown the Change System Value display. Type 1 for the value and press the Enter key.
6. Press F12 to cancel the Work with System Values display.

Setting the QVFYOBJRST system value to restrict restore operations

You can learn the procedure of how to set the QVFYOBJRST system value to restrict restore operations.

1. Type `WRKSYSVAL  QVFYOBJRST` and press the Enter key.
2. You are shown the Work with System Values display.

   Type 2 (Change) in the Opt column next to QVFYOBJRST and press the Enter key.
3. You are shown the Change System Value display.

   Type the value you wrote down in step 3 in How to Set the QVFYOBJRST System Value to Allow Complete Recovery topic. Press the Enter key.

Minimizing object scans after restore operations

You can learn how to prevent scanning of objects during the first access of the objects after they are restored.

IBM i provides the capability to scan objects in the integrated file system for things like viruses and file usage. As part of this support, you can specify a scan attribute for integrated file system objects so that they are never scanned, or are scanned only after the object is changed. The default scan attribute value is that the objects will be scanned if object scanning is implemented for the integrated file system. Additionally, the Scan file systems control (QSCANFSCTL) system value was introduced in conjunction with this function. One of the default options for QSCANFSCTL is that all objects will be scanned on first access after the object is restored even if the scan attribute for the object is to never be scanned or to scan only if the object is changed.

If you are restoring objects from trusted sources or are restoring from saves where a scan was done as part of the save, and the trusted source or you have set the scan attributes to never be scanned or to scan only if the objects are changed, you might want to take action to prevent the scans during the first access of the objects after they are restored. To prevent the scan that occurs after the restore, you must

specify *NOPOSTRST on the QSCANFSCTL system value before the restores are started. After your restore is complete, you should consider removing the *NOPOSTRST option on the QSCANFSCTL system value so that on subsequent restores objects will be scanned on first access after an object is restored, regardless of what attribute is specified.

Setting the QSCANFSCTL system value to minimize object scans after restore operations

1. Type `WRKSYSVAL  QSCANFSCTL` and press the Enter key.
2. You are shown the Work with System Values display.

   Type 5 (Display) in the Opt column next to QSCANFSCTL and press the Enter key.
3. You are shown the Display System Value display.

   Write down the current settings for use after you complete your recovery. If the value *NOPOSTRST is already specified, you do not need to change it for your recovery. Skip to step 6.
4. Press F12 to return to the Work with System Values display.

   Type 2 (Change) in the Opt column next to QSCANFSCTL and press the Enter key.
5. You are shown the Change System Value display.

   Type *NOPOSTRST. If *NONE is specified currently, replace the *NONE value. If any other values are specified, add *NOPOSTRST as an additional value. Press the Enter key.
6. Press F12 to cancel the Work with System Values display.

Setting the QPFRADJ system value for a different processor or memory

You can set the QPFRADJ system value for a different processor or memory.

1. Type WRKSYSVAL QPFRADJ and press the Enter key.
2. You are shown the Work with System Values display.

   Type 2 (Change) in the option column next to QPFRADJ and press the Enter key.
3. You are shown the Change System Value display.

   Change the value to 2 (Adjustment at IPL and automatic adjustment) and press the Enter key.
4. Press F12 to cancel the Work with System Values display.

Locked objects while restoring

In general, an object is locked to prevent a read operation while it is being restored. If the system cannot obtain a lock on an object within the specified time, that object is not restored and a message is sent to the job log.

Table 7 on page 38 shows the type of lock that is required to restore objects:

Table 7. Lock type needed for restore operation

Object type	Lock type needed for restore operation
Most object types	*EXCL
Job queue	*SHRRD
Output queue	*SHRRD
Message queue	*EXCLRD
Library, during RSTLIB command	*SHRUPD
Library, when object is being restored to it	*SHRUPD

If you restore an Original Program Model (OPM) program that is running, the program might end abnormally.

Verifying whether objects are restored successfully

You can use the job log or an output file to determine which objects are restored successfully.

Note: The system does not restore files to libraries QGPL and QUSRSYS if the file names begin with QAPZ. No diagnostic message is sent indicating that these files are not restored.

Using the job log: The restore commands send the following messages stating the number of objects restored and not restored:

CPC370E
Number of objects restored from media.

CPC370F
Number of objects restored from a save file.

CPC3703
Number of objects restored to a library.

CPC371C
Number of deferred objects completed.

CPC9023
Number of document library objects restored from media.

CPC9053
Number of document library objects restored from a save file.

CPF37C1
Number of deferred objects completed and not completed.

CPF3773
Number of objects restored and not restored to a library.

CPF3839
Number of objects restored and not restored from media.

CPF383E
Number of objects restored and not restored from a save file.

CPF9003
Number of document library objects restored and not restored.

CPF909B
Number of document library objects restored and not restored from a save file.

CPI373B
All deferred objects completed.

CPI373C
Number of deferred objects remaining.

These messages tell the number of objects that are restored and the number of objects that are not restored. An object is counted only if it fits the selection values you specified. For example, assume that library LIB1 has 75 objects. The names of 74 of those objects begin with the characters ORD. You specify `RSTOBJ OBJ(ORD*) OBJTYPE(*ALL) SAVLIB(LIB1)`. If all objects restored successfully, the completion message says that 74 objects were restored to library LIB1. You are not notified that 1 object was not restored.

For example, assume that library DB1 has 20 objects with dependent files. You specify `RSTLIB SAVLIB(DB1) RSTLIB(DB2) DFRID(D1)`. If 5 of the objects cannot be restored because they have dependent files, message CPI373C says that 5 deferred objects remain to be restored.

A diagnostic message is sent if one of following situations happen:

An object could not be restored
The reasons why the system cannot restore an object successfully are as follows:

- The object exists on the system and is being used. Restoring an object requires an exclusive lock for most object types.

- The object is being saved or restored by another job.
- The object on the media is damaged.
- The user does not have the necessary authority to restore the object.
- The object does not have a valid signature.
- The object type is not supported in an independent auxiliary storage pool (ASP).
- The user profile does not exist on the system.
- The system found a validation error.
- The deferred files cannot be restored.

Security information was changed

Under some conditions, the system might do the following tasks:

- Revoke public and private authority
- Change object ownership.
- Change the object's primary group.
- Not link to the authorization list.

An integrity change occurred

- Journaling cannot be started for an object that was being journaled at the time of the save operation or when the journaled library inheritance rules indicate journaling should be started for the object.
- A logical file is restored over a deleted and recreated physical file.
- The QAUDJRN (audit) journal was created by the system. You receive message CPF7088. If you restore the operating system and the QAUDLVL system value is not *NONE, the system creates the QAUDJRN if it does not exist. This ensures that security auditing is restarted for your system.

Using An output file: Most restore commands create output that shows what was restored. You can direct this output to a printer (OUTPUT(*PRINT)), a database file (OUTPUT(*OUTFILE)), a stream file, or a user space. The default for restore commands is not to create output. You must request it each time you run the restore command. Or you can change the default for the OUTPUT parameter for restore commands by using the Change Command Default (CHGCMDDFT) command.

You can print the output and save it. Or you can create a program to analyze and report on the information in the output file.

You can use the OUTPUT parameter with these commands:

RST	RSTDLO	RSTOBJ	RSTDFROBJ
RSTCFG	RSTLIB	RSTUSRPRF	

See the information center topics for each restore commands for the names of the model database outfiles that they use for output. The RSTLIB and RSTOBJ commands and the QSRRSTO API use the QASRRSTO restore output file.

Notes:

- The output file you specify is in use throughout the restore operation. Therefore, the system cannot restore it as part of the operation. Depending on how you perform your restore operation, you might see a CPF379D message in the job log for the output file. If you want to restore the output file after your restore operation completes, use the Restore Object (RSTOBJ) command.
- The RST command places output in a stream file or a user space, rather than an output file.
- The Restore Library (RSTLIB), RSTOBJ, and RST commands have an information type (INFTYPE) parameter to specify how much detail you want in the output file.

Related concepts

Deferring the restore of dependent objects

You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Restoring journaled objects
This section describes options for restoring journaled objects (for example, libraries, data areas, data queues, database files, and integrated file system objects).

Related reference

Restore operation output file information
This table shows the format of the QASRRSTO (restore operation) output file. The Restore Library (RSTLIB), Restore Object (RSTOBJ), Restore Configuration (RSTCFG), Restore Deferred Objects (RSTDFROBJ), and Restore User Profiles (RSTUSRPRF) commands and the Restore Object List (QSRRSTO) API use the QASRRSTO output file.

Related information

Interpret output from save (SAV) and restore (RST)
Restore Object List (QSRRSTO) API

Recovery from an unsuccessful restore operation

A restore operation can be unsuccessful either because an error was encountered when trying to restore an object or because the operation was interrupted. If the object existed on the system before the restore operation, it might be damaged by the unsuccessful restore operation.

An object is not restored when an error is encountered. The error is either recoverable or not.

Restore operation error is recoverable: If an object cannot be restored and the error is recoverable, the following situations will occur:

- A diagnostic message is sent to the job log for each object that is not restored. The message ID can vary, depending on why the object was not restored.
- Each object that is associated with the errors is not restored. However, other objects not associated with the errors but involved in the same restore operation are restored.
- Only the save and restore status information for the objects that were successfully restored is updated.
- A count of the number of objects successfully restored and a count of the number of objects not restored are sent to the user in a diagnostic message.

Restore operation error is not recoverable: If the error is not recoverable, the following situations occur:

- Diagnostic messages are sent to the job log for each object.
- The save and restore status information for each object is not updated.
- A diagnostic message that identifies the error condition is sent to the user.
- The restore command ends immediately. No other objects are restored.

Recovering from an error while restoring libraries

Here are the basic library recovery steps for a restore operation when errors occur.

Some errors that occur during a restore operation stop the operation. Here are two examples of this type of error:

- An unrecoverable media error.
- Exceeding the maximum storage specified in the user profile of the user running the restore operation or in a user profile that owns objects that are being restored.

If an error stops the restore operation, you can correct the error condition and then start the restore operation where it ended. For example, if the maximum storage is exceeded, you can increase the MAXSTG parameter in the user profile.

You can use the STRLIB parameter on the RSTLIB command to restart the restore operation. The STRLIB parameter is valid only when *NONSYS, *ALLUSR, or *IBM is specified for the restore operation,

The basic recovery steps for a restore operation are as follows:

1. Check the job log to determine the library where the previous RSTLIB SAVLIB(*NONSYS, *IBM, or *ALLUSR) failed.

   Find the last library that is restored which is indicated by a successful restore completion message.
2. Load the media volume of the SAVLIB LIB(*NONSYS, *ALLUSR, or *IBM) media.
3. Type the following and press the Enter key:

```
RSTLIB SAVLIB(*NONSYS, *IBM or *ALLUSR)
       DEV(media-device-name) ENDOPT(*LEAVE)
       STRLIB(library-name)
       OMITLIB(library-name)
```

   - If the restore operation stopped because of a media error that you cannot correct, the *library-name* for the STRLIB and the OMITLIB parameters should be the library where the restore operation failed. This causes the operation to start with the library after that library.
   - If the failure was not related to a media error, the *library-name* for the STRLIB and the OMITLIB parameters should be the name of the last library that was successfully restored. This causes the operation to start with the library that caused the error.
4. You will be asked to load the volume that contains the starting library.
5. After the restore operation is complete, restore the library that failed using the media from a previous save operation.

   Note: Consider eliminating the media volume with the media error from the next save.

Recovering from an error while restoring document library objects

Here are the basic document library objects (DLOs) recovery steps for a restore operation when errors occur.

Some errors that occur during a restore operation stop the operation. Here are two examples of this type of error:

- An unrecoverable media error.
- Exceeding the maximum storage specified in the user profile of the user running the restore operation or in a user profile that owns objects that are being restored.

If an error occurs that stops the restore operation, you can correct the error condition and then start the restore operation where it ended. For example, if the maximum storage is exceeded, you can increase the MAXSTG parameter in the user profile.

If an unrecoverable error occurs when running the RSTDLO DLO(*ALL) SAVFLR(*ANY) command, you must determine where the failure occurred and continue the restore operation.

Check the job log to determine if the failure occurred on a distribution object or a folder. The job log might identify where the failure occurred. If the failure occurred on a folder, go to Recovering document and folders topic.

Recovering documents and folders

If an unrecoverable error occurs during the RSTDLO procedure, you can restart the procedure by using the SAVFLR parameter on the Restore Document Library Object (RSTDLO) command.

The basic recovery steps for a restore operation are as follows:

1. Check the job log to determine where the previous RSTDLO DLO(*ALL) command failed.

   The job log identifies which folder failed to restore.

   Note: If the failure occurred during the restore of mail, you need to restore all documents and folders.
2. Find the first folder after the folder that failed to restore.

Use the list that was created during the last SAVDLO OUTPUT(*PRINT or *OUTFILE) operation or use the DSPTAP DATA(*SAVRST) command to determine which first-level folder is next. To find the first-level folders, find the object type *FLR. Look at the *Document or Folder Information* column. The name of a first-level folder does not contain a forward slash (/).

3. Load the first media volume of the SAVDLO DLO(*ALL) save media.

   Note: You must always start with the first volume of the SAVDLO media for each set of 300 first-level folders. You must load each volume in the set of SAVDLO save media in sequence.

4. For each first-level folder, type the following command and press the Enter key:

```
RSTDLO DLO(*ALL) SAVFLR(folder-name-list)
       DEV(media-device-name)
```

   Where the *folder-name-list* has the names of the first-level folders identified from the list described in step 2. You can specify a limit of 300 first-level folders.

   Repeat this step for each set of 300 first-level folders.

Recovering from unsuccessful system signon

If you are using additional workstations to run concurrent restore operations after IBM i has been recovered, but before an initial program load (IPL) has been performed,, you might be unable to sign on to these workstations.

You might have received message CPF4101,

`File xxx in library xxx not found or inline data file missing.`

If this library has already been restored to the system, you can resolve the problem by editing the QSYSLIBL system value. Use the CHGSYSVAL command and add the missing library to QSYSLIBL (even if it is already listed). Perform another operation to remove the library you just added.

As an alternative, you can perform an IPL of the system to resolve the signon problem, if the libraries listed in the QSYSLIBL and QUSRLIBL system values have been restored.

Performing a normal initial program load

Perform a normal initial program load (IPL) at the end of any recovery before allowing users to resume normal activity.

To perform a normal IPL, follow these steps:

1. Place the system in Normal mode.

   If you are using logical partitions with a Hardware Management Console (HMC), perform these steps on the console of the logical partition on which you want to perform an IPL:

   a) Expand your **Management Environment in your HMC**.

   b) Expand your Server.

   c) Expand **Server and Partition**.

   d) Click **Server Management**.

   e) Expand **Partitions**.

   f) Select a partition.

   g) Right-click the partition and select **Properties**.

   h) In the Partition Properties window, click the **Settings** tab.

   i) Under the Boot category, change the **Keylock position to Normal**.

   j) Click the **OK** button.

2. Ensure that no users are signed on and no jobs are active.
3. Type the following command and press the Enter key:

```
PWRDWNSYS OPTION(*IMMED) RESTART(*YES)
```

4. When the IPL is complete, sign on the system.
5. Start any other subsystems that need to be started, such as QTCP or QSNADS.

```
STRSBS SBSD(subsystem-name)
```

Parallel restore operations

You can perform restore operations while using more than one device simultaneously. The data that you restore in this manner must have been saved in parallel format.

You can use the Restore Library (RSTLIB), Restore Object (RSTOBJ), and Restore (RST) commands with a media definition to perform a parallel restore. You can use a media definition with the RSTLIB command to restore *ALLUSR, *IBM, and *NONSYS libraries that were saved with a media definition.

Restoring from a parallel save is typically faster. This is usually true when restoring an entire library or restoring any objects that were saved near the start of a given tape file than restoring from a serial save operation. However, there are cases where restoring individual objects from a parallel save can take significantly longer than restoring from a serial save operation. This situation can occur if the object to be restored is near the end of the tape file

You can restore from a parallel save if you are using fewer devices than the save operation used. However, IBM does not suggest this, due to the amount of volume switching that you will need to do. IBM also does not suggest this due to performance reasons. If you are restoring files with constraints, you must use the same number of devices for the restore as you used for the save. Restore operations that use fewer drives should only be used to restore individual objects. Restore operations that use fewer drives should never be used as part of a system recovery strategy, or to restore large amounts of data. Whenever possible, the same number of devices that were used during the save operation should be used during a restore operation.

A Display Tape (DSPTAP) command displays the list of objects that the system saves across all media files. You only need one media file to display all of the objects that the system saved during a parallel save operation. This list also displays the number of media files that you need to restore data. However, you need all of the media files to restore any of the objects that the system saved. This might include multiple volumes.

IBM suggests that you use the same media definition object when saving and restoring the same objects. If you use a different media definition object when restoring, make sure that the same number of media files are defined within that media definition object. If the number of media file definitions is different from the number that exists on the storage media, you will receive an error message.

Related information

Create Media Definition (QSRCRTMD, QsrCreateMediaDefinition) API

Restore operation output file information

This table shows the format of the QASRRSTO (restore operation) output file. The Restore Library (RSTLIB), Restore Object (RSTOBJ), Restore Configuration (RSTCFG), Restore Deferred Objects (RSTDFROBJ), and Restore User Profiles (RSTUSRPRF) commands and the Restore Object List (QSRRSTO) API use the QASRRSTO output file.

Unused fields, fields that are not set, contain a value of zero for numeric fields and blanks for character fields.

Table 8. QASRRSTO (restore output file) information

Identifier	Type (in bytes)	Field
RSRECR	CHAR(1)	Record type

Table 8. QASRRSTO (restore output file) information (continued)

Identifier	Type (in bytes)	Field
RSOSTA	CHAR(1)	Status
RSOMSG	CHAR(7)	Message ID
RSOSEC	CHAR(1)	Security messages
RSNAME	CHAR(10)	Object name
RSMNAM	CHAR(10)	Member name
RSTYPE	CHAR(8)	Object type
RSATTR	CHAR(10)	Object attribute
RSSIZE	ZONED(15)	Size
RSOOWN	CHAR(10)	Owner
RSOTXT	CHAR(50)	Text
RSRPLV	CHAR(704)	Replacement variables
RSOLIB	CHAR(10)	Saved library name
RSOASP	ZONED(2)	Library ASP number
RSOSEQ	ZONED(4)	Sequence number
RSLABL	CHAR(17)	File label
RSSVDT	CHAR(13)	Save date/time
RSRSDT	CHAR(13)	Restore date/time
RSVLID	CHAR(450)	Volume identifiers
RSODEV	CHAR(40)	Device names
RSOSVF	CHAR(10)	Save file name
RSOSFL	CHAR(10)	Save file library name
RSOCMD	CHAR(10)	Restore command
RSSYSN	CHAR(8)	System name
RSORRL	CHAR(6)	Restore release
RSOSRL	CHAR(6)	Save release
RSOPGP	CHAR(10)	Primary group
RSOSQ2	ZONED(10)	Large sequence number
RSONLB	CHAR(10)	Restore library name
RSOFMT	CHAR(1)	Save format
RSOMFN	ZONED(3)	Media file number
RSOTMF	ZONED(3)	Total media files
RSOMDN	CHAR(10)	Media definition name
RSOMDL	CHAR(10)	Media definition library name
RSOVLC	ZONED(3)	Volume count

Table 8. QASRRSTO (restore output file) information (continued)		
Identifier	**Type (in bytes)**	**Field**
RSOVLL	ZONED(3)	Volume length
RSOVLD	CHAR(2400)	Volume identifiers (complete)
RSOOPT	CHAR(256)	Optical file
RSOAS1	CHAR(10)	ASP name
RSOAS2	ZONED(5)	ASP number
RSOPRT	CHAR(1)	Partial transaction exists
RSOJN	CHAR(10)	Journal name
RSOJL	CHAR(10)	Journal library name
RSOJRN	CHAR(10)	Journal receiver name
RSOJRL	CHAR(10)	Journal receiver library name
RSOJRA	CHAR(10)	Journal receiver ASP
RSOPFL	CHAR(10)	Spooled file name
RSOPFN	ZONED(6)	Spooled file number
RSOPJB	CHAR(10)	Spooled file job name
RSOPUN	CHAR(10)	Spooled file user name
RSOPJN	CHAR(6)	Spooled file job number
RSOPJS	CHAR(8)	Spooled file job system name
RSOPCD	CHAR(7)	Spooled file creation date
RSOPCT	CHAR(6)	Spooled file creation time
RSOPQN	CHAR(10)	Spooled file output queue name
RSOPQL	CHAR(10)	Spooled file output queue library
RSOPUD	CHAR(10)	Spooled file user data
RSOPFT	CHAR(10)	Spooled file form type
RSOPPG	PACKED(11)	Spooled file pages
RSOPCP	ZONED(3)	Spooled file copies
RSOPSZ	PACKED(15)	Spooled file size
RSOPXD	CHAR(7)	Spooled file expiration date
RSOPDR	CHAR(7)	Spooled file restored creation date
RSOPTR	CHAR(6)	Spooled file restored creation time
RSOPQR	CHAR(10)	Spooled file restored to output queue name
RSOPLR	CHAR(10)	Spooled file restored to output queue library
RSOPVA	CHAR(1)	Private authorities
RSOPVS	ZONED(10)	Number of private authorities saved
RSOPVR	ZONED(10)	Number of private authorities restored

Table 8. QASRRSTO (restore output file) information (continued)		
Identifier	**Type (in bytes)**	**Field**
RSODFR	CHAR(1)	Defer status
RSODFI	CHAR(10)	Defer ID

Related concepts

Verifying whether objects are restored successfully

You can use the job log or an output file to determine which objects are restored successfully.

Field descriptions for restore output file

This table shows the possible values for the restore (RST) output fields in the QASRRSTO (restore operation) output file.

ASP name

The auxiliary storage pool (ASP) device name of the object when it was restored. The possible values are:

***SYSBAS**

System and basic auxiliary storage pools.

device name

Name of the independent auxiliary storage pool.

ASP number

The auxiliary storage pool (ASP) of the object when it was restored. The possible values are:

1

System ASP.

2-32

Basic user ASPs

33-255

Independent ASPs

Defer ID

The identifier used to defer the restore of database files.

Defer status

Indicates whether the object is deferred. The possible values are:

'0'

The object is not deferred.

'1'

The object is deferred. It is dependent on a missing file, missing journal, or both.

'2'

The file dependency of the deferred object is complete.

'3'

The object is deferred. It is dependent on a missing file.

'4'

The journal dependency of the deferred object is complete. The object is now journaled.

'5'

The object is deferred. It is not journaled because it is dependent on a missing journal.

Device names

The name of the devices used to perform the restore operation. The field contains a list of device names. Each device name is CHAR(10) and there can be 1–4 devices listed.

File label

The file label of the media file used by the restore operation. If a save file was used, this field is blank.

Journal library name
The name of the library that contains the journal to which the object was journaled when it was saved.

Journal name
The name of the journal to which the object was journaled when it was saved.

Journal receiver ASP
The name of the auxiliary storage pool (ASP) that contains the earliest journal receiver needed for applying journal changes when recovering the object.

Journal receiver library name
The name of the library that contains the earliest journal receiver needed for applying journal changes when recovering the object.

Journal receiver name
The name of the earliest journal receiver needed for applying journal changes when recovering the object.

Large sequence number
The sequence number of the file on media. The value will be 0 if the restore media is not tape.

Library ASP number
The auxiliary storage pool (ASP) of the library to which the object was restored. The possible values are:

1
System ASP.

2-32
Basic user ASPs

-1
Independent ASPs. See the **ASP number** field for the independent ASP number.

Media definition library name
The name of the library that contains the media definition used in the restore operation.

Media definition name
The name of the media definition used in the restore operation.

Media file number
A number to identify this media file when a library is saved in parallel format. This field is only valid if the **Save format** field is '1' (save format is parallel). The value will be 0 if the save media is not tape.

Member name
The name of the database file member that was restored. This field will be blank if the object is not a database file, or if INFTYPE(*MBR) is not specified, or if the record is the summary record for the database file.

Message ID
The message identifier of an error message that was issued for this object or an error or completion message that was issued for this library.

Number of private authorities restored
The number of private authorities restored for the object.

Number of private authorities saved
The number of private authorities saved with the object.

Object attribute
The attribute of the object that was restored.

Object name
The name of the object that was restored.

Object type
The type of the object.

Optical file
The name of the optical file used by the restore operation. For a restore operation that does not use optical, this field is blank. This field is a variable-length field.

Owner
The name of the user profile that owns the restored object.

Partial transaction exists
Indicates whether this object was saved with one or more partial transactions. If you restore an object that was saved with partial transactions, you cannot use the object until you apply or remove journal changes. To apply or remove journal changes you will need the journal identified by the **Journal name** field and the journal receivers starting with the one identified by the **Journal receiver name** field. The possible values are:

'0'
The object was saved with no partial transactions.

'1'
The object was saved with one or more partial transactions.

Primary group
The name of the primary group for the restored object.

Private authorities
Indicates whether saved private authorities were requested to be restored. The possible values are:

'0'
PVTAUT(*NO) was specified.

'1'
PVTAUT(*YES) was specified.

Record type
The type of outfile record. The possible values are:

'1'
Object information.

'2'
Library information.

'3'
Restore operation. This record type appears at the end of a multiple library restore operation.

Replacement variables
The message data for the message identified in the Message ID field. This field is a variable-length field.

Restore command
The command that was used when the operation was performed. The possible values are:

RSTCFG
Restore configuration operation.

RSTDFROBJ
Restore deferred objects.

RSTLIB
Restore library operation.

RSTOBJ
Restore object operation.

RSTUSRPRF
Restore user profiles operation.

Restore date/time
The time at which the object was restored in CYYMMDDHHMMSS format.

Restore library name
The name of the library to which the object was restored.

Restore release
The release level of the operating system on which the objects were restored. This field has a VvRrMm format, containing the following:

Vv
The character V is followed by a 1-character version number.

Rr
The character R is followed by a 1-character release number.

Mm
The character M is followed by a 1-character modification number.

Save date/time
The time at which the object was saved in CYYMMDDHHMMSS format.

Save file name
The name of the save file used in the restore operation.

Save file library name
The name of the library that contains the save file used in the restore operation.

Save format
Indicates whether the data was saved in serial or parallel format. The possible values are:

'0'
The save format is serial.

'1'
The save format is parallel.

Save release
The target release level of the operating system for which the objects were saved. This field has a VvRrMm format, containing the following:

Vv
The character V is followed by a 1-character version number.

Rr
The character R is followed by a 1-character release number.

Mm
The character M is followed by a 1-character modification number.

Saved library name
The name of the library that contained the objects when they were saved.

Security messages
Indicates whether any security messages were sent for this object. The possible values are:

'0'
No security messages were sent for this object.

'1'
One or more security messages were sent for this object.

Sequence number
The sequence number of the tape file on media. The possible values are:

-5
The sequence number is larger than 9999. Use the value in the **Large sequence number** field.

0
The media used was not tape.

1-9999
The sequence number of the tape file.

Size
The size of the object.

Spooled file copies
The number of copies for the spooled file.

Spooled file creation date
The date when the spooled file was created.

Spooled file creation time
The time when the spooled file was created.

Spooled file expiration date
The expiration date of the restored spooled file.

Spooled file form type
The form type of the spooled file.

Spooled file job name
The name of the job that owns the spooled file.

Spooled file job number
The number of the job that owns the spooled file.

Spooled file job system name
The name of the system where the job that owns the spooled file ran.

Spooled file name
The name of the spooled file.

Spooled file number
The number of the spooled file in the job that owns it.

Spooled file output queue library
The name of the output queue library that contained the spooled file when it was saved.

Spooled file output queue name
The name of the output queue that contained the spooled file when it was saved.

Spooled file pages
The number of pages for the spooled file.

Spooled file restored creation date
The date when the restored spooled file was created.

Spooled file restored creation time
The time when the restored spooled file was created.

Spooled file restored to output queue library
The name of the output queue library to which the spooled file was restored.

Spooled file restored to output queue name
The name of the output queue to which the spooled file was restored.

Spooled file size
The size of the spooled file.

Spooled file user data
The user data of the spooled file.

Spooled file user name
The name of the user who owns the spooled file.

Status
Indicates whether the object restored successfully. The possible values are:

'0'
The object did not restore successfully.

'1'
The object restored successfully.

'2'
The object was excluded from the restore operation.

'3'
The object was restored but not converted.

'4'
Only private authorities were restored for the object.

System name
The name of the system on which the restore operation was performed.

Text
The text description of the object.

Total media files
The total number of media files created for a library saved in parallel format. This field is only valid if the **Save format** field is '1' (save format is parallel). The value will be 0 if the save media is not tape.

Volume count
The number of volume identifiers in the **Volume identifiers (complete)** fields.

Volume identifiers
The list of volume identifiers that are used during this restore operation. The list can contain from one to 75 six-character volume identifiers. See the **Volume count** field to tell how many volume identifiers are in the list. This field is a variable-length field.

Volume identifiers (complete)
The list of volume identifiers that are used during this restore operation. The list can contain from one to 75 volumes. See the **Volume count** field to tell how many volume identifiers are in the list. This field is a variable-length field.

Volume length
The length of each volume identifier in the **Volume identifiers (complete)** field.

Chapter 3. Selecting the appropriate recovery strategy

Use this information to determine the correct procedure for recovering your system.

Before beginning your recovery, you must perform these tasks.

- Make sure you understand what has caused the problem. Understanding the cause helps you choose the correct recovery steps.
- Plan your recovery. Use Table 9 on page 57 in the Choosing the recovery procedure for a disk failure or disk errors topic to find the appropriate recovery checklist for your situation.
- Make a copy of the checklist and check off each step as you complete it.
- Keep the checklist for future reference.
- Keep a record of what you have already done and what you do for the rest of your recovery. This record is important if you need help later.
- If your problem requires hardware or software service, make sure that you understand what the service representative did. Do not be afraid to ask questions, such as:
  - Was a disk unit replaced? If yes, which one?
  - Was the Licensed Internal Code restored? If yes, what option from the Install Licensed Internal Code (LIC) menu was used?
  - Did the disk configuration need to be recovered? Was it successful?
  - Can the failed disk unit be pumped? How successfully?

Some common recovery terminology

You might need to understand these terms when you are discussing your situation with your service representative or software support.

Term	Definition
abnormal end (abend)	A system failure or operator action that causes the system to end without being able to end all jobs and close all files. Your system might end abnormally because of a power failure or a problem with certain hardware or software components.
auxiliary storage pool	A group of disk units defined from all the disk units that make up auxiliary storage. Auxiliary storage pools (ASPs), also known as disk pools in IBM Navigator for i, isolate objects on one or more specific disk units. Isolation of objects can reduce the loss of data due to a disk media failure. Usually, only data that is stored on disk units in the affected ASP is lost.
ASP group	An auxiliary storage pool (ASP) group is made up of a primary ASP and zero or more secondary ASPs. Related to data storage, each ASP is independent but in the ASP group they combine to act as one entity. If one ASP is made available or unavailable, the rest of the ASPs in the group are also made available or unavailable at the same time. In a clustered environment, all the ASPs in a group switch to another node at the same time. The primary and secondary ASPs also use the same database.
basic ASP	A user ASP created by grouping a physical set of disk units and assigning them a number between 2 and 32. A basic ASP is always available for data storage and access.

© Copyright IBM Corp. 1997, 2025

Term	Definition
dedicated service tools (DST)	A set of tools to work with the system when the operating system is not available or not working.
disk configuration	An internal system table that tells how the physical disk units are arranged on your system. The disk configuration is used to assign units to an ASP. The disk configuration is stored on the load source unit.
disk pump	A commonly used term for the procedure used by a service representative to attempt to copy data from a disk unit that failed.
encrypting tape drive	A tape drive, such as the IBM System Storage TS1120 or IBM Ultrium 4 tape drive that provides data encryption and key management for backup data. An encrypting tape drive uses the IBM Encryption Key Manager (EKM) to manage the encryption keys.
encrypted ASP	An independent ASP or user ASP that is encrypted uses cryptographic services.
independent ASP	A user ASP that can be made available (varied on) and made unavailable (varied off) without restarting the system. An independent ASP can be either switchable between multiple systems or privately connected to a single system.
library user ASP	A user ASP that contains libraries, directories, and folders and all the objects associated with them.
Licensed Internal Code	The layer of IBM i architecture that is next to the hardware. You must have the Licensed Internal Code on your system before you can restore the operating system.
load source unit	The first unit (unit 1) in the system ASP. It contains the Licensed Internal Code and the disk configuration for your system.
nonlibrary user ASP	A user ASP that might contain journals, journal receivers, and save files. The libraries that are associated with these objects are in the system ASP. A non-library user ASP is sometimes called an *old-style ASP* because it was the only type of user ASP available before Version 1 Release 3 of the OS/400® licensed program.
primary ASP	An independent auxiliary storage pool (ASP) that defines a collection of directories and libraries. There might also be secondary ASPs that are associated with the primary ASP. A primary ASP defines a database for itself and other ASPs that can be added in its ASP group.
secondary ASP	An independent ASP that defines a collection of directories and libraries and must be associated with a primary ASP in an ASP group. A possible use for a secondary ASP is to store journal receivers for objects that are journaled in the primary ASP.
software encryption	An encryption method that uses backup software encryption to encrypt data to a tape device. For example, BRMS uses keystore files from cryptographic services during encrypted backups.
system ASP	The system creates the system ASP and it is always configured. The system ASP (ASP 1) contains the Licensed Internal Code, licensed programs, and system libraries. The system ASP might also contain user libraries, folders, and directories. The system ASP contains all configured disk units that are not assigned to a user ASP.
system service tools (SST)	A subset of the DST tools. The tools available through SST, such as displaying the disk configuration, can be used while the operating system is running and other users are on the system.

Term	Definition
user ASP	A basic or independent ASP that is created by grouping a physical set of disk units. You can assign a basic ASP a number between 2 and 32. When you create an independent ASP, you specify a name; the system then assigns the independent ASP a number between 33 and 255. ASP 1 is always reserved as the system ASP.
UDFS ASP	An independent ASP that contains only user-defined file systems (UDFS). It cannot be a member of an ASP group unless it is converted to a primary or secondary ASP.

Related information

Planning for single system independent disk pools

Recovery procedure for a power failure

If your system stops because power is lost, you need to follow special procedures when you start the system again.

If you experience frequent power outages, consider using an uninterruptible power supply for your system.

If power loss to workstations causes your system to perform constant error recovery, you should modify your applications to handle losing communications to workstations.

Related information

Examples: Techniques and programming for backup and recovery

Recovery procedure for a system failure

A system failure is a problem either with hardware (other than disk) or with operating system software that causes your system to end abnormally.

After your service representative has corrected the problem, follow the procedure to start your system after an abnormal end.

If the service representative replaced a disk unit, use the information in the "Choosing the recovery procedure for a disk failure or disk errors" on page 56 topic to determine the correct recovery procedure.

Related concepts

Starting the system after it ends abnormally
Use this information to start the system after it ends abnormally.

Recovery procedure for a program failure or human error

You might need to recover objects because a program updated them incorrectly or because a user deleted them.

See Chapter 9, "Restoring specific types of information," on page 209 topic for the kind of objects you are restoring. Some objects have special considerations or need to be restored in a particular sequence.

If you are restoring an object that does not exist on the system, private authorities for the object are not restored with it. You can do one of the following tasks:

- Reconstruct the private authorities manually, using the Edit Object Authority (EDTOBJAUT) display.
- Restore private authorities by using this procedure:
  1. Restore all user profiles from your most recent SAVSYS or SAVSECDTA tape. Type: RSTUSRPRF. Restoring user profiles requires a restricted state.
  2. Restore the objects you need to recover.

3. Restore authorities. Type: RSTAUT. Only one Restore Authority (RSTAUT) command can run on the system at any given time.

Recovering the vital product data information if the partition does not IPL in A-mode or B-mode

The vital product data (VPD) information contains information about the disk unit which is the load-source disk unit.

If the VPD information is lost, the partition might not IPL in A-mode or B-mode. Use the Work with Disk Units service function in D-mode to refresh the VPD information.

1. Enter Option 8 (Update system vital product data) on the Work with Disk Units display. The following confirmation displays:

```
                Update system vital product data

The disk unit information below identifies the load source whose
location will be written to the system vital product data.

Press enter to confirm the vital product data should be written.

             Serial
ASP   Unit   Number       Type    Model
  1      1   $$$$$$$$$$   $$$$    $$$
```

   The load-source disk unit is identified in the display above.
2. Press Enter to update the VPD.
3. Press F3 to return to the Use Dedicated Service Tools (DST) menu.
4. On the Use Dedicated Service Tools (DST) menu, select option 7 (Start a service tool).
5. On the Start a Service Tool menu, select option 7 (Operator panel functions) .
6. On the Operator Panel Functions display, select the IPL source and IPL mode. Press F8 to restart the partition or press F10 to power off the partition.

```
                  Operator Panel Functions
                                                          System:   XXXX
IPL source:  2   (1=A, 2=B or 3=D)
IPL mode:    1   (1=Manual, 2=Normal, 3=Secure or 4=Auto)

Press Enter to change the IPL attributes and return
  to the main DST menu.
Press F8 to set the IPL attributes and restart the system.
  Machine processing will be ended and the system will be
  restarted.
Press F10 to set the IPL attributes and power off the system.
  Machine processing will be ended and the system will be
  powered off.
Press F12 to return to the main DST menu without changing
  IPL attributes.
```

After the VPD information is updated, you can IPL the partition in A-mode or B-mode.

Related information

Installing IBM i and related software on a new system or logical partition

Choosing the recovery procedure for a disk failure or disk errors

Use this information to determine the actions that you need to take when you recover your system because of a failed or damaged disk unit.

Note: If you receive a system reference code (SRC) code that indicates a disk problem, do not perform an initial program load (IPL) before your service representative arrives. If you perform an IPL, your service representative might not be able to recover data from the damaged disk unit.

The steps you follow to recover from a disk failure depend on the following items:

- What unit failed.
- Whether disk protection, such as device parity protection or mirrored protection is active.
- Whether you have user auxiliary storage pools (ASPs) configured.
- Whether some or all of the sectors on the disk are damaged. If a disk unit must be replaced, a service representative normally tries to copy the information from the disk unit when it is replaced. This procedure is sometimes referred to as a disk pump.

Use Table 9 on page 57 to determine what recovery procedure you should follow, based on the failure that has occurred on your system. To find your situation on the chart, ask your service representative whether data was copied successfully (the results of the disk pump).

Service representative terminology	Terminology in recovery charts
Full pump	None of the data is lost
Partial pump	Some of the data is lost
Could not pump	All of the data is lost

Recovery for disk errors that do not require disk replacement: Some types of disk units automatically recover from errors without needing to be replaced. In some cases, however, sectors are damaged before the disk unit reassigns them and some object damage occurs. If you receive a message indicating that object damage has occurred and disk sectors have been reassigned, consider this to be the value Some for the column *Data loss on the failed unit* in Table 9 on page 57.

If you are recovering from disk errors but you did not need a service representative to replace the disk unit, you might need to perform tasks that are normally performed by a service representative. Make a copy of the appropriate checklist and mark it as follows:

1. Begin at the task immediately following "Attach the new disk unit".
2. If the checklist contains a task called, "Restore the disk unit data", skip that task.

Table 9. Choosing the correct recovery procedure for disk media failure

Type of unit that failed	Data loss on failed unit	Availability protection on failed unit	User APSs configured?	Procedure to follow
Any	N/A	Mirrored protection	N/A [1]	“Checklist 14: Actions for non-load source disk unit failure” on page 87
Any	N/A	Device parity protection	N/A [1]	“Checklist 15: Actions for non-load source disk unit failure” on page 87
Load source unit	None	None	N/A [1]	“Checklist 1: Actions for load-source disk unit failure” on page 59
Load source unit	Some[2]	None	N/A [1]	“Checklist 2: Actions for load-source disk unit failure” on page 60
Load source unit	All	None	No	“Checklist 3: Actions for load-source disk unit failure” on page 61
Load source unit. No basic ASPs in overflowed status[3]	All	None	Yes	“Checklist 4: Actions for load-source disk unit failure” on page 62

Table 9. Choosing the correct recovery procedure for disk media failure (continued)

Type of unit that failed	Data loss on failed unit	Availability protection on failed unit	User APSs configured?	Procedure to follow
Load source unit. One or more basic ASPs in overflowed status[3].	All	None	Yes	"Checklist 5: Actions for load-source disk unit failure" on page 66
Non-Load source unit in system ASP[4]	None	None	N/A [1]	"Checklist 6: Actions for non-load source disk unit failure or disk units in basic user auxiliary storage pool disk failure" on page 70
Non-Load source unit in system ASP[4]	Some[2]	None	N/A [1]	"Checklist 7: Actions for non-load source disk unit failure" on page 71
Non-Load source unit in system ASP[4]	All	None	No	"Checklist 8: Actions for non-load source disk unit failure" on page 72
Non-Load source unitin system ASP[4]. No basic ASPs in overflowed status[3].	All	None	Yes	"Checklist 9: Actions for non-load source disk unit failure" on page 73
Non-Load source unit in system ASP[4]. One or more basic ASPs in overflowed status[3].	All	None	Yes	"Checklist 10: Actions for non-load source disk unit failure" on page 77
Disk unit in basic ASP	None	None	Yes	"Checklist 6: Actions for non-load source disk unit failure or disk units in basic user auxiliary storage pool disk failure" on page 70
Disk unit in basic ASP	Some[2]	None	Yes	"Checklist 11: Actions for a failure in a basic auxiliary storage pool disk unit" on page 81
Disk unit in basic ASP. Failed unit not in overflowed status[3].	All	None	Yes	"Checklist 12: Actions for a failure in a basic auxiliary storage pool disk unit" on page 82
Disk unit in basic ASP. Failed unit in overflowed status[3].	All	None	Yes	"Checklist 13: Actions for a failure in a basic auxiliary storage pool disk unit" on page 84
Disk unit in independent ASP	None	None	Yes	"Checklist 17: Actions for independent auxiliary storage pool disk failure" on page 89
Disk unit in independent ASP	Some[2]	None	Yes	"Checklist 18: Actions for a failure in an independent auxiliary storage pool disk unit" on page 90

Table 9. Choosing the correct recovery procedure for disk media failure (continued)

Type of unit that failed	Data loss on failed unit	Availability protection on failed unit	User APSs configured?	Procedure to follow
Disk unit in independent ASP	All	None	Yes	"Checklist 19: Actions for a failure in an independent auxiliary storage pool disk unit" on page 91
Cache storage in input/output processor (IOP)	Some	N/A	N/A[1]	"Checklist 23: Actions for a failed cache card" on page 104

1

The recovery procedure is the same whether user ASPs are configured.

2

If the service representative was partially successful in saving data from a failed disk unit, you should consider treating the situation as a complete data loss on the failed unit.

3

Step 4 in Resetting an overflowed user ASP without an IPL topic describes how to determine whether a user ASP is in overflowed status.

4

If a unit in your system ASP fails and a replacement is not immediately available, you can use the procedure in "Checklist 16: Actions for non-load source disk unit failure" on page 88 topic. This procedure allows you to return your system to operation. You will have less disk storage and you will need to recover all the data in the system ASP.

Checklist 1: Actions for load-source disk unit failure

This checklist shows the sequence of steps you should use to recover after a load-source disk unit failure when no data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Load source unit

Data loss:
None

User ASP configured:
N/A

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Note: When your load source unit is failed and you are recovering from distribution media with LAN Console configured, you might be required to use another console type for your recovery.

Table 10. Recovery checklist for disk failure–Checklist 1

Task	What to do	Where to read more about it
Actions to be performed by the service representative		

Table 10. Recovery checklist for disk failure–Checklist 1 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 1	Save the disk unit data.	
___ Task 2	Attach the new disk unit.	
___ Task 3	Install the Licensed Internal Code using option 4 (Install Licensed Internal Code and Restore Disk Unit Data).	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128.
___ Task 4	Restore the disk unit data.	
Actions to be performed by the user		
___ Task 5	You must perform an initial program load (IPL) at this time. Follow the procedure for starting the system after it ends abnormally.	Chapter 6, "Starting the system after it ends abnormally," on page 159, task 1 through task 4.

Related information

Console Advanced Topics

Checklist 2: Actions for load-source disk unit failure

This checklist shows the sequence of steps you should use to recover after a load-source disk unit failure when some data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Load source unit

Data loss:
Some

User ASP configured:
N/A

Most steps in the checklist include references to other tasks in this topic. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Note: When your load source unit is failed and you are recovering from distribution media with LAN Console configured, you might be required to use another console type for your recovery.

Table 11. Recovery checklist for disk failure–Checklist 2		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Save the disk unit data.	
___ Task 2	Attach the new disk unit.	
___ Task 3	Install the Licensed Internal Code using option 4 (Install Licensed Internal Code and Restore Disk Unit Data).	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128.
___ Task 4	Restore the disk unit data.	
Actions to be performed by the user		

Table 11. Recovery checklist for disk failure–Checklist 2 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 5	You must perform an initial program load (IPL) at this time. Follow the procedure for starting the system after it ends abnormally.	Chapter 6, "Starting the system after it ends abnormally," on page 159, task 1 through task 4.
___ Task 6	Restore the operating system. You are performing an abbreviated install operation. **Note:** You might have some objects that are damaged. You might be required to perform a complete restore of the operating system.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.
___ Task 7	If you restored the operating system using distribution tapes, some system information, such as access path recovery times and the system reply list, might have been reset to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 8	Reclaim storage.	"Reclaim Storage (RCLSTG) command" on page 28.
___ Task 9	Evaluate the extent of the damage. Determine whether you will attempt to recover damaged objects or restore the entire system. Do not skip this step.	"Task 4: Recovering from damaged objects and unreadable sectors" on page 165.
___ Task 10	If you have decided to do a complete restore operation, determine the correct procedure for restoring user information.	"Choosing the procedure to recover user information" on page 105.
___ Task 11	If you have decided to attempt to recover damaged objects, perform the tasks described in "Task 4: Recovering from damaged objects and unreadable sectors" on page 165.	

Related information

Console Advanced Topics

Checklist 3: Actions for load-source disk unit failure

This checklist shows the sequence of steps you should use to recover after a load-source disk unit failure after a complete data loss.

This checklist should be used for the following problem situation:

Failed unit:
Load source unit

Data loss:
All

User ASP configured:
No

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Note: When your load source unit is failed and you are recovering from distribution media with LAN Console configured, you might be required to use another console type for your recovery.

Table 12. Recovery checklist for disk failure–Checklist 3

Task	What to do	Where to read more about it
Actions to Be Performed by the Service Representative		
___ **Task 1**	Attach the new disk unit.	
___ **Task 2**	Prepare to load the Licensed Internal Code.	"Preparing for loading the Licensed Internal Code" on page 122.
___ **Task 3**	Install the Licensed Internal Code using option 3 (Install Licensed Internal Code and Recover Configuration).	"Loading the Licensed Internal Code" on page 128.
___ **Task 4**	Recover the disk configuration (assignment of disks to auxiliary storage pools (ASPs) and protection).	"Recovering your disk configuration" on page 136.
Actions to be performed by the user		
___ **Task 5**	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are performing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.
___ **Task 6**	If you restored the operating system using distribution media, some system information, such as access path recovery times and the system reply list, was returned to default values. Set these values correctly.	"Recovering system information" on page 209.
___ **Task 7**	Use the "Choosing the procedure to recover user information" on page 105 topic to determine the correct procedure for recovering user information. You will need to recover all user data.	

Related information

Console Advanced Topics

Checklist 4: Actions for load-source disk unit failure

The checklist shows the sequence of steps to use to recover after a failure on a load-source disk unit with complete data loss. A user auxiliary storage pool was configured.

This checklist is used for the following problem situation:

Failed unit
Load source unit

Data loss
All

User ASP is configured
Yes

Basic user ASP overflowed
No

Attention: When you replace a disk unit in your system auxiliary storage pool (ASP), the system loses addressability to the objects in your basic user ASPs. Recovering object ownership for objects, other than document library objects (DLOs), you must manually assign ownership for every object in every basic user ASP. You might want to treat this situation as a total recovery and restore all your information from your save media if the following conditions are true:

- You have many objects in your basic user ASPs
- You backed up your system

Run the steps that are described in "Checklist 20: Recovering your entire system after a complete system loss" on page 93 topic to recover your system.

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative run the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects, if they do not apply in your situation.

Note: When your load source unit fails and you are recovering from distribution media, you might not be able to use the LAN Console for the recovery.

Table 13. Recovering from a disk failure–Checklist 4

Task	What to do	Where to read more about it
Actions run by the service representative		
___ Task 1	Attach the new disk unit.	
___ Task 2	Prepare to load the Licensed Internal Code[1] using the most recent Save System (SAVSYS) media.	"Preparing for loading the Licensed Internal Code" on page 122.
___ Task 3	Install the Licensed Internal Code by selecting option 3 (Install Licensed Internal Code and Recover Configuration).	"Loading the Licensed Internal Code" on page 128.
___ Task 4	Recover the disk configuration (assignment of disks to ASPs and protection).	"Recovering your disk configuration" on page 136.
Actions run by the user		
___ Task 5	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 6	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are doing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.

Table 13. Recovering from a disk failure–Checklist 4 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 7	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 8	If you restored the operating system from distribution media, some system information was reset to default values, such as access path recovery times and the system reply list. Set these values correctly.	"Recovering system information" on page 209.
___ Task 9	If necessary, change the QALWOBJRST system value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 10	If necessary, change the QVFYOBJRST system value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 11	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 12	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 13	Describe or diagram, as much as possible, the contents of your user ASPs before the failure.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 14	Recover user profiles, configuration, libraries in the system ASP, and the contents of your basic user ASPs. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP" on page 172, task 1 through task 11.
___ Task 15	Restore document library objects.	"Restoring documents and folders" on page 260.
___ Task 16	Restore your last complete save of directories.[1]	"Restoring objects in directories" on page 265.

Table 13. Recovering from a disk failure–Checklist 4 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 17	If you have user-defined file systems in user ASPs that do not restore correctly, you might need to run more recovery steps.	“Task 7: Restoring user-defined file systems to the user auxiliary storage pool” on page 178.
___ Task 18	Restore changed objects and apply journaled changes.	Chapter 13, “Restoring changed objects and apply journaled changes,” on page 277, task 1 through task 7.
___ Task 19	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	Chapter 13, “Restoring changed objects and apply journaled changes,” on page 277.
___ Task 20	If the Save System Information (SAVSYSINF) was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved with the Save System Information (SAVSYSINF) command.	Chapter 12, “Restoring system information,” on page 275.
___ Task 21	Restore authority. Type - RSTAUT	“Restoring object authorities” on page 217.
___ Task 22	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, “Restoring program temporary fixes (PTFs),” on page 273.
___ Task 23	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 24	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 25	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.

Table 13. Recovering from a disk failure–Checklist 4 (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 26	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 27	Do an IPL now.	“Performing a normal initial program load” on page 43.

___ Task 28	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 29	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 30	Review job logs or output from your restore operations to verify that all objects were restored successfully.	“Verifying whether objects are restored successfully” on page 39.

Related information

Console Advanced Topics

Checklist 5: Actions for load-source disk unit failure

The checklist shows the sequence of steps to use to recover after a failure on a load-source disk unit with complete data loss. A user auxiliary storage pool was configured, and the basic user ASP overflowed.

The checklist is used for the following problem situation -

Failed unit
Load source unit

Data loss
All

User ASP is configured
Yes

Basic user ASP overflowed
Yes

Attention: When you replace a disk unit in your system auxiliary storage pool (ASP), the system loses addressability to the objects in your basic user ASPs. Recovering object ownership for objects other than document library objects (DLOs) requires manually assigning ownership for every object in every basic user ASP. You might want to treat this situation as a total recovery and restore all your information from your save media if the following conditions are true:

- You have many objects in your user ASPs
- You thoroughly backed up your system

Run the steps that are described in "Checklist 20: Recovering your entire system after a complete system loss" on page 93 to recover your system.

Table 14. Recovering from a disk failure–Checklist 5

Task	What to do	Where to read more about it
Actions run by the service representative		
___ Task 1	Attach the new disk unit.	
___ Task 2	Prepare to load the Licensed Internal Code by using the most recent Save System (SAVSYS) media.	"Preparing for loading the Licensed Internal Code" on page 122.
___ Task 3	Install the Licensed Internal Code by selecting option 3 (Install Licensed Internal Code and Recover Configuration).	"Loading the Licensed Internal Code" on page 128.
___ Task 4	Recover the disk configuration (assignment of disk to ASPs and protection).	"Recovering your disk configuration" on page 136.
Actions run by the user		
___ Task 5	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 6	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are running a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.
___ Task 7	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.

Table 14. Recovering from a disk failure–Checklist 5 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 8	If you restored the operating system from distribution media, some system information might be reset to default values, such as access path recovery times and the system reply list. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 9	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 10	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 11	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 12	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 13	Describe or diagram, as much as possible, the contents of your user ASPs before the failure.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 14	Recover user profiles, configuration, libraries in the system ASP, and the contents of the user ASPs that were not in overflowed status. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP" on page 172, task 1 through task 11.
___ Task 15	Recover the objects in the user ASPs that were overflowed.	"Recovering a damaged basic auxiliary storage pool" on page 186, task 1 through task 9.
___ Task 16	Restore document library objects to the system ASP and to any overflowed user ASPs with DLOs.	"Restoring documents and folders" on page 260.
___ Task 17	Restore your last complete save of directories.[1]	"Restoring objects in directories" on page 265.
___ Task 18	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to run other recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.

Table 14. Recovering from a disk failure–Checklist 5 (continued)

Task	What to do	Where to read more about it
___ Task 19	Restore changed objects and apply journaled changes.	Chapter 13, “Restoring changed objects and apply journaled changes,” on page 277.
___ Task 20	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	Chapter 13, “Restoring changed objects and apply journaled changes,” on page 277.
___ Task 21	If the Save System Information (SAVSYSINF) command was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved by using the Save System Information (SAVSYSINF) command.	Chapter 12, “Restoring system information,” on page 275.
___ Task 22	Restore authority. Type - RSTAUT	“Restoring object authorities” on page 217.
___ Task 23	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, “Restoring program temporary fixes (PTFs),” on page 273.
___ Task 24	If necessary, use the WRKSYSVAL command to change the QALWOBJRST system value back to its original value.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 25	If necessary, use the WRKSYSVAL command to change the QVFYOBJRST system value back to its original value.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 26	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.

___ Task 27	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

Table 14. Recovering from a disk failure–Checklist 5 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 28	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

___ Task 29	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 30	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 31	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 6: Actions for non-load source disk unit failure or disk units in basic user auxiliary storage pool disk failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit failure but no data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Non-load source unit in system ASP or disk unit in basic user ASP

Data loss:
None

User ASP Configured:
N/A

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 15. Recovery checklist for disk failure–Checklist 6		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Save the disk unit data.	
___ Task 2	Attach a new disk unit.	
___ Task 3	Restore data to the new disk unit.	
Actions to be performed by the user		

Table 15. Recovery checklist for disk failure–Checklist 6 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 4	Perform an initial program load (IPL). Follow the procedure for starting the system after it ends abnormally.	Chapter 6, "Starting the system after it ends abnormally," on page 159, task 1 through task 4.

Checklist 7: Actions for non-load source disk unit failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit failure but some data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Non-load source unit in system auxiliary storage pool (ASP)

Data loss:
Some

User ASP configured:
N/A

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 16. Recovery checklist for disk failure–Checklist 7		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Save the disk unit data.	
___ Task 2	Attach the new disk unit.	
___ Task 3	Restore the disk unit data.	
Actions to be performed by the user		
___ Task 4	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 5	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are performing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.

Table 16. Recovery checklist for disk failure–Checklist 7 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 6	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography information center topic.
___ Task 7	If you restored the operating system using distribution media, some system information, such as access path recovery times and the system reply list, might have been reset to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 8	Reclaim storage.	"Reclaim Storage (RCLSTG) command" on page 28.
___ Task 9	Evaluate the extent of the damage. Determine whether you will attempt to recover damaged objects or restore the entire system. Do not skip this step.	"Task 4: Recovering from damaged objects and unreadable sectors" on page 165.
___ Task 10	If you have decided to do a complete restore operation, determine the correct procedure for recovering user information.	"Choosing the procedure to recover user information" on page 105.
___ Task 11	If you have decided to attempt to recover damaged objects, perform the tasks in "Task 4: Recovering from damaged objects and unreadable sectors" on page 165.	

Checklist 8: Actions for non-load source disk unit failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit failure after a complete data loss.

This checklist should be used for the following problem situation:

Failed unit:
Non-load source unit in system auxiliary storage pool (ASP)

Data loss:
All

User ASP configured:
No

Most steps in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 17. Recovery checklist for disk failure–Checklist 8		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Attach the new disk unit.	

Table 17. Recovery checklist for disk failure–Checklist 8 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 2	Delete the ASP data.	
___ Task 3	Restore the Licensed Internal Code using option 1 (Restore Licensed Internal Code) from the latest Save System (SAVSYS) media. If user ASPs are configured, they remain intact.	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128.
Actions to be performed by the user		
___ Task 4	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 5	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are performing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.
___ Task 6	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 7	If you restored the operating system using distribution media, some system information, such as access path recovery times and the system reply list, might have been reset to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 8	Reclaim storage.	"Reclaim Storage (RCLSTG) command" on page 28.
___ Task 9	Use "Choosing the procedure to recover user information" on page 105 to determine the correct procedure for recovering user information.	

Checklist 9: Actions for non-load source disk unit failure

The checklist shows the sequence of steps to use to recover after a non-load-source disk unit failure after a complete data loss. A user auxiliary storage pool (ASP) is configured.

This checklist is used for the following problem situation:

Failed unit
Non-load source unit in system auxiliary storage pool (ASP)

Data loss
All

User ASP is configured
Yes

User ASP is overflowed
No

Attention: When you replace a disk unit in your system ASP, the system loses addressability to the objects in your user ASPs. Recovering object ownership for objects other than document library objects (DLOs) requires manually assigning ownership for every object in every user ASP. You might want to treat this situation as a total recovery and restore all your information from your save media if the following conditions are true:

- You have many objects in your user ASPs
- You backed up your system

- You have many objects in your user ASPs
- You backed up your system

Run the steps that are described in "Checklist 20: Recovering your entire system after a complete system loss" on page 93 to recover your system.

Table 18. Recovering from a disk failure–Checklist 9

Task	What to do	Where to read more about it
Actions run by the service representative		
___ Task 1	Delete the data in the ASP that contains the failed unit.	
___ Task 2	Replace the failed disk unit.	
___ Task 3	Configure the replacement disk unit by adding it to the correct ASP.	
___ Task 4	Restore the Licensed Internal Code by selecting option 1 (Restore Licensed Internal Code) from the most recent Save System (SAVSYS) media. If user ASPs are configured, they remain intact.	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128.
Actions run by the user		
___ Task 5	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 6	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are running a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.

Table 18. Recovering from a disk failure–Checklist 9 (continued)

Task	What to do	Where to read more about it
___ Task 7	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 8	If distribution media was used to restore the operating system, some system information, such as access path recovery times and the system reply list, might be reset to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 9	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 10	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 11	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 12	After you change the system values, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 13	Describe or diagram, as much as possible, the contents of your user ASPs before the failure.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 14	Recover user profiles, configuration, libraries in the system ASP, and the contents of your basic user ASPs. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP" on page 172, task 1 through task 11.
___ Task 15	Restore document library objects to the system ASP.	"Restoring documents and folders" on page 260.
___ Task 16	Restore your last complete save of directories.[1]	"Restoring objects in directories" on page 265.

Table 18. Recovering from a disk failure–Checklist 9 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 17	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to run other recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 18	Restore changed objects and apply journaled changes.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277, task 1 through task 7.
___ Task 19	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277.
___ Task 20	If the Save System Information (SAVSYSINF) was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved with the Save System Information (SAVSYSINF) command.	Chapter 12, "Restoring system information," on page 275.
___ Task 21	Restore authority. Type - RSTAUT	"Restoring object authorities" on page 217.
___ Task 22	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 23	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 24	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 25	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.

Table 18. Recovering from a disk failure–Checklist 9 (continued)

Task	What to do	Where to read more about it

___ Task 26	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 27	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

___ Task 28	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 29	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 30	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 10: Actions for non-load source disk unit failure

The checklist shows the sequence of steps to use to recover after a non-load-source disk unit failure after a complete data loss. A user auxiliary storage pool (ASP) was configured and the basic user ASP overflowed.

The checklist is used for the following problem situation:

Failed unit
Non-load source unit in system auxiliary storage pool (ASP)

Data loss
All

User ASP is configured
Yes

Basic user ASP overflowed
Yes

Attention: When you replace a disk unit in your system ASP, the system loses addressability to the objects in your user ASPs. Recovering object ownership for objects other than document library object (DLO) requires manually assigning ownership for every object in every user ASP. You might want to treat this situation as a total recovery and restore all your information from your save media if the following conditions are true:

- You have many objects in your user ASPs
- You thoroughly backed up your system

Run the steps that are described in "Checklist 20: Recovering your entire system after a complete system loss" on page 93 to recover your system.

Table 19. Recovering from a disk failure–Checklist 10

Task	What to do	Where to read more about it
Actions run by the service representative		
___ Task 1	Physically remove the failed disk unit from the system.	
___ Task 2	Delete the data in the ASP that contains the failed unit. When you delete the data in the system ASP, the system also deletes data in any basic user ASPs that have a status of overflowed.	
___ Task 3	Install the replacement disk unit.	
___ Task 4	Configure the replacement disk unit by selecting the "Replace configured unit" function on the Work with Disk Units display.	
___ Task 5	Restore the Licensed Internal Code by selecting option 1 (Restore Licensed Internal Code) from the most recent Save System (SAVSYS) media. If user ASPs are configured and are not overflowed, they remain intact.	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128.
Actions run by the user		
___ Task 6	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 7	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are doing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.

Table 19. Recovering from a disk failure–Checklist 10 (continued)

Task	What to do	Where to read more about it
___ Task 8	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 9	If you restored the operating system from distribution media, some system information might be reset to default values, such as access path recovery times and the system reply list. Verify that these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 10	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 11	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 12	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 13	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 14	Describe or diagram, as much as possible, the contents of your user ASPs before the failure.	"Describing the contents of your user auxiliary storage pools" on page 171.
___ Task 15	Recover user profiles, configuration, libraries in the system ASP, and the contents of any basic user ASPs that were not in overflowed status. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP" on page 172, task 1 through task 11.
___ Task 16	Recover the objects in the basic user ASPs that were overflowed.	"Recovering a damaged basic auxiliary storage pool" on page 186, task 1 through task 9.
___ Task 17	Restore document library objects to the system ASP and to any overflowed basic user ASPs with DLOs.	"Restoring documents and folders" on page 260.

Table 19. Recovering from a disk failure–Checklist 10 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 18	Restore your last complete save of directories.[1]	"Restoring objects in directories" on page 265.
___ Task 19	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to run more recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 20	Restore changed objects and apply journaled changes.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277, task 1 through task 7.
___ Task 21	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	
___ Task 22	If the Save System Information (SAVSYSINF) command was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that were saved with the Save System Information (SAVSYSINF) command. **Note:** If you have independent ASPs, do not run this step now.	Chapter 12, "Restoring system information," on page 275.
___ Task 23	Restore authority. Type - RSTAUT	"Restoring object authorities" on page 217.
___ Task 24	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 25	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 26	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 27	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.

Table 19. Recovering from a disk failure–Checklist 10 (continued)

Task	What to do	Where to read more about it

___ Task 28	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 29	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

___ Task 30	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 31	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 32	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 11: Actions for a failure in a basic auxiliary storage pool disk unit

This checklist shows the sequence of steps you should use to recover after basic auxiliary storage pool (ASP) has failed and some data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Basic ASP

Data loss:
Some

Most steps in the checklist include references to other tasks in this topic. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 20. Recovery checklist for disk failure–Checklist 11		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Save the disk unit.	
___ Task 2	Attach the new disk unit.	
___ Task 3	Restore the disk unit data.	
Actions to be performed by the user		
___ Task 4	You must perform an initial program load (IPL) at this time. Follow the procedure for starting the system after it ends abnormally.	Chapter 6, "Starting the system after it ends abnormally," on page 159, task 1 through task 4.
___ Task 5	Reclaim storage.	"Reclaim Storage (RCLSTG) command" on page 28.
___ Task 6	Evaluate the extent of the damage. Determine whether you will attempt to recover damaged objects or restore the entire system. Do not skip this step.	"Task 4: Recovering from damaged objects and unreadable sectors" on page 165.
___ Task 7	If you have decided to do a complete restore operation, determine the correct procedure for recovering user information.	"Choosing the procedure to recover user information" on page 105.
___ Task 8	If you have decided to attempt to recover damaged objects, perform the tasks in "Task 4: Recovering from damaged objects and unreadable sectors" on page 165.	

Checklist 12: Actions for a failure in a basic auxiliary storage pool disk unit

This checklist shows the sequence of steps you should use to recover after basic auxiliary storage pool (ASP) has failed and all data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Basic ASP not in overflow status

Data loss:
All

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 21. Recovery checklist for disk failure–Checklist 12		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Physically remove the failed disk unit from the system.	

Table 21. Recovery checklist for disk failure–Checklist 12 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 2	Delete the data in the ASP that contains the failed unit.	
___ Task 3	Install the replacement disk unit.	
___ Task 4	Configure the replacement disk unit by selecting the "Replace configured unit" function on the Work with Disk Units display.	
Actions to be performed by the user		
___ Task 5	You must perform an initial program load (IPL) at this time. Follow the procedure for starting the system after it ends abnormally.	Chapter 6, "Starting the system after it ends abnormally," on page 159, task 1 through task 4.
___ Task 6	Restore user profiles: RSTUSRPRF USRPRF(*ALL) DEV(TAP01)	"Restoring user profiles" on page 211.
___ Task 7	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 8	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 9	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____	The System values topic in the IBM i documentation.
___ Task 10	After changing the system values, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 11	Recover the objects in the basic ASP.	"Recovering a damaged basic auxiliary storage pool" on page 186, task 1 through task 9.
___ Task 12	If you have user-defined file systems (UDFSs) in basic ASPs that do not restore correctly, you might need to perform additional recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.

Table 21. Recovery checklist for disk failure–Checklist 12 (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 13	Perform either a:
	`SIGNOFF *LIST`
	or a
	`DSPJOBLOG * *PRINT`
	Check the job log to ensure all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to tell you how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 14	Restore changed objects to the basic ASP. Apply journaled changes to objects in the basic ASP.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277, task 1 through task 7.
___ Task 15	Restore authority. Type: RSTAUT	"Restoring object authorities" on page 217.
___ Task 16	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 17	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 18	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.
___ Task 19	You must perform an initial program load (IPL) at this time.	"Performing a normal initial program load" on page 43.
___ Task 20	Review job logs or output from your restore operations to ensure that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 13: Actions for a failure in a basic auxiliary storage pool disk unit

This checklist shows the sequence of steps you should use to recover after basic auxiliary storage pool (ASP) has overflowed and all data has been lost.

This checklist should be used for the following problem situation:

Failed unit:
Basic ASP in overflow status

Data loss:
All

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 22. Recovery checklist for disk failure–Checklist 13

Task	What to do	Where to read more about it
Actions to be performed by the service representative		
___ **Task 1**	Physically remove the failed disk unit from the system.	
___ **Task 2**	Delete the data in the ASP that contains the failed unit.	
___ **Task 3**	Install the replacement disk unit.	
___ **Task 4**	Configure the replacement disk unit by selecting the "Replace configured unit" function on the Work with Disk Units display.	
Actions to be performed by the user		
___ **Task 5**	You must perform an initial program load (IPL) at this time. Follow the procedure for starting the system after it ends abnormally.	Chapter 6, “Starting the system after it ends abnormally,” on page 159.
___ **Task 6**	Reclaim storage.	“Reclaim Storage (RCLSTG) command” on page 28.
___ **Task 7**	Delete the overflowed objects.	“Deleting overflowed objects during recovery” on page 186.
___ **Task 8**	If necessary, change the QALWOBJRST system value. Write the old value here: _____	“Controlling restoration of security-sensitive objects” on page 34.
___ **Task 9**	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	“Controlling restoration of security-sensitive objects” on page 34.
___ **Task 10**	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.

Table 22. Recovery checklist for disk failure–Checklist 13 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 11	After changing the system values, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 12	Recover the objects in the basic ASP.	"Recovering a damaged basic auxiliary storage pool" on page 186, task 1 through task 9.
___ Task 13	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to perform additional recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 14	Restore changed objects to the basic ASP. Apply journaled changes to objects in the basic ASP.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277, task 1 through task 7.
___ Task 15	Restore authority. Type: RSTAUT	"Restoring object authorities" on page 217.
___ Task 16	If necessary, change the QALWOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 17	If necessary, change the QVFYOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 18	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.

___ Task 19	Perform either a:
	`SIGNOFF *LIST`
	or a
	`DSPJOBLOG * *PRINT`
	Check the job log to ensure all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to tell you how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 20	You must perform a normal IPL at this time.	"Performing a normal initial program load" on page 43.

Table 22. Recovery checklist for disk failure–Checklist 13 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 21	Review job logs or output from your restore operations to ensure that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 14: Actions for non-load source disk unit failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit with mirrored protection has failed.

This checklist should be used for the following problem situation:

Failed unit:
Any

Mirrored protection:
Yes

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these tasks if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Note: For many failures, the system does not have to stopped and started again. The service representative can repair the failed component while the system continues to run.

Table 23. Recovery checklist for disk failure–Checklist 14		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Replace the failed disk unit.	
___ Task 2	Resume mirrored protection.	
Actions to be performed by the user		
___ Task 3	Ensure that disk configuration is correct.	Chapter 24, "Working with mirrored protection," on page 431.

Related concepts

Mirrored protection recovery actions

On a system with mirrored protection, errors and failures have different effects. When a failure occurs on a system with mirrored protection, the recovery procedure is affected by the level of protection that is configured.

Checklist 15: Actions for non-load source disk unit failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit with device parity protection has failed.

This checklist should be used for the following problem situation:

Failed unit:
Any

Device parity protection:
Yes

Most steps in the checklist include references to other tasks in this topic. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Note: For many failures, the system does not have to stopped and started again. The service representative can repair the failed component while the system continues to run.

Table 24. Recovery Checklist for disk failure–Checklist 15

Task	What to do	Where to read more about it
Actions to be performed by the service representative		
___ Task 1	Attach the new disk unit.	
___ Task 2	Rebuild the failed device parity disk unit data.	
Actions to be performed by the user		
___ Task 3	Ensure that disk configuration is correct.	Chapter 23, “Working with device parity protection,” on page 419.

Related information
Working with device parity protection

Checklist 16: Actions for non-load source disk unit failure

This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit in the system auxiliary storage pool (ASP) has failed.

This checklist should be used for the following problem situation:

Failed unit:
Non-load source unit in system auxiliary storage pool (ASP)

Data loss:
N/A. Disk unit cannot be replaced immediately

Before you begin your recovery, make a copy of this checklist. Fill in the appropriate areas as you and the service representative perform the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 25. Recovery checklist for disk failure–Checklist 16

Task	What to do	Where to read more about it
Actions to be performed by the user		
___ Task 1	Remove the failed disk unit from the configuration.	“Removing a disk unit from an auxiliary storage pool” on page 403.

Table 25. Recovery checklist for disk failure–Checklist 16 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 2	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 3	Restore the Licensed Internal Code using option 1 (Restore Licensed Internal Code) from the latest Save System (SAVSYS) media.	"Preparing for loading the Licensed Internal Code" on page 122 and "Loading the Licensed Internal Code" on page 128
___ Task 4	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are performing a complete restore operation.	Chapter 5, "Restoring the operating system," on page 141, task 1 through task 6.
___ Task 5	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 6	If you restored the operating system using distribution media, some system information, such as access path recovery times and the system reply list, might have been reset to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 7	Use "Choosing the procedure to recover user information" on page 105 to determine the correct procedure for recovering user information.	

Checklist 17: Actions for independent auxiliary storage pool disk failure

The checklist shows the sequence of steps to use to recover after an independent auxiliary storage pool (ASP) failed, but no data was lost. Independent ASPs are used in a high availability environment.

This checklist can be used for the following problem situation:

Failed unit
Disk unit in independent ASP

Data loss
None

User ASP is configured
Yes

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative run the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects, if they do not apply in your situation.

Table 26. Recovering from a disk failure–Checklist 17

Task	What to do	Where to read more about it
The service representative runs the following actions		
___ Task 1	Save the disk unit data.	
___ Task 2	Attach a new disk unit.	
___ Task 3	Restore data to the new disk unit.	
The user runs the following actions		
___ Task 4	Vary on the independent ASP.	Use the Vary Configuration (VRYCFG) command or the IBM Navigator for i interface to vary on the independent ASP.

Checklist 18: Actions for a failure in an independent auxiliary storage pool disk unit

The checklist shows the sequence of steps that you can use to recover after an independent auxiliary storage pool (ASP) failed with some data loss.

This checklist can be used for the following problem situation:

Failed unit
Independent ASP

Data loss
Some

User ASP configured
Yes

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects, if they do not apply in your situation.

Table 27. Recovering from a disk failure - Checklist 18

Task	What to do	Where to read more about it
Actions run by the service representative		
___ Task 1	Save the disk unit.	
___ Task 2	Attach the new disk unit.	
___ Task 3	Restore the disk unit data.	

Table 27. Recovering from a disk failure - Checklist 18 (continued)		
Task	**What to do**	**Where to read more about it**
Actions run by the user		
___ Task 4	Vary on the independent ASP.	Use the Vary Configuration (VRYCFG) command or the IBM Navigator for i interface to vary on the independent ASP.
___ Task 5	Reclaim storage.	"Reclaim Storage (RCLSTG) command" on page 28.
___ Task 6	Evaluate the extent of the damage. Determine whether to attempt to recover damaged objects or restore the entire system. Do not skip this step.	"Task 4: Recovering from damaged objects and unreadable sectors" on page 165.
___ Task 7	If you decided to proceed, continue with the restore operation for independent ASP data.	"Recovering an independent auxiliary storage pool" on page 192.
___ Task 8	Restore changed objects and apply journaled changes to the independent ASP.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277.
___ Task 9	Restore authority. Type - RSTAUT	"Restoring object authorities" on page 217.

Checklist 19: Actions for a failure in an independent auxiliary storage pool disk unit

The checklist shows the sequence of steps that you use to recover after an independent auxiliary storage pool (ASP) failed and all data was lost.

This checklist is used for the following problem situation:

Failed unit
Independent ASP

Data loss
All

User ASP is configured
Yes

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative run the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects, if they do not apply in your situation.

Table 28. Recovering from a disk failure–Checklist 19		
Task	**What to do**	**Where to read more about it**
Actions run by the service representative		
___ Task 1	Physically remove the failed disk unit from the system.	

Table 28. Recovering from a disk failure–Checklist 19 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 2	Delete the data in the ASP that contains the failed unit.	
___ Task 3	Install the replacement disk unit.	
___ Task 4	Configure the replacement disk unit by selecting the "Replace configured unit" function on the Work with Disk Units display.	
Actions run by the user		
___ Task 5	Vary on the independent ASP.	Use the Vary Configuration (VRYCFG) command or the IBM Navigator for i interface to vary on the independent ASP.
___ Task 6	If necessary, change the QALWOBJRST system value. Write the old value here: _____	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 7	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 8	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____	System values.
___ Task 9	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 10	Recover the objects in the independent ASP.	“Recovering an independent auxiliary storage pool” on page 192.
___ Task 11	Restore changed objects and apply journaled changes to the independent ASP.	Chapter 13, “Restoring changed objects and apply journaled changes,” on page 277.
___ Task 12	Restore authority. Type - RSTAUT	“Restoring object authorities” on page 217.

Table 28. Recovering from a disk failure–Checklist 19 (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 13	Run on of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 14	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 15	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	“Controlling restoration of security-sensitive objects” on page 34.
___ Task 16	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values.
___ Task 17	Review job logs or output from your restore operations to verify that all objects were restored successfully.	“Verifying whether objects are restored successfully” on page 39.

Checklist 20: Recovering your entire system after a complete system loss

Use this checklist if you need to restore your entire system to a system that is running the same version of the IBM i licensed program. Do not use this checklist if you are doing an upgrade.

Note: If the system you must recover contains an independent auxiliary storage pool (ASP), see the rzarmrstiasp.dita topic.

Attention: If you use the Hardware Management Console (HMC), you must recover the HMC before you compete the following tasks in your recovery.

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative do the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the topics if you need more information about how to do a particular step. You might not need to complete some steps, such as restoring changed objects, if they do not apply in your situation.

Table 29. Recovery checklist for complete system loss Checklist 20		
Task	**What to do**	**Where to read more about it**
Actions the user must take		
___ Task 1	Starting in V5R4, you no longer need to manually reset the LAN Console service tools device ID password on the client. The client automatically attempts to resynchronize this password on all connections where a mismatch might occur. However, manually resetting the service tools device ID password on the system might still be needed depending on your restore data.	To resynchronize a previous release of Access Client Solutions, see Console Advanced Topics.
___ Task 2	If you are using LAN Console, disconnect from your system and close Access Client Solutions (ACS). Restart ACS and reconnect by using the user ID of QSECOFR and password of QSECOFR to sign on to your system.	
___ Task 3	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm.
___ Task 4	Prepare to load the Licensed Internal Code by using the most recent Save System (SAVSYS) media.	rzarmsrlic.dita.
___ Task 5	Install the Licensed Internal Code by using option 2 (Install Licensed Internal Code and Initialize System)[1].	rzarmilic2.dita.
___ Task 6	Configure the disk units (assign to ASP and set up disk protection). If you saved any user-defined file systems (UDFSs), you must configure your user ASPs or UDFSs might fail to restore. Specify whether you want to create encrypted or unencrypted user ASPs.	rzarmwrkdsk.dita and Creating a disk pool.

Table 29. Recovery checklist for complete system loss Checklist 20 (continued)

Task	What to do	Where to read more about it
___ Task 7	To allow security system values to change during your recovery, follow these steps: 1. On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)). Press Enter to continue. 2. Sign on to DST by using your DST service tools username and password. 3. Select option 13 (Work with system security). 4. Type 1 in the **Allow system value security changes** field and press Enter. 5. Press F3 or F12 to get back to the 'IPL or Install the System' display.	
___ Task 8	Restore the operating system, beginning with rzarmstrrst.dita. You are doing a complete restore operation.	rzarmrcpfs.dita.
___ Task 9	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	rzarmencryptrecoversw.dita and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.
___ Task 10	If you restored the operating system by using distribution media, some system information might be reset to default values, such as access path recovery times and the system reply list. Verify these values and correct them if necessary.	rzarmrcvsi.dita.
___ Task 11	Recover user information from your save media. Restore changed objects and apply journal entries. If you are restoring to a different system, or a different logical partition, you must specify `ALWOBJDIF(*ALL)` or `ALWOBJDIF(*COMPATIBLE)` on the **RST*xxx*** commands. **Note:** If you use Backup Recovery and Media Services (BRMS), refer to your BRMS recovery report to recover your user information.	rzarmrcvru.dita.

Table 29. Recovery checklist for complete system loss Checklist 20 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 12	If you are not sure what the password is for the QSECOFR profile that was restored from tape, change it before you sign off by running the following command. `CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)`	rzarmrstusrw.dita.
___ Task 13	If the operating system was restored from distribution media, restore your system information to the correct settings.	rzarmrcvsi.dita.

___ Task 14	Do either a
	`SIGNOFF *LIST`
	or a
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the error, and then restore those objects from the media.	

___ Task 15	Do a normal IPL now.	rzarmnormipl.dita.

___ Task 16	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by running the following commands.
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 17	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run the following command. `CALL QCCADEV/QCCAELOAD`	

Table 29. Recovery checklist for complete system loss Checklist 20 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 18	Review job logs or output from your restore operations to verify that all objects were restored successfully.	rzarmvfyrst.dita.

Related information

Upgrading or replacing IBM i and related software

Managing partitions (logical partitioning)

Backup and recovery for clusters

Checklist 21: Recovering your entire system after a complete system loss including independent auxiliary storage pools

Use this checklist if you need to restore your entire system that includes an independent ASP to a system that is running the same version of the IBM i licensed program. Do not use this checklist if you are doing an upgrade.

Note: If you are restoring a clustered system with independent ASPs, see the Scenario: Performing backups in geographic mirroring environment topic along with this checklist.

Attention: If you use the Hardware Management Console (HMC), you must recover the HMC before you do the following tasks in your recovery.

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative do the recovery steps. This checklist provides an important record of your recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the topics if you need more information about how to do a particular step. You might not need to complete some steps, such as restoring changed objects, if they do not apply in your situation.

Table 30. Recovering from a complete system loss Checklist 21		
Task	**What to do**	**Where to read more about it**
Actions the user must take		
___ Task 1	Starting in V5R4, you no longer need to manually reset the Operations Console (LAN) service tools device ID password on the client. The client automatically attempts to resynchronize this password on all connections where a mismatch might occur. However, manually resetting the service tools device ID password on the server might still be needed depending on the data you restored.	To resynchronize a previous release of Access Client Solutions, see Console Advanced Topics.
___ Task 2	If you are using LAN Console, disconnect from your server and close Access Client Solutions (ACS). Restart ACS and reconnect by using the user ID of QSECOFR and password of QSECOFR to sign on to your server.	

Table 30. Recovering from a complete system loss Checklist 21 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 3	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm.
___ Task 4	Prepare to load the Licensed Internal Code by using the most recent Save System (SAVSYS) media.	rzarmsrlic.dita
___ Task 5	Install the Licensed Internal Code by using option 2 (Install Licensed Internal Code and Initialize System).	rzarmilic2.dita
___ Task 6	Configure the disk units (assign to ASP and set up disk protection). If you saved any user-defined file systems (UDFSs), you must configure your user ASPs or the UDFSs might fail to restore. **Note:** Configure and restore independent ASPs in a later step.	rzarmwrkdsk.dita and rzarmwasp.dita.
___ Task 7	To allow security system values to be changed during your recovery, follow these steps: 1. On the 'IPL or Install the System' display, select 3 (Use Dedicated Service Tools (DST)). Press Enter to continue. 2. Sign on to DST by using your DST service tools username and password. 3. Select option 13 (Work with system security). 4. Type 1 in the **Allow system value security changes** field and press Enter. 5. Press F3 or F12 to get back to the IPL or Install the System display.	
___ Task 8	Restore the operating system, beginning with rzarmstrrst.dita. You are doing a complete restore operation.	rzarmrcpfs.dita.
___ Task 9	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	rzarmencryptrecoversw.dita and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.

Table 30. Recovering from a complete system loss Checklist 21 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 10	If you restored the operating system by using distribution media, some system information might be reset to default values, such as access path recovery times and the system reply list. Verify these values and correct them if necessary.	rzarmrcvsi.dita.
___ Task 11	Recover user information from your save media. Restore changed objects and apply journal entries. If you are restoring to a different system, or a different logical partition, you must specify `ALWOBJDIF(*ALL)` or `ALWOBJDIF(*COMPATIBLE)` on the RST*xxx* commands and SRM(*NONE) on the Restore Configuration (RSTCFG) command. **Note:** You might want to wait to restore authorities until after you configure the independent ASPs. You can also prompt on RSTAUT and specify `SAVASPDEV(*SYSBAS)` to allow users to access objects in the system and basic ASPs while you complete recovery procedures for the independent ASPs. Both options can eliminate restricting your server again. **Note:** If you use Backup Recovery and Media Services (BRMS), refer to your BRMS recovery report to recover your user information.	rzarmrcvru.dita.
___ Task 12	If you are not sure what the password is for the QSECOFR profile that was restored from tape, change it before you sign off. `CHGUSRPRF USRPRF(QSECOFR) PASSWORD(`*new-password*`)`	rzarmrstusrw.dita.
___ Task 13	If the operating system was restored from distribution media, restore your system information to the correct settings.	rzarmrcvsi.dita.

Table 30. Recovering from a complete system loss Checklist 21 (continued)

Task	What to do	Where to read more about it

___ Task 14	Do either a
	`SIGNOFF *LIST`
	or a
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the error, and then restore those objects from the media.	

___ Task 15	Configure the Service Tools Server for IBM i to access disk management functions.	The Configuring the service tools server IBM i documentation topic.
___ Task 16	Do a normal IPL now.	rzarmnormipl.dita.
___ Task 17	Choose one of the following actions to configure independent ASPs: • Use the 'Configure Device ASP (CFGDEVASP)' command. • Use IBM Navigator for i. **Important:** If you have the Encrypted ASP Enablement (5770-SS1, Option 45) licensed program is installed, specify whether you want to encrypt the independent ASP.	Creating an independent disk pool
___ Task 18	Verify the RESOURCE and make the independent ASP available now. A directory for the independent ASP is created and the default UDFS is automatically mounted over the directory.	Use the **Vary Configuration (VRYCFG)** command in the character-based interface, or the **Make Available** function in Navigator for i.
___ Task 19	Restore independent ASP data.	rzarmrasp3.dita and rzarmencryptrestoreasp.dita.
___ Task 20	Update program temporary fix (PTF) information for all PTF save files in library QGPL by running. `UPDPTFINF` .	rzarmrcvcj.dita.

Table 30. Recovering from a complete system loss Checklist 21 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 21	If the **Save System Information (SAVSYSINF)** command was used, then run the **Restore System Information (RSTSYSINF)** command. The **Restore System Information (RSTSYSINF)** command restores a subset of the system data and objects that are saved by the **Save System Information (SAVSYSINF)** command.	rzarmrst.dita
___ Task 22	Run the following command to restore authority: RSTAUT	rzarmrstaut.dita.

___ Task 23	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands.
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 24	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 25	Review job logs or output from your restore operations to verify that all objects were restored successfully.	rzarmvfyrst.dita.
___ Task 26	When you are ready to start clustering on the node you just recovered, you must start clustering from the active node. Clustering propagates the most current configuration information to the recovered node.	Starting nodes in the Implementing high availability IBM i documentation topic.

Related tasks

Recovering from an encrypted backup using software encryption
Backup Recovery and Media Services (BRMS) provides you with the ability to encrypt your data to a tape device. This method is called *software encryption*, because you do not need to use an encryption device. The BRMS interface asks for the encryption key information and what items you want encrypted. BRMS saves the key information, so for restoring, BRMS knows what key information is needed to decrypt on the restore.

Related information

Upgrading or replacing IBM i and related software
Managing partitions (logical partitioning)
Backup and recovery for clusters

Checklist 22: Restoring a logical partition to another logical partition

The checklist is used if you need to restore one logical partition (LPAR) to another logical partition.

If you have independent auxiliary storage pools (ASPs) configured in your logical partition, refer to the “Checklist 21: Recovering your entire system after a complete system loss including independent auxiliary storage pools” on page 97 topic as well.

Before you begin your recovery, make a copy of this checklist. Complete the appropriate areas as you and the service representative run the recovery steps. This checklist provides an important record of your

recovery actions. It can help you to diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects, if they do not apply in your situation.

Table 31. Recovering from a complete system loss -Checklist 22

Task	What to do	Where to read more about it
Actions taken by the user		
___ Task 1	If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation. The EKM contains the encryption keys that are needed for the recovery operation. Skip this step if you are not using an encrypting tape drive.	rzarmencryptrecovertape.htm
___ Task 2	Prepare to load the Licensed Internal Code[1] from the most recent Save System (SAVSYS) media.	"Preparing for loading the Licensed Internal Code" on page 122.
___ Task 3	Install the Licensed Internal Code by using option 3 (Install Licensed Internal Code and Recover Configuration).	"Loading the Licensed Internal Code" on page 128
___ Task 4	Configure the disk units (assign to ASP and set up disk protection). If you saved any user-defined file systems (UDFSs), you must configure your user ASPs or the UDFS cannot be restored. Specify whether you want to create encrypted or unencrypted user ASPs.	Chapter 21, "Configuring disks and disk protection," on page 369 and Creating a disk pool.
___ Task 5	Restore the operating system, beginning with "Task 1: Starting to restore the operating system" on page 143. You are running a complete restore operation.	"Tasks to restoring the operating system" on page 143.
___ Task 6	If you are restoring from an encrypted backup, the save/restore master key on the target system must match the save/restore key on the source system. If the save/restore master key does not match, set the target system's save/restore master key to recover all the master keys.	"Recovering from an encrypted backup using software encryption" on page 307 and Loading and setting save/restore master key in the Cryptography IBM i documentation topic.

Table 31. Recovering from a complete system loss -Checklist 22 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 7	If you restored the operating system from distribution media, some information, such as access path recovery times and the system reply list, might be set to default values. Verify these values and correct them if necessary.	"Recovering system information" on page 209.
___ Task 8	Recover user information from your save media. Restore changed objects and apply journal entries. If you are restoring to a different system, or a different logical partition, you must specify `ALWOBJDIF(*ALL)` on the *RSTxxx* commands.	"Choosing the procedure to recover user information" on page 105.
___ Task 9	If you are not sure what the password is for the QSECOFR profile that was restored from tape, change it before you sign off - `CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)`	"What happens when you restore user profiles" on page 213.
___ Task 10	If distribution media was used for the restore, set your system information to the correct settings.	"Recovering system information" on page 209.

___ Task 11	Run on of the following command -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 indicates that a library was successfully restored. Message CPF3773 indicates how many objects were not restored for a library. Objects are not restored for various reasons. Check for any error messages, correct the error, and then restore those objects from the media.	

___ Task 12	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

Table 31. Recovering from a complete system loss -Checklist 22 (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 13	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 14	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3 installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 15	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 23: Actions for a failed cache card

Use this checklist to recover storage after the cache card has failed.

This checklist should be used for the following problem situation:

Failed unit:
Cache storage in input/output processor (IOP)

Data loss:
Some

Most steps in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular step. You might not need to perform some steps, such as restoring changed objects, if they do not apply in your situation.

Table 32. Recovery checklist for failed cache card–Checklist 23		
Task	**What to do**	**Where to read more about it**
Actions to be performed by the service representative		
___ Task 1	Reclaim IOP cache storage.	
Actions to be performed by the user		

Table 32. Recovery checklist for failed cache card–Checklist 23 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 2	Evaluate the extent of the damage. It might be fastest to recover the entire ASP. Choose from the following actions to identify damaged objects in critical files where you suspect damage: • Use the Copy File (CPYF) command. • Use the Display Object Description (DSPOBJD) command. • Perform a save of your critical data. Choose from the following actions to identify damaged objects at the system level: • Use the Retrieve Disk Information (RTVDSKINF) command. • Use the Display Object Description (DSPOBJD) command and specify DSPOBJD OBJ(*ALL) OBJTYPE(*ALL). • Perform a full system save using GO SAVE option 21. • Run the Reclaim Storage (RCLSTG) procedure. Running the procedure is described in "Reclaim Storage (RCLSTG) command" on page 28.	
___ Task 3	If you have decided to do a complete restore operation, use Table 33 on page 106 to determine the correct procedure for recovering user information.	
___ Task 4	If you have decided to attempt to recover damaged objects, perform the tasks in "Task 4: Recovering from damaged objects and unreadable sectors" on page 165.	

Choosing the procedure to recover user information

When your system is running normally, you are ready to recover user information.

Your first step in a recovery is to return your system to a normal operating condition, which might require the following tasks:

- Replace hardware
- Restore or install the Licensed Internal Code
- Perform an initial program load (IPL) after the system ends abnormally

After your system is running normally, you can recover user information. Use Table 33 on page 106 to determine the procedure you should follow. In the table, N/A in a column means that the recovery procedure is the same, whether you respond yes or no.

Table 33. Choosing the correct recovery procedure for user information

Are you recovering all *SYSBAS ASPs?	Save procedure used	Do you have SAVCHGOBJs or journals to apply?	Do you want to use menu options to recover?	Recovery procedure to follow
Yes	Commands	N/A	See note 1.	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 21	No	Yes	"Checklist 25: Using option 21 from the restore menu" on page 110
Yes	Save menu option 21	Yes	N/A	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 21	No	No	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 22 Save menu option 23	No	Yes	"Checklist 26: Using options 22 and 23 from the restore menu" on page 113
Yes	Save menu option 22 Save menu option 23	Yes	N/A	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 22 Save menu option 23	No	No	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 21 Save menu option 23	No	Yes	"Checklist 26: Using options 22 and 23 from the restore menu" on page 113
Yes	Save menu option 21 Save menu option 23	Yes	N/A	"Checklist 24: Recovering user information using commands" on page 107
Yes	Save menu option 21 Save menu option 23	No	No	"Checklist 24: Recovering user information using commands" on page 107
Yes	Operational Assistant bckup[2]	N/A	N/A	"Checklist 27: Recovering user information using tapes from Operational Assistant backup" on page 116
No	Any	N/A	N/A	"Checklist 24: Recovering user information using commands" on page 107

Table 33. Choosing the correct recovery procedure for user information (continued)

Are you recovering all *SYSBAS ASPs?	Save procedure used	Do you have SAVCHGOBJs or journals to apply?	Do you want to use menu options to recover?	Recovery procedure to follow

1

If you save using commands rather than menu options, you should recover using commands.

2

You have saved using either the Run Backup (RUNBCKUP) command or the Run Backup menu.

Related tasks

Using restore menu options 21, 22, and 23
To restore information using option 21, 22, or 23 from the Restore menu, follow these steps.

Checklist 24: Recovering user information using commands

The checklist shows the sequence of steps to use to recover user information with commands. You might need to run some tasks more than one time.

The correct steps for your situation depend on the following conditions:

- How you saved your information?
- Do you use journaling or whether applications you purchased use journaling?
- Do you have document library objects?
- Do you save changed objects?

Before you begin recovering user information, make a copy of this checklist. Complete the appropriate areas as you run the recovery steps. This checklist provides an important record of your recovery actions. It can help you diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects and applying journal changes, if they do not apply in your situation.

Restoring to a different system or different logical partition?

- Specify `ALWOBJDIF(*ALL)` or `ALWOBJDIF(*COMPATIBLE)` on the RST*xxx* commands.
- Specify `SRM(*NONE)` on the Restore Configuration (RSTCFG) command.
- Network attributes can be reset to the IBM-supplied defaults.

Table 34. Recovering user information by using commands

Task	What to do	Where to read more about it
___ Task 1	If your system is not already in a restricted state, verify that all users are off the system and that all jobs are ended. Then type `ENDSBS SBS(*ALL) OPTION(*CNTRLD) DELAY(600)`[1,2].	"Putting your system in a restricted state" on page 27.
___ Task 2	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.

Table 34. Recovering user information by using commands (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 3	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 4	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 5	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 6	If you are restoring to a system with a different processor or memory, verify that the QMCHPOOL, QBASPOOL, and QPFRADJ system values are correct by using the WRKSYSVAL command.	"Setting the QPFRADJ system value for a different processor or memory" on page 38.
___ Task 7	Prevent messages that are not related to the recovery from interrupting by typing the following command: `CHGMSGQ MSGQ(QSYSOPR) DLVRY(*NOTIFY) SEV(99)`	
Determine what ENDOPT value to use - When you are restoring from tape, you tell the system whether to rewind the tape. If you are using tape in the tasks that follow, specify `ENDOPT(*LEAVE)` when you have more steps. Specify `ENDOPT(*REWIND)` for your last step.		
___ Task 8	Restore user profiles: `RSTUSRPRF DEV(TAP01) USRPRF(*ALL)`	"Restoring user profiles" on page 211.
___ Task 9	Restore device configuration: `RSTCFG OBJ(*ALL) OBJTYPE(*ALL) DEV(TAP01)`	"Restoring configuration objects" on page 227.
___ Task 10	Restore the libraries to each auxiliary storage pool (ASP) that you are recovering. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Restoring libraries" on page 231.
___ Task 11	Restore the ownership for document library objects (DLOs) in the user ASPs you are restoring.	"Task 8: Reclaiming document library objects" on page 179.
___ Task 12	Restore your last complete save of document library objects to each user ASP you are recovering.	"Restoring documents and folders" on page 260.

Table 34. Recovering user information by using commands (continued)

Task	What to do	Where to read more about it
___ Task 13	Restore your last complete save of directories.[3]	"Restoring objects in directories" on page 265.
___ Task 14	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to do more recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 15	Restore changed objects and apply journaled changes.	Chapter 13, "Restoring changed objects and apply journaled changes," on page 277, task 1 through task 7.
___ Task 16	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	
___ Task 17	If the Save System Information (SAVSYSINF) was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved by using the Save System Information (SAVSYSINF) command. **Note:** If you have independent ASPs, do not run this step now.	Chapter 12, "Restoring system information," on page 275
___ Task 18	Restore authority. Type - RSTAUT **Note:** If you are restoring independent ASPs, you can prompt on RSTAUT and specify SAVASPDEV(*SYSBAS). SAVASPDEV(*SYSBAS) allows users access to objects in the system and basic ASPs, while you complete recovery procedures for the independent ASPs.	"Restoring object authorities" on page 217.
___ Task 19	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 20	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 21	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 22	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values
___ Task 23	If you are recovering from a complete system loss, return to the appropriate checklist. Continue with the tasks on that checklist.	

Table 34. Recovering user information by using commands (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 24	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 25	Do a normal initial program load (IPL) now.	“Performing a normal initial program load” on page 43.

___ Task 26	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 27	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 28	Review job logs or output from your restore operations to verify that all objects were restored successfully.	“Verifying whether objects are restored successfully” on page 39.

1 Your system must be in a restricted state to restore user profiles. Other steps in the recovery might not require a restricted state. However, to verify the success of your recovery and better performance when you are restoring information, a restricted state is suggested.

2 For the delay parameter, specify the number of seconds to allow your system time to bring most jobs to a normal end. On a large, busy system, you might need a longer delay.

Checklist 25: Using option 21 from the restore menu

The checklist shows the sequence of steps to run to recover user information by using option 21 from the Restore menu. Option 21 restores your system to your last complete save.

Before you begin recovering user information, make a copy of this checklist. Complete the appropriate areas as you run the recovery steps. This checklist provides an important record of your recovery actions.

It can help you diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects and applying journal changes, if they do not apply in your situation.

Restoring to a different system?

If you are restoring to a different system or to a different logical partition, be aware of the following considerations -

- Specify `ALWOBJDIF(*ALL)` or `ALWOBJDIF(*COMPATIBLE)` on the restore commands.
- Specify `SRM(*NONE)` on the RSTCFG command.
- Network attributes can be reset to the IBM-supplied defaults.

Note: An option is available on the restore menu that indicates that you are restoring to a different system. If you selected this option, the system automatically specifies the first two items that are previously listed. Also specify this option if you are restoring to a different logical partition.

Table 35. Recovering user information with option 21

Task	What to do	Where to read more about it
___ Task 1	If necessary, change the QALWOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 2	If necessary, change the QVFYOBJRST system value. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 3	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	System values
___ Task 4	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 5	If you are restoring to a system with a different processor or memory, verify that the QMCHPOOL, QBASPOOL, and QPFRADJ system values are correct by using the WRKSYSVAL command.	"Setting the QPFRADJ system value for a different processor or memory" on page 38.

Table 35. Recovering user information with option 21 (continued)

Task	What to do	Where to read more about it
___ Task 6	Select option 21 from the Restore menu. Use your most recent tapes from running option 21 on the Save menu. If you are recovering by using "Checklist 20: Recovering your entire system after a complete system loss" on page 93 and restoring to a different system, use the "Restore to different system" option on the Specify Command Defaults display. Also use this option if you are restoring to a different logical partition. This option automatically specifies ALWOBJDIF(*ALL) or ALWOBJDIF(*COMPATIBLE) on the RSTxxx commands and SRM(*NONE) on the RSTCFG command.[1]	"Using restore menu options 21, 22, and 23" on page 203.
___ Task 7	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to do other recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 8	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	
___ Task 9	If the Save System Information (SAVSYSINF) was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved with the Save System Information (SAVSYSINF) command.	Chapter 12, "Restoring system information," on page 275
___ Task 10	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 11	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 12	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 13	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	The System values topic in the IBM i documentation.
___ Task 14	If you are recovering from a complete system loss, return to "Checklist 20: Recovering your entire system after a complete system loss" on page 93. Refer to the appropriate task in the checklist.	
___ Task 15	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

Table 35. Recovering user information with option 21 (continued)		
Task	**What to do**	**Where to read more about it**

___ Task 16	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 17	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 18	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 26: Using options 22 and 23 from the restore menu

The checklist shows the sequence of steps to use to recover user information by using option 22 and 23 from the restore menu. Option 22 restores your IBM-supplied libraries to your last save. Option 23 restores your user libraries to your last save.

Before you begin recovering user information, make a copy of this checklist. Complete the appropriate areas as you run the recovery steps. This checklist provides an important record of your recovery actions. It can help you diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects and applying journal changes, if they do not apply in your situation.

Table 36. Recovering user information by using options 22 and 23		
Task	**What to do**	**Where to read more about it**
___ Task 1	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 2	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 3	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.

Table 36. Recovering user information by using options 22 and 23 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 4	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 5	If you are restoring to a system with a different processor or memory, verify that the QMCHPOOL, QBASPOOL, and QPFRADJ system values are correct by using the WRKSYSVAL command.	"Setting the QPFRADJ system value for a different processor or memory" on page 38.
___ Task 6	Select option 22 from the **Restore** menu to restore IBM-supplied libraries and directories. Use your most recent tapes, which were created by running either option 21 or option 22 on the **Save** menu. If you are recovering by using "Checklist 20: Recovering your entire system after a complete system loss" on page 93 and restoring to a different system, use the "Restore to different system" option on the Specify Command Defaults display. Also use this option if you are restoring to a different logical partition. This option automatically specifies ALWOBJDIF(*ALL) or ALWOBJDIF(*COMPATIBLE) on the restore commands and SRM(*NONE) on the RSTCFG command.[1]	"Using restore menu options 21, 22, and 23" on page 203.
___ Task 7	Select option 23 from the **Restore** menu to restore user libraries and user directories. Use your most recent tapes by running either option 21 or option 23 on the **Save** menu. If you are recovering by using "Checklist 20: Recovering your entire system after a complete system loss" on page 93 and restoring to a different system, use the "Restore to different system" option on the Specify Command Defaults display. Also use this option if you are restoring to a different logical partition. This option automatically specifies ALWOBJDIF(*ALL) or ALWOBJDIF(*COMPATIBLE) on the restore commands and SRM(*NONE) on the RSTCFG command.[1]	"Using restore menu options 21, 22, and 23" on page 203.
___ Task 8	If you have user-defined file systems (UDFSs) in user auxiliary storage pools (ASPs) that do not restore correctly, you might need to run more recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 9	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: UPDPTFINF.	

Table 36. Recovering user information by using options 22 and 23 (continued)

Task	What to do	Where to read more about it
___ Task 10	If the Save System Information (SAVSYSINF) command was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved by using the Save System Information (SAVSYSINF) command.	Chapter 12, "Restoring system information," on page 275
___ Task 11	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 12	If necessary, change the QALWOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 13	If necessary, change the QVFYOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 14	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values

___ Task 15	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 16	Do a normal initial program load (IPL) now.	"Performing a normal initial program load" on page 43.

___ Task 17	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

Table 36. Recovering user information by using options 22 and 23 (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 18	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 19	Review job logs or output from your restore operations to verify that all objects were restored successfully.	"Verifying whether objects are restored successfully" on page 39.

Checklist 27: Recovering user information using tapes from Operational Assistant backup

The checklist shows the sequence of steps to use to recover user information from an Operational Assistant backup.

These procedures assume that your entire backup was done by using Operational Assistant and you did not mix Operational Assistant backup with other save methods.

Before you begin recovering user information, make a copy of this checklist. Complete the appropriate areas as you run the recovery steps. This checklist provides an important record of your recovery actions. It can help you diagnose any problems that occur after the recovery. It might also be useful in evaluating your backup strategy.

Most steps in the checklist include references to other topics in this document. Refer to the following topics if you need more information about how to run a particular step. You might not need to run some steps, such as restoring changed objects and applying journal changes, if they do not apply in your situation.

Restoring to a different system?

If you are restoring to a different system or to a different logical partition, be aware of the following considerations -

- Specify `ALWOBJDIF(*ALL)` or `ALWOBJDIF(*COMPATIBLE)` on the RST*xxx* commands.
- Specify `SRM(*NONE)` on the Restore Configuration (RSTCFG) command.
- Network attributes are reset to the IBM-supplied defaults.

Table 37. Recovering user information from Operational Assistant backup tapes		
Task	**What to do**	**Where to read more about it**
___ Task 1	If your system is operational and the QUSRSYS library is on the system, print the Backup Status and the Backup History by typing: `DSPBCKSTS OUTPUT(*PRINT)`.	
___ Task 2	If your system is operational and the QUSRSYS library is on the system, print the Backup List by typing: `DSPBCKUPL OUTPUT(*PRINT)`.	
___ Task 3	If your system is not already in a restricted state, verify that all users are off the system. Then type `ENDSBS SBS(*ALL) OPTION(*CNTRLD) DELAY(600)`[1,2].	"Putting your system in a restricted state" on page 27.

Table 37. Recovering user information from Operational Assistant backup tapes (continued)

Task	What to do	Where to read more about it
___ Task 4	If necessary, change the QVFYOBJRST system value back to its original value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 5	If necessary, change the QALWOBJRST system value back to its original value by using the WRKSYSVAL command. Write the old value here: _____	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 6	If necessary, change the system value that controls whether the job log wraps when it is full. Use the Work with System Values command: WRKSYSVAL QJOBMSGQFL. Write down the current value here: _____. Then change the value to *PRTWRAP.	The System values topic in the IBM i documentation.
___ Task 7	After the system values are changed, sign off by using the command SIGNOFF *LIST. Then, using a newly created password, sign back on as QSECOFR for the new values to take effect.	
___ Task 8	If you are restoring to a system with a different processor or memory, verify that the QMCHPOOL, QBASPOOL, and QPFRADJ system values are correct by using the WRKSYSVAL command.	"Setting the QPFRADJ system value for a different processor or memory" on page 38.
___ Task 9	Prevent messages that are not related to the recovery from interrupting by typing the following command: `CHGMSGQ MSGQ(QSYSOPR) DLVRY(*NOTIFY) SEV(99)`	
___ Task 10	Restore user profiles: `RSTUSRPRF DEV(TAP01) USRPRF(*ALL)`.	"Restoring user profiles" on page 211.
___ Task 11	Restore device configuration: `RSTCFG OBJ(*ALL) OBJTYPE(*ALL) DEV(TAP01)`	"Restoring configuration objects" on page 227.
___ Task 12	Restore the libraries to each auxiliary storage pool (ASP) that you are recovering. If you choose not to restore all of your libraries now, verify that you restore the QGPL and QUSRSYS libraries along with the libraries that you are restoring.	"Restoring your libraries" on page 304
___ Task 13	Restore the ownership for document library objects (DLOs) in the user ASPs that you are restoring.	"Task 8: Reclaiming document library objects" on page 179.
___ Task 14	Restore your last complete save of document library objects to each user ASP that you are recovering.	"Restoring documents and folders" on page 260.

Table 37. Recovering user information from Operational Assistant backup tapes (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 15	Restore your last complete save of directories[3].	"Restoring objects in directories" on page 265.
___ Task 16	Restore incremental backups of libraries.	"Restoring libraries that you saved by using a backup list" on page 305.
___ Task 17	Restore changed objects.	"Restoring changed objects that you saved by using Operational Assistant" on page 305.
___ Task 18	Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing - UPDPTFINF.	
___ Task 19	If the Save System Information (SAVSYSINF) command was used, then run the Restore System Information (RSTSYSINF) command. The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects that are saved by using the Save System Information (SAVSYSINF) command. **Note:** If you have independent ASPs, do not run this step now.	Chapter 12, "Restoring system information," on page 275
___ Task 20	Restore authority. Type - RSTAUT	"Restoring object authorities" on page 217.
___ Task 21	Reapply any PTFs that were applied since your last SAVSYS operation.	Chapter 11, "Restoring program temporary fixes (PTFs)," on page 273.
___ Task 22	If you have user-defined file systems (UDFSs) in user ASPs that do not restore correctly, you might need to run more recovery steps.	"Task 7: Restoring user-defined file systems to the user auxiliary storage pool" on page 178.
___ Task 23	If necessary, change the QALWOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 24	If necessary, change the QVFYOBJRST system value.	"Controlling restoration of security-sensitive objects" on page 34.
___ Task 25	If necessary, change the QJOBMSGQFL system value back to its original value by using the WRKSYSVAL command.	System values

Table 37. Recovering user information from Operational Assistant backup tapes (continued)

Task	What to do	Where to read more about it

___ Task 26	Run one of the following commands -
	`SIGNOFF *LIST`
	`DSPJOBLOG * *PRINT`
	Check the job log to verify that all objects were restored. The job log contains information about the restore operation. To verify that all objects were restored, spool the job log for printing along with the job's remaining spooled output, if any.
	Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to indicate how many objects were restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.	

___ Task 27	Do a normal initial program load (IPL) now.	“Performing a normal initial program load” on page 43.

___ Task 28	If IBM Content Manager OnDemand for i is installed, complete journaling for Content Manager OnDemand by typing the following commands -
	`CALL QRDARS/QRLCSTRJ PARM('RLC')`
	`CALL QRDARS/QRLCSTRJ PARM('RLR')`	

___ Task 29	If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3, installed, run this command: `CALL QCCADEV/QCCAELOAD`	
___ Task 30	Review job logs or output from your restore operations to verify that all objects were restored successfully.	“Verifying whether objects are restored successfully” on page 39.

1 Your system must be in a restricted state to restore user profiles. Other steps in the recovery might not require a restricted state. However, to verify the success of your recovery and better performance when you are restoring information, a restricted state is suggested.

2 For the delay parameter, specify the number of seconds to wait before the system ends so most jobs can be ended normally. On a large, busy system, you might need a longer delay.

Related concepts

Restoring your system using Operational Assistant tapes
According to the different situations, perform one of these three tasks to restore information from Operational Assistant backup tapes.

Chapter 4. Recovering the Licensed Internal Code

Licensed Internal Code is the layer of IBM i architecture just above the hardware.

You must have the Licensed Internal Code on your system before you can restore the operating system. You must use the control panel to start the recovery of the Licensed Internal Code. The *control panel* can be one of the following:

- Physical control panel.
- IBM i remote control panel.
- IBM i virtual control panel.
- System partitions using system service tools (SST) and dedicated service tools (DST).
- Hardware Management Console (HMC) if the system is managed by the HMC.

The Install Licensed Internal Code (LIC) menu provides several methods for loading the Licensed Internal Code to your system. Table 38 on page 121 describes the options and how they are used:

Table 38. Options from the Install the Licensed Internal Code (LIC) Menu

Option number	Description	Purpose
1	Restore	Restores the Licensed Internal Code without removing other information that is on the system. Option 1 is normally used in the following situations: • You are encountering problems with the operating system, such as damaged objects. You sometimes need to restore the Licensed Internal Code before restoring the operating system. • The software support center suggests it. • You have replaced a failed disk unit other than unit 1 in the system auxiliary storage pool (ASP). • You are updating your system to a new release. See Installing, upgrading, or deleting IBM i and related software for the procedures to install a new release of the IBM i operating system.
2	Install the Licensed Internal Code and Initialize system	Installs the Licensed Internal Code and removes all data from all disk units. Option 2 is normally used in the following situations: • You are restoring to another system to recover from a complete system loss. • You are recovering with SAVSYS media that is at a previous release than what is currently installed on the system. If a system is configured to use LAN Console, and that system undergoes a backup and recovery cycle, you might have to perform the following steps: 1. Perform an initial program load (IPL) in Manual mode. 2. Use dedicated service tools (DST) to reconfigure the system so that it will detect the PC console when you perform an IPL in Normal mode.

© Copyright IBM Corp. 1997, 2025

Table 38. Options from the Install the Licensed Internal Code (LIC) Menu (continued)

Option number	Description	Purpose
3	Install Licensed Internal Code and Recover Configuration	Installs the Licensed Internal Code and prompts you to begin the procedure to recover information about how the disks were configured on your system (including ASP assignments and protection). Option 3 is normally used in the following situations: • You have replaced the load source unit. • The software support center suggests it.
4	Install Licensed Internal Code and Restore Disk Unit Data	Installs the Licensed Internal Code and restores data to a replacement disk unit. This option is used only by a load source unit after data was successfully saved (pumped) from a failed load source disk unit.

The recovery checklists in Chapter 3, "Selecting the appropriate recovery strategy," on page 53 specify which procedures in this chapter are required for your situation.

Attention: Make sure you use the correct procedure for your situation. Some of the procedures in this chapter will remove all data from your system.

Related concepts

The relationship between save and restore commands
You can learn which store commands can be used, based on how the objects were saved.

Related information

Console Advanced Topics
Accessing service tools
Managing the control panel functions

Preparing for loading the Licensed Internal Code

Perform these tasks to start to load the Licensed Internal Code.

- Find the right media and documentation.
- Stop your system, if it is running.
- Perform an initial program load (IPL) from an alternate device, either tape or optical media.

Check off each step in these tasks as you complete it.

Task 1: Getting ready to load the Licensed Internal Code

You can learn the things that you need to get ready to load the Licensed Internal Code.

Find these things before you begin loading Licensed Internal Code:

- Your most recent SAVSYS media. One of the following methods creates a SAVSYS media image:
  - Running the Save System (SAVSYS) command.
  - Using option 21 from the **Save** menu.
  - Using option 22 from the **Save** menu.
  - Using option 11 from the **Run Backup** menu.

Attention: Do not use a media volume that you created through Dedicated Service Tools (DST) by using option 5=Save Licensed Internal Code, unless you were instructed to do so by Software Services. A media volume that is created by using this process does not contain the Licensed Internal Code Program Temporary Fix (PTF) inventory information, cryptographic

services master keys, or the IBM i operating system. If you do the recovery by using this media volume, you might not be able to IPL, use your system, or load PTFs. In this case, you need to reinstall the Licensed Internal Code from either a SAVSYS media volume or from your distribution media.

- The manual for your alternate IPL device. An *alternate IPL device*, also known as the *restart device*, is a tape or optical device that you can use to start your system. The manual describes how to operate your device and system reference code (SRC) codes that you might see.
- If you enabled your device as an alternate installation device, you need the Licensed Internal Code DVD. (An *alternate installation device* is an alternate IPL device that is connected to a bus other than the system bus (bus 1).)
- If you do not have current SAVSYS media or they are damaged, you need the following items:
  - The distribution media (optical media or tape) that is supplied by IBM.
  - All the optical media for PTFs that you previously applied. Use the distribution media only if you do not have SAVSYS media. If you use the distribution media to restore the Licensed Internal Code, some of your system information is lost, such as the PTFs you were applied.
- The list of all the PTFs applied to your system at the time you saved the entire system. This list is attached to your backup log or with your SAVSYS media.

Note: If you saved to virtual media, you also saved your Licensed Internal Code and operating system to physical media from the virtual images. Physical media is needed to begin the recovery process. You must also have access to all of your user data on physical media.

Before you begin loading the Licensed Internal Code.

- Clean the read and write heads of the tape unit if you are using a tape device.
- If your source system (the system that was saved and needs recovery) is operational, print a list of all the PTFs currently on the system. Type the following and press the Enter key:

```
DSPPTF LICPGM(*ALL) OUTPUT(*PRINT)
```

- If you are using LAN Console, disconnect from your system and close the Access Client Solutions (ACS) application. Restart ACS and reconnect by using the user ID of QSECOFR and password of QSECOFR to sign on to your system.

Related concepts

Using an alternate installation device
You can learn how to set up, disable, verify, and select an alternate installation device.

Task 2: Powering down the system

You need to power down the system that has activities.

If your system is already powered off or if you are recovering to a system at an IBM Business Recovery Services Center, skip this task and begin with "Task 3: Preparing to start the system from an alternate IPL device" on page 124. It is not necessary to restrict a system that has no activity on the system.

1. Sign on the system as QSECOFR.
2. Change the QSYSOPR message queue:

```
CHGMSGQ MSGQ(QSYSOPR) DLVRY(*BREAK)
        SEV(60)
```

3. Bring your system to a restricted state:

```
ENDSBS SBS(*ALL) OPTION(*CNTRLD)
       DELAY(600)
```

Note: Specify a number of seconds for the delay parameter that is long enough for your system to bring most jobs to a normal end. On a large, busy system, you might need more time.

The system displays a message that subsystem ending is in progress. The system displays another message when all subsystems have ended and the system is in a restricted state. After the subsystems have ended, continue with the next step.

4. Power down the system:

```
PWRDWNSYS OPTION(*IMMED)
```

When the Power On light goes off or the partition on the HMC indicated **Not Activated**, continue with the next task.

Task 3: Preparing to start the system from an alternate IPL device

To start the system from an alternate IPL device, set your control panel for IPL type D and manual mode. You may also need to specify which I/O processor to use. The steps vary based on the type of system and control panel.

If you are using logical partitions with Hardware Management Console (HMC), perform these steps on the console of the logical partition on which you want to perform an IPL.

1. In your HMC, expand **Systems Management** > **Servers**.
2. Select the system.
3. Select the check box for the partition.
4. From the menu for the partition, click **Properties**.
5. In the Partition Properties window, click the **Settings** tab.
6. Under the Boot category, change the **IPL source** to D and the **Keylock position** to Manual.
7. Click **OK**.
8. From the menu for the partition, click **Configuration** > **Manage Profiles**.
9. In the Managed Profiles pane, click the profile name.
10. In the Logical Partition Profile Properties pane, click the **Tagged I/O** tab.
11. From the **Alternate restart device** list, select the slot with the I/O processor that is attached to the restart device.
12. Click **OK**.
13. In the Managed Profiles pane, click **Close**.
14. If the partition is running, from the menu for the partition, click **Operations** > **Shut Down**.

    Wait for the partition to reach Not Activated status.
15. From the menu for the partition, click **Operations** > **Activate**.

Related information

Basic system operations

Task 4: Loading the Licensed Internal Code from media

You need to load the Licensed Internal Code from media to complete the preparation.

1. Find the Licensed Internal Code tape or optical media.

   It is the first volume of the most current set of SAVSYS media or the first volume of the distribution optical media.

   Notes:

   - Do not use save media that was created through dedicated service tools (DST) by using option 5=Save Licensed Internal Code from the initial program load (IPL) or Install the System menu unless you have been instructed to do so by Software Services. The save media that you create

through this process does not contain the Licensed Internal Code program temporary fix (PTF) inventory information or the IBM i operating system. If you perform the recovery process using this save media, you will have to reinstall the Licensed Internal Code from either a SAVSYS media or from your distribution media before you can load any PTFs onto the system.

- Use the distribution media only if no SAVSYS media volume exists. If you use the distribution media, some system information will be lost. This includes, but is not limited to, PTFs and PTF packages. If you use the distribution media, you must reinstall all cumulative PTF packages and individual PTFs applied after the initial installation of your system.

2. Place the media volume in your alternate IPL device, or place the optical media in the optical disk unit.

   When you start the IPL, the system searches the alternate IPL devices for the correct media on the system bus (bus 1). For more information about loading the tape or optical media, see the setup manual for the device.

   Notes:

   - If you cannot load your alternate IPL device when the power is off, continue with the next step. The system prompts you later with a system reference code (SRC) code for the tape device or optical device.
   - If you use a tape device that you enabled as an alternate installation device, you must load both the Licensed Internal Code DVD and your tape media. (An alternate installation device is an alternate IPL device that is connected to a bus other than the system bus (bus 1).)

3. Turn on the power to the system.
4. If you could not load your media volume in step 2, load the first media volume into the device that you use for IPL. Make the device ready and then continue with the next step.
5. Ensure that the tape device or optical device is online or ready.

   No action is required for devices that perform this step automatically (such as the tape cartridge unit).

6. Ensure that the console display is turned on.

   After a delay, you should see the Install Licensed Internal Code menu. The length of the delay varies, depending on your system configuration and the speed of your alternate IPL device. The delay is typically between 5 minutes and 30 minutes. When you see this menu, continue with step 7.

   If the system attention light appears and one of the SRC codes that are shown in Table 39 on page 125 is displayed in the Data display, complete the instructions for that SRC code.

   Note: If you are using logical partitions, the SRC codes will be shown on the HMC under the operator panel value.

Table 39. SRC codes when loading the Licensed Internal Code

SRC code	Why it appears	What you do
A1xx 1933 A12x 1933 ('x' is any character)	The device for the alternate IPL is not ready.	Make sure that you loaded the correct media volume. Make the device unit ready. Wait for the System Attention light to go off. Then continue with the next step. If the System Attention light stays on for more than 5 minutes, check to see if you have the correct tape loaded in the device for the alternate IPL and make the device ready. Then continue with the next step.

Table 39. SRC codes when loading the Licensed Internal Code (continued)		
SRC code	**Why it appears**	**What you do**
B1xx 1803 **B1xx 1806** **B1xx 1938**	The device for the alternate IPL was not found or was not ready.	Make sure that you powered on the device, that you loaded the correct media volume, and the media volume is ready. Then continue with the next step.
B1xx 1934	The wrong media volume is loaded. Or the high-speed feature is enabled on the 2440 Tape Unit.	Load the correct media volume and make the device ready. Then continue with the next step or disable the high-speed feature on the 2440 Tape Unit.
B608 1105	This SRC occurs when you exit from the automatic installation because an alternate installation device attached to the system is enabled but is otherwise not ready. The device might not be ready because the media volume is not loaded or the device might not be enabled as an alternate installation device. Or, there might be a problem with the device.	Power down the system. If necessary, fix the device. Power on the system. Verify that the alternate installation device is enabled. Load the media volume in the alternate installation device and start the installation procedure again.
2507 0001 **2642 0001** **2643 0001**	A media volume is not loaded in the device for the alternate IPL.	Make sure the correct media volume is loaded in the correct device. Then continue with the next step.
2644 3136	The device is assigned to another system.	Go to the other system and vary off the device. Make sure that you loaded the correct media volume. Then continue with the next step.
Any other SRC	The system encountered a problem loading the Licensed Internal Code.	

7. You are shown the Install Licensed Internal Code display.

```
                      Install Licensed Internal Code

Select one of the following:

        1. Install Licensed Internal Code
        2. Work with Dedicated Service Tools (DST)
        3. Define alternate installation device
```

   a) If you have an alternate installation device attached to the system, perform steps 8 through 10. If you do not have an alternate installation device attached to the system, type a 1 and press the Enter key.

   Note: You are now ready to load your Licensed Internal Code. Consult your recovery checklist before continuing. The checklist tells you the correct option to select from the Install Licensed Internal Code display.

8. If you have an alternate installation device attached to the system, type a 3 to verify its address and determine whether it is enabled or disabled.

9. The Select Alternate Installation Device Bus display appears.

```
                  Select Alternate Installation Device Bus
                                                          System: YOURSYS
Type Option, press Enter.
   1=Select

Option      Bus Number      Selected
  _             2
  _             3              *
  _             4
  _             5
  _             6
  _             7
  _             8
  _             9
  _             A
  _             B
  _             C
  _             D
                                                               More.....
F2=Deselect device    F3=Exit    12=Cancel
```

The `Selected` column shows the bus where the alternate load source is currently defined. You can use option 1 (Select) to verify that the correct device has been selected. Select your media type and then press the Enter key. If the correct device has been selected, press F3 to exit.

You can use the F2 key to deselect the current bus and then use option 1 to select another. All buses that exist on the system are displayed.

After pressing the Enter key, there will be a brief delay (up to 10 minutes) while the bus is initialized. Following this delay, the Select Alternate Installation Device display appears.

```
                  Select Alternate Installation Device
                                                          System: YOURSYS
Type option, press Enter.
  1=Select   5=Details

            Resource                              Serial
Option      Name           Type       Model      Number          Selected
  _         TAP01          6380       001        00-1017187
  _         TAP08          3287       030        32-234333
  _         TAP02          6380       001        00-2017187
  _         TAP05          3287       030        72-234333          *
  _         TAP09          6380       001        00-1015187
  _         TAP16          3287       030        22-234633


F2=Deselect device    F3=Exit    F5=Refresh    F12=Cancel
```

Type a 1 in the **Option** field to select the device you want to use, and press the Enter key.

Note: When installing from an alternate installation device, be sure that only one device contains valid install media. This will prevent the wrong version of the Licensed Internal Code from being installed.

10. Press F3 to exit.
11. The Install Licensed Internal Code display appears. Type a 1 and press the Enter key.
12. Press the Enter key on the Confirm Alternate Installation Device display.

You are now ready to load your Licensed Internal Code. Consult your recovery checklist before continuing. The checklist tells you the correct option to select from the Install Licensed Internal Code display.

Related concepts

Using an alternate installation device

You can learn how to set up, disable, verify, and select an alternate installation device.

Loading the Licensed Internal Code

Use these steps to load the Licensed Internal Code to your system during a recovery.

If you receive an error display -

If you are using an alternate installation device and you receive an error display, it might be due to one of the following conditions -

- You are trying to use an alternate installation device that is not enabled.

See the "Setting up an alternate installation device" on page 361 and the "Disabling an alternate installation device" on page 364 topics and run the appropriate procedure.

Note: You might find that the address information is not available, or that the system configuration was changed so that the address information is not correct. If one of these situations occurs, you must determine the address information by a physical inspection of your system configuration. This inspection can be difficult and might vary depending on your system model and the specific configuration of your input/output (I/O) buses. IBM suggests that you call your next level of support for assistance in determining the addresses that you need for the alternate installation device. A service agreement might be needed for this type of assistance.

To complete the procedure for loading the Licensed Internal Code to your system during a recovery, follow these steps:

1. If a valid functional disk unit is not detected, the Select Load Source Device display is displayed. Select the disk that you want to be the load source disk unit. If the Select Load Source Device display is not shown, the system found a valid load source disk unit. Go to step 2.

```
                                Select Load Source Device
Type 1 to select, press Enter.

                                  Sys  Sys   I/O       I/O
Opt  Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
     68-0CCDDA23      6718   050     12    11       0      0     7     0
     68-0DDA4212      6718   050     12    11       0      0     6     0
     68-0C211074      6718   050     12    11       0      0     3     0
     68-0DDE1201      6718   050     12    11       0      0     1     0
     68-0CDDEA44      6718   050     55    11       0      0     4     0
```

   a) The following confirmation display is shown:

```
                                Confirm Load Source Device

You have chosen the following device to be the load source.
Press F10 to confirm your choice.
                                Sys  Sys   I/O       I/O
Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
68-0CDDEA44      6718   050     55    11       0      0     4     0
```

2. See the Install Licensed Internal Code (LIC) display.

   Select the correct option and press the Enter key.

```
                    Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
    Serial Number     Type     Model     I/O Bus      Controller     Device
     xx-xxxxxxx       xxxx      xxx          x              x            x

Select one of the following:

     1. Restore Licensed Internal Code
     2. Install Licensed Internal Code and Initialize system
     3. Install Licensed Internal Code and Recover Configuration
     4. Install Licensed Internal Code and Restore Disk Unit Data
     5. Install Licensed Internal Code and Upgrade Load Source
```

Attention: Be sure that you consult the correct recovery checklist before you select an option from the Install Licensed Internal Code (LIC) display. Some options remove all data from your system.

3. If an alternate installation device is defined and enabled, the Confirm Alternate Installation Device display appears.

   a) To recover from the alternate installation device, press the Enter key.

4. One of the following Install LIC and Initialize System – Confirmation displays is shown after you choose option 2, 3, 4, or 5 from the LIC installation display. Press F10 to continue the installation; pressing the enter key just redisplays the confirmation display.

```
               Install LIC and Initialize System - Confirmation

Warning:
    All data on this system will be destroyed and the Licensed
    Internal Code will be written to the selected disk if you
    choose to continue the initialize and install.

    Return to the install selection display and choose one of the
    other options if you want to perform some type of recovery
    after the install of the Licensed Internal Code is complete.



    Press F10 to continue the install.
    Press F12 (Cancel) to return to the previous display.
    Press F3 (Exit) to return to the install selection display.
```

```
               Install LIC and Recover Configuration - Confirmation

Warning:
    All data on the selected disk will be destroyed and the Licensed
    Internal Code will be written to this disk if you choose to
    continue the install.  When the install is complete, an IPL
    will be done and you will be prompted to continue the recovery
    of the disk configuration.



    Press F10 to continue the install.
    Press F12 (Cancel) to return to the previous display.
    Press F3 (Exit) to return to the install selection display.
```

```
               Install LIC and Restore Disk Unit Data - Confirmation

Warning:
    All data on the selected disk will be destroyed and the Licensed
    Internal Code will be written to this disk if you choose to
    continue the install.  When the install is complete, an IPL
    will be done and you will be prompted to restore the disk unit
    data that you previously saved.




    Press F10 to continue the install.
    Press F12 (Cancel) to return to the previous display.
    Press F3 (Exit) to return to the install selection display.
```

```
               Install LIC and Upgrade Load Source - Confirmation

Warning:
    All data on the selected disk will be destroyed and the Licensed
    Internal Code will be written to this disk if you choose to
    continue the install.  When the install is complete, an IPL
    will be done and you will be prompted to complete the upgrade.




     Press F10 to continue the install.
     Press F12 (Cancel) to return to the previous display.
     Press F3 (Exit) to return to the install selection display.
```

If you chose option 2, 3, 4, or 5 on the Install Licensed Internal Code (LIC) menu and pressed F10 on the confirmation display, the Initialize the Disk–Status display screen is shown. The actual time to initialize the disk can be considerably less than the estimated time, depending on the current state of the disk.

```
             Initialize the Disk - Status

The load source disk is being initialized.


Estimated time to initialize in minutes :   ___

Elapsed time in minutes . . . . . :   ___._
```

5. You are shown the Install Licensed Internal Code Installation Status display. You do not need to respond to this display. The system shows this display for approximately 30 minutes.

```
                  Licensed Internal Code Installation Status

Installation of the Licensed Internal Code in progress.


                    +----------------------------------------+
Percent complete    | XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX       |
                    +----------------------------------------+
                    0         20        40        60        100


Elapsed time in minutes  . . . . . . . . :     x.x




 Please wait.


 ------------------------------------------------------------------------
```

6. If an error occurs, you might see a display that requires a response.
7. You might see the Accept Console display. If it appears, press F10 to accept the current console. On the Accept And Set New Console Type On This IPL display, press Enter to continue.

Note: You finished loading your Licensed Internal Code.

Related information

Licensed Internal Code installation error displays

Recovering the vital product data information if the partition does not IPL in A-mode or B-mode

The vital product data (VPD) information contains information about the disk unit which is the load-source disk unit.

If the VPD information is lost, the partition might not IPL in A-mode or B-mode. Use the Work with Disk Units service function in D-mode to refresh the VPD information.

1. Enter Option 8 (Update system vital product data) on the Work with Disk Units display. The following confirmation displays:

```
                 Update system vital product data

The disk unit information below identifies the load source whose
location will be written to the system vital product data.

Press enter to confirm the vital product data should be written.

             Serial
ASP   Unit    Number       Type    Model
  1      1   $$$$$$$$$$   $$$$    $$$
```

   The load-source disk unit is identified in the display above.
2. Press Enter to update the VPD.
3. Press F3 to return to the Use Dedicated Service Tools (DST) menu.
4. On the Use Dedicated Service Tools (DST) menu, select option 7 (Start a service tool).
5. On the Start a Service Tool menu, select option 7 (Operator panel functions) .
6. On the Operator Panel Functions display, select the IPL source and IPL mode. Press F8 to restart the partition or press F10 to power off the partition.

```
                  Operator Panel Functions
                                                          System:   XXXX
IPL source:  2   (1=A, 2=B or 3=D)
IPL mode:    1   (1=Manual, 2=Normal, 3=Secure or 4=Auto)

Press Enter to change the IPL attributes and return
  to the main DST menu.
Press F8 to set the IPL attributes and restart the system.
  Machine processing will be ended and the system will be
  restarted.
Press F10 to set the IPL attributes and power off the system.
  Machine processing will be ended and the system will be
  powered off.
Press F12 to return to the main DST menu without changing
  IPL attributes.
```

After the VPD information is updated, you can IPL the partition in A-mode or B-mode.

Setting up your disk configuration after installing the Licensed Internal Code

Perform these steps to set up your disk configuration after installing the Licensed Internal Code and initializing the system.

When you install the Licensed Internal Code by using option 2 from the Install Licensed Internal Code (LIC) menu, the system does the following tasks:

- The system clears disk unit 1. Disk unit 1 contains information about how all the other disk units on your system are configured. If you are using logical partitioning, disk unit 1 also contains your partitioning configuration data.
- The system prepares to delete all data in the system auxiliary storage pool (ASP). The system ASP is not actually cleared until you perform the initial program load (IPL) after installing the Licensed Internal Code.

1. When you complete installing the Licensed Internal Code, you are shown the Disk Configuration Attention Report display on the A or B mode IPL:

```
                  Disk Configuration Attention Report

Type option, press Enter
  5=Display Detailed Report

OPT    Warning
___    New disk configuration
```

2. If you type a 5 in the option column (OPT), you are shown the following display:

```
The current configuration indicates a single unit system. You
can choose to accept it or do one of the following:

Following are the possible causes and recovery procedures:

• You can define a new configuration by adding units.

• Press F3 to exit to dedicated service tools (DST) and if
necessary, take the right option to get to
'Use DST' display.
On the 'Use DST' display,
- Select option 4, Work with disk units.
- Select option 1, Work with disk configuration.
- Select option 3, Work with ASP configuration.
- Select option 3, Add units to ASPs.

• If you are performing 'Recover mirror load source'
utility, press F3 to exit to Dedicated Service Tools (DST).
If necessary, take the right option to get to 'UseDedicated Service Tools
 (DST)' display.
On the 'Use Dedicated Service Tools (DST)' display,
- Select option 4, Work with disk units.
- Select option 2, Work with disk unit recovery.
- Select option 16, Recover mirrored load source.

Press Enter to accept the current configuration and continue.
```

3. Press F10 or Enter to accept the new disk configuration and continue.
4. Perform the following steps:
   a) Initialize all of the non-load source disk units on your system.
   b) Define which ASP each disk unit is configured in.
   c) Determine which ASPs to start mirrored protection on.

   See “Checklist 1: Configuring disks on a new system” on page 370 for a checklist on adding disk units to the system ASP, adding disk units to the user ASPs, starting mirrored protection on ASPs, and starting device parity protection. When you have completed the checklist, press F12 to cancel the DST menu. You are shown the IPL or Install the System menu.

   You have now completed your disk configuration. Continue with the next step on your recovery checklist, which is restoring the operating system.

Using IBM Navigator for i to set up your disk configuration after installing Licensed Internal Code

After you install Licensed Internal Code, you need to set up your disk configuration. Use the dedicated service tools (DST) debug mode to access disk management functions in Navigator for i where you can configure disk units in system, basic, and independent auxiliary storage pools (ASPs).

When you install the Licensed Internal Code by using option 2 from the Install Licensed Internal Code (LIC) menu, the system does the following tasks:

- The system clears disk unit 1. Disk unit 1 contains information about how all the other disk units on your system are configured. If you are using logical partitioning, disk unit 1 also contains your partitioning configuration data.
- The system prepares to delete all data in the system ASP. The system ASP is not cleared until you do the initial program load (IPL) after you install the Licensed Internal Code.

When you are finished configuring disk units, you can complete the step-mode IPL.

Note: Configure the Service Tools Network Interface to do these steps.

1. You might receive a Disk Configuration Attention Report like the following one after you load the Licensed Internal Code.

   If so, press F10 to accept the problems and continue.

```
                    DISK CONFIGURATION ATTENTION REPORT

TYPE OPTION, PRESS ENTER.
  5=DISPLAY DETAILED REPORT

PRESS F10 TO ACCEPT ALL THE FOLLOWING PROBLEMS AND CONTINUE.
THE SYSTEM WILL ATTEMPT TO CORRECT THEM.

OPT  PROBLEM
     NEW DISK CONFIGURATION
```

2. From the IPL or Install the System menu, select option 3 (Use Dedicated Service Tools (DST).

```
                         IPL or Install the System

   Select one of the following:
            1. Perform an IPL
            2. Install the operating system
            3. Use DST
            4. Perform automatic installation of the operating system
            5. Save Licensed Internal Code
```

3. On the DST Sign On display, sign on with the QSECOFR service tools ID and the password for the QSECOFR service tools user ID.

```
                   DST Sign On

Type choices, press Enter.

  Service tools user . . . . . . . . . . . QSECOFR
  Service tools password . . . . . . . . .        _
```

4. Change the password for the QSECOFR user profile on the resulting display, as the password is expired after the first use.

```
                         Change Service Tools User Password

Service tools user profile name . . . . . :  QSECOFR
Password last changed . . . . . . . . . . :  02/05/01

Type choices, press Enter.
  Current password . . . . . . . . . . . .   _

  New password . . . . . . . . . . . . . .

  New password (to verify) . . . . . . . .
```

5. On the Use Dedicated Service Tools (DST) menu, select option 6, Select DST console mode.

```
                      Use Dedicated Service Tools (DST)

 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system

     10. Work with remote service support

     13. Work with system security
     14. End batch restricted state
```

6. On the Select DST Console Mode display, select option 2 (Start DST debug mode on IPL).

On the Select DST Console Mode display, select option 2 (Start DST debug mode on IPL).

```
                              Select DST Console Mode

   Attention: Incorrect use of DST debug mode can cause damage
   to data in this system.  Contact your service representative
   for assistance.

   Select one of the following:

        1. Exit DST on IPL
        2. Start DST debug mode on IPL
```

7. On the IPL or Install the System menu, choose option 1 (Perform an IPL).

```
                           IPL or Install the System

   Select one of the following:
            1. Perform an IPL
            2. Install the operating system
            3. Use DST
            4. Perform automatic installation of the operating system
            5. Save Licensed Internal Code
```

8. On the Add All Disk Units to the System display, select option 1 (Keep the current disk configuration).

```
                      Add All Disk Units to the System

   Select one of the following:
    1. Keep the current disk configuration
    2. Perform disk configuration using DST
    3. Add all units to the system auxiliary storage pool (ASP)
    4. Add all units to the system ASP and balance data
```

9. On the Install Required for Operating System display, press Enter to continue.

```
                      Install Required for Operating System

   The system ASP has been cleared, which requires an install
   of the operating system.

   To install the operating system , do the following:
    -  Load the install media in the device used to install the
       operating system and make the device ready.
    -  Press Enter when the device is ready to install the
       operating system.

   -OR-

   Press F11 to display the dedicated service tools sign on or
     main menu and not install the operating system.
```

10. On the Select Type of IPL display, select option 2 (Step-mode IPL).

```
                      Select Type of IPL

   Select one of the following:
        1. Normal IPL
        2. Step-mode IPL
```

11. Step through the IPL by pressing Enter.

    The last IPL step before you configure disk units is Storage Management Recovery. Press Enter on the Storage Management Recovery display.

```
                      Licensed Internal Code IPL in Progress

IPL:
  Type . . . . . . . . . . . . . . . . . . :  Attended
  Start date and time . . . . . . . . . . .:  00/00/00 00:00:00
  Previous system end . . . . . . . . . . .:  Abnormal

IPL step . . . . . . . . . . . . . . . . . :  Storage Management Recovery
```

12. Stop at the resulting IPL step Start LIC Log display.

```
                    Licensed Internal Code IPL in Progress

IPL:
  Type . . . . . . . . . . . . . . . . . . :  Attended
  Start date and time . . . . . . . . . . .:  00/00/00 00:00:00
  Previous system end . . . . . . . . . . .:  Abnormal

IPL step . . . . . . . . . . . . . . . . . :  Start LIC Log
```

13. Open Navigator for i to begin disk unit configuration.
14. In the Environmental tasks window, click **Open Navigator for i service tools window**.
15. Enter the system name or IP address on the resulting window.
16. Expand ***your system*** > **Disk Units**.
17. Right-click **All Disk Units** and select **Add Disk Unit**.

    The wizard guides you through the process of adding disk units to disk pools and starting disk protection. Refer to a printed copy of your disk configuration to create all the necessary disk pools.
18. When your disk unit configuration is complete, continue through the step mode IPL. Press Enter on all the following IPL steps through Start the Operating System.

    After the IPL steps complete, the Install the Operating System menu appears:

```
                        Install the Operating System

Type options, press Enter.

     Install
       option . . . .  __              1=Take defaults (No other
                                         options are displayed)
                                       2=Change install options

Date:
  Year . . . . .  __                   00-99
  Month  . . . .  __                   01-12
  Day  . . . . .  __                   01-31

Time:
  Hour . . . . .  __                   00-23
  Minute . . . .  __                   00-59
  Second . . . .  __                   00-59
```

You completed the disk unit configuration. Continue with the next step on your recovery checklist, which is restoring the operating system.

Recovering your disk configuration

Every disk unit on your system contains information about how it is configured. Dedicated services tools (DST) provides an option to recover the disk configuration on your system by using this information. The system reads each disk, assigns it to the correct ASP, and rebuilds the disk configuration information on unit 1.

When you install the Licensed Internal Code by using option 3 from the Install Licensed Internal Code (LIC) menu, the system performs the following operations:

- Clears disk unit 1. Disk unit 1 contains information about how all the other disk units on your system are configured.
- Prepares to delete all data in the system auxiliary storage pool (ASP). The system ASP is not actually cleared until you perform the initial program load (IPL) after installing the Licensed Internal Code.

In many cases, you can recover your disk configuration and avoid having to reload all your user ASPs. To recover your disk configuration, follow these steps:

1. When you complete installing the Licensed Internal Code, decide whether you want to install the Licensed Internal Code and initialize the system, or install Licensed Internal Code and recover the configuration.

- If you select option 2 (Install the Licensed Internal Code and Initialize the System), when the system does an A-mode or B-mode IPL, you are shown the following display:

```
                  DISK CONFIGURATION ATTENTION REPORT

TYPE OPTION, PRESS ENTER
  5=DISPLAY DETAILED REPORT

PRESS F10 TO ACCEPT ALL THE FOLLOWING PROBLEMS AND CONTINUE.
THE SYSTEM WILL ATTEMPT TO CORRECT THEM.

OPT  PROBLEM
     NEW DISK CONFIGURATION
```

  Press F3 to continue, do not press F10 to confirm.

- If you select option 3 (Install the Licensed Internal Code and Recover Configuration), when the system does an A-mode or B-mode IPL, you are shown the following display:

```
                  Disk Configuration Error Report

Type option, press Enter
  5=Display Detailed Report

OPT    Error
___    Missing disk configuration
```

  a) If you type a 5 in the option column (OPT), you are shown Missing Disk Configuration display:

```
                  Missing Disk Configuration

The system disk configuration has been lost.  The IPL
cannot be continued.

The following are the suggested recovery procedures:
o  If the original system has more than one disk unit
configured and you want to keep the configuration
currently on the system, use Recover configuration
under Work with Disk Units.

Press F3 to exit to dedicated service tools (DST) and
if necessary, take the right option to get to the 'Use
dedicated service tools' (DST) display.  On the 'Use
dedicated service tools' (DST) display:
-  Select option 4, Work with disk units.
-  Select option 2, Work with disk unit recovery.
-  Select option 5, Recover configuration.
o  If the original system had only one disk unit configured
or you do not want to save the data currently on the
system, re-install the Licensed Internal Code, then
re-create the configuration, restore your data.

F3=Exit to use DST   F12=Cancel
```

  b) From either display, press F3 (Exit to Use dedicated service tools (DST)). The Use dedicated service tools Sign On display is shown:

```
                    DST Sign On

Type choice, press Enter.

  Service tools user . . . . . .
  Service tools password . . . . . .
```

2. Sign on to dedicated service tools.

   The system displays the Use Dedicated Service Tools menu.
3. From the Use Dedicated Service Tools (DST) menu, select option 4 (Work with disk units).
4. From the Work with Disk Units menu, select option 2 (Work with disk unit recovery).

5. From the Work with Disk Unit Recovery menu, select option 5 (Recover disk configuration). You might see the following display:

```
                  Problem Report

Note: Some action for the problems listed below might need to
be taken.  Please select a problem to display more detailed
information about the problem and to see what possible
action might be taken to correct the problem.

Type option, press Enter.
  5 = Display Detailed Report

OPT   Problem
 _    Load Source has been re-built
 _    ASPs will be cleared
```

   Press F10 to ignore problems and continue.

```
                       Confirm Recover Configuration

ATTENTION: There are problems in the system that could cause
some of the system data to be destroyed.

Press F11 to display the problems.
Press F10 to confirm your choice to recover configuration.
Press F12 to return to change your choice.

Possible configuration found in the system records:

            Serial                      Resource
ASP  Unit   Number      Type  Model     Name           Status
--   ----   ----------  ----   ---      -------------  ----------------------
--   ----   ----------  ----   ---      -------------  ----------------------
--   ----   ----------  ----   ---      -------------  ----------------------
--   ----   ----------  ----   ---      -------------  ----------------------
--   ----   ----------  ----   ---      -------------  ----------------------
--   ----   ----------  ----   ---      -------------  ----------------------
                                                                    More...

F10=Confirm recover         F12=Cancel
--------------------------------------------------------------------------------
```

6. Check the configuration of disk units on the display.

   The display shows the disk units that are assigned to each user ASP and to the system ASP (ASP 1). The warning on the display means that the system will clear all data on disk units in the system ASP.

   If this configuration is not correct, contact a service representative or software support for assistance. Do not proceed further without getting help.

   If the configuration that is shown is correct, press F10 to confirm the configuration. The system builds the configuration information and returns to the DST menu.
7. Press F12 to cancel the DST menu. You are shown the IPL or Install the System menu.

You have now completed disk configuration. Continue with the next step on your recovery checklist, which is restoring the operating system.

Starting your system after restoring the Licensed Internal Code

If you do not need to restore the operating system, perform this procedure to start your system.

After you complete loading the Licensed Internal Code, you should see the IPL or Install the System display.

```
                        IPL or Install the System

Select one of the following:
            1. Perform an IPL
            2. Install the operating system
            3. Use dedicated service tools (DST)
            4. Perform automatic installation of the operating system
            5. Save Licensed Internal Code
```

To start your system after restoring it, follow these steps:

1. Select option 1 (Perform an IPL) on the IPL or Install the System menu. When the initial program load (IPL) completes, the Sign On display is shown.
2. If your operator panel has a keylock switch, turn the key in the keylock switch to the normal position.
3. Sign on the system as QSECOFR.
4. If the Select Product to Work with PTFs display appears, press F3 (Exit) to continue the IPL.
5. Press the Enter key in response to any messages that are displayed.
6. When you are shown the IPL options display, type your choices and press the Enter key.

```
                               IPL Options

Type choices, press Enter.

System date  . . . . . . . . . . . . . .   08 / 01 / 04
System time  . . . . . . . . . . . . . .   12 : 00 : 00
System time zone . . . . . . . . . . . .   Q0000UTC
Clear job queues . . . . . . . . . . . .   N
Clear output queues  . . . . . . . . . .   N
Clear incomplete job logs  . . . . . . .   N
Start print writers  . . . . . . . . . .   Y
Start system to restricted state . . . .   N


Set major system options . . . . . . . .   Y
Define or change system at IPL . . . . .   N
```

You have now completed starting your system after recovering the Licensed Internal Code. Consult your recovery checklist to determine the next step in your recovery process.

Chapter 5. Restoring the operating system

Use this information to recover the operating system.

The procedure that is described here assumes that the Licensed Internal Code is already installed on your system. Normally, the Licensed Internal Code is installed. However, if the Licensed Internal Code is not on your system or is damaged, see Chapter 3, “Selecting the appropriate recovery strategy,” on page 53 to determine the correct recovery procedure for your situation.

Why you restore the operating system: You might need to restore the operating system for several reasons, such as -

- You are encountering problems with the operating system, such as damaged objects.
- The software support center suggests it.
- You replaced a disk unit in the system auxiliary storage pool (ASP).
- You are updating your system to a new release.

Find these things before you begin -

- Your most recent SAVSYS media. You can choose any one of the following actions to create a SAVSYS media volume:
  - Run the Save System (SAVSYS) command.
  - Use option 21 from the **Save** menu.
  - Use option 22 from the **Save** menu.
  - Use option 11 from the **Run Backup** menu.

  Attention: Do not use a media volume that you created through dedicated service tools (DST) by using option 5=Save Licensed Internal Code unless you are instructed to do so by a Software Services representative. This process creates a media volume that does not contain the Licensed Internal Code program temporary fix (PTF) inventory information or the IBM i operating system. If this media is used for the recovery, the Licensed Internal Code needs to be reinstalled from either SAVSYS media or distribution media before any PTFs can be loaded on the system.
- If you do not have current SAVSYS media or they are damaged, you need the following media -
  - The distribution media that is supplied by IBM
  - All the media for program temporary fixes (PTFs) that you applied.

    Attention: Use the distribution media only if you do not have SAVSYS media. If you use the distribution media to restore the operating system, the version you restore does not include any of the PTFs that you applied. In addition, the restore sets the following default values for the IBM i licensed program:

    - System information
    - Network attributes
    - Configuration lists
    - Edit descriptions
    - Reply list entries
    - IBM-supplied subsystem descriptions
    - Passwords for IBM-supplied profiles
- The list of all the PTFs applied to your system at the time you saved the entire system. Attach this list to your backup log or keep it with your SAVSYS media.
- The DST password for the QSECOFR service tools user ID.
- The QSECOFR password that is associated with the SAVSYS media that you are using.

© Copyright IBM Corp. 1997, 2025

Do these things before you begin -

- Clean the read and write heads of the tape unit if you are using a tape unit.
- If your source system (the system that needs recovery) is operational, print a list of all the PTFs currently on the system, unless you printed the list before the Licensed Internal Code was restored. Type the following command and press the Enter key:

```
DSPPTF LICPGM(*ALL) OUTPUT(*PRINT)
```

Related concepts

The relationship between save and restore commands
You can learn which store commands can be used, based on how the objects were saved.

Related information

Installing, upgrading, or deleting IBM i and related software
Console Advanced Topics

Choosing the right procedure for restoring the operating system

You can restore the operating system in several different ways.

At several points during the restore process, you need to make decisions based on which of these operations you are using:

Complete Restore
Use a complete restore operation if the operating system is not on your system or if it has damaged objects. This restores all the IBM-supplied objects in the QSYS library and the language libraries from media.

Abbreviated Install
Use an abbreviated installation operation to replace parts of the operating system or system information, such as system values or the system reply list.

Use the recovery checklist that you selected in Chapter 3, “Selecting the appropriate recovery strategy,” on page 53 topic to determine the correct procedure for your situation. You also need to know whether you are restoring from SAVSYS media or from the IBM-supplied distribution media. Use the distribution media only if you do not have usable SAVSYS media.

How to begin restoring the operating system: The steps you have already performed determine where you start:

- If you have just restored or installed the Licensed Internal Code, you are doing a complete restore operation of the IBM i operating system. You should see the initial program load (IPL) or Install the System display. Start with the steps described in “Tasks to restoring the operating system” on page 143 topic.
- If restoring the operating system is the first step or only step in your recovery process, start by performing a manual IPL.

Loading the operating system using a manual initial program load

Follow the following steps to begin loading the operating system.

Use these steps only if you did not restore the Licensed Internal Code, as described in Chapter 4, “Recovering the Licensed Internal Code,” on page 121 topic.

To do a manual IPL, follow these steps:

1. Verify that the tape unit or optical drive is ready.
   For more information about loading the tape or optical media see the setup manual for the device.
2. Load the first media volume from your most recent SAVSYS media in the appropriate device.

If you do not have SAVSYS media or they are unusable, load the first volume from the distribution media. Use the distribution media only if no SAVSYS media exists.

Attention: If you were not instructed to do so by Software Services, do not use media that was created through DST by selecting option 5=Save Licensed Internal Code. Media that was created through this process does not contain the Licensed Internal Code program temporary fix (PTF) inventory information or the IBM i operating system. If you use this media for the recovery process, the Licensed Internal Code must be installed from either SAVSYS media or from your distribution media before you can load any PTFs onto the system.

3. Place the system in manual mode.
4. Verify that all jobs are ended and all users are signed off.
5. Power off the system.

```
PWRDWNSYS OPTION(*IMMED) RESTART(*YES) IPLSRC(B)
```

6. Continue with "Tasks to restoring the operating system" on page 143.

Tasks to restoring the operating system

You are ready to begin these steps when you have completed an initial program load (IPL). Either you have just restored the Licensed Internal Code or you have just performed a manual IPL from your alternate IPL device.

Note: If you use LAN Console, perform the following steps to reset your console:

1. On the IPL or Install the System display, select option 3 (Use dedicated service tools (DST)). Press Enter to continue.
2. Sign on to DST using a service tools user profile that has security officer authority and the assigned password.
3. On the *Use DST* display, select option 5 (Work with DST environment). Press Enter to continue.
4. On the Work with DST Environment display, select option 2 (System Devices). Press Enter to continue.
5. On the Work with System Devices display, select option 6 (Console Mode). Press Enter to continue.
6. On the Select Console Type display, select option 3 (Operations Console (LAN)). Press Enter to continue.

   Note: If you replace the load source disk unit you might need to use the console service functions (65+21) to specify your console type, so that you can perform the restore operations.
7. Press F3 or F12 to get back to the IPL or Install the System display screen.

The IPL or Install the System display will appear.

```
                        IPL or Install the System

Select one of the following:
         1. Perform an IPL
         2. Install the operating system
         3. Use DST
         4. Perform automatic installation of the operating system
         5. Save Licensed Internal Code
```

Task 1: Starting to restore the operating system

Use these procedures to prepare for restoring the operating system.

1. If you are using an encrypting tape drive, verify that the Encryption Key Manager (EKM) is running and connected to the tape library before you begin the recovery operation.

   The EKM contains the encryption keys that are needed for the recovery operation. If you are not using an encrypting tape drive, skip this step.
2. Load the first media volume from your most recent SAVSYS media in the appropriate device.

- If you do not have SAVSYS media or they are unusable, load the first volume from the distribution media. Use the distribution media only if no SAVSYS media exists.
- If you are using encrypted user auxiliary storage pools (ASPs) or independent ASPs, verify that you are using the SAVSYS media that contains the encryption keys for the ASPs.

Attention: Do not use media that you created through dedicated service tools (DST) unless you were instructed to do so by Software Services. Media that was created through this process does not contain the Licensed Internal Code program temporary fix (PTF) inventory information or the IBM i operating system. If this media is used to do a recovery, you must reinstall the Licensed Internal Code from either SAVSYS media or from distribution media before you can load any PTFs onto the system.

3. From the IPL or Install the System display, select option 2 (Install the operating system).

   Note: Do not use option 4 (Perform automatic installation of the operating system) to restore the operating system. Option 4 is used to install the system. Option 4 is not used for system recovery.

4. Make a selection on the Install Device Type Selection display and press Enter.

   Note: Depending on the installation device type selected, other input might be needed.

5. Press the Enter key.

   The Confirm Install of the Operating System screen is displayed.

```
                    Confirm Install of Operating System

Press Enter to confirm your choice to install the operating
   system.
Press F12 to return and cancel your choice to install the
   operating system.
```

6. Press the Enter key.

   If you see the Dedicated Service Tools (DST) Sign On display, continue with step 7. If you see the Select a Language Group display, skip to step 8.

7. If your system is set up to prevent unauthorized installation of the operating system, the Dedicated Service Tools (DST) Sign On screen is displayed.

```
                  Dedicated Service Tools (DST) Sign On

   Type choice, press Enter.
   Service tools user . . . . . .   ______
   Service tools password . . . . . .   ______
```

   Type the DST service tools user ID and the DST service tools password and press the Enter key. You see the Select a Language Group display.

   Notes:

   - The DST service tools user ID and DST service tools password are case-sensitive.
   - Sign on to DST by using a service tools user profile that has security officer authority and the assigned password. The service tools user for security level DST is QSECOFR. The service tools profile QSECOFR password is expired after the first use. At the Change Service Tools User Password display, enter in all uppercase letters the current password QSECOFR and a new password, along with the verification password.
   - If your current DST password does not work, the password might be reset to the default value. Try QSECOFR as the DST password.
   - For more information about securing the installation of the operating system, see Restoring the operating system.

8. You are shown the Select a Language Group display.

   This display shows the primary language currently on the save media that you are restoring.

```
                          Select a Language Group

Note:  The language feature shown is the language feature
installed on the system.

Type choice, press Enter.

Language feature . . . . . .  2924
```

If this value does not match the value that is already on your system, verify that you have the correct save media. If you change the value on the display, you are prompted to insert different media to load a different language feature. Press the Enter key. You are shown the Confirm Language Feature Selection display.

```
                    Confirm Language Feature Selection

Language feature  . . . . . :   2924

Press Enter to confirm your choice for language feature.
   Installing the system will continue.
Press F12 to return to change your choice for
   language feature.
```

9. Press the Enter key to confirm the information.

   Note: If you must change your system's primary language, see Changing the primary language of your system or logical partition for more information.

   If you see the Add All Disk Units to the System display, continue with step 10. If you see an IPL status message display, skip to step 11.

10. The Add All Disk Units to the System display is shown only if disk units are in nonconfigured status.

```
                    Add All Disk Units to the System
                                                          System: XXXX
   Select one of the following:
    1. Keep the current disk configuration
    2. Perform disk configuration using DST
    3. Add all units to the system auxiliary storage pool
    4. Add all units to the system ASP and balance data
```

    Disk units might be in nonconfigured status for these reasons:

    - The Licensed Internal Code was installed by using option 2 or option 3.
    - The recover disk configuration procedure was not run.
    - All disk units except unit 1 appear in nonconfigured status.
    - You have new or spare disk units that are not configured on your system.

    Use the information in Table 40 on page 145 to determine how to respond to this display:

Table 40. Configuring disk while installing the operating system

Your recovery situation	How to respond to the display
Restoring your entire system to a different system or to an upgraded system.	• If you plan to have user ASPs or mirrored protection, you can select option 2 to configure your disks now. Or, you can select option 1 now and configure the disks after you restore the operating system. If you plan to configure disk protection or user ASPs, use the instructions in the Chapter 21, “Configuring disks and disk protection,” on page 369 topic. • If you want all disks in the system ASP and do not plan to use mirrored protection, select option 3.

Table 40. Configuring disk while installing the operating system (continued)

Your recovery situation	How to respond to the display
The Licensed Internal Code was installed on your system during an option 2 or option 3 recovery.	• If you want all disks in the system ASP and do not plan to use mirrored protection, select option 3. • If you had user ASPs or mirrored protection on your system before the failure, you can select option 2 to reconfigure your disks. Option 2 removes any data from disks that show as not configured. • You can select option 1 and configure your disks later. However, the system cannot recover the data on the disks that are not configured.

If option 3 is selected, you receive an Attention Report display. Take the directed action for more detailed information, if necessary. Otherwise, press F10 to accept the problems and continue.

If you installed the Licensed Internal Code by selecting option 2, you receive an Attention report display. Take the directed action for more detailed information, if necessary. Otherwise, press F10 to accept the problems and continue.

11. This screen is an example of a status display.

    These status displays do not require any action by the user.

```
                        Licensed Internal Code IPL in Progress

IPL:
   Type . . . . . . . . . . . . . . . :    Attended
   Start date and time  . . . . . . . :    xx/xx/xx  xx:xx:xx
   Previous system end  . . . . . . . :    Abnormal
   Current step / total . . . . . . . :       16      16
   Reference code detail  . . . . . . :    C6004065

IPL step                                   Time Elapsed   Time Remaining
  Commit Recovery                            xx:xx:xx        xx:xx:xx
  Data Base Initialization                   xx:xx:xx        xx:xx:xx
  Journal IPL Clean up                       xx:xx:xx        xx:xx:xx
  Commit Initialization                      xx:xx:xx        xx:xx:xx
 >Start the operating system                 xx:xx:xx        xx:xx:xx

Item:
  Current / Total . . . . . . . :

Sub Item:
  Identifier . . . . . . . . . . :    Searching tape for install media
  Current / Total . . . . . . . :

Wait for next display or press F16 for DST main menu
```

The following list shows some of the IPL steps that are shown on the Licensed Internal Code IPL in Progress display:

- Authority Recovery
- Journal Recovery
- Database Recovery
- Journal Synchronization
- Start the Operating System

Some of the IPL steps might take a long time.

While the system is being restarted, system reference codes (SRCs) are displayed on the control panel of the system unit or logical partition to indicate what step is in progress. The System reference code list describes these SRCs. If the same SRC is displayed for a long time in solid (not flickering) lights, your system might not be able to complete the IPL. Contact software support.

The system might prompt you for more volumes of your SAVSYS media or the distribution media. Follow the instructions on the display.

After the IPL steps are completed, the Install the Operating System menu is displayed.

```
                        Install the Operating System

Type options, press Enter.

     Install
       option . . . .  __              1=Take defaults (No other
                                         options are displayed)
                                       2=Change install options

Date:
  Year . . . . .  __                   00-99
  Month  . . . .  __                   01-12
  Day  . . . . .  __                   01-31

Time:
  Hour . . . . .  __                   00-23
  Minute . . . .  __                   00-59
  Second . . . .  __                   00-59
```

Task 2: Selecting the installation options

You can learn how to select the installation options.

1. Type your choice for the *Install option* prompt based on the following situations:
   - If you are doing a complete restore operation, select option 1 (Take defaults). This restores the entire operating system. Use this option if any of the following conditions is true:
     - You are recovering from a failure of the load source unit.
     - You are restoring your entire system.
     - You are upgrading to a new system.
     - You are recovering damaged objects in your operating system.

     Note: If you are doing a complete restore operation and restoring a primary language other than English or if you have changed some of the shipped values of international system values, such as date and time (QDATFMT) or coded character set ID (QCCSID), you must select option 2 (Change install options). This ensures the language-dependent system values are restored correctly.

     Restoring to a different system?

     If you are restoring to a different system (with a different serial number or different logical partition) and you want your network attributes to be restored, select option 2 (Change install options). This will allow you to select to restore your network attributes from your save media.
   - If you are doing an abbreviated install operation, select option 2 (Change install options). This allows you to specify which parts of the operating system you want to restore. Use this option if you are recovering damaged system information, such as system values. Network attributes will not be restored.
2. If the date and time are not correct, type new values.
3. Press the Enter key.

   If you selected install option 1 (Take defaults), skip to step 14.

   If you selected install option 2 (Change install options), you are shown the Specify Install Options display.

```
                          Specify Install Options

Type choices, press Enter.

     Restore option . . . . __          1=Restore programs and language objects
                                          from current media set
                                        2=Do not restore programs or
                                          language objects
                                        3=Restore only language objects
                                          from current media set
                                        4=Restore only language objects
                                          from a different media set using the
                                          current install device

     Job and output
       queues option . . . . __         1=Clear,  2=Keep

     Distribute IBM i on
       available disk units  __         1=Yes,  2=No
```

Continue with step 4.

4. Type your choice for the *Restore option* prompt based on the following options:

   Note: If you are performing a complete system recovery, you must select option 1.

   1 = Restore programs and language objects
   This option restores system objects from the media that you are using. Use this option if damage to a system user profile is found. If you select the option, you will be prompted to specify whether to restore system information, edit descriptions, or the system reply list.

   Notes:

   - If your system had access path recovery times for user auxiliary storage pool (ASP)s and the use ASPs have not yet been reconfigured, the system cannot restore access path recovery times for the ASPs. Use the Edit Recovery Times for Access Paths (EDTRCYAP) command to set the times after you have reconfigured your ASP configuration.
   - If you are restoring a primary language other than English, you must select option 1.

   2 = Do not restore programs or language objects
   This option leaves the current programs and language objects on the system. Select this option to do an abbreviated install of the operating system.

   When you select this option, the system does the following tasks:

   - Nothing is restored from media. Any damaged objects that the system deletes and creates again are empty.
   - All libraries on the system are checked for damage. Damaged libraries are deleted and then created again.
   - All system libraries (including QSYS) are created if they do not exist.
   - Information associated with the user profiles is created if it does not exist or is damaged.
   - The system entry-point table is created again.

   3 = Restore only language objects from current media
   This option loads only those objects that make up the national language. The search for the language files begins on the current media. Select this option only if you need to change your primary language because you are recovering using the distribution media.

   4 = Restore only language objects from a different media
   This option loads only those objects that make up the national language. The system prompts you to insert the language media. Select this option only if you need to change your primary language because you are recovering using the distribution media.

   Note: If you want to change the primary language, see Installing, upgrading, or deleting IBM i and related software topic. You should avoid changing the primary language during a recovery.

5. Type your choice for the *Clear Job and Output Queues* prompt based on the following options:

1 = Clear
If you do not want to keep any spooled files or entries on job queues after the installation or if you know they are damaged, select this option. The system removes all jobs on job queues and spooled files. It re-creates any internal objects associated with them. You should select this option if you are performing an abbreviated installation of the operating system. This option also resets the counter for assigning unique job numbers.

2 = Keep
Any entries on job queues and output queues remain after the installation. This is the normal option for this prompt.

6. Type your choice for the Distribute operating system on available disk units prompt based on the following options:

   1 = Yes

   This option will take more time. You should only use this option in specific recovery situations as directed by your service representative. This option specifies to distribute operating system objects on available disk units in the system ASP during the installation process.

   2 = No
   Specifies to not distribute operating system objects on available disk units. This option will restore operating system objects from the media over the existing objects on the system.

7. Type your choices on the Specify Restore Options display and press the Enter key.

   If you did not specify 1 for the *Restore option* prompt, skip to step 14.

   If you specified 1 for the *Restore option* prompt, you are shown the Specify Restore Options display.

```
                        Specify Restore Options

Type choices, press Enter.

     Restore from the installation media:

        System information. . . _    1=Restore,  2=Do not restore
        Edit descriptions . . . _    1=Restore,  2=Do not restore
        Message reply list. . . _    1=Restore,  2=Do not restore
        Job descriptions. . . . _    1=Restore,  3=Keep customization
        Subsystem descriptions. _    1=Restore,  3=Keep customization
```

   Continue with step 8.

8. Type your choice for the *System information* prompt based on the following options:

   Note: If you are performing a complete system recovery to a different system (with a different serial number or different logical partition), and you want to restore the network attributes from the save media, you must select option 1 (Restore).

   1 = Restore
   The system restores the system values and system management objects, such as access path recovery times, from media. Select this option if any of the following situations is true:

   - You received a message during a previous initial program load (IPL) stating that the system value object was created again.
   - You want to restore them to their values from your last save operation.
   - You restored the operating system using option 2 or option 3 from the Install Licensed Internal Code (LIC) display.
   - You are restoring a primary language other than English, or if you changed the shipped value of the QDATFMT system value.
   - Your system is in a distributed relational database with unlike systems, and the coded character set ID (QCCSID) system value is different from the shipped value.

   If you have changed your primary language since your last save operation, the system might change some language-dependent system values during the restore operation.

2 = Do not restore
The system values and system management objects that are currently on the system are not changed.

A system value object must always be present on an operational system. If the system value object does not exist, the system restores it, even if you select option 2.

9. Type your choice for the *Edit descriptions* prompt based on the following options:

   1 = Restore
   The system restores the edit descriptions from media. Select this option if:

   - The edit descriptions are damaged.
   - You want to restore them to their values from your last save operation.
   - You installed the Licensed Internal Code by using option 2 or option 3 from the Install Licensed Internal Code (LIC) display.

   2 = Do not restore
   The edit descriptions currently on the system are not changed.

10. Type your choice for the *Message Reply List* prompt based on the following options:

    1 = Restore
    The system restores the reply list from media. Select this option if:

    - The message reply list is damaged.
    - You want to restore it to its values from your last save operation.
    - You installed the Licensed Internal Code by using option 2 or option 3 from the Install Licensed Internal Code (LIC) display.

    2 = Do not restore
    The message reply list currently on the system is not changed.

    The default value for these options is 2 if the operating system is loaded on the system. The default value is 1 if the operating system is not already loaded.

11. Type your choice for the *Job descriptions* prompt based on the following options:

    1 = Restore
    The system restores the job descriptions from media.

    3 = Keep customization
    The system restores the objects from media and customizes them with the values from the same objects that are already on the system.

12. Type your choice for the *Subsystem descriptions* prompt based on the following options:

    1 = Restore
    The system restores the subsystem descriptions from media.

    3 = Keep customization
    The system restores the objects from media and customizes them with the values from the same objects that are already on the system.

13. Type your choices on the display and press the Enter key.
14. The IBM i Installation Status display indicates how many program or language objects have been restored.

    They are for your information only and require no response. The display is blank for a time between stage 4 and stage 5.

```
Message ID. . . :   CPI2070                      IBM i Installation Status

Stage 2    +----------------------------------------------------+
           |XXXXXXX                                             |
           +----------------------------------------------------+
           0       20         40         60         80         100

Installation                                                          Objects
Stage                                                  Completed      Restored

   1  Creating needed profiles and libraries . . . . :      X             XXXX

>> 2  Restoring programs to library QSYS . . . . . . :                   XXXX

   3  Restoring language objects to library QSYS . . :

   4  Updating  program table. . . . . . . . . . . . :

   5  Installing database files. . . . . . . . . . . :

   6  Installing base directory objects. . . . . . . :

   7  Completing installation. . . . . . . . . . . . :
```

15. Continue loading media in sequence when messages are shown that ask you to load additional media.

    The system searches through the media and loads the necessary programs and language information. After processing all the system save media or distribution media, the system might display the following message at the bottom of a blank display:

    ```
    Operating system has been installed. IPL in progress.
    ```

    When the IPL is complete, the IPL Sign On display is shown and the system in ready to complete the IPL.

Related information

System values

Journal management

Task 3: Selecting initial program load options

You can learn how to select initial program load (IPL) options.

1. Sign on as QSECOFR.

   The password for QSECOFR depends on the recovery steps you have performed:

   - If you restored the operating system without first restoring the Licensed Internal Code, the QSECOFR password was not changed during the restore process.
   - If you used option 1 to restore the Licensed Internal Code before you restore the operating system, the system associates the QSECOFR password with your SAVSYS media.
   - If you used option 2 or option 3 to install the Licensed Internal Code before you restore the operating system, the system requires no password at this time. The system displays the Change Password display. The system sets the QSECOFR user profile to *EXPIRED, and the system sets the password to QSECOFR. The system restores the system security level after you install the operating system and the IPL completes. When the IPL completes, the QSECOFR password is the password associated with the SAVSYS media you used.

   Note: If you do not know the QSECOFR password, you can use dedicated service tools (DST) to reset the password to its shipped value of QSECOFR.

```
                     Change Password

Password last changed . . . . . xx/xx/xx

Type choices, press Enter.

   Current password . . . . . QSECOFR

   New password . . . . . _______

   New password (to verify) . . . _______
```

2. Press the Enter key.

   Informational messages are displayed.

3. If the Select Product to Work with PTFs display appears, press F3 (Exit) to continue.

```
                Select Product to Work with PTFs

Position to . . . . . .____________ Product


Type options, press Enter.  Press F21 to select all.
   1=Select

              Product
Opt  Product  Option  Release
 _   5770999  *BASE   V7R1M0
 _   5770SS1  *BASE   V7R1M0
```

4. You are shown the IPL Options display.

```
                              IPL Options

Type choices, press Enter.

System date  . . . . . .   xx / xx / xx     MM / DD / YY
System time  . . . . . .   xx : xx : xx     HH : MM : SS
System time zone . . . . . .   Q0000UTC     F4 for list
Clear job queues . . . . . .   N             Y=Yes, N=No
Clear output queues  . . . . . .   N         Y=Yes, N=No
Clear incomplete job logs  . . . . . .   N   Y=Yes, N=No
Start print writers  . . . . . .   Y         Y=Yes, N=No
Start system to restricted state . . . .   N  Y=Yes, N=No


Set major system options . . . . . .   Y     Y=Yes, N=No
Define or change system at IPL . . . . .   Y  Y=Yes, N=No
```

   The values that appear as defaults depend on the recovery steps you have performed.

5. If the system date, time, and time zone are not correct, type the correct values.

   If you installed the Licensed Internal Code using option 2 or option 3, the date and time might be blank. The system date must have a year value in the range of 87 to 99, or 00 to 22.

6. Type your choice for the *Start print writers* prompt based on the following options:

   N = No
   Select this value if you are going to restore user profiles, device configuration objects, user libraries, and authorities.

   Y = Yes
   Select this value if you have completed your recovery.

7. Type your choice for the *Start system to restricted state* prompt based on the following options:

   Y = Yes
   Select this value if you are going to restore user profiles, device configuration objects, user libraries, and authorities. Only the console is started (varied on).

N = No
Select this value if you have completed your recovery. All devices are started.

8. Type Y (Yes) for the *Set major system options* prompt.
9. Type Y (Yes) at the *Define or change system at IPL* prompt.
10. Type your choices on the display and press the Enter key.

Task 4: Setting major system options

You can learn how to set major system options.

1. You are shown the Set Major System Options display.

```
                        Set Major System Options

Type choices, press Enter.

Enable automatic configuration . . . . . Y            Y=Yes, N=No
Device configuration naming  . . . . . . *NORMAL      *NORMAL, *S36,
                                                                    *DEVADR
Default special environment. . . . . . . *NONE        *NONE, *S36
```

2. If you choose to enable automatic configuration, the system will create a device description for every device that is attached to your system.

   The device description will be named according to the value that you specify for Device configuration naming. You might need to change the names and descriptions of these device descriptions later.

   If you choose not to enable automatic configuration, you will need to configure at least one device later in your recovery. You must configure the device after you finish restoring the operating system and before you restore any other information. You might also need to correct the device configuration for the system console and respond to SRC A900 2000 (step 10 in Task 6–Completing the IPL topic). The instructions to recover from SRC A900 2000 are provided.
3. Type your choices and press the Enter key.
4. If you specified Y for the Define or change system at the *IPL* prompt in step 4 of Task 3, continue with "Task 5: Defining or changing the system at initial program load (IPL)" on page 153. If you specified N for the *Define or change system at IPL* prompt in step 4 of Task 3, skip to "Task 6: Completing the initial program load" on page 155.

Task 5: Defining or changing the system at initial program load (IPL)

You can learn how to define or change the system at IPL.

1. If you specified Y for Enable automatic configuration on the Set Major System Options display, skip to step 3. If you specified N, continue with step 2.
2. If you have chosen not to enable automatic configuration, you must change the QIPLTYPE system value. Follow these steps:
   a) From the Define or Change the System at IPL menu, select option 3 (System value commands). Press the Enter key.
   b) Select option 3 (Work with system values) and press the Enter key.
   c) Type a 2 in the Option column next to the system value QIPLTYPE and press the Enter key.
   d) Change the value to 2 and press the Enter key.
   e) Press F12 until you return to the Define or Change the System at IPL menu.
3. When you are recovering your system, some system values must be set correctly to prevent the recovery from failing.

   Also, if you restore the system from distribution media, your system values reset to the IBM-supplied defaults.

Use whatever documentation you have to set the system values to the correct settings for your installation.

Attention!

If you are restoring to a system with a different processor or memory, you must ensure that the QMCHPOOL, QBASPOOL, and QPFRADJ system values are correct.

As a general rule, if the main storage size is 64 MB or larger, change the QMCHPOOL system value to be 15 percent of the main storage size. If the main storage size is less than 64 MB, change the QMCHPOOL system value to be 20 percent of the main storage size. For a more precise setting of the QMCHPOOL system value, see the Types of memory pools IBM i documentation topic.

The QBASPOOL system value should be equal to 5 to 10 percent of the main storage size.

The QPFRADJ system value should be set to 2.

a. From the Define or Change the System at IPL menu, select option 3 (System value commands) and press the Enter key.
b. Select option 3 (Work with system values) and press the Enter key.

If you are restoring to the same system from your SAVSYS media, skip to 6.

4. Update the following system values.

   Write down the existing values so you can update them after the recovery, if necessary.

   - QALWOBJRST to *ALL
   - QFRCCVNRST to 0
   - QIPLTYPE to 2
   - **QINACTITV to *NONE**
   - QJOBMSGQFL to *PRTWRAP
   - QJOBMSGQMX to minimum of 30 with a suggestion of 64
   - **QLMTDEVSSN to 0**
   - **QLMTSECOFR to 0**
   - QLOGOUTPUT to *JOBLOGSVR
   - **QMAXSIGN to *NOMAX**
   - **QPFRADJ to 2**
   - **QPWDEXPITV to *NOMAX**
   - QSCANFSCTL add *NOPOSTRST
   - QVFYOBJRST to 1

5. If you are restoring to a different system with a different serial number or different logical partitions, and you selected install option 1 (`Take defaults`) on the Install the Operating System menu, the following network attributes are reset to the shipped values.

   If you selected install option 2 (`Change install options`) on the Install the Operating System menu, and selected option 1 (`Restore`) at the **System Information** field on the Specify Restore Options display, the network attributes will be restored.

   - System name
   - Local network ID
   - Local control point name
   - Default local location name
   - Default node
   - Default type
   - Maximum number of intermediate sessions
   - Route addition resistance

- Network node servers
- Alter primary focal point
- Alert default focal point

If you are restoring from distribution media and have previously changed the network attributes from the IBM-supplied defaults, you need to reset them. Follow these steps:

a. From the Define or Change the System at IPL menu, select option 4 (Network attributes commands) and press the Enter key.

b. Select option 2 (Change network attributes). Press the Enter key to display a list of network attributes.

c. Change the values to the correct network attributes and press the Enter key.

d. Press F12 (Cancel) to return to the Define or Change the System at IPL menu.

6. If you are partially restoring (only some libraries), continue with step 7. Otherwise, skip to "Task 6: Completing the initial program load" on page 155.
7. If you are partially restoring and do not plan to restore certain libraries that are listed in the QSYSLIBL and QUSRLIBL system values, you need to edit the QSYSLIBL and QUSRLIBL system values to remove those libraries. Follow these steps:

a) From the Define or Change the System at IPL menu, select option 3 (System value commands). Press the Enter key.

b) Select option 3 (Work with system values) and press the Enter key.

c) Type a 2 in the Option column next to the system values you want to change and press the Enter key.

d) Change the values to the correct values and press the Enter key.

e) Press F12 to return to the Define or Change the System at IPL menu.

Task 6: Completing the initial program load

You can learn how to complete the initial program load (IPL).

1. Press F3 to continue the IPL.
2. The Edit Rebuild of Access Paths display is shown during the IPL process (attended mode) when system access paths are marked for rebuild.

   The "Task 2: Editing the rebuilding of access paths" on page 163 topic describes how to interpret and update this display.

```
                  Edit Rebuild of Access Paths
                                                      05/12/90  13:49:34

IPL threshold . . . . . . . . . . . . . 50   0-99

Type sequence, press Enter.
Sequence:  1-99, *OPN, *HLD

                  ------------Access Paths----------  Unique  Rebuild
Seq   Status      File        Library     Member      Keyed   Time
25__  IPL         QAPZSYM2    QSYS        QAPZSYM2    NO      00:00:01
```

   A status message is sent to notify the user that the system is performing access path recovery.

3. Make any changes and press the Enter key.

   If you have made changes, the Edit Rebuild of Access Paths display is shown again confirming your changes or showing your error messages. Repeat this step until the Display Access Path Status display is shown or the IPL continues.
4. The Display Access Path Status display is updated every 5 seconds while the system is rebuilding access paths.

```
                    Display Access Path Status

IPL Threshold . . . . . : 50

          ----------Access Paths----------   Rebuild     Current
Status    File        Library    Member      Build Time  Run Time
RUN       QAPZSYM2    QSYS       QAPZSYM2    00:00:01    00:00:01
JRN       QAPZREQ2    QSYS       QAPZREQ2    00:00:01
SYS       QASULE03    QSYS       QASULE03    00:00:01
IPL       QASULE01    QSYS       QASULE01    00:00:01
```

If you want to make changes, press F12 (Cancel) to return to the Edit Rebuild of Access Paths display. If all access paths are rebuilt or you no longer want to see the display, press F3 (Exit and continue IPL).

Note: The rebuild of access paths requires system memory. It is suggested that you avoid a high level of system activity so that the rebuild of access paths can complete.

5. The Edit Check Pending Constraints display is shown if referential constraints need to be verified.

   "Task 3: Editing check pending constraints" on page 164 topic describes how to interpret and update this display.

```
                       Edit Check Pending Constraints
                                                          xx/xx/xx  10:09:27
Type sequence, press Enter.
Sequence:  1-99, *HLD

                      ----------Constraints-----------  Verify    Elapsed
Seq   Status           Cst       File        Library     Time      Time

75__  AFTIPL       CSTF1       FILE567890  LIB4567890  00:00:56  00:00:00
```

6. Make any changes and press the Enter key.

   The Edit Check Pending Constraints display is shown again confirming your changes or showing your error messages if you have made changes. Repeat this step until the Display Constraint Status display is shown or the IPL continues.

7. The Display Constraint Status display is updated every 5 seconds while the system is verifying constraints.

```
                    Display Constraint Status

IPL Threshold . . . . . : 50

          ----------Constraints-----------   Verify      Elapsed
Status    Constraint  File       Library     Time        Time
RUN       CUST1       CUSTMAST   CUSTLIB     00:00:04    00:00:01
RUN       CUST2       CUSTMAST   CUSTLIB     00:00:05    00:00:01
IPL       ORDHST1     ORDHIST    ORDLIB      00:00:23    00:00:00
```

   If you want to make changes, press F12 (Cancel) to return to the Edit Check Pending Constraints display. Press F3 (Exit and continue IPL) if all constraints are verified or you no longer want to see the display.

8. Press the Enter key if QSYSOPR messages are displayed.
9. You might receive A900 2000 on the control panel for the system or logical partition, or message CPF0975, `Console did not vary on`, on the console display.

   This occurs if your system configuration was lost and you have disabled automatic configuration. The system has created device description QCONSOLE to allow you to continue the restore operation. You might also receive SRC A900 2000 if you perform an IPL when the QIPLTYPE system value is set to 2. Do not create a user-defined device description for the console display. This can cause unpredictable results.

   If you receive this message, perform the steps that are described in "Recovering from SRC A900 2000" on page 157 before continuing.

10. If you restored from the distribution media by using a 1/4-inch cartridge tape drive, the light on the tape drive might still be on.

After the system has finished restoring the operating system, you can remove the tape while the light is on.

Recovering from SRC A900 2000

When you restore the operating system, you might see SRC A900 2000. This happens when you use option 2 or option 3 to install the Licensed Internal Code and automatic configuration is not active while you are restoring the operating system.

Before you can continue your recovery operations, you must create a device description and possibly a controller description to finish the restore operation. Do not create a user-defined device description for the console display.

When the Sign On display appears, you have completed restoring the operating system. Consult your recovery checklist for the next step in your recovery process.

Creating a configuration for other tape units

If you are not using a 34*xx* tape unit, or want to create a 34*x*x (3490 Model E or F) tape unit without a controller, you need to know how to create a configuration for other tape units.

To create a configuration for other tape units, follow these steps:

1. Use the Work with Hardware Resource (WRKHDWRSC) command to determine tape controller name.

   ```
   WRKHDWRSC TYPE(*STG)
   ```

2. Locate the tape controller on the Work with Hardware Resources display.
3. Type a 9 (Work with resource) next to tape controller name and press the Enter key.
4. Locate the resource name for the tape unit (for example, TAP01).
5. Enter a 5 (Work with Configuration Descriptions) in the *Opt* column next to the tape resource name and press the Enter key.

   You are shown the Work with Configuration Descriptions display.
6. Type a 1 (Create) in the **Option** field and a device description name (for example, TAP01) in the **Description** field and press the Enter key.

   You are shown the Create Device Description display.
7. Change any values that you want to change, then press the Enter key (twice) to create the device description.

   You are shown the Work with Configuration Descriptions display again. The device that you created should appear on the display.
8. Type an 8 (Work with configuration status) in front of the new device description.

   You are shown the Work with Configuration Status display.
9. Type a 1 (Vary on) in front of the new device.

   If the status does not change to Varied On, wait a few minutes. Then press F5 (Refresh). If the status still does not change to Varied On, follow normal problem determination procedures for the device.
10. Press F3 until you return to the main menu.

SRC A900 2000 remains displayed on the control panel for the system or logical partition throughout the remaining restore operations. When the final IPL of the system is complete, SRC A900 2000 disappears. The user-defined device description for the console display will be restored when the Restore Configuration (RSTCFG) command is run later in the recovery.

Chapter 6. Starting the system after it ends abnormally

Use this information to start the system after it ends abnormally.

When your system stops normally, it does the following actions:

- Writes changed pages of information from memory to auxiliary storage.
- Closes access paths and files.
- Ends programs and jobs at the natural stopping points.

If your system stops without having time to do these things, it is called an *abnormal end*. Your system might end abnormally for the following reasons:

- A power failure.
- A disk failure, if you do not have mirrored protection or device parity protection.
- A failure in the processor.
- Failure of a critical operating system program.
- Operator action (forced power down).

Related concepts

Recovery procedure for a system failure
A system failure is a problem either with hardware (other than disk) or with operating system software that causes your system to end abnormally.

What happens when your system stops

This list describes the circumstances that cause your system to stop unexpectedly and what happen when your system stops.

- **Power failure with uninterruptible power supply:** When the system loses normal power, the uninterruptible power supply system takes over and keeps the system running. The system detects this change and sends a message to your power-monitoring program. Your program can decide whether to keep the system running until power returns or to begin an orderly shutdown.
- **Power failure with no protection:** If your system does not have an uninterruptible power supply feature and the power fails, your system stops immediately. The contents of main memory are lost. The system must reconstruct information when power returns. This can be very time-consuming. Whether the system starts automatically depends on how you have set the QPWRRSTIPL system value.
- **Disk failure with device parity protection or mirrored protection:** In many cases, the system can continue running without full disk protection until the failed unit is replaced.
- **Disk failure without disk protection:** This is like a power failure without protection. The system stops immediately. The system must reconstruct information about jobs that were running and files that were open after the disk is repaired or replaced.
- **Failure of a critical operating system program:** The system will stop immediately, just as it does if an unprotected power failure or disk failure occurs. The system attempts to copy the contents of main memory so that the problem can be analyzed. This is called a *main storage dump*. When the system stops, you will see the Main Storage Dump Manager Occurred display.

© Copyright IBM Corp. 1997, 2025

Using the Disk Configuration Error Report display

When your system starts, it checks to ensure that it can access all of the disk units that are configured. If it cannot access one or more disk units, you are shown the Disk Configuration Error Report display.

```
                    Disk Configuration Error Report

Type option, press Enter.
  5=Display Detailed Report

OPT  Error
 _   Missing disk units in the configuration
```

Following a temporary power outage, you might see this display because power has been restored to the processor but not to the peripheral devices. Wait to respond to the display until power is restored to all the disk units. The system's ability to access all the disk units when the system is starting is important for a successful recovery. If disk units are not available, the system might not be able to recover changed pages of memory. This can lengthen the time it takes to perform the initial program load (IPL).

This display might also be presented:

- After abnormal termination if the system is unable to activate all of the disk on the re-IPL.
- During any system IPL that has a similar problem, even if normal system shutdown had taken the system down last.

Using the Main Storage Dump Occurred display

If your system encounters a serious software problem, you are shown the Main Storage Dump Occurred display.

```
                     Main Storage Dump Occurred
                                                          S/N xxxxxxxx

The system has failed.  Report the following information to
your IBM service representative.

              SRC word 1 . . . . . . . :   A1D03000
              SRC word 2 . . . . . . . :   69B0015F
              SRC word 3 . . . . . . . :   0000308F
              SRC word 4 . . . . . . . :   3FFFDE00
              SRC word 5 . . . . . . . :   0C211008
              SRC word 6 . . . . . . . :   00000000
              SRC word 7 . . . . . . . :   00000000
              SRC word 8 . . . . . . . :   00D5A400
              SRC word 9 . . . . . . . :   00CDA400
              Type/Model/Feature . . . :   xxxx xxxx xxxx

Warning: The Main Storage Dump (MSD) must be copied for service.
         Failure to copy the Main Storage Dump will limit
         the ability to diagnose the failure.

Press Enter to copy the MSD for service or view the MSD.

F3=Exit     F12=Cancel
```

Follow the instructions of your service provider in responding to this display. In most cases, you should make a copy of the main storage dump. Save it either to save media or to auxiliary storage (disk), to assist with diagnosing the problem.

Related information

Performing a main storage dump

Restarting your system

When you have solved whatever problem caused your system to stop, you must start it again.

In some cases, you start the initial program load (IPL) yourself. In other cases, such as a power loss, the system starts automatically. When you start your system again after it ends abnormally, the system tries to put things back in order. It closes files that were in use, rebuilds access paths that were open, and verifies file constraints. This process can take a long time.

If you want the system to determine when to rebuild and verify, perform a normal (automatic) IPL to restart your system. If you want to view and change the schedules for rebuilding access paths and verifying referential constraints, complete the tasks in this topic collection.

Task 1: Performing an attended initial program load

You can do an attended initial program load (IPL) so that you have the opportunity to change rebuild options.

Note: If your service representative already started the IPL, skip to the step in this task for the display that is shown on your system.

To do an attended IPL, you must use the control panel on the system unit or logical partition. The steps vary slightly based on the type of system unit that you have. If you are unsure of the procedures for your system, see Basic system operations for more information.

To do an attended IPL, follow these steps:

1. Place the system in Manual mode.
2. Verify that any switches for all disk units are in the On position.
3. If your system is running, verify that all users are signed off and all jobs are ended.

   Type the following command:

   PWRDWNSYS OPTION(*CNTRLD) DELAY(600) RESTART(*YES)

   For the delay parameter, specify the number of seconds to allow your system time to bring most jobs to a normal end. On a large, busy system, you might need a longer delay.
4. If your system is not running, power on your system.
5. When you see the IPL or Install the System display, select option 1 (Perform an IPL).

   An example of a status display. The status displays do not require any action by the user.

```
                        Licensed Internal Code IPL in Progress

IPL:
   Type . . . . . . . . . . . . . . . :    Attended
   Start date and time  . . . . . . . :    xx/xx/xx  xx:xx:xx
   Previous system end  . . . . . . . :    Abnormal
   Current step / total . . . . . . . :       16      16
   Reference code detail  . . . . . . :    C6004065

Ipl step                                   Time Elapsed   Time Remaining
  Commit Recovery                            xx:xx:xx        xx:xx:xx
  Data Base Initialization                   xx:xx:xx        xx:xx:xx
  Journal IPL Clean up                       xx:xx:xx        xx:xx:xx
  Commit Initialization                      xx:xx:xx        xx:xx:xx
 >Start the operating system                 xx:xx:xx        xx:xx:xx

Item:
  Current / Total . . . . . . . :

Sub Item:
  Identifier . . . . . . . . . . :    Searching tape for install media
  Current / Total . . . . . . . :

Wait for next display or press F16 for DST main menu
```

The following list shows some of the IPL steps that are shown on the Licensed Internal Code IPL in Progress display:

- Authority Recovery
- Journal Recovery
- Database Recovery
- Journal Synchronization
- Start the Operating System

Some of the IPL steps might take a long time.

During the IPL, system reference codes (SRCs) are displayed on the control panel of the system unit to indicate what step is in progress. If the same SRC is displayed for a long time in solid (not flickering) lights, your system might have a problem with the IPL.

6. Press the Entered key. Informational messages are displayed.
7. If the **Select Product to Work with PTFs** display appears, press F3 (Exit) to continue.

```
                Select Product to Work with PTFs

Position to . . . . . .______________ Product


Type options, press Enter.  Press F21 to select all.
   1=Select

              Product
Opt  Product  Option  Release
 _    5770999  *BASE   V7R1M0
 _    5770SS1  *BASE   V7R1M0
```

8. You are shown the IPL Options display:

```
                              IPL Options

Type choices, press Enter.

System date  . . . . . . . . . . . . . .   xx / xx / xx
System time  . . . . . . . . . . . . . .   12 : 00 : 00
System time zone . . . . . . . . . . . .   QN0600UTCS
Clear job queues . . . . . . . . . . . .   N
Clear output queues  . . . . . . . . . .   N
Clear incomplete job logs  . . . . . . .   N
Start print writers  . . . . . . . . . .   Y
Start system to restricted state . . . .   N


Set major system options . . . . . . . .   Y
Define or change system at IPL . . . . .   Y
```

The values that appear as defaults depend on the recovery steps you used.

9. If the system date, time, and time zone are not correct, type the correct values.

   If you installed the Licensed Internal Code by using option 2 or option 3, the date and time might be blank. The system date must have a year value in the range of 87 to 99, or 00 to 22.

10. Specify these responses for the prompts on the display:

Clear job queues	N
Clear output queues	N
Clear incomplete job logs	N
Start print writers	Y
Start system to restricted state	Y
Set major system options	N
Define or change system at IPL	N

11. Enter your choices and press the Enter key.

Task 2: Editing the rebuilding of access paths

Use this information to edit the rebuilding of access paths if they are marked for rebuilding.

If access paths are marked for rebuilding, you are shown the following display:

```
                  Edit Rebuild of Access Paths                  SYSTEMA
                                                      05/12/90  13:49:34

IPL threshold . . . . . . . . . . . . . 50   0-99

Type sequence, press Enter.
Sequence:  1-99, *OPN, *HLD

                  ------------Access Paths----------  Unique  Rebuild
Seq   Status      File        Library     Member      Keyed   Time
25__  IPL         QAPZSYM2    QSYS        QAPZSYM2    NO      00:00:01
25__  IPL         QAPZREQ2    QSYS        QAPZREQ2    NO      00:00:01
25__  IPL         QAPZPTF3    QSYS        QAPZPTF3    NO      00:00:01
25__  IPL         QAPZPTF2    QSYS        QAPZPTF2    NO      00:00:01
25__  IPL         QAPZOBJ2    QSYS        QAPZOBJ2    NO      00:00:01
*OPN  OPEN        QTWALL      QSYS        QTWALL      NO      00:00:06
*OPN  OPEN        QASULTEL    QSYS        QASULTEL    NO      00:00:01
*OPN  OPEN        QASULE05    QSYS        QASULE05    NO      00:00:01
*OPN  OPEN        QASULE03    QSYS        QASULE03    NO      00:00:01
*OPN  OPEN        QASULE01    QSYS        QASULE01    NO      00:00:01
More...
F5=Refresh   F11=Display member text   F13=Change multiple  F15=Sort by
F16=Repeat position to   F17=Position to
```

Note: No access paths are listed if all access paths that are marked for rebuild have a status of SYS, JRN, or SMAPP.

While you are working with this display, the system is rebuilding access paths. You can use this display to perform the following tasks:

- Change the sequence in which access paths are rebuilt.
- Delay rebuilding some access paths until after the initial program load (IPL).

1. If you do not want to make changes to this display, press the Enter key.

   Skip to step 5. If you want to make changes, continue with step 2.
2. You can change the value of the IPL threshold.

   All access paths with a sequence (SEQ) less than or equal to the threshold are rebuilt during the IPL. Access paths with a greater sequence number are rebuilt after the IPL completes. The default threshold is 50.
3. You can change the sequence (SEQ) column on the display for specific access paths.

   Initially, the sequence numbers are set in the following way:

   25
   Files with MAINT(*IMMED) and RECOV(*IPL)

   75
   Files with MAINT(*IMMED) and RECOV(*AFTIPL)

   ***OPN**
   Files with MAINT(*DLY)

   Within a group (same sequence numbers), the system rebuilds access paths according to rebuild time, starting with the longest rebuild time.

   Rebuild time is an estimate, based on the file size and key length. For journaled access paths (status JRN) and access paths protected by system-managed access-path protection (status SMAPP), the rebuild time shows as 0. The system uses the journal entries to recover these access paths rather than rebuilding them. The time that is required is minimal.

   The estimate for rebuild time assumes that the rebuild job is not contending for resources. If an access path is rebuilt after the IPL, the rebuild will probably take longer.

4. Type your changes and press the Enter key.

   You are shown the Edit Rebuild of Access Paths display again. You see error messages if the system cannot make some of the changes you requested. For example, you might have tried to change the sequence number for an access path that the system rebuilt while you were using the display.

   If you have errors, return to step 2.
5. When you have finished the display, press the Enter key without making changes.

   You are shown the Display Access Path Status display:

```
                    Display Access Path Status

IPL Threshold . . . . . : 50

          ----------Access Paths----------   Rebuild     Current
Status    File        Library    Member      Build Time  Run Time
RUN       QAPZSYM2    QSYS       QAPZSYM2    00:00:01    00:00:01
JRN       QAPZREQ2    QSYS       QAPZREQ2    00:00:01
SYS       QASULE03    QSYS       QASULE03    00:00:01
IPL       QASULE01    QSYS       QASULE01    00:00:01
```

   This display is updated every 5 seconds while the system is rebuilding access paths.
6. If you want to make changes to the IPL threshold or the sequence for rebuilding access paths, press F12 (Cancel) to return to the Edit Rebuild of Access Paths display.

   Repeat steps 2 through 5.

   If you do not want to make changes, you can allow the Display Access Path Status display to continue updating or you can press F3 (Exit and continue IPL). In either case, the system completes rebuilding access paths before continuing to the next step of the IPL.

Task 3: Editing check pending constraints

You can define required attributes for physical files on the system. The descriptions for these values are referential constraints or other types of constraints.

When you perform an initial program load (IPL) after the system ends abnormally or when you restore database files, the system checks the validity of file constraints. For more information, see Constraints.

If database constraints are marked for verification, you are shown the following display:

```
                      Edit Check Pending Constraints               SYSTEMA
                                                        03/30/94  10:09:27
IPL threshold . . . . . . . . . . . . . 50_  0-99

Type sequence, press Enter.
Sequence:  1-99, *HLD

                       ----------Constraints----------  Verify    Elapsed
Seq   Status           Cst       File        Library     Time      Time

75__  AFTIPL       CSTF1       FILE567890  LIB4567890  00:00:56  00:00:00
75__  AFTIPL       CSTF2       FILE567890  LIB4567890  00:00:56  00:00:00
*HLD  INVAP        CSTF5       FILE567890  LIB4567890  10:30:06  00:00:00
*HLD  CHKPND       CSTF6       FILE567890  LIB4567890  09:30:06  00:00:00
*HLD  HELD         CSTF7       FILE567890  LIB4567890  08:30:06  00:00:00
                                                          More...
```

You can use this display to do the following tasks:

- Change the sequence in which constraints are verified.
- Have the system verify some constraints before the IPL completes.
- Delay verification for some constraints until you specifically request it.

1. If you do not want to make changes to this display, press the Enter key and skip to step 5.

   If you want to make changes, continue with step 2.
2. You can change the value of the IPL threshold.

All constraints with a sequence (SEQ) less than or equal to the threshold are verified during the IPL. Constraints with a greater sequence number are verified after the IPL completes. The default threshold is 50.

3. You can change the sequence (SEQ) column on the display for specific constraints.

   Initially, the sequence for all constraints is set to 75.

   Within a group (same sequence numbers), the system verifies constraints according to verify time, starting with the longest estimated time.

   Verify that time is an estimate. The estimate assumes that the verify job is not contending for resources. If a constraint is verified after the IPL, the verify might take longer.

   If you want to delay the verification of a constraint indefinitely, specify *HLD for the sequence. You can use the Edit Check Pending Constraint (EDTCPCST) command later to set a sequence and have the system verify the constraint.

4. Type your changes and press the Enter key.

   You are shown the Edit Check Pending Constraint display again. You see error messages if the system cannot make some of the changes you requested. For example, you might have tried to change the sequence number for a constraint that the system verified while you were using the display.

   If you have errors, return to step 2.

5. When you have finished the display, press the Enter key without making changes.

   You are shown the Display Constraint Status display:

```
                    Display Constraint Status

IPL Threshold . . . . . : 50

          ----------Constraints-----------   Verify      Elapsed
Status    Constraint  File       Library     Time        Time
RUN       CUST1       CUSTMAST   CUSTLIB     00:00:04    00:00:01
RUN       CUST2       CUSTMAST   CUSTLIB     00:00:05    00:00:01
IPL       ORDHST1     ORDHIST    ORDLIB      00:00:23    00:00:00



F3=Exit and continue IPL   F12=Cancel
```

   This display is updated every 5 seconds while the system is verifying constraints.

6. If you want to make changes to the IPL threshold or the sequence for verifying constraints, press F12 (Cancel) to return to the Edit Check Pending Constraints display.

   If you want to make changes to the IPL threshold or the sequence for verifying constraints, press F12 (Cancel) to return to the Edit Check Pending Constraints display. Repeat steps 2 through 5.

   If you do not want to make changes, you can allow the Display Constraint Status display to continue updating or you can press F3 (Exit and continue IPL). In either case, the system completes verifying constraints before continuing to the next step of the IPL.

7. When the IPL completes, continue with "Task 4: Recovering from damaged objects and unreadable sectors" on page 165.

Task 4: Recovering from damaged objects and unreadable sectors

If your system stops without warning or disk errors occur, some object description information might not be updated correctly. If this happens, the object is considered damaged.

When you perform an initial program load (IPL), the system attempts to locate damaged objects and write them to the object recovery list. It writes a message (CPI18xx) to the history (QHST) log for each damaged object that it finds. If any damaged objects are added to the object recovery list during the IPL, message CPI8197 is sent to the QSYSOPR message queue.

Note: Some damage, such as damage to the contents of a database file, might not be detected until the object is used. If you suspect that a large number of objects on your system have been damaged, contact your service representative for advice on how to recover.

To check and recover damaged objects, follow these steps:

1. Display the QHST (history) log by typing DSPLOG and pressing F4 (Prompt).
2. Display the QHST (history) log by typing DSPLOG and pressing F4 (Prompt).
3. On the display, type *PRINT for the *Output* prompt and press the Enter key.
4. Type: WRKSPLF.

   You are shown a list of spooled files for your job.
5. Locate the spooled file for the DSPLOG command.

   Use option 3 to hold the spooled file.
6. Use option 5 to display the spooled file.
7. Look for entries for damaged objects that are not synchronized.

   You can use the Find function to search for lines that have these character strings: damage and sync. Following are some examples of messages you might see:

   CPF3113
   Member damaged

   CPF3175
   File is not synchronized

   CPF3176
   Data space is partially damaged

   CPF3171
   Journal is damaged

   CPF3172
   Objects are not synchronized with the journal

   CPF3173
   Journal receiver is damaged

   CPF3174
   Journal receiver is partially damaged

   CPF700C
   Object of type *object-type cannot be synchronized with journal.

   CPF81xx
   General messages about object damage
8. Write down the names and types of the objects you find.

   Consult Table 41 on page 166 for the correct recovery procedure, based on the type of object that is damaged.

Table 41. Recovery for damaged objects by object type

Type of object	Recovery procedure
Operating system object in QSYS library	Contact software support for assistance. You might need to install the operating system again.
IBM-supplied user profile	Perform an abbreviated installation of the operating system.
Job description that is specified on the workstation entry for the console in the controlling subsystem	If no other workstation entries exist for the controlling subsystem, the system is not usable. Contact software support for assistance.
Job queue	Perform an IPL. Restore or re-create the damaged job queue. All entries are lost.

Table 41. Recovery for damaged objects by object type (continued)

Type of object	Recovery procedure
Output queue	Perform an IPL. If the output queue is the default output queue for a printer, it is re-created and its entries are rebuilt. Other output queues must be restored or re-created. Their entries are not recovered.
Damaged file whose name starts with QA0SS	Delete the file. Restore it from a backup copy. Run the RCLDLO DLO(*DOCDTL) command.
Journal	See "Recovering a damaged journal" on page 167.
Journal receiver	See "Recovering a damaged journal receiver" on page 168.
Journaled object	See "Recovering a journaled object that is damaged or not synchronized" on page 169.
All others	See "Recovering other types of damaged objects" on page 170.

9. Watch for additional indications that objects have been damaged.

   Here are some indications:

   - You cannot start the system because auxiliary storage is full.
   - The system has ended abnormally several times since the last time you ran the Reclaim Storage (RCLSTG) procedure.
   - You see objects on the Work with Objects by Owner display that have no library associated with them.
   - The system status display shows an unexpectedly high percentage of auxiliary storage that is used.
   - You cannot access the data in a database file because a member is damaged. Message CPF8113 indicates this.
   - You cannot access objects because a damaged authorization list or authority holder secures them.

   If you see these indications on your system, the following actions can help you identify damaged objects:

   a. Choose from the following actions to identify damaged objects in critical files where you suspect damage:

      - Use the Copy File (CPYF) command
      - Use the Display Object Description (DSPOBJD) command
      - Perform a save of your critical data

   b. Choose from the following actions to identify damaged objects at the system level:

      - Use the Retrieve Disk Information (RTVDSKINF) command
      - Use the Display Object Description (DSPOBJD) command and specify DSPOBJD OBJ(*ALL) OBJTYPE(*ALL)
      - Perform a full system save using GO SAVE option 21
      - Run the RCLSTG procedure. Running the procedure is described in "Reclaim Storage (RCLSTG) command" on page 28 topic.

   If you see these indications after a disk unit was replaced and the data was restored from a partial pump, you should recover the entire ASP that contained the failed disk unit. See the appropriate checklist.

Recovering a damaged journal

To recover a damaged journal, perform these steps.

1. Type WRKJRN.
2. On the prompt display, type the name of the journal.

   You are shown the Work with Journals display:

```
                            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
     JRNACC      DSTA1                  JOURNAL FOR ACCOUNTS
```

3. Select option 6 (Recover damaged journal).
4. Type: WRKJRNA JRN(*library-name*/*journal-name*) OUTPUT(*PRINT).

   You receive a listing that shows all the objects that are currently being journaled.
5. Start journaling for any physical files that should be journaled but are not on the list by using the Start Journal Physical File (STRJRNPF) command.
6. Start journaling for any access paths that should be journaled but are not on the list by using the Start Journal Access Path (STRJRNAP) command.
7. Start journaling for any integrated file system objects that should be journaled but are not on the list by using the Start Journal (STRJRN) command.
8. Start journaling for any libraries that should be journaled but are not on the list using the Start Journal Library (STRJRNLIB) command.
9. Start journaling for any other object types that should be journaled but are not on the list by using the Start Journal Object (STRJRNOBJ) command.
10. If you previously had any remote journals that were associated with the damaged journal, add those remote journals again.

    You can add the remote journals by using the Add Remote Journal (QjoAddRemoteJournal) API or the Add Remote Journal (ADDRMTJRN) command.
11. Save each journaled object. You should always save objects after you start journaling them.

Recovering a damaged journal receiver

This topic provides information about how to recover a damaged journal receiver that was attached to a local journal.

Note: The following steps only apply to recovering a damaged journal receiver that was attached to a local journal.

1. Type WRKJRN.
2. On the prompt display, type the name of the journal that is associated with the damaged journal receiver.

   You are shown the Work with Journals display:

```
                            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
     JRNACC      DSTA1                  JOURNAL FOR ACCOUNTS
```

3. Select option 7 (Recover damaged journal receivers).

Recovering a journaled object that is damaged or not synchronized

For a journaled object that is damaged, either select option 2=Forward Recovery from the Work with Journals display (WRKJRN command), or perform these steps.

1. Find your most recent saved copy of the object.
2. Delete the object.
3. Load the save media and restore the object.
   a. For journaled database physical files, data areas, or data queues, type the following command:

```
RSTOBJ OBJ(object-name) OBJTYPE(*object-type)
       SAVLIB(library-name)
       DEV(media-device-name)
```

   b. For journaled libraries, type the following command:

```
RSTLIB SAVLIB(library-name)  DEV(media-device-name)
```

   c. For journaled integrated file system objects, type the following command:

```
RST DEV('device-path-name')
    OBJ('object-path-name')
    SUBTREE(*ALL)
```

4. Restore any journal receivers that are needed to recover the object, if they are not already on the system.
5. Use the Apply Journaled Changes (APYJRNCHG) command to apply journaled changes.

Perform the following steps for a journaled object that cannot be synchronized:

1. Restore the object from your most recent saved copy.
2. Apply journaled changes to bring the object up to date.

   Note: You might need to perform this procedure for all the objects if there are other objects related to the object that is not synchronized. Otherwise, the objects might not be synchronized with each other.

Related information

Journal management

Recovering damaged objects in the integrated file system

To recover damaged objects in the integrated file system, run the Reclaim Storage (RCLSTG) command. To recover damaged objects without putting your system in a restricted state, you can use the Reclaim Object Links (RCLLNK) command. Perform these recovery steps if the object was not journaled.

If you are unable to recover your integrated file system objects using the RCLSTG or RCLLNK commands, use the Restore Object (RST) command to restore the objects.

```
RST  DEV('/QSYS.LIB/media-device-name.DEVD')
     OBJ('/mydir/myfile')
```

If the damaged integrated file system object is journaled, follow the steps for recovering a damaged journaled object.

Related concepts

Reclaim Storage (RCLSTG) command

You can use the RCLSTG command to recover the addressability of lost or damaged objects. This allows you to identify and then restore those objects that were damaged.

Reclaiming object links

The Reclaim Object Links (RCLLNK) command reclaims objects in the "root" (/), QOpenSys, and mounted user-defined file systems (UDFSs).

Recovering other types of damaged objects

Use these procedure to recover most damaged objects on the system.

Table 41 on page 166 shows which types of objects require special procedures.

1. Find your most recent saved copy of the damaged object.

   Note: If the damaged object is in the QSYS library, you might need to restore the operating system. Contact software support for assistance.
2. Delete the object.
3. Load the save media and restore the object.

   Type the following command:

```
RSTOBJ OBJ(object-name)
       OBJTYPE(object-type)
       SAVLIB(library-name)
       DEV(media-device-name)
```

Chapter 7. Recovering information in a user auxiliary storage pool

When you have user auxiliary storage pools (ASPs) on your system, you assign specific libraries or objects to certain physical disk devices. One reason for having user ASPs is to limit the amount of information you need to recover if a disk device must be replaced.

The basic process for recovering a user ASP is:

1. Understand what was in the user ASP.
2. Choose the correct recovery procedure.
3. Plan your recovery.
4. Perform the recovery steps.

Describing the contents of your user auxiliary storage pools

To choose the correct procedure for recovering the information on your user auxiliary storage pools (ASPs)s, you must understand what they looked like before the failure.

Figure 4 on page 171 shows an example of a user ASP configuration. This example is used throughout the explanations that follow. You might want to begin by drawing a similar picture of your configuration.

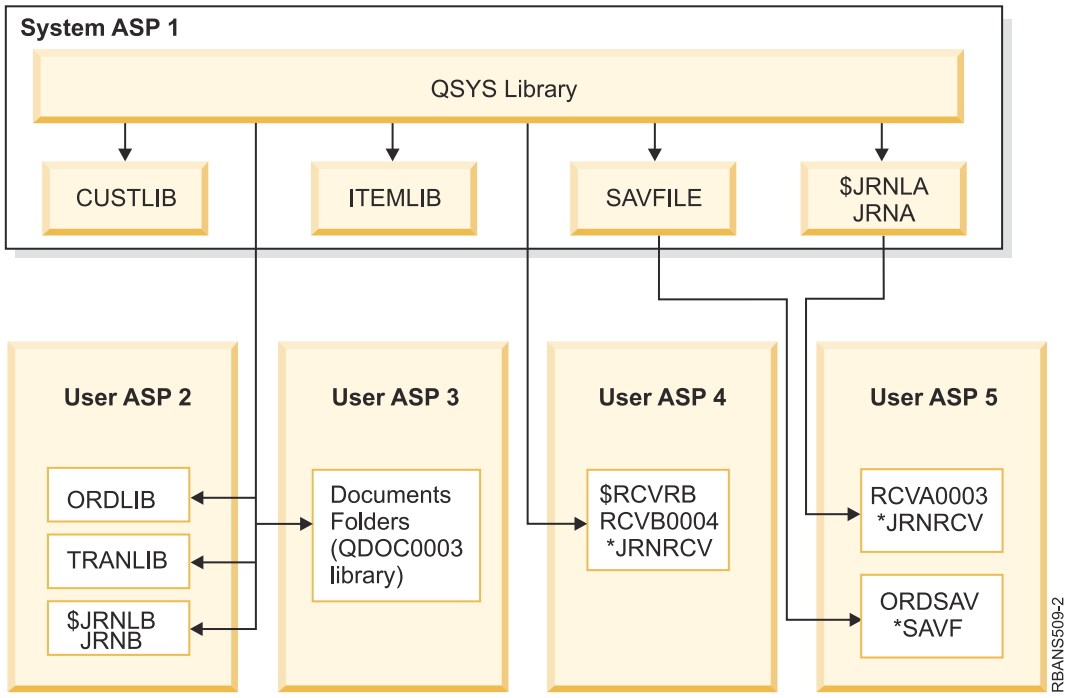

Figure 4. User ASP configuration before failure

In the example:

- ASP 2 is a library user ASP. It contains these libraries: ORDLIB, TRANLIB, and $JRNLB.
- Files in the ORDLIB library and TRANLIB library are journaled to the JRNB journal in the $JRNLB library.
- The journal receivers for the JRNB journal are in the $RCVRB library in ASP 4.
- ASP 3 is a library user ASP containing documents and folders.
- ASP 4 is a library user ASP. It contains the $RCVRB library.

© Copyright IBM Corp. 1997, 2025

- ASP 5 is a nonlibrary user ASP. It contains the ORDSAV save file. ORDSAV is in the SAVFLIB library, which is in the system ASP. It also contains the RCVA0003 journal receiver, which is in the $JRNLA library.

Before the failure, the receiver directory for the JRNA journal looks as follows:

```
                        Work with Receiver Directory

Journal  . . . . . :   JRNA            Library  . . . . . :   $JRNLA

Total size of receivers (in kolobytes). . . . . . . . . . . :   155648

Type options, press Enter.
  4=Delete   8=Display attributes

                                      Attach                        Save
Opt  Receiver    Library     Number   Date      Status              Date
 _   RCVA0001    $JRNLA      00001    06/08/0x  SAVED               06/08/9x
 _   RCVA0002    $JRNLA      00002    06/09/0x  SAVED               06/09/9x
 _   RCVA0003    $JRNLA      00003    06/09/0x  ATTACHED            00/00/00
```

Choosing the procedure to recover user auxiliary storage pools (ASPs)

Here are the basic situations that might require you to recover information in a user ASP.

- You have replaced a disk unit in the system ASP. Although the data is still in your user ASPs, you must recover the system's ability to locate that data (addressability). This process is described in “Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP” on page 172.
- You have replaced a disk unit in a basic user ASP. All the information that was in the user ASP must be recovered. If this is your situation, follow the procedure described in “Recovering a damaged basic auxiliary storage pool” on page 186.
- You have replaced a disk unit in the system ASP. One of your basic user ASPs was in overflowed status. You must recover the addressability to information in the user ASPs that were not overflowed by using the procedure described in “Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP” on page 172. You must also recover the information in the user ASP that was overflowed by using the procedure described in “Recovering a damaged basic auxiliary storage pool” on page 186.
- The service representative has replaced a failed disk unit in an independent ASP. When you lose a disk unit in an ASP, you must recover all of the information in that ASP. The information in other ASPs on your system is not affected. Follow the procedure described in “Recovering an independent auxiliary storage pool” on page 192.

Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP

After restoring your Licensed Internal Code and the operating system, you need to perform the tasks in this topic.

When you replace a unit in your system ASP, the system loses addressability to the objects in your basic user ASPs. The system in the example will look as follows after the operating system has been restored:

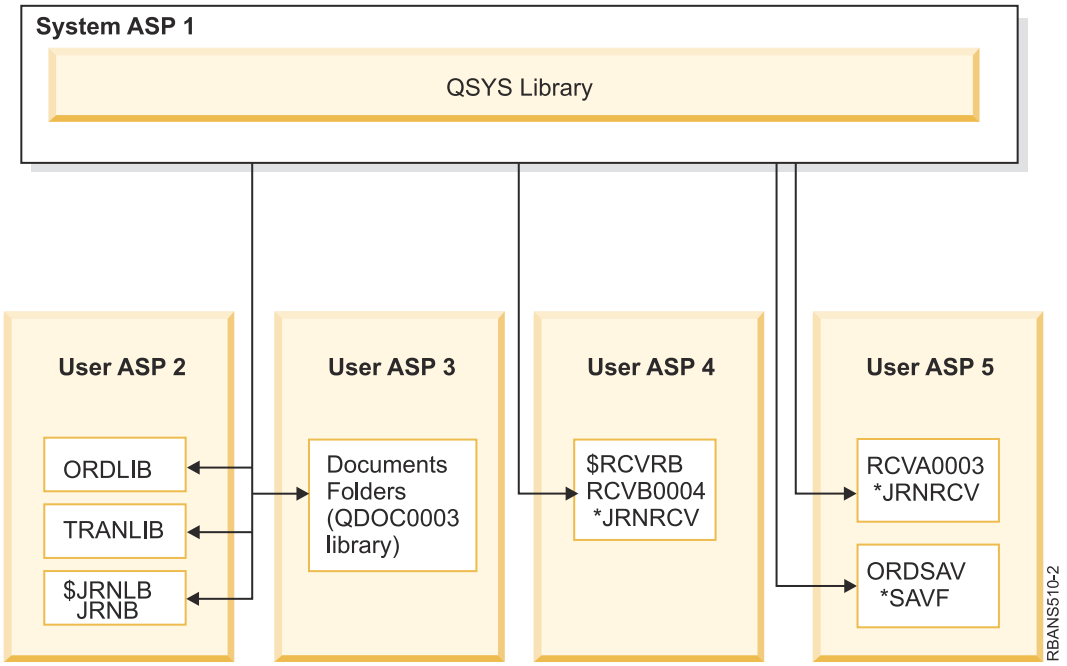

Figure 5. Basic user ASP configuration after restoring operating system

The libraries and objects in the basic user ASPs are not known to the system.

Use the following procedures described in this topic to recover the objects in your basic user ASPs. However, the system cannot recover ownership to the objects other than document library objects (DLO) in the basic user ASPs because the addresses for all user profiles are changed when you restore them. All object types except DLOs use the address of the user profile to identify the owner.

Recovering object ownership for objects other than DLOs requires manually assigning ownership for every object in every basic user ASP.

Related tasks

Restoring encrypted auxiliary storage pools
If you have an encrypted user or independent auxiliary storage pool (ASP), you must run special steps to verify that the data in these ASPs can be recovered.

Task 1: Reclaiming storage

Use these steps to reclaim storage.

1. Sign on the system with a user profile that is authorized to the Reclaim Storage (RCLSTG) command. Either sign on at the console or use the Transfer Job (TFRJOB) command to transfer your job to the controlling subsystem.
2. Type `DSPSYSVAL QALWUSRDMN`. If the current value does not include the QRCL (Reclaim Storage) library or does not specify *ALL, use the CHGSYSVAL command to add QRCL to the list of libraries for this system value. Write the current value here: _____
3. Type `DSPSYSVAL QCTLSBSD` to display the name of your controlling subsystem. Write the current value here: _____
4. Ensure your system is in a restricted state. If it is not, follow the procedure in “Putting your system in a restricted state” on page 27 topic.
5. Use these commands to reclaim storage in various ways:

RCLSTG	Reclaim storage of the entire system.
RCLSTG SELECT(*DBXREF)	Reclaim storage of the database cross-reference table.
RCLSTG SELECT (*DIR)	Reclaim storage of the directory recovery portion.

RCLSTG OMIT(*DBXREF)	Reclaim storage of the entire system except the database cross-reference table.
RCLSTG ASPDEV(*SYSBAS)	Reclaim the system ASP and all basic ASPs.
RCLSTG OMIT(*DIR)	Reclaim storage of the entire system except the directory recovery portion.

6. Use the Change System Value (CHGSYSVAL) command to set the QALWUSRDMN system value back to its original setting. (You wrote the setting in step 2.)
7. When the reclaim storage procedure finishes, start the controlling subsystem by typing the following command:

```
STRSBS SBSD(controlling-subsystem)
```

   (You wrote the name of the controlling subsystem in step 3.)

After the reclaim storage procedure, the example system looks as follows:

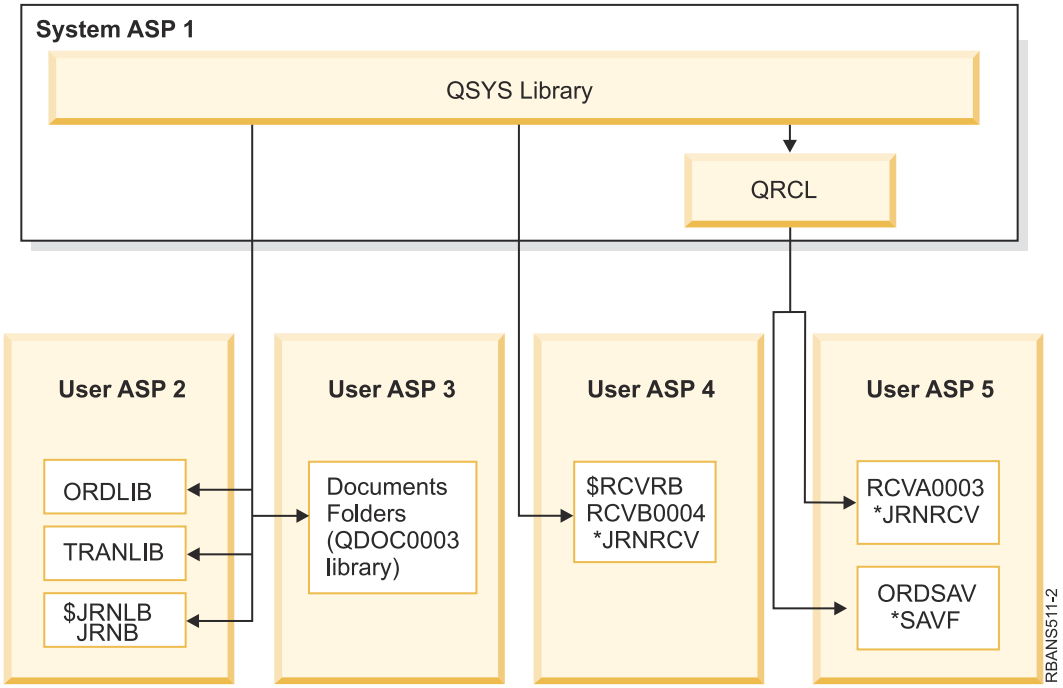

Figure 6. User ASP configuration after reclaiming storage

The system recovers addressability to the objects in ASP 5, but it cannot recover their original library assignments. They are placed in the QRCL (Recovery) library. The objects in all user ASPs are owned by the QDFTOWN (Default Owner) user profile.

Related concepts

Reclaim Storage (RCLSTG) command

You can use the RCLSTG command to recover the addressability of lost or damaged objects. This allows you to identify and then restore those objects that were damaged.

Task 2: Restoring user profiles

Use these steps to restore user profiles.

1. Sign on as QSECOFR.
2. In a clustered environment, use the End Cluster Resource Group (ENDCRG) command and the End Cluster Node (ENDCLUNOD) command to stop clustering before ending subsystem QSYSWRK and continuing to restore the independent disk pool.

3. Ensure the system is in a restricted state.
4. Find the most recent save media that has your user profiles.

   It might be a SAVSYS media volume or a SAVSECDTA media volume. The file on the save media volume is called QFILEUPR.
5. If you are using a SAVSYS media volume, type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*LEAVE)
```

   If you are using a SAVSECDTA media volume, type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*UNLOAD)
```

The time that this takes can vary significantly. "What happens when you restore user profiles" on page 213 topic describes what the system does when you restore user profiles.

Related tasks

Putting your system in a restricted state
Many recovery procedures require that your system have no other activity on it. When no subsystems except the controlling subsystem are active on your system, it is in a *restricted state*.

Task 3: Restoring the configuration

Use these steps to restore the configuration.

1. Find the most recent save media that has your configuration.

   It might be a SAVSYS media volume or a SAVCFG media volume. The file on the save media volume is called QFILEIOC.
2. If you are using a SAVSYS media volume, type the following command:

```
RSTCFG OBJ(*ALL) DEV(media-device-name)
       OBJTYPE(*ALL)
       ENDOPT(*LEAVE)
```

   If you are using a SAVCFG media volume, type the following command:

```
RSTCFG OBJ(*ALL) DEV(media-device-name)
       OBJTYPE(*ALL)
       ENDOPT(*UNLOAD)
```

Task 4: Recovering journals and journal receivers in the QRCL library

To recover journals and journal receivers in the QRCL library, follow these steps.

1. Determine if any objects are in the QRCL library. Type: `DSPLIB QRCL`.

   You see the Display Library display.
2. If no objects are listed on the display, skip to "Task 5: Restoring libraries to the system auxiliary storage pool" on page 176.
3. If the QRCL library contains objects, save them before proceeding with your recovery.

   Load a scratch media volume. Type the following command:

```
SAVLIB LIB(QRCL) DEV(media-device-name)
       ENDOPT(*UNLOAD)
```

4. If the QRCL library does not contain journals or journal receivers, skip to "Task 5: Restoring libraries to the system auxiliary storage pool" on page 176.

5. Create one or more libraries in the system ASP for the journals and journal receivers from the QRCL library.

   The libraries you create must have the same names as the original libraries that contained the journals and journal receivers.

   In the example shown in Figure 6 on page 174, the QRCL library contains the ORDSAV save file and the RCVA0003 journal receiver. At this point you need to create the $JRNLA library. You type: `CRTLIB LIB($JRNLA)`.
6. Move the journals and journal receivers to the newly created libraries.

   This is the only circumstance in which you can move journals and journal receivers between libraries. You must use the MOVOBJ command. You cannot use save and restore commands. The MOVOBJ command leaves the journal or journal receiver in the user ASP but establishes its association with the correct library.

   For the example shown in Figure 6 on page 174, you type the following command:

   ```
   MOVOBJ OBJ(QRCL/RCVA0003) OBJTYPE(*JRNRCV) TOLIB($JRNLA)
   ```
7. Delete the QRCL library by typing `DLTLIB QRCL`.

   Note: If the QRCL library contains save files, you will recover them in "Task 9: Recovering save files from the QRCL library" on page 179. When you recover them, you will use the media volume that you created in step 3.

At this point, the system in the example looks as follows:

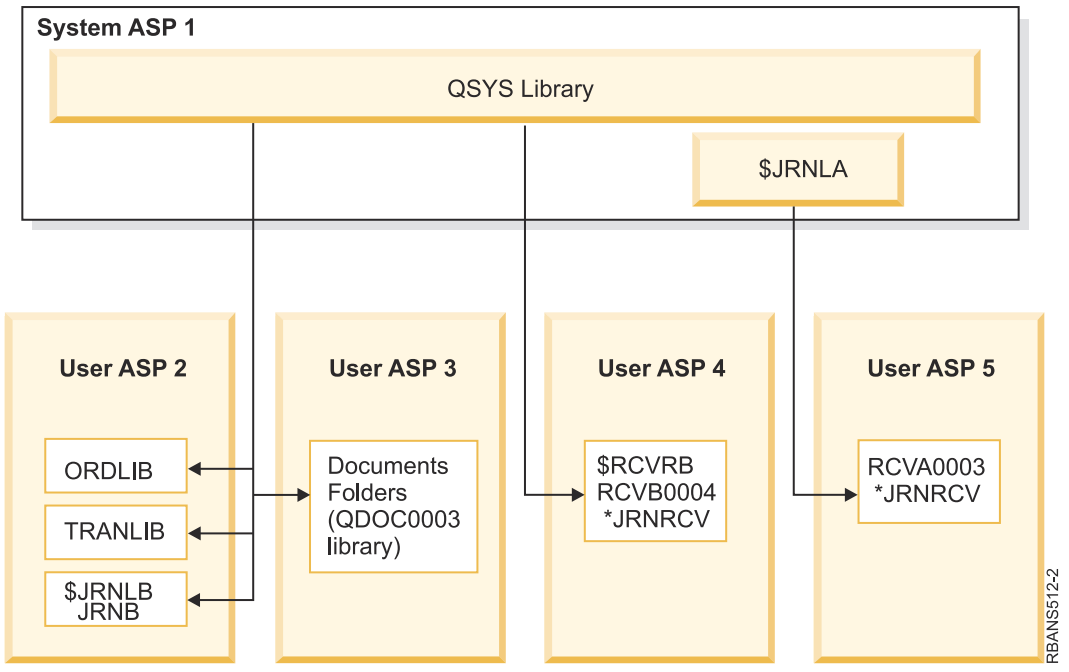

Figure 7. User ASP configuration after recovering isolated journal receiver

Task 5: Restoring libraries to the system auxiliary storage pool

Before you restore libraries to the system auxiliary storage pool (ASP), decide which libraries to restore.

1. You should restore only the libraries in your system ASP. Do not restore the libraries that are already on your system in user ASPs.

   If you are not sure which libraries are currently on your system, type `DSPOBJD OBJ(*ALL) OBJTYPE(*LIB)`.

Note: When you install the operating system, the system creates the QGPL library and the QUSRSYS library. You should still restore these libraries to restore the data from your saved copy.

2. Plan your restore sequence.

   The system restores objects in the same library in the correct order. But if related objects are in different libraries, you must restore them in the correct order or defer the restore of dependent objects. For more information, see the Sequence for restoring related objects topic and the Deferring the restore of dependent objects topic.

3. Choose the commands or menu options you will use.

   You can restore libraries by name or in a group, such as *NONSYS.

   If you restore libraries in a group, omit the libraries in your user ASPs.

4. Type the restore commands or menu options that you have chosen.

   In the example shown in Figure 4 on page 171, libraries were saved using SAVLIB(*ALLUSR). One way to restore them is to type the following command:

```
RSTLIB SAVLIB(*ALLUSR) DEV(media-device-name)
       OMITLIB(ORDLIB TRANLIB $JRNLB $RCVRB)
```

If a media error occurs:

If you have an unrecoverable media error when you are restoring multiple libraries, see "Recovering from an error while restoring libraries" on page 41 topic for more information.

Related concepts

Sequence for restoring related objects
Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

The relationship between save and restore commands
You can learn which store commands can be used, based on how the objects were saved.

Deferring the restore of dependent objects
You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Task 6: Restoring document library objects to the system auxiliary storage pool

To restore document library objects (DLOs) to the system ASP, follow these steps.

1. Find your most recent save media volume that you used to save all of the documents in the system ASP.

   You might have specified ASP(1) or ASP(*ANY) for the save operation. The media volume should have the library QDOC on it.

2. Use the following command to restore the DLOs:

```
RSTDLO DLO(*ALL) FLR(*ANY) ASP(1)
```

If a media error occurs...

If you have an unrecoverable media error when you are restoring DLOs, see "Recovering from an error while restoring document library objects" on page 42 topic for more information.

Task 7: Restoring user-defined file systems to the user auxiliary storage pool

If you are journaling, you need to plan your restore sequence. If you restore in the wrong sequence, your journaling environment might not be started again or some objects might not restore successfully.

For example, journals must be restored before the journaled objects. If the objects are integrated file system objects, you must restore them in the correct order.

Choose one of the three methods based on the way in which your user-defined file systems (UDFSs) were saved.

Related concepts

Sequence for restoring related objects

Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Recovery steps for unmounted user-defined file systems in a user ASP

Use these steps to recover unmounted user-defined file systems (UDFSs) in a user auxiliary storage pool (ASP).

1. Load the media volume with the most recent backup of your UDFS that was saved unmounted.
2. To restore a user auxiliary storage pool (ASP), type `RST OBJ(('/DEV/QASP`*xx*`'))` where *xx* is the ASP number.

   This step restores all the unmounted UDFSs that were saved from QASPxx.

Note: In step 2, you can omit certain objects with either the OBJ or PATTERN parameter to reduce the recovery window when you restore the unmounted UDFS. For example:

```
RST OBJ(('/DEV/QASPxx/udfs_name.udfs') ('*.TEMP' *OMIT))
```

Recovery steps for a mounted user-defined file systems if data is not restored

Use these recovery steps for mounted user-defined file systems (UDFS) in a user auxiliary storage pool (ASP), if the data is not already restored. You can optionally restore both the objects in the UDFS and the file-system information in the UDFS.

To recover the mounted UDFS, perform the following step:

```
RST  OBJ(('/directory_mounted_over '))
     RBDMFS(*UDFS)
```

The mounted UDFS is rebuilt during the restore operation. However, if you omit the RBDMFS parameter or specify RBDMFS(*NONE), only the objects contained in the directory are restored and none of the file-system information is restored.

Recovery steps for a mounted user-defined file system if data is restored

Use these recovery steps for mounted user-defined file systems (UDFSs) if data is already restored.

Attention: The previous method, “Recovery steps for a mounted user-defined file systems if data is not restored” on page 178, is the recommended method for recovering a mounted UDFS. Use this method only if the data is already restored.

1. Create the UDFS exactly as they were before the recovery by using the Create User-Defined File System (CRTUDFS) command.
2. Create temporary directories to use as mount points, using the Create Directory (CRTDIR) command.

3. Mount the UDFS over the temporary directories, using the Add Mounted File System (MOUNT) command. This now becomes your UDFS in the user auxiliary storage pool (ASP).
4. Move or copy the objects in the new UDFS, using the Move Object (MOV) or Copy Object (CPY) commands.
5. Unmount the UDFS, using the Remove Mounted File System (UNMOUNT) command.

Task 8: Reclaiming document library objects

Use these steps to reclaim document library objects (DLOs).

1. If you do not have DLOs in any user auxiliary storage pools (ASPs), skip to "Task 9: Recovering save files from the QRCL library" on page 179.
2. Type the following command:

```
RCLDLO DLO(*ALL) ASP(*ANY)
```

This procedure rebuilds the association between the DLOs in the user ASP and the search index records. It also attempts to assign the DLOs to the correct owner.

Related tasks

Task 11: Restoring object ownership
The Reclaim Storage (RCLSTG) command assigns ownership of all the objects in your user auxiliary storage pool (ASP) to the QDFTOWN user profile.

Task 9: Recovering save files from the QRCL library

Use these steps to recover save files from the QRCL library.

If you did not have any save files in the QRCL library, skip to "Task 10: Associating journal receivers with journals" on page 179.

Note: You displayed the QRCL library in "Task 4: Recovering journals and journal receivers in the QRCL library" on page 175.

1. Load the scratch media volume that you created in "Task 4: Recovering journals and journal receivers in the QRCL library" on page 175.
2. Ensure that the original libraries for the save files were restored in "Task 5: Restoring libraries to the system auxiliary storage pool" on page 176.

   You can check by typing DSPOBJD OBJ(*library-name*) OBJTYPE(*LIB).
3. Restore each save file from the scratch media volume to the correct library and user ASP.

   In the example shown in Figure 4 on page 171, you type the following command:

```
RSTOBJ OBJ(ORDSAV) SAVLIB(QRCL)
       RSTLIB(SAVFLIB)
       OBJTYPE(*FILE) RSTASP(5)
```

Task 10: Associating journal receivers with journals

Whenever you do a recovery involving journals and journal receivers, you must ensure that your journal receivers are associated with the journal.

If you do not have any journals or journal receivers involved in the recovery, skip to "Task 11: Restoring object ownership" on page 181.

Based on the steps performed so far, the receiver directory for journal JRNA in the example looks as follows:

```
                        Work with Receiver Directory

Journal  . . . . . :   JRNA            Library  . . . . . :   $JRNLA

Total size of receivers (in kilobytes). . . . . :   155648

Type options, press Enter.
  4=Delete   8=Display attributes

                                    Attach                      Save
Opt  Receiver    Library     Number  Date      Status            Date
 _   RCVA0003    $JRNLA      00001   06/08/9x  ONLINE            00/00/00
 _   RCVA1002    $JRNLA      01001   06/09/9x  ATTACHED          00/00/00
```

Notice that when JRNA was restored, the system created a new journal receiver called RCVA1002 and attached it. The receiver name is based on the name of the journal receiver that was attached when the journal was saved.

To associate journals and journal receivers, follow these steps:

1. Type WRKJRN on a command line and press the Enter key.
2. On the prompt display, type the name of the journal and the library name.
3. On the Work with Journals display, type a 9 (Associate receivers with journal) in the Opt column next to the journal that you want to work with.
4. Press the Enter key.
5. If you want to limit the number of libraries being searched to be limited, type the name of the library that contains the journal receivers on the prompt display.

   A generic name can be specified for the journal receivers.

If any of the journal receivers in the user auxiliary storage pool (ASP) were created before V3R1, using option 9 from the Work with Journals display might not associate them in the correct sequence. If you have journal receivers from a prior release or if any of the journal receivers you need are not online, follow these steps:

1. Save the journal receivers that are on the system to a scratch media volume:

   ```
   SAVOBJ OBJ(*ALL) LIB(library-name)
          DEV(media-device-name) OBJTYPE(*JRNRCV)
          VOL(*MOUNTED) ENDOPT(*UNLOAD)
   ```

2. After you ensure that the receivers were saved successfully, delete the journal receivers from the library:

   a. Type WRKLIB *library-name* and press the Enter key. You are shown the work with library display.

   b. Type a 12 (Work with Objects) in the Opt column.

   c. Type a 4 (Delete) in the Opt column for each journal receiver you want to delete.

   d. Press the Enter key.

3. Restore the journal receivers you need from the scratch media volume and from your save media volumes. Restore them from newest to oldest by typing the following command for each journal receiver:

   ```
   RSTOBJ OBJ(receiver-name)
          LIB(library-name) DEV(media-device-name)
          OBJTYPE(*JRNRCV) VOL(*MOUNTED)
          ENDOPT(*UNLOAD)
   ```

   The receivers are reassociated with the journal.

At this point the receiver directory for JRNA looks as follows:

```
                        Work with Receiver Directory

Journal  . . . . . :   JRNA           Library  . . . . . :   $JRNLA

Total size of receivers (in kilobytes). . . . . :   155648

Type options, press Enter.
  4=Delete   8=Display attributes

                                    Attach                    Save
Opt  Receiver    Library     Number  Date      Status          Date
 _   RCVA0001    $JRNLA      00001   06/08/9x  SAVED           06/08/9x
 _   RCVA0002    $JRNLA      00002   06/09/9x  SAVED           06/09/9x
 _   RCVA0003    $JRNLA      00003   06/08/9x  ONLINE          00/00/00
 _   RCVA1002    $JRNLA      01002   06/09/9x  ATTACHED        00/00/00
```

Task 11: Restoring object ownership

The Reclaim Storage (RCLSTG) command assigns ownership of all the objects in your user auxiliary storage pool (ASP) to the QDFTOWN user profile.

In "Task 8: Reclaiming document library objects" on page 179, you transferred ownership of document library objects (DLOs) to the correct user profiles. To transfer ownership of the other objects to the correct user profiles, follow these steps:

1. Type `WRKOBJOWN USRPRF(QDFTOWN)` and press the Enter key.

   The Work with Objects by Owner display is shown:

```
                         Work with Objects by Owner

User profile . . . . . :   QDFTOWN

Type options, press Enter.
 2=Edit authority        4=Delete   5=Display author
 8=Display description   9=Change owner

Opt  Object        Library       Type      Attribute
 9   ORDRCV001     JRNLIB        *JRNRCV
 9   ORDHDR        ORDLIB        *FILE
 9   ORDDTL        ORDLIB        *FILE
 9   ORDHST        ORDLIB        *FILE
 9   ORDSAV        SAVFLIB       *SAVF
 9   TRAN01        TRANLIB       *FILE

:
Parameters or command
===> NEWOWN(OWNORD)
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve
F18=Bottom
```

   Note: If you see document library objects on this list (type *DOC or *FLR), one of the following situations has occurred:

   - You forgot to run the Reclaim Document Library Object (RCLDLO) command.
   - The user profile that owns the DLO has not been restored. Restore the user profile. Then run the RCLDLO command.
   - The DLO was owned by the QDFTOWN profile when it was saved. Determine the correct owner for the DLO and transfer ownership.

2. To transfer ownership of objects individually, follow these steps:
   a) Type a 9 in the Opt column for the object and press the Enter key.

      The Change Object Owner display is shown.

   b) Type the name of the correct owner in the *New owner* prompt and press the Enter key.

   c) Repeat steps a and b for each object on the display.

3. To transfer ownership of multiple objects that should have the same owner, use the technique shown in the display:

a) Type 9 in the Opt column.

b) Type NEWOWN(*owner-name*) on the parameter line at the bottom of the display.

c) Press the Enter key.

The system transfers ownership of each object you specified to the new owner.

You have completed the recovery of information in your user ASPs. Consult your recovery checklist for the next step in your recovery process.

Related tasks

Task 8: Reclaiming document library objects
Use these steps to reclaim document library objects (DLOs).

Recovering an overflowed user auxiliary storage pool

Reset a user auxiliary storage pool (ASP) in overflowed status as soon as possible. An overflowed ASP affects system performance. It also makes recovery more difficult and might increase the amount of data that is lost if a failure occurs.

When the disk units allocated to a user ASP become full, the user ASP is in overflowed status. The system sends message CPI0953 to the QSYSOPR message queue warning you that an ASP is approaching its storage threshold. The system sends a message CPI0954 when the storage threshold is exceeded and the ASP is in overflowed status.

To recover an overflowed user ASP, follow the procedures in the "Resetting an overflowed user auxiliary storage pool (ASP) without an initial program load (IPL)" on page 182 topic.

Note: To simplify overflow recovery operations, you can enable automatic overflow recovery for basic user ASPs with the IBM Navigator for i Disk Management function.

Related concepts

Working with auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) to work with auxiliary storage pools (ASPs). ASPs are also called *disk pools*. You can create new encrypted or unencrypted user ASPs, or add disk units to an existing ASP.

How the system responds to disk unit full 448 reference code A6xx 0277
The system bases its response to the disk unit full condition on the type of input/output (I/O) operation that caused the condition.

Resetting an overflowed user auxiliary storage pool (ASP) without an initial program load (IPL)

You can learn the detailed steps of how to reset an overflowed user ASP without an IPL.

To reset a user ASP in overflowed status, follow these steps:

1. Determine which objects in the ASP have overflowed.

   Use one of the following methods:

   - Use the Display Object Description (DSPOBJD) command to create an output file. Then run a query against that output file:

     a. For the first library in the user ASP, type the following command:

     ```
     DSPOBJD OBJ(library-name/*ALL) OBJTYPE(*ALL)
            DETAIL(*FULL) OUTPUT(*OUTFILE)
            OUTFILE(library-name/file-name)
     ```

     b. For each additional library in the user ASP, type the following command:

     ```
     DSPOBJD OBJ(library-name/*ALL) OBJTYPE(*ALL)
            DETAIL(*FULL) OUTPUT(*OUTFILE)
     ```

```
        OUTFILE(library-name/file-name)
        OUTMBR(*FIRST *ADD)
```

c. Create a query against the output file. Look for objects that have a 1 (Yes) in the **ODOASP (overflowed ASP)** field.

- For a user ASP that contains only DLOs, use the Query Document Library (QRYDOCLIB) command. It has a parameter to display overflowed DLOs.

2. Save each overflowed object to a scratch media volume.
3. Delete each overflowed object.

   Some objects, such as journals and physical files, require that you take certain actions before deleting them. Table 42 on page 183 shows what to do before deleting these objects.

Table 42. Object types that require special procedures for deleting

Object type	Do this before deleting
Journal	“Deleting a journal” on page 253
Journal receiver	“Deleting a journal receiver” on page 256
Physical file	“Deleting a physical file” on page 252

4. Ensure that the ASP is no longer in overflowed status.

   You should have received a message in the QSYSOPR message queue that the overflow condition has been recovered. You can also use system service tools (SST) to check:

   a. Type STRSST. You are shown the System Service Tools (SST) menu.

   b. Select the option to work with disk units.

   c. Select the option to display disk configuration.

   d. Select the option to display disk configuration capacity. You are shown the Display Disk Configuration Capacity display:

```
                    Display Disk Configuration Capacity

 --Protected--  --Unprotected--
ASP  Unit  Type  Model  Threshold  Overflow   Size  %Used    Size  %Used
1                         90%     No           0   0.00%   1400   8.22%
       1   9332   400                          0   0.00%    200   17.97%
       2   9332   400                          0   0.00
2                                 Yes          0   0.00%    200   99.99%
       8   9332   200      90%                 0   0.00%    200   99.99%
```

   This display shows whether any ASPs are in overflowed status.

   If the user ASP is still overflowed, follow the procedure described in “Resetting an overflowed user auxiliary storage pool (ASP) during an initial program load (IPL)” on page 184 topic.

5. Before you can restore the overflowed objects from a media volume, you must make additional space available in the user ASP.

   Do one or more of the following tasks:

   - Delete objects from the ASP if you no longer need them.
   - Move one or more libraries to a different ASP.

     Note: You cannot use the Move Object (MOVOBJ) command to do this. You must save the library, delete it, and restore it to a different ASP.
   - Move one or more folders to a different ASP by saving the folder, deleting it, and restoring it to a different ASP.
   - Add additional disk units to the ASP.

6. After you have made additional space available in the ASP, restore the objects you saved in step 2.

7. Check to make sure the user ASP has sufficient space and is not overflowed. Repeat the procedure described in step 4.

Related tasks

Deleting overflowed objects during recovery
To recover a user auxiliary storage pool (ASP) that was in overflowed status, follow these steps.

Resetting an overflowed user auxiliary storage pool (ASP) during an initial program load (IPL)

You can learn how to reset an overflowed user ASP during an IPL.

Sometimes you are not able to find all the overflowed objects in a user ASP. If you did the steps in the "Resetting an overflowed user auxiliary storage pool (ASP) without an initial program load (IPL)" on page 182 topic and the user ASP is still overflowed, you can do a manual IPL to reset the user ASP. Do the following tasks -

1. Verify that you have enough space to reset the overflowed user ASP. Follow these steps:
   a) Type STRSST. You are shown the System Service Tools (SST) menu.
   b) Select the option to work with disk units.
   c) Select the option to display disk configuration.
   d) Select the option to display disk configuration capacity.
      You are shown the Display Disk Configuration Capacity display:

```
                    Display Disk Configuration
Capacity

 --Protected--  --Unprotected--
ASP  Unit  Type  Model  Threshold  Overflow   Size  %Used    Size  %Used
1                          90%    No            0   0.00%   1400   8.22%
1          9332   400                          0   0.00%    200  17.97%
       2   9332   400                          0   0.00
2                                 Yes           0   0.00%    200  99.99%
       8   9332   200       90%                 0   0.00%    200  99.99%
```

      This display shows whether any ASPs are in overflowed status.
   e) Press F9 (Display ASP Overflow information) to display the overflow amount and the additional amount of storage that is needed in the ASP to recover the overflowed objects.

```
                     Display ASP Overflow Information
 Overflow      ----Amount Needed to Recover---- ASP  Threshold    Amount
               To Capacity     To Threshold
  2                  90%           14                0              0
  3                  90%           25               25             45
```

   f) If the amount in the **To Capacity** field is greater than zero, the ASP remains overflowed when the recovery completes.
      The user ASP does not have enough available space to contain the overflowed data.
   g) If you do not have enough space, repeat the instructions in step 5 of the Resetting an overflowed user ASP without an initial program load IPL topic to make more space available.
2. To put your system in a restricted state, follow these steps:
   a) Before you put your system in a restricted state, verify that all users are signed off and all jobs are ended.
   b) To receive notification that the subsystems are ended, type the following command and press the Enter key:

```
CHGMSGQ MSGQ(QSYSOPR) DLVRY(*BREAK)
        SEV(60)
```

c) To end all subsystems, type the following command:

```
ENDSBS SBS(*ALL) OPTION(*CNTRLD)
       DELAY(600)
```

Note: For the delay parameter, specify the number of seconds to allow your system time to bring most jobs to a normal end. On a large, busy system, you might need a longer delay.

A message is sent that indicates that the procedure for ending subsystems is in progress. A final message is sent when the system is in a restricted state.

3. Do a manual IPL and access dedicated service tools (DST):

   Use this procedure to start DST. If the **IPL** or **Install the System** menu is already displayed, start with step 5 in Task 1–Accessing dedicated service tools topic.

   a. Place the system in manual mode.

   b. Power off the system -

```
PWRDWNSYS OPTION(*CNTRLD) DELAY(600)
          RESTART(*YES) IPLSRC(B)
```

   Note:

   If you are sure that no jobs are running on your system, you can specify `OPTION(*IMMED)` when you power off the system. Otherwise, specify a delay time that is sufficient to allow jobs to end normally.

   c. When the IPL completes, the **IPL** or **Install the System** menu appears.

```
                          IPL or Install the System

   Select one of the following:
            1. Perform an IPL
            2. Install the operating system
            3. Use dedicated service tools (DST)
            4. Perform automatic installation of the operating system
            5. Save Licensed Internal Code
```

4. Select option 1 (Perform an IPL).

   You are shown the Disk Configuration Attention Report:

```
                    Disk Configuration Attention Report

Type option, press Enter.
     5=Display Detailed Report
Press F10 to accept all the following problems and continue.
The system will attempt to correct them.

Opt  Problem
     Overflowed ASPs
```

   If you type 5 in the **Option** field, the following display is shown, listing the User ASPs that are overflowed.

```
                         Recover Overflowed User ASP

The following user ASPs are overflowed.


ASP
 2
 3
```

5. Press the F10 key to request recovery of the overflowed user ASPs.

The recovery takes place during IPL storage management recovery. The operation takes from several minutes to a few hours, depending on the number of objects on the system and the amount of data to be recovered.

6. When the IPL of the system is complete, the Sign On display is shown.
7. Sign on and verify the results by checking the messages in the QSYSOPR message queue.

Deleting overflowed objects during recovery

To recover a user auxiliary storage pool (ASP) that was in overflowed status, follow these steps.

1. After running the Reclaim Storage (RCLSTG) command, display the contents of the QRCL library by typing: `DSPLIB QRCL`
2. Write down the names of the objects in the library.

   These objects were overflowed into the system ASP at the time of the failure.

   Although the initial disk extents for these objects might have been allocated in the system ASP, portions of the objects might still have been lost. The integrity of these objects cannot be predicted. They should be deleted and recovered.
3. Delete the overflowed objects.

   You must take special action before deleting certain types of objects.
4. When you run the RCLSTG command, any documents from the lost user ASP that were in overflowed status are placed in the user ASP again. (The system creates a new QDOC*nnnn* library, where *nnnn* is the number of the lost ASP, and places the overflowed DLOs in it.)

   Assuming you have not yet restored document library objects (DLOs) to the user ASP, type the following command to delete the overflowed DLOs:

   ```
   DLTDLO DLO(*ALL) FLR(*ANY) ASP(n)
   ```

   where *n* is the number of the ASP whose data was lost.

Related tasks

Resetting an overflowed user auxiliary storage pool (ASP) without an initial program load (IPL)
You can learn the detailed steps of how to reset an overflowed user ASP without an IPL.

Recovering a damaged basic auxiliary storage pool

You can use this procedure to recover a damaged basic auxiliary storage pool (ASP).

Perform this procedure if one of the following is true:

- The service representative has replaced a failed disk unit in a basic user ASP. If a disk unit is not parity protected or mirrored, when you lose that disk unit in an ASP, you must recover all of the information in that ASP. The information in other ASPs on your system is not affected.
- The system has reassigned sectors on a disk unit, but object damage has occurred.
- You have replaced a disk unit in the system ASP and one or more user ASPs was overflowed.

Task 1: Restoring user profiles

Even though user profiles are not lost when you replace a unit in a user auxiliary storage pool (ASP), they must be restored to prepare for restoring authority to objects in the user ASP.

To restore user profiles, follow these steps:

1. Sign on with the QSECOFR user profile.
2. End all subsystems with the End Subsystem (ENDSBS) command and go to a restricted state.
3. Load your most recent SAVSYS or SAVSECDTA media volume.
4. Restore all user profiles.

Type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*UNLOAD)
```

5. If you know what libraries or objects were in the user ASP that was lost, skip to "Task 3: Determining tasks to restore objects" on page 188.

   If you do not know what was in the user ASP, continue with "Task 2: Determining the contents of the lost auxiliary storage pool" on page 187.

Task 2: Determining the contents of the lost auxiliary storage pool

If you are not sure what was on the basic auxiliary storage pool (ASP), follow these steps to determine the contents of the lost ASP.

If your system has a small number of libraries and is well-documented, like the one in Figure 4 on page 171, your task is relatively simple. In the example, if a disk unit in ASP 2 is replaced, the user must recover the ORDLIB, TRANLIB, and $JRNLB libraries. If a disk unit in ASP 5 is replaced, the user must recover all the journal receivers in the $JRNLA library and the ORDSAV save file in the SAVFLIB library.

1. Sign on with a user profile that has *ALLOBJ special authority so that your listings show all libraries.
2. Print a list of the libraries that are on the lost basic ASP by following these steps:
   a) Create a list of all the libraries in an output file:

   ```
   DSPOBJD OBJ(QSYS/*ALL) OBJTYPE(*LIB) ASPDEV(*ALLLVL)
           DETAIL(*FULL) OUTPUT(*OUTFILE)
           OUTFILE(library-name/file-name)
   ```

   b) Use a program or a query tool to display or print the output file. Select all entries that have an **ASP** field that matches the ASP that is lost.

   Notes:

   - When you lose a basic ASP, you lose the contents of any libraries in the ASP, not the libraries themselves. The library objects are in the QSYS library, which is in the system ASP.
   - If you had documents in the basic ASP, you should have a library on your listing for the ASP. The library name is QDOC*nnnn*, where *nnnn* is the number of the ASP.
3. If you have determined what must be recovered, continue with "Task 3: Determining tasks to restore objects" on page 188. If you have not found any libraries to recover, continue with step 4.
4. If you did not find any libraries to recover in step 2, the ASP was probably a nonlibrary user ASP.

   A nonlibrary user ASP can contain only save files, journals, and journal receivers.

   Determining the objects that were in a nonlibrary user ASP can be very time consuming. The following steps are one method. This method works only if you have not yet run Reclaim Storage (RCLSTG) command after losing the user ASP.

   a. Type the following command:

   ```
   DSPOBJD OBJ(*ALL/*ALL)
           OBJTYPE(*LIB *FILE *JRN *JRNRCV)
           OUTPUT(*OUTFILE)
           OUTFILE(library-name/file-name)
   ```

   b. Use a program or query tool to list all the objects in the output file that are in the ASP that is damaged.
5. When you have determined the objects that need to be recovered, continue with "Task 3: Determining tasks to restore objects" on page 188.

Task 3: Determining tasks to restore objects

Use this information to determine the appropriate tasks to perform to restore basic ASP objects and other types of objects, such as libraries and user-defined file systems.

1. Use Table 43 on page 188 to determine how to recover the objects in your basic auxiliary storage pool (ASP).

   It shows the recovery tasks you must perform, based on the contents of the basic ASP you are recovering.
2. If you have different types of objects to recover, such as libraries and user-defined file systems, perform the tasks in the order shown in the table.

Table 43. Tasks for restoring basic ASP objects

Type of ASP	Contents	Recovery tasks
Library user ASP	Libraries	"Task 4: Restoring libraries to a basic auxiliary storage pool" on page 188
Nonlibrary user ASP	Journals	"Task 5: Restoring journals to a basic auxiliary storage pool" on page 188
Library user ASP	Documents	"Task 6: Restoring documents to a basic auxiliary storage pool" on page 190
Library user ASP	User-defined file systems (UDFS)	"Task 7: Restoring user-defined file systems to a basic auxiliary storage pool" on page 190
Nonlibrary user ASP	Journal receivers	"Task 8: Restoring journal receivers to a basic auxiliary storage pool" on page 191
Nonlibrary user ASP	Save files	"Task 9: Restoring save files to a basic auxiliary storage pool" on page 191

Task 4: Restoring libraries to a basic auxiliary storage pool

Use these steps to restore libraries to a basic auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. For each library you need to recover, load the correct volume from your latest save media volumes.
3. Type the following command:

```
RSTLIB SAVLIB(library-name) DEV(media-device-name)
       ENDOPT(*LEAVE)
```

   Note: You should restore changed objects and apply journaled changes for all the ASPs included in your recovery at the same time. These steps appear on the recovery checklist at the appropriate point.
4. Continue with the next task shown in Table 43 on page 188. If you have completed all the appropriate tasks in the table, continue with the next task in the recovery checklist in Chapter 3, "Selecting the appropriate recovery strategy," on page 53.

Task 5: Restoring journals to a basic auxiliary storage pool

Use these steps to restore journals to a basic auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. For each journal you need to recover, load the correct save media volume and type the following command:

```
RSTOBJ OBJ(journal-name) SAVLIB(library-name)
       DEV(media-device-name) OBJTYPE(*JRN)
```

When you restore the journal, the system automatically creates and attaches a new journal receiver. The Journal management topic describes how the system names the journal receiver that is created when you restore a journal.

3. Establish your journaling environment again.
   a) For each database physical file that was journaled to the restored journal, type the following command:

```
STRJRNPF FILE(library-name/file-name)
         JRN(library-name/journal-name)
```

   Note: To determine what options you specified for the file when you last journaled it, you can use the Display File Description (DSPFD) or Display Object Description (DSPOBJD) commands for the file to find out.

   b) For each access path that was journaled to the restored journal, type the following command:

```
STRJRNAP FILE(library-name/file-name)
         JRN(library-name/journal-name)
```

   c) For each integrated file system object that was journaled to the restored journal, type the following command:

```
STRJRN OBJ ('object-path-name')
       JRN('journal-path-name')
```

   Note: To determine what options you specified for the object when you last journaled it, you can use the Display Link (DSPLNK) command.

   d) To start journaling a library, type the following command:

```
STRJRNLIB LIB(library-name)
JRN(library-name/journal-name)
```

   Note: To determine what options you specified for the object when you last journaled it, you can use the Display Object Description (DSPOBJD) command. Use the Display library description (DSPLIBD) command to determine the inheritance rules that were associated with the library when it was previously journaled.

   e) For all other object types that were journaled, type the following command:

```
STRJRNOBJ OBJ(library-name/object-name)
          OBJTYPE(object-type)
          JRN(library-name/journal-name)
```

   f) Save each object that you started journaling.
4. If you need to restore journal receivers for the journals, skip to “Task 8: Restoring journal receivers to a basic auxiliary storage pool” on page 191.
5. Associate journal receivers with the journals you restored.
   a) Type WRKJRN on a command line and press the Enter key.
   b) On the prompt display, type the name of the journal and the library name.
   c) On the Work with Journals display, type a 9 (Associate receivers with journal) in the Opt column next to the journal that you want to work with.
   d) Press the Enter key.
   e) If you want to limit the number of libraries being searched for needs to be limited, type the name of the library that contains the journal receivers on the prompt display.

      A generic name can be specified for the journal receivers.

6. Continue with the next task shown in Table 43 on page 188. If you have completed all the appropriate tasks in the table, continue with the next task in the recovery checklist from Chapter 3, “Selecting the appropriate recovery strategy,” on page 53.

Related information

Displaying and retrieving library descriptions

Task 6: Restoring documents to a basic auxiliary storage pool

Use these steps to restore documents to a basic auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. Load the media volume with your last complete save of documents in the user ASP.
3. Restore the documents to the user ASP by typing the following command:

```
RSTDLO DLO(*ALL) SAVASP(ASP-number)
       RSTASP(ASP-number)
```

   This command restores the documents and makes any changes necessary to the search index database files.
4. Use the Query Document Library (QRYDOCLIB) command to locate any documents that were created on the user ASP since the last save operation. Query by ASP number and creation date. Inform your users that these documents were lost and develop a plan to re-create them.
5. Continue with the next task in the recovery checklist from Chapter 3, “Selecting the appropriate recovery strategy,” on page 53.

Task 7: Restoring user-defined file systems to a basic auxiliary storage pool

Based on the way in which your user-defined file systems (UDFS) were saved, choose one of these three methods to restore UDFS to a basic auxiliary storage pool (ASP).

Recovery steps for unmounted user-defined file systems

Use these steps to recover unmounted user-defined file systems (UDFSs) in a basic auxiliary storage pool (ASP).

1. Load the media volume with the most recent backup of your UDFS were saved unmounted.
2. To restore a basic auxiliary storage pool (ASP), type `RST OBJ(('/DEV/QASPxx'))` where *xx* is the ASP number.

   This step restores all the unmounted UDFSs that were saved from QASPxx.

Note: In step 2, you can omit certain objects with either the OBJ or PATTERN parameter to reduce the recovery window when you restore the unmounted UDFS. For example:

```
RST OBJ(('/DEV/QASPxx/udfs_name.udfs') ('*.TEMP' *OMIT))
```

Recovery steps for a mounted user-defined file system if data is not restored

Use these recovery steps for mounted user-defined file systems (UDFSs) in a basic auxiliary storage pool (ASP), if the data is not already restored. You can optionally restore both the objects in the UDFS and the file-system information in the UDFS.

To recover the mounted UDFS, perform the following step:

```
RST  OBJ(('/directory_mounted_over '))
     RBDMFS(*UDFS)
```

The mounted UDFS is rebuilt during the restore operation. However, if you omit the RBDMFS parameter or specify RBDMFS(*NONE), only the objects contained in the directory are restored and none of the file-system information is restored.

Recovery steps for a mounted user-defined file system if data is restored

Use these steps to recover mounted user-defined file systems (UDFSs) if the data is already restored.

Attention: This method is not recommended for recovery of UDFSs. It is listed only as a means of recovery if the data is already restored. The previous method, "Recovery steps for a mounted user-defined file systems if data is not restored" on page 178, is the suggested method.

1. Because the UDFS information is not saved or restored if the UDFS were saved mounted, re-create the UDFS information exactly as they were before the recovery by using the Create User-Defined File System (CRTUDFS) command.
2. Create a temporary directory to use as a mount point, using the Create Directory (CRTDIR) command.
3. Mount the UDFS over the temporary directory, using the Add Mounted File System (MOUNT) command. This now becomes your UDFS in the user ASP.
4. Create the directories that are currently in the restored mounted UDFS in the UDFS that you created in the previous three steps. This tree structure must exist in order to move or copy the objects.
5. Move or copy the objects in the new UDFS, using the Move Object (MOV) or Copy Object (CPY) commands.
6. Unmount the UDFS, using the Remove Mounted File System (UNMOUNT) command.

Task 8: Restoring journal receivers to a basic auxiliary storage pool

Use these steps to restore journal receivers to a basic auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. For each journal receiver you need to recover, load the correct save media volume and type the following command:

```
RSTOBJ OBJ(receiver-name) SAVLIB(library-name)
       DEV(media-device-name) OBJTYPE(*JRNRCV)
```

3. Continue with the next task shown in Table 43 on page 188. If you have completed all the appropriate tasks in the table, continue with the next task in the recovery checklist in Chapter 3, "Selecting the appropriate recovery strategy," on page 53.

Task 9: Restoring save files to a basic auxiliary storage pool

Use these steps to restore save files to a basic auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. For each save file you need to recover, load the correct save media volume and type the following command:

```
RSTOBJ OBJ(save-file-name) SAVLIB(library-name)
       DEV(media-device-name) OBJTYPE(*FILE)
```

   Note: This command restores the description of the save file and its contents, if you specified SAVFDTA(*YES) when you saved the save file. If you specified SAVFDTA(*NO) when you saved the save file, this command restores only the save file description.
3. Continue with the next task in the recovery checklist from Chapter 3, "Selecting the appropriate recovery strategy," on page 53.

Recovering an independent auxiliary storage pool

Run this procedure if one of the following situations is true.

Independent auxiliary storage pools (ASPs) are called independent disk pools in IBM Navigator for i.

- The service representative replaced a failed disk unit in an independent ASP. If a disk unit is not parity protected or mirrored, when you lose that disk unit in an ASP, you must recover all the information in that ASP. The information in other ASPs on your system is not affected.
- The system reassigned sectors on a disk unit but object damage occurred.
- You are running a full system recovery and were directed to this procedure from a recovery checklist.

The independent ASP that you recover must be in the Available state to run the restore.

Related tasks

Restoring encrypted auxiliary storage pools
If you have an encrypted user or independent auxiliary storage pool (ASP), you must run special steps to verify that the data in these ASPs can be recovered.

Task 1: Restoring user profiles

Even though user profiles are not lost when you replace a unit in an independent auxiliary storage pool (ASP), they must be restored to prepare for restoring authority to objects in the independent ASP.

Note: If you are performing a full system recovery and chose not to perform a Restore Authority (RSTAUT) command before the initial program load (IPL), or if you chose to only restore authorities for the system and basic ASPs, you can skip this step.

1. Sign on with the QSECOFR user profile.
2. End all subsystems with the End Subsystem (ENDSBS) command and go to a restricted state.
3. Load your most recent SAVSYS or SAVSECDTA media volume.
4. Restore all user profiles. Type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*UNLOAD) SECDTA(*PVTAUT)
```

5. If you are restoring a user-defined file system (UDFS) ASP, skip to "Task 4: Restoring user-defined file systems to the independent auxiliary storage pool" on page 194.

Task 2: Determining tasks to restore objects to an independent auxiliary storage pool

If you are performing a full system recovery, you need to perform this task for each independent auxiliary storage pool (ASP).

Your restore process will be most efficient if you restore independent ASPs and their contents in the order that they were saved. Independent ASPs are saved alphabetically. Secondary ASPs are saved along with their primary.

Table 44. Example of restore order for independent ASPs saved with GO SAVE: Option 21 or 23

Restore order	Independent ASP name	Independent ASP type	What is restored	Command
1	Apples	Primary	Libraries	RSTLIB SAVLIB (*ALLUSR)
	Cantaloupe	Secondary		
2	Apples	Primary	User-defined file systems	RST OBJ(('/DEV/ iasp-name'))
	Cantaloupe	Secondary		

Table 44. Example of restore order for independent ASPs saved with GO SAVE: Option 21 or 23 (continued)

Restore order	Independent ASP name	Independent ASP type	What is restored	Command
3	Bananas	UDFS	User-defined file systems	RST OBJ(('/DEV/ iasp-name'))

1. Use Table 45 on page 193 to determine how to recover the objects in your independent ASP.

   It shows the recovery tasks you must perform, based on the contents of the independent ASP you are recovering.
2. If you have different types of objects to recover, such as libraries and documents, perform the tasks in the order shown in the table.

Table 45. Tasks for restoring independent ASP objects

Contents	Recovery Tasks
Libraries	"Task 3: Restoring libraries to the independent auxiliary storage pool" on page 193
User-defined file systems	"Task 4: Restoring user-defined file systems to the independent auxiliary storage pool" on page 194

Task 3: Restoring libraries to the independent auxiliary storage pool

Use these steps to restore libraries to the independent auxiliary storage pool (ASP).

1. Sign on with a user profile that has *SAVSYS and *JOBCTL special authority.
2. Specify the independent ASP group by typing the following command:

```
SETASPGRP(iasp-group-name)
```

3. To recover libraries within the independent ASP group, load the correct volume from your latest save media volumes.

   Make sure you are at the correct spot in the save media. You might need to specify a sequence number to access the correct library data for your independent ASP.

   - To restore libraries that were saved by a GO SAVE: Option 21 or 23 command, type the following command:

```
RSTLIB SAVLIB(*ALLUSR) DEV(media-device-name)
       ENDOPT(*LEAVE)
```

   - To restore an individual library to the independent ASP, type the following command:

```
RSTLIB SAVLIB(library-name) DEV(media-device-name)
       ENDOPT(*LEAVE)
```

   Note: If you are restoring the library from optical media, you must also specify the path name:

```
RSTLIB SAVLIB(library-name) DEV(optical-device-name)
       OPTFILE('QSRSAVIASP/iasp_name/*')
```

   The default value for the Restore Library (RSTLIB) command is to restore the library to the independent ASP from which it was saved. If you want to restore the library to another independent ASP, you can use the RSTASPDEV parameter. If you want to restore the library to a basic ASP or the system ASP instead of the independent ASP you can use the RSTASP parameter. It is possible to

restore the same library to two different independent ASPs. However, you cannot restore the same library to an independent ASP as well as the system ASP or basic ASP.

If you are restoring libraries to an independent ASP with the same name but a different ASP number, the following libraries will automatically be renamed when restored:

- QSYS2*nnnnn*
- QRCL*nnnnn*
- SYSIB*nnnnn*

where *nnnnn* is the number of the independent ASP.

Note: You should restore changed objects and apply journaled changes for all the ASPs included in your recovery at the same time. These steps appear on the recovery checklist at the appropriate point.

4. Continue with the next task shown in Table 45 on page 193. If you have completed all the appropriate tasks in the table, continue with the next task in the recovery checklist from Chapter 3, "Selecting the appropriate recovery strategy," on page 53.

Task 4: Restoring user-defined file systems to the independent auxiliary storage pool

Based on the way in which your user-defined file systems (UDFS) were saved, choose one of these three methods.

Recovery steps for unmounted user-defined file systems

Use these steps to recover unmounted user-defined file systems (UDFSs) in an independent auxiliary storage pool (ASP).

1. Load the media volume with the most recent backup of your UDFS that was saved when it was unmounted.
2. Unmount any QDEFAULT user-defined file systems in the independent ASP.
3. To restore all the UDFSs into an independent ASP, type `RST OBJ(('/DEV/iasp-name'))` where *iasp-name* is the name of the independent ASP.

   If you are restoring multiple independent ASPs in the order you saved them, you can also specify `RST OBJ(('/DEV/*'))` to restore all user-defined file systems for each independent ASP.

   Note: If you are restoring file systems from optical media, you must also specify the device and optical file path name for the current primary ASP group that is associated with the job -

```
RST  DEV('/qsys.lib/optical-device-name') OBJ(('/DEV/*'))
     OPTFILE('QSRSAVIASP/primary-ASP-name/*')
```

4. Mount any QDEFAULT user-defined file systems that were unmounted in step 2.

Recovery steps for a mounted user-defined file system if data is not restored

Use these recovery steps for mounted user-defined file systems (UDFSs) in an independent auxiliary storage pool (ASP), if the data is not already restored. You can optionally restore both the objects in the UDFS and the file-system information in the UDFS.

To recover the mounted UDFS, perform the following step:

```
RST  OBJ(('/directory_mounted_over '))
     RBDMFS(*UDFS)
```

The mounted UDFS is rebuilt during the restore operation. However, if you omit the RBDMFS parameter or specify RBDMFS(*NONE), only the objects contained in the directory are restored and none of the file-system information is restored.

If you omit the RBDMFS(*UDFS) parameter from the RST command, you need to perform the following steps to recover the mounted UDFS:

1. Create the UDFS exactly as they were before the recovery by using the Create User-Defined File System (CRTUDFS) command. Be sure to include the authorities and object auditing.
2. Create the directory that each UDFS was mounted over at the time of the save by using the Create Directory (CRTDIR) command.
3. Mount the UDFS over the directory by using the Add Mounted File System (MOUNT) command.

   Note: If you were instructed to refer to these steps from another checklist, return to that checklist now.
4. Restore the UDFS by using the following command:

```
RST OBJ(('/directory_mounted_over '))
```

Recovery steps for a mounted user-defined file system if data is restored

Use these steps to recover mounted user-defined file systems (UDFSs) if the data is restored.

Attention: This method is not suggested for the recovery of UDFSs. It is listed only as a means of recovery if the data is already restored. The previous method, "Recovery steps for a mounted user-defined file systems if data is not restored" on page 178, is the suggested method.

Because the UDFS information was not restored with the data, you need to re-create this information in step 1.

1. Create the UDFS information exactly as it was before the recovery by using the Create User-Defined File System (CRTUDFS) command.
2. Create a temporary directory to use as a mount point, using the Create Directory (CRTDIR) command.
3. Mount the UDFS over the temporary directory, using the Add Mounted File System (MOUNT) command. This now becomes your UDFS in the user ASP.
4. Create the directories that are currently in the restored mounted UDFS in the UDFS that you created in the previous three steps. This tree structure must exist in order to move or copy the objects.
5. Move or copy the objects in the new UDFS, using the Move Object (MOV) or Copy Object (CPY) commands.
6. Unmount the UDFS, using the Remove Mounted File System (UNMOUNT) command.

Task 5: Restoring authority for an independent auxiliary storage pool

These scenarios are the most common restore methods for restoring authority for an independent auxiliary storage pool (ASP).

- Save from an independent ASP and restore it to an independent ASP with an identical name.
- Save from *SYSBAS and restore it to a new independent ASP.
- Save from an independent ASP and restore it to an independent ASP with a different name.
- Save from an independent ASP and restore it to *SYSBAS.

The "Restoring authorities for independent auxiliary storage pool data" on page 225 topic explains the scenarios in more detail with command examples.

Removing a failed disk unit from the system auxiliary storage pool

By removing a failed disk unit from the system auxiliary storage pool (ASP), you return your system to operation if a replacement disk unit is not immediately available. However, this action removes all data from your system and requires a complete restore operation.

Before you remove a failed disk unit from the system ASP, ensure that the remaining 2800-001 storage units in your system ASP are large enough for a main storage dump. Consult software support or refer to the Chapter 22, “Working with auxiliary storage pools,” on page 391 topic.

After you remove a failed disk unit from the system ASP, your system will have less disk capacity. You might not be able to restore all user information until you have installed and configured a replacement disk unit.

Task 1: Accessing dedicated service tools

You can use this procedure to start dedicated service tools (DST).

If the **Initial Program Load (IPL)** or **Install the System** menu is already displayed, start with step 5.

1. Place the system in manual mode.
2. Power off the system -

```
PWRDWNSYS OPTION(*CNTRLD) DELAY(600) RESTART(*YES)
          IPLSRC(B)
```

Note:

If you are sure that no jobs are running on your system, you can specify `OPTION(*IMMED)` when you power off the system. Otherwise, specify a delay time that is sufficient to allow jobs to end normally.

3. When the IPL completes, the **IPL** or **Install the System** menu appears.

```
                        IPL or Install the System

  Select one of the following:
           1. Perform an IPL
           2. Install the operating system
           3. Use dedicated service tools (DST)
           4. Perform automatic installation of the operating system
           5. Save Licensed Internal Code
```

4. Select option 3 (Use Dedicated Service Tools (DST)) and press the Enter key.

   The Dedicated Service Tools Sign On display is shown.

```
                   Dedicated Service Tools (DST) Sign On
  Type choice, press Enter.
    Service tools user . . . . . .   _______
    Service tools password . . . . . .   _______
```

5. In the **Service tools user** field, type QSECOFR.

   In the **Service tools password** field, type your DST service tools password. On a new system, the password is QSECOFR. The password is case-sensitive; type in all capital letters. The service tools profile QSECOFR password is expired after the first use. At the Change Service Tools User Password display, enter in all uppercase letters the current password QSECOFR and a new password, along with the verification password.

   The Use Dedicated Service Tools (DST) menu is shown.

```
                    Use Dedicated Service Tools (DST)

 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system

     10. Work with remote service support

     13. Work with system security
     14. End batch restricted state
```

Related information

Service tools user IDs and passwords

Task 2: Deleting the auxiliary storage pool data

Use these steps to delete the auxiliary storage pool (ASP) data.

1. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
2. Select option 4 (Delete ASP data) on the Work with ASP Configuration display.

   Note: Selecting this option deletes all data in the system ASP. Do not use this procedure unless you have a failed disk unit and there is no immediate replacement for the disk unit.

```
                      Select ASP to Delete Data From
   Type options, press Enter

4=Delete ASP data
Option

ASP  Threshold  Overflow    --Protected--  --Unprotected
                             Size  %Used     Size  %Used
1       90%       No         0.00  0.00%     1200  74.84%
2       90%      Yes         0.00  0.00%      200  99.99%
3       90%      Yes         0.00  0.00%      200  99.99%
```

3. Type a 4 in the Option column to select the ASP that you want to delete the data from and press the Enter key.

   The following display is shown.

```
                          Confirm Delete ASP Data

   Warning:  All data will be deleted from the selected ASPs.
             You have selected to delete data from ASP 1.  This will prevent
             you from changing the disk configuration in some ways until
             the system is IPLed again to DST.

             Press F10 to confirm your choice for 4=Delete ASP data.
             Press F12=Cancel to return to change your choice.

--Protected--  --Unprotected--
Option  ASP  Threshold  Overflow     Size  %Used     Size  %Used
4         1      90%       No           0   0.00     1200    *
```

4. Press F10 (Confirm) to confirm your choice to delete the ASP data.
5. When the delete of the ASP data is complete, you return to the Use Dedicated Service Tools (DST) menu.

Task 3: Removing the disk unit from the auxiliary storage pool configuration

To remove the disk unit from the auxiliary storage pool (ASP) configuration, follow these steps.

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
3. The Remove Units From Configuration display is shown.

```
                      Remove Units from Configuration

Type options, press Enter.
  4=Remove unit from configuration

                 Serial                       Resource
OPT  Unit  ASP   Number      Type   Model     Name            Status
      2     1    10-00A7529  9332   400       DD010           Configured
      3     1    10-00A4936  9332   400       DD012           Configured
      4     1    10-00A4936  9332   400       DD014           Configured
 4    5     1    10-00A7498  9332   400       DD015           Configured
 4    6     1    10-00A7498  9332   400       DD017           Configured
      7     1    10-00A7530  9332   400       DD018           Configured
      8     1    10-00A7530  9332   400       DD021           Configured
```

4. Type a 4 (Remove unit from configuration) in the OPT column for each unit that you want to remove and press the Enter key.

   If the remove operation would leave the ASP with insufficient storage, you receive an error message.

   If you see the Confirm Remove Disk Units display, skip to 6.

   The Confirm Continuation display might be shown before the Confirm Remove Disk Units display if the storage management directories are not usable.

```
                              Confirm Continuation

To proceed, the system must perform internal processing
that might take several minutes during which the system might
appear inactive.

Press Enter to continue.
Press F12=Cancel to return and change your choice.
```

5. Determine whether you want to cancel the procedure or continue.

   If you want to continue, press the Enter key.
6. The Confirm Remove Disk Units display is shown:

```
                            Confirm Remove Disk Units

Removing disk units will take several minutes.

Press Enter to confirm remove of disk units.
Press F9=Capacity information to display the capacity information.
Press F12=Cancel to return to change your choice.


                 Serial                       Resource
OPT   Unit  ASP  Number      Type   Model     Name            Status
 4      5    1   10-00A7498  9332    400      DD010           Configured
 4      6    1   10-00A7498  9332    400      DD012           Configured
```

   Press F9 (Capacity information to display the resulting capacity.

```
                                Resulting Capacity

The configuration change that you requested would result in the
following ASP capacities.

Press Enter to continue.
               -----------Current----------  ----------Modified--------
               --Protected--  -Unprotected-  --Protected-- -Unprotected-
ASP   Threshold  Size  %Used   Size  %Used   Size   %Used   Size   %Used
 1        90%       0   0.00%  1600  52.70%     0    0.00%  1200   70.26%
```

7. Press the Enter key to return to the Confirm Remove Disk Units display.
8. Press the Enter key on the Confirm Remove Disk Units display to remove the selected units.

   The system moves the data off the units selected to be removed to the remaining units in the source ASP. The remove can take several minutes or several hours during which the system appears inactive.

   Notes:

   - The time it takes to remove a unit depends on the disk unit type and model.
   - If the data on the unit being removed is severely fragmented and the amount of storage used is high, the remove operation might take several hours.
9. When the remove operation is complete, you return to the Work with ASP Configuration display.

Press F3 until you return to the Use Dedicated Service Tools (DST) display.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Chapter 8. The restore menu

The restore menu provides many options for recovering information.

Figure 8 on page 201 shows the menu. Options marked with a plus sign (+) require the system to be in a restricted state. When your system is in a restricted state that does not prevent client workstations from attempting to access information.

```
RESTORE                            Restore

Select one of the following:

  Restore Data
     1. Files
     2. Libraries
     3. Documents and folders
     4. Programs
     5. Other objects
     6. Licensed programs
     7. Configuration
 +   8. User profiles
     9. Objects in directories
```

Figure 8. Restore menu–first display

You can page down on the **Restore** menu to see other options:

Restore System and User Data
+ 21. System and user data
+ 22. System data only
+ 23. All user data
Restore Office Data
30. All documents, folders, and mail
31. Documents and folders
32. Mail only
33. Calendars
Restore Libraries
+ 40. All libraries other than system library
41. All IBM libraries other than system library
42. All user libraries
Restore from Different Systems
50. Restore from System/36 format

What the restore menu options do

Here are the commands that the system runs for the menu options that restore either the system, the system data only, or all user data.

The name of the control language (CL) program that the system runs is in parentheses () following the description of the menu option. You can change this CL program if you need different values than the system-supplied defaults.

© Copyright IBM Corp. 1997, 2025

The options that place your system into a restricted state (options 21, 22, 23, and 40) also end TCP before ending the subsystems. The following commands are run as part of these options:

- ENDTCPSVR
- ENDHOSTSVR
- DLYJOB
- ENDTCP
- DLYJOB
- ENDSBS SBS(*ALL)

The two Delay job (DLYJOB) commands allow time for TCP processing to end. You can choose the amount of time that each DLYJOB command waits. The default value is 300 seconds for a total delay of 10 minutes before the End subsystem (ENDSBS) command is run.

Option number	Description and commands

21	System and user data (QMNRSTE):
	`ENDTCPSVR`
	`ENDHOSTSVR`
	`DLYJOB`
	`ENDTCP`
	`DLYJOB`
	`ENDSBS SBS(*ALL) OPTION(*IMMED)`
	`RSTUSRPRF USRPRF(*ALL)`
	`RSTCFG OBJ(*ALL)`
	`RSTLIB SAVLIB(*NONSYS)`
	`RSTDLO DLO(*ALL) SAVFLR(*ANY)`
	`RST DEV('/QSYS.LIB/media-device-name.DEVD')`
	`OBJ(('/*') ('/QSYS.LIB' *OMIT)`
	`('/QDLS' *OMIT))`
	`RSTAUT`
	`STRSBS SBSD(`*controlling subsystem*`)`

22	System data only (QSRRSTI):
	`ENDTCPSVR`
	`ENDHOSTSVR`
	`DLYJOB`
	`ENDTCP`
	`DLYJOB`
	`ENDSBS SBS(*ALL) OPTION(*IMMED)`
	`RSTUSRPRF USRPRF(*ALL)`
	`RSTCFG OBJ(*ALL)`
	`RSTLIB SAVLIB(*IBM)`
	`RST DEV('/QSYS.LIB/media-device-name.DEVD')`
	`OBJ(('/QIBM/ProdData')`
	`('/QOpenSys/QIBM/ProdData'))`
	`STRSBS SBSD(`*controlling subsystem*`)`

23	All user data (QSRRSTU):
	`ENDTCPSVR`
	`ENDHOSTSVR`
	`DLYJOB`
	`ENDTCP`
	`DLYJOB`
	`ENDSBS SBS(*ALL) OPTION(*IMMED)`
	`RSTUSRPRF USRPRF(*ALL)`
	`RSTCFG OBJ(*ALL)`
	`RSTLIB SAVLIB(*ALLUSR)`
	`RSTDLO DLO(*ALL) SAVFLR(*ANY)`
	`RST DEV('/QSYS.LIB/media-device-name.DEVD')`
	`OBJ(('/*') ('/QSYS.LIB' *OMIT)`
	`('/QDLS' *OMIT)`
	`('/QIBM/ProdData' *OMIT)`
	`('/QOpenSys/QIBM/ProdData' *OMIT))`
	`RSTAUT USRPRF(*ALL)`
	`STRSBS SBSD(`*controlling subsystem*`)`

Using restore menu options 21, 22, and 23

To restore information using option 21, 22, or 23 from the Restore menu, follow these steps.

The basic steps are the same for each menu option. Which option or options you use depends on which save menu option was used and what other procedures, if any, you use to save information.

Before you begin

- Clean the read and write heads of the tape unit if you are restoring from a tape drive.

1. Sign on the system using a user profile with sufficient authority to do the restore operation (such as QSECOFR).
2. Ensure that you load the correct volume of your last set of save media and make the device ready. The save media should contain the file labeled QFILEUPR.
   a) If you use tape media, run the Display Tape (DSPTAP) command and specify DATA(*LABELS) to find the file labeled QFILEUPR.
   b) If you use DVD-RAM optical media, perform the following steps:
      i) From a command line run this command: `DSPOPT VOL(*MOUNTED) DEV(OPT01) DATA(*FILATR) PATH('QFILEUPR').`
      ii) If the file is on the media, page down on your display to verify that the file is on the first volume. If the display says `Continued from previous volume...NO`, then the file QFILEUPR is on the first volume of your save media set.
3. Ensure that any device configuration objects not used in the restore operation are varied off.

   You can use the Work with Configuration Status (WRKCFGSTS) command to display the status of devices.
4. Ensure that the devices that you are using for the restore operation (workstations, devices, and device controllers) are varied on.

   These configuration objects are excluded from the restore operation (message CPF379C in the job log).
5. Display the Restore menu: `GO RESTORE`.
6. If you want to do an attended restore, skip to step “7” on page 203.

   In most cases, you should do an attended restore operation to monitor for messages and correct any problems that occur. This helps your system return to operation as quickly as possible. If you want to do an unattended restore, do the following steps. This prevents your restore operation from stopping because of unanswered messages:
   a. Display the reply list sequence numbers to find what numbers are available for use:

      ```
      WRKRPYLE
      ```

   b. If MSGID(CPA3709) is not already in your reply list, add it. For *xxxx*, substitute an unused sequence number from 1 through 9999:

      ```
      ADDRPYLE SEQNBR(xxxx) MSGID(CPA3709) RPY('G')
      ```

   c. Change your job to use the reply list:

      ```
      CHGJOB INQMSGRPY(*SYSRPYL) BRKMSG(*NOTIFY)
      ```
7. Select the option (21, 22, or 23) from the Restore menu.

```
                          Specify Command Defaults

Type choices, press Enter.

  Devices  . . . . . .   TAP01        Names
                         ----------
                         ----------
                         ----------

  Prompt for commands  . . . . .   Y            Y=Yes, N=No

  Message queue delivery . . . .   *BREAK       *BREAK, *NOTIFY

  Restore to different system      N            Y=Yes, N=No

  Spooled file data  . . . . . .   *NEW         *NEW, *NONE

  End TCP/IP wait time . . . . .   300          *NONE, 1-9999 seconds


  Print joblog . . . . . .   Y            Y=Yes, N=No

  Start controlling subsystem      Y            Y=Yes, N=No
```

8. Type your choices for the *Devices* prompt.

   You can specify up to four device names. If you specify more than one device, the system automatically switches to the next device after it finishes reading the current save media.
9. Type your choice for the *Prompt for commands* prompt.

   Specify N (No) if you want to run an unattended restore. Specify Y (Yes) if you want to change the defaults on the RSTxxx commands.
10. Type your choice for the *Message Queue Delivery* prompt.

    The *Message Queue Delivery* prompt only applies if messages are sent to the QSYSOPR message queue. Specify *NOTIFY if you want to do an unattended restore. This prevents communications messages from stopping the restore procedure. For example, messages that request a new volume to be loaded interrupt the restore operation because they are associated with the job, they have nothing to do with setting the delivery mode of the message queue. You cannot continue until you reply to these messages. Specify *BREAK, if you want to be interrupted for severity 99 messages that are sent to xxxx that require a reply.

    Specify *BREAK, if you want to be interrupted for severity 99 messages that require a reply.
11. Type your choice for the *Restore to different system* prompt. If you specify Y (Yes), the following values will be specified.

    The system requires these values in order to perform a system recovery to a different system or to a different logical partition.

    - SRM(*NONE) will be specified on the Restore Configuration (RSTCFG) command
    - ALWOBJDIF(*ALL) or ALWOBJDIF(*COMPATIBLE) will be specified on the restore commands
    - MBROPT(*ALL) will be specified on the Restore Library (RSTLIB) command
12. Type your choice in the *Spooled file data* prompt.

    Specify *NEW to restore spooled file data that was saved and does not already exist on the system. Specify *NONE if you do not want to restore spooled file data.
13. Type your choice for the *End TCP/IP wait time* prompt. Specify the number of seconds to delay the job to allow TCP/IP processing to end before ending all subsystems. Specify *NONE to end all subsystems without ending TCP/IP processing first.
14. Type your choice for the *Print joblog* prompt. Specify Y (Yes) if you want to print the joblog when the menu operation ends.
15. Type your choice for the *Start controlling subsystem* prompt. Specify Y (Yes) if you want to start the controlling subsystem when the menu operation ends.
16. After you type your choices, press the Enter key.
17. If you responded Y to the *Prompt for commands* prompt, you are shown the End Subsystem display after TCP has ended.

Type any changes and press the Enter key. While the system is ending subsystems, you see and respond to these messages:

- `CPF0994  ENDSBS(*ALL) command being processed`

  . Press the Enter key.
- `CPF0968  System ended to restricted condition`

  . Press the Enter key.

If you responded `N` to the *Prompt for commands* prompt, skip to step "19" on page 206.

18. When the system is ready to perform each major step in the restore process, you are shown the prompt display for that step.

    The time between displays might be quite long.

    For option 21, you are shown these displays:

    - `ENDSBS SBS(*ALL) OPTION(*IMMED)`
    - `RSTUSRPRF USRPRF(*ALL)`
    - `RSTCFG OBJ(*ALL)`
    - `RSTLIB SAVLIB(*NONSYS)`
    - `RSTDLO DLO(*ALL) SAVFLR(*ANY)`
    - `RST DEV('/QSYS.LIB/media-device-name.DEVD') OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))`
    - `RSTAUT`
    - `STRSBS SBSD(`*`controlling subsystem`*`)`

    For option 22 (System data only) you are shown these displays:

    - `ENDSBS SBS(*ALL) OPTION(*IMMED)`
    - `RSTUSRPRF USRPRF(*ALL)`
    - `RSTCFG`
    - `RSTLIB SAVLIB(*IBM)`
    - `RST DEV('/QSYS.LIB/media-device-name.DEVD') OBJ(('/QIBM/ProdData') ('/QOpenSys/QIBM/ProdData'))`
    - `STRSBS SBSD(controlling-subsystem)`

    For option 23 (All user data) you are shown these displays:

    - `ENDSBS SBS(*ALL) OPTION(*IMMED)`
    - `RSTUSRPRF USRPRF(*ALL)`
    - `RSTCFG`
    - `RSTLIB SAVLIB(*ALLUSR)`
    - `RSTDLO DLO(*ALL) SAVFLR(*ANY)`
    - `RST DEV('/QSYS.LIB/media-device-name.DEVD') OBJ(('/*') ('/QSYS.LIB' *OMIT) ('QDLS' *OMIT) ('/QIBM/ProdData' *OMIT) ('/QOpenSys/QIBM/ProdData' *OMIT))`
    - `RSTAUT`
    - `STRSBS SBSD(controlling-subsystem)`

    Type your changes, if any, when the display is shown and press the Enter key.

    Note: The Restore Authority (RSTAUT) command will run immediately after the Restore Object (RST) commands when you use option 21 or option 23. If you use option 22 only, you must run the RSTAUT command. If you have additional restore operations to run, you might need to restore security data and restore authority again after those restore operations.

19. When the system sends a message asking you to load the next volume, load the next media volume and respond to the message.

    If a media error occurs...

    If an error occurs during the restore operation, see "Recovery from an unsuccessful restore operation" on page 41.

    If an unrecoverable error occurs when running the RSTDLO DLO(*ALL) SAVFLR(*ANY) command, see "Recovering from an error while restoring document library objects" on page 42.
20. If you used the distribution media to restore the operating system, some information was not restored.

    If you are restoring to a different system or different logical partitions, the network attributes might have been reset to the IBM-supplied defaults. You must create or change this information again. You should have lists of this information that were created at the time you performed your save operation.

    The following might need to be created or changed:

    - Configuration lists
    - Network attributes
    - Edit descriptions
    - Reply list entries
    - IBM-supplied subsystem descriptions

    a. For the configuration lists, perform the following task:

       Use the Work With Configuration Lists (WRKCFGL CFGL(*ALL)) command to create the configuration lists to match the information in your list.

    b. For network attributes, perform the following task:

       Use the Change Network Attributes (CHGNETA) command to change the network attributes to match the information in your list.

    c. For edit descriptions, perform the following task:

       Use the Work with Edit Descriptions (WRKEDTD EDTD(*ALL)) command to create edit descriptions to match the information in your list.

    d. For reply list entries, perform the following task:

       Use the Add Reply List Entry (ADDRPYLE) command to add reply list entries to match the information in your list.

    e. For IBM-supplied subsystem descriptions, perform the following task:

       Use the Work with Subsystem Descriptions (WRKSBSD SBSD(*ALL)) command to change the IBM-supplied subsystem descriptions to match the information in your list.
21. If you need to restore independent ASPs, see "Recovering an independent auxiliary storage pool" on page 192. Skip this step if you use "Checklist 21: Recovering your entire system after a complete system loss including independent auxiliary storage pools" on page 97.
22. This completes the restore operation.
23. If you are unsure what the QSECOFR password is, change it now.

    To see if the password has expired, type the following command:

    ```
    DSPUSRPRF QSECOFR
    ```

    The passwords from your save media are now the current passwords. If the password expiration is active for the QSECOFR user profile, you will see the expiration date on the Date password expired field. If the date is the current system date or prior, change the password for the QSECOFR user profile.
24. Check the job log to ensure all objects were restored.

The job log contains information about the restore operation. To verify that all objects were restored, you should spool the job log for printing, along with the job's remaining spooled output, if any.

```
DSPJOBLOG * *PRINT
```

or

```
SIGNOFF *LIST
```

Message CPC3703 is sent to the job log for each library that was successfully restored. Message CPF3773 is sent to tell you how many objects were restored. It also tells you how many objects were not restored. Objects are not restored for various reasons. Check for any error messages, correct the errors, and then restore those objects from the media.

Related concepts

Choosing the procedure to recover user information
When your system is running normally, you are ready to recover user information.

Chapter 9. Restoring specific types of information

This topic describes procedures for restoring particular types of information on the system. It also describes considerations when you restore particular types of information, whether you restore the information by using menu options or commands. The topics are presented in the same order as the recovery operations should occur.

Recovering system information

You can customize some system information, such as edit descriptions and network attributes. When you run the Save System (SAVSYS) command this system information is saved. Whenever library QUSRSYS is saved with the Save Library (SAVLIB) or Save Changed Objects (SAVCHGOBJ) functions, the current system information will be retrieved and saved. The system information saved will be the same as that retrieved by the Retrieve System Information (RTVSYSINF) function.

If you have SAVSYS media and need to restore system information, follow the procedure that is described in Chapter 5, “Restoring the operating system,” on page 141 topic. Do an abbreviated installation of the operating system.

If a system needs to be reloaded from distribution media, the system information can be updated by restoring the user information and then running UPDSYSINF LIB(QUSRSYS).

If a system needs to be reloaded from distribution media and you do not have a current save of your QUSRSYS library, you need to reconstruct system information. Find the most recent system information listings you have. Table 46 on page 209 shows the commands for changing the system information to the correct values:

Table 46. Commands for changing system information

Information type	Command
Access path recovery times[1]	EDTRCYAP
Configuration lists	WRKCFGL
Edit descriptions	WRKEDTD
IBM-supplied subsystem descriptions	WRKSBSD
Network attributes	CHGNETA
Reply list entries	ADDRPYLE
Service attributes	CHGSRVA
System values	WRKSYSVAL

1

When you reset your access path recovery times, ensure that the auxiliary storage pool (ASP) configuration matches the configuration at the time that you printed the recovery times. If it does not, make a note to reset your access path recovery times after ASP configuration.

Related concepts

Recovering TCP configuration information

TCP/IP configuration information consists of a number of different physical files, logical files, validation lists, stream files, and environment variables. Whenever library QUSRSYS is saved with the Save Library (SAVLIB) or Save Changed Objects (SAVCHGOBJ) functions, the current TCP/IP configuration information is retrieved and saved.

Related information

Printing system information

© Copyright IBM Corp. 1997, 2025

Recovering TCP configuration information

TCP/IP configuration information consists of a number of different physical files, logical files, validation lists, stream files, and environment variables. Whenever library QUSRSYS is saved with the Save Library (SAVLIB) or Save Changed Objects (SAVCHGOBJ) functions, the current TCP/IP configuration information is retrieved and saved.

The TCP configuration information that is saved is the same as that retrieved by the Retrieve TCP/IP Information (RTVTCPINF) function.

If a system needs to be reloaded from distribution media, the TCP configuration information can be updated by restoring the user information and then running UPDTCPINF LIB(QUSRSYS).

Note: The TCP/IP configuration information update should be performed after you have installed IBM TCP/IP Connectivity Utilities for i (5770-TC1).

If a system needs to be reloaded from distribution media and you do not have a current save of your QUSRSYS library, you need to reconstruct TCP configuration information.

Related tasks

Recovering system information

You can customize some system information, such as edit descriptions and network attributes. When you run the Save System (SAVSYS) command this system information is saved. Whenever library QUSRSYS is saved with the Save Library (SAVLIB) or Save Changed Objects (SAVCHGOBJ) functions, the current system information will be retrieved and saved. The system information saved will be the same as that retrieved by the Retrieve System Information (RTVSYSINF) function.

Sequence for restoring security information

Recovering your system often requires restoring data and associated security information. It is essential that you restore security information in the correct sequence. Otherwise, object ownership and authority information is not restored correctly and your applications might not run correctly.

Security information on your system consists of the following items:

- User profiles and group profiles
- Authorization lists
- Authority holders
- Authority information that is stored with objects:
  - Owner
  - Owner authority
  - Primary group
  - Primary group authority
  - Public authority
  - Authority list
- Private authorities
- Digital Certificate Manager (DCM) data
- Function usage information

The recovery checklists include the correct sequence of steps for restoring security information. If you are developing your own restore procedure, restore security information in the following sequence:

1. Restore user profiles.

   The user profile that owns an object must exist before the object can be restored.

   If you restore all user profiles (RSTUSRPRF USRPRF(*ALL)), you also restore authorization lists, authority holders, and other security information. Authorization lists and authority holders must also be on the system before you restore objects.

2. Restore objects (RSTCFG, RSTLIB, RSTOBJ, RSTDLO or RST).
   This restores ownership and the authority information that is stored with the object.
3. Use the Restore Authority (RSTAUT) command to restore the private authorities to objects.

Note: Alternatively, you can restore private authorities for objects by specifying the PVTAUT(*YES) parameter on the restore command.

Related information

Backup and recovery of security information

Restoring user profiles

You can restore a single user profile, a list of user profiles, or all user profiles. You restore a user profile to move a user from one system to another system and to recover a damaged user profile.

You can use the *NEW value on the USRPRF parameter of the Restore User Profiles (RSTUSRPRF) command to restore only user profiles that are new to your system.

If you restore all user profiles, their password and group connections are restored. If you restore new user profiles or individual user profiles, you can specify SECDTA(*PWDGRP) to restore their passwords and group connections.

You might also find the *NONE value beneficial if you only want to restore the data needed to verify signatures, and not all of the actual user profiles.

The OMITUSRPRF parameter allows you to limit the number of user profiles you restore. You can specify a list of up to 300 specific or generic user profile values that will not be restored. This value is helpful if you are restoring a subset of user profiles.

The SAVASPDEV parameter allows you to limit the private authorities that are restored based on auxiliary storage pools.

When you restore all user profiles, the Digital Certificate Manager (DCM) data and function usage information are restored unless you specify a value on the OMITSECDTA parameter. If you want to omit DCM data from the restore, specify the *DCM value on the OMITSECDTA parameter on the RSTUSRPRF command. To omit authority lists, specify the *AUTL value on the OMITSECDTA parameter. To omit function usage information, specify *FCNUSG on the OMITSECDTA parameter.

The following values are useful if you are merging user profiles from multiple systems onto a single system:

- *NEW value on the USRPRF parameter
- *PWDGRP value on the SECDTA parameter
- *DCM, *AUTL, *FCNUSG values on the OMITSECDTA parameter

Note: You cannot delete an IBM-supplied user profile if it is damaged. You must restore the operating system again by way of an abbreviated install to recover a damaged IBM-supplied user profile.

Table 47. How user profiles are restored

Method	Restricted state?
RSTUSRPRF command [1,3]	No
Restore menu option 8 [1,3]	No
Restore menu option 21 [1,2]	Yes
Restore menu option 22 [1,2]	Yes
Restore menu option 23 [1,2]	Yes

Table 47. How user profiles are restored (continued)

Method	Restricted state?
1 You must have *SAVSYS special authority. You must have *ALLOBJ special authority to specify a value other than *NONE on the ALWOBJDIF parameter.	
2 These menu options restore all user profiles.	
3 You need to put the system in a restricted state if you specify USRPRF(*ALL).	

Do this to restore all user profiles

1. Sign on as QSECOFR.
2. Ensure that the system is in a restricted state.
3. Find the most recent media volume that has your user profiles. It might be a SAVSYS volume or a SAVSECDTA volume. The name of the file on the media volume is QFILEUPR.
4. If you are using a SAVSYS media volume, type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*LEAVE)
```

   If you are using a SAVSECDTA media volume, type the following command:

```
RSTUSRPRF DEV(media-device-name) USRPRF(*ALL)
          ENDOPT(*UNLOAD)
```

In this example, the user profiles are copied from the source system to the target system without affecting the Digital Certificate Manager (DCM) setup and signature verification store on the target system.

Table 48. Copying user profiles without affecting the Digital Certificate Manager information

Source system	Target system
SAVSECDTA	RSTUSRPRF USRPRF(*ALL) OMITSECDTA(*DCM)

Related concepts

Guidelines for restoring information from the development system
Follow these guidelines when you restore information from the development system.

Related tasks

Putting your system in a restricted state
Many recovery procedures require that your system have no other activity on it. When no subsystems except the controlling subsystem are active on your system, it is in a *restricted state*.

Related information

Restoring user profiles
Backup and recovery considerations for DCM

What happens when you restore user profiles

When you restore a user profile, you restore all the attributes of the profile that you see on the Display User Profile display. The system builds a working table that holds that user's private authorities to objects.

You must use the Restore Authority (RSTAUT) command to restore the user's private authorities. If you specify SECDTA (*PVTAUT), then only the working tables that hold the user's private authorities are restored. The user profiles themselves are not restored.

Some values in a user profile might be changed when it is restored. Table 49 on page 213 shows the actions the system takes when you restore user profiles:

Table 49. Results of restoring user profiles

	Restore procedure used		
User profile attributes	**Restore *ALL user profiles**	**Restore individual user profile that already exists on system**	**Restore individual user profile that does not exist on system**
Group profile (GRPPRF)	Value is restored from media	Value on system is not changed [1]	Value is set to *NONE [1]
Owner (OWNER) of new objects	Value is restored from media	Value on system is not changed	Value is set to *USRPRF
Group authority (GRPAUT) to new objects	Value is restored from media	Value on system is not changed [1]	Value is set to *NONE [1]
Password	Value is restored from media	Value on system is not changed [1]	Value is set to *NONE [1]
Document Password	Value is restored from media	Value on system is not changed [1]	Value is set to *NONE [1]
Date password was last changed	Value is restored from media	Value on system is not changed [2]	The current date is used.
Owner of user profile	See "How the system establishes ownership for restored objects" on page 216.	Value on system is not changed	Value is restored from the media. If the owning profile does not exist, ownership is assigned to the QDFTOWN user profile.
Primary group of user profile	See "How the system establishes the primary group for restored objects" on page 216	Value on system is not changed.	Value is restored from the media. If the primary group does not exist, the value in the user profile is set to *NONE.
*ALLOBJ special authority	See "What you should know about restoring user profiles" on page 214.	See "What you should know about restoring user profiles" on page 214.	See "What you should know about restoring user profiles" on page 214.
User identification number (*UID*)	Results depend upon whether the user profile already exists on the system or not at the time it is restored. Actions are the same as for individual user profiles.	Value on system is not changed.	Value is restored from the media unless it is a duplicate of a *UID* on the system. In that case, a new *UID* is generated.

Table 49. Results of restoring user profiles (continued)

	Restore procedure used		
User profile attributes	**Restore *ALL user profiles**	**Restore individual user profile that already exists on system**	**Restore individual user profile that does not exist on system**
Group identification number (*GID*)	Results depend upon whether the user profile already exists on the system or not at the time it is restored. Actions are the same as for individual user profiles.	Value is restored from the media unless: • The user profile on the media has a GID of 0 (*NONE). • The user profile on the system has a GID that is not 0. • The user profile on the system is a member of a group. • The GID for the user profile on the media is a duplicate of a GID for a different user profile on the system.	Value is restored from the media unless it is a duplicate of a *GID* on the system. In that case, a new *GID* is generated.
User expiration fields	See "What you should know about restoring user profiles" on page 214.	See "What you should know about restoring user profiles" on page 214.	See "What you should know about restoring user profiles" on page 214.

1
If you specify SECDTA (*PWDGRP), the value is restored from media.

2
If you specify SECDTA(*PWDGRP), the current date is used.

Related concepts

Restoring object authorities

When you restore user profiles, the system builds authority reference tables for each user profile that you restore. The authority reference tables temporarily hold the user's private authorities to objects.

What you should know about restoring user profiles

When you are restoring user profiles from a source system to a target system, you must make sure that the password level values (QPWDLVL) are compatible.

For example, restoring a user profile from the source system with a password value of 2 might result in a non-valid password on the target system with a password value of 0 or 1. Password level 2 allows more characters than password level 0 or 1.

Restoring all profiles: When you restore all profiles, the system does not first delete all profiles, authorization lists, and authority holders on the system. Therefore, the result is both of the following:

- All the profiles, authorization lists, and authority holders on the media.
- Any profiles, authorization lists, and authority holders on the system that were not on the save media.

Restoring all profiles is the only way to restore authorization lists and authority holders. If an object does not exist on the system, the link between an authorization list and the object is restored when you restore the object. However, objects in library QSYS are restored before the user profiles and the authorization lists. Therefore the links between authorization lists and objects in library QSYS are not restored when the objects are restored. You can restore the links for objects in QSYS if you saved security data from IBM i 7.3 or higher. To restore these links, run RSTUSRPRF USRPRF(*ALL) or RSTUSRPRF USRPRF(*NEW) and then run RSTAUT.

Security note: If the IBM-supplied user profiles have the default passwords on your save media, they will again have default passwords after you perform the restore operation. This is a security exposure. After a restore operation, verify that the IBM-supplied user profiles do not have the default passwords.

Restoring *ALLOBJ special authority: *ALLOBJ special authority is removed from user profiles being restored to a system at security level 30 or higher in either of these situations:

- The profile was saved from a different system, and the person doing the restore does not have *ALLOBJ and *SECADM authority.
- The profile was saved from the same system or a different system at security level 10 or 20.

The systems keeps *ALLOBJ special authority for the following system user profiles:

- QSYS
- QSECOFR
- QLPAUTO
- QLPINSTALL

Moving users to another system: To transfer user profiles and their authorities to another system, perform these tasks:

1. Save the user profiles and authorities by using the Save Security Data (SAVSECDTA) command.
2. Restore the user profiles by using RSTUSRPRF USRPRF(*ALL) ALWOBJDIF(*ALL). The following are some considerations for restoring authority information for user profiles:
   - Use the USRPRF(*NEW) parameter to restore only user profiles which do not currently exist on the target system.
   - Use the OMITUSRPRF parameter to omit user profiles that you do not want to restore. To omit Digital Certificate Manager (DCM) data, specify the *DCM value on the OMITSECDTA parameter on the RSTUSRPRF command. To omit authority lists, specify the *AUTL value on the OMITSECDTA parameter. To omit function usage information, specify *FCNUSG on the OMITSECDTA parameter. These values are useful if you are merging user profiles from multiple systems onto a single system.
3. Restore the needed objects by using the Restore Library (RSTLIB), Restore Object (RSTOBJ), Restore Object (RST), or Restore Document Library Object (RSTDLO) commands by specifying ALWOBJDIF(*ALL) or ALWOBJDIF(*COMPATIBLE). For the RSTLIB and RSTOBJ commands, it is preferable to specify ALWOBJDIF(*COMPATIBLE).
4. Restore the private authorities of the user profiles by using the Restore Authority (RSTAUT) command.

User expiration fields: The following rules apply to the user expiration functions when profiles are restored:

- When a user profile is restored (all profiles or individual profile) that exists on the system, the restore operation cannot change the existing user expiration fields.
- When a user profile is restored (all profiles or individual profile) that does not yet exist on the system, all fields in the user profile are restored from the save media, including the user expiration interval and user expiration date fields:
  - If the profile is enabled and the user expiration date is past, the user profile is set to disabled and CPF2271 diagnostic message is sent.
  - If the profile is enabled and the user expiration date has not past, the job scheduler entry is added.

Related concepts

Release-to-release support

The release-to-release support on the IBM i environments allows you to move data from the current release to a previous-release system. This support also allows you to move data from a previous-release system to a current-release system.

Related information

Security reference PDF

How the system establishes ownership for restored objects

Most objects on the system have an owner.

When you restore an object, the system determines what profile owns the restored object by using the following rules:

- Ownership is restored to that profile if the profile that owns the object is on the system.
- If the owner profile does not exist on the system, ownership of the object is given to the QDFTOWN (default owner) user profile.
- If the owner on the system is different from the owner on the save media, the object is only restored if ALWOBJDIF(*OWNER), ALWOBJDIF(*COMPATIBLE), or ALWOBJDIF(*ALL) is specified. In that case, the object is restored and the owner on the system is used.
- See "How the system restores programs" on page 256 for more considerations for restoring programs.

Related concepts

What happens when you restore objects
When you restore an object, the system takes different actions depending on several conditions.

Related tasks

Comparing file attributes during a restore operation
When you restore a database file or member that exists on the system, the system expects the creation dates for the system copy and the media copy to be the same. If they are not the same, the system cannot ensure that the contents of the saved copy match the format of the copy on the system.

How the system establishes the authorization list for a restored object

This topic shows what happens when you restore an object that already exists if the object is linked to an authorization list. These rules do not apply when you are restoring a document or a folder.

Table 50. Restoring an object linked to an authorization list

Authorization list on system and media	Value of ALWOBJDIF parameter	Result
Same	ANY	Data restored; link to authorization list not changed.
Different	*NONE	Object is not restored
Different	*ALL, *AUTL, *COMPATIBLE	Data restored; linked to authorization list associated with existing object.

Related concepts

What happens when you restore objects
When you restore an object, the system takes different actions depending on several conditions.

How the system establishes the primary group for restored objects

Many objects on the system have a primary group. When you restore an object, the system determines the primary group for the object by using a list or rules.

The list of rules that the system uses is as follows:

- If the profile that is the primary group for the object is on the system, that profile is the primary group for the object.
- If the profile that is the primary group for the object is not on the system, the primary group is set to *NONE. Message CPI380E is sent to the job log.
- If the object exists on the system and the primary group on the system is different from the primary group on the save media, the system does not restore the object unless you specify ALWOBJDIF(*PGP), ALWOBJDIF(*COMPATIBLE) or ALWOBJDIF(*ALL). In that case, the system restores the object with the primary group on the system.

Restoring object authorities

When you restore user profiles, the system builds authority reference tables for each user profile that you restore. The authority reference tables temporarily hold the user's private authorities to objects.

Possible method	Restricted state?
RSTAUT command [1]	No
PVTAUT parameter on the restore command [2]	No
Restore menu option 21 [1]	Yes
Restore menu option 22 [1]	Yes
Restore menu option 23 [1]	Yes

[1] You must have *SAVSYS special authority.
[2] You must have *ALLOBJ special authority.

Related concepts

What happens when you restore user profiles

When you restore a user profile, you restore all the attributes of the profile that you see on the Display User Profile display. The system builds a working table that holds that user's private authorities to objects.

Overview of restoring authorities

When you run the Restore Authority (RSTAUT) command, the system restores authority for a specific user profile, a list of user profiles, or all user profiles.

If you restore authority for all users, the RSTAUT command restores authority by using every authority reference table that it finds on the system. If you restore a single user profile, you can also use RSTAUT and specify that profile name to restore authorities for that user profile.

Note: Do not run the Reclaim Storage (RCLSTG) command between the Restore User Profile (RSTUSRPRF) command and the RSTAUT command. The RCLSTG command deletes authority reference tables.

You can also restore authority for a specific profile or a list of profiles. For example, if you restored a single user profile to the system, you can also use the RSTAUT command and specify that profile name.

When you run RSTAUT USRPRF(*ALL), you receive the status message CPI3821 that indicates the number of user profiles for which authorities are restored.

You can run the RSTAUT command regardless of whether a system is in a restricted state. Running RSTAUT on systems in a restricted state versus nonrestricted state have differences with system performance, job log appearance, and object availability.

During recovery, restoring authorities is the last thing that you do before you do an initial program load (IPL). If you restore authorities with default parameter SAVASPDEV(*ALLAVL), but you have more restore operations to run, you might need to restore security data and authority again.

Related information

Restoring authority

Restoring authorization lists

Restoring private authorities in a full-system recovery

The recommended method for restoring private authorities in a full-system recovery is to use the Restore Authority (RSTAUT) command.

For security data that was saved by the Save System (SAVSYS) or Save Security Data (SAVSECDTA) command, use the Restore User Profiles (RSTUSRPRF) and Restore Authority (RSTAUT) commands to restore the private authorities along with the data. This method is recommended for recovering an entire system.

To restore private authorities in a full-system recovery, follow these steps:

1. Restore user profiles.

   The user profile that owns an object must exist before the object can be restored.

   If you restore all user profiles (RSTUSRPRF USRPRF(*ALL)), you also restore authorization lists, authority holders, and other security information. Authorization lists and authority holders must also be on the system before you restore objects.
2. Restore objects (RSTCFG, RSTLIB, RSTOBJ, RSTDLO or RST).

   This restores ownership and the authority information that is stored with the object.
3. Use the Restore Authority (RSTAUT) command to restore the private authorities to objects.

Note: To save and restore private authorities as you save and restore objects, specify *YES for the PVTAUT parameter on the save and restore commands. You should not use this method for full-system or large-scale recovery of user data for performance reasons.

Restoring private authorities for selected objects

If you need to recover or migrate a few objects when the user profiles already exist on the target system, the quickest method for restoring private authorities for selected objects is to specify PVTAUT(*YES) with any of the save and restore commands.

Prerequisite: You need save system (*SAVSYS) or all object (*ALLOBJ) special authority to save private authorities. You need *ALLOBJ special authority to restore private authorities.

To save and restore private authorities as you save and restore objects, specify *YES for the PVTAUT parameter on the save and restore commands. Although saving private authorities increases the amount of time it takes to save the objects, it simplifies the recovery of the objects. Using the PVTAUT(*YES) parameter is recommended for restoring specific objects, but it is *not* recommended for recovering the entire system or a large-scale recovery of user data.

This procedure saves the object data and the private authorities for objects in the user-defined file system (UDFS) on an independent ASP, and restores them to another independent ASP. Use this method only if you are saving and restoring a few objects:

1. Save the data and private authorities for the objects:

```
SAV  DEV('device-path-name')
     OBJ('/DEV/asp1/udfs-name.udfs') PVTAUT(*YES)
```

2. Restore the data and private authorities to the new independent disk pool:

```
RST  DEV('device-path-name') OBJ(('/DEV/asp1/udfs-name.udfs'
     *INCLUDE '/DEV/asp2/udfs-name.udfs')) PVTAUT(*YES)
```

In this example, the objects and their private authorities are saved and restored. The user profiles already exist on the target system.

Table 51. Saving and restoring private authorities using the PVTAUT parameter	
Source system	**Target system**
SAVLIB ... PVTAUT(*YES)	RSTLIB ...PVTAUT(*YES)
SAVOBJ ... PVTAUT(*YES)	RSTOBJ ... PVTAUT(*YES)
SAVCHGOBJ ... PVTAUT(*YES)	RSTOBJ ... PVTAUT(*YES)
SAV ... PVTAUT(*YES)	RST ... PVTAUT(*YES)
SAVRSTLIB ... PVTAUT(*YES)[1]	
SAVRSTOBJ ... PVTAUT(*YES)[1]	
SAVRSTCHG ... PVTAUT(*YES)[1]	
SAVRST ... PVTAUT(*YES)[1]	
[1] The SAVRST*xx* commands perform the save and restore operation with one command.	

Related tasks

Transferring a library to a different auxiliary storage pool
Use this procedure to move a library to a different auxiliary storage pool (ASP).

Restoring authority on a system in a nonrestricted state

The Restore Authority (RSTAUT) command uses prestarted jobs in order to process more than one user's authorities at a time. The prestarted jobs that are used by RSTAUT use subsystem description QSYSWRK in library QSYS, program QSRRATBL in library QSYS, and class QINTER in library QGPL.

There are several advantages to running the RSTAUT command on a system in a nonrestricted state. These advantages are as follows:

- Because more than one user's authority reference tables are being processed at a time, the RSTAUT command on a system in a nonrestricted state is up to 30% faster in most cases than the same command on a system in a restricted state. Generally, the more user profiles for which RSTAUT is being run, the greater the performance gain for the RSTAUT command overall.
- Subsystems do not have to be terminated when one or more user profiles are restored without a full system recovery.
- Authority reference tables are not always deleted after RSTAUT is run for a user profile. If all private authorities are successfully granted or an abnormal error occurs then the authority reference table is deleted. Authority reference tables are also deleted if you create a data area named QSRCLRAUTS and it exists in the library list. However, if some of the private authorities are not granted for any reason such as 'object not found' or 'object in use', then the entries for those private authorities that were not granted are kept in the authority reference table, and the RSTAUT command can be run again for the user profile to try granting the failed private authorities before the next restore of the user profile.

There are also some limitations to running the RSTAUT command on a system in a nonrestricted state. These limitations are as follows:

- Because the system is not in a restricted state, all objects must be locked by RSTAUT. This means that several objects might be in use during the processing of any authority reference table. If the RSTAUT command is unable to lock an object, a diagnostic message of CPF3736 or CPD3776 is sent to the job log of the prestarted job for each object that could not have authority granted. This is most likely to occur when the object is a user profile or a message queue. Because private authorities that are not granted are kept in the authority reference table, the RSTAUT command might be run again to grant authorities to objects that were in use.

  You might also receive CPD3776 if you use a product that has objects or directories with the "Allow save" attribute set to "No." If this scenario is true, then the CPD3776 messages can be ignored.

- If you are running RSTAUT for a large group of user profiles that have private authorities to the same few objects, you must put the system in a restricted state before running the RSTAUT command. This action minimizes the number of objects in use and consequently minimize the number of objects that are found locked by the RSTAUT command.
- Only one RSTAUT command can be run on a system at a time.

What you should know before running the Restore Authority command

There are some general facts you should be aware of when running the Restore Authority (RSTAUT) command on a system in a nonrestricted state.

- This command might take a long time depending on how many private authorities that you have on your system.
- During system recovery, you should not start all subsystems and allow all users to sign on and use the system while RSTAUT is being run. The only subsystem that is needed to run RSTAUT in a nonrestricted state is QSYSWRK. Allowing all users access to the system before RSTAUT completes might cause many more objects to be locked, preventing a larger number of the private authorities from being restored.
- All private authorities for all authority reference tables that are processed by one prestarted job might not be successfully regranted. If they are regranted, then the message logging level that is used for that prestarted job will be the same as the logging level that is used by the user's main job.

  If one or more private authorities for an authority reference table are not successfully regranted, then LOG(*SECLVL) will be used for the message logging for that prestarted job.

  For example, you can run the RSTAUT command with the system default logging level of LOG(4 0 *NOLIST). All prestarted jobs that are run by RSTAUT that have all private authorities regranted successfully will use the LOG(4 0 *NOLIST) message logging level. The job log will not remain on the system after the prestarted job completes. All prestarted jobs that are run by RSTAUT, that have one or more private authorities that are not regranted, will instead use the LOG(4 0 *SECLVL) logging level. The job log will remain on the system after the prestarted job completes.

  Do not cancel any of the prestarted jobs RSTAUT has started. Doing so will cause the entire RSTAUT command to cancel, similar to cancelling a RSTAUT command on a restricted state system.
- One user's authority reference tables are always processed by a single prestarted job.
- If the authority reference tables are small, one prestarted job might process more than one user's authorities.

Subsystem QSYSWRK must be started in order for the prestarted jobs to start. The RSTAUT command will start several prestarted jobs at once, and assign the restore of authorities for one or more user profiles to each of the prestarted jobs. During the RSTAUT command, when the prestarted jobs are running, an entry will appear for each prestarted job on the Work with Active Jobs display.

```
                           Work with Active Jobs                MYSYSTEM
                                                          05/01/97  16:02:05
CPU %:   26.5     Elapsed time:  00:00:31      Active jobs: 94

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

Opt  Subsystem/Job  User        Type  CPU %  Function        Status
__   QSYSWRK        QSYS         SBS     .1                  DEQW
__   QSRRATBL       QUSER          PJ    1.2                  RUN
__   QSRRATBL       QUSER          PJ    1.0                  RUN
```

If subsystem QSYSWRK is active but the prestarted jobs cannot be started for any reason, you should receive messages in your job log, including escape message CPF386D, stating why the prestarted jobs cannot be started.

Job log considerations

The job logs that are generated by a Restore Authority (RSTAUT) command running on a system in a nonrestricted state are significantly different than those for a system in a restricted state.

When the RSTAUT command runs on a system in a restricted state, one job log is generated. When the RSTAUT command runs on a system in a nonrestricted state, each prestarted job (run by RSTAUT) generates a job log that contains some of the information that is found in the single job log that is produced on a system running in a nonrestricted state. If only one user profile is being restored, only one job log is produced.

You might encounter a situation where job logs that contain diagnostic messages from prestarted jobs that ran during the RSTAUT get deleted. If this occurs, you can rerun the RSTAUT command at any time before running another Restore User Profiles (RSTUSRPRF) command or Reclaim Storage (RCLSTG) command. The system will attempt to regrant the failed private authorities, and will generate new job logs.

Figure 9 on page 221 through Figure 11 on page 221 show a sample job log and message information for a RSTAUT USRPRF(QPGMR) command run on a system in a restricted state.

```
>RSTAUT USRPRF(QPGMR)
Authority not restored for user QPGMR.
Some authorities not restored for user profile QPGMR.
Not all user profiles had all authorities restored.
```

Figure 9. Sample job log for RSTAUT on a system in a restricted state

The expanded text for message CPF3736 appears as follows:

```
              Additional Message Information

Message ID......:  CPF3736        Severity......:  20
Message type....:  Diagnostic
Date sent ......:  04/24/97     Time sent......:  19:35:17

Message....:  Authority not restored for user QPGMR.
Cause......:  Private authority for DTAARA DTAARA1 in library QGPL was
 not restored. Either the object does not exist, is damaged, or was not
 available at the time authority was being granted.
Recovery...:
  Do one of the following:
 --If the system was dedicated while the RSTAUT command was running,
 display the description of the object (DSPOBJD command). If the object was
 damaged or not found, restore the user profile (RSTUSRPRF command), restore
 the object (RSTOBJ command), and restore the authorizations (RSTAUT command).
 If the object exists and is not damaged, report the problem (ANZPRB command).
```

Figure 10. Expanded text for message CPF3736

The expanded text for message CPF3845 appears as follows:

```
              Additional Message Information

Message ID......:  CPF3845        Severity......:  20
Message type....:  Diagnostic
Date sent ......:  04/24/97     Time sent......:  19:35:17

Message....:  Some authorities not restored for user profile QPGMR.
Cause......:  1434 authorities were restored and 1 authorities were not
 restored for user profile QPGMR at 04/24/97 19:21:36.
   The prestart job name used to restore private authorities for this user
 profile is *N. The prestart job name that contains messages about
 authorities not restored is *N.
  --If the job name above is *N then a prestart job was not used to restore
  authorities for this user profile.
  --If a job name other than *N is listed above, then a prestart job was
  used to restore private authorities for this user profile and messages can
  be found in the joblog for the job name listed. Use one of the following
  commands to display the joblog for the prestarted job:
```

Figure 11. Expanded text for message CPF3845

In the case of a system in a restricted state, all messages appear in the user's main job log. When the name of the prestarted job that is used in message CPF3845 is *N, then no prestarted job was used.

Figure 12 on page 222 and Figure 13 on page 222 show a sample job log and message information for a RSTAUT USRPRF(QPGMR QUSR) command run on a system in a nonrestricted state.

```
>RSTAUT USRPRF(QPGMR QUSR)
Start of prestart jobs in progress.
Some authorities not restored for user profile QPGMR.
Private authorities restored for user profile QUSR.
End of prestart jobs.
Not all user profiles had all authorities restored.
```

Figure 12. Sample job log for RSTAUT on a system in a nonrestricted state

The expanded text for message CPF3845 appears as follows:

```
              Additional Message Information

Message ID......:  CPF3845       Severity......:  20
Message type....:  Diagnostic
Date sent ......:  04/24/97     Time sent......:  19:35:17

Message....:  Some authorities not restored for user profile QPGMR.
Cause......:  1433 authorities were restored and 2 authorities were not
 restored for user profile QPGMR at 04/24/97 19:21:36.
   The prestart job name used to restore private authorities for this user
 profile is 010648/QUSER/QSRRATBL. The prestart job name that contains
 messages about authorities not restored is 010648/QUSER/QSRRATBL.
  --If the job name above is *N then a prestart job was not used to restore
  authorities for this user profile.
  --If a job name other than *N is listed above, then a prestart job was
  used to restore private authorities for this user profile and messages might
  be found in the joblog for the job name listed. Use one of the following
  commands to display the joblog for the prestarted job:
```

Figure 13. Expanded text for message CPF3845

In Figure 13 on page 222, the name of the prestarted job used is 010648/QUSER/QSRRATBL, and it appears in the CPF3845 message. The CPF3736 message for the data area DTAARA1 in library QGPL whose authority was not restored, does not appear in the user's main job log. Instead, all messages that are related to restoring of individual private authorities are in the job log for the prestarted job. To view these messages, you would run the command DSPJOB JOB(010648/QUSER/QSRRATBL), and then select option 4 to view the job log for the prestarted job. The expanded message text for CPF3736 appears in that job log.

You should pay special attention to any CPF3845 message that states that *N authorities were not restored. This might indicate a problem such as damaged objects or a function check. Any CPF3845 message with *N authorities that are not restored should be investigated further by checking the job log of the prestarted job that is named.

If all authorities in an authority reference table were restored successfully, then message CPC3706 is sent for the user profile instead of CPF3845. CPC3706 will also contain the name of the prestarted job that is used to restore authorities for the user profile. If all authorities restored from a prestarted job were restored successfully, then the job log for the prestarted job will contain only job start and end messages.

The order of CPC3706 and CPF3845 messages depends on whether you run the RSTAUT command on a system that is in a restricted or nonrestricted state. These messages are for user profiles with restored private authorities. The order of these messages is as follows:

Restricted state system
The order will generally be alphanumeric because only one authority table is being restored at a time, in alphanumeric order

Nonrestricted state system
The order will generally be that these messages will appear first for the user profiles with fewer private authorities, then later for the user profiles with many private authorities. This is because multiple authority reference tables are being restored at once and the smaller authority reference tables typically complete first.

Restoring authority on a system in a restricted state

The Restore Authority (RSTAUT) command on a system running in a restricted state restores authorities for each authority reference table, one table at a time. No prestarted jobs are used.

When processing is complete for an authority reference table, the table is deleted regardless of whether all private authorities were successfully restored or not.

Examples: How the system restores authority

When you run the Restore Authority (RSTAUT) command, the system grants all the private authorities that it finds in each authority reference table.

The user's private authorities after the command are both of the following items:

- The authorities from the temporary authority reference table.
- Any authorities that are granted to the user since the save operation.

How the system restores authority–Example 1: Assume that the authority to PRICES looks like this at the time of the save operation:

```
                     Display Object Authority

Object . . . . . :   PRICES          Owner  . . . . .
  Library  . . . . :     CONTRACTS     Primary group  .
Object type  . . . :   *FILE
Object secured by authorization list  . . . . . .

                       Object
User        Group      Authority
OWNCP                  *ALL
DPTSM                  *CHANGE
DPTMG                  *CHANGE
WILSONJ                *USE
*PUBLIC                *EXCLUDE
```

Note: Your display looks different when your user profile has a user option setting of *EXPERT.

After you save security information, you grant and revoke several authorities to the PRICES file. Just before the restore operation, the authority looks as follows:

```
                     Display Object Authority

Object . . . . . :   PRICES          Owner  . . . . .
  Library  . . . . :     CONTRACTS     Primary group  .
Object type  . . . :   *FILE
Object secured by authorization list  . . . . . .

                       Object
User        Group      Authority
OWNCP                  *ALL
DPTSM                  *USE
DPTMG                  *CHANGE
WILSONJ                *EXCLUDE
ANDERSP                *USE
*PUBLIC                *EXCLUDE
```

If authority is restored for all users, the authority to the PRICES file looks as follows:

```
                     Display Object Authority

Object . . . . . :   PRICES          Owner  . . . . .
  Library  . . . . :     CONTRACTS     Primary group  .
Object type  . . . :   *FILE

Object secured by authorization list  . . . . . .

                        Object
User        Group       Authority
OWNCP                   *ALL
DPTSM                   *CHANGE
DPTMG                   *CHANGE
WILSONJ                 *USE
ANDERSP                 *USE
*PUBLIC                 *EXCLUDE
```

Authorities for DPTSM and WILSONJ are restored to the values they have on the save media. The authority for ANDERSP remains, even though it did not exist on the save media.

How the system restores authority–Example 2: Assume that the authority for the PRICES file looks like this just before the restore operation:

```
                     Display Object Authority

Object . . . . . :   PRICES          Owner  . . . . .
  Library  . . . . :     CONTRACTS     Primary group  .
Object type  . . . :   *FILE

Object secured by authorization list  . . . . . .

                        Object
User        Group       Authority
OWNCP                   *ALL
DPTMG                   *CHANGE
WILSONJ                 *CHANGE
*PUBLIC                 *USE
```

If authority is restored for all users, the authority to the PRICES file looks as follows:

```
                     Display Object Authority

Object . . . . . :   PRICES          Owner  . . . . .
  Library  . . . . :     CONTRACTS     Primary group  .
Object type  . . . :   *FILE

Object secured by authorization list  . . . . . .

                        Object
User        Group       Authority
OWNCP                   *ALL
DPTSM                   *CHANGE
DPTMG                   *CHANGE
WILSONJ                 *CHANGE
*PUBLIC                 *USE
```

Notice that WILSONJ still has *CHANGE authority. The authority from the save media (*USE) is granted to WILSONJ, but the authority WILSONJ already has is not revoked. *USE authority is added to *CHANGE authority, so WILSONJ has *CHANGE authority.

Notice also that *PUBLIC authority is not affected by this process. Public authority is stored with the object and is handled when the object is restored. If public authority on the system is different from public authority on the save media, the public authority on the system is used.

Authority is restored to the object with the same name in the same library. In some cases, this might result in restoring authority to a different object.

If you restore authorities for an independent auxiliary storage pool (ASP), you can use the SAVASPDEV to limit the authorities that you restore. For example, you can limit the authorities to a specific independent

ASP, or an ASP group. Authority is restored to the object with the same name in the same library and the same independent ASP, unless you specify another value for the RSTASPDEV parameter. A user will have separate authority reference tables for each independent ASP to which he is authorized.

Assume that you delete program PGMA in library CUSTLIB. You create a new program with the same name but different function. If you restore authority, users who were authorized to the original PGMA are now authorized to the new PGMA.

Related concepts

How the system restores programs

Restoring programs to your system represents a security exposure. A restored program might be altered to do functions that you do not intend, or the program might adopt the authority of a powerful user profile.

Restoring authorities for independent auxiliary storage pool data

You can use the methods shown in these figures to restore authority for an independent auxiliary storage pool (ASP).

Note: These scenarios are examples of the most common restore options.

To restore authorities for data that was saved from an independent ASP and that will be restored to an independent ASP with an identical name, use this command:

```
RSTAUT
```

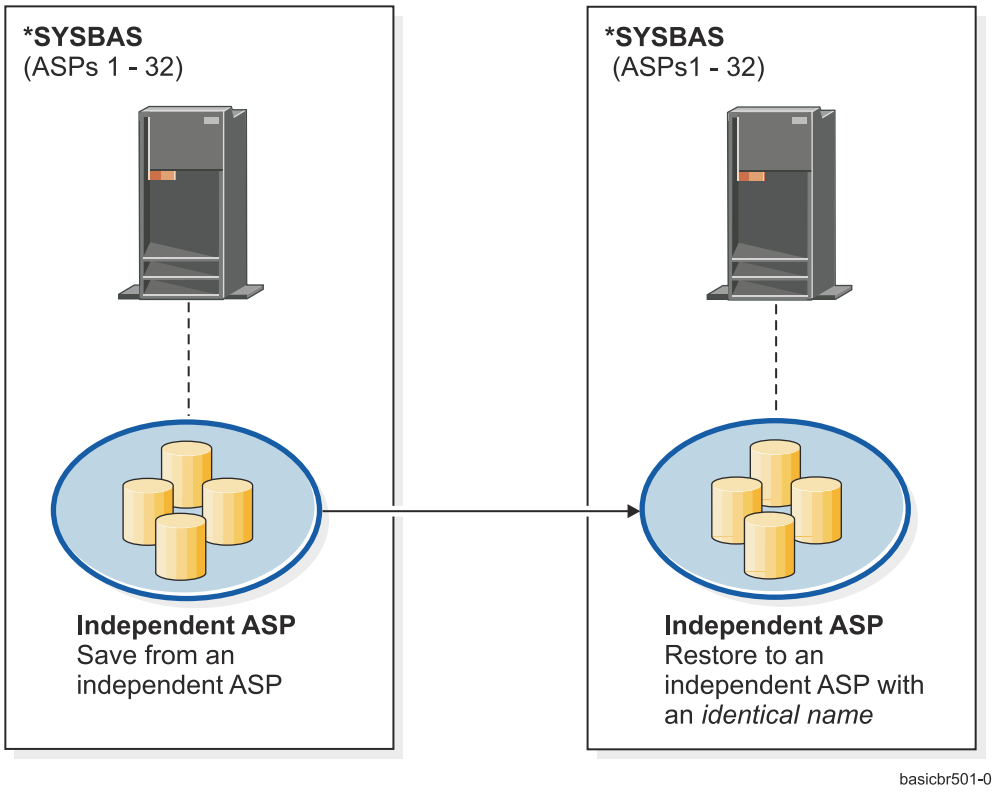

Figure 14. Save data from an independent ASP and restore it to an independent ASP with an identical name

To restore authorities for data that was saved from *SYSBAS and that will be restored to an independent ASP, use this command:

```
RSTAUT SAVASPDEV(*SYSBAS) RSTASPDEV(IASPABC)
```

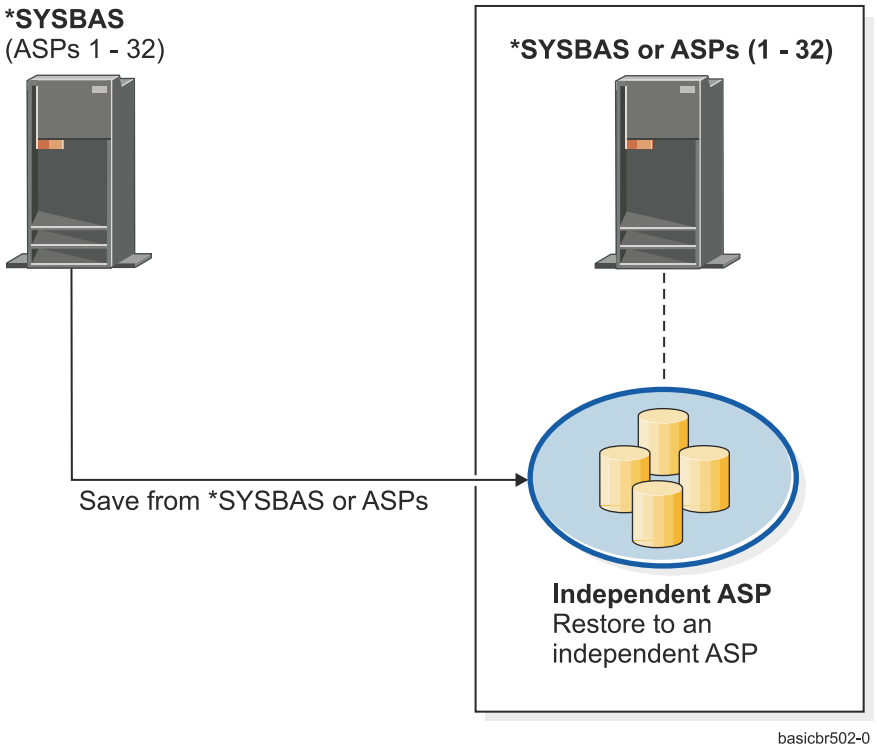

*Figure 15. Save data from *SYSBAS and restore it to an independent ASP*

To restore authorities for data that was saved from an independent ASP and that will be restored to an independent ASP with a different name, use this command:

```
RSTAUT SAVASPDEV(IASPABC) RSTASPDEV(IASPXYZ)
```

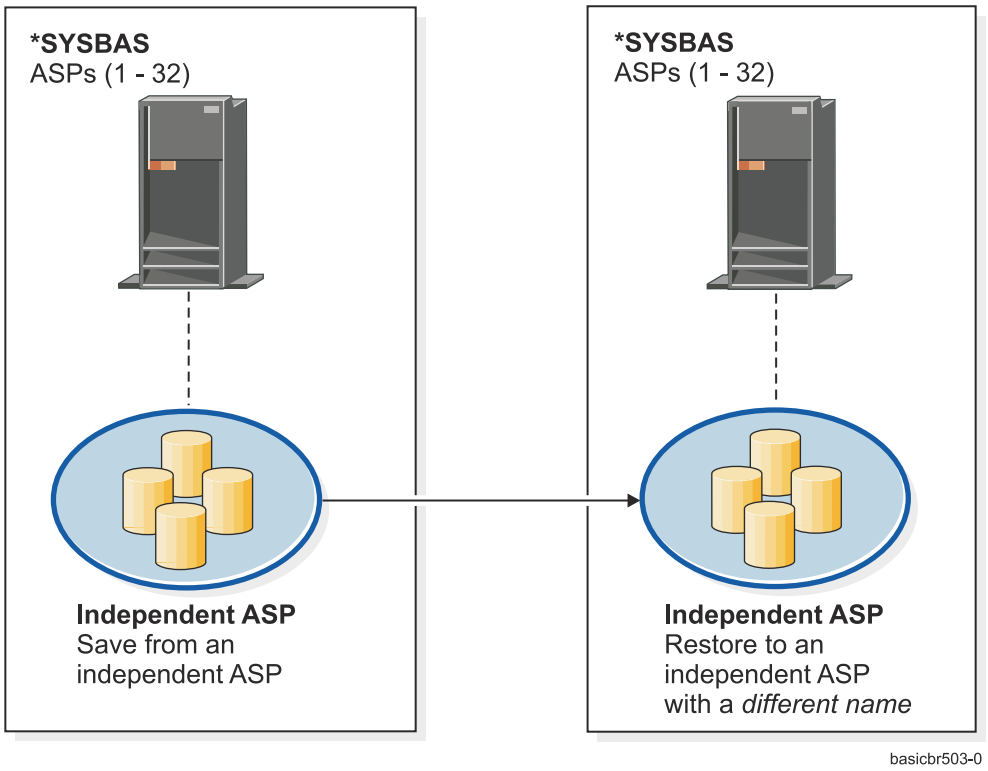

Figure 16. Save data from an independent ASP and restore it to an independent ASP with a different name

To restore to *SYSBAS the private authorities for objects that were previously restored, use this command:

```
RSTUSRPRF DEV('device-path-name') USRPRF(*ALL)
RSTAUT SAVASPDEV(iasp-name) RSTASPDEV(*SYSBAS)
```

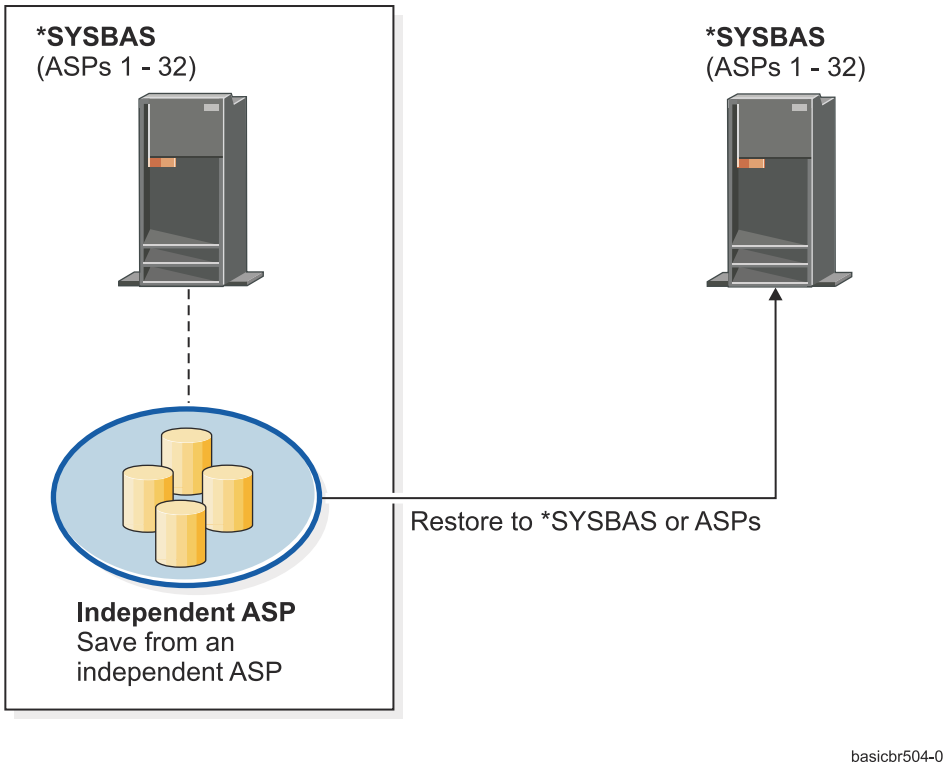

*Figure 17. Save data from an independent ASP and restore it to *SYSBAS*

Related tasks

Transferring a library to a different auxiliary storage pool
Use this procedure to move a library to a different auxiliary storage pool (ASP).

Restoring configuration objects

Before you restore a configuration object, you must vary it off.

You can restore the following items:

- All configuration objects
- A group of configuration objects by generic name
- Only specific types of configuration objects, such as line descriptions or connection lists.
- System resource management information

If you run the restore configuration (RSTCFG) command against a printer device description and the output queue associated with that device description contains no spooled files, the system will re-create the output queue. Any changes made to the output queue before running the RSTCFG command will be lost.

Table 52. How configuration objects are restored

Possible method	Restricted state?
RSTCFG command [1]	No
Restore menu option 7	No

Table 52. How configuration objects are restored (continued)

Possible method	Restricted state?
Restore menu option 21	Yes
Restore menu option 22	Yes
Restore menu option 23	Yes

1

You must have *ALLOBJ special authority to specify ALWOBJDIF(*ALL)

To restore all configuration objects, follow these steps:

1. Create the independent auxiliary storage pool (ASP), if necessary. If the independent ASP is unavailable, you cannot restore configuration objects.
2. Find the most recent media volume that has your configuration. It might be a SAVSYS volume or a SAVCFG volume. The name of the file on the volume is QFILEIOC.
3. If you are using a SAVSYS media volume, type the following command:

```
RSTCFG OBJ(*ALL) DEV(media-device-name)
       OBJTYPE(*ALL)
       ENDOPT(*LEAVE)
```

If you are using a SAVCFG media volume, type the following command:

```
RSTCFG OBJ(*ALL) DEV(media-device-name)
       OBJTYPE(*ALL)
       ENDOPT(*UNLOAD)
```

Restoring to a different system

You must specify ALWOBJDIF(*ALL) when you restore the configuration to a different system. (An option is available on the restore menu that indicates that you are restoring to a different system or different logical partitions. If you selected this option, the system automatically specifies ALWOBJDIF(*ALL) for you.)

The restoring of configuration objects to a different system whose configuration objects exist overlays the existing configuration. In some cases, the configuration description might not match the hardware on the system.

Do not restore system resource management objects to another system. This might cause problems that can be fixed only by a service representative. When you use the RSTCFG command to another system, specify SRM(*NONE).

Correcting problems with the system resource management information

The system resource management (SRM) information provides a link between the hardware on your system and the software descriptions of that hardware (the configuration).

When you restore your configuration to a different system, you should not restore the SRM information because it will not match the hardware on the target system. Sometimes during a system upgrade, you are instructed to restore the SRM information to your system even though some of your hardware has changed.

If you have restored the SRM information and the hardware configuration does not match, use the following procedure to correct the SRM information:

1. Type STRSST and press the Enter key to access system service tools (SST).
2. Select Option 1 (Start a service tool) from the System Service Tools menu and press Enter.
3. Select Option 7 (Hardware service manager) from the Start a Service Tool menu and press Enter.

4. Select Option 2 (Logical hardware resources) from the Hardware Service Manager menu and press Enter.
5. Select Option 1 (System bus resources) from the Logical Hardware Resources menu and press Enter.
6. Select F10 (Non-reporting resources) to display any nonreporting resources.

   Any hardware resources that did not report during the last initial program load (IPL) or that were created during the last Restore Configuration (RSTCFG) command will be displayed.
7. Type a 4 (Remove) in the Option column to delete any entry you are certain is not valid for this system's configuration.

Recovering devices that will not vary on

If you have a problem with your devices, such as not being able to vary on a device, it might be because the system resource management (SRM) database that was restored does not match the device descriptions on the system.

To correct the problem for a tape unit or a tape controller, perform these tasks:

1. Type `WRKHDWRSC TYPE(*STG)`. You are shown the Work with Storage Resources display.
2. Type a 9 (Work with resource) in the Opt column next to the resource name that would not vary on. The Work with Storage Controller Resources display is shown.
3. Write down the valid resource name for the device type and model that you tried to vary on.
4. Press F12 (Cancel) until you return to a display with a command line. If a problem exists with the tape unit, skip to step 7.
5. Type `WRKCTLD CTLD(`*`controller-name`*`)`. You are shown the Work with Controller Descriptions display.
6. Type 2 (Change) in the Opt column next to the controller that would not vary on and press the Enter key. The Change Controller Description display is shown. Skip to step 9.
7. Type `WRKDEVD DEVD(`*`media-device-name`*`)`. The Work with Device Descriptions display is shown.
8. Type 2 (Change) in the Opt column next to the device description that you want to change and press the Enter key. The Change Device Description display is shown.
9. Change the name in the *Resource name* prompt to the correct name for the resource and press the Enter key. You return to the Work with Device Descriptions display or the Work with Controller Descriptions display.
10. Type 8 (Work with status) in the Opt column next to the device or controller that you changed and press the Enter key. The Work with Configuration Status display is shown.
11. Type 1 (Vary on) in the Opt column next to the device description name or the controller description name. Press the Enter key to vary it on.

Local workstation controller: To correct the problem for a workstation, follow these steps:

1. Type the following command and press the Enter key to display the Work with Local Workstation Resources display.

   ```
   WRKHDWRSC TYPE(*LWS)
   ```
2. Find the correct controller description for the device that would not vary on.
3. Type a 5 (Work with configuration description) in the Opt column next to the controller description name and press the Enter key.

   The Work with Configuration Description display is shown.
4. Type a 5 (Display) in the Opt column to display the valid resource name for the workstation controller.
5. Press F12 (Cancel) until you return to a display with a command line.
6. Type the following command and press the Enter key to display the device description for the device that would not vary on.

```
WRKCTLD CTLD(controller-name)
```

The Work with Controller Descriptions display is shown.

7. Type a 2 (Change) in the Opt column next to the controller description that you want to change and press the Enter key.

   The Change Controller Description display is shown.

8. Change the name in the *Resource name* prompt to the correct name for the resource and press the Enter key.

   You will return to the Work with Controller Descriptions display.

9. Type an 8 (Work with status) in the Opt column next to the controller description that you changed and press the Enter key.

   The Work with Configuration Status display is shown.

10. Type a 1 (Vary on) in the Opt column next to the controller description name and press the Enter key to vary on the device.

    Note: It is possible that another device description is varied on for this resource. Vary off the device first and then vary on the changed device description. This situation can happen to the console device.

Recovering when you change the console type

When you restore your information to a different system or an upgraded system, you might have a different console type on the target system.

After you have restored user information, you need to create a new controller and device description. Follow these steps:

1. Type `WRKHDWRSC *LWS` and press the Enter key.

   You are shown the Work with Local Work Station Resources display.

2. Type a 5 (Work with controller descriptions) in the Opt column next to the name of the first workstation controller and press the Enter key.

   You are shown the Work with Controller Descriptions display.

   Note: The first workstation controller might not be CTL01.

3. Type a 1 in the Opt column and press the Enter key.

   You are shown the Create Controller Description display.

4. For the *New controller description* prompt, type the name that you want for the console. Press the Enter key.

   Note: If you want to use the name that you had on your old system, you must first delete the device configuration name and then re-create it.

5. Use the Create Device Description (Display) (CRTDEVDSP) command to create a device description for the console.

Recovering the System/36 environment configuration

If you are experiencing a problem with the System/36 environment after restoring the system, it might be caused by the locking rules used during the installation process. The QS36ENV configuration object in library #LIBRARY might have been locked by the System/36 environment.

This object contains the System/36 environment names for the workstation, printer, tape and diskette units on the system, and default System/36 environment values used for all users. This object might have been modified by the Change S/36 Environment Configuration (CHGS36) command to customize the System/36 environment.

When the first subsystem is started on the system after the installation process is complete, a new #LIBRARY and a new QS36ENV object in #LIBRARY are created with the system defaults. In addition to

creating the new objects, each subsystem holds a lock on the QS36ENV configuration object to ensure that it is not deleted. This lock will not allow the saved QS36ENV configuration object to be restored.

If the QS36ENV configuration object did not restore, start with step 1. If the configuration object did restore but you are experiencing problems with the System/36 environment configuration, go to step 5.

1. Rename the newly created #LIBRARY to something else (for example, #LIBNEW).

   The locks held on the QS36ENV object remain with the renamed library. This allows the saved System/36 environment configuration object to be restored.
2. Restore the saved copy of library #LIBRARY: `RSTLIB SAVLIB(#LIBRARY)`
3. Perform an initial program load (IPL) of the system.

   The QS36ENV object in the restored copy of #LIBRARY is the System/36 environment configuration again.
4. Delete the earlier renamed version of #LIBRARY (for example, #LIBNEW).
5. Use the Change S/36 Environment Configuration (CHGS36) command to refresh the configuration object.
   a) Select each of the device types that you want to change.
      - Workstation devices
      - Printer devices
      - Tape devices
      - Diskette devices

   b) For each device type that you want to change, perform these tasks:
      i) Press the F5 key to ensure the configuration object matches the device descriptions on the system.
      ii) If any System/36 names are not specified, do one of the following:
         - Press the F10 key to use the defaults for the System/36 names for those devices.
         - Update the System/36 names manually.

   c) Save the changes to the configuration object.

   See the topic on configuring the System/36 environment in the *Concepts and Programmer's Guide for the System/36 Environment* for more information about configuring the System/36 environment.

Restoring logical partitions

Use this information to help recover logical partitions.

If your logical partition is managed by Integrated Virtualization Manager, see PowerVM® Editions Operations Guide for information about how to recover partition configuration data.

For more information about logical partitions when you are using a Hardware Management Console (HMC) to manage the logical partitions, refer to the following information in the IBM Power Systems Hardware documentation:

- Logical Partitioning Guide

Related information

IBM i client logical partition considerations

Restoring libraries

Restoring a single library or a group of libraries is a common way to recover user information.

Use the Restore Library (RSTLIB) command to restore a single saved library or a group of libraries. The RSTLIB command restores the entire library, including the library description and the objects in the

library. This command also restores status information for programming temporary fixes (PTFs) that were in the library at the time the library was saved.

When you use the RSTLIB command, you can use the OPTION parameter to specify which objects in a library are restored:

Possible values for the OPTION parameter of the RSTLIB command:

*ALL	Old objects are replaced and new objects are added to a library. *ALL is the default value.
*OLD	Only old objects that already exist on the system are replaced in a library.
*NEW	Only objects not found on the system are added to a library. The old objects are not replaced.
*FREE	Only those objects that have their storage freed on the system are restored.

Important: If you are restoring to a different system or logical partition, specify ALWOBJDIF(*COMPATIBLE) when you are restoring libraries.

Restoring a library from a previous release

When you are restoring a library that was saved on a system at an earlier release, you can use the Force object conversion (FRCOBJCVN) parameter to specify whether programs are translated when they are restored.

This can significantly impact the time it takes to restore the library.

Related concepts

Restoring programs to a different release
This topic provides information about how to restore programs to a different release.

Sequence for restoring related objects
Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

How the system restores programs
Restoring programs to your system represents a security exposure. A restored program might be altered to do functions that you do not intend, or the program might adopt the authority of a powerful user profile.

Restoring multiple libraries

You can use the Restore Library (RSTLIB) command to restore multiple libraries.

Use the RSTLIB command to restore libraries in these groups:

***NONSYS**
All libraries that were saved with the SAVLIB LIB(*NONSYS) command, including the IBM-supplied libraries QSYS2, QGPL, QUSRSYS, and licensed program libraries.

***ALLUSR**
All user libraries that were saved with SAVLIB LIB(*ALLUSR) or SAVLIB LIB(*NONSYS).

***IBM**
All IBM-supplied libraries that were saved with SAVLIB LIB(*IBM) or SAVLIB(*NONSYS). Only IBM-supplied libraries that contain IBM objects are restored.

***ANY**
The first version of all saved libraries that are found on the tape are restored, beginning with the sequence number specified for the SEQNBR parameter. When you restore from optical media, all

saved libraries are restored that are found in the directory specified for the Optical file (OPTFILE) parameter.

***SELECT**
The libraries that are specified for the Select (SELECT) parameter are restored.

User-specified (up to 300 repetitions)
generic-name – You can specify the generic name of the library to be restored. A generic name is specified as a character string that contains one or more characters followed by an asterisk (*). If a generic name is specified, then all objects that have names with the same prefix as the generic object name are selected.

name – You can specify the names of the libraries to be restored. The names of the libraries being restored must be the same as the names that were used when the libraries were saved.

Figure 2 on page 20 shows which libraries are saved and restored in these groups. If you are restoring any of the following libraries, QSYS2, QGPL, QUSRSYS, or QSYS2*nnnnn*, where *nnnnn* specifies an independent auxiliary storage pool (ASP) number, you must restore them before restoring any other user libraries. If you use the special values (*ALLUSR or *NONSYS), the system restores these libraries in the correct sequence.

When you restore a group of libraries, you can omit up to 300 libraries by using the OMITLIB parameter. You can specify to omit specific libraries or you can specify to omit generic names for the libraries to be omitted. The libraries you omit are not restored from the save media. You can also use the OMITOBJ parameter to omit up to 300 specific object names or generic object names. You can also use the SELECT parameter to select libraries and objects to include or omit.

When you use a media definition to restore libraries that were saved in parallel with one of the following groups specified, *ALLUSR, *IBM, *NONSYS, or a generic value such as X*, you might have to perform some involved recovery operations. You must first load each drive with the volume that contains the QFILE so that the system can verify that each library resulted from the same save operation. You will then have to answer an inquiry message for each drive to position it to the correct volume, if you are starting the recovery on a library other than *FIRST.

Note: If you have related objects, such as physical and logical files or journals and journaled objects, in different libraries, you must ensure that you restore them in the correct sequence.

Related concepts
Deferring the restore of dependent objects
You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Considerations and restrictions of the Restore Library command

These restrictions and considerations apply to the Restore Library (RSTLIB) command.

- You cannot restore a QDOC*nnnn* (Document) library by using the RSTLIB command. Use the Restore Document Library Object (RSTDLO) command to restore documents.
- You cannot restore the QSYS (System) library by using the RSTLIB command. Use the procedures for restoring the operating system in Chapter 5, “Restoring the operating system,” on page 141 to restore QSYS.
- A RSTLIB command can run concurrently with a Restore Object (RSTOBJ) or Save Object (SAVOBJ) command that uses the same library.
- You cannot run multiple concurrent RSTLIB commands that use the same library.
- A RSTLIB and SAVLIB command cannot run concurrently using the same library name, even if the libraries are in different independent ASPs.

Restoring all libraries from a single save operation

Use these steps to restore all libraries that were saved with a single command or menu option.

1. Sign on with a user profile that has *SAVSYS special authority.

Using *SAVSYS special authority ensures that you will not have authority problems during the restore procedure and improves restore performance.

2. Ensure that the system is in a restricted state.
3. Find your most recent save media.
4. Use “Task 4: Restoring libraries to a basic auxiliary storage pool” on page 188.
   Type your choice and press F4 (Prompt).

Table 53. Methods for restoring all libraries–single save operation

How your libraries were saved	Type this to restore them
Save Menu option 21	`RSTLIB SAVLIB(*NONSYS)`
SAVLIB LIB(*NONSYS)	`RSTLIB SAVLIB(*NONSYS)`

5. Type your choices for other parameters, such as device and whether to rewind the tape in a tape device. Press the Enter key.
6. If you receive messages to load a media volume, load the correct media volume and respond to the messages.
7. When the restore operation completes, check your job log to see which libraries were restored and whether any objects were not restored.

Related tasks

Putting your system in a restricted state
Many recovery procedures require that your system have no other activity on it. When no subsystems except the controlling subsystem are active on your system, it is in a *restricted state*.

Restoring all libraries from multiple save operations

Use these steps to restore all libraries if they were saved with multiple menu options or commands.

Adapt the examples to fit your own save procedures and recovery situation. Before restoring multiple libraries, make sure that you read about “Sequence for restoring related objects” on page 25.

1. Sign on with a user profile that has *SAVSYS special authority.
2. Ensure that the system is in a restricted state.
3. Find your most recent save media.
4. Use Table 54 on page 234 repeat this step and step 5 for each command.
   Type your choice and press F4 (prompt).

Table 54. Methods for restoring all libraries–multiple save operations

How your libraries were saved	Type this to restore them
Save Menu options 22 and 23	`RSTLIB SAVLIB(*IBM)` `RSTLIB SAVLIB(*ALLUSR)`
Save Menu options 21 and 23	`RSTLIB SAVLIB(*IBM)` `RSTLIB SAVLIB(*ALLUSR)`

SAVLIB *NONSYS followed by
`SAVLIB LIB(LIBA LIBB LIBC)`
	`RSTLIB SAVLIB(*NONSYS)`
	`OMITLIB(LIBA LIBB LIBC)`
	`RSTLIB LIB(LIBA)`
	`RSTLIB LIB(LIBB)`
	`RSTLIB LIB(LIBC)`

5. Type your choices for other parameters, such as device and whether you want to rewind the tape in a tape device or not. Press the Enter key.
6. If you receive messages to load a media volume, load the correct media volume and respond to the messages.

7. When the restore operation completes, check your job log to see which libraries were restored and whether any objects were not restored.

Restoring objects

You can use the Restore Object (RSTOBJ) command to restore one or more objects to a library.

You can use the RSTOBJ command to restore objects that were saved by the Save Library (SAVLIB), Save Object (SAVOBJ), or Save Changed Objects (SAVCHGOBJ) commands. You must restore objects to an existing library. When you restore all objects from data that was saved by the SAVLIB command, the library description is not restored.

There are several parameters on the RSTOBJ command that allow you to specify which objects to restore. Parameters that allow you to select objects from the save media include Objects (OBJ), Saved Library (SAVLIB), Object types (OBJTYPE), Libraries to omit (OMITLIB), Objects to omit (OMITOBJ), Select (SELECT), and File member (FILEMBR). The Restore to library (RSTLIB), Restore to ASP device (RSTASPDEV), and Restore to ASP number (RSTASP) parameters allow you to specify where to restore the objects. The Option (OPTION), Data base member option (MBROPT), and Allow object differences (ALWOBJDIF) parameters allow you to specify whether to restore objects that already exist in the destination library.

You can run multiple concurrent RSTOBJ commands for one or more libraries, if the saved objects are available on different save media.

Attention: Do not use RSTOBJ to restore licensed programs to library QSYS. Unpredictable results can occur.

Example: This example restores the most recently saved version of data area CLUSTERDTA in library LIBX, using tape device TAP01.

```
RSTOBJ  OBJ(CLUSTERDTA)  SAVLIB(LIBX)  DEV(TAP01)
        OBJTYPE(*DTAARA) VOL(*SAVVOL)
```

When you restore individual objects from a tape file that contains many saved objects, you can improve the performance of the restore operation if you specify the following information:

1. Tape file sequence number.

   The tape file sequence number is the quickest way to locate the tape file that contains the wanted objects.

2. Specific object name and type.

   The specific object name and specific object type, when specified, causes the restore operation to stop searching when the wanted object is found. If *ALL, a generic name or more than one object type is specified, the restore operation searches through the entire tape file, trying to find additional matching objects.

3. Starting position.

   The starting position in the tape file for each object is provided by save operations when output to an outfile is requested.

Example: This example restores the object Z0002OFR in library TESTLIB, using device TAP02 and specifying the saved object located at tape file sequence number 3199 and at the starting position '00000000000027FC0000000000004964'.

```
RSTOBJ  OBJ(Z0002OFR)  SAVLIB(TESTLIB)  DEV(TAP02)
        SEQNBR(3199) POSITION(00000000000027FC0000000000004964)
```

Restoring journaled objects

This section describes options for restoring journaled objects (for example, libraries, data areas, data queues, database files, and integrated file system objects).

You can journal the following types of objects:

- Access paths
- Data areas
- Database files
- Data queues
- Libraries
- Materialized query tables
- Stream files
- Directories
- Symbolic links

If the journal exists on the system before the journaled objects are restored, all objects that were saved while being journaled are journaled again if one of the following situations is true:

- The objects are not on the system at restore time.
- The objects are on the system and journaling was not ended for the objects.

If the journal is not on the system when a journaled object is restored, the restore operation for the object sends a warning message and journaling is not resumed. This warning message causes an escape message to be sent at the end of the restore operation. If you specify a value for the Defer ID (DFRID) parameter on the restore command, you can use the Restore Deferred Objects (RSTDFROBJ) command when the journal is available to start journaling for the objects.

If you do not want to start journaling for the restored objects, specify *NO for the Start journaling (STRJRN) parameter on the restore command.

Journaled libraries

You can journal libraries like any other object. Journaling can start automatically for objects such as database files, data areas, and data queues that are created, moved into, or restored into a journaled library. The library's inheritance rules determine which objects to start journaling automatically and with what journaling attributes. The inheritance rules are saved and restored with the library. When you restore a library that does not already exist on the system, you also restore the inheritance rules for that library.

You can replay changes to libraries by using the Apply Journaled Changes (APYJRNCHG) command. To start journaling for a library, use the Start Journal Library (STRJRNLIB) command.

When you restore an object in a library that was being journaled at the time of the save operation, an entry is written to the journal used during save time to indicate that the object was restored. When you restore a new library, journaling information also is restored. After all objects in the library are restored, journaling is started for the library.

There are a few considerations to keep in mind when saving and restoring journaled libraries:

- You can save and restore journaling information for libraries, as well as other object types.
  - If the library is restored over an existing library, the journaling state and attributes of the library do not change.
  - If the library is restored to a different name and the library was journaled when it was saved, the system starts journaling the library to the same journal when it is restored.
  - If the library was not journaled when it was saved, any past journaling information is also restored.
  - Libraries are restored first, but do not start journaling until after all objects have been restored into the library.

- You can save and restore inheritance rules for libraries.
- If you save a library to IBM i V5R4, or earlier, the library's journaling information is not saved.

Related concepts

Sequence for restoring related objects
Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Verifying whether objects are restored successfully
You can use the job log or an output file to determine which objects are restored successfully.

Related tasks

Applying journaled changes
With this synchronization method, you use system journaling capabilities to synchronize the information in journaled objects on your test and existing systems.

Related information

Automatically starting journaling
Journaling libraries
Saving journaled objects and libraries
Start Journal Library (STRJRNLIB) command

What happens when you restore journaled objects to a different library or directory

The system assigns a unique internal journal identifier (JID) to every object that is journaled.

If you restore a journaled object to a library or directory other than the original library or directory, and the object still exists on the system and continues to be journaled to the same journal, the JID of the restored object is changed. Message CPF70CB — "Journal identifier &2 is currently in use" is sent to the job log to confirm the changed JID of the restored object.

All the journal entries associated with the media copy of the object have the original JID. You cannot apply these journal entries to the object that was restored to a different library or directory because it has a different JID. For this reason, you should avoid restoring a journaled object to a different library or directory.

For example, in Figure 18 on page 237, the original object FILEA in LIBX library has an internal journal identifier of Z that is recorded with every journal entry associated with FILEA in LIBX. When FILEA is restored from the media to LIBC library, it is assigned the journal identifier of Y because FILEA still exists in LIBX and continues to be journaled.

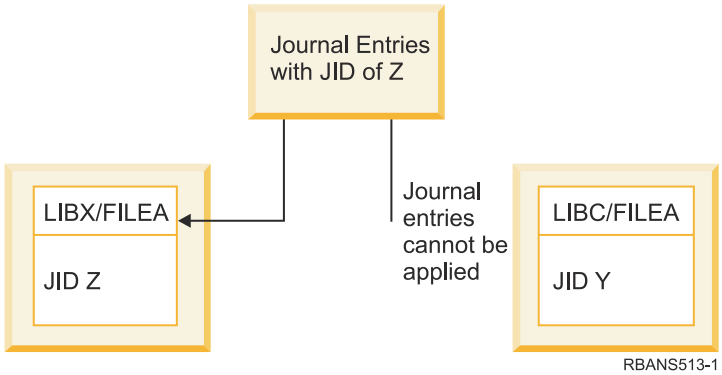

Figure 18. Example: Restoring a journaled object to a different library

Any journal operation that references an object by name and involves using journal entries requires that the journal identifier of the object and the journal identifier recorded in the journal entries be the same. Because FILEA in LIBC has journal identifier Y, journal entries with journal identifier Z are not associated with the restored FILEA in LIBC. As a result, journal changes that are recorded for FILEA in LIBX cannot be applied to FILEA in LIBC. For the same reason, if you are referencing FILEA in LIBC on the Display Journal (DSPJRN), Receive Journal Entry (RCVJRNE), or Retrieve Journal Entry (RTVJRNE) commands, or on the Retrieve Journal Entries (QjoRetrieveJournalEntries) API, the entries for FILEA in LIBX are not returned.

To display or retrieve the journal entries of the original object, follow these steps:

1. Save and then delete the existing object on the system.
2. Restore the original object to the system.
3. Display or retrieve the journal entries.
4. Delete the original object.
5. Restore the existing object back to the system.

Restoring user-defined file systems

Use this information to restore an unmounted user-defined file systems (UDFS), an individual object from an unmounted UDFS, and a mounted UDFS.

Restoring an unmounted user-defined file system

Use this information to restore an unmounted user-defined file system (UDFS).

To restore an unmounted UDFS, specify the following command:

```
RST OBJ(('/dev/QASP02/udfs_name.udfs'))
```

If the UDFS does not exist on the system, the system creates the block special file (*BLKSF) object that represents the UDFS. If the UDFS does exist, objects from the save media overlay objects on the system.

If you perform a disaster recovery, you must create the auxiliary storage pools (ASPs) that contain the UDFS before you attempt the restore operation. If you do not create the ASPs, the system does not restore the UDFSs.

You can optionally omit certain objects with either the OBJ or PATTERN parameter to reduce the recovery window when restoring the unmounted UDFS. For example:

```
RST OBJ(('/DEV/QASPxx/udfs_name.udfs') ('*.TEMP' *OMIT))
```

Restrictions while you restore an unmounted user-defined file systems

Here are the restrictions while you restore an unmounted user-defined file systems (UDFS).

- You cannot restore individual objects to unmounted UDFS.
- You cannot view or work with objects in an unmounted UDFS. Therefore, you cannot determine the amount of storage or time required by the restore operation once you unmount the UDFS.

Restoring an individual object from an unmounted user-defined file systems

You can restore individual objects from a save media volume that contains unmounted user-defined file systems (UDFS).

To perform the task, give a new name to the object that you restore. The parent directory of the new name must exist in an accessible file system.

For example, use the following save command to save the unmounted UDFS /dev/QASP01/*udfs_name.udfs* that contains object payroll:

```
SAV OBJ('/dev/QASP01/udfs_name.udfs')
```

To restore the object payroll from the unmounted UDFS to an existing directory /home/JON, use the following command:

```
RST OBJ(('/DEV/QASP01/udfs_name.udfs/payroll' +
         *INCLUDE +
         '/home/JON/payroll'))
```

Restoring a mounted user-defined file system

You can restore either just the objects within the mounted user-defined file system (UDFS), or restore both the UDFS information (/dev/*asp*/*udfs_name*.udfs) and the objects within the UDFS.

A UDFS, in the /dev/*asp* directory, is mounted over another directory. When you save the directory that the UDFS is mounted over, both the UDFS information and objects within the UDFS are saved.

In this example, the UDFS '/dev/qasp01/a.udfs' was mounted over directory '/appl/dir1' when directory '/appl/dir1' was saved. If you only want to restore the objects within the mounted UDFS, specify the following command:

```
RST OBJ(('/appl/dir1'))
```

The system restores the objects into the file system of the parent directory to which the objects are restored. The system does not restore UDFS 'dev/qasp01/a.udfs'.

When you recover from a disaster and if you saved your UDFS as mounted, use the following command to re-create the UDFS information during the restore operation:

```
RST  OBJ(('/directory_mounted_over'))
     RBDMFS(*UDFS)
```

Use the RBDMFS(*UDFS) parameter on the RST command to rebuild the mounted UDFS during the restore operation. However, if you omit the RBDMFS parameter or specify RBDMFS(*NONE), only the objects contained in the directory are restored, and none of the UDFS information is restored.

Restore actions for mounted user-defined file systems

This table describes what actions occur during the restore operation when RBDMFS(*UDFS) is specified on the RST command. The user-defined file systems (UDFS) might not be mounted, or might be mounted over the same directory or a different directory.

	UDFS that was mounted over the directory when saved			
Directory being restored over	**UDFS does not exist**	**UDFS not mounted**	**UDFS mounted over same directory**	**UDFS mounted over different directory**
Directory being restored over does not exist	• Create directory • Create UDFS • Mount UDFS • Restore tree	• Create directory • Mount UDFS • Restore tree	N/A	The restore operation fails with message CPD37B8.

Directory being restored over	UDFS that was mounted over the directory when saved			
	UDFS does not exist	**UDFS not mounted**	**UDFS mounted over same directory**	**UDFS mounted over different directory**
No UDFS is mounted over the directory	• Create UDFS • Mount UDFS • Restore tree	• Mount UDFS • Restore tree	N/A	The restore operation fails with message CPD37B8.
Same UDFS is mounted over the directory	N/A	N/A	Restore tree[1]	N/A
Different UDFS is mounted over the directory	• Create UDFS • Mount UDFS • Restore tree	• Mount UDFS • Restore tree	N/A[1]	The restore operation fails with message CPD37B8.
[1] Suppose the target UDFS is mounted over the target directory and one or more UDFSs are mounted over it. This situation is treated as if the target UDFS is mounted over a different directory. The restore of that mount point will fail with a message that indicates that the UDFS is in use.				

Restoring database files

You can restore one or more database files or one or more members of database files by using the Restore Object (RSTOBJ) command.

Figure 19 on page 240 shows, conceptually, how a database file with two members looks to the system. It has multiple parts:

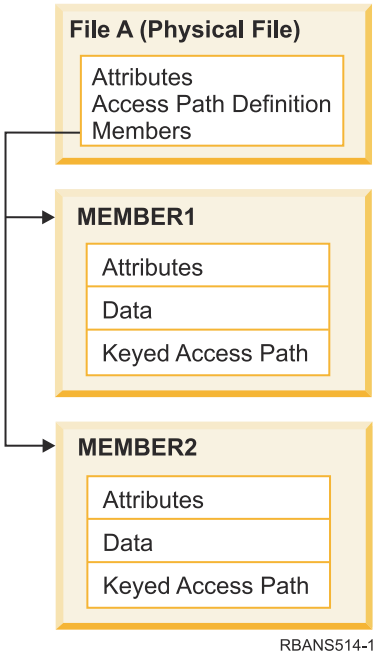

Figure 19. Example of a database file with two members

If FILEA exists on the system and you restore it, the system restores the data and access paths for FILEA's two members. The attributes for the file and its members are not changed on the system.

If you want to restore the file attributes as they existed at the time of the save operation, delete the file, then restore it. If you want to restore the member attributes, remove the member (RMVM) and then restore it, specifying MBROPT(*NEW).

When you restore a database file, the system uses information that is stored with the file and the parameters you specify to make decisions. The topics that follow describe special considerations when restoring database files and members.

Unique file identification: You can restore a file only to itself. A saved version and a copied version of the same file are not the same and cannot be used interchangeably in a restore operation. Figure 20 on page 241 illustrates this procedure:

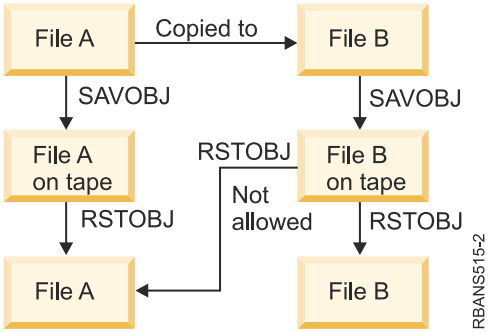

Figure 20. Restoring a copy of a file

File locking during restore operation: When you restore a file, no member of the file can be used during the restore operation, even through logical files. The file is exclusively locked during the restore operation.

Comparing file attributes during a restore operation

When you restore a database file or member that exists on the system, the system expects the creation dates for the system copy and the media copy to be the same. If they are not the same, the system cannot ensure that the contents of the saved copy match the format of the copy on the system.

If you specify ALWOBJDIF(*NONE) on the restore command, the system does not restore the file or member if the creation dates do not match. A message is sent to the user to indicate that the file or member could not be restored from the media. ALWOBJDIF(*NONE) is the default value.

The creation dates on the system and media might be different because of the following reasons:

- A file or member was deleted and created again after the save operation.
- The file or member on the media was created on another system, but it has the same name as an existing file or member.

If you really want to restore a file or member whose creation date differs from the system version, you have some choices:

- Delete the file or member from the system. Then restore.
- Specify ALWOBJDIF(*COMPATIBLE) on the restore command. This value provides the *FILELVL function and allows all other object differences. The *COMPATIBLE value allows all object differences that are compatible with existing database files.
- Specify ALWOBJDIF(*FILELVL) on the restore command. This value lets you attempt to restore physical file data even though its creation date is different than the system copy creation date.

  You can specify *AUTL, *OWNER, and *PGP with *FILELVL on the ALWOBJDIF parameter. You can use one or more of these values to allow for differences in authorization lists (*AUTL), object ownership (*OWNER), and primary group (*PGP). These values let you filter which types of differences are allowed on the restore operation. If you use all four values, you achieve the equivalent of ALWOBJDIF(*COMPATIBLE).
- Specify ALWOBJDIF(*ALL) on the restore command. However, this can cause problems. You should be aware of what the system does when you specify ALWOBJDIF(*ALL).

How the system restores database files with ALWOBJDIF(*ALL): Figure 21 on page 242 shows what the system does when creation dates for a database file are different on the system and media copies:

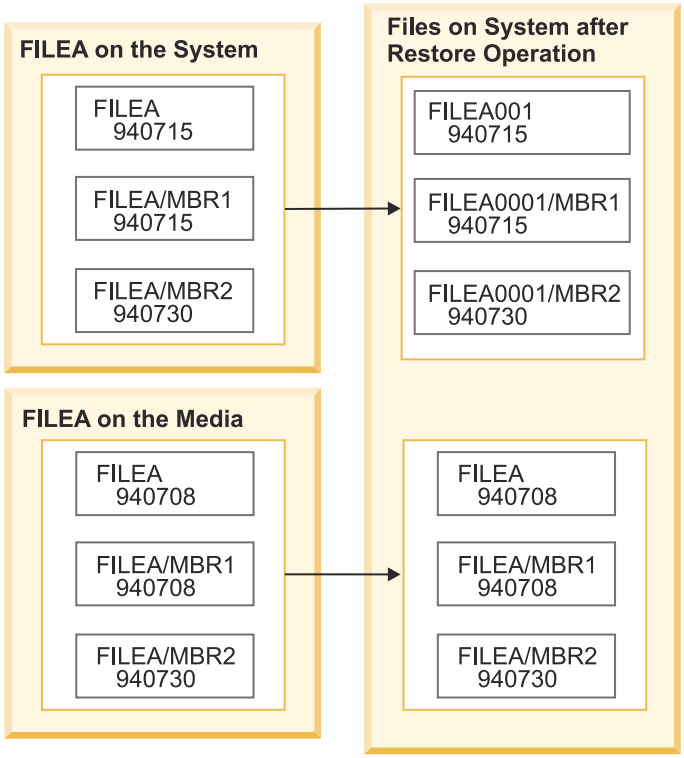

Figure 21. Restoring database files with different creation dates

The file on the system is renamed. The media version is restored. A message is sent to the user.

Figure 22 on page 242 shows what the system does when the creation date for one of the members in the file is different:

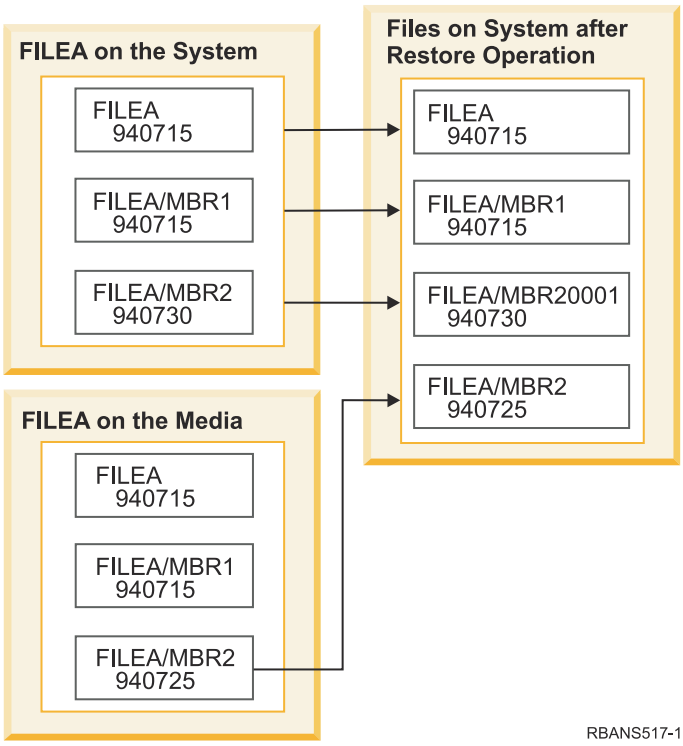

Figure 22. Restoring database files with different creation dates

The member on the system is renamed. All members from the media are restored. A message is sent to the user.

When you specify ALWOBJDIF(*ALL) and additional members are created during a restore operation, the system ignores the MAXMBRS (maximum members) parameter for the file. After the restore operation, you might have more than the allowed members in the file.

If a logical file is associated with a file or member that is renamed, the logical file is still associated with the renamed file or member, not the restored member.

In both examples, specifying ALWOBJDIF(*ALL) can result in duplicate information, additional files, and additional members. Your system becomes cluttered and your applications might produce unexpected results. If you specify ALWOBJDIF(*ALL), carefully check the messages you receive and analyze your files and members after the restore operation.

Notes:

- The ALWOBJDIF parameter also affects object ownership.
- When you specify MBROPT(*MATCH) on a restore command, you cannot specify ALWOBJDIF(*ALL).

Related concepts

What happens when you restore objects
When you restore an object, the system takes different actions depending on several conditions.

How the system establishes ownership for restored objects
Most objects on the system have an owner.

How the system matches file members during a restore operation
When you are restoring to an existing database file, you use the Member option (MBROPT) parameter on either the Restore Object (RSTOBJ) command or the Restore Library (RSTLIB) command to tell the system what to do if the members do not match.

How the system matches file members during a restore operation

When you are restoring to an existing database file, you use the Member option (MBROPT) parameter on either the Restore Object (RSTOBJ) command or the Restore Library (RSTLIB) command to tell the system what to do if the members do not match.

The choices are as follows:

***MATCH**
If the set of members on the save media and on the database are not identical, the restore operation fails. *MATCH is the default value.

***ALL**
All members on the save media are restored, whether they exist on the system copy.

***NEW**
Only those members on the save media that do not exist in the database file are restored.

***OLD**
Only those members on the save media that already exist in the database file are restored.

Note: The ALWOBJDIF parameter determines what the system does if creation dates on the members do not match.

Related tasks

Comparing file attributes during a restore operation

When you restore a database file or member that exists on the system, the system expects the creation dates for the system copy and the media copy to be the same. If they are not the same, the system cannot ensure that the contents of the saved copy match the format of the copy on the system.

Restoring members to a file

You can restore a list of members for a database file by using the FILEMBR parameter of the Restore Object (RSTOBJ) command.

This list might consist of specifically named members, generically named members, or a combination of both specifically and generically named members.

The FILEMBR parameter is used to specify the following items:

- A list of file members (specific or generic) for a specific database file
- The same group of members from more than one file

The default value *ALL causes all file members of files specified with the OBJ parameter to be restored.

Restrictions on the file member (FILEMBR) parameter

This topic provides information about the restrictions on the FILEMBR parameter.

The following restrictions apply to the FILEMBR parameter:

- Each database file that is specified in the FILEMBR parameter must also be specified in the OBJ parameter by its complete name, a generic name, or *ALL.
- Generic names are not valid for the database file name.
- Generic names are valid for the member name.

  If a generic file member name is used, and the file does not have members that fit the generic name, the file is not restored. If all files specified by the FILEMBR parameter are not restored, a diagnostic message is sent and the restore operation ends with an escape message giving the number of files not restored.

  If a name that is not generic is used, the specific members must exist in the file for any part of the file to be restored.
- The OBJTYPE must be *ALL or include *FILE.
- The MBROPT parameter must not have the *MATCH value.

Restoring logical files

When you restore a logical file, the system uses the description for the logical file to establish its relationship with the based-on files. For example, database views or indexes are logical files.

If the based-on files do not exist, you can do a deferred restore of those logical files. You also can restore physical and logical files in any order.

You can restore a logical file to a library different than the library for the associated physical file. However, the associated physical file must remain in or be restored to its original library location.

If you try to restore a logical file to a library in which it does not exist, the restore operation fails if any of the associated physical files have had their storage freed.

When a logical file is restored, it must be dependent on the same physical files as it was when it was saved.

- The logical file is created over the physical files in the library where they are being restored if any of the following situations occur:
  - The logical file and the associated physical files existed in the same library at the time of the save operation.

  - The logical file and the associated physical files are present in the library where the files are being restored.
  - The logical file and the associated physical files are being restored to the same library.
- If the files are not present in the restore library, then the logical files are created over the physical files in the original saved library.

  Note: When deferred files are restored, the system looks for the based-on files in the SAVLIB library, but not in the RSTLIB library.
- If the correct physical files are not found in either library, the restore operation of the logical file fails. To correct the problem, run the Restore Object (RSTOBJ) command again and specify OBJ(*NEW). If the restore operation is successful, an informational message (CPF3291) is sent to indicate which library was used for associated physical files.

The creation dates of the physical files must not have changed because the logical file was saved. If the date has changed, an informational message (CPF3293) is sent indicating that the physical file has been changed since the save operation, but the restore operation continues.

Restore physical or logical files with dependent logical files before the dependent logical files, unless the physical and logical files already exist on the system, or a defer ID is specified on the restore operation. The following considerations apply to restoring logical files:

- If the dependent physical or logical files are in the same library, the system provides the proper sequencing.
- If the files are in different libraries, you must restore the libraries in order, so that the physical or logical files that have logical files built on them are restored first, or you must specify a defer ID for the restore operation.
- If the depended-on physical or logical files are not restored before you attempt to restore the logical files, and if no defer ID is specified on the restore command, restoring the logical files fails.
- This sequencing also applies to other requirements between files, such as shared formats. You can restore those logical files that failed by using the RSTOBJ command.
- If the logical files and MQTs contain missing based-on files, you can set up a deferred restore by specifying the Defer ID (DFRID) parameter on the RSTOBJ or Restore Library (RSTLIB) command. When the based-on files become available, you can complete the deferred restore by using the Restore Deferred Objects (RSTDFROBJ) command.

Related concepts

Sequence for restoring related objects

Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Deferring the restore of dependent objects

You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Related information

Restore Deferred Objects (RSTDFROBJ) command

Restore Library (RSTLIB) command

Restore Object (RSTOBJ) command

How the system restores access paths

The description for a database file contains a description of its access path, if it has one. When you save a database file, you can save the access path with the file. This depends on the type of file, the type of access path, and how you performed the save operation.

When you restore a file, the system either restores the access path with the file or rebuilds the access path based on the information in the file description. The process of rebuilding the access path for a large database file can take a long time. This topic describes when the system restores access paths and when it cannot. If possible, you should plan your save operations to avoid having to rebuild access paths during a restore operation.

The system always restores the access path for a keyed physical file of type *DATA unless the access path was not saved. The access path for a keyed physical file is always saved, unless the access path is not valid at the time of the save.

Normally, source physical files are not keyed. The default value for the CRTSRCPF is to create a non-keyed file. When you restore a keyed source physical file, the access path is rebuilt after the restore operation.

Access paths that are owned by logical files are restored if all of the following conditions are true:

- The system saved the access path. Although this seems obvious, the system only saves access paths if the certain conditions are met.
- All based-on physical files are in the same library and are being restored at the same time on the same restore command.
- If the logical file exists on the system, it does not specify MAINT(*REBLD).
- The logical file owned the access path at the time it was saved.
- If the logical file is created again by the restore operation and it shares an access path that already exists, the key length for the access path must be equal to the maximum key length of the logical file or you receive an error.

If you meet these conditions, you minimize the rebuilding of access paths. However, during the restore operation, the system checks the integrity of each access path. If it detects any discrepancy, the access path is rebuilt.

In a few cases, the system might decide to rebuild access paths even though they were saved. For example, you might have defined a new logical file that specified the same key as the physical file but also specified UNIQUE. The based-on physical file was in use at the time that the logical file was created. Therefore, the system had to create a new access path for the logical file. Assume that you save these two files with a single command. If you restore them with a single command, the system will determine that they can share a single access path. Instead of restoring the two access paths, it builds a new, shared access path for the two files.

A database *index*, which is a type of logical file, cannot be restored if the associated physical file is missing. If the physical file is restored first, the index must be rebuilt, which can be time consuming. However, if the index is restored first and you specify a defer ID, the index is deferred and can be restored later, along with the physical file's data space. A deferred restore eliminates the need to rebuild the index. When you restore both the logical and physical files, specify the same Defer ID (DFRID) value on the Restore Object (RSTOBJ) or Restore Library (RSTLIB) command.

Related concepts

Sequence for restoring related objects
Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Related information

Backing up your server

Examples: Restoring a file network

This topic provides several examples of how to restore a file network.

Figure 23 on page 247 shows a physical file and two logical files:

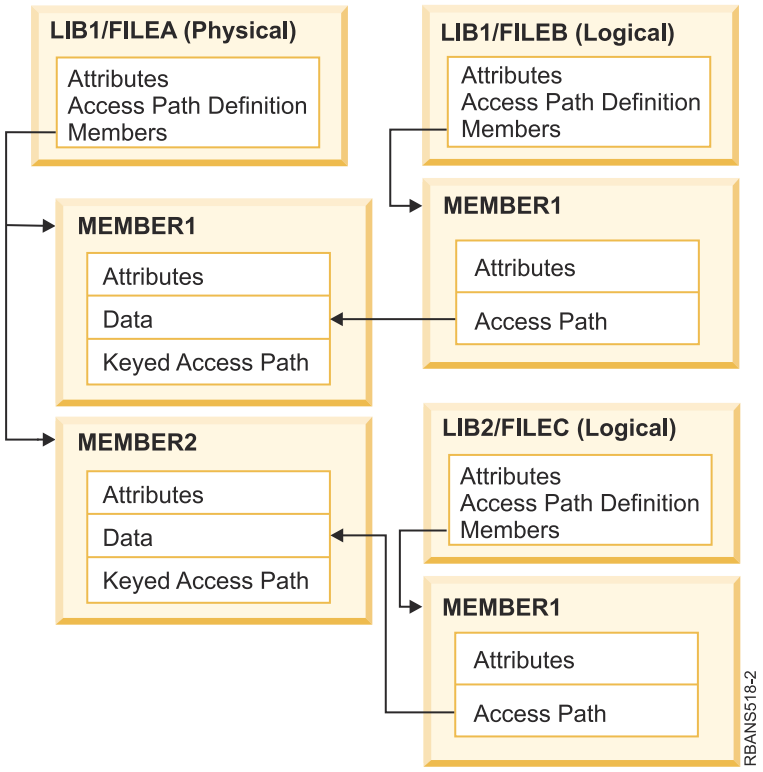

Figure 23. Restoring access paths

Assume that these files were saved with this command:

```
SAVLIB LIB(LIB1 LIB2) ACCPTH(*YES)
```

The save media volume contains all three files (FILEA, FILEB, and FILEC) and three access paths, each owned by a different file. Table 55 on page 247 shows what the system does when you restore these libraries by using different methods. These examples assume that none of the files are on the system when the system restores them.

Table 55. Restoring a file network

Sequence of restore commands	What the system does
Example 1: 1. `RSTLIB SAVLIB(LIB1)` 2. `RSTLIB SAVLIB(LIB2)`	Results of Example 1: 1. FILEA and FILEB are restored. The access paths owned by FILEA and FILEB are restored. 2. FILEC is restored. The access path owned by FILEC is rebuilt.
Example 2: 1. `RSTLIB SAVLIB(LIB2)` 2. `RSTLIB SAVLIB(LIB1)`	Results of Example 2. 1. FILEC is not restored because FILEA is not on the system. 2. FILEA and FILEB are restored. The access paths owned by FILEA and FILEB are restored.

Table 55. Restoring a file network (continued)

Sequence of restore commands	What the system does
Example 3: 1. RSTLIB SAVLIB(LIB2) DFRID(ABC) 2. RSTLIB SAVLIB(LIB1) DFRID(ABC) 3. RSTDFROBJ DFRID(ABC)	Results of Example 3. 1. Because you specified a Defer ID (DFRID) on the restore commands, the restore of FILEC is deferred. 2. FILEA and FILEB are restored. The access paths owned by FILEA and FILEB are restored. 3. FILEC is restored using the RSTDFROBJ command. The access path for FILEC is restored.

These examples highlight the problems that can occur when logical files and based-on physical files are in different libraries. Access paths are restored when physical files are restored because they are built over data that is contained in the physical file. In the first example, FILEC owned the access path but FILEC was not on the system when the physical file was restored. Therefore, the access path was not restored. In the second example, FILEC cannot be restored because its based-on physical file (FILEA) was not on the system.

The third example uses a deferred restore operation. FILEC depends on FILEA. When FILEA is restored, FILEC can be restored.

Preventing the system from rebuilding a large access path

This topic provides information about how to prevent the system from rebuilding a large access path.

If the situation that is shown in Table 55 on page 247 occurs on your system and you want to prevent the system from rebuilding a large access path, specify a defer ID on the restore command, or follow the steps below:

1. Restore the physical file or the library that contains the physical file.

   In the case of example 2, restore FILEA or LIB1.
2. Restore the logical file (FILEC) using the Restore Object (RSTOBJ) command.
3. Immediately after restoring the logical file, type EDTRBDAP.

   You are shown the Edit Rebuild of Access Paths display.
4. Change the value in the Seq column for the logical file to *HLD.
5. Restore the physical file (FILEA) again using the RSTOBJ command.

   Because the logical file (FILEC) is now on the system, the system will restore the access path that is owned by FILEC.
6. Type EDTRBDAP.

   You are shown the Edit Rebuild of Access Paths display.
7. Change the sequence number for FILEC to a value from 1 through 99 to remove the access path from the display.

How the system restores files with shared formats

When a database file is restored and that file, before it was saved, had shared the record format of another file, an attempt is made to find the file whose format was shared, and reestablish the original format sharing.

The search for restoring the shared format starts in the library to which the restored file is directed and continues in the library from which the restored file was saved. The results of the search is as follows:

- If the sharing file is found and has not been changed (level check) since the save, then no new format is created for the restored file.
- If the sharing file is not found, or it is found but fails the level check, then a new format for the restored file is created with the same definition as the one it initially shared.

- If a format sharing file has been renamed, deleted, or moved to a library other than the save or restore library, a new format is created for the dependent file when the dependent file is restored.

How the system restores files with referential constraints

Information about DB2 database files is kept in system cross-reference files. This includes information about constraints that are defined.

When you define a referential constraint, you specify that a record with a certain primary key must exist in the parent file before a record with the same values in a foreign key can exist in the dependent file. For example, you cannot add an order to the order file (dependent file) unless a record exists for the customer in the customer file (parent file).

A referential constraint is defined, stored, and saved with the dependent file. Each referential constraint has a name, which must be unique for the library that contains the dependent file. When you restore a file that has a referential constraint name that already exists in the library, the system generates a new name for the referential constraint that is being restored.

When you restore a database file that already exists on the system, the referential constraints that are defined for the system copy of the file are used. If the saved version of the file has additional referential constraints that are not on the system copy, these additional constraints are not restored.

When you restore a database file that does not exist, you should ensure that any referential constraints that were not on the saved copy are reestablished. Otherwise, you lose the data integrity checking that was on your system before a failure occurred.

Files that are related by referential constraints form a database network similar to the network formed by logical files and the based-on physical files. You should try to save an entire referential constraint network in one operation. If this is not possible, you should at least save the files with consecutive operations where no activity occurs in between. This ensures that the files are synchronized.

If you journal database files, you should journal all physical files that are part of a referential constraint network. This ensures that your referential constraints remain valid after you have applied journaled changes.

Related information

Journal management

Example: Referential constraint network

This figure shows an example of a referential constraint network.

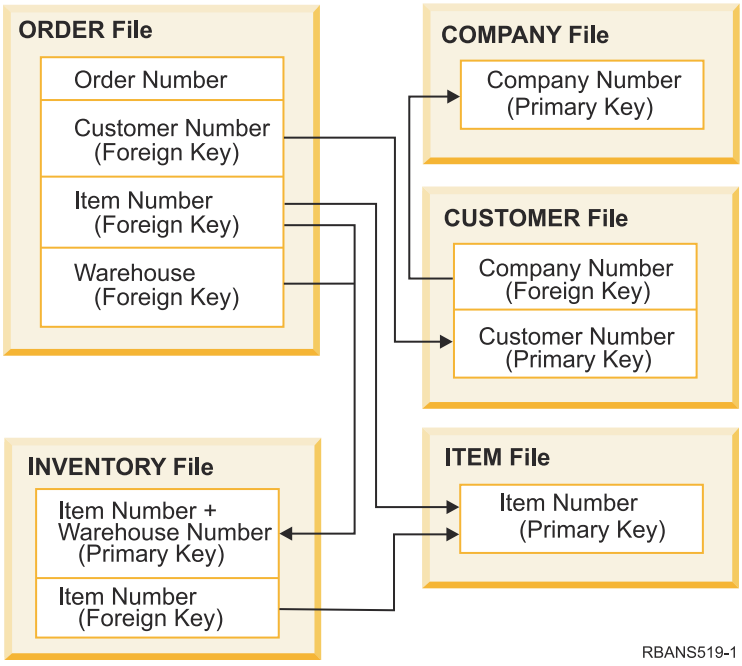

Figure 24. Restoring a referential constraint network

You can restore the files in this network in any sequence. When you restore the files, the system reestablishes the relationships and attempts to determine whether the constraints are still valid.

For example, if you restore both the ITEM file and the INVENTORY file, the system checks the internal information stored with the files to determine whether the indexes for the two files are synchronized.

If the internal information does not match, the system validates the constraint for the INVENTORY file. It does this by reading every record in the INVENTORY file and ensuring that a record with that item number exists in the ITEM file. If this process is successful, the constraint is valid. If this process is not successful, the status of the constraint is set to Check pending. When the status of a constraint is check pending, you must take action to correct the situation, either by restoring one of the files or using a program to update the files.

If you restore one of the files, the system again attempts to validate the constraint. If you use a program to update the information, you must use the Edit Check Pending Constraints (EDTCPCST) command to force the system to revalidate the constraint. The topic “Task 3: Editing check pending constraints” on page 164 describes how to determine the status of files that have referential constraints.

How the system restores files with triggers

This topic provides information about how the system restores files with triggers.

You can define one or more trigger programs for a file. When a certain event occurs in the file, the system calls the trigger program. When you save a file that has trigger programs, you are saving only the definitions of the trigger programs, not the programs themselves. You must ensure that the programs are also saved, perhaps by placing them in the library with the file.

When you restore a database file that already exists, the system does not restore any trigger program definitions from the save media. When you restore a database file that does not exist, you should ensure that any definitions for trigger programs that were not on the saved copy are reestablished. Otherwise, you lose the data integrity checking that was on your system before a failure occurred.

The system does not stop restoring a database file if its trigger programs cannot be found. Therefore, you must ensure that files and trigger programs are saved and restored correctly. Otherwise, an error might occur.

Table 56 on page 251 shows examples of actions the system takes when you restore the physical file FILEA and the trigger program PGMA:

Table 56. Restoring files that have trigger programs

Save procedure that is used	Restore procedure that is used	How the trigger program is defined after the restore operation
FILEA is saved from LIBX. PGMA is saved from LIBX. The trigger is defined as LIBX/PGMA.	PGMA is restored to LIBY. FILEA is restored to LIBX.	The trigger is defined as LIBX/PGMA. When an event occurs that causes this trigger, the program will not be found.
FILEA is saved from LIBX. PGMA is saved from LIBX. The trigger is defined as LIBX/PGMA.	PGMA is restored to LIBY. FILEA is restored to LIBY.	The trigger is defined as LIBY/PGMA.
FILEA is saved from LIBX. PGMA is saved from LIBY. The trigger is defined as LIBY/PGMA.	PGMA is restored to LIBZ. FILEA is restored to LIBZ.	The trigger is defined as LIBY/PGMA. When an event occurs that causes this trigger, the program will not be found.

The Journal management topic describes special considerations when you journal database files that have triggers defined. You must make special provisions to ensure the integrity of your data because trigger programs are not called when you apply journaled changes.

Restoring SQL materialized query tables

SQL materialized query tables (MQTs) are physical files, like logical files, that have dependencies on other files. You can defer the restore of logical files and MQTs whose based-on files or members are missing.

A deferred MQT is created, so that its data spaces can be restored from the media on the first pass. A deferred MQT is created, but until the restore operation is complete, you cannot perform functions on the MQT that require access to the based-on files.

In this example, you would use the Restore Deferred Object (RSTDFROBJ) command to complete the restore of a deferred MQT. The logical files and MQT, which have multiple based-on files, are created using database commands. Then libraries A and B are saved using the Save Library (SAVLIB) command, deleted, and restored using the Restore Library (RSTLIB) command. The logical files and MQT being restored are deferred, because their based-on files do not exist and because a defer ID is specified on the RSTLIB command. An *implicit restore* operation occurs only if you save and restore the same files using the SAVLIB and RSTLIB commands. Otherwise, you must use the RSTDFROBJ command to complete the restore operation for the deferred files.

```
CRTLIB A
CRTLIB B
CREATE TABLE B/T1 (C1 INT)
CREATE TABLE B/T2 (C1 INT)
CREATE VIEW A/M1
    AS (SELECT * FROM B/TI UNION SELECT * FROM B/T2)
    DATA INITIALLY DEFERRED
    REFRESH DEFERRED
    ENABLE QUERY OPTIMIZATION
    MAINTAINED BY USER
SAVLIB LIB(A B) DEV(TAP01)
DLTLIB LIB(A)
DLTLIB LIB(B)
RSTLIB SAVLIB(A B) RSTLIB(*SAVLIB) DFRID(FLOW4) DEV(TAP01)
```

If the completion message indicates that not all of the deferred files were restored, issue the RSTDFROBJ DFRID(FLOW4) command to complete the restore operation for the deferred files.

The Remove Defer ID (RMVDFRID) command removes information about objects that were deferred during the restore operation. Use this command to remove the deferral information for the MQT if, for example, the deferred MQT is deleted and recreated with new based-on files.

```
RMVDFRID DFRID(FLOW4)
```

Related concepts

Deferring the restore of dependent objects
You can defer the restore processing for some objects that depend on related objects when the related objects are not available.

Related information

Restore Deferred Objects (RSTDFROBJ) command
Materialized query table optimization

Deleting a physical file

In some situations, you must delete a physical file as part of your recovery.

For example, the physical file might be damaged. Or a physical file in a user auxiliary storage pool (ASP) might have overflowed into the system ASP. You cannot delete a physical file if other files are dependent on it, such as logical files or files that share the record format.

Before deleting a physical file, follow these steps:

1. Use the Display Database Relationships (DSPDBR) command to list all the files that are dependent on the physical file.
2. Save and delete each file that is dependent on the physical file.

After you have recovered the physical file, restore all the dependent files.

Restoring physical files that are temporal

Physical files that contain ROW BEGIN, ROW END, TRANSACTION START ID fields become temporal physical files when you perform an ALTER TABLE ADD VERSIONING against the system-period temporal table and specify its associated history table. A temporal file and its associated history table are in the same relationship network. If possible, restore all the files in a relationship network in a single restore operation.

If you restore files that are in a relationship network during separate restore operations, the system must verify that the relationships are still valid and current before implicit versioning is allowed. A CPD32FE warning is signaled when you restore files that are in the same relationship network separately. You can avoid this warning and improve restore performance if you save and restore relationship networks in a single operation.

Restoring journals and journal receivers

Typically, you can restore journals or journal receivers only to the same library from which they were saved. The one exception is journal receivers that are part of a remote journal network. These receivers can often be restored to the remote receiver library as well.

Use the Restore Object (RSTOBJ) and Restore Library (RSTLIB) commands to restore journals and journal receivers. When you are restoring multiple objects with one of these commands, journals and journaled objects are restored before the journal receivers.

When you use several commands to restore several objects, restore the objects in the following order:

1. Journals
2. Based-on physical files

3. Other journaled objects associated with those journals
4. Dependent logical files
5. Journal receivers

   Journal receivers can be restored at any time after the journals. They do not have to be restored after the journaled objects.

Note: Materialized query tables can be dependent on both physical and logical files.

Related concepts

Sequence for restoring related objects

Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Related information

Recovery operations for journal management

Considerations for save and restore operations with remote journals

Restoring journals

When you restore a journal, the system creates a new journal receiver and attaches it.

The characteristics of the new journal receiver are based on the journal receiver that was attached when the journal was saved:

- The system creates a name that is not likely to conflict with other journal receivers that might be on the system.
- The system attempts to assign the same owner and to create the journal receiver in the same library. If the owner of the receiver is not found, the receiver is assigned to the default owner (QDFTOWN) user profile. If the library is not found, the journal receiver is placed in the journal's library.
- The system starts a new receiver chain.

Note: At the time a new journal receiver is created and attached, private authorities have not been restored on the system. Therefore, private authorities will not be assumed by the new journal receiver. After the Restore Authority (RSTAUT) command is run, users will receive private authority to the receiver that was attached before the restore operation. Users will not receive private authority to the new receiver. Users must be manually granted private authority to the new receiver.

You cannot restore a journal to a library that contains the same journal. If a journal must be restored (because of damage) to a library, the existing journal must be deleted first.

Related information

Naming conventions for journal receivers

Naming conventions for journals

Journal receiver chains

Deleting a journal

In some situations, you must delete a journal as part of your recovery.

For example, the journal might be damaged, or a journal in a basic auxiliary storage pool (ASP) might have overflowed into the system ASP. You cannot delete a journal while objects are being journaled to it.

You use the Delete Journal (DLTJRN) command to delete a journal. Before deleting a journal, try to do the following steps. You might not be able to perform these steps successfully if the journal is damaged.

1. Type

```
WRKJRNA JRN(library-name/journal-name)
OUTPUT(*PRINT)
```

and press the Enter key. You receive a listing that shows all the objects that are currently being journaled.

2. End journaling for all the access paths assigned to the journal by typing the following command:

```
ENDJRNAP FILE(*ALL)
JRN(library-name/journal-name)
```

3. End journaling for all the physical files assigned to the journal by typing the following command:

```
ENDJRNPF FILE(*ALL)
JRN(library-name/journal-name)
```

4. End journaling for all the integrated file system objects assigned to the journal by typing the following command:

```
ENDJRN OBJ(*ALL)
JRN(/QSYS.LIB/library-name.LIB/journal-name.JRN)
```

5. Save the inheritance rules associated with all libraries assigned to the journal by using the following command:

```
DSPLIBD  LIB(library-name) OUTPUT(*PRINT)
```

6. End journaling for all the libraries assigned to the journal by typing the following command:

```
ENDJRNLIB  LIB(*ALL) JRN(library-name/journal-name)
```

7. End journaling for all other object types assigned to the journal by typing the following command:

```
ENDJRNOBJ OBJ(*ALL) OBJTYPE(*ALL)
JRN(library-name/journal-name)
```

8. Deactivate any remote journals that are associated with the journal by using the Change Journal State (QjoChangeJournalState) API or the Change Remote Journal (CHGRMTJRN) command.

When you try to delete the journal, you might receive message CPF7021 indicating that the journal is being used for commitment control. If this occurs, end the jobs that are using commitment control and then try to delete the journal again. To see the commitment control uses of the journal use the Work with Journal Atrributes (WRKJRNA) command, function key 19 (F19=Display journaled objects), and option 6 (6=Commitment Definitions). You can use the End Job (ENDJOB) command or you can use the *End* option from the Work with Active Jobs (WRKACTJOB) display.

After you restore the journal or create it again, you must start journaling again for each object. Use the following commands to begin journaling for each object type listed below:

- Database physical files — STRJRNPF
- Access paths — STRJRNAP
- Integrated file system objects — STRJRN
- Libraries — STRJRNLIB (Use the inheritance rules that you obtained earlier with the DSPLIBD command.)
- All other object types — STRJRNOBJ

You should save the objects after you start journaling, in case the system assigned a new journal identifier (JID) to an object. If you previously had any remote journals associated with the journal, add them again by using the Add Remote Journal (ADDRMTJRN) command or the Add Remote Journal

(QjoAddRemoteJournal) API. If you added any remote journals, you should save the journal in order to preserve that information.

Restoring journal receivers

The system does not restore a journal receiver over the journal receiver that is currently attached. The system does not restore a journal receiver over an existing journal receiver that contains more entries.

If you use the Save Changed Objects (SAVCHGOBJ) command to save journal receivers, this is likely to occur. The journal receiver that is attached at the time of the save operation is a changed object and is saved by the command. When you restore, you receive message CPF3706 and the system continues with the next journal receiver.

If your save procedure saves the currently attached journal receiver, you can try to restore a journal receiver with fewer entries than the journal receiver on the system. For example, assume that you save your journal receivers when receiver RCVR0006 is attached. RCVR0006 has 1500 entries. Later, you use the CHGJRN command to create and attach a new receiver. Now receiver RCVR0007 is attached. Receiver RCVR0006 is still on the system and has 4300 entries. If you attempt to restore receiver RCVR0006 from your save media volume, the operation fails because the saved copy has only 1500 entries.

If the library you specify on the restore command for a journal receiver does not exist, the system restores the journal receiver to the library that contains the journal. If you specify RSTASP(*SAVASP) and the ASP does not exist, the system typically restores the journal receiver to the same auxiliary storage pool (ASP) as the library that contains the journal.

Placing journal receivers in the correct auxiliary storage pool: If the attached journal receiver is not in the ASP that you want after the restore operation, follow these steps:

1. Create a journal receiver in the ASP that you want.

   Follow your existing naming convention and use the same journal receiver attributes.
2. Use the Change Journal (CHGJRN) command to attach the new journal receiver to the journal.

Resolving name conflicts when restoring journal receivers

When you restore a journal, the system creates and attaches a new journal receiver. The system attempts to name this journal receiver so that a name conflict does not occur.

However, in rare cases, the new journal receiver that the system creates after you restore a journal might have a name that matches the name of a journal receiver that you need to restore. If this occurs, follow these steps:

1. Create a new journal receiver with a name separate from your normal naming convention.
   For example, type: `CRTJRNRCV JRNRCV(`*`library-name`*`/TMP0001)`.
2. Use the Change Journal (CHGJRN) command to attach the temporary journal receiver: `CHGJRN JRN(`*`library-name/journal-name`*`) JRNRCV(library-name/TMP0001)`.
3. Delete the journal receiver that has the name conflict.

   This journal receiver should not have any entries you need for your recovery because it was created when the journal was restored.
4. Restore the journal receivers.
5. Create a journal receiver that continues your naming convention and has the same journal receiver attributes.
6. Use the CHGJRN command again to attach the journal receiver that you created in step 5.

Correcting the journal receiver directory

Before you begin a recovery by using journal receivers, you should ensure that this receiver directory is current and correct.

Every journal has a directory of journal receivers. The sequence of journal receivers is called the *receiver chain*.

1. Type `WRKJRNA JRN(`*`library-name/journal-name`*`)` and press the Enter key.
2. From the Work with Journal Attributes display, press F15 (Work with receiver directory).

   You are shown the Work with Receiver Directory display.
3. If the receiver directory is not correct, follow these steps:
   a) Type `WRKJRN` and press the Enter key.
   b) On the prompt display, enter the name of the journal.
   c) On the Work with Journals display, type a 9 (Associate receivers with journal) in the option column in front of the journal. The system establishes the receiver chain for the journal.

Deleting a journal receiver

In some situations, you must delete a journal receiver as part of your recovery.

For example, the journal receiver might be damaged. Or a journal receiver in a user auxiliary storage pool (ASP) might have overflowed into the system ASP.

You cannot delete a journal receiver that is currently attached to a local journal. You also cannot delete a journal receiver if later journal receivers in the receiver chain are still on the system, unless one of the following conditions exists:

- The receiver being deleted is damaged
- The journal is a remote journal
- The journal is system-managed and the system is deleting journal receivers

If you need the journal receiver for recovery, you should not delete it without first saving it. If you do, the system warns you but does not prevent you from deleting the journal receiver.

Before deleting a journal receiver, follow these steps:

1. If the journal receiver is attached, detach it by typing the following command:

   ```
   CHGJRN JRN(library-name/journal-name)
   JRNRCV(*GEN)
   ```

   Note: If the current journal receiver is damaged, you cannot specify `JRNRCV(*GEN)`. Use the Create Journal Receiver (CRTJRNRCV) command to create a new journal receiver that follows your naming convention and has the same attributes. Specify that receiver name on the Change Journal (CHGJRN) command.
2. If earlier journal receivers are on the system, save them and delete them.

   You can print the receiver chain by typing `WRKJRNA JRN(`*`library-name/journal-name`*`) OUTPUT(*PRINT)`.

How the system restores programs

Restoring programs to your system represents a security exposure. A restored program might be altered to do functions that you do not intend, or the program might adopt the authority of a powerful user profile.

When the QSECURITY (security level) system value on your system is 40 or higher, the system checks for restricted instructions in all programs that are restored. You can use the QALWOBJRST system value to allow or prevent the restoration of certain types of objects on your system. You can also set the QVFYOBJRST (verify object on restore) system value to specify how the system verifies program-object signatures during a restore operation.

For an object to be successfully restored, it must satisfy the three system values that work together during a restore operation.

- Verify object signatures during restore (QVFYOBJRST).
- Force conversion on restore (QFRCCVNRST).

- Allow restore of security-sensitive objects (QALWOBJRST).

However, if the Force object conversion (FRCOBJCVN) parameter on the restore command is not compatible with the QFRCCVNRST system value, then no conversion occurs and nothing is restored.

The system stores a validation value for all programs. When a program is restored, the system calculates the validation value and compares it to the value on the media. You can use the QFRCCNVRST and QALWOBJRST system values to control the action to take if programs are restored with validation errors. You can choose to have the system re-create the program. If re-creation is successful, the program is restored and the validation error is corrected. If re-creation is not successful, the program is not restored. As an alternative, you can choose not to have the system attempt re-creation to prevent any program with a validation error from being restored. A third alternative is to not attempt re-creation and have the system restore the program with the validation error (which might be a security risk). The system contains all the information necessary to re-create an IBM i program.

The system can now create *MODULE, *PGM, and *SRVPGM objects that take advantage of the most recent hardware features of the system where they are created. New hardware features can be used immediately. In V5R4 and previous releases, the system did not use new processor features until all models supported by the current release of the operating system had the new feature. For more information, see the "Adaptive Code Generation" section in the "Advanced Optimization Techniques" chapter of ILE Concepts .

Programs that use the most recent hardware features might require conversion when moved to a system that is running on an older version of the hardware. The program might need to be converted during restore because it uses at least one feature the system does not support. To restore a program without converting it, set the Force conversion on restore system value (QFRCCVNRST) to 0 before to restoring the program and specify FRCOBJCVN(*NO) on the restore command.

When an external procedure associated with an Integrated Language Environment® (ILE) external program is created, an attempt is made to save the attributes for the procedure in the associated program object. If the *PGM object is saved and then restored to this or another system, the catalogs are automatically updated with those attributes.

The attributes can be saved for external procedures subject to the following restrictions:

- The external program library must not be QSYS or QSYS2.
- The external program must exist when the CREATE PROCEDURE statement is issued.
- The external program must be an ILE *PGM object.
- The external program must contain at least one Structured Query Language (SQL) statement.
- If the object cannot be updated, the procedure is still created.

During the restore operation of the procedure, the following items might occur:

- If the procedure name that is being restored is not unique, an error is issued.
- If the specific name was not specified, a unique name is generated, if necessary.
- If the procedure name and number of parameters are not unique, the procedure cannot be registered, and an error is issued.

Related concepts

What happens when you restore objects
When you restore an object, the system takes different actions depending on several conditions.

Examples: How the system restores authority
When you run the Restore Authority (RSTAUT) command, the system grants all the private authorities that it finds in each authority reference table.

Controlling restoration of security-sensitive objects
You can use two different system values to control the restoration of security-sensitive objects.

Related tasks

Restoring a library from a previous release

When you are restoring a library that was saved on a system at an earlier release, you can use the Force object conversion (FRCOBJCVN) parameter to specify whether programs are translated when they are restored.

Related information

Effects of system value settings on restore operations

ILE Concepts PDF

Restoring programs to a different release

This topic provides information about how to restore programs to a different release.

The internal object format for program objects is different between IBM i V6R1 and earlier releases. When you move a program object (*MODULE, *PGM, *SRVPGM, *SQLPKG) between a release earlier than V6R1 and V6R1 or later, the system must convert the program object using the information that is stored with the program. If you are restoring an older format program onto V6R1 or later, the settings of the Force conversion on restore (QFRCCVNRST) system value and the Force object conversion (FRCOBJCNV) parameter on the restore command control whether the program object gets converted during the restore operation. If you set the QFRCCVNRST system value to 2 or greater, or use the FRCOBJCVN (*YES *ALL) parameter on the restore command, the program object gets converted on the restore.

In contrast, a V6R1-format program object is always converted when restored to a system running a release earlier than V6R1, regardless of the system value and restore command settings.

Object conversion occurs at one of the following times:

- When the object is used for the first time. This is the default behavior.
- When you use the Start Object Conversion (STROBJCVN) command to convert objects. This command operates on an entire library.
- When you restore the object, subject to the conditions described above.

You also can use the Change Module (CHGMOD), Change Program (CHGPGM), and Change Service Program (CHGSRVPGM) commands to re-create the object. The object is re-created in the format required by the system.

For information about how the system can take advantage of the latest hardware features of the system, see the How the system restores programs topic.

Related tasks

Restoring a library from a previous release

When you are restoring a library that was saved on a system at an earlier release, you can use the Force object conversion (FRCOBJCVN) parameter to specify whether programs are translated when they are restored.

Restoring save file data

You can save a save file to tape or optical media with the Save Save File Data (SAVSAVFDTA) command.

When you restore the save file, it appears as though the data originally came from the same type of save media. You can use the Restore Object (RSTOBJ), Restore Library (RSTLIB), Restore Object (RST), Restore Document Library Object (RSTDLO), Restore Configuration (RSTCFG), or Restore User Profiles (RSTUSRPRF) commands to restore the data.

You can save file data to tape or optical media by using the Save Library (SAVLIB), Save Object (SAVOBJ), or Save Changed Object (SAVCHGOBJ) commands. If you specified SAVFDTA(*YES) on the save command, you must restore the save file before you can restore the objects in the save file.

Saving and restoring spooled files

You can use any of the methods that are described here to save and restore spooled files.

Table 57. Save and restore spooled files

Save Methods	Restore Methods	Spooled file attributes preserved
SAVLIB, SAVOBJ, SAVCHGOBJ, SAVRSTLIB, SAVRSTOBJ, SAVRSTCHG commands	RSTLIB, RSTOBJ, SAVRSTLIB, SAVRSTOBJ, SAVRSTCHG commands	Data and all attributes
QSRSAVO API	QSRRSTO API	Data and all attributes
Save menu options 21-23	Restore menu options 21-23	Data and all attributes
QSPOPNSP, QSPGETSP, QUSRSPLA APIs	QSPCRTSP, QSPPUTSP, QSPCLOSP APIs	Data, but not all attributes[1]
CPYSPLF, SAVOBJ commands	CPYF command	Textual data only[1]

[1]For releases before IBM i V5R4, you must use one of these methods for saving and restoring spooled files. These methods do not preserve all of the spooled file attributes.

When you save an output queue with the SAVLIB or SAVOBJ commands, SAVE menu, or QSRSAVO API, you can choose to save all its spooled files. You can do this by specifying *ALL for the Spooled file data (SPLFDTA) command parameter, menu prompt, or API key.

When you save an output queue with the SAVCHGOBJ command, you can choose to save its new spooled files that were created after the reference date and time. You can do this by specifying *NEW for the Spooled file data (SPLFDTA) command parameter.

When you restore output queues with the restore commands, menu, or QSRRSTO API, you can choose to restore any saved spooled files that do not already exist on the system. You can do this by specifying *NEW for the SPLFDTA parameter, prompt, or key.

With the QSRSAVO and QSRRSTO APIs, you can also choose to save or restore spooled files by using a set of selection criteria. If you save spooled files with the QSRSAVO API using selection criteria and the *SPLF special library value, then you must use the QSRRSTO API with the *SPLF special library value to restore the spooled files.

The QSRSAVO and QSRRSTO APIs also allow you to change the expiration date of the spooled files. The QSRRSTO API also allows you to restore spooled files to a different output queue and allows you to create an additional copy of existing spooled files.

This example describes how to save spooled files:

1. Create an output queue to store the spooled files.

```
CRTOUTQ  OUTQ(lib-name/que-name)
```

2. Use the Work with Spooled File (WRKSPLF) command to list the spooled files.
3. Use option 2, Change Spooled File Attributes (CHGSPLFA) command to move the spooled files you want to save to the output queue you created.
4. Use the Save Object (SAVOBJ) command to save the spooled file data.

```
SAVOBJ  OBJ(que-name) LIB(lib-name) DEV(dev-name) OBJTYPE(*OUTQ) SPLFDTA(*ALL)
```

This example describes how to restore spooled files:

1. Restore spooled files that are not already on the system.

```
RSTOBJ  OBJ(que-name) SAVLIB(lib-name) DEV(dev-name)
        OBJTYPE(*OUTQ) SPLFDTA(*NEW)
```

Using spooled file APIs

If your source or target system is earlier than V5R4, you can use spooled file APIs as an indirect method to save and restore spooled files. This method preserves the spooled file data stream but not all of the attributes.

To save spooled files:

- The spooled files are opened using the Open Spooled File (QSPOPNSP) API.
- The spooled file data is retrieved using the Get Spooled File Data (QSPGETSP) API.
- The spooled file attributes are retrieved using the User Spooled File Attributes (QUSRSPLA) API.

To restore spooled files:

1. The spooled files are created using the Create Spooled File (QSPCRTSP) API.
2. The spooled file data is written to a new spooled file using the Put Spooled File Data (QSPPUTSP) API.
3. The spooled file is closed using the Close Spooled File (QSPCLOSP) API.

You can find an example and a tool for using these APIs in the QUSRTOOL library in the TSRINFO member of the QATTINFO file.

Copying spooled files to database files

If your source or target system is earlier than V5R4, you can copy data between spooled files and database files as an indirect method to save and restore spooled files. This method copies textual data only and not advanced function attributes such as graphics and variable fonts. This method does not provide a complete solution for saving your spooled files.

The Copy Spooled File (CPYSPLF) command saves spooled file data to a database file. The Copy File (CPYF) command can copy data from a database file to a spooled file.

Related tasks

Saving the object for the previous release

Save the object on the current release by using the target release (TGTRLS) parameter before you restore it on the previous release or previous modification. The object is saved in a format that is known to the previous release or previous modification.

Restoring licensed programs

You can use the Restore Licensed Program (RSTLICPGM) command to add or replace licensed programs on the system.

Related information

Using the Restore Licensed Programs command to install IBM or non-IBM licensed programs

Restoring documents and folders

You can use the Restore Document Library Object (RSTDLO) command to restore documents, folders, and mails.

To use this command most efficiently, you should know how documents were saved. To determine this, use the output that was printed for the Save Document Library Object (SAVDLO) procedures, the Display Tape (DSPTAP) command, or the Display Optical (DSPOPT) command. RSTDLO performance is also better if you have *SAVSYS special authority.

Restore Document Library command options

The Restore Document Library (RSTDLO) command provides many options.

You can restore any of the following options:

- A specific document or system object whose name you specify.
- All the documents and folders you saved by typing: `RSTDLO DLO(*ALL) SAVFLR(*ANY)`. If you saved DLOs from more than one auxiliary storage pool (ASP), you must specify `SAVASP(*ANY)`. You must also specify the sequence numbers (SEQNBR parameter) for the files on the save media.

  Note: When you use RSTDLO DLO(*ALL), this includes the folders that are used by IBM-supplied programs, such as IBM i Access. Ensure that you saved these folders from the current release, or you might need to install the licensed programs again.
- 1 to 300 documents from the same media file by specifying the names of the documents or the system object names.
- 1 to 300 folders from the same media file.
- All filed documents that are not in any folder on the save media.

Related tasks

Restoring folders

To restore a folder object, the entire folder (the folder object plus all document and folder objects within it) must also be restored.

Using multiple concurrent document library object commands

Multiple concurrent Save Document Library Object (SAVDLO) or Restore Document Library Object (RSTDLO) commands can be used in specific situations.

No two of the following commands can be run on one system at the same time:

- RCLDLO DLO(*ALL)
- RCLDLO DLO(*DOCDTL)
- RCLDLO DLO(*INT)
- DLTDLO DLO(*ALL)
- RNMDIRE

An attempt to run these commands at the same time results in the message CPF8A47:

```
Internal system objects are in use.
```

An attempt to run a SAVDLO or RSTDLO operation while one of these commands is running also results in message CPF8A47 and no objects are saved or restored.

Output from the Restore Document Library Object Command

You can use the OUTPUT parameter on the Restore Document Library Object (RSTDLO) command to show information about the restored documents, folders, and mails.

You can either print the output (OUTPUT(*PRINT)) or save it to a database file (OUTPUT(*OUTFILE)).

If you print the output, you should be aware of device dependencies:

- The heading information in the output is device-dependent. All information does not appear for all devices.
- The print file for the RSTDLO command uses a character identifier (CHRID) of 697 500. If the printer you are using does not support this character identifier, you will receive message CPA3388. To print the RSTDLO output and not receive message CPA3388, specify the following before specifying *PRINT on the RSTDLO command:

```
CHGPRTF FILE(QSYSOPR/QPRSTDLO) CHRID(*DEV)
```

For more information about character identifiers (CHRID), see the *Printer Device Programming* book.

If you use an output file, the system uses the file format QSYS/QAOJRSTO.OJRDLO. The file layout is described in the *Office Services Concepts and Programmer's Guide* book.

Considerations and restrictions

You should be aware of these additional factors when using the Restore Document Library Object (RSTDLO) command.

Moving documents

When you restore documents, you can rename them, restore them to a different folder, or have the system assign new system object names.

The folder for a document determines its auxiliary storage pool (ASP) location. You can move a document to a different ASP by following these steps:

1. Save the document.
2. Delete it with the Delete Document Library Object (DLTDLO) command.
3. Restore it into a folder in a different ASP.

Searching tape files

When you restore documents or folders from a list and specify SEQNBR(*SEARCH), the system restores from the first tape file that contains any of the documents or folders that you specified.

If the tape file does not contain all the documents and folders in your list, the system does not search other tape files for the additional documents and folders. You can specify SEQNBR(*starting-sequence ending-sequence*) to search more than one tape file.

Selecting files from DVD-RAM optical media

The OPTFILE and SAVASP parameters control which files that the system uses.

If you specify a file path name, the system uses that file. If you specify the default value of OPTFILE('*') or OPTFILE('*directory-path-name*/*'), the system uses files named QDOC or QDOC*nnnn* in the directory that you specify, depending on the SAVASP value.

Searching index database errors

When you restore document library objects (DLOs), the system updates the search index database information for the DLOs.

If you receive error messages during the restore procedure because the information in the database does not match the DLOs, run the Reclaim Document Library Object (RCLDLO) command. Then try the restore procedure again.

Note: The message tells you if the RCLDLO procedure is necessary. Use RCLDLO only if you are instructed by a message or by the recovery checklist you are using.

Authority required to restore document library objects

If you are restoring document library objects (DLOs) into a folder, you must have authority to the folder.

If you are restoring existing DLOs, you must have authority to those DLOs. Certain combinations of the Restore Document Library Object (RSTDLO) command require additional authority. The Document library object commands IBM i documentation topic provides information about the specific authorities that are required for the RSTDLO command.

How the system restores new document library objects

When you restore new document library objects (DLOs), the system files them.

The DLO is treated as new to the system if any of the following situations is true:

- It has been previously deleted.
- It is being restored to a different system.
- It is being restored with the NEWOBJ(*NEW) parameter.

How the system restores existing document library objects

When you are restoring an existing document library object (DLO), the system skips the DLO and continues with the next one if either of these situations is true.

- The DLO is in use.
- You do not have the necessary authority.

If the existing document is damaged, some of the security information might be lost. The restore operation continues and a message is sent informing you that the document is damaged and some of the security information is lost.

Size limitations when restoring document library objects

This topic provides information about the size limitations when restoring document library objects (DLOs).

You cannot restore more than 1 000 000 DLOs to a single user auxiliary storage pool (ASP).

Related information

Save and restore limits

Restoring folders

To restore a folder object, the entire folder (the folder object plus all document and folder objects within it) must also be restored.

However, if the specific folder being restored was stored in other folders at the time it was saved, those higher level folders do not have to be restored to restore the specific folder.

When you restore a folder, the fully qualified folder path name you are restoring must exist unless you are restoring a first-level folder. For example, if you save folder A and then delete it, you can enter RSTDLO DLO(*ALL) SAVFLR(A) and restore folder A in addition to all the documents and folders in it. However, if you want to restore folder A/B/C/D, you must create folder A, then folder B in folder A, then folder C in folder A/B, before you can restore folder D in folder C. You only have to create the folders that comprise the A/B/C path, and you do not have to create folder D in folder A/B/C before you can restore it.

If you try to restore a folder that is in use, the system bypasses restoring the folder and all the document library objects (DLOs) in it.

If you try to restore into an existing folder but the folder is damaged and cannot be reclaimed, you receive a message informing you that the folder is damaged and not restored. The folder and all documents and folders in it are not restored.

Related concepts

Restore Document Library command options

The Restore Document Library (RSTDLO) command provides many options.

Renaming documents when restoring

You can use the RENAME parameter to give documents a different name when they are restored. You can also place them in a different folder by using the RSTFLR parameter.

If renaming a document when it is restored would result in a duplicate name in a folder, the system does the following actions:

- If ALWOBJDIF(*NONE) is specified, the document is not restored.
- If ALWOBJDIF(*ALL) is specified, the document is restored and replaces the existing document in the folder.

You can specify more than one value for the RENAME parameter. The system matches the RENAME values with the DLO values until it runs out of values for one or the other. Assume that you specify the following command:

```
RSTDLO DLO(A B C D) SAVFLR(X) RENAME(J K L) RSTFLR(Y)
```

After the restore operation, you have the following documents:

- Document J in folder Y
- Document K in folder Y
- Document L in folder Y
- Document D in folder Y

How the system restores descriptive information for document library objects

When restoring descriptive information for document library objects (DLOs), the system sets the creation date, file date, and revision date for restored documents and folders.

- The creation date of the document or folder on the save media is restored with the document or folder.
- When the RSTDLO command replaces a document or folder, the file date of the document or folder being replaced on the system is used.
- The object revision date is set to the current date when the document or folder is restored.
- The content revision date of the document on the save media is restored with the document.
- The content revision date is set to the current date when replacing a folder.
- The content revision date of the folder on the save media is restored with the folder if the folder is new.

How the system restores authority and ownership for document library objects

The system restores authority and ownership for document library objects (DLOs) in the same way as it handles authority and ownership for restored objects, with these additions.

- If the user profile that owns a DLO is not in the system distribution directory, ownership is assigned to the QDFTOWN user profile.
- When you restore a DLO that does not exist on the system, any access codes and explicit users are removed. If you have restored user profiles and you later run the Restore Authority (RSTAUT) command, the private authorities to the DLO are restored. The access codes are not restored.

See the "How the system establishes ownership for restored objects" on page 216 and "Restoring object authorities" on page 217 topics for how the system handles ownership and authority when restoring objects.

When to run the Rename Directory Entry command

When you need to run the Rename Directory Entry (RNMDIRE) command for a local user, schedule it before you perform these operations.

- Saving mail
- Saving the system distribution directory

If the rename operation is performed just before saving the mail and the directory, the changed information is saved and the information will be the same as what is on the system. If the information on the media does not match the information on the system, the mail will not be restored during the restore operation.

When to run the Rename Document Library Object command

When you need to run the Rename Document Library Object (RNMDLO) command, schedule it just before you back up document library objects.

If the rename operation is performed just before saving the document library object, the changed name is saved and the information on the media will be the same as what is on the system.

If you rename a document library object after a save operation, the document library object name on the system is different than the name on the media. However, the system object names remain the same. The restore operation fails because the system thinks that the document library object already exists. Message CPF90A3 or CPF909C is sent indicating that the document or folder already exists.

- To create a new document or folder, specify NEWOBJ(*NEW).
- To replace an existing document, specify RENAME(*document-name*), where *document-name* is the name that is given to the document by the RNMDLO command.
- To replace an existing folder, specify RSTFLR(*folder-name*), where *folder-name* is the name that is given to the folder by the RNMDLO command.

Restoring objects in directories

You can use the Restore Object (RST) command to restore objects that you saved with the Save Object (SAV) command.

These commands are most commonly used to save and restore objects from the "root" file system and the QOpenSys file system.

Note:

- If related objects exist, such as journals and journaled objects, you must verify that you restore them in the correct sequence.
- If you are restoring to a different system or different logical partitions, specify ALWOBJDIF(*ALL) when you are restoring directories.

You can use the RST command to restore the following items:

- A specific object
- A directory or subdirectory
- An entire file system
- Objects that meet search criteria
- A list of object path names

You can also restore the items in the preceding list by using the Restore Object (QsrRestore) API.

For example, to restore all objects (or changed objects) in directories, use the following command:

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

Note: This example is the same restore (RST) command that is issued under option 21 on the **Restore** menu.

When you restore individual objects from a tape file that contains many saved objects, you can improve the restore performance by specifying the following information -

1. Sequence number for a tape media file.

   The sequence number is the quickest way to locate the tape file that contains the correct objects.

2. Starting position.

   The starting position within the tape file for each object is provided by save operations when output to a stream file or user space is requested.

This example restores the object FXX from tape volume LN4381 on device NET3590 and specifies sequence number 4488 and starting position '00000000000190490000000000019CAA'.

```
RST DEV('/QSYS.LIB/NET3590.DEVD')
    OBJ(('/JS200000/FXX')) VOL(LN4381) ENDOPT(*LEAVE)
    SEQNBR(4488) POSITION(00000000000190490000000000019CAA)
```

You can rename an object or restore it to a different directory by using the new-name element of the object (OBJ) parameter.

The OBJ parameter on the RST command supports the use of wildcard characters and the directory hierarchy. To restore a subset of similar objects in a directory subtree, specify the Name pattern (PATTERN) parameter. The PATTERN parameter refines the list of objects that are selected with the OBJ parameter.

Some file systems allow the same physical object to have multiple names by using hard links and symbolic links. For examples of objects with multiple link names and how those objects are saved, see the Backing up your system topic.

In the example of Figure 25 on page 266, FILEA in the JCHDIR directory and FILEB in the DRHDIR directory are both hard links to the same file. They point to the same object. They can have the same name or different names for the objects.

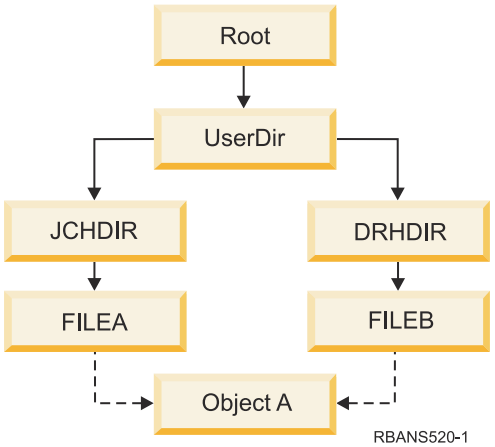

Figure 25. An object with hard links–Example

Table 58 on page 267 shows several examples of how the system restores these objects. These examples assume that you use this SAV command: SAV `OBJ('/UserDir/*')`. The media volume contains OBJECT A and both hard links that point to the object.

Table 58. Restoring objects that have hard links

Object parameter on RST command	Objects that are on the system before the RST command	Objects after the RST command
`RST OBJ('/UserDir/*')`	JCHDIR/FILEA	The object with links FILEA and FILEB from the media are restored over the existing object. Links FILEA and FILEB both point to the restored object.
`RST OBJ('/UserDir/DRHDIR/*')`	JCHDIR/FILEA	The object with link FILEB from the media is restored. Link FILEB points to the restored object. FILEA is unchanged; it points to the same object it pointed to before the restore.
`OBJ('/UserDir/*')`, or `OBJ('/UserDir/JCHDIR/*')`, or `OBJ('/UserDir/DRHDIR/*')`	JCHDIR/FILEA, DRHDIR/FILEB	The object with links FILEA and FILEB from the media are restored over the existing object. Links FILEA and FILEB both point to the restored object.
RST	CRTPRNDIR	If the parent directories of restored objects do not exist, the CRTPRNDIR parameter specifies whether to create them.
RST	PRNDIROWN	The PRNDIROWN parameter specifies the name of an existing user profile to own the parent directories created by the restore operation.

Figure 26 on page 267 shows that the symbolic link called `customer` points to file CUSTMAS in the CUSTLIB library.

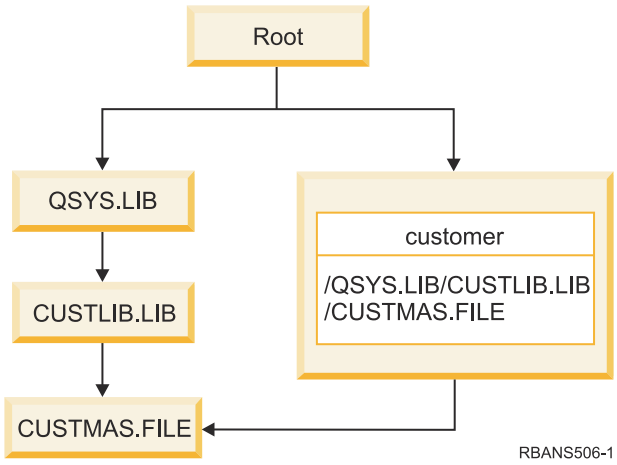

Figure 26. An object with a symbolic link–Example

If you restore the customer object (`RST OBJ('/customer')`), you are restoring only the fact that it points to the CUSTMAS file, not the CUSTMAS file itself. If the CUSTMAS file does not exist, the restore operation succeeds. However, if you try to use the customer object, you receive an error message. If you restore the CUSTMAS file or re-create it, the relationship between the symbolic link customer and the CUSTMAS file is reestablished.

Related concepts

Sequence for restoring related objects

Some objects depend on other objects. For example, SQL materialized query tables (MQTs) are physical database files, but like logical files and SQL views, have dependencies on other files. When related objects are in the same library or directory, the system restores them in the correct order. If related objects are in different libraries or directories, you must restore them in the correct order or perform additional recovery steps after they are restored.

Related information

Object naming rules

Chapter 10. Restrictions when using the Restore command

The Restore (RST) command can be used to restore objects to any file system. This topic describes restrictions that apply when using the RST command.

Restrictions when restoring objects to multiple file systems: When you use the RST command to restore objects to more than one file system at the same time and the file systems include the QSYS.LIB file system or the QDLS file system, the following restrictions apply:

- Different file systems support different types of objects and different methods of naming objects. Therefore, when you restore objects from more than one file system with the same command, you cannot specify object names or object types. You can restore all objects from all file systems, or you can omit some file systems. These combinations are valid:
  - Restoring all objects on the system: `OBJ('/*')`

    Note: Using this command is not the same as using option 21 from the Restore menu. The differences between RST OBJ('/*') and option 21 are as follows:

    - RST OBJ('/*') does not put the system in a restricted state.
    - RST OBJ('/*') does not start the controlling subsystem when it finishes.
    - RST OBJ('/*') does not provide prompting to change default options.
  - Restoring all objects in all file systems except the QSYS.LIB file system and the QDLS file system: `OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))`
  - Restoring all objects in all file systems except the QSYS.LIB file system, the QDLS file system, and one or more other file systems: `OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT) ('/`*other values*`' *OMIT))`
- Values for other parameters of the RST command are supported only for some file systems. You must choose values that are supported by all file systems. Specify the following parameters and values:

  OPTION
  *ALL

  ALWOBJDIF
  *NONE or *ALL

  LABEL
  *SEARCH

  OUTPUT
  *NONE

  SUBTREE
  *ALL

  SYSTEM
  *LCL

  DEV
  (Must be a tape device or an optical device)

  VOL
  *MOUNTED
- When you specify RST OBJ('/*'), the following situations apply:
  - The system restores only objects that are saved by SAV OBJ('/*').
  - The system must be in a restricted state.
  - You must have *SAVSYS or *ALLOBJ special authority.

© Copyright IBM Corp. 1997, 2025

– You cannot specify the save file for the DEV parameter.
– You must specify SEQNBR(*SEARCH).

Note: RST OBJ('/*') is not the suggested method for restoring the entire system. Chapter 3, “Selecting the appropriate recovery strategy,” on page 53 describes how to determine the recovery procedure for your situation.

Restrictions when restoring objects to the QSYS.LIB file system: When you use the RST command to restore objects to the QSYS.LIB (library) file system, the following restrictions apply:

- The OBJ parameter must have only one name.
- You specify objects in the same way that you specify them on the RSTOBJ command and the RSTLIB command. Table 59 on page 270 shows the valid options for the Object (OBJ) parameter when restoring objects to the QSYS.LIB file system and the equivalent Restore Object (RSTOBJ) or Restore Library (RSTLIB) command:

Table 59. Using the RST Command for QSYS.LIB Objects

Object parameter on RST command	Equivalent RSTxxx command
OBJ('/QSYS.LIB/*library-name*.LIB')	RSTLIB SAVLIB(*library-name*)
OBJ('/QSYS.LIB/*library-name*.LIB/*')	RSTOBJ SAVLIB(*library-name*) OBJ(*ALL) OBJTYPE(*ALL)
OBJ('/QSYS.LIB/*library-name*.LIB/*.*object-type*')	RSTOBJ SAVLIB(*library-name*) OBJ(*ALL) OBJTYPE(*object-type*) FILEMBR(*ALL)
OBJ('/QSYS.LIB/*library-name*.LIB/*object-name*.*object-type*')	RSTOBJ SAVLIB(*library-name*) OBJ(*object-name*) OBJTYPE(*object-type*)
OBJ('/QSYS.LIB/*library-name*.LIB/*file-name*.FILE/*')	RSTOBJ SAVLIB(*library-name*) OBJ(*file-name*) OBJTYPE(*FILE)
OBJ('/QSYS.LIB/*library-name*.LIB/*file-name*.FILE/*.MBR')	RSTOBJ SAVLIB(*library-name*) OBJ(*file-name*) OBJTYPE(*FILE)
OBJ('/QSYS.LIB/*library-name*.LIB/*file-name*.FILE/*member-name*.MBR')	RSTOBJ SAVLIB(*library-name*) OBJ(*file-name*) OBJTYPE(*FILE) FILEMBR((*ALL) (*member-name*))

- You can specify only object types that are allowed on the RSTOBJ command. For example, you cannot use the RST command to restore user profiles because OBJTYPE(*USRPRF) is not allowed on the RSTOBJ command.
- Some libraries in the QSYS.LIB file system cannot be saved with the SAVLIB command or restored with the RSTLIB command. You cannot use the RST command to restore these entire libraries. The QDOC, QSYS, and QTEMP libraries are examples.
- You can use the new-name element of the object parameter to rename an object in a directory, restore an object to a different directory, or restore an object to a different library. Table 60 on page 271 shows some examples:

Table 60. New name options on the RST command–Examples

Object parameter on RST command	Results
`OBJ(('/DBSDIR/FILEB' *INCLUDE '/DBSDIR/FILEX'))`	FILEX is created in the DBSDIR directory. The data that was saved with FILEB is restored to FILEX. If FILEB still exists on the system, it is not changed.
`OBJ(('/DBSDIR/FILE*' *INCLUDE '/LMSDIR))`	Restores all objects from the DBSDIR whose names begin with FILE to the LMSDIR directory.
`OBJ(('/QSYS.LIB/LIB1.LIB' *INCLUDE '/QSYS.LIB/LIB2.LIB'))`	Library LIB1 (and all objects) are restored as library LIB2.
`OBJ(('/QSYS.LIB/LIB1.LIB/*' *INCLUDE '/QSYS.LIB/LIB2.LIB'))`	All the objects of library LIB1 are restored into library LIB2.
`OBJ(('QSYS.LIB/LIB1.LIB/*.type' *INCLUDE '/QSYS.LIB/LIB2.LIB/*.type'))`	All objects of type 'type' from library LIB1 are restored into library LIB2.

- For database file members, OPTION(*NEW) restores members for new files only.
- Other parameters must have these values:

  SUBTREE
  *ALL

  SYSTEM
  *LCL

  OUTPUT
  *NONE

  ALWOBJDIF
  *ALL or *NONE
- You can only rename the library, you cannot rename the object. The new name must be *SAME or

  ```
  /QSYS.LIB/libname.LIB
  ```

  where the library that is specified by *libname* must exist.

Restrictions when restoring objects to the QDLS file system: When you use the RST command to restore objects to the QDLS (document library services) file system, the following restrictions apply:

- The OBJ parameter must have only one name.
- The OBJ and SUBTREE parameters must be one of the following items:
  - `OBJ('/QDLS/path/folder-name') SUBTREE(*ALL)`
  - `OBJ('/QDLS/path/document-name') SUBTREE(*OBJ)`
- Other parameters must have these values:

  SYSTEM
  *LCL

  OUTPUT
  *NONE

  ALWOBJDIF
  *ALL or *NONE

  OPTION
  *ALL

Chapter 11. Restoring program temporary fixes (PTFs)

If you have restored the Licensed Internal Code or the operating system, you need to ensure that the PTFs on your system are current.

1. Print a list of all the PTFs currently on the system.

   Type the following command and press the Enter key:

   ```
   DSPPTF LICPGM(*ALL) OUTPUT(*PRINT)
   ```

2. Compare this list of PTFs with the list you printed when you saved the system.

   If the lists are the same, return to your recovery checklist. If PTFs are missing from the list you printed in step 1, you must apply them. Continue with the next step.

3. Find the most recent cumulative program temporary fix media.

   This package might be on distribution media or on a stand-alone media volume.

   Note: If you do not have the PTFs you need, order them and apply them later. Continue with your recovery checklist.

4. You can use option 8 (Install program temporary fix package) on the Program Temporary Fix menu.

   All of the PTFs in the cumulative PTF package will be installed for the licensed programs you have installed on your system. Refer to the *IBM i System PTF Shipping Information Letter* for special instructions that are required.

© Copyright IBM Corp. 1997, 2025

Chapter 12. Restoring system information

The Restore System Information (RSTSYSINF) command restores a subset of the system data and objects saved by the Save System Information (SAVSYSINF) command.

The RSTSYSINF command is not to be used for system upgrades or migrations.

The RSTSYSINF command does not restore the following items:

- Security-related system values are not restored if they have been locked.
- Objects that are in use when the RSTSYSINF command is running will be restored. This command can be run while the system is either restricted or not restricted. There will be less objects in use if the command is run while the system is restricted.

  Note: You will need to determine if the objects on the system are usable or if it is necessary to perform a recovery on the objects.
- The QPWDLVL (Password level) system value will not be restored. See the Planning password level changes topic before changing the QPWDLVL system value. Also, based on the current setting of the QPWDLVL system value, QPWDMAXLEN (Maximum password length), QPWDMINLEN (Minimum password length), and QPWDVLDPGM (Password validation program) system values might not be restored.

These are examples to restore your system using the RSTSYSINF command:

Example 1: System information restored from TAP01

```
RSTSYSINF DEV(TAP01)
```

This command restores the system information from the tape put on the TAP01 tape drive.

Example 2: System information restored from SAVF and Printed Output Generated

```
RSTSYSINF DEV(*SAVF) SAVF(QGPL/SAVF) OUTPUT(*PRINT)
```

This command restores the system information from the save file named SAVF in library QGPL. The information that was restored will be written to a spooled file.

Related information

Locking and unlocking security-related system values

Save System Information (SAVSYSINF) command

© Copyright IBM Corp. 1997, 2025

Chapter 13. Restoring changed objects and apply journaled changes

The procedures in this topic are designed to recover activity that has occurred since the last complete save operation.

Figure 27 on page 277 shows a typical time line for the system.

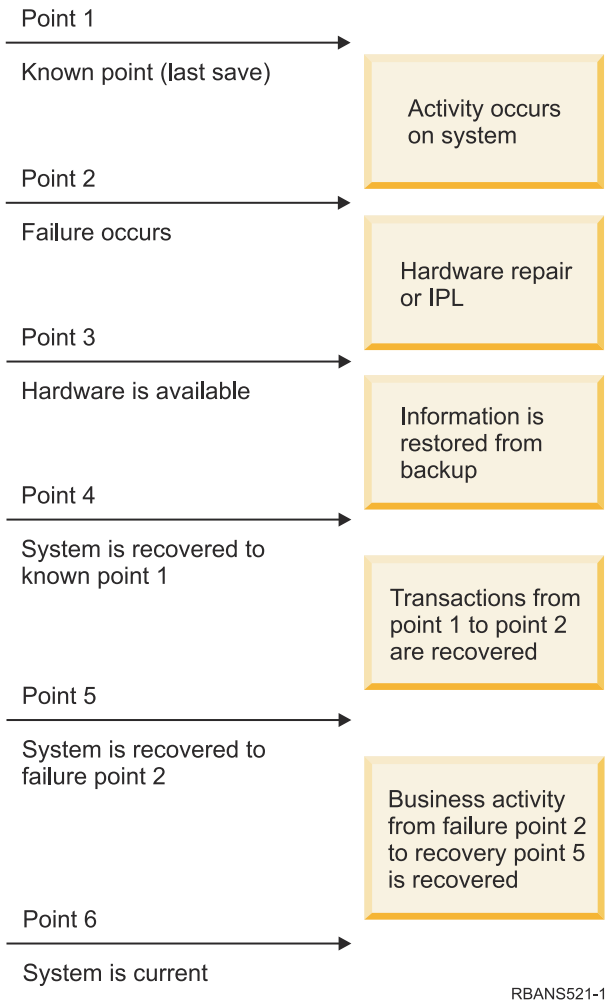

Figure 27. Sample recovery time line

Chapter 4, “Recovering the Licensed Internal Code,” on page 121 through Chapter 9, “Restoring specific types of information,” on page 209 describe what is required to reach point 4 in the time line. This returns your system to the point of the last complete save operation.

This topic collection describes two procedures that are available to reach point 5 in the time line:

- Restoring changed objects.
- Applying journal changes.

Task 1: Restoring changed objects

Decide whether you want to restore changed objects individually, or by library.

The Saving changed objects topic describes different methods of saving changed objects. Table 61 on page 278 shows the two methods and the correct restore procedures for each:

© Copyright IBM Corp. 1997, 2025

Table 61. Restore procedures for changed objects

Save Method	Description	Restore procedure
Note: The Save Changed Objects (SAVCHGOBJ) command does not apply to objects in directories. If you are restoring changed objects in directories, go to "Task 2: Restoring changed objects in directories" on page 279 for restore instructions from both the cumulative and not cumulative save methods.		
Cumulative	You save all changes since the last complete save operation.	"Restoring changed objects by library" on page 278
Not cumulative	You save changes since the last SAVCHGOBJ operation.	"Restoring changed objects individually" on page 278

If you save journal receivers using the SAVCHGOBJ command, read the "Restoring journal receivers" on page 255 topic for special considerations that might apply when restoring them.

Restoring changed objects by library

Use these steps to restore changed objects by library.

1. Load the Save Changed Object (SAVCHGOBJ) media volume.
2. Type `DSPTAP DEV(`*`media-device-name`*`) OUTPUT(*PRINT)` for tape media.

   Type `DSPOPT VOL(*MOUNTED) DEV(OPT01) DATA(*SAVRST) PATH(*ALL) OUTPUT(*PRINT)` for DVD-RAM optical media. Press the Enter key.
3. Restore the data with this command:

```
RSTOBJ OBJ(*ALL) DEV(media-device-name) SAVLIB(*ANY)
OBJTYPE(*ALL) MBROPT(*ALL)
```

4. Choose one of the following options:
   - If you have journaled changes to apply, continue with "Task 4: Determining what journal receivers to use" on page 280.
   - If you do not need to apply journaled changes, skip to "Task 8: Restoring changed documents and folders" on page 285.
   - If you are not sure whether you need to apply journaled changes, continue with "Task 3: Determining whether you need to apply journaled changes" on page 280.

Restoring changed objects individually

If your method for saving changed objects is not cumulative, you might have the same object on more than one set of Save Changed Object (SAVCHGOBJ) save media.

You can choose to restore each set of SAVCHGOBJ save media completely, starting from the oldest media volume. This is the simplest method. However, it might be time-consuming if you have the same large objects on more than one SAVCHGOBJ media volume.

If you want to restore each set of SAVCHGOBJ save media completely, follow the procedure described in "Restoring changed objects by library" on page 278 for each set of save media. If you want to restore each object only once, follow this procedure:

1. Load each SAVCHGOBJ media volume.
2. Type `DSPTAP DEV(`*`media-device-name`*`) OUTPUT(*PRINT)` and press the Enter key.
3. Compare the listings and find the most recent saved copy of each object.
4. For each object, load the correct media volume and type the following command:

```
RSTOBJ OBJ(object-name)DEV(media-device-name)
```

```
       SAVLIB(library-name) OBJTYPE(*ALL)
       ENDOPT(*LEAVE) MBROPT(*ALL)
```

Repeat this step for each object you need to restore.

5. Choose one of the following options:
   - If you have journaled changes to apply, continue with "Task 4: Determining what journal receivers to use" on page 280.
   - If you do not need to apply journaled changes, skip to "Task 8: Restoring changed documents and folders" on page 285.
   - If you are not sure whether you need to apply journaled changes, continue with "Task 3: Determining whether you need to apply journaled changes" on page 280.

Task 2: Restoring changed objects in directories

Perform this task if you saved changed objects in directories. If you do not need to do this task, continue with the next step in your recovery checklist.

If you use a cumulative method when you saved changed objects from directories (your save media contain all objects that have changed since the last complete save operation), follow these steps:

1. Mount your most recent save media from saving changed objects in directories.
2. Type the following command:

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

3. Choose one of the following options:
   - If you have journaled changes to apply, continue with "Task 4: Determining what journal receivers to use" on page 280.
   - If you do not need to apply journaled changes, skip to "Task 8: Restoring changed documents and folders" on page 285.
   - If you are not sure whether you need to apply journaled changes, continue with "Task 3: Determining whether you need to apply journaled changes" on page 280.

If your save media from saving changed objects in directories is not cumulative (your save media contains only objects that have changed since the last save operation), repeat the following steps for each set of save media because your last complete save operation. Start with the oldest save media volumes and end with the most recent volumes.

1. Mount each volume with the cumulative save operations. Start with the oldest save media volumes and end with the most recent volumes.
2. Mount the media volume.
3. Type the following command:

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

4. Choose one of the following options:
   - If you have journaled changes to apply, continue with "Task 4: Determining what journal receivers to use" on page 280.
   - If you do not need to apply journaled changes, skip to "Task 8: Restoring changed documents and folders" on page 285.
   - If you are not sure whether you need to apply journaled changes, continue with "Task 3: Determining whether you need to apply journaled changes" on page 280.

Task 3: Determining whether you need to apply journaled changes

You might have set up journaling yourself, or you might be using applications that use journaling.

For example, IBM i Access uses the QUSRSYS/QAOSDIAJRN journal. Some applications provided by software vendors also use journaling. If you have objects that are restored with partial transactions (message CPI3731), you must apply journaled changes in order to bring these objects to a usable state.

If you are not sure about whether you have journaled changes to apply, follow these steps:

1. Type `DSPOBJD OBJ(*ALL/*ALL) OBJTYPE(*JRN) OUTPUT(*PRINT)` and press the Enter key.
   This command prints a list of all the journals on your system.
2. For each journal on the list, follow these steps:
   a) Type: `WRKJRNA JRN(`*`library-name`*`/`*`journal-name`*`)`.
      You are shown the Work with Journal Attributes display.
   b) Press F19 to display the objects being journaled.
   c) Press F12 to return to the Work with Journal Attributes display.
   d) Press F15 to display the receiver directory.
      Note the attach and detach times for the journal receivers in relation to your journaled objects change dates. Additionally, you can use option 8 to display specifics about each journal receiver.
   e) Press F12 to return to the Work with Journal Attributes display.
   f) From the information you have seen, you should be able to determine whether any objects are being journaled and whether any journal entries exist that are more recent than your most recent saved copies of the objects.
      You can also determine which receivers are on the system for the journal. Repeat these steps for each additional journal.
3. If you need to apply journaled changes, continue with "Task 4: Determining what journal receivers to use" on page 280. If you do not need to apply journaled changes, skip to "Task 8: Restoring changed documents and folders" on page 285.

Task 4: Determining what journal receivers to use

You need all journal receivers that were attached to the journal for the length of time for which journaled changes need to be applied to the restored files.

1. Ensure that all the journal receivers required for the apply journaled changes operation are available on the system.
   If you have restored objects with partial transactions, you might also need earlier receivers. Refer to any CPI3731 messages that you received during the restore to find the earliest receiver that you need. You can also use the Display File Description (DSPFD) command for the files you just restored to determine the earliest receiver that is required.
2. Restore all necessary journal receivers that are not already on the system.
   Use the Display Journal Receiver Attributes (DSPJRNRCVA) command to determine when a journal receiver was attached to and detached from a journal.
3. Determine the name of the last journal receiver (the last receiver restored) and whether there are chain breaks by printing the receiver chain:
   a) Type `WRKJRNA JRN(`*`library-name`*`/`*`journal-name`*`) OUTPUT(*PRINT)` and press the Enter key.
      You receive a listing that shows the receiver directory and all the objects being journaled.
   b) Look at the receiver directory part of the listing.
      If you saved the currently attached journal receiver, your journal receiver directory should look similar to Figure 28 on page 281. The journal receiver that was attached during the save

procedure shows a status of Partial. The following example shows the displayed version of the receiver directory:

```
                              Receiver Directory
Total size of receivers (in kilobytes). . . . . :   1507
                                Attach     Save                    Size
Number   Receiver    Library    Date       Date        Status       (K)
00001    RCVA0001    DSTJRN     06/08/xx   06/08/xx    SAVED         42
00002    RCVA0002    DSTJRN     06/09/xx   06/09/xx    SAVED        900
00003    RCVA0003    DSTJRN     06/09/xx   06/09/xx    PARTIAL       92
01001    RCVA1003    DSTJRN     06/10/xx   00/00/00    ATTACHED     473
```

Figure 28. Receiver directory–Saving attached receivers

If you save only detached journal receivers, your receiver directory should look similar to Figure 29 on page 281:

```
                              Receiver Directory
Total size of receivers (in kilobytes). . . . . :   1507
                                Attach     Save                  Size
Number   Receiver    Library    Date       Date        Status     (K)
00001    RCVA0001    DSTJRN     06/08/xx   06/08/xx    SAVED       42
00002    RCVA0002    DSTJRN     06/09/xx   06/09/xx    SAVED      900
00003    RCVA0003    DSTJRN     06/09/xx   06/09/xx    SAVED       92
01001    RCVA1003    DSTJRN     06/10/xx   00/00/00    ATTACHED   473
```

Figure 29. Receiver directory–Saving detached receivers

4. On the listing, mark the name of the last receiver with a status of SAVED or PARTIAL.
5. Determine the chain of receivers to be used in the Apply Journaled Changes (APYJRNCHG) command from the Work with Receiver Directory listing.

   Mark the first and last receiver that you need, based on the date that you saved the objects being recovered. Notice that the first and last receiver are the same if only one journal receiver was restored.

   Note: While looking at the receiver directory, you should also look for any receiver chain breaks. You can determine a chain break by looking at the first two digits in the Number column on the Work with Receiver Directory display. You cannot apply journaled changes across receiver chain breaks. Therefore, you must write down the beginning and ending receiver names for each receiver chain. Then you need to run a series of apply journaled changes operations, one for each chain using these receivers. A chain break might mean that you are missing all or part of a journal receiver. (It was on the system and was not saved before the failure occurred.) Evaluate how applying journaled changes across a change break might affect the integrity of your data.
6. Look at the part of the listing that shows what objects are currently being journaled. (You printed the listing in step 3a.)

   Compare it to your records of what objects should be journaled. Follow the procedures in the Printing system information topic before you save your system.
7. For each physical file that should be journaled and does not appear on the current listing, type the following command:

   ```
   STRJRNPF FILE(library-name/file-name)
            JRN(library-name/journal-name)
   ```
8. For each access path that should be journaled and does not appear on the current listing, type the following command:

   ```
   STRJRNAP FILE(library-name/file-name)
            JRN(library-name/journal-name)
   ```
9. For each integrated file system object that should be journaled and does not appear on the current listing, type the following command:

```
STRJRN OBJ('object-path-name')
       JRN('journal-path-name')
```

10. For each library that should be journaled and does not appear on the current listing, type the following command:

```
STRJRNLIB LIB(library-name)
          JRN(library-name/journal-name)
```

11. For all other object types that should be journaled and do not appear on the current listing, type the following command:

```
STRJRNOBJ OBJ(library-name/object-name)
          OBJTYPE(object-type)
          JRN(library-name/journal-name)
```

12. The journal receiver that is currently attached might not match your naming conventions.

    Typically this is because the journal receiver was created when you restored the journal. If this is the case, create a new receiver that follows the same naming convention and receiver attributes as the last receiver but assign it a number of one greater. In the example shown on the Work with Receiver Directory display, you need to type the following command:

```
CRTJRNRCV JRNRCV(DSTJRN/RCVA0004)
```

13. Use the Change Journal (CHGJRN) command to detach the current receiver and attach the journal receiver you just created.

    In the example, you need to type the following command:

```
CHGJRN JRN($JRNLA/JRNA)
       JRNRCV(DSTJRN/RCVA0004)
```

Related information

Journal management

Task 5: Applying journaled changes for user journals

If you need to apply journaled changes from user journals, perform this task.

If you do not need to apply journaled changes, skip to "Task 6: Applying journaled changes for the QAOSDIAJRN journal" on page 284.

1. If you have a single receiver chain for the journal entries that you need to apply and the status of the last receiver you are using is SAVED, perform one of the following tasks:

   a) To apply journaled changes to objects in a library, type the following command:

```
APYJRNCHG JRN(library-name/journal-name)
          OBJ((library-name/*ALL *ALL))
          RCVRNG(*LASTSAVE)
          FROMENT(*LASTSAVE) TOENT(*LAST)
```

   If journal entries exist for the creation of a file, data area, or data queue, those objects are re-created. Use the APYLF parameter to specify whether to apply changes to the logical files that are built over the changed physical files.

   To apply entries to *all* objects in all libraries that are journaled to the journal, type the following command:

```
APYJRNCHG JRN(library-name/journal-name)
          OBJ((*ALLJRNOBJ))
```

```
          RCVRNG(*LASTSAVE)
          FROMENT(*LASTSAVE) TOENT(*LAST)
```

b) For objects in directories, type the following command:

```
APYJRNCHG JRN(jrnlib/jrnname)
          OBJPATH('object-path-name')
          RCVRNG(*LASTSAVE)
          FROMENT(*LASTSAVE) TOENT(*LAST)
```

Note: If you want to apply journaled changes to library and directory objects using the same command, you can use both the OBJ and OBJPATH parameters in one Apply Journaled Changes (APYJRNCHG) command invocation.

c) If you have restored objects with partial transactions, but you do not have any journal receivers more recent than the receiver that contains the save entries, you have the option to remove journal changes to resolve the partial transaction.

The following command removes the changes in journal JRN1 from all the members of OBJ1:

```
RMVJRNCHG JRN(JRN1) FILE(LIB1/OBJ1)
          FROMENT(*LASTSAVE) TOENT(*COMMITSTART)
          RCVRNG(*LASTSAVE)
```

Starting with the last save journal entry, only changes for journal entries for any partial transactions are removed, back to the start of the commit transaction.

d) If you are unable to complete the previous steps for objects with partial transactions because the journal receivers are unavailable, you can use the Change Journaled Object (CHGJRNOBJ) command to get the object to a usable state.

The Partial Transactions (PTLTNS) parameter allows the object to be used, but does not complete the transactions. In the example below, the object, BRKNOBJ, still has changes caused by the partial transactions, but you are able to open the file.

Attention: Only use the following command as a last resort. You will lose data if you use this command. You should only use this command for the following reasons:

- You have objects with partial transactions as a result of the terminations of a long-running rollback and you have no saved version to restore.
- You have objects with partial transactions as a result of a save-while-active operation, and the journal receivers required to apply or remove journaled changes have been lost, destroyed, or damaged beyond repair.

```
CHGJRNOBJ OBJECT(LIB1/BRKNOBJ *FILE) ATR(*PTLTNS) PTLTNS(*ALWUSE)
```

2. If you have determined that this journal had receiver chain breaks, you must decide whether you are actually missing journal receivers and necessary journal entries or whether the chain breaks were caused by something else.

   You should evaluate how applying journaled changes across a chain break might affect the integrity of your data.

   If you decide to apply journal entries across chain breaks, you must use a APYJRNCHG command for each chain. Type the APYJRNCHG command and use these values in place of the values shown in step 1.

   For the first (earliest) receiver chain, the following parameters apply:

   RCVRNG
   First and last receivers in this chain

   FROMENT
   *LASTSAVE

   TOENT
   *LAST

For each middle receiver chain, the following parameters apply:

RCVRNG
First and last receivers in this chain

FROMENT
*FIRST

TOENT
*LAST

For the last receiver chain, the following parameters apply:

RCVRNG
First and last receivers in this chain

FROMENT
*FIRST

TOENT
*LAST

Task 6: Applying journaled changes for the QAOSDIAJRN journal

If you have document library objects, you might need to apply journaled changes from the receivers associated with the QAOSDIAJRN journal.

If you are not sure, determine when you last saved the QUSRSYS library. Then perform the steps through step 1c to determine whether you have any journal entries for the QAOSDIAJRN journal that are more recent than your save media for the QUSRSYS library.

You cannot apply all journaled changes in the QAOSDIAJRN journal in library QUSRSYS. You must specify individual files on the FILE parameter instead of *ALL. Do not apply journal changes to the document and folder search index database files (QAOSSS10 through QAOSSS15, QAOSSS17, and QAOSSS18) for journal QAOSDIAJRN in library QUSRSYS.

1. Display the receiver chain for the QAOSDIAJRN journal by doing the following steps:
   a) Type: `WRKJRNA JRN(QUSRSYS/QAOSDIAJRN)` and press the Enter key.
   b) From the Work with Journal Attributes display, press F15 (Work with receiver directory). Examine the receiver directory to determine whether any chain breaks exist. (See step 5 in the Task 4–Determining what journal receivers to use topic.)
   c) Press F12 two times to return to a command line.
2. If no chain breaks exist, type the following command to apply journaled changes for the QAOSDIAJRN journal to individual files:

```
APYJRNCHG JRN(QUSRSYS/QAOSDIAJRN)
          FILE((QUSRSYS/QAOKPLCA) (QUSRSYS/QAOSAY05)
               (QUSRSYS/QAOKPX4A) (QUSRSYS/QAOSAY07)
               (QUSRSYS/QAOKP01A) (QUSRSYS/QAOKP02A)
               (QUSRSYS/QAOKP03A) (QUSRSYS/QAOKP04A)
               (QUSRSYS/QAOKP05A) (QUSRSYS/QAOKP06A)
               (QUSRSYS/QAOKP08A) (QUSRSYS/QAOKP09A))
          RCVRNG(lib-name/first-receiver
                 lib-name/last-receiver)
          FROMENT(*LASTSAVE)
          TOENT(*LAST)
```

3. If chain breaks exist, you must determine whether any journal receivers are missing and how that might affect the integrity of your recovery.

   If you decide to apply journaled changes, use the command shown in step 2. Repeat the command for each receiver chain, substituting the correct receiver range, from-entry, and to-entry parameters. Step 2 in the Task 5–Applying journaled changes for user journals topic describes how to use these parameters.

Task 7: Replaying journaled changes in a disaster recovery

If an object becomes damaged or is not usable, you can recover the object using the Apply Journaled Changes (APYJRNCHG) command. You also can replay changes to libraries by using the APYJRNCHG command in a disaster recovery situation.

If you restore an object that was saved with partial transactions, then you must apply journaled changes to that object before it is usable. You can apply journaled changes to all objects that are journaled to the journal by specifying OBJ(*ALLJRNOBJ) on the APYJRNCHG command.

To replay journaled changes in a disaster recovery situation, perform the following steps:

1. Use the Display Object Description (DSPOBJD) command to determine the earliest journal receiver that needs to be on the system to recover the objects.
2. After the journal, journal receivers, and journaled objects have all been restored, issue the APYJRNCHG command to apply the journaled changes to all the objects, as the following example shows:

```
APYJRNCHG JRN(JRN2) OBJ(*ALLJRNOBJ)
FROMENTLRG(*LASTSAVE) TOENTLRG(*LAST)
RCVRNG(*LASTSAVE)
```

Related information

Applying journaled changes

Apply Journaled Changes (APYJRNCHG) command

Task 8: Restoring changed documents and folders

If you save changed documents and folders, perform this task.

If you do not need to do this task, skip to "Task 2: Restoring changed objects in directories" on page 279.

1. If your procedure for saving changed document library objects (DLOs) is cumulative, load the last daily Save Document Library Object (SAVDLO) media volume.

   If your procedure is not cumulative, start with your earliest daily save volume and repeat these steps for each set of SAVDLO save media.
2. If you have documents in user auxiliary storage pools (ASPs), display the save media volumes to find the sequence numbers for each ASP.

   Type `DSPTAP DEV(media-device-name) OUTPUT(*PRINT)` for tapes. Mark the names and sequence numbers of the files on the listing. They will be named QDOC for the system ASP and QDOC*nnnn* for each user ASP that contains DLOs, where *nnnn* is the number of the ASP.
3. To restore the DLOs to a single ASP, type the following command:

```
RSTDLO DLO(*ALL) DEV(media-device-name) SAVFLR(*ANY)
       SAVASP(ASP-number) RSTASP(*SAVASP)
```

4. To restore the DLOs to all ASPs, type the following command:

```
RSTDLO DLO(*ALL) DEV(media-device-name) SAVFLR(*ANY)
       SAVASP(*ANY) RSTASP(*SAVASP)
```

5. If an unrecoverable error occurs when running the RSTDLO DLO(*ALL) SAVFLR(*ANY) command, see the "Recovering from an error while restoring document library objects" on page 42 topic.

Chapter 14. Mirrored protection recovery actions

On a system with mirrored protection, errors and failures have different effects. When a failure occurs on a system with mirrored protection, the recovery procedure is affected by the level of protection that is configured.

In considering aspects of recovery, you need to distinguish between *errors* and *failures* in the disk subsystem.

A disk *error* refers to an unexpected event during an input/output (I/O) operation which can cause the loss or corruption of the data that is being transferred. Most disk errors are caused by a failure in some part of the component chain from the I/O processor to the disk surface. Environmental effects such as power abnormalities or severe electrostatic discharges can also cause disk errors. Included in the definition of disk errors is a failure of the Licensed Internal Code that controls the disk subsystem.

When the system detects an error, generally the occurrence is logged and the operation is attempted again. Temporary errors are those from which the system can recover and complete the I/O operation successfully. When the error is so severe that the I/O operation cannot succeed, it is a permanent error.

When the system detects a permanent error, it classifies it as a *failure* in that hardware subsystem. In an ASP that does not have mirrored protection, a failure causes the system to become unusable. The system displays an error message which contains a System Reference Code (SRC) of A6*xx* 0244, A6*xx* 0255, or A6*xx* 0266 where *xx* is incremented every minute. During this time, the system tries the operation again in which the failure occurred. If the condition that caused the failure can be corrected (for example, by powering on a disk unit or replacing an electronic component), then normal system operations are resumed.

Related tasks

Checklist 14: Actions for non-load source disk unit failure
This checklist shows the sequence of steps you should use to recover after a non-load-source disk unit with mirrored protection has failed.

System actions for permanent errors

When a permanent error occurs and mirroring is active, the system attempts to recover.

Device error

If the system detects a device, input/output (I/O) processor, or bus failure on a mirrored unit, it does the following actions:

1. The system disables the failing unit and suspends mirroring for the pair. If the other unit in the pair is also failing or is already suspended, the first unit is considered unprotected.
2. The system sends a message that identifies the failing unit and indicates that mirroring has been suspended. You can use problem analysis on this message for more information.
3. When a disk unit is suspended following an error, the system records all updates that are performed on the active unit of the mirrored pair. If the suspended disk unit becomes usable within a short period of time, the system automatically synchronizes the data between the mirrored units.
4. After the failed unit is replaced, the system synchronizes the pair and resumes mirrored protection. The system sends a message that indicates that mirrored protection has been resumed.

Read error

1. The system reads from the other storage unit of the mirrored pair. If the permanent read error occurs on the other storage unit as well, the original read request completes with a permanent read error.

© Copyright IBM Corp. 1997, 2025

2. If the read operation from the other storage unit is successful, the data is written back to the first unit of the mirrored pair, assigning an alternate sector. Only then does the system signal that the original read request is complete.

Connection failure

If the system cannot communicate with a device, it does the following actions:

1. The system attempts to recover from the communications error. Any job requesting the disk unit waits during the period that the system is attempting recovery.
2. If the recovery is successful, normal system operations continues.
3. If the system cannot recover within the time limit for the reset command, the unit is considered to have a device error. The system performs the steps that are described in Device error section.

Load source unit failure

If an error occurs on the load source unit before the Storage Management Recovery portion of the initial program load (IPL), the system does the following actions:

1. The system determines if the other mirrored unit in the load source mirrored pair is usable. If not, the system fails.
2. If the system is able to continue, it starts an IPL from the remaining usable unit in the load source mirrored pair.

Suspending mirrored units

If you have to suspend a mirrored unit, you can do so using the Suspend Mirrored Protection option on the Work With Disk Unit Recovery display using system service tools (SST) or dedicated service tools (DST).

To suspend mirrored protection, follow these steps:

1. Type:STRSST
2. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 3 (Work with disk unit recovery) on the Work with Disk Units display.
3. Select option 3 (Suspend mirrored protection) on the Work with Disk Unit Recovery display and press the Enter key.

```
                    Suspend Mirrored Protection

Type option, press Enter.
  1=Suspend Mirrored Protection

                  Serial                        Resource
OPT   Unit   ASP  Number       Type   Model     Name      Status
 _       1    1   00-31297     6109    030      DD002     Resuming
 _       3    1   00-0184097   6602    050      DD011     Active
 _       3    1   00-0125986   6602    050      DD005     Active
```

4. Type a 1 (Suspend Mirrored Protection) in the Option column for each unit that you want to suspend mirrored protection.

   You can suspend protection only on units that have both units in an Active or Resuming status. If one of the units is in Resuming status, then it is the only unit that can be selected to suspend. It takes several minutes to suspend a resuming unit that is using SST.

If you suspend a mirrored unit that is using SST, the system starts to keep a list of disk pages that are changed. If you resume mirrored protection on the suspended mirrored unit before this list becomes full, the system uses this list to copy data from only those disk pages that were changed instead of copying an entire disk.

Resuming mirrored units

If you have to resume a mirrored unit, you can do so using the Resume Mirrored Protection option on the Work With Disk Unit Recovery display using system service tools (SST) or dedicated service tools (DST).

To resume mirrored protection, follow these steps:

1. Type: STRSST
2. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 3 (Work with disk unit recovery) on the Work with Disk Units display.
3. Select option 4 (Resume mirrored protection) on the Work with Disk Unit Recovery display and press the Enter key.

```
                    Resume Mirrored Protection
   Type option, press Enter.
1=Resume Mirrored Protection
                          Serial
Resource OPT  Unit  ASP   Number      Type  Model  Name      Status
       _        2    3    00-59681F7  6602   050   DD004     Suspended
```

4. Type a 1 (Resume Protection) in the Option column for each unit that you want to resume mirrored protection. You can select only a unit in Suspended status to resume.

Replacing a mirrored unit

A unit selected to replace the failed mirrored unit must satisfy all of the mirrored protection configuration rules and restrictions when it is paired with the remaining unit in the mirrored pair.

You can replace mirrored units by using the Replace Disk Unit option in either dedicated service tools (DST) or system service tools (SST). To do this, you need to have a spare storage unit that can be paired with the mirrored unit of the storage unit that is being replaced. The storage unit being replaced can have a status of active or suspended. However, one of the storage units in the pair must be suspended. The results of the replace operation is different for each status. Replacing a suspended storage unit causes that storage unit to go to a status of resuming after the replace operation. Replacing an active unit causes data in the auxiliary storage pool (ASP) to be lost, so you must first delete the data in ASP (using the DST Delete ASP Data option). The storage unit being replaced can also be missing or not missing. To replace a unit with a status of resuming, you must suspend it. If the status of unit 1 is unknown, replace operations are not allowed until the status of the mirrored units for unit 1 is known. The unit selected to replace another mirrored unit must satisfy all of the mirrored protection configuration rules and restrictions when it is paired with the remaining unit in the mirrored pair.

If a storage unit fails, and if the same storage unit that failed has been repaired, it is not necessary to replace it. The failed disk should have a status of suspended and can be resumed after the repair is complete.

If the storage unit that is being replaced is active, it can only be replaced at DST before the initial program load (IPL) to the IBM i licensed program. It should never be necessary to replace an active unit unless both units of the mirrored pair have failed. If this situation does occur, the service representative should first try to recover the data from the failed units using the Save Disk Unit Data option on the Work with Disk Unit Recovery display. When an active unit is replaced, the last good copy of the data has been lost. The data in the ASP that contains the unit being replaced must be deleted using the DST Delete ASP Data option before an active unit is allowed to be replaced.

Replacing unit 1 requires special handling. If the system ASP has mirrored protection, one of the units in the mirrored pair for unit 1 is selected as the IPL device. It is the only unit that is used until the system performs an IPL to the IBM i licensed program. Until then, it cannot be replaced or even suspended. However, its mirrored unit can be both suspended and replaced. After the IPL to the IBM i licensed program, the IPL device can be suspended and then replaced.

Replacing a unit might cause the level of protection for a mirrored pair to change. If a lower level of protection results from a replacement operation, it shows a warning display. At certain times, especially when missing units are involved in the replace operation, the system might not be able to accurately calculate the level of protection and the same warning display is shown.

To replace a disk unit by using SST, follow these steps:

1. Type: STRSST
2. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 3 (Work with disk unit recovery) on the Work with Disk Unit display.
3. Select option 1 (Replace configured unit) on the Work with Disk Unit Recovery display and press the Enter key.

   The Select Configured Unit to Replace display is shown.

```
                     Select Configured Unit to Replace

   Type option, press Enter.
1=Select
                         Serial
Resource OPT  Unit  ASP  Number      Type  Model  Name       Status
  _            1     1   00-0163477  6602   030   DD019      Suspended
  1            2     1   00-17900    6602   030   DD002      Suspended
```

4. Type a 1 in the Option column on the Select Configured Unit to Replace display and press the Enter key.

```
                       Select Replacement Unit
                    Serial
Resource Unit  ASP  Number      Type  Model  Name       Status
    2     1        00-17900    6602   030   DD002      Suspended

Type option, press Enter.   1=Select

                Serial
Resource Option Number      Type  Model  Name       Status
   _            00-0330477  6602   030   DD005      Non-configured
   1            00-0323200  6602   030   DD033      Non-configured
```

5. Type a 1 in the Option column on the Select Replacement Unit display and press the Enter key.

```
                       Confirm Replace of Configured Unit

   This screen allows the confirmation of the configured unit to be replaced
with the selected replacement unit.  Press Enter to confirm your choices for
Replace Press F12 to return to change your choices.   The configured unit
being replaced is:
                   Serial
Resource Unit ASP  Number      Type  Model  Name      Status
           2    1  00-17900    6602   030   DD002     Suspended

The replacement unit will be:
                   Serial
Resource Unit ASP  Number      Type  Model  Name      Status
           2    1  00-0323200  6602   030   DD033     Resuming
```

6. Press Enter to confirm.
7. The replacement function runs for several minutes. Wait until the replacement function completes.

Related concepts

Mirrored protection: Configuration rules

Here are the configuration rules for mirrored protection.

Using spare nonconfigured units for replacement

If mirrored units become suspended as a result of a hardware failure, the system continues to run. However, one or more storage units will be suspended and therefore unprotected until your service representative can repair or replace the failed hardware.

If you have spare nonconfigured units, you might be able to resume mirrored protection before the repair actions are done. Call your service representative. You might be directed to examine the Service Action Log for information that is concerning the failure. Use the Display Disk Configuration Status option by using system service tools (SST) or the Work with Disk Status (WRKDSKSTS) command to determine which units are suspended. If all disk units under an input/output (I/O) processor are suspended, the I/O processor probably has failed. If you have enough spare units of the right type and model, and if the spare units are not on the I/O processor that has failed, you might be able to use the spare nonconfigured units to resume mirrored protection.

After your service representative repairs a failed storage unit, you might want to use it instead of the spare to restore the previous level of protection. To use the repaired unit, follow these steps:

1. Suspend the active storage unit that was previously used as the spare by typing the following on a command line and pressing the Enter key: STRSST
2. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 3 (Work with dik unit recovery) on the Work with Disk Units display.
3. Select option 3 (Suspend mirrored protection).

```
                     Suspend Mirrored Protection

   Type option, press Enter.
   1=Suspend Mirrored Protection
                           Serial
Resource OPT  Unit  ASP   Number      Type  Model  Name     Status
  _              1   1    00-0193825  6602   030   DD001    Active
  _              1   1    00-0184097  6602   030   DD019    Active
  _              2   1    00-0125986  6602   030   DD036    Active
  _              2   1    00-0125986  6602   030   DD002    Active
```

4. Type a 1 (Suspend Mirrored Protection) in the Option column.
   The original spare unit is the same disk type and model as the repaired disk unit.
5. Return to the Work with Disk Unit Recovery display by pressing F12 (Cancel).

```
               Work with Disk Unit Recovery

   Select one of the following:
      1. Replace configured unit
      2. Disk unit problem recovery procedures
      3. Suspend mirrored protection
      4. Resume mirrored protection
      5. Delete disk unit data
      6. Rebuild disk unit data
```

6. Select option 1 (Replace configured unit).

```
                     Select Configured Unit to Replace
   Type option, press Enter.
   1=Select
                        Serial
Resource OPT  Unit  ASP  Number      Type  Model  Name       Status
  _              1    1  00-0163477  6602   030   DD019      Suspended
  1              2    1  00-17900    6602   030   DD002      Suspended
```

7. Type a 1 in the Option column on the Select Configured Unit to Replace display and press the Enter key.

```
                    Select Replacement Unit
                  Serial
Resource Unit  ASP  Number      Type  Model  Name       Status
    2            1  00-17900    6602   030   DD002      Suspended

 Type option, press Enter.   1=Select
              Serial
Resource Option Number      Type  Model  Name       Status
   _            00-0330477  6602   030   DD005      Non-configured
   1            00-0323200  6602   030   DD033      Non-configured
```

8. Type a 1 in the Option column on the Select Replacement Unit display and press the Enter key.

```
                       Unit
   This screen allows the confirmation of the configured unit to be replaced
with the selected replacement unit.  Press Enter to confirm your choices for
Replace Press F12 to return to change your choices.   The configured unit
being replaced is:
                  Serial
Resource Unit ASP  Number      Type  Model  Name      Status
            2   1  00-17900    6602   030   DD002     Suspended

   The replacement unit will be:
                  Serial
Resource Unit ASP  Number      Type  Model  Name      Status
            2   1  00-0323200  6602   030   DD033     Resuming
```

9. Press Enter to confirm.
10. The replacement function runs for several minutes. Wait until the replacement function completes.

Mirrored protection recovery actions performed by the service representative

Here are the steps and considerations involved in disk unit repair in the mirrored environment, which are performed by your service representative.

Actions when concurrent maintenance is possible

This topic provides actions to take when concurrent maintenance is possible.

1. Perform problem analysis on the failing storage unit.

   Performing problem analysis can cause mirrored protection to be suspended on the failing storage unit, and in some cases, on additional storage units.
2. Power off the failing storage unit.
3. Repair or replace the failing storage unit.
4. If the Replace Configured Unit option is necessary, the new storage unit is formatted and initialized, and mirrored protection is automatically resumed.
5. Resume mirrored protection on the repaired unit, if necessary, and on any other units that were suspended as part of the repair action.

   Synchronization begins on the resuming storage units immediately, and a message is sent to the QSYSOPR message queue when synchronization completes.

Actions when concurrent maintenance is not possible

This topic provides actions to take when concurrent maintenance is not possible.

1. Power off the system.
2. If unit 1 has failed, see the “Mirrored protection: Configuration rules” on page 431 topic for restrictions that apply.
3. Perform an attended initial program load (IPL) to dedicated service tools (DST).
4. Perform problem analysis on the failing storage unit.

Performing problem analysis might cause mirrored protection to be suspended on the failing storage unit, and in some cases, on additional storage units.

5. Power off the failing storage unit.
6. Repair or replace the failing storage unit.
7. If the Replace Configured Unit option is necessary, the new storage unit is formatted and initialized, and mirrored protection automatically resumes.
8. Resume mirrored protection on the repaired unit, if necessary, and on any other storage units that were suspended as part of the repair action.
9. Continue the IPL to command entry. Synchronize the resuming storage units during the IPL.

Other recovery considerations for mirrored protection

This topic provides information about other recovery considerations for mirrored protection.

Message handling

When a system with mirrored protection experiences a disk failure, the only external indication of the failure is a message that is sent to the system operator message queue (QSYSOPR). If there is a message queue that is named QSYSMSG in the QSYS library, the messages are sent there also.

When suspended units exist, the system sends a message every hour to the QSYSOPR message queue as a reminder.

You should have a method of bringing these messages to the attention of the system administrator. If the interactive job at the console allocates the QSYSMSG message queue and places it in break mode, it notifies you of any problems.

Synchronization

When the system is synchronizing (resuming) a disk unit, the system response time becomes longer.

When mirrored protection is resumed on a suspended disk unit at dedicated service tools (DST), the synchronization to the IBM i licensed program is done during the initial program load (IPL).

Related information

Using QSYSMSG message queue

Mirrored protection disk-error handling

Mirrored protection handles disk errors before the initial program load (IPL) to the IBM i operating system.

Unrecoverable device error

1. The system suspends the failing storage unit and mirrored protection is suspended for the mirrored pair.
2. The system continues operations by using the operating storage unit of the mirrored pair.
3. A message sent to the QSYSOPR message queue identifies the failing storage unit. It informs you that mirrored protection is suspended for the mirrored pair.

Permanent read error

1. The system reads from the other storage unit of the mirrored pair. If the permanent read error occurs on the other storage unit as well, the original read request completes with a permanent read error.

2. If the read from the other storage unit is successful, the data is written back to the first unit of the mirrored pair, assigning an alternate sector. Then the system signals that the original read request is complete.

Not an operational storage unit

1. The system attempts recovery. If it is successful, normal system operations continue with mirrored protection and without suspending or synchronizing the unit.
2. If recovery is unsuccessful, the unit is considered to have an unrecoverable device error, which is processed as described earlier.

Timeout

1. The system attempts to recover from the timeout. If it is successful, normal system operations continue with mirrored protection and without suspending or synchronizing this unit.
2. If recovery is unsuccessful, the storage unit is considered to have an unrecoverable device error, which is processed as described earlier.

I/O processor or bus failure

1. The system suspends each disk unit that is attached to the failing input/output (I/O) processor or bus in the same way it is done for an unrecoverable error.
2. The system saves a copy of the failing I/O processor's storage so the problem can be diagnosed. The system continues without the failing I/O processor.

Disk-related failure of unit 1 before the initial program load (IPL) to the IBM i operating system

See rzarmmprules.dita for restrictions that apply.

Missing disk units

If a disk unit, a controller, or an input/output (I/O) processor fails during an initial program load (IPL), the system detects the failure and does one of these actions.

- Displays a system reference code (SRC) on the control panel if the keylock switch is not in the Manual position.
- Displays the Missing Disk Unit display on the console if the keylock switch is in the Manual position.

If the failing unit has mirrored protection and its mirrored unit is active, the following display is shown.

```
              Disk Configuration Warning Report

Type option, press Enter.
 5=Display Detailed Report

Press F10 to accept all the warnings and continue the IPL.
The system will attempt to correct the warnings.

OPT Warning
 5  Missing mirror protected units in the configuration
```

If you type a 5 in the option column and press the Enter key, the following display is shown.

```
                  Suspend Missing Disk Units

The following disk units are missing from the disk configuration:
                                          Serial
Resource Reference ASP  Unit  Type  Model  Number      Name     Code
                     1     2  6602    030  00-0190494  DD036    1713
```

You can suspend mirrored protection on the affected units and continue the IPL. An entry is written in the problem log. You can run problem analysis on the failing unit at a later time. The **type** and **reference code** fields can be used with the unit reference code guide to determine the cause of the problem. If the keylock switch is not in the Manual position, a system reference code is displayed on the control panel. If the affected units do not report to the system within six minutes, the system automatically suspends mirrored protection on the affected units and continues the IPL.

If the suspended disk units become ready before the system is powered down, the system will automatically resume mirrored protection on these units.

Saving a unit

The system allows you to save data from storage units that use the dedicated service tools (DST) Save Disk Unit Data option.

The following rules apply to saving units on a system with mirrored protection:

- Only configured units can be saved.
- The save operation is not allowed when both mirrored units of a mirrored pair are active. Only one of the mirrored units can be saved. Therefore, one mirrored unit must be suspended.
- Only the active unit of a mirrored pair can be saved because the active unit contains the current data.
- If multiple failures cause the state of unit 1 to be unknown, saving of any storage unit is not allowed.

Restoring a unit

In the mirrored environment, the system allows you to restore data to storage units.

The following rules apply to restoring units on a system with mirrored protection:

- The restore is possible only for an active device.
- This option can restore to either a configured or nonconfigured disk unit.
- The restore operation requires that the unit restored to is as large as or larger than the unit saved.
- The restore operation is not permitted if the state of a unit is unknown. You can restore unit 1 only to the initial program load (IPL) device.
- After a unit is restored, the system performs an IPL to dedicated service tools (DST).
- The unit being restored must satisfy all mirrored protection configuration rules and restrictions.

Active mirrored load source failure

If unit 1 is mirrored, the system attempts to load the initial program from a load source mirrored unit that contains Licensed Internal Code and system data. The mirrored unit state of that storage unit will be active.

System cannot find active mirrored load source for initial program load

If the system cannot find a load source unit that contains current data and can only find a load source unit that is suspended or resuming, the system will load the initial program on the suspended or resuming unit.

This suspended or resuming unit contains data at the previous level. The system cannot be used until it finds or repairs the active mirrored load source.

If the system has not been able to IPL on an active mirrored load source, it is presumed to be broken in some way, and the following displays are shown.

```
                      Disk Configuration Error Report

Type option, press Enter.
  5=Display Detailed Report

OPT   Error
 5    Load source failure
```

Type a 5 in the option column and press the Enter key.

```
                          Display Load Source Failure

The system could not use the load source disk unit that
contains correct data.

The following disk unit contains the correct data:

Disk unit:
  Type . . . . . . . . . . . . . . . . . . :   6603
  Model  . . . . . . . . . . . . . . . . . :   030
  Serial number  . . . . . . . . . . . . . :   00-0193825
  Resource name  . . . . . . . . . . . . . :   DD001

Press Enter to use dedicated service tools (DST).
```

Active mirrored load source being used for initial program load failure

If the system is loading the initial program on an active mirrored load source and that storage unit fails during the initial program load (IPL) to dedicated service tools (DST) or at dedicated service tools (DST), the system attempts to perform a directed IPL to the other storage unit (tries to reload the initial program on the remaining load source).

- If the directed IPL fails, the system ends abnormally and displays a system reference code.
- If the remaining storage unit of the load source mirrored pair is active and the original load source is still broken on the re-IPL, the broken load source is treated like any other missing mirrored unit, and the following are displayed:

```
              Disk Configuration Warning Report
   Type option, press Enter.
     5=Display Detailed Report

   Press F10 to accept all the warnings and continue the IPL.
   The system will attempt to correct the warnings.

   OPT Warning
    5  Missing mirror protected units in the configuration
```

  Type a 5 in the option column and press the Enter key.

```
                    Suspend Missing Disk Units
The following disk units are missing from the disk configuration:
                                          Serial
Resource Reference ASP  Unit  Type  Model  Number      Name     Code
                    1      2  6602   030   00-0190494  DD036    1713
```

- If the remaining storage unit of the load source mirrored pair does not contain current data (it is suspended or resuming), it is treated as if the system cannot find an active mirrored load source for IPL, as described earlier. The IPL is not allowed to continue past DST until the active load source is found or repaired.

Active mirrored load source fails late in the initial program load (IPL) or at run time

When an active mirrored load source fails after storage management recovery is completed, it is treated like a failure in any other mirrored pair.

- If the other storage unit in the mirrored pair is present and active, the failing unit is suspended and the system continues to run using data the data on the remaining active unit of the pair.
- If the failing storage unit is the last active unit of the mirrored pair (the other unit of the pair is either suspended or resuming), the system displays a Disk Attention system reference code and becomes unusable.

Cannot read system configuration data from active mirrored load source

If the system cannot read the system configuration data from the active mirrored load source that is being used for the initial program load (IPL), you can see one of these displays.

```
                    Accept Load Source Warning Report

Some of the configuration information on the load source is
missing.  The system can rebuild this information using
the default values.

Press Enter to let the system rebuild the configuration
information on the load source.

If you were performing any disk unit recovery actions, go to
Work with Disk Units and complete those actions.
```

```
                  Disk Configuration Warning Report
   Type option, press Enter.
     5=Display Detailed Report

   Press F10 to accept all the warnings and continue the IPL.
   The system will attempt to correct the warnings.

   OPT  Warning
    5   Bad load source configuration
```

Unknown unit 1 status

This topic provides information about unknown unit 1 status.

If both the service processor and one unit of the mirrored pair for unit 1 have failed, the following display is shown.

```
                     Disk Configuration Error Report

Type option, press Enter.
  5=Display Detailed Report

OPT   Error
 5    Unknown load source status
```

Type a 5 in the option column and press the Enter key.

```
                        Display Unknown Mirrored Load Source Status

The system cannot determine which disk unit of the load
source mirrored pair contains the correct level of data.

The following disk unit is not available:

Disk unit:
    Type . . . . . . . . . . . . . . . . . . :   6603
    Model  . . . . . . . . . . . . . . . . . :   030
    Serial number  . . . . . . . . . . . . . :   00-0193825
    Resource name  . . . . . . . . . . . . . :   DD001

Press Enter to continue.
```

If the keylock switch is not in the Manual position, the control panel displays a system reference code.

The missing unit must be repaired or the state of the unknown load source recovered. If the missing unit can be repaired without losing the data on the unit, then the state of the load source will become known when the system is IPLed. If the missing unit cannot be repaired or if the data on it is lost, then you can possibly recover the state of the unknown load source and avoid restoring the entire system.

You should only attempt to recover the state of the unknown load source when you know its mirrored unit state was active before the failures that caused the state to become unknown. Because the state is unknown, the system cannot verify that your choice is correct. If you recover the state of the unknown load source when the actual state of the disk unit used to initial program load (IPL) was not active, you will cause data loss or damaged objects on your system.

Recovering the state of the unknown load source

Use these steps to recover the state of the unknow load source.

1. From the DST main menu, select option 4, Work with disk units.
2. From the Work with disk units menu, select option 2, Work with disk unit recovery.
3. From the Work with disk unit recovery menu, selection option 15, Recover unknown load source.

   This will display a confirmation screen, showing the disk configuration and mirrored unit states the system will have after the recovery.
4. If the configuration and states are as you expect, press Enter to confirm.

   The state of the load source mirrored pair is changed so that the load source just used to IPL is active and the other (missing) load source is suspended.

If you cannot recover the state of the unknown load source and if the missing unit cannot be repaired, you must install the Licensed Internal Code and restore the entire system.

Displaying incorrect Licensed Internal Code installation

When the Licensed Internal Code is restored on a mirrored unit for unit 1, one of the mirrored units might have the incorrect level of data stored on it.

If this condition occurs, and the disk unit that contains the correct data is not available, the Licensed Internal Code is restored to the disk unit with the incorrect data. When a disk preforms an initial program load (IPL) and the correct disk unit is available, the following display is shown. If the keylock switch is not in the Manual position, system reference code (SRC) is shown on the control panel for the system or logical partition.

```
              Display Incorrect Licensed Internal Code Install

Licensed Internal Code has been installed on the incorrect
disk unit of the load source mirrored pair.

If you continue the IPL, the previously installed Licensed Internal
Code installed on the incorrect disk unit of the mirrored
load source pair will be deleted.  The Licensed Internal Code
will be replaced by the Licensed Internal Code from the correct
disk unit.

The following disk unit is the correct disk unit.

Disk unit:
  Type  . . . . . . . . . . . . . . . . . :     6602
  Model   . .  . . . . . . . . . . . . . . :     030
  Serial number. . . . . . . . . . . . . . :     00-0163477_
  Resource name  . . . . . . . . . . . . . :     DD019

Press Enter to continue.
```

Recovering remote load source mirroring

Use this information to recover remote load source mirroring.

Recovering after the remote load source has failed

If the load source is not attached to its controlling input/output processor (IOP), a non-load source disk unit failure will occur.

After the disk unit has been suspended, it can be repaired or replaced using concurrent or deferred maintenance, depending of the type of failure. While the remote load source is suspended, the system can still load the initial program by using the load source.

Recovering after the local load source has failed on the system

When the load source that is attached to its controlling input/output processor (IOP) has failed, the system can continue to run using the data on the other load source.

However, if the system is powered down, either for the repair or for another reason, it cannot load the initial program again without repairing the broken load source, because, the system can only load the initial program from the load source attached to its controlling IOP. If the local load source can be repaired or replaced using concurrent maintenance procedures, without powering down the system, then system use will not be interrupted. If the IPL load source has failed in such a way that concurrent maintenance is not possible, or if the IOP that controls the load source or Bus 1 have failed, then the system must be powered down, the broken hardware repaired or replaced, and the mirrored load source recovered.

Use normal system problem analysis to determine if concurrent maintenance is possible. If it is, use concurrent maintenance to repair or replace the broken load source. The repaired or replaced load source will be synchronized with the other load source, which the system has continued to use. When the repaired load source is synchronized, the load source is fully protected again and can be used for IPL and main storage dumps.

When concurrent maintenance is not possible, you can sometimes defer maintenance until a convenient time. However, the load source and your system will be running unprotected until all service and mirrored synchronization have been completed. If the other load source fails, data might be lost and the system might have to be reloaded.

There are two scenarios for dedicated recovery after a local load source failure.

Dedicated recovery of local load source – Local system is still functional

To recover from a local load source failure that uses dedicated maintenance when the local system unit and site are functional, follow these steps.

Note: If the input/output processor (IOP) that controls the load source supports 520 or 522 byte sector disk units, then you can move the remote load source disk units to it.

1. Power off the system.
2. Repair or replace the broken load source.
3. Install the Licensed Internal Code on the repaired/replaced load source.

   After the install completes, the system will automatically initial program load (IPL) to dedicated service tools (DST). When the system reaches DST, all disk units and the system disk configuration will be missing.

   Note: Press F3 to go to DST and perform the remote load source recovery actions. Do not press F10 to accept the New Disk Configuration warning. If you do, you must restart the dedicated recovery from remote disk at Step 2 above.
4. Use the Recover Mirrored Load Source function to recover the disk configuration and data.

   When recovery is complete, the system will re-IPL automatically.

Dedicated recovery from remote disk units – After a local site disaster

To recover the system after a local site disaster, using the set of remote disk, follow these steps.

Note: If your input/output processor (IOP) that controls your load source supports 520 or 522 byte sector disk, then you can move the remote load source disk to the replacement system.

1. Attach a new system unit to the remote disk.

   The new system unit must contain a disk unit that will become the new load source.
2. Install the Licensed Internal Code to the disk unit in the new system unit.

   After the install completes, the system will automatically initial program load (IPL) to dedicated service tools (DST). When the system reaches DST, all disk units and the system disk configuration will be missing.

   Note: Do not accept the New Disk Configuration warning. If you do, you must restart at Step 2.
3. Use Recover Mirrored Load Source to recover the disk configuration and data.

   When recovery is complete, the system will automatically re-IPL.

Using the recover mirrored load source function

The recover mirrored load source function is used after a local load source failure to recover both the system and user data from the remote load source and to copy the data to the new local load source.

Recover mirrored load source assumes that the load source used for the initial program load (IPL) is new and has just been installed. This should be true whether recovering after a site disaster or after a local load source failure. After the new load source is installed and the system IPLs to dedicated service tools (DST), recover mirrored load source finds the remote load source and then synchronizes the new load source from the remote load source. When the synchronization is complete, the system re-IPLs automatically, because the Licensed Internal Code that was copied onto the IPL load source during synchronization might be different from the code used to IPL to DST.

Recover mirrored load source requires the following conditions:

- The current load source must have just been installed.
- The system must be able to find a valid configuration on the other disk units that are attached to the system.
- In the recovered configuration, remote load source mirroring must have been enabled.
- In the recovered configuration, the system ASP and the load source must have been mirrored.

- The remote load source must be attached to the system, it must be functional, and its mirrored unit state must be active; that is, it must contain current data.

To use the recover mirrored load source function, follow these steps:

1. From the Dedicated Service Tools (DST) Main Menu, selection option 4, Work with disk units.
2. From the Work with disk units menu, select option 2, Work with disk unit recovery.
3. From the Work with disk units recovery menu, select option 16, Recover mirrored load source.

   The system checks to see if it is possible to recover the mirrored load source. If recovery is possible, the system finds the best recovered configuration, finds the remote load source to use for recovery, and verifies that the remote load source and other disk units of the configuration are functional.

   - If the mirrored load source can be recovered, the system will show a confirmation display that indicates the configuration that will be recovered.
   - If the mirrored load source cannot be recovered, the system will display an error message. Follow the recovery actions that are listed for the error, and resolve the problem, if possible. If it is not possible to recover the mirrored load source, the data in the system auxiliary storage pool (ASP) will be lost. If other ASPs in the system are intact, you can use the Recover Configuration function to recover the system configuration and the data in the complete ASPs.
   - If the mirrored load source can be recovered, but there is not an active, usable disk unit present for each configured logical unit of the system (at least one active unit of each mirrored pair present and all device parity and unprotected units present), then all data will be lost in ASPs that have missing units.

   The system will display a warning message for the ASPs that will lose data. Follow the recovery actions and try to resolve the missing unit situations if possible before continuing. If you cannot make the missing units present and active, then you can still continue, but all data will be lost in the ASPs where units are missing.
4. Press Enter on the confirmation display to begin to recover the mirrored load source.

   The recovery will use the configuration and remote load source that are found, verified, and displayed on the confirmation display. During recovery, the following actions occur:

   - The system copies the recovered load source onto the load source used for the IPL. All data on the disk , including Licensed Internal Code, system data, and user data, are copied, except for the few pages that uniquely identify each individual disk unit and help distinguish between the two disk units of a mirrored pair. Those pages are built separately for the recovered load source.

     During the recovery the system displays the System Reference Code C6 XX 4205 in the control panel on the system or logical partition. The percent complete, in decimal, is displayed in the third and fourth digits of the reference code (the XX).
   - When the recovery of the load source data is complete, the system performs a directed IPL, using the load source attached to the input/output processor (IOP) that controls the load source, which now contains the recovered data.
5. If any errors (for example, I/O errors or hardware failures) occur after the copy phase of recovery has begun, you must start the recovery process over again, beginning with the Licensed Internal Code install of the new load source.

Chapter 15. Restoring your system using Operational Assistant tapes

According to the different situations, perform one of these three tasks to restore information from Operational Assistant backup tapes.

The descriptions assume that you are recovering all the data on your system. If you are recovering a single library or a single auxiliary storage pool (ASP), adapt the procedures to your situation.

Figure 30 on page 303 shows the parts of your system and how they are saved with Operational Assistant.

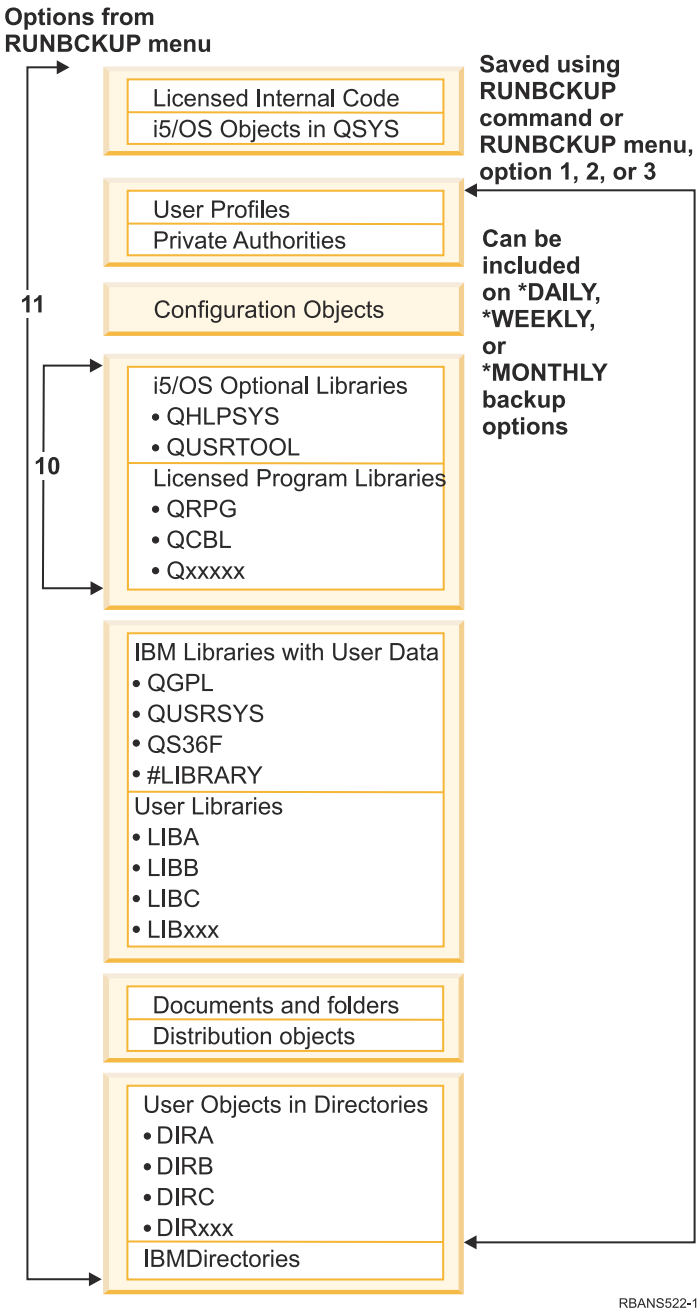

Figure 30. How the system is saved with Operational Assistant backup

© Copyright IBM Corp. 1997, 2025

Related tasks

Checklist 27: Recovering user information using tapes from Operational Assistant backup
The checklist shows the sequence of steps to use to recover user information from an Operational Assistant backup.

Restoring your libraries

To recover your entire system, you must restore IBM-supplied libraries and user libraries.

To restore IBM-supplied libraries, follow these steps:

1. Find the tapes that you used most recently to save IBM-supplied libraries. You saved them by using one of the following methods:
   - Option 10 from the Run Backup menu.
   - Option 11 from the Run Backup menu.
   - The `SAVLIB LIB(*IBM)` command.
   - The `SAVLIB LIB(*NONSYS)` command.
   - Option 21 from the Save menu.
   - Option 22 from the Save menu.
   - Option 41 from the Save menu.
2. Mount the first tape and type: `RSTLIB SAVLIB(*IBM) DEV(media-device-name)`. Press the Enter key.

To restore user libraries, follow these steps:

1. Find the tapes that you used most recently to save all user libraries.

   You saved them by using one of the following methods:
   - Option 1, 2, or 3 from the Run Backup menu and specifying 2 (All) for the User libraries option.
   - Option 11 from the Run Backup menu.
   - The `SAVLIB LIB(*ALLUSR)` command.
   - The `SAVLIB LIB(*NONSYS)` command.
   - Option 21 from the Save menu.
   - Option 23 from the Save menu.
   - Option 40 from the Save menu.
   - Option 42 from the Save menu.

   If you are not sure which tapes have your user libraries on them, perform the following tasks for each possible tape:

   a. Mount the tape.
   b. Type `DSPTAP DEV(media-device-name)`
   c. Page through the displays, looking for the file called QFILE.
   d. When you find the tape with the QFILE file on it, write down the sequence number for that file on the tape.
   e. Leave the tape in the tape unit and type:

   ```
   DSPTAP DEV(media-device-name) LABEL(QFILE)
   SEQNBR(sequence-number) DATA(*SAVRST) OUTPUT(*PRINT)
   ```

   f. If the listing contains user libraries, it was created by either the SAVLIB(*NONSYS) command or the SAVLIB(*ALLUSR) command. The libraries from the tape can be restored using the RSTLIB SAVLIB(*ALLUSR) command.
2. Mount the first tape that has user libraries and type: `RSTLIB SAVLIB(*ALLUSR) DEV(media-device-name)`. Press the Enter key.

You have now restored all the libraries on your system to the point where they were all saved completely. Return to "Checklist 27: Recovering user information using tapes from Operational Assistant backup" on page 116.

Restoring libraries that you saved by using a backup list

Use this procedure to restore libraries that you saved using a backup list, either for daily or weekly backup. It assumes that you save all libraries for your monthly backup.

Use this procedure if all of the following conditions are true:

- You have an Operational Assistant backup that is more recent than the last time you saved the entire system or all libraries.
- You specified 1 (Selected from list) for the User libraries option for your Operational Assistant backup.
- You specified N (No) for the Save changed objects only option for your Operational Assistant backup.

If you have both a weekly and a daily backup that meet these conditions, perform the following tasks:

- If your daily backup and weekly backup both save exactly the same libraries from the backup list, perform steps 2 through 4 once, using your most recent set of tapes (daily or weekly).
- If your daily backup saves fewer libraries than your weekly backup, perform the following tasks:
  - If your most recent backup is a weekly backup, perform steps 2 through 4 once, using your most recent set of weekly tapes.
  - If your most recent backup is a daily backup, perform steps 2 through 4 once, using your most recent set of weekly tapes. Repeat step 2 through 4, using your most recent set of daily tapes.

1. Mount the first tape.
2. Find the printed copy of the backup list associated with the save tapes. If you have the list, skip to step 4
3. If you do not have the list, display the contents of the save tapes by typing: `DSPTAP DEV(media-device-name) OUTPUT(*PRINT) DATA(*SAVRST)`.
4. Use the listing from step 2 or step 3.

   For each library that was saved, follow these steps:

   a. Type: `RSTLIB SAVLIB(`*`library-name`*`) DEV(media-device-name)`.

   b. Check off the library name on the list.

   Note: Restore the user libraries to each user auxiliary storage pool (ASP) that you are recovering. If you are restoring the QGPL library and the QUSRSYS library and doing partial recovery, restore these libraries before any other libraries.

Restoring changed objects that you saved by using Operational Assistant

If you save only changed objects for either your weekly or daily backup, use this procedure.

If you save changed objects both weekly and daily, use your most recent set of tapes. If you save complete libraries on your weekly backup and changes on your daily backup, perform this procedure only if your daily backup is more recent than your weekly backup.

1. Mount the first tape from your most recent backup of changed objects.
2. Determine if any objects are on the tape for libraries that do not exist on the system:
   a) Print a list of libraries on the system by typing: `DSPBCKUPL OUTPUT(*PRINT)`.
   b) Print the contents of the tape by typing: `DSPTAP DEV(media-device-name) OUTPUT(*PRINT) DATA(*SAVRST)`.
   c) Compare the two lists. Mark any libraries on the DSPTAP listing (from step 2b) that do not appear on the DSPBCKUPL listing (from step 2a).

d) For any libraries you marked in step 2c, type the following command: `CRTLIB LIB(`*`library-name`*`)`.

3. Restore the changed objects from the tapes.

   For each library that appears on the DSPTAP listing (from step 2b), type the following command:

```
RSTOBJ OBJ(*ALL) SAVLIB(library-name)
       OBJTYPE(*ALL) DEV(media-device-name)
```

Chapter 16. Restoring encrypted backups and encrypted auxiliary storage pools

You can either encrypt backup data to a tape, or encrypt data stored in user auxiliary storage pools (ASPs) and independent ASPs. If you lose the data on your system, you need to be able to restore the encrypted data.

Recovering from an encrypted backup using software encryption

Backup Recovery and Media Services (BRMS) provides you with the ability to encrypt your data to a tape device. This method is called *software encryption*, because you do not need to use an encryption device. The BRMS interface asks for the encryption key information and what items you want encrypted. BRMS saves the key information, so for restoring, BRMS knows what key information is needed to decrypt on the restore.

The key used to encrypt the data on tape is stored in a cryptographic services keystore file called QUSRBRM/Q1AKEYFILE. All key values in the keystore file are encrypted under a master key. If the master key is not set correctly or is missing, or if the keystore file is missing, or the key record in the keystore file is missing, you cannot recover the encrypted data off the tape.

If you are restoring the encrypted backup on another system, ensure that the keystore file QUSRBRM/Q1AKEYFILE exists. If not, perform one of the following methods:

- Move the keystore file from the source system to the target system.

  See Distributing keys.

  a) Use the Save Object (SAVOBJ) CL command to copy the keystore file from the source system.

  b) Transfer the media to the target system.

  c) Use the Restore Object (RSTOBJ) CL command to copy the keystore file to the target system.
- Creating a new keystore file.

  If you create the keystore file, or if it already exists, you must move the key used to encrypt the data to tape into the keystore file. For information about moving keys from a keystore file on one system to a keystore file on another system, see Distributing keys.

  Note: The master key value that encrypts the key values in QUSRBRM/Q1AKEYFILE must be identical on both systems.

If you must restore the master key (for example, the Licensed Internal Code was reinstalled, or you are restoring on another system), use one of the following methods:

- Reload the individual passphrases and set the master key.
- Restore the master keys from a Save System (SAVSYS) tape. In this situation, you must ensure that the save/restore master key on the target system matches the save/restore master key on the source system.

For information about using BRMS to encrypt your data to a tape device, see "Software encryption using BRMS" in *Backup, Recovery, and Media Services for i.*

Related tasks

Checklist 21: Recovering your entire system after a complete system loss including independent auxiliary storage pools

Use this checklist if you need to restore your entire system that includes an independent ASP to a system that is running the same version of the IBM i licensed program. Do not use this checklist if you are doing an upgrade.

Related information

Creating a new keystore file

© Copyright IBM Corp. 1997, 2025

Decrypting your data
Saving and restoring master keys
Loading and setting master keys

Recovering from an encrypted backup using an encrypted tape

Hardware tape encryption uses tape devices with data encryption capabilities and key management software to encrypt your data. IBM i supports only library-managed encryption. Use these steps to recover data that you backed up using an encrypting tape drive or tape library.

To restore from an encrypted backup using an encrypting tape drive or tape library, follow these steps:

1. Ensure that the key management software is running and connected to the system where you plan to restore the data.

   The key manager contains the encryption keys that are needed for the recovery operation.
2. Restore the data from the most recent backup tape. When the data is restored, it is decrypted.

   When you share tapes with another company, the key manager writes the tape with the other company's public key. They can decrypt and read the tape by using their private key.

Attention: It is important to preserve your keystore data, which is stored in the key manager. Without access to your keystore data, you cannot decrypt your encrypted tapes during a restore operation. Back up the keystore data so that you can recover it as needed. You also can have two key managers that are mirror images of each other with built-in backup of the critical keystore information, as well as a failover if one key manager becomes unavailable. When you configure your tape device, you can point it to two key managers. If one key manager becomes unavailable for any reason, your device uses the alternate key manager.

Related tasks

Backing up the Encryption Key Manager
Hardware tape encryption uses tape devices with data encryption capabilities and key management software to encrypt your data. Use these steps to back up and restore the key manager data. If you lose the encryption keys in the key manager, you will not be able to decrypt your tapes in a system recovery.

Related information

Tape encryption

Backing up the Encryption Key Manager

Hardware tape encryption uses tape devices with data encryption capabilities and key management software to encrypt your data. Use these steps to back up and restore the key manager data. If you lose the encryption keys in the key manager, you will not be able to decrypt your tapes in a system recovery.

Key management software assists IBM encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information being written to, and decrypt information being read from, tape media. The key manager operates on IBM i, and many other system platforms. The key manager can serve numerous IBM encrypting tape drives, regardless of where those drives reside. The key manager uses a keystore to hold the certificates and keys required for all encryption tasks. You can have multiple copies of the key manager on the network.

The key manager uses the following method to handle save requests.

- The tape library receives a save request with a volume serial that is marked for encryption.
- The tape library asks key manager to generate a random data key.
- The key manager generates the data key for this tape. This data key is used to encrypt the data.
- The key manager uses the public key to encrypt the data key that is ready to be stored on the tape.
- The tape library writes the encrypted data key on the cartridge in both the cartridge memory and on the tape.
- The tape library uses the session key to encrypt the data as it writes it to the tape.

During a restore, the key manager decrypts the key using the public/private pair. The library uses the data key to decrypt the data as it reads it from the tape.

Important: Due to the critical nature of the keys in your keystore, it is highly recommended that you back up this data so that you can recover it as needed, and be able to read the tapes that were encrypted using those certificates associated with that tape drive or library.

Use any of the following methods to back up this keystore information in the key manager:

- Keep a copy of all certificates loaded in the keystore.
- Use system backup capabilities, such as save/restore commands or BRMS commands, to create a backup copy of this keystore information.

  Be careful not to encrypt this copy using the encrypting tape drives, as it would be impossible to decrypt it for recovery.
- Maintain a primary and secondary key manager and keystore copy for backup, as well as for high availability.

  You can have two key managers that are mirror images of each other with built-in backup of the critical keystore information. When you configure your tape device, you can point it to two key managers. If one key manager becomes unavailable for any reason, your device will use the alternate key manager.
- If you are using a JCEKS (UNIX System Services file-based) keystore, copy the keystore file and store the clear (unencrypted) copy in a secure location, such as a vault.

  Be careful not to encrypt this copy using the encrypting tape drives, as it would be impossible to decrypt it for recovery.

It is important to test your recovery strategy carefully. At the primary site, run multiple key manager servers so that backups can continue to run while one key manager server is down. Export and synchronize keys on all key manager servers each time the keys change. Keep an offsite backup of the key manager data. At the disaster recovery site, have an encryption-capable tape drive or library with access to the key manager server. Do not encrypt the key manager server. Run the key manager on a system or logical partition where none of the save operations are encrypted.

Related tasks

Recovering from an encrypted backup using an encrypted tape
Hardware tape encryption uses tape devices with data encryption capabilities and key management software to encrypt your data. IBM i supports only library-managed encryption. Use these steps to recover data that you backed up using an encrypting tape drive or tape library.

Related information

IBM Security Key Lifecycle Manager

Restoring encrypted auxiliary storage pools

If you have an encrypted user or independent auxiliary storage pool (ASP), you must run special steps to verify that the data in these ASPs can be recovered.

To use disk encryption, you must have 5770-SS1 Option 45 - Encrypted ASP Enablement installed. The option to enable encryption is available when you create a user auxiliary storage pool (ASP) or independent ASP in IBM Navigator for i. Set the ASP master key before you can create an encrypted independent auxiliary storage pool. The data keys for independent ASPs are kept with the storage pool and are protected with the ASP master key.

The ASP master key is not needed to create an encrypted user ASP.

After you create an encrypted ASP or change the data key for an encrypted ASP, run a Save System (SAVSYS) operation so that the media has the correct encryption keys. The encryption keys are stored in the system ASP and saved during the SAVSYS operation.

If disk encryption is used in a clustering environment, you must set the master key manually on each system within the device domain.

Important: If you are using encrypted user ASPs and the system ASP fails, install the system ASP with the SAVSYS media that contains the encryption keys. If the SAVSYS media does not contain the keys, the encrypted ASPs are unusable because the encryption keys do not exist on the system. If the encrypted user ASP is not usable, the system will not IPL.

If you are using encrypted independent ASPs and the system ASP fails, install the Licensed Internal Code from SAVSYS media that contains the ASP master key. If you do not have SAVSYS media that contains the ASP master key, you can manually set the ASP master key after the installation completes. The encrypted independent ASPs cannot be varied on to the system until the ASP master key is set correctly.

Remember: If you restore the Licensed Internal Code from the save media after a scratch installation, encrypted ASPs that are already configured are functional. However, you must IPL to activate the Encryption ASP Enablement option so that you can create new encrypted ASPs.

Recovering an encrypted user ASP - If you have an encrypted user ASP, choose one of the following methods to recover the data in the encrypted user ASP:

- Reinstall the operating system by using the most recent SAVSYS media.

  Reinstalling the operating system is only necessary if the system ASP is lost. If a user ASP failed, the keys are still set in the system ASP.
- Delete and re-create the user ASP.
- Clear the user ASP. If a bad disk is the reason for needing to recover the data in the user ASP, remove or replace the failing drive.

Recovering an encrypted independent ASP - If you have an encrypted independent ASP, choose one of the following methods to recover the data in the independent ASP:

- Reinstall the operating system by using the most recent SAVSYS media. Reinstalling the operating system is only necessary if the system ASP is lost. If an independent ASP failed, the key is already be set in the system ASP.
- Delete and re-create the encrypted independent ASP.
- Clear the independent ASP. If a bad disk is the reason for needing to recover the data in the independent ASP, remove or replace the failing drive.
- Manually load and set the ASP master key. Run this step only if you were unable to restore the SAVSYS media with the most recent master keys.

Related concepts

Recovering a basic user auxiliary storage pool (ASP) after recovering the system ASP
After restoring your Licensed Internal Code and the operating system, you need to perform the tasks in this topic.

Recovering an independent auxiliary storage pool
Run this procedure if one of the following situations is true.

Related information

Backing up encrypted auxiliary storage pools

Loading and setting auxiliary storage pool master key

Creating an independent disk pool

Managing master keys

Part 3. Release-to-release support

© Copyright IBM Corp. 1997, 2025

Chapter 17. Release-to-release support

The release-to-release support on the IBM i environments allows you to move data from the current release to a previous-release system. This support also allows you to move data from a previous-release system to a current-release system.

Releases of licensed programs on the IBM i environments have a three-part name that consists of a version, a release, and a modification. For example, the current version is version 7, release 1, modification 0. The short form of the current release name is i 7.1.

Related concepts

What you should know about restoring user profiles
When you are restoring user profiles from a source system to a target system, you must make sure that the password level values (QPWDLVL) are compatible.

Restrictions when going from previous release to current release
Previous releases of IBM i might support products that are not supported at the current release.

Current release to previous release support

This support enables objects (such as programs and files) that are created and saved on the current release to be restored and used on a previous release.

Object compatibility is provided for in many languages and most object types are supported on both release levels if the objects use only functions from a previous release.

You can enable current release-to-previous release support by using the target release (TGTRLS) parameter on create and save commands.

Table 62 on page 313 illustrates the TGTRLS parameter and values available for the current and previous releases. Refer to this table to determine the valid values for the release currently on your system.

The TGTRLS values must use the form VvRrMm.

Table 62. Values for TGTRLS parameter

Current IBM i release	*CURRENT	*PRV	Other valid values
V7R6M0	V7R6M0	V7R5M0	V7R4M0
V7R5M0	V7R5M0	V7R4M0	V7R3M0
V7R4M0	V7R4M0	V7R3M0	V7R2M0
V7R3M0	V7R3M0	V7R2M0	V7R1M0
V7R2M0	V7R2M0	V7R1M0	V6R1M0
V7R1M0	V7R1M0	V6R1M0	V5R4M0
V6R1M0	V6R1M0	V5R4M0	V5R3M0
V5R4M0	V5R4M0	V5R3M0	V5R2M0

This support is useful to the following situations:

- A network enterprise with some current release systems and previous release systems.
- An application development business with a single system on the current release that supports customers who might still be on the previous release.

Current release-to-previous release support provides great savings and productivity improvements to application developers. By using this support, most network enterprises and application development businesses no longer need to maintain two development systems. (For example, two development

© Copyright IBM Corp. 1997, 2025

systems might contain a previous-release system that contains previous-release objects, or a current-release system that contains current-release objects.) Usually this support enables previous release and current release objects to exist on one development system.

Creating the object for the previous release

If you want to run control language (CL) programs for a previous release, you must install option 9 (*PRV CL Compiler Support) from IBM i.

The following object types must be created specifically for a target release:

- Program (*PGM)
- Service program (*SRVPGM)
- Module (*MODULE)
- C locale description (*CLD)
- SQL package (*SQLPKG)
- Java program (*STMF)

Create the object on the current release by using the appropriate create command with the target release (TGTRLS) parameter. All other object types can skip this step.

Note: To determine what release an existing object was created on, use the Display Object Description (DSPOBJD) command with DETAIL(*SERVICE) to display the **System level** field.

It is suggested that previous-release programs and current-release programs be stored in separate libraries to simplify maintenance. Using the Display Program (DSPPGM) command, the **Earliest release that program can run** field shows if a program can be saved to the previous release. This suggestion also applies to previous-release and current-release modules and service programs. To determine whether *MODULE objects can be saved to a previous release, use the Display Module (DSPMOD) command. To determine whether *SRVPGM objects can be saved to a previous release, use the Display Service Program (DSPSRVPGM) command.

To determine whether Java class file (*STMF) or Java jar file (*STMF) objects with attached Java programs can be saved to a previous release, use the Display Java Program (DSPJVAPGM) command.

Table 63 on page 314 shows the languages and commands that support the TGTRLS parameter.

Table 63. Language support for the target release parameter

Language	Command
ILE C	CRTBNDC
	CRTCMOD
	CRTCLD
ILE C++	CRTBNDCPP
	CRTCPPMOD
CL	CRTBNDCL
	CRTCLMOD
	CRTCLPGM
ILE COBOL	CRTBNDCBL
	CRTCBLMOD
	CRTCBLPGM
	CRTS36CBL

Table 63. Language support for the target release parameter (continued)

Language	Command
ILE RPG	CRTBNDRPG
	CRTRPGMOD
	CRTRPGPGM
	CRTRPTPGM
	CRTS36RPG
	CRTS36RPT
SQL	CRTSQLCI
	CRTSQLCBL
	CRTSQLCBLI
	CRTSQLCPPI
	CRTSQLPLI
	CRTSQLRPG
	CRTSQLRPGI
Java	CRTJVAPGM
Other	CRTPGM
	CRTSRVPGM

Saving the object for the previous release

Save the object on the current release by using the target release (TGTRLS) parameter before you restore it on the previous release or previous modification. The object is saved in a format that is known to the previous release or previous modification.

Use communication lines or removable storage media (tape or optical media volume to move the objects from the current-release system).

It is suggested that you store previous-release and current-release objects in separate libraries to simplify maintenance.

The following save commands support the TGTRLS parameter:

- Save (SAV)
- Save Changed Objects (SAVCHGOBJ)
- Save Document Library Objects (SAVDLO)
- Save DLO using BRM (SAVDLOBRM)
- Save Folder List using BRM (SAVFLRLBRM)
- Save Library (SAVLIB)
- Save Library using BRM (SAVLIBBRM)
- Save Licensed Program (SAVLICPGM)
- Save Media Information using BRM (SAVMEDIBRM)
- Save Object (SAVOBJ)
- Save Object using BRM (SAVOBJBRM)
- Save Object List using BRM (SAVOBJLBRM)

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

Object compatibility is provided for most object types that are supported on both levels on the condition that the only previous-release functions are used by the object.

Table 64 on page 316 shows what object types can and cannot be created or saved for a previous release. IBM does not support saving IBM-supplied objects (such as system commands and programs) from the current release and restoring them on a previous-release system. Refer to Table 62 on page 313 for a list of supported TGTRLS values.

Table 64 on page 316 uses the following values:

- All means that the object can be saved to all TGTRLS values that are supported on the current version of the operating system.
- VvRrMm indicates the earliest release to which an object can be saved. However, you might need to refer to Table 62 on page 313 to find the earliest TGTRLS value that is supported on the current version of the operating system.
- *CURRENT means that the object can be saved only to the current release, TGTRLS(*CURRENT).
- None means that the object is saved by using a command that does not support the TGTRLS parameter, such as Save Security Data (SAVSECDTA) or Save Configuration (SAVCFG).

Table 64. Previous-release support by object types

Object type	**Earliest target release**
*ALRTBL	All
*AUTHLR	None
*AUTL	None
*BLKSF	All
*BNDDIR	All
*CFGL	None
*CHTFMT	All
*CHRSF	V5R1M0
*CLD	All
*CLS	All
*CMD	All
*CNNL	None
*COSD	None
*CRG	*CURRENT
*CRQD	All
*CSI	All
*CTLD	None
*DDIR	All

Table 64. Previous-release support by object types (continued)	
Object type	**Earliest target release**
*DEVD	None
*DIR	All
*DOC	All
*DSTMF	All
*DTAARA	All
*DTADCT	All
*DTAQ	All[5]
*EDTD	All
*EXITRG	All
*FCT	All
*FIFO	V5R1M0
*FILE (database, device, save) [9]	All
*FLR	All
*FNTRSC	All
*FNTTBL	V3R7M0
*FORMDF	All
*FTR	All
*GSS	All
*IGCDCT	All
*IGCSRT	All
*IGCTBL	All
*IMGCLG	V5R2M0
*IPXD	None
*JOBD	All
*JOBQ [10]	All
*JOBSCD	All
*JRN	All
*JRNRCV [4]	All
*LIB	All
*LIND	None
*LOCALE	V3R7M0
*MEDDFN	V4R4M0
*MENU	All
*MGTCOL	V4R4M0

Table 64. Previous-release support by object types (continued)	
Object type	**Earliest target release**
*MODD	None
*MODULE[1]	
• ILE C	All
• ILE C++	All
• ILE CL	All
• ILE COBOL	All
• ILE RPG/400®	All
*MSGF	All [6]
*MSGQ	All
*NODGRP	V3R2M0
*NODL	All
*NTBD	None
*NWID	None
*NWSCFG	V5R4M0
*NWSD	None
*OUTQ[8]	All
*OVL	All
*PAGDFN	All
*PAGSEG	All
*PDFMAP	V5R3M0
*PDG	All
*PGM[1]	
• BASIC	*CURRENT
• CL (S/38 environment)	*CURRENT
• CL (IBM i environment)	All
• COBOL (IBM i environment)	All
• COBOL/74 (S/38 environment)	*CURRENT
• COBOL/74 (S/36 environment)	All
• ILE C	All
• ILE C++	All

Table 64. Previous-release support by object types (continued)	
Object type	**Earliest target release**
• ILE CL	All
• ILE COBOL	All
• ILE RPG	All
• PASCAL	*CURRENT
• PL/I	*CURRENT
• RPG/II (S/36 environment)	All
• RPG/III (S/38 environment)	*CURRENT
• RPG (IBM i environment)	All
*PNLGRP	All
*PRDAVL	*CURRENT
*PRDDFN	All
*PRDLOD	All
*PSFCFG	V3R2M0
*QMFORM	All
*QMQRY	All
*QRYDFN	All
*RCT	*CURRENT
*SBSD	All
*SCHIDX	All
*SOCKET	None
*SPADCT	All
*SQLPKG	All
*SQLUDT	V4R4M0
*SQLXSR	V7R1M0
*SRVPGM[2]	
• ILE C	All
• ILE C++	All
• ILE CL	All
• ILE COBOL	All
• ILE RPG/400	All

Table 64. Previous-release support by object types (continued)	
Object type	**Earliest target release**
*SSND	All
*STMF[3]	All
*SVRSTG	V3R2M0
*SYMLNK	All
*S36	*CURRENT
*TBL	All
*TIMZON	V5R3M0
*USRIDX	All
*USRPRF	V5R1M0[7]
*USRQ	All
*USRSPC	All
*VLDL	V4R1M0
*WSCST	All

Table 64. Previous-release support by object types (continued)

Object type	Earliest target release
1 Integrated Language Environment (ILE) program (*PGM) objects are created by using the Create Program (CRTPGM) command. CRTPGM binds module (*MODULE) objects together. The *MODULE objects that are used to create an ILE program can be created by using different ILE compilers. The target release of the *PGM object is based on the target release of the *MODULE objects that were used to create the *PGM object. If the *MODULE objects' target release values are different, the most recent target release value is used. **2** Integrated Language Environment (ILE) service program (*SRVPGM) objects are created by using the Create Service Program (CRTSRVPGM) command. CRTSRVPGM binds module (*MODULE) objects together. The *MODULE objects that are used to create an ILE service program can be created by using different ILE compilers. The target release of the *SRVPGM object is based on the target release of the *MODULE objects that were used to create the *PGM object. If the *MODULE objects' target release values are different, the most recent target release value is used. **3** In V4R3M0, support was added for *STMF sizes up to 4 gigabytes - 1 byte. A *STMF larger than 2 gigabytes - 1 byte cannot be saved to a release before V4R3M0. In V4R4M0, support was added for *STMF sizes greater than 4 gigabytes - 1 byte. A *STMF larger than 4 gigabytes - 1 byte cannot be saved to releases earlier than V4R4M0. **4** If a journal receiver was attached to a journal when RCVSIZOPT(*MAXOPT1) was in effect, then it cannot be saved or restored to a release earlier than V4R5M0. Also, it cannot be replicated to any remote journals on any systems at release earlier than V4R5M0. If a journal receiver was attached to a journal when RCVSIZOPT(*MAXOPT2) was in effect, then it cannot be saved or restored to a release earlier than V5R1M0. Also, it cannot be replicated to any remote journals on any systems at releases earlier than V5R1M0. If a journal receiver was attached to a journal when any MINENTDTA options were in effect, then it cannot be saved or restored to a release earlier than V5R1M0. Also, it cannot be replicated to any remote journals on any system earlier than V5R1M0. If a journal receiver was attached to a journal when RCVSIZOPT(*MAXOPT3) was in effect, then it cannot be saved or restored to a release earlier than V5R3M0. Also, it cannot be replicated to any remote journals on any systems at releases earlier than V5R3M0. **5** V4R5M0 is the earliest release for *DTAQ if the SIZE and AUTORCL parameters on Create Data Queue (CRTDTAQ) did not contain the default values. Data queues cannot be saved for releases earlier than V5R4 if QDTA(*DTAQ) is specified. V6R1M0 is the earliest release if the data queue was changed to enforce data queue locks.	
6 V4R5M0 is the earliest release if *UBIN or *BIN 8 is specified for the format of a message description within the message file. V6R1M0 is the earliest release if *UTC, *UTCD, or *UTCT is specified for the format of a message description in a message file. **7** *USRPRF can be saved only to an earlier target release by using the Save Object List (QSRSAVO) application programming interface (API). **8** Spooled file data on output queues cannot be saved or restored for releases earlier than V5R4M0. **9** A *SAVF file that is larger than 1 terabyte cannot be saved to a release before V6R1. **10** Starting in V6R1M0, you can restore a *JOBQ object in an independent ASP.	

Related reference

Saving and restoring spooled files
You can use any of the methods that are described here to save and restore spooled files.

Related information

Save Object List (QSRSAVO) API

Testing the object on the current release

After the object has been created and saved using the target release (TGTRLS) parameter, it can be tested on the current-release system.

Thus, it is no longer necessary to support and maintain two development systems (one running the current release and one running the previous version). Testing this object should be like testing any other object. Make sure that all the objects that are to be used on the previous-release system have been saved using the TGTRLS parameter, restored onto the current-release system, and tested as a group on the current-release system.

Restoring and using the object on the previous release

After testing on the current-release system is completed, it is suggested that the object be distributed on a limited basis to previous-release systems or previous modification systems. Thus, if problems arise, they can be quickly corrected and contained with minimal impact to users.

Restrictions for current release to previous release support

These restrictions apply when you create and save objects on the current release, and then restore and use them on the previous release.

- The System/38 environment compilers (CL, COBOL/74, and RPG/III) do not support the TGTRLS parameter. Programs that are created using these compilers cannot be saved to, restored, or run on a previous-release system.
- The only way to save an object for a previous-release system is to use the TGTRLS parameter. If the TGTRLS parameter is not specified on the save command, and you attempt to restore the object on the previous-release system, the object is not restored.
- IBM does not support saving IBM-supplied objects (such as system commands, system programs, IBM spelling aid dictionaries, and so forth) from a current-release system and restoring and using them on a previous-release system or previous modification. As a result, the TGTRLS parameter is not supported on a SAVLIB command that specifies *ALLUSR, *IBM, or *NONSYS on the LIB parameter.
- IBM does not support new function from the current release to be used on a previous-release system or previous modification.
- When saving document library objects for a previous release, only folders and filed documents can be saved. Other items, such as mail or documents that are not filed cannot be saved using TGTRLS values other than *CURRENT.
- If a current-release program temporary fix (PTF) save file is sent from a current release-system to a previous-release system for distribution to another current-release system, object distribution must be used. The Copy PTF (CPYPTF) command or any save file command, such as DSPSAVF, cannot process the PTF save file.

Previous release to current release support

This topic describes considerations when you are moving specific types of information from an earlier release to the most current release.

Considerations when moving your data to the current release

Before you begin a data migration, see the Data migrations topic.

Generally, the system to which you are restoring objects must be at the same or a higher release level than the system from which the objects were saved, unless you specified a target release value when you saved. When moving data to a higher level release, you should only move user data. This might include user libraries, user directories, user profiles, user objects in IBM-supplied libraries, document library objects (DLOs), and mail. IBM-supplied libraries and IBM-supplied directories should not be restored to a higher release because these are handled during the licensed program install process. The target system should have the current-level release installed. This includes the Licensed Internal Code, IBM i operating system, IBM-supplied libraries QGPL and QUSRSYS, IBM i optional libraries, and any licensed programs that are purchased.

Alternatively, you can obtain preventive service planning (PSP) information from your software service provider. You should check the PSP before beginning your system upgrade or data migration. The PSP identifier for information that pertains to upgrading your system or migrating your data is SF98206. Information within this PSP describes known high-impact and pervasive problems related to upgrades and migrations. To receive this PSP information by using electronic customer support, type this command: `SNDPTFORD SF98206`.

Related information

Installing, upgrading, or deleting IBM i and related software
Preventative Service Planning - PSP

Considerations when moving system customization information

Some system customization information that is stored in the QSYS library cannot be saved.

This information includes network attributes, system values, the system reply list, and configuration information. You must manually re-create this information on your new or upgraded system. In addition, you will not be able to recover your problem log and your question and answer database. Use the procedure that is described in the Printing system information topic to print your current values.

When performing a disaster recovery or migrating to a new system, you need to maintain the authority of customized output queues, if they are on the system. When a printer device is created, an output queue by the same name as the device is also created. For example, a user might create a remote output queue and then create a virtual printer device to point to that output queue. If an output queue object being restored already exists on the system, the output queue object is overwritten by the object on the tape if no differences exist between the two versions. To ensure that the output queue object is restored from tape, specify ALWOBJDIF(*COMPATIBLE) on the Restore Library (RSTLIB) command to ensure that object differences, such as ownership, do not prevent the output queue from tape being restored over the output queue on the system.

Restoring previous release user data to a new system

The preferred method to restore previous release user data on your new, target system is to use the migration method.

The migration method asks you to first install the new, current release onto your old, source system. After which, you save your old system and then you perform a full system recovery onto the new, target system. Only use these instructions if it is not possible to perform the preferred migration process referenced in the Data migrations topic.

This section provides instructions to restore user data from a previous-release system to a later release system. Look for supported releases for software upgrades in the Installing, upgrading, or deleting IBM i and related software topic for your new, target release. This information explains for which releases you can use these instructions to restore your previous release user data to your new system.

The restore procedure involves two save steps, and four recovery steps. The save steps include printing your system information and completely backing up your old, source system.

The recovery steps on the new, target system involve the following steps:

1. Install Licensed Internal Code and IBM i on the target system using new release distribution media.

Note: If you order a new system (feature code 0205), the system arrives with the Licensed Internal Code and the base IBM i operating system preloaded on its disk units. Therefore, you can skip this step unless you are configuring a system with logical partitions. Then you will need to install the Licensed Internal Code and IBM i on each secondary logical partition. Allow at least two additional hours to load the Licensed Internal Code and base IBM i. If your system has more than Licensed Internal Code and IBM i on it you will also need to perform this step.

2. Restore system and user data to target system using the option 21 save of the source system.
3. Update the system information on the target system.
4. Install QGPL, QUSRSYS, Base options, and licensed program products (LPPs) using new release distribution media on target system.

   This converts restored source data to new target release.

Verify the prerequisites and proceed to the instructions to restore previous release user data to a new system.

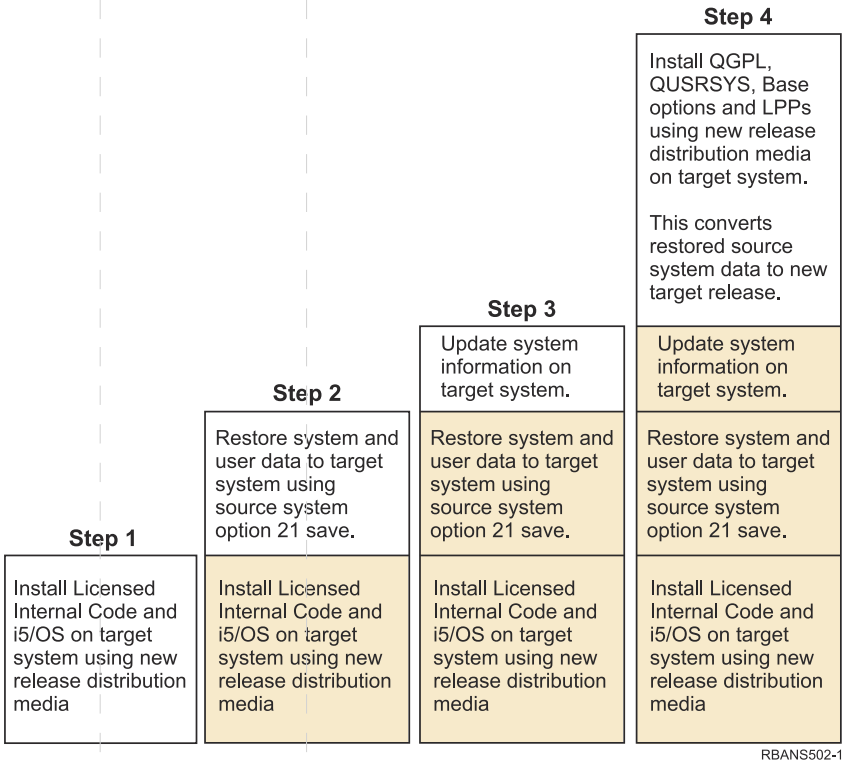

Figure 31. Recovery steps for restoring previous release user data to a new system

Prerequisites for the recovery

These instructions are sometimes used for a system upgrade when you install a replacement processor.

You must perform the following prerequisite tasks before you start the recovery portion of these instructions:

- Run the Retrieve System Information (RTVSYSINF) command on the source system. Some releases of IBM i do not support the RTVSYSINF command. When you run the RTVSYSINF command, the system asks you what library to use. Typically, specify the QUPGRADE library. If the QUPGRADE library does not exist, create it with the Create Library (CRTLIB) command.
- Print system information using the Print System Information (PRTSYSINF) command on the source system. Some releases of IBM i do not support the PRTSYSINF command.
- If necessary, save spooled files. For instructions on how to save spooled files, see the "Saving spooled files" on page 338 topic.

  Note: Job scheduler entries will not be restored. If necessary, make a note of your current job scheduler entries by using the Work with Job Schedule Entries (WRKJOBSCDE) command and selecting the print option. Then you can re-create the job scheduler entries manually on the new system.
- Run Save menu option 21 of the GO SAVE command on the source system. Make sure that you specify the following options:
  - `Vary off the network servers...*ALL`
  - `Unmount file systems.....Y`
- Make sure that you have a copy of the distribution media for the target system.
- If you want to restore data to logical partitions on your target system, make sure that you have already configured the logical partitions.
- If required, set up device parity protection and load source mirroring on the target system. If you have not done so on the new system, consult a IBM i specialist to configure your disk for your level of mirroring and parity protection before you continue with these instructions.

Restoring previous release user data to a new system

Complete these steps on the new (target) system to restore previous release user data to it. First, you install only the Licensed Internal Code and the IBM i operating system from the I_Base distribution media for the target release. Do not install any of the base options or licensed programs now.

Note: If you ordered a new system with feature code 0205, refer to the Installing, upgrading, or deleting IBM i and related software topic, and then go to step "31" on page 335.

Check off each item as you complete the task on the target system.

1. If you use the LAN Console, you might need to manually reset the service tools device ID password to match the IBM i device ID.

   After you begin the step, you must continue until you complete the initial program load (IPL) from Save System (SAVSYS) media. If you are using LAN Console, only the service tools device ID password must be reset on the system. The client attempts to sign on using the reset device ID password for LAN Console.
2. If you are using LAN Console, disconnect from your system and close Access Client Solutions (ACS). Restart ACS and reconnect by using the user ID of QSECOFR and password of QSECOFR to sign on to your system.
3. IPL the system from the first distribution media.
   a) Mount the first distribution media on the **`alternate IPL`** device. Wait for the READY status.
   b) At the central processing unit (CPU) control screen, place the system in MANUAL mode.
   c) Press the Function Select switch to display 02 (IPL) in the Function display.
   d) Press Enter.
   e) Press the Function Select switch to display D (IPL from tape or DVD) in the Data display.

f) Press Enter.

g) If the system is powered down, press the power button on the system to power on the system. Skip to step “6” on page 327. Otherwise, continue with step “3.h” on page 327.

h) If the system is powered on, press the Function Select switch to display 03 (continue the IPL) in the Function display.

i) Press Enter.

If you are using an Hardware Management Console (HMC) to do an IPL of the system from alternate media, complete the steps in rzarmprp.dita.

4. You see one of the following two displays. This shows the language on the distribution media.

   Press the Enter key to accept this language value.

```
                          Select a Language Group

 The language feature shown is the service tools
 language feature installed on the system.

 Attention: To keep the same language, ensure that the
 media you use for installing the operating system
 matches the language feature shown. If the operating
 system  media does not match what is shown, the
 installation process will attempt to install the
 operating system in a different language feature than
 Licensed Internal Code.

 Type choice, press Enter.

 Language feature . . . . . .  2924
```

```
LANGUAGE FEATURE ===>   2924
```

5. Press the Enter key to confirm the language group.
6. On the 'Install Licensed Internal Code (LIC)' display, select option 1 (Install Licensed Internal Code).
7. If a valid and functional load source disk unit was not detected on the partition, you are presented with the 'Select Load Source Device' display. Select the disk unit that you want to become the load source. If you do not see the 'Select Load Source Device' display, continue with step “9” on page 328.

```
                                 Select Load Source Device

Type 1 to select, press Enter.

                                   Sys  Sys   I/O       I/O
Opt  Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
     68-0CCDDA23      6718   050     12    11       0      0     7     0
     68-0DDA4212      6718   050     12    11       0      0     6     0
     68-0C211074      6718   050     12    11       0      0     3     0
     68-0DDE1201      6718   050     12    11       0      0     1     0
     68-0CDDEA44      6718   050     55    11       0      0     4     0
```

8. Press F10 to confirm the selection of the disk that you want to become the load-source disk unit.

```
                                Confirm Load Source Device

You have chosen the following device to be the load source.
Press F10 to confirm your choice.

                               Sys  Sys   I/O       I/O
Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
68-0CDDEA44      6718   050     55    11       0      0     4     0
```

9. Select option 2 (Install Licensed Internal Code and Initialize system) to start the installation of the Licensed Internal Code.

```
                     Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
   Serial Number    Type     Model      I/O Bus     Controller   Device
   xx-xxxxxxx       xxxx     xxx        x           x            x

Select one of the following:

   1. Restore Licensed Internal Code
 =>2. Install Licensed Internal Code and Initialize System
   3. Install Licensed Internal Code and Recover Configuration
   4. Install Licensed Internal Code and Restore Disk Unit Data
   5. Install Licensed Internal Code and Upgrade Load Source

Selection
2
```

10. On the 'Installed LIC and Initialize System - Confirmation' display, press F10 to confirm the initialization and to continue the installation.

```
                   Install LIC and Initialize System - Configuration

Warning:

   All data on this system will be destroyed and the Licensed
   Internal Code will be written to the selected disk if you
   choose to continue the initialize and install.

   Return to the install selection screen and choose one of the
   other options if you want to perform some type of recovery
   after the install of the Licensed Internal Code is complete.


   Press F10 to continue the install.
   Press F12 (Cancel) to return to the previous screen.
   Press F3 (Exit) to return to the install selection screen.
```

a. The 'Initialize the Disk - Status' display appears.

```
                             Initialize the Disk - Status

The load source disk is being initialized.

Estimated time to initialize in minutes:  xx
Elapsed time in minutes . . . . . .:  0.0
```

b. The 'Install Licensed Internal Code - Status' display appears.

```
                       Install Licensed Internal Code - Status


Install of the Licensed Internal Code in progress.

            +--------------------------------------+
Percent     |             XX %                     |
complete    +--------------------------------------+

Elapsed time in minutes . . . . . .:    x.x

Please wait.
```

11. You might see the 'Accept Console' display.

If it appears, press F10 to accept the current console. On the 'Accept And Set New Console Type On This' IPL display, press Enter to continue.

12. On the 'Disk Configuration Attention Report' display, press F10 to accept any problems and to continue.

```
                          Disk Configuration Attention Report

Type option, press Enter
   5=Display Detailed Report


Press F10 to accept all the following problems and continue.
The system will attempt to correct them.


OPT Problem
_ New disk configuration



F3=Exit                   F10=Accept the problems and continue
```

13. On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)).

```
                            IPL or Install the System
                                                                  System: XXXX
Select one of the following:

   1. Perform an IPL
   2. Install the operating system
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the operating system
   5. Save Licensed Internal Code


Selection
3
```

14. Sign on to DST as service tools user QSECOFR, with the password for the QSECOFR service tools user ID.

    The password is case-sensitive; use all capital letters.

```
                       Dedicated Service Tools (DST) Sign On

Type choices, press Enter.

Service tools user . . . . . . . . QSECOFR
Service tools password . . . . . . QSECOFR
```

    a. Select option 4 (Work with Disk Units).
    b. Select option 1 (Work with Disk Configuration).
    c. Select option 3 (Work with ASP Configuration).
    d. Select option 3 (Add Units to ASPs).

15. On the 'Add Units to ASPs' display, enter 3 to add units to existing ASPs, and add disk units to the system ASP.

```
                            Add Units to ASPs

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs
```

16. On the 'Specify New ASPs to Add Units to' display, enter the ASP number 1 to add disk units to the system ASP (1).

    a) Enter the ASP number on each selected unit.

```
                    Specify New ASPs to Add Units to

Specify the new ASP to add each unit to.

Specify  Serial                                    Resource
  ASP    Number          Type  Model  Capacity    Name
    1    21-6C597        4327   050      70564   DD007
    1    50-128840F      2107   A84      70564   DD004
    1    50-128940F      2107   A84      70564   DD005
         50-128A40F      2107   A85      35165   DD011
         50-128B40F      2107   A85      35165   DD003
         68-0C8BA12      6717   050       8589   DD008
         68-0C9D209      6717   050       8589   DD009
         68-606E0        6718   050      17548   DD006
         68-0CDAB10      6718   050      17548   DD010
         21-05348        4327   050      70564   DD012
         21-05322        4327   050      70564   DD013

F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

b) After you complete all units, press Enter.

c) If the list of units is correct, press Enter to start initializing the units.

17. On the 'Problem Report' display, press F10. Ignore the problems and continue.

```
                                Problem Report

   Note: Some action for the problems listed below might need to
   be taken. Please select a problem to display more detailed
   information about the problem and to see what possible
   action can be taken to correct the problem.


Type option, press Enter.
   5=Display Detailed Report


OPT Problem
   _ Unit possibly configured for Power PC AS

F3=Exit          F10=Ignore problems and continue              F12=Cancel
```

18. On the 'Confirm Add Units' display, press Enter to confirm the selected units.

```
                            Confirm Add Units

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data on units.
Press F12=Cancel to return and change your choice.

           Serial                   Resource
ASP  Unit  Number      Type Model  Name       Protection
  1                                           Unprotected
       1   68-0CDAB35  6718  050   DD001      Unprotected
       3   21-6C597    4327  050   DD007      Unprotected
       4   50-128840F  2107  A84   DD004      Unprotected
       5   50-128940F  2107  A84   DD005      Unprotected
       2   50-100240F  2107  A82   DD002      Unprotected
       6   50-128A40F  2107  A85   DD011      Unprotected

F9=Resulting Capacity              F10=Add and Balance
F11=Display Encryption Status      F12=Cancel
```

a. The 'Function Status' display shows the percentage of completion.

b. The `Selected units have been added successfully` message appears when the system completes the 'Add Units' process.

Note: If you require more than ASP 1 configured, you can return to step “15” on page 329 and create more ASPs by selecting 1 to create ASPs that are not encrypted. You can create multiple

user ASPs at the same time by specifying ASP numbers 2 through 32 on the selected disk units. You cannot create encrypted ASPs or independent ASPs now.

c. Press F12 to return to the 'Work with Disk Configuration' display.

d. If your system requires mirrored protection, continue to step “18.e” on page 331. If your system does not require mirrored protection, press F3 until you return to the 'Dedicated Service Tools (DST)' display, then select option 1 to exit the Dedicated Service Tools (DST).

e. To start mirrored protection for your system, follow these instructions:

i) On the 'Work with Disk Configuration' display, select option 4 (Work with Mirrored Protection).

ii) On the 'Work with Mirrored Protection' display, select option 2 (Start Mirrored Protection).

iii) Select an ASP by placing a 1 next to it. Press Enter to start mirrored protection.

iv) On the 'Confirm Continuation' display, press Enter.

v) On the 'Confirm Start Mirrored Protection' display, press Enter.

vi) The 'Function Status' display shows the completion status of the 'Start Mirrored Protection' request.

vii) The

```
Start mirroring completed successfully
```

message appears on the Disk Configuration Information Report display.

viii) Press Enter to continue.

f. If you use the Operations Console, follow these instructions to switch from the current console, or to verify the console mode value:

i) On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)). Press Enter to continue.

ii) Sign on to DST by using a service tools user profile that has security officer authority and the assigned password.

iii) On the 'Use Dedicated Service Tools (DST)' display, select option 5 (Work with DST environment). Press Enter to continue.

iv) On the 'Work with DST Environment' display, select option 2 (System Devices). Press Enter to continue.

v) On the 'Work with System Devices' display, select option 6 (Console Mode). Press Enter to continue.

vi) On the 'Select Console Type' display, select option 2 (Operations Console) or option 3 (Operations Console (LAN)). Press Enter to continue.

vii) If you select option 3 (Operations Console (LAN)), press F11 (Keywords) to configure the LAN console.

viii) Press F3 or F12 to get back to the 'IPL or Install the System' display.

19. Load the first volume of installation media that contains IBM i. (This volume is labeled B_GROUPx_01, where 'x' is the language group.)

    When you are installing from optical media, wait until the In Use indicator goes out before you continue.

20. On the 'IPL or Install the System' display, select option 2 (Install the Operating System).

```
                      IPL or Install the System

Select one of the following:
                                                                  System: XXXX
   1. Perform an IPL
   2. Install the Operating System
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the Operating System
   5. Save Licensed Internal Code

Selection
2
```

a. Make a selection on the 'Install Device Type Selection' display and press Enter.

   Note: Depending on the installation device type selected, more input might be needed.

b. On the 'Confirm Install of the Operating System' display, press Enter.

c. The 'Select a Language Group' display shows the primary language feature that is on your save media. Press Enter to accept this value.

```
                         Select a Language Group

Note:  The language feature shown is the language feature
installed on the system.

Type choice, press Enter.

Language feature . . . . . .  2924
```

d. On the 'Confirm Language Feature Selection' display, press Enter.

21. Make your selection on the 'Add All Disk Units to the System' display.

    If you have any unconfigured and unprotected disk units that support device parity protection, the 'Add All Disk Units to the System' display is shown as follows.

```
                        Add All Disk Units to the System
                                                                  System: XXXX
Non-configured device parity capable disk units are attached
to the system. Disk units cannot be added automatically.
It is more efficient to device parity protect these
units before adding them to the system.
These disk units might be parity enabled and added at SST.
Configured disk units must have parity enabled at DST.

Select one of the following:

1. Perform any disk configuration at SST
2. Perform disk configuration using DST

Selection
```

Note: This display does not appear if you selected all disk units that are known to the system on step “16” on page 329.

The 'Add All Disk Units to the System' display might also look as follows.

```
                        Add All Disk Units to the System

Select one of the following:
                                                                  System: XXXX
1. Keep the current disk configuration
2. Perform disk configuration using DST
3. Add all units to the system auxiliary storage pool
4. Add all units to the system ASP and balance data

Selection
1
```

22. The 'Licensed Internal Code IPL in Progress' display shows the IPL progress.

```
                        Licensed Internal Code IPL in Progress

IPL:
   Type . . . . . . . . . . . . . . . :    Attended
   Start date and time  . . . . . . . :    xx/xx/xx  xx:xx:xx
   Previous system end  . . . . . . . :    Abnormal
   Current step / total . . . . . . . :       16      16
   Reference code detail  . . . . . . :    C6004065

IPL step                                   Time Elapsed   Time Remaining
  Commit Recovery                            xx:xx:xx        xx:xx:xx
  Data Base Initialization                   xx:xx:xx        xx:xx:xx
  Journal IPL Clean up                       xx:xx:xx        xx:xx:xx
  Commit Initialization                      xx:xx:xx        xx:xx:xx
 >Start the operating system                 xx:xx:xx        xx:xx:xx

Item:
  Current / Total . . . . . . :

Sub Item:
  Identifier . . . . . . . . . . :    Searching tape for install media
  Current / Total  . . . . . . . :

Wait for next display or press F16 for DST main menu
```

23. On the 'Install the Operating System' display, select option 1 (Take defaults).

    Make sure that the values for Date and Time are correct. Press Enter to continue.

```
                             Install the Operating System

Type options, press Enter.

Install
   option . . . . . 1      1=Take defaults (No other
                             options are displayed)
                           2=Change install options

Date
   Year . . . . . . xx     00-99
   Month. . . . . . xx     01-12
   Day. . . . . . . xx     01-31

Time
   Hour . . . . . . xx     00-23
   Minute . . . . . xx     00-59
   Second . . . . . xx     00-59
```

24. The 'IBM i Installation Status' display shows the installation status of the IBM i installation profiles and libraries.

```
Message ID . . : CPI2070                                  IBM i Installation Status

        +-------------------------------------------------------------------------------+
Stage 1 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-------------------------------------------------------------------------------+
        0              20              40             60             80              100

Installation                                                                    Objects
Stage                                                          Completed       Restored

1  Creating needed profiles and libraries. . . :
2  Restoring programs to library QSYS  . . . . :
3  Restoring language objects to library QSYS. :
4  Updating program table. . . . . . . . . . . :
5  Installing database files . . . . . . . . . :
6  Installing base directory objects . . . . . :
```

25. The system installs the remaining IBM i objects.

```
Message ID . . : CPI2070                                    IBM i Installation Status

        +-----------------------------------------------------------------------------------+
Stage 4 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-----------------------------------------------------------------------------------+
        0              20              40              60              80              100

Installation                                                    Objects
Stage                                              Completed    Restored

   2  Restoring programs to library QSYS  . . . . :      x          09
   3  Restoring language objects to library QSYS. :      x          03
>> 4  Updating program table  . . . . . . . . . :      x
   5  Installing database files . . . . . . . . . :      x
   6  Installing base directory objects . . . . . :
   7  Completing installation . . . . . . . . . . :
```

Note: When the operating system is being installed, you might get a message to load the next media volume.

26. On the 'Sign On' display, logon as user QSECOFR. You do not need to enter a password now.
27. On the 'IPL Options' display, enter the correct values for the system date, time, and time zone. Only the following options are set to Y:
    - Start system to restricted state
    - Set major system options
    - Define or change the system at IPL

```
                                  IPL Options

Type choices, press Enter.

   System date  . . . . . . . . . . . . . . . xx/xx/xx   MM / DD / YY
   System time  . . . . . . . . . . . . . . . xx:xx:xx   HH : MM : SS
   System time zone . . . . . . . . . . . . .  Q0000UTC   F4 for list
   Clear job queues . . . . . . . . . . . . . N          Y=Yes, N=No
   Clear output queues  . . . . . . . . . . . N          Y=Yes, N=No
   Clear incomplete job logs  . . . . . . . . N          Y=Yes, N=No
   Start print writers  . . . . . . . . . . . N          Y=Yes, N=No
   Start system to restricted state . . . . . Y          Y=Yes, N=No

   Set major system options . . . . . . . . . Y          Y=Yes, N=No
   Define or change system at IPL . . . . . . Y          Y=Yes, N=No
```

On the 'Set Major System Options' display, select **N** to not enable automatic configuration.

```
                        Set Major System Options

Type choices, press Enter.

Enable automatic configuration . . . . . . N            Y=Yes, N=No
Device configuration naming  . . . . . . . *NORMAL      *NORMAL, *S36,
                                                        *DEVADR
Default special environment. . . . . . . . *NONE        *NONE, *S36
```

28. The 'Define or Change the System at IPL' display appears.
    a) Select option 3 (System value commands).
    b) On the 'System Value Commands' display, select option 3 (Work with System Values).
    c) On the 'Work with System Values' display, select the system value that you plan to change by placing a 2 next to it.

       Press Enter only after you select all the values.

       Update the following system values. Write down the existing values so you can update them after the recovery, if necessary.

       - QALWOBJRST to *ALL

- QFRCCVNRST to 0
- QIPLTYPE to 2
- QINACTITV to *NONE
- QJOBMSGQFL to *PRTWRAP
- QJOBMSGQMX to a minimum of 30 with a suggestion of 64
- QLMTDEVSSN to 0
- QLMTSECOFR to 0
- QLOGOUTPUT to *JOBLOGSVR
- QMAXSIGN to *NOMAX
- QPFRADJ to 2
- QPWDEXPITV to *NOMAX
- QSCANFSCTL add *NOPOSTRST
- QVFYOBJRST to 1

d) After the system changes the system values, press **F3** twice to return to the 'Define or Change the System at IPL' display.

e) On the 'Define or Change the System at IPL' display, press **F3** to exit, and continue the IPL.

29. On the 'Change Password' display, type QSECOFR as the current password. Enter a new password. Reenter the password to verify and press Enter. (New password cannot be QSECOFR.)

```
                                Change Password

Password last changed . . . . . xx/xx/xx

Type choices, press Enter.

   Current password . . . . . . . . . QSECOFR
   New password . . . . . . . . . . . _______
   New password (to verify) . . . . . _______
```

30. On the 'Work with Software Agreements' display, type **5** to display the software agreements for *MCHCOD (which includes Licensed Internal Code) and the IBM i Operating System (5770-SS1). Read and accept these agreements. On the 'Confirm Acceptance of Software Agreements' display, press Enter.
31. Because you did not select automatic configuration, you must manually configure a tape drive to complete the following restore step for your user data and related system data.

    To configure tape units that are not models 34*xx*, use the following instructions:

    a) Use the Work with Hardware Resource (WRKHDWRSC) command to determine the tape controller name: `WRKHDWRSC TYPE(*STG)`

    b) Locate the tape controller on the 'Work with Storage Resources' display.

    c) Type a 9 (Work with resource) next to the tape controller name and press the Enter key. If the tape controller is not listed on the display, you must select other resources, such as disk storage controllers. For some system models, tape units are now attached through combined-function IOPs. Look through the resources until you find the tape unit you want.

    d) Locate the resource name for the tape unit (usually it is TAP01).

    e) Enter a 5 (Work with Configuration Descriptions) in the **Opt** column next to the tape resource name and press the Enter key. You are shown the 'Work with Configuration Descriptions' display.

    f) Type a 1 (Create) in the **Opt** field and a device description name for a tape drive, such as TAP01, in the **Description** field. Press the Enter key. You are shown the 'Create Device Description (Tape)' display.

    g) Change any values that you want to change, then press the Enter key (twice) to create the device description. You are shown the 'Work with Configuration Descriptions' display again. The device that you created is displayed.

h) Type an 8 (Work with configuration status) in front of the new device description. You are shown the 'Work with Configuration Status' display.

i) Type a 1 (Vary on or Make available) in front of the new device. If the status does not change to Varied On or Available, wait a few minutes. Then press F5 (Refresh). If the status still does not change to Varied On or Available, follow normal problem analysis for the device.

j) Press F3 until you return to the main menu.

```
                                IBM i Main Menu

Select one of the following:

   1. User tasks
   2. Office tasks
   3. General system tasks
   4. Files, libraries, and folders
   5. Programming
   6. Communications
   7. Define or change the system
   8. Problem handling
   9. Display a menu
  10. Information Assistant options
  11. IBM i Access tasks

  90. Sign off

Selection or command
=>
```

32. Use the save media from option 21 (save from the source system) to complete the following steps that restore the user data and related system data to the target system.

a) Type the **End Subsystem (ENDSBS)** command as follows.

```
ENDSBS SBS(*ALL) OPTION(*IMMED)
```

b) Prevent messages that are not necessary by typing the **Change Message Queue (CHGMSGQ)** command as follows.

```
CHGMSGQ MSGQ(QSYSOPR) DLVRY(*NOTIFY) SEV(99)
```

c) Type the **Restore User Profiles (RSTUSRPRF)** command as follows.

```
RSTUSRPRF DEV(media-device-name) ALWOBJDIF(*ALL) ENDOPT(*LEAVE)
```

d) To restore your output queue information for printers, type the **Restore Object (RSTOBJ)** command as follows.

```
RSTOBJ OBJ(*ALL) SAVLIB(QUSRSYS) DEV(media-device-name) OBJTYPE(*OUTQ)
       ENDOPT(*REWIND) MBROPT(*ALL) ALWOBJDIF(*ALL)
```

e) Type the Restore Configuration (RSTCFG) command as follows.

```
RSTCFG OBJ(*ALL) DEV(media-device-name) OBJTYPE(*ALL)
       SRM(*NONE) ALWOBJDIF(*ALL) ENDOPT(*LEAVE)
```

Note: If the saved configuration does not match the new system, consider skipping this step.

f) Type the Restore Library (RSTLIB) command as follows.

```
RSTLIB SAVLIB(*NONSYS) DEV(media-device-name) OPTION(*NEW)
       ALWOBJDIF(*ALL) MBROPT(*ALL) FRCOBJCVN(*NO) ENDOPT(*LEAVE) OMITOBJ(QOFC)
```

Note: This command does not restore the QAUDJRN and QACGJRN objects and any job scheduler entries.

g) Type the Restore Document Library Object (RSTDLO) command as follows.

```
RSTDLO DLO(*ALL) DEV(media-device-name) SAVFLR(*ANY)
       ALWOBJDIF(*ALL) ENDOPT(*LEAVE)
```

Note: If you have DLOs in any of your user ASPs, use the following command to restore DLOs to each user ASP.

```
RSTDLO DLO(*ALL) ALWOBJDIF(*ALL) SAVASP(ASP-number) RSTASP(ASP-number)
```

h) Type the Restore Object (RST) command as follows.

```
RST DEV('/QSYS.LIB/media-device-name.DEVD') OBJ(('/*')
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) ALWOBJDIF(*ALL) OPTION(*NEW)
```

i) If you copied spooled files to database files before you saved them on your source system, follow these steps to restore the spooled files:

Note: Before release V5R4, to save spooled files, you must copy the spooled files to database files first. For release V5R4 and later, you can save spooled files directly. Follow these steps only if you saved spooled files with the older method.

i) In the rzarmsavespooledfiles.dita topic, you saved spooled files to database files in a library. If that library was not restored to your upgraded system, restore it now by using the **RSTLIB** command.

Note: Use the **RSTLIB** command only if you used the **Save Library (SAVLIB)** command to save the objects. If you used the **Save Object (SAVOBJ)** command, you must use the **Restore Object (RSTOBJ)** command.

ii) For each spooled file that you need to restore, follow these steps:

a) On the printout that you used when you saved the spooled files, locate the name of the printer file that was used to create the spooled file. It appears in the File column on the left side of the printout.

b) To override the control character value for the QPRINTS file, type the following command: `OVRPRTF FILE(QPRINTS) CTLCHAR(*FCFC)`

c) Copy the database file to the printer file by typing the following command: `CPYF FROMFILE(LIBSPLF/SPOOLF1) TOFILE(QGPL/QPRINTS) MBROPT(*REPLACE)`

In this example, a spooled file that is named QPRINTS is placed on the output queue that is associated with the definition for the QPRINTS printer file.

d) Delete the database file by using the Delete File (DLTF) command.

iii) After you restore all spooled files, type the following to delete the override for the QPRINTS file. `DLTOVR FILE(QPRINTS)`

j) It is suggested that you use the **Start Object Conversion (STROBJCVN)** command to convert user objects from the format of the previous version, release, and modification level of an IBM-supported operating system to the current operating system format.

Note: It might take a long time to convert user objects. You can complete this step now or you can complete this step later.

33. If you used the **Retrieve System Information (RTVSYSINF)** command on the source system, run the **Update System Information (UPDSYSINF)** command now to update the system information on the target system.

    If your source system does not support the **RTVSYSINF** command, use the printed system information to update all system information, such as system values and message reply lists entries, on the target system.

    The **RTVSYSINF** command does not update all system information. Use the printed system information to update system information such as the time it takes to recover access paths, subsystem descriptions, and RJE configurations.
34. Install the base options (including QGPL and QUSRSYS) and other licensed programs by using the distribution media for the target system (current release) and the GO LICPGM command.

To verify that your optical installation is varied on, begin with the Verifying completion of manual installation step. Do not continue with the Checklist: Completing the upgrade or replacement until instructed to do so.

35. Run the Restore Authority (RSTAUT) command.
36. If necessary, change the following system values back to their original values by using the Work with System Value (WRKSYSVAL) command.
    - QALWOBJRST
    - QIPLTYPE
    - QJOBMSGQFL
    - QJOBMSGQMX
    - QPFRADJ
    - QVFYOBJRST
37. If you do not know the password for the restored QSECOFR profile, change the password before you sign off.

    Type the following command:

    ```
    CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)
    ```

38. Type the command: DSPJOBLOG *PRINT.

    Check the job log to verify that the job restored all objects. To verify whether the job restored all objects, spool the job log for printing along with any remaining spooled job output. Check for any error messages. Correct the errors and restore those objects from media. If you are using this checklist as part of a data migration, return to the Completing the migration topic.
39. Complete the installation.

    Follow the instructions in Checklist: Completing the IBM i installation When you install the cumulative program temporary fix (PTF) package, accept the default option to do an IPL of the system.
40. Add job scheduler entries with the Add Job Schedule Entry (ADDJOBSCDE) command by using the information that you printed from your source system.

    Use the Work with Job Schedule Entries (WRKJOBSCDE) command and select the print option.
41. In a previous step #rzarmrstprerlsuserdatanewsys/userdata, library QUSRSYS was restored. If this library was saved from IBM i 7.2 or a later release, it includes the TCP/IP Configuration information at the time of the save of QUSRSYS. Type the following command to restore the TCP/IP configuration to the system now.

    ```
    UPDTCPINF LIB(QUSRSYS) TYPE(*ALL)
    ```

Related information

Console Advanced Topics

Service tools user IDs and passwords

Saving spooled files

When you save a library that contains output queues, you save the descriptions of the output queues but not their contents.

Note: Use this method if you are saving spooled files on a release earlier than V5R4. For release V5R4 and later, you can save the output queues with their contents by specifying SPLFDTA(*ALL) when you save the output queues.

If you have critical spooled files that you will need after the upgrade procedure, you can use the following procedure to save them:

1. Create a library to hold copies of the critical spooled files by using the Create Library (CRTLIB) command.

For example, to create a library called LIBSPLF to hold your spooled files, type the following command: `CRTLIB LIB(LIBSPLF)`

Note: If the spooled files contain confidential information, specify `AUT(*EXCLUDE)` on the CRTLIB command.

2. Use the Work with Output Queue command to locate the spooled files that you want to save from a designated output queue: `WRKOUTQ OUTQ(lib-name/queue-name) OUTPUT(*PRINT)`

   Note: If you do not have special output queues that are designated for critical spooled files, type `WRKOUTQ OUTQ(*ALL)`
3. Print out and retrieve the listing of the spooled files that you need to save.
4. On the printout, mark the spooled files that you need to save.
5. For each spooled file, follow these steps:
   a) Choose a name (8 characters or less) for the spooled file that will help you to identify it. Each file should have a unique name.
   b) Create a database file to hold the contents of the spooled file by typing the following command: `CRTPF FILE(LIBSPLF/file-name) RCDLEN(133)`

      Notes:

      - For file name, substitute the name that you assigned in step 5a.
      - Use an appropriate record length for the spooled file that you are copying. The record length must be at least 1 character larger than the spooled data to allow for the control character.
      - If you are copying a large spooled file, specify SIZE(*NOMAX) when you create the database file.

   c) Copy the contents of the spooled file to the database file that you created by typing the following command: `CPYSPLF FILE(spooled-file) TOFILE(LIBSPLF/file-name) JOB(job-number/user-name/job-name) SPLNBR(nnn) CTLCHAR(*FCFC)`

      Note: For spooled file, substitute the value from the File column on the listing that you created in step 2.

   d) You might receive message CPA3312 if the spooled file contains special attributes. Respond with G (GO) to continue saving the contents of the spooled file.
6. Repeat step 5, steps 5a through step 5d, for each spooled file that you need to save.
7. If you have additional output queues to process, return to step 2.
8. Use the Save Library (SAVLIB) command to save the library that contains the copies of your spooled files.

Restrictions when going from previous release to current release

Previous releases of IBM i might support products that are not supported at the current release.

For a list of products that will no longer work with the current release, see the Support for IBM System i Web site (www.ibm.com/support/pages/node/668131). To avoid losing data that depends on these products, migrate that data from your system to an accessible location before upgrading to the current release.

Related concepts

Release-to-release support
The release-to-release support on the IBM i environments allows you to move data from the current release to a previous-release system. This support also allows you to move data from a previous-release system to a current-release system.

Chapter 18. System synchronization: Planning and procedures

System synchronization is the second part of moving data between two systems.

You need to perform system synchronization if one of the following situations occurs:

- You have purchased a new system.
- You have moved all of your data to the new system.
- Or one of the following situations applies:
  - The existing system is a production system that needs to stay up and running while you convert to the new system.
  - You want to test the new system before making the switch over.

The existing system is a production system, so changes occur to the existing system do not get applied to the new system. Therefore after you have loaded the new system, you need to synchronize the new system and the existing system before you can put the new system into production.

You can perform system synchronization only if one the following conditions is true:

- The new system and the existing system are at the same release. You have fully reloaded the new system from the existing system, using "Checklist 20: Recovering your entire system after a complete system loss" on page 93.
- The new system is at a newer release than the existing system. You have fully reloaded the new system from the source system using one of the procedures in the Chapter 17, "Release-to-release support," on page 313 topic.

The method that you use to synchronize the two systems is the side-by-side method. The underlying principal of the side-by-side method is that you will run your existing system and your new system in parallel for a test period. During that test period, you will periodically perform activities to synchronize the data on your new system with the data on your existing system. At the end of the test period, you will perform final synchronization activities before moving your production work to your new system. When you complete your final synchronization, the software environment on the two systems should be identical.

The topics that follow discuss several different approaches that you can take for performing synchronization. In all cases, synchronization requires careful planning and monitoring. It also requires a good understanding of your applications and the library structure on your system. Running two systems in parallel also requires strong change-control practices. This chapter focuses primarily on synchronizing data.

If possible, during the synchronization period you should carefully limit other changes on your existing system, such as adding or changing user profiles or changing the system distribution directory. When this type of change to system customization occurs on your existing system, you need to manually perform the same updates on your new system.

You might find the security auditing function helpful for keeping track of changes to system information on your existing system. If you are not familiar with security auditing, see Auditing security on IBM i, which describes how to set up security auditing and which values to choose to get the entries that you need.

You can print the entries in the audit journal receiver and enter the same information on your new system.

© Copyright IBM Corp. 1997, 2025

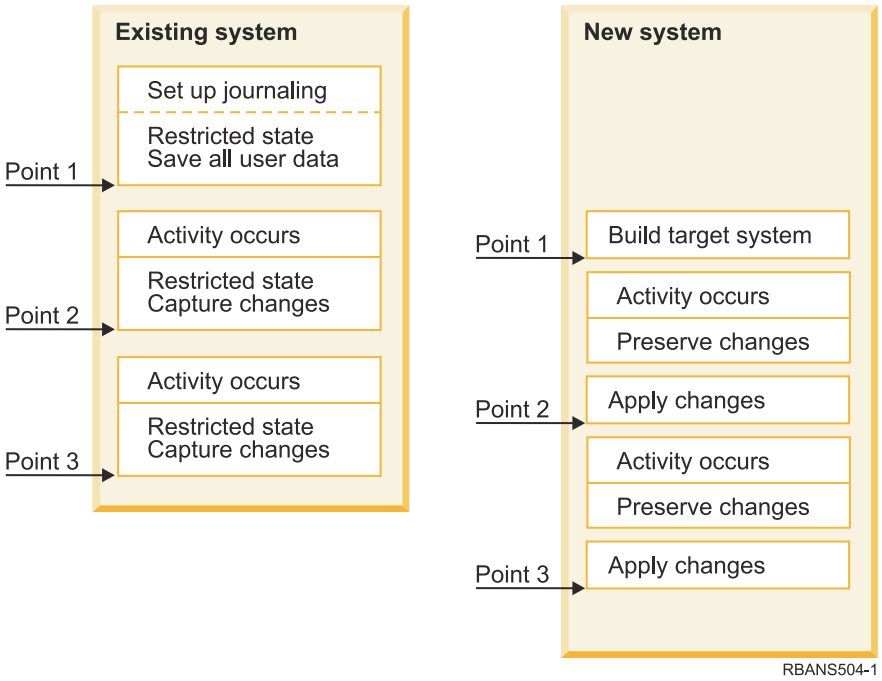

Figure 32. Overview of synchronization process

Figure 32 on page 342 provides an overview of the synchronization process. The process starts with building an initial test environment that matches a known point on your existing system (Point 1). Periodically, you establish a new known point (synchronization point) on your existing system. Point 2 and Point 3 are examples of synchronization points. After preserving any work that you have performed on your new system, such as object conversion, you use one of the synchronization methods to bring your new system to the same level as your existing system.

While you read and plan, consider how the options for synchronization relate to your current procedures (such as regular save procedures and change control procedures). By using your existing procedures as a starting point, you can reduce the level of disruption and build on your existing base of knowledge. For example, if you currently use object journaling, then object journaling might be a logical part of your synchronization strategy. If no one in your organization has experience with object journaling, then it might not be the best solution for you.

Start with a valid test environment

All of the synchronization methods in this topic collection assume that you start by loading an entire copy of your existing system to your new system. To create this initial test environment, you must follow the sequence in your reload checklist to ensure that the system correctly merges your user data with licensed program data. The reload checklist is “Checklist 20: Recovering your entire system after a complete system loss” on page 93 if you are restoring to the same release, or the procedures in Chapter 17, “Release-to-release support,” on page 313 if you are restoring to a higher release.

Synchronization methods: Overview

Here are the common methods for synchronizing test and existing systems.

Moving changed objects
With this method, you periodically save everything on your existing system that has changed since your last synchronization point. You then restore these changed objects to the new system.

Moving libraries
With this method, you periodically copy entire libraries from your existing system to your new system. This method works best when your programs are in separate libraries from your database files. You synchronize the libraries that contain database files.

Moving individual objects
With this method, you periodically copy specific objects, such as database files, from your existing system to your new system.

Applying journaled changes
With this method, you journal objects on your existing system. You move the journal receivers from your existing system to your new system. You apply the journaled changes to the test objects on your new system. This method is sometimes used in conjunction with moving changed objects.

Refreshing new system
With this method, you periodically refresh your new system by restoring an entire copy of your existing system.

Table 65 on page 343 provides a comparison of these synchronization methods. It also shows where you can read more about each method. You can use these methods individually or in combination.

Table 65. Comparison of synchronization methods

Method	Complexity	Risk	Time required for synchronization	Frequency used	Where to read more about it
Moving changed objects	High	Medium to high	Medium	High	"Moving changed objects" on page 343
Moving libraries	Low	Low to medium	Medium	Medium to high	"Moving entire libraries" on page 348
Moving objects	Medium to high	Low to medium	Low to medium	Medium	"Moving individual objects" on page 349
Applying journaled changes	High	High	Medium [1]	Low	"Applying journaled changes" on page 350
Refreshing new system	Low	Low	High	Low	"Refreshing your new system" on page 352

1 The time required to apply journaled changed depends on how many journal entries need to be applied, and how many objects you need to apply those entries to. The greater these two numbers are, the longer synchronization time is.

Moving changed objects

With this method, you periodically save everything that has changed since your last synchronization point.

You then restore those changed objects to your new system. The suggested method when you save changed objects is to specify an exact reference date and time that corresponds to your last synchronization point. This ensures that the contents of your save tapes are not affected by any intermediate save operations that might have occurred since your last synchronization point.

Here is an example of the saving and restoring procedures when you use this method. You will need to change these sample steps to meet the needs of your own situation. This example assumes that the last synchronization point was at 1800 hours (6 p.m.) on July 27, 1998.

Saving changed objects

Use these steps to save the changed objects on your existing system.

1. To avoid any problems with inadequate authority, sign on as the security officer (QSECOFR).
2. Place your system in a restricted state to ensure that you get a stable copy of the changed objects on your existing system.
3. Use the Save Security Data (SAVSECDTA) command to save all user profiles.

   You use this information to correctly synchronize ownership and authority for any new objects that you move.
4. To save objects that have changed since your last synchronization point, use the Save Changed Object (SAVCHGOBJ) command.

This command example saves objects in libraries (the QSYS.LIB file system):

```
SAVCHGOBJ OBJ(*ALL) LIB(*ALLUSR) DEV(tape-device)
   OBJJRN(*YES) REFDATE('07/27/xx') REFTIME(180000)
   ACCPTH(*YES)
```

Note: If you are using the SAVCHGOBJ method in conjunction with applying journaled changes, specify OBJJRN(*NO).

5. If you have user libraries whose names begin with Q, save the changed objects in those libraries.

   If you have user libraries whose names begin with Q, save the changed objects in those libraries. Repeat step 4 and substitute the name of your Q library in place of *ALLUSR.

   Note: The online information for the LIB parameter tells which Q libraries are included when you specify *ALLUSR.
6. To save document library objects that have changed since your last synchronization point, use the Save Document Library Object (SAVDLO) command:

   ```
   SAVDLO DLO(*SEARCH) DEV(tape-device)
      REFCHGDATE('07/27/xx') REFCHGTIME(180000)
      SRCHTYPE(*ALL) OWNER(*ALL)
   ```
7. To save mail, use the SAVDLO command as follows:

   ```
   SAVDLO DLO(*MAIL)
   ```

   You cannot only save changed mail. You must save all mail, if necessary.
8. To save objects in directories that have changed since your last synchronization point, follow these steps:

   a. If you have any network server descriptions (NWSDs), you must vary them off before starting the save procedure. Use the command WRKCFGSTS CFGTYPE(*NWS) (Work with Configuration Status) to display the configured NWSDs on your system.

      Select option 2 (Vary off) on this display to vary off the NWSDs.

      Note: Alternatively, use the Vary Configuration command to vary off a NWSD:

      ```
      VRYCFG CFGOBJ(XXX) CFGTYPE(*NWS) STATUS(*OFF)
      ```

   b. Use the Save Object (SAV) command to save changed objects:

      ```
      SAV DEV('/QSYS.LIB/tape-device.DEVD')
        OBJ(('/*' *INCLUDE) ('/QSYS.LIB' *OMIT)
        ('/QDLS' *OMIT)) CHGPERIOD('07/27/xx' 180000)
        UPDHST(*YES)
      ```
9. To vary on the network servers, use the Work with Network Server Status (WRKNWSSTS) command and select option 14.

   Use the WRKCFGSTS *NWS command to display all network servers and select option 1 to vary on any that were previously varied off.
10. Use the Start Subsystem (STRSBS) command to restart the controlling subsystem to return your system to productive use.
11. To display the log of what changed objects the system saved, use the Display Job Log (DSPJOBLOG) command.
12. Print the job log and highlight each library that was saved.

    You will need this information to restore changed objects.

Restoring changed objects

Use these steps to restore changed objects on your existing system.

Testing objects to preserve

This example assumes that you are not making permanent changes to volatile objects (such as database files) on your new system. When you restore, you will be overlaying test objects.

It also assumes that after you build your initial new system, you will not restore programs from the existing system to the new system during synchronization (because those programs are already converted on your new system).

If you need to preserve test objects or if programs change on your existing system, you need to make special plans for your restore procedures.

To restore the changed objects that you saved, perform these steps on your test system:

1. To avoid any problems with inadequate authority, sign on as the security officer (QSECOFR).
2. Place your system in a restricted state.
3. To restore the saved user profiles, use the Restore User Profile (RSTUSRPRF) command:

   ```
   RSTUSRPRF USRPRF(*ALL) DEV(tape-device)
       ENDOPT(*LEAVE)
   ```

4. Locate the printout of the job log from your save operation.

   Use it to determine which libraries the system saved. If you do not have the job log, you can use the Display Tape (DSPTAP) command to display the contents of the save tapes:

   ```
   DSPTAP DATA(*SAVRST) OUTPUT(*PRINT)
   ```

5. Use one of the following methods to restore objects from the Save Changed Objects tape. The first method is easier and quicker:

   - For all libraries on the save tapes, type the following command:

     ```
     RSTOBJ(*ALL) DEV(tape-device) SAVLIB(*ANY)
     OMITLIB(QGPL QUSRSYS) ENDOPT(*REWIND) MBROPT(*ALL)
     ```

     If you have changes in the QGPL and QUSRSYS libraries, type the following command for QGPL and QUSRSYS:

     ```
     RSTOBJ OBJ(*ALL) DEV(tape-device) OBJTYPE(*ALL)
        SAVLIB(library-name) ENDOPT(*LEAVE) MBROPT(*NEW)
     ```

   - For each library on the save tapes, type the following command:

     ```
     RSTOBJ OBJ(*ALL) DEV(tape-device) OBJTYPE(*ALL)
        SAVLIB(library-name) ENDOPT(*LEAVE) MBROPT(*ALL)
     ```

   Notes:

   - For both the QGPL library and the QUSRSYS library, you must specify `MBROPT(*NEW)`. This prevents overlaying new system files with old system files.
   - Do not specify `ALWOBJDIF(*ALL)`. Normal restore processing with the default ALWOBJDIF value protects you from accidentally overlaying critical information. ALWOBJDIF(*ALL) is intended only when you are initially loading information from one system to another.
   - If your new system has a different auxiliary storage pool (ASP) configuration from your existing system, you might need to specify the SAVASP and RSTASP parameters.

6. To restore document library objects that you saved in step 6 in the Steps for saving changed objects topic, type the following command:

   ```
   RSTDLO DLO(*ALL) DEV(tape-device) ENDOPT(*LEAVE)
   ```

   Notes:

- Do not use this restore command unless your save tapes contain only changed document library objects (DLOs). If you restore all DLOs from your existing system, you will overlay IBM-supplied objects that are used for IBM i Access.
- This command does not restore mail that has changed. Mail gets restored in step 8.
- Changes to calendars are restored when you restore the QUSRSYS library.
- If you have DLOs in more than one ASP, you need to run the Restore Document Library Object (RSTDLO) command for each ASP. You specify the SAVASP and RSTASP parameters.

7. To restore mail that you saved in step 7 in the Steps for saving changed objects topic, use the RSTDLO command as follows:

```
RSTDLO DLO(*MAIL)
```

8. To restore the changed directory objects that you saved in step 8b in the Steps for saving changed objects topic, type the following command:

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

9. Use the Display Job Log to print your job log:

```
DSPJOBLOG OUTPUT(*PRINT)
```

   Review it carefully. Whenever you restore changed objects, you are likely to encounter situations that you will need to recover manually. If you plan to synchronize your system several times, you might find it useful to create a log that describes synchronization problems and their resolutions. This will help to reduce your synchronization time in the future.

10. After you resolve any problems that occurred with restored objects, use the Restore Authority (RSTAUT) command to restore private authorities.

    Note: You should wait to restore authority until after you resolve any problems because some problem resolution steps might require you to restore new objects.

11. Restart the controlling subsystem and make the system available for more testing.

Problems when restoring changed objects

Because you specify ALWOBJDIF(*NONE) when you restore changed objects, the system compares heading information on the system copy of an object with heading information on the media copy.

When the system detects discrepancies that might indicate mismatches, the system does not restore an object.

Related information

Restoring changed objects

Problems restoring journal receivers

If you are journaling objects on both your existing system and your new system, you might have cases where two different journal receivers exist with the same name.

Typically the coexistence of two different journal receivers occurs because you change journal receivers on both systems. The systems automatically generate the same next-receiver name.

In some environments, you might not need the information that is in the journal receivers on your new system. Use the Change Journal (CHGJRN) command to create and attach a new journal receiver with a unique name. Then you can save and delete the journal receivers that you do not need (on your new system).

Note: This strategy applies when you are using a change-object synchronization method. If you plan to apply journaled changes to synchronize systems, you need to devise a method for naming and changing journal receivers that allows you to successfully restore journal receivers.

Related information

Journal management

Problems restoring database file members

When you delete and re-create a database file, that database file has a different file level identifier than the previously saved version.

Therefore, the file level identifier for the database file on your save media that you attempt to restore will not match the file level identifier on the new database file. When this type of mismatch occurs, the system will not restore the database file using the default values on the restore command parameters. The same problem occurs when you delete and re-create individual members for a database file, causing the member level identifiers of a file to mismatch.

Here are several options for handling this problem. Choose the correct option for your situation. Base your decision on your synchronization requirements and your application architecture. Always make sure that you have a good backup of your new system.

Recovery option 1–Allowing file level identifier differences on the restore operation: If you specify *FILELVL on the ALWOBJDIF parameter of the restore command, the file and/or member level identifier of a physical file can be different. If a database physical file already exists on the system with a different file, or a member level identifier, or both than the saved physical file, then an attempt will be made to restore the physical file data. The file, or the member level identifiers of the physical file, or both, on the system remain unchanged. Objects other than physical files will be restored as if ALWOBJDIF(*NONE) is specified.

Recovery option 2–Restoring entire library: A simple solution is to restore the entire library from your existing system to your new system. To do this, you will need to first clear the library on your new system. To use this option, you might need to change your save strategy. For libraries where you regularly delete and re-create database files or members, you might not be able to use the Save Changed Object (SAVCHGOBJ) approach.

Recovery option 3–Deleting files or members before restoring: Another option is to delete (from the new system) the database file or the database file member that causes the problem. When you use this method, you must make provisions for any logical files that are dependent on the files or members that you plan to delete.

1. To determine whether dependent logical files exist, use the Display Database Relationships (DSPDBR) command on your test system:

   ```
   DSPDBR FILE(library-name/file-name) MBR(member-name)
   ```

   Note: Specify the member only when you need to delete individual members instead of the whole file.
2. If no database dependencies exist, continue with step 5.
3. On your existing system, use the Save Object (SAVOBJ) command to save each dependent file to tape.
4. On your test system, use the Delete File (DLTF) command to delete each dependent file.
5. On your test system, delete the physical files or file members.
6. From your SAVCHGOBJ tape, use the Restore Object (RSTOBJ) command to restore the physical files or physical file members that you could not previously restore.
7. If you saved dependent files in step 3, use the RSTOBJ command to restore them.

Recovery option 4–Using a temporary library: A variation of the previous option is to use a temporary library on your new system.

1. On the test system, use the Create Library (CRTLIB) to create a temporary library.

2. Restore the physical files or file members from your SAVCHGOBJ tape to the temporary library. (Use the RSTOBJ command with the SAVLIB and RSTLIB parameters.)
3. To determine whether dependent logical files exist, use the Display Database Relationships (DSPDBR) command on your test system:

```
DSPDBR FILE(library-name/file-name) MBR(member-name)
```

   Note: Specify member only when you need to delete individual members instead of the whole file.
4. If no database dependencies exist, continue with step 7.
5. On the test system, use the Copy File (CPYF) command to copy dependent files from the original libraries to the temporary library.
6. Delete the dependent files from the original libraries.
7. Delete the physical files from the original libraries.
8. Copy the physical files from the temporary library to the original libraries.
9. If the temporary library contains any dependent files, copy them to the original libraries.
10. Use the Delete Library (DLTLIB) command to delete the temporary library.

Problems with object authority or ownership

To protect you from someone who might attempt to restore an illicit program to your system, the system makes authority or ownership changes during some restore operations.

Review the job log to determine whether these changes occurred. You can use the Change Object Owner (CHGOBJOWN) command to transfer ownership to the correct user profile. You can use the Edit Object Authority (EDTOBJAUT) command to make changes to public authority for an object.

Moving entire libraries

When your library structure is organized to separate volatile information (for example, database files) from static information (for example, programs), this synchronization method might be simple and effective. You can periodically restore entire database file libraries from your existing system to your new system.

Note: Do not use this method for IBM-supplied (Q*xxx*) libraries, for example, libraries QSYS2, QGPL, and QUSRSYS.

Here is an example of the steps for moving a library:

1. On your existing system, sign on with the QSECOFR user profile to avoid authority problems.
2. Place your existing system in a restricted state to ensure that you get a fixed (static) copy of the database files.
3. Use the Save Security Data (SAVSECDTA) command to save all user profiles.

   You use this information to correctly synchronize ownership and authority for any new objects that you move.
4. Use the Save Library (SAVLIB) command to save the libraries to tape:

```
SAVLIB LIB(library-name) DEV(tape-device)
   ENDOPT(*LEAVE) ACCPTH(*YES)
```

   Notes:

   - Specify ENDOPT(*REWIND) when you save the last library.
   - You can specify multiple libraries on the SAVLIB command.
5. Restart the controlling subsystem on your existing system.
6. On your test system, sign on with the QSECOFR user profile to avoid authority problems.

7. Place your new system in a restricted state to ensure that you do not have restore problems because of object-locking conflicts.
8. Use the Clear Library (CLRLIB) for each library that you plan to restore.

   This eliminates any potential problems with objects not restoring because of mismatches between the media version and the system version.

   Note: If you restore a library that contains Structured Query Language (SQL) collections that contain *DTADCT objects, for each of these libraries use the Delete Library (DLTLIB) command. (Use the DLTLIB command rather than the Clear Library (CLRLIB) command.) SQL collections that contain *DTADCT objects will fail during the Restore Library (RSTLIB) operation unless you first delete the library.
9. To restore the saved user profiles, use the Restore User Profiles (RSTUSRPRF) command:

   ```
   RSTUSRPRF USRPRF(*ALL) DEV(tape-device) ENDOPT(*LEAVE)
   ```
10. For each library that you saved, use the RSTLIB command:

    ```
    RSTLIB SAVLIB(library-name) DEV(tape-device) MBROPT(*NEW)
       ENDOPT(*LEAVE) OPTION(*NEW) ALWOBJDIF(*ALL)
    ```

    Notes:

    - If you have a different auxiliary storage pool (ASP) organization on your new system, you might need to specify the SAVASP and RSTASP parameters.
    - You specify ALWOBJDIF(*ALL) because you might be restoring to a system with a different serial number. ALWOBJDIF(*ALL) links the authorization lists back with the objects. You should only specify ALWOBJDIF(*ALL) when you are restoring to an empty library or the library does not exist on the system.
    - When you restore the last library, specify `ENDOPT(*REWIND)` unless you have additional objects to restore from the tape.
11. Check your job log to ensure that the system successfully restored the libraries.
12. Use the Restore Authority (RSTAUT) command to restore private authorities to objects.

Considerations for moving entire libraries

Consider these items when you use the synchronization method of moving entire libraries.

- You might want to combine this method with the Save Changed Object (SAVCHGOBJ) method. For example, you might move several large libraries that contain database files. You can use the SAVCHGOBJ command for other libraries (by using the OMITLIB parameter on the SAVCHGOBJ command to omit libraries that you are moving in their entirety). Structured Query Language (SQL) collections containing *DTADCT objects will fail during the Restore Library (RSTLIB) operation unless you first delete the library.
- When you use this method, you need to decide how to handles document library objects (DLOs) and objects in directories. You might use a save-changed method for those objects. Or, you might consider moving entire folders or directories if that method will work with your folder and directory organization.
- In all cases, it is essential that you have a complete copy of your existing system on tape before you cut over to your new system. This provides for recovery if you discover you have neglected to synchronize some critical objects.

Moving individual objects

With this method, you periodically copy specific objects (or example, database files) from your existing system to your new system.

The moving individual objects method is most often used in the following two situations:

- When you have a short test period, careful change control, and a very well-defined set of database files that change frequently.
- When you plan to completely rebuild your new system at the end of the test period. In this case, you might periodically move individual objects to create a more current set of test data on your new system.

Here is an example of the procedure for moving individual objects:

1. On your existing system, sign on with the QSECOFR user profile to avoid authority problems.
2. Place your existing system in a restricted state to ensure that you get a fixed (static) copy of the database files.
3. Use the Save Security Data (SAVSECDTA) command to save all user profiles.

   You use this information to correctly synchronize ownership and authority for any new objects that you move.
4. Use the Save Object (SAVOBJ) command to save individual objects that you want to synchronize:

   ```
   SAVOBJ OBJ(object-name)
      LIB(library-name)
      OBJTYPE(object-type)
      DEV(tape-device) ENDOPT(*LEAVE)
   ```

   Notes:

   - Specify `ENDOPT(*REWIND)` for the last object.
   - On the same SAVOBJ command, you can save multiple objects of the same type from the same library.
5. Restart the controlling subsystem on your existing system.
6. Place your new system in a restricted state.
7. On the new system, use the Restore User Profiles (RSTUSRPRF) command:

   ```
   RSTUSRPRF USRPRF(*ALL) DEV(tape-device)
       ENDOPT(*LEAVE)
   ```
8. For each library that contains objects that you saved, use the RSTOBJ command:

   ```
   RSTOBJ OBJ(*ALL) SAVLIB(library-name)
      DEV(tape-device) ENDOPT(*LEAVE) OBJTYPE(*ALL)
   ```

   Notes:

   - If you have a different auxiliary storage pool (ASP) organization on your new system, you might need to specify the SAVASP and RSTASP parameters.
   - When you restore the last object, specify `ENDOPT(*REWIND)`.
9. Check your job log to ensure that the system successfully restored the objects.
10. Use the Restore Authority (RSTAUT) command to restore private authorities to objects.
11. Restart the controlling subsystem on your new system.

Applying journaled changes

With this synchronization method, you use system journaling capabilities to synchronize the information in journaled objects on your test and existing systems.

The applying journaled changes method is most commonly used either by installations that already use and understand journaling or by installations that have very large journaled objects.

The advantage of this method is that you save and restore only the changes that occur to a journaled object, not the entire object. The disadvantage of this method is its complexity.

Journaling before saving

You must set up journaling on your existing system before you perform the save operation from which you build your new systems. When you start journaling for an object, the system assigns a journal ID number (JID) to that object. An object must have a JID before you save it from your existing system and restore it to your new system or you will not be able to apply journaled changes to the test version.

Conceptually, when you use journaling to synchronize objects, you perform the following steps to establish a synchronization point:

1. On your existing system, follow these steps:
   a. Save security data to get a current copy of user profiles and private authorities.
   b. Save new journal receivers (that contain entries that you have not yet applied on your new system).
   c. Save any new objects (that do not yet exist on your new system).
2. On your new system, follow these steps:
   a. Restore user profiles (to get any new profiles and current authority information).
   b. Restore any new objects.
   c. Restore journal receivers.
   d. Apply journaled changes from the journal receivers.
   e. Restore authority to any new objects.

Here is an example of the steps for using journaled changes to synchronize systems:

1. To establish a checkpoint on your existing system, follow these steps:
   a) Use the Work with Journal Attributes (WRKJRNA) command to determine what journal receivers have been created since your last checkpoint.
   b) Write down the names of the new journal receivers.
   c) Determine whether any files, data areas, data queues, or libraries have been added to any journaled objects since your last checkpoint.

      You can use the Display Journal (DSPJRN) command to do this by typing the following command:

      ```
      DSPJRN JRN(journal-name) JRNCDE((D) (E) (Q) (Y)) ENTTYP(JF EG QB YB)
             RCVRNG(first-receiver *CURRENT)
      ```

      For *first-receiver*, use the name of the first receiver after the most recent checkpoint.

      Note: If you are journaling integrated file system objects, and your directories are not using the inherit journaling attribute, look for new integrated file system objects by adding B to the JRNCDE parameter, and JT to the ENTTYP parameter.
   d) Write the new object names on a list. (You will need to save them later.)
   e) If you have other journals on your system, repeat step 1a through step 1c for each additional journal.
   f) For each journal on your system, use the Change Journal (CHGJRN) command to detach the current journal receivers and attach new journal receivers.
   g) Use the Save Object (SAVOBJ) command or Save Object (SAV) command to save any newly journaled objects that you listed in step 1d and step 1c.

      Use the Save Library (SAVLIB) command to save journaled libraries that you listed in step 1d.

      Note: The system needs an exclusive lock on an object to save it. You might need to stop certain application activity on your system to be able to save the newly journaled objects.
   h) Use the SAVOBJ command to save the journal receivers that you listed in step 1b.
   i) If you do not have a current copy of your user profiles on tape, use the Save Security Data (SAVSECDTA) command to save them to tape.
   j) You have completed establishing a new checkpoint (such as Point 2) on your existing system.

2. To synchronize the journaled objects on your new system with the existing versions, follow these steps:
   a) Place your new system in a restricted state.
   b) On the new system, use the Restore User Profiles (RSTUSRPRF) command:

```
RSTUSRPRF USRPRF(*ALL) DEV(tape-device)
    ENDOPT(*LEAVE)
```

   c) Use the Restore Object (RSTOBJ) or Restore Object (RST) command to restore any objects that you saved in step 1g.
   d) Use the RSTOBJ command to restore the journal receivers that you saved in step 1h.
   e) Use the Apply Journaled Changes (APYJRNCHG) command to bring your journaled objects to the checkpoint level:
      - For the starting receiver, specify the journal receiver that was created and attached when you established your most recent checkpoint on the existing system. For the ending receiver, specify *CURRENT.
      - For the starting sequence number, specify *FIRST.
      - For the ending sequence number, specify *LAST.

      See the Journal management topic to understand what journal entries might either be skipped or might stop the apply process.
   f) Use the Restore Authority (RSTAUT) command to restore private authorities for any new objects that you restored.
   g) Restart the controlling subsystem on your new system.

Related concepts

Restoring journaled objects

This section describes options for restoring journaled objects (for example, libraries, data areas, data queues, database files, and integrated file system objects).

Related information

Journal management

Refreshing your new system

When you use this method, you periodically rebuild your entire new system from the current information on your existing system.

To make sure that user data and system data are merged correctly, follow the same procedure that you used when you initially built your new system:

1. Install the Licensed Internal Code from scratch.
2. Install the IBM i operating system.
3. Restore user data.
4. Install licensed programs.

Note: The preceding list is a conceptual view of the sequence. Use the checklists for the complete list of steps.

Before you refresh your new system, be sure to save the work that you did on your new system. In particular, save any program objects that you converted. After you rebuilt the new system, restore these converted objects.

Additional synchronization tips

Consider these items when you plan to keep you test and existing systems synchronized during a test period.

- Synchronization (especially the first few times) can be difficult. You should always save your entire new system before you begin your synchronization efforts. With these save tapes available, you have the option of restoring your entire new system and starting the synchronization again (or changing synchronization methods).
- To synchronize correctly, you need to understand how to save and restore the authority to objects. When you use the Save Security Data (SAVSECDTA) command, the system saves user profiles, authorization lists, and private authorities.

  When you restore user profiles (RSTUSRPRF *ALL), the system restores user profiles and authorization lists. It puts private authority information in work files on the system. After you restore objects, use the Restore Authority (RSTAUT) command to reestablish the private authorities that are in the work files.
- When you are ready to perform your final synchronization before you move your new system to productive use, be sure to plan both for disaster recovery and for verification. If possible, keep your former existing system functional during the verification period in case you discover objects that you did not properly synchronize.

  In addition, make sure that you save an entire copy of your former existing system to tape before you dismantle it. You might need the objects on these tapes if you discover errors in your synchronization procedures.

  Finally, print listings from your former existing system that you can use as a basis for verifying the validity of information on your new existing system.
- Good synchronization requires careful change-control procedures. You must plan for information that you cannot easily synchronize automatically, such as changes to system information. As much as possible, you should freeze changes to this type of information during your test period.
- If you do not use the Save Changed Objects (SAVCHGOBJ) command as part of your synchronization strategy, you might need to use special procedures to synchronize mail.

  The steps for moving mail from your existing system to your new system are as follows:

  1. On your existing system, use the Save Document Library Object (SAVDLO) command:

     ```
     SAVDLO DLO(*MAIL) DEV(tape-device)
     ```

  2. On your new system, use the Restore Document Library Object (RSTDLO) command:

     ```
     RSTDLO DLO(*MAIL) DEV(tape-device)
     ```

- To synchronize the Backup Recovery and Media Services (BRMS) licensed program, follow these steps:

  Note: Use the following tip only when no other BRMS activity exists on any system that might be in a BRMS network with this system.

  1. On your existing system, stop all activity that might place locks on objects in the BRMS libraries. If you have scheduled jobs that use BRMS, you need to hold them.
  2. Mount a tape that is compatible with the tape unit on your new system.
  3. Type the following command:

     ```
     SAVLIB LIB(QUSRBRM) DEV(tape-device)
     ```

     Note: If you want, you can use save files and transfer the library electronically.
  4. On the new system, follow these steps:

     a. Stop all activity that might place locks on objects in the BRMS libraries. If you have scheduled jobs that use BRMS, you need to hold them.

b. Save a copy of the current BRMS product; enter the following command:

```
SAVLICPGM LICPGM(57nnBR1) DEV(tape-device)

(Replace nn with the appropriate number for your release,
 for example, DSPPTF 5761BR1 for V6Rx.)
```

c. Delete the version of BRMS that has outdated information; enter the following command:

```
DLTLICPGM LICPGM(57nnBR1)
```

d. Delete library QUSRBRM.

e. Mount the tape that you created in step 3.

f. Restore the BRMS libraries; enter the following command:

```
RSTLIB SAVLIB(QUSRBRM) DEV(tape-device)
```

g. Load the tape that you created using SAVLICPGM in step 4b.

h. Restore the current version of BRMS; enter the following command:

```
RSTLICPGM LICPGM(57nnBR1) DEV(tape-device)
```

i. To set up BRMS again, see Backup, Recovery, and Media Services (BRMS).

Part 4. Tips for merging two or more IBM i operating systems

© Copyright IBM Corp. 1997, 2025

Chapter 19. Tips for merging two systems into a single system

Merging systems is a complex process. Various factors can affect merging systems that include the software release, system names, or IBM-supplied objects.

The following information is provided as a guideline to assist you in the planning process of merging systems. You should consult the documentation for your additional products to ensure proper migration.

Objects that can migrate from one system only are the following items:

- Directory entries
- Office enrollments
- Distribution lists
- Other IBM-supplied files or libraries
- System values
- Network attributes
- Access path recovery times
- Communications configurations
- System reply list
- Edit descriptions

Determine which system to restore all of these objects from and restore that system first. If you have a production system and a development system, restore the production system first, then follow the guidelines below to restore the information from the development system.

Guidelines for restoring information from the development system

Follow these guidelines when you restore information from the development system.

1. Determine which system to restore first. To assist in determining this, you might need to answer several questions: Which system is more complex? Which has more users? Which system is more critical to your operations? If you are choosing between merging a production system and merging a development system, selecting the production system is recommended.

   Restore the production system by following the steps in Table 29 on page 94.
2. User profiles and associated objects that are the same on both systems will not be merged. However, they are affected in the following manner:
   - Object owners, authorization lists, and primary groups will be restored from the production system.
   - Passwords and group connections will be restored from the development system.
   - Merging of private authorities is an AND operation. The object authorities and data authorities will be added together from both systems. The resulting authorities on the merged system will be the higher of the matching authorities from the production and development systems.
   - USRPRF (*NEW) and OMITUSRPRF are parameters that might be useful when you consolidate systems. They allow you to restore only new user profiles or omit certain user profiles. If you want to preserve the Digital Certificate Manager information when you combine the systems, specify the OMITUSRPRF(*DCM) parameter on the RSTUSRPRF command. See “Restoring user profiles” on page 211.
3. Groups of configurations that are needed from the development system can be restored with the Restore Configuration (RSTCFG) command:

© Copyright IBM Corp. 1997, 2025

```
RSTCFG OBJ(workstation) OBJTYPE(*DEVD)
       SRM(*NONE)
```

Automatic configuration can also be enabled to recover the groups of configurations from the development system.

4. User libraries can be restored with the Restore Library (RSTLIB) command.

   Note: Be sure to omit any IBM-supplied libraries such as QGPL and QUSRSYS.

   If there are libraries that are the same on both systems, you should consider using the OPTION(*NEW) parameter to restore only the new objects:

   ```
   RSTLIB SAVLIB(user_library) OPTION(*NEW)
   ```

   Then determine which objects you want from each system and restore those objects individually. If there are objects in QGPL or QUSRSYS that are unique to either system, those objects should be restored individually as well.

5. Documents and folders can be restored with the Restore Document Library Object (RSTDLO) command. When saving documents and folders to be restored, any IBM-supplied folders should be omitted when using the Save Document Library Object (SAVDLO) command:

   ```
   SAVDLO DLO(*ALL) OMITFLR(Q*)
   ```

   If any IBM-supplied folders are restored, the original information might be overwritten. Additional considerations will need to be made if any of the DLOs are from a previous release.

6. The integrated file system objects can be restored with the following command:

   ```
   RST OPTION(*NEW)
   ```

7. Once the preceding instructions have been completed, run the Restore Authorities (RSTAUT) command.
8. Once the RSTAUT command completes, perform a normal initial program load (IPL).

Related tasks

Restoring user profiles

You can restore a single user profile, a list of user profiles, or all user profiles. You restore a user profile to move a user from one system to another system and to recover a damaged user profile.

Part 5. Alternate Installation device

© Copyright IBM Corp. 1997, 2025

Chapter 20. Using an alternate installation device

You can learn how to set up, disable, verify, and select an alternate installation device.

Related tasks

Task 1: Getting ready to load the Licensed Internal Code
You can learn the things that you need to get ready to load the Licensed Internal Code.

Task 4: Loading the Licensed Internal Code from media
You need to load the Licensed Internal Code from media to complete the preparation.

Alternate installation device: Overview

You can perform installation and recovery procedures by using the primary installation device along with an alternate installation device that you define.

The alternate installation device can be a tape device or an optical device. Some systems, typically with faster devices attached, might see a performance improvement when using an alternate installation device. You can use the alternate installation device for any installation or recovery that requires replacing Licensed Internal Code. Some models might require that you set up an alternate installation device.

When you use an alternate installation device, the system might use the primary installation device to load a small portion of the Licensed Internal Code during an IPL-type D and then continues the installation by using the media in the alternate installation device. The media in the alternate installation device can be SAVSYS media or distribution media created by a central site.

Notes:

- If you use an alternate installation device, you need to ensure that you set up the device and that you enable the device. You also need to have the optical media or tape media that contains the Licensed Internal Code and your save media.

Setting up an alternate installation device

Before you use an alternate installation device, you need to ensure that you define it on a bus, and you must enable the device.

You need to record and keep the logical address of the system bus and system card at which the alternate installation device is attached. If you have the device set up on a different bus and you do not have this information available, you will not be able to complete installations or recoveries.

It is helpful to record the tape or optical drive's serial number that you plan to use as the alternate installation device.

If you change the configuration of your system, you need to ensure that this information is correct before you start to use the alternate installation device.

To set the addresses and enable the alternate installation device, perform the following steps:

Note: You need to know the password for dedicated service tools (DST) to perform this procedure.

1. Place the media that contains the current release Licensed Internal Code into the primary installation device.
2. Set the system mode to Manual.

   Then perform an initial program load (IPL) using this command:

   ```
   PWRDWNSYS OPTION(*IMMED) RESTART(*YES) IPLSRC(B)
   ```

   Note: The default value for the Power Down System (PWRDWNSYS) command is to display a confirmation. This can be changed by specifying *NO for the Confirm (CONFIRM) parameter: `CONFIRM(*NO)`.

© Copyright IBM Corp. 1997, 2025

Alternative: An alternative to this step is to select function 21 (Dedicated service tools). If you use this alternative, the next step is step 4. You will not need to perform step 12.

Logical partition users:

- If you are using Hardware Management Console (HMC), see the Managing the Hardware Management Console (HMC) topic.

3. When the IPL or Install the System display appears, select option 3 (Use dedicated service tools (DST)) and press the Enter key.
4. The Dedicated Service Tools (DST) Sign On display appears.

```
                  Dedicated Service Tools (DST) Sign On
                                                        System:   SYSTEMA
Type choices, press Enter.
   Service tools user . . . . . .  QSECOFR
   Service tools password . . . . . .  _______
```

Sign on using the QSECOFR user profile.

Note: Do not sign on with a profile other than QSECOFR.

5. The Use Dedicated Service Tools (DST) menu appears.

From the Use Dedicated Service Tools (DST) menu, perform the following tasks:

a. Select option 5 (Work with DST environment).
b. Select option 2 (System devices) on the Work with DST Environment display.
c. Select option 5 (Alternate installation device) on the Work with System Devices display.
d. Select your device type on the Select Media Type display.

6. The Select Alternate Installation Device display lists available devices.

```
                  Select Alternate Installation Device
                                                      System:   SYSTEMA

Additional devices may be available for use.  Press F5 to see
if any additional devices have been detected.

Type option, press Enter.
  1=Select   5=Display details

         Bus      Resource                    Serial
Option   Number   Name         Type   Model   Number          Selected
           16     TAP22        3592   01A     13-1300283
           16     TAP23        3592   E05     13-1350492
           16     TAP24        3592   01A     13-1300291
           16     TAP25        3592   E05     13-1350477
           16     TAP26        3580   002     00-0110666
           16     TAP27        3590   E1A     13-C1610
           16     TAP28        3590   E1A     13-C1599
           16     TAP29        3590   01A     13-AB733
           16     TAP30        3590   01A     13-AB388
           16     TAP31        3580   003     00-19944
                                                                More...
F2=Deselect device   F3=Exit   F5=Refresh   F12=Cancel
```

From this display, type a 5 (Display details) next to the resource you want to view details and press the Enter key.

7. The Display Device Details display appears.

```
                          Display Device Details
                                                              System:   SYSTEMA
Resource                                    Serial
Name            Type        Model           Number
TAP23           63A0        001             13-1355555

Location    . . . . . . . . . . . . . . :   UTMP3.C01.01B7A4G-DB3-D41


Logical address:
 Bus:
  System bus  . . . . . . . . . :      16
  System board  . . . . . . . . :       0
  System card . . . . . . . . . :      32

 Storage:
  I/O bus number  . . . . . . . :       0
  Controller  . . . . . . . . . :      14
  Device address  . . . . . . . :       0



F3=Exit   F12=Cancel
```

You need to have a record of the addresses assigned to the alternate installation device selected for installing a system. Record the following information:

- Type/Model: _____
- System bus: _____
- System card: _____
- Serial number: _____

Notes:

- You might want to repeat this process to record the addresses for all alternate installation devices that appear in step 6. You should store this information in a safe place, such as the location where your recovery information and recovery media are stored.
- If more than one alternate installation device is defined, only one can be enabled.
- You should ensure that only one device contains valid installation media. This prevents you from loading the wrong version of the Licensed Internal Code.

Press the Enter key.

8. The Select Alternate Installation Device display appears.

   Type 1 (Select) next to the resource you want and press the Enter key.
9. You should see the following message at the bottom of the display:

   ```
   Alternate installation device selected
   ```
10. Press F3 (Exit) to return to the Use Dedicated Service Tools (DST) display.
11. Press F3 (Exit) again.

    The Exit Dedicated Service Tools (DST) display appears.

```
                        Exit Dedicated Service Tools (DST)
                                                              System:   SYSTEMA
Select one of the following:

     1. Exit DST
     2. Resume dedicated service tools
```

Type 1 (Exit Dedicated Service Tools (DST)) and press the Enter key.

12. The next display you see is the IPL or Install the System display.

    Type 1 (Perform an IPL) and press the Enter key to complete the procedure.

Disabling an alternate installation device

If you need to disable an alternate installation device, follow these steps.

You might need to disable an alternate installation device for one of the following reasons:

- To continue with an installation using the primary installation device.
- To enable a different device as an alternate installation device.
- To correct the logical address if hardware has been moved or changed.

1. Use the system or logical partition control panel to set the mode to `Manual`.

   Then perform an attended initial program load (IPL) using the command: `PWRDWNSYS OPTION(*IMMED) RESTART(*YES) IPLSRC(B)`.

   Note: Beginning in V6R1, the default value for the Power Down System (PWRDWNSYS) command is to display a confirmation. This can be changed by specifying *NO for the Confirm (CONFIRM) parameter: `CONFIRM(*NO)`.

   Alternative: An alternative to this step is to select function 21 (Dedicated service tools). If you use this alternative, skip steps "2" on page 364 and "9" on page 364.

   Logical partition users: If you are using Hardware Management Console (HMC), see the Managing the Hardware Management Console (HMC) topic.
2. When the IPL or Install the System display appears, select option 3 (Use dedicated service tools (DST)) and press the Enter key.
3. The Dedicated Service Tools (DST) Sign On display appears. Sign on using the QSECOFR user profile.
4. The Use Dedicated Service Tools (DST) menu appears.

   From the Use Dedicated Service Tools (DST) menu, perform the following tasks:

   a. Select option 5 (Work with DST Environment).

   b. Select option 2 (System devices) on the Work with DST Environment display.

   c. Select option 5 (Alternate installation device) on the Work with System Devices display.
5. At the Select Alternate Installation Device display, press F2 (Deselect device).
6. You should see the following message at the bottom of the display:

   ```
   Alternate installation device deselected
   ```
7. Press F3 (Exit) to return to the Use Dedicated Service Tools (DST) display.
8. Press F3 (Exit) again. The Exit Dedicated Service Tools (DST) display appears.

   Type 1 (Exit DST) and press the Enter key.
9. The next display you see is the IPL or Install the System display.

   Type 1 (Perform an IPL) and press the Enter key to complete the procedure.

Verifying and selecting an alternate installation device during a manual installation

To verify and select an alternate installation device during a manual installation, follow these steps.

1. The Install Licensed Internal Code display appears.

```
                      Install Licensed Internal Code

Select one of the following:

        1. Install Licensed Internal Code
        2. Work with Dedicated Service Tools (DST)
        3. Define alternate installation device
```

Type 3, and press Enter. This option provides the ability to define an alternative installation device as the install source for the Licensed Internal Code.

2. Verifying and selecting alternate installation device: The Select Alternate Installation Device Bus display appears.

```
                  Select Alternate Installation Device Bus
                                                                 System: XXXX
Type Option, press Enter.
   1=Select

Option      Bus Number          Selected
  _             1
  _             2
  1             3
  _             4
  _             5
  _             6
  _             7
  _             8
  _             9
  _             A
  _             B
  _             C
  _             D
                                                               More.....
F2=Deselect Device  F3=Exit  F5=Refresh    F12=Cancel
F13=Bypass bus select
```

   a. Verify that the selected device is on the correct system bus. Type 1 in the **Options** field next to the selected bus, and press Enter.

   b. Select the media type on the Select Media Type display and press Enter.

```
                    Select Media Type
                                                                 System: XXXX
Select one of the following:

  1. Tape
  2. Optical
```

      This might take several minutes. If you see the message

      `No alternate installation device configured`

      , wait one minute and refresh the display.

   c. The Select Alternate Installation Device display appears. Use this display to verify the resource name, type, model, and serial number for the device.

```
                     Select Alternate Installation Device
                                                          System:   SYSTEMA

Additional devices might be available for use.  Press F5 to see
if any additional devices have been detected.

Type option, press Enter.
  1=Select   5=Display details

          Bus      Resource                    Serial
 Option   Number   Name         Type   Model   Number          Selected
   _         1     TAP01        6384   001     00-0270004
   _         3     TAP19        3570   B11     13-22222


F2=Deselect device    F3=Exit    F5=Refresh    F12=Cancel
```

   d. Type 1 to select the device.

   e. The message

      `Alternate installation device selected`

appears. Press F3 to return to the Install Licensed Internal Code display.

f. Type 1 and press Enter to install the Licensed Internal Code.

End of verifying and selecting alternate installation device.

If you are using an Hardware Management Console (HMC) to perform an IPL of the system from alternate media, complete the steps in Task 3: Preparing to start the system from an alternate IPL device.

Part 6. Disk configuration and protection

© Copyright IBM Corp. 1997, 2025

Chapter 21. Configuring disks and disk protection

The topic describes the procedures for using the character-based interface to configure and protect disks on your system.

The topic includes checklists for running the disk configuration tasks in the correct sequence. You can also configure disks and disk protection through IBM Navigator for i.

You can use system service tools (SST) to do some disk configuration procedures while your system is active. For other procedures, you must stop your system and use dedicated service tools (DST). The topic provides information about both SST and DST.

Related information

Disk pools

Disk protection

Choosing the right procedure for configuring disks

This topic contains several checklists for performing configuration procedures.

Use Table 66 on page 369 to determine which checklist to use for your situation.

Table 66. Choosing the right disk procedure

Task description	Procedure to follow	Requires DST?
Configure your system for the first time	Checklist 1 on page "Checklist 1: Configuring disks on a new system" on page 370.	Yes
Add one or more disk units that will not have device parity protection. This checklist applies to disk units that are capable of device parity protection if you do not plan to start device parity protection for the disks.	Checklist 2 on page "Checklist 2: Adding disk units without device parity protection" on page 371.	No
Add one or more disks to an existing input/output adapter (IOA). Use this checklist if you plan to protect some or all of the new disks with device parity protection.	Checklist 3 on page "Checklist 3: Adding disk units to an existing input/output adapter" on page 372.	No
Add a new IOA. Use this checklist if you plan to protect some or all of the new disks with device parity protection.	Checklist 4 on page "Checklist 4: Adding a new input/output adapter" on page 373.	Yes
Move disk units between existing auxiliary storage pools (ASPs) without mirrored protection.	Checklist 5 on page "Checklist 5: Moving disk units between non mirrored auxiliary storage pools" on page 375.	Yes
Move disk units between existing ASPs with mirrored protection.	Checklist 6 on page "Checklist 6: Moving disk units between mirrored auxiliary storage pools" on page 376.	Yes
Delete a user ASP.	Checklist 7 on page "Checklist 7: Deleting an auxiliary storage pool" on page 377.	Yes

© Copyright IBM Corp. 1997, 2025

Table 66. Choosing the right disk procedure (continued)		
Task description	**Procedure to follow**	**Requires DST?**
Remove one or more disk units without device parity protection.	Checklist 8 on page "Checklist 8: Removing disk units without device parity protection" on page 378.	Yes
Remove one or more disk units from an IOA. Use this checklist if device parity protection is started for some or all of the disk units that are attached to the IOA and they are in ASPs without mirrored protection.	Checklist 9 on page "Checklist 9: Removing disk units that have device parity protection from an auxiliary storage pools without mirrored protection" on page 379.	Yes
Remove one or more disk units from an IOA. Use this checklist if device parity protection is started for some or all of the disks units that are attached to the IOA and they are in ASPs with mirrored protection.	Checklist 10 on page "Checklist 10: Removing disk units that have device parity protection from an auxiliary storage pool with mirrored protection" on page 380.	Yes

Checklist 1: Configuring disks on a new system

This checklist shows the sequence of tasks that you use to configure disks on a new IBM i environment.

Whether you need to perform all the tasks depends on the disk protection that you want on your system.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 67. Configuring disks on a new system–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Start dedicated service tools (DST).	"Starting dedicated service tools" on page 382.
___ Task 2	Display your disk configuration. Currently, all of your disk units except the load source unit appear as nonconfigured.	"Displaying your disk configuration" on page 385.
___ Task 3	If you plan to have device parity protection on any of your disk units, start it using the procedure for the types of disk units that you have.	"Starting device parity protection for an input/output adapter" on page 419
___ Task 4	Add nonconfigured disk units to the correct auxiliary storage pools (ASPs).	"Adding disk units to an existing auxiliary storage pool" on page 392.

Table 67. Configuring disks on a new system–Tasks (continued)

Task	What to do	Where to read more about it
___ Task 5	The default storage threshold for each ASP is 90%. If you want a different storage threshold for any ASP, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 6	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Changing the storage threshold for the system auxiliary storage pool" on page 400.
___ Task 7	If you plan to have mirrored protection for any ASPs, start it.	"Starting mirrored protection" on page 431.
___ Task 8	If you started mirrored protection, wait until the system completes the initial program load. Then sign on and start system service tools (SST).	"Starting system service tools" on page 384.
___ Task 9	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 10	End DST or SST.	"Stopping dedicated service tools" on page 383 or "Stopping system service tools" on page 384.

Related information

Planning for disk protection

Checklist 2: Adding disk units without device parity protection

This checklist shows the sequence of tasks that you use to add one or more disk units to your system when you do not plan to protect the new disk units with device parity protection.

You can use either dedicated service tools (DST) or system service tools (SST) to perform the tasks in this checklist. If you use SST, you can perform the tasks while your system is active. If you use DST, you must stop your system to perform the tasks in this checklist.

Adding to an auxiliary storage pool (ASP) with mirrored protection?

You can add disk units to an ASP that has mirrored protection without stopping and starting mirrored protection. You must add disk units in pairs with equal capacities. The added units will always be paired with each other. You might want to choose a later time, when your system can be unavailable for several hours, to stop and start mirrored protection. When you start mirrored protection again, the system evaluates the pairing for all disk units on your system. This might provide a higher level of availability for failures that affect a controller, an input/output adapter (IOA), or a bus.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 68. Adding disk units without device parity protection–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Physically attach disk units.	
___ Task 2	Start DST or SST.	"Starting dedicated service tools" on page 382 or "Starting system service tools" on page 384.
___ Task 3	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 4	Add nonconfigured disk units to the correct ASPs. See note 1 and note 2.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 5	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 6	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Changing the storage threshold for the system auxiliary storage pool" on page 400.
___ Task 7	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 8	End DST or SST.	"Stopping dedicated service tools" on page 383 or "Stopping system service tools" on page 384.

1
You can add the disk units to an existing ASP or you can add them to a new ASP.

2
If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.

Checklist 3: Adding disk units to an existing input/output adapter

Use this checklist if you plan to protect some or all of the new disk units with device parity protection.

If you do not plan to protect any of the new disk units, use "Checklist 2: Adding disk units without device parity protection" on page 371.

You can use this procedure whether or not you have mirrored protection on your system because you start device parity protection before you add the disk units to an auxiliary storage pool (ASP). You can use either dedicated service tools (DST) or system service tools (SST) to perform the tasks in this checklist. If you use SST, you can perform the tasks while your system is active. If you use DST, you must stop your system to perform the tasks in this checklist.

Before you begin, make a copy of this checklist. Fill in the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in Table 69 on page 373 include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 69. Adding disk units to an existing IOA–Tasks

Task	What to do	Where to read more about it
___ Task 1	Physically attach disk units. This is normally done by a service representative.	
___ Task 2	Start DST or SST.	"Starting dedicated service tools" on page 382 or "Starting system service tools" on page 384.
___ Task 3	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 4	For the disks that you want to parity protect, do one of the following steps: • Start device parity protection. • Include the disk units in an existing device parity set.	"Starting device parity protection" on page 419 and "Including a disk unit in device parity protection" on page 427.
___ Task 5	Add nonconfigured disk units to the correct ASPs. See note 1 and note 2.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 6	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 7	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Changing the storage threshold for the system auxiliary storage pool" on page 400.
___ Task 8	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 9	End DST or SST.	"Stopping dedicated service tools" on page 383 or "Stopping system service tools" on page 384.

1
You can add the disk units to an existing ASP or you can add them to a new ASP.

2
If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.

Checklist 4: Adding a new input/output adapter

Use this checklist if you plan to protect some or all of the new disks with device parity protection.

You can use this procedure whether you have mirrored protection on your system because you start device parity protection before you add the disk units to an auxiliary storage pool (ASP). If you do have mirrored protection and you are adding disks that do not have device parity protection, you must add them in pairs that have equal capacities. You can use either dedicated service tools (DST) or system service tools (SST) to perform the tasks in this checklist. If you use SST, you can perform the tasks while your system is active. If you use DST, you must stop your system to perform the tasks in this checklist.

Note: If you do not plan to start device parity protection for any of the new disks, use the procedure in checklist 2 to add them.

Before you begin, make a copy of this checklist. Fill in the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 70. Adding a new IOA–Tasks

Task	What to do	Where to read more about it
___ Task 1	Install the new input/output adapter in the system. This is normally done by a service representative.	
___ Task 2	Physically attach disk units to the new IOA. This is normally done by a service representative.	
___ Task 3	Start DST.	"Starting dedicated service tools" on page 382.
___ Task 4	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 5	Start device parity protection for the IOA.	"Starting device parity protection for an input/output adapter" on page 419.
___ Task 6	Add nonconfigured disk units to the correct ASPs.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 7	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 8	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Displaying your disk configuration" on page 385.
___ Task 9	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 10	End DST.	"Stopping dedicated service tools" on page 383.

Notes:

- You can add the disk units to an existing ASP or you can add them to a new ASP.
- If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.

Checklist 5: Moving disk units between non mirrored auxiliary storage pools

Use this checklist to perform a sequence of tasks to move one or more disk units from one auxiliary storage pool （ASP） to another ASP when the ASPs do not have mirrored protection.

You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Note: You cannot move disk units to or from an independent ASP.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It can help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 71. Moving disk units between ASPs–Tasks

Task	What to do	Where to read more about it
___ **Task 1**	Print your current disk configuration.	“Displaying your disk configuration” on page 385.
___ **Task 2**	Calculate the space requirements for both the source and target ASPs for the disk units.	“Calculating space requirements for an auxiliary storage pool” on page 406.
___ **Task 3**	Use option 21 from the Save menu to save your entire system.	
___ **Task 4**	Start DST.	“Starting dedicated service tools” on page 382.
___ **Task 5**	Move the disk units.	“Moving a disk unit to a different auxiliary storage pool” on page 402.
___ **Task 6**	If you created a new ASP on your system when you moved disk units, the system set the storage threshold for the ASP to 90%. If you want a different threshold, change it.	“Changing the storage threshold for an auxiliary storage pool” on page 399.
___ **Task 7**	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	“Changing the storage threshold for the system auxiliary storage pool” on page 400.
___ **Task 8**	Verify that your disk configuration is correct and print a copy for your records.	“Displaying your disk configuration” on page 385.
___ **Task 9**	End DST.	“Stopping dedicated service tools” on page 383.
___ **Task 10**	If necessary, move objects between ASPs.	“Transferring objects between auxiliary storage pools” on page 408.

Checklist 6: Moving disk units between mirrored auxiliary storage pools

Use this checklist to perform a sequence of tasks to move one or more disk units from one auxiliary storage pool （ASP） to another ASP when one or more of the ASPs involved in the move have mirrored protection.

You cannot use the move unit procedure when mirrored protection is active. Instead, you need to remove mirrored pairs from the source ASP and add them to the target ASP. You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Note: You cannot move disk units to or from an independent ASP.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It can help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 72. Moving disk units between ASPs with mirrored protection–Tasks

Task	What to do	Where to read more about it
___ **Task 1**	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ **Task 2**	Calculate the space requirements for the ASPs that are involved in moving disk units.	"Calculating space requirements for an auxiliary storage pool" on page 406.
___ **Task 3**	Use option 21 from the Save menu to save your entire system.	GO SAVE: Option 21 (saving the entire system) in the Backing up your system IBM i documentation topic
___ **Task 4**	Start DST.	"Starting dedicated service tools" on page 382.
___ **Task 5**	Remove disk units that you plan to add to a different ASP.	"Removing a disk unit from an auxiliary storage pool" on page 403.
___ **Task 6**	Add nonconfigured disk units to the correct ASPs. See note 1.	
___ **Task 7**	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ **Task 8**	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Changing the storage threshold for the system auxiliary storage pool" on page 400.
___ **Task 9**	If you created any new ASPs and you want those ASPs to have mirrored protection, start mirrored protection now.	"Starting mirrored protection" on page 431.
___ **Task 10**	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.

Table 72. Moving disk units between ASPs with mirrored protection–Tasks (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 11	End DST.	"Stopping dedicated service tools" on page 383.
___ Task 12	If necessary, move objects between ASPs.	"Transferring objects between auxiliary storage pools" on page 408.

1 If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.

Checklist 7: Deleting an auxiliary storage pool

This checklist shows the sequence of tasks that you use to delete a user ASP. You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Before you begin, make a copy of this checklist. Fill in the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It can help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation. Also note that when an ASP is deleted, all data remaining in that ASP is lost.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 73. Deleting an user ASP–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 2	Calculate the space requirements for the remaining ASPs.	"Calculating space requirements for an auxiliary storage pool" on page 406.
___ Task 3	Use option 21 from the Save menu to save your entire system.	
___ Task 4	Remove objects from the ASP that you are deleting or move the objects to a different ASP.	"Transferring objects between auxiliary storage pools" on page 408.
___ Task 5	Start DST.	"Starting dedicated service tools" on page 382.
___ Task 6	Delete the ASP. This procedure places all of the disks that were assigned to the deleted ASP in nonconfigured status.	"Deleting an auxiliary storage pool" on page 405.
___ Task 7	Add nonconfigured disk units to the correct ASPs. See note 1.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 8	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.

Table 73. Deleting an user ASP–Tasks (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 9	Specify the storage threshold for the system ASP. If you use the QSTGLOWLMT and QSTGLOWACN system values, you can prevent the system ASP from filling to capacity and causing an abnormal shutdown.	"Changing the storage threshold for the system auxiliary storage pool" on page 400.
___ Task 10	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 11	End DST.	"Stopping dedicated service tools" on page 383.
___ Task 12	If necessary, move objects between ASPs.	"Transferring objects between auxiliary storage pools" on page 408.
1 If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have identical capacities.		

Checklist 8: Removing disk units without device parity protection

Use this checklist to perform a sequence of tasks to remove one or more disk units from your system when the disk units do not have device parity protection.

Use these tasks when you are permanently removing disk units from your system. Do not use these tasks when you are repairing or replacing a failed disk unit. You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It can help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 74. Removing disk units that do not have device parity protection–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 2	Calculate the space requirements for the auxiliary storage pools (ASPs) that are involved in removing disk units.	"Calculating space requirements for an auxiliary storage pool" on page 406.
___ Task 3	Use option 21 from the Save menu to save your entire system.	
___ Task 4	Start DST.	"Starting dedicated service tools" on page 382.
___ Task 5	Remove disk units that you plan to remove from the system.	"Removing a disk unit from an auxiliary storage pool" on page 403.

Table 74. Removing disk units that do not have device parity protection–Tasks (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 6	Verify that your disk configuration is correct and print a copy for your records.	“Displaying your disk configuration” on page 385.
___ Task 7	End DST.	“Stopping dedicated service tools” on page 383.

Note: This checklist works only if at least one unit remains in the ASP and there is enough capacity remaining.

Checklist 9: Removing disk units that have device parity protection from an auxiliary storage pools without mirrored protection

Use this checklist to perform a sequence of tasks to remove one or more disk units from an input/output adapter (IOA). These tasks apply when the auxiliary storage pools (ASPs) containing the disk units do not have mirrored protection and when device parity protection is started for the IOA.

Use these tasks when you are permanently removing disk units from your system. Do not use these tasks when your are repairing or replacing a failed hard disk. You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 75. Removing disk units from an IOA and a non-mirrored ASP–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Print your current disk configuration.	“Displaying your disk configuration” on page 385.
___ Task 2	Calculate the space requirements for the ASPs that are involved in removing disk units.	“Calculating space requirements for an auxiliary storage pool” on page 406.
___ Task 3	Use option 21 from the Save menu to save your entire system.	
___ Task 4	Start DST.	“Starting dedicated service tools” on page 382.
___ Task 5	Remove disk units that you plan to remove from the system.	“Removing a disk unit from an auxiliary storage pool” on page 403.
___ Task 6	Exclude the disk units from device parity protection. If you were successful in excluding the disk units, skip to task 8. Otherwise, continue with task 7.	“Excluding a disk unit from device parity protection” on page 429.
___ Task 7	Stop device parity protection for the IOA.	“Stopping device parity protection on an input/output adapter” on page 421.

Table 75. Removing disk units from an IOA and a non-mirrored ASP–Tasks (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 8	Power off the system and remove the disk units.	"Removing disk units that have device parity protection for System i while powered on" on page 422. "Removing disk units that have device parity protection for System i while powered off" on page 424. "Removing disk units that have device parity protection for a IBM i" on page 426. .
___ Task 9	Start device parity protection for the IOA.	"Starting device parity protection for an input/output adapter" on page 419.
___ Task 10	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ Task 11	End DST.	"Stopping dedicated service tools" on page 383.

Checklist 10: Removing disk units that have device parity protection from an auxiliary storage pool with mirrored protection

Use this checklist to perform a sequence of tasks to remove one or more disk units from an input/output adapter (IOA). These tasks apply when the auxiliary storage pools (ASPs) that contain the disk units have mirrored protection and when the disk units have device parity protection.

Use these tasks when your are permanently removing disk units from your system. Do not use these tasks when you are repairing or replacing a failed disk unit. You must stop your system and use dedicated service tools (DST) to perform the tasks in this checklist.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 76. Removing disk units from an IOA and a mirrored ASP–Tasks		
Task	**What to do**	**Where to read more about it**
___ Task 1	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 2	Calculate the space requirements for the ASPs that are involved in removing disk units.	"Calculating space requirements for an auxiliary storage pool" on page 406.
___ Task 3	Use option 21 from the Save menu to save your entire system.	
___ Task 4	Start DST.	"Starting dedicated service tools" on page 382.

Table 76. Removing disk units from an IOA and a mirrored ASP–Tasks (continued)

Task	What to do	Where to read more about it
___ **Task 5**	Remove disk units that you plan to remove from the system.	"Removing a disk unit from an auxiliary storage pool" on page 403.
___ **Task 6**	Exclude the disk units from device parity protection. If you were successful in excluding the disk units, skip to task 9. Otherwise, continue with task 7.	"Excluding a disk unit from device parity protection" on page 429.
___ **Task 7**	Stop mirrored protection for the ASPs that will have disk units removed. When you stop mirrored protection, one disk unit from each mirrored pair becomes nonconfigured. See note 1.	"Stopping mirrored protection" on page 434.
___ **Task 8**	Stop device parity protection for the IOA.	"Stopping device parity protection on an input/output adapter" on page 421.
___ **Task 9**	Power off the system and remove the disk units. If you stopped device parity protection in task 8, then continue with task 10. If you did not stop device parity protection, then skip to task 14.	
___ **Task 10**	Start device parity protection for the IOA.	"Starting device parity protection for an input/output adapter" on page 419.
___ **Task 11**	Add nonconfigured disk units to the correct ASPs. These disks became nonconfigured when you stopped mirrored protection in task 7.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ **Task 12**	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ **Task 13**	Start mirrored protection for the ASPs that had mirrored protection stopped in task 7.	"Starting mirrored protection" on page 431.
___ **Task 14**	Verify that your disk configuration is correct and print a copy for your records.	"Displaying your disk configuration" on page 385.
___ **Task 15**	End DST.	"Stopping dedicated service tools" on page 383.

1

You need to stop mirrored protection only if the ASP contains other disk units that are attached to the IOA and have device parity protection.

Using system service tools and dedicated service tools

This topic describes how to start and end system service tools (SSTs) and dedicated service tools (DSTs) and also includes a list of the options available through these tools.

Several backup and recovery procedures, including disk storage management, require the use of DST or SST.

Starting dedicated service tools

To start dedicated service tools (DST), follow these steps.

If the **Initial Program Load (IPL)** or **Install the System** menu is already displayed, start with step 5.

1. Place the system in manual mode.
2. Power off the system -

```
PWRDWNSYS OPTION(*CNTRLD) DELAY(600)
          RESTART(*YES) IPLSRC(B)
```

Note:

If you are sure that no jobs are running on your system, you can specify `OPTION(*IMMED)` when you power off the system. Otherwise, specify a delay time that is sufficient to allow jobs to end normally.

3. When the IPL completes, the **IPL** or **Install the System** menu appears.

```
                        IPL or Install the System
  Select one of the following:
           1. Perform an IPL
           2. Install the operating system
           3. Use dedicated service tools (DST)
           4. Perform automatic installation of the operating system
           5. Save Licensed Internal Code
```

4. Select option 3 (Use Dedicated Service Tools (DST)) and press the Enter key. The Dedicated Service Tools (DST) Sign On display is shown.

```
                    Dedicated Service Tools (DST) Sign On

Type choice, press Enter.

     Service tools user . . . . . .   _______
     Service tools password . . . . . .   _______
```

5. In the **Service tools user** field, type QSECOFR.

   In the **Service tools password** field, type your service tools password. On a new system, the password is QSECOFR. The password is case-sensitive; use all capital letters. The service tools profile QSECOFR password is expired after the first use. At the Change Service Tools User Password display, enter in all uppercase letters the current password QSECOFR and a new password, along with the verification password. Record the new password for future reference.

   The Use Dedicated Service Tools (DST) menu is shown.

```
                      Use Dedicated Service Tools (DST)

 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system

     10. Work with remote service support

     13. Work with system security
     14. End batch restricted state
```

Related tasks

Task 3: Removing the disk unit from the auxiliary storage pool configuration

To remove the disk unit from the auxiliary storage pool (ASP) configuration, follow these steps.

Displaying disk hardware configuration (service tools method)
To display your disk hardware configuration using the service tools method, follow these steps.

Displaying your disk configuration (software view)
When you display your software disk configuration, you see how disk units are assigned to auxiliary storage pools (ASPs) and how they are protected. A separate display shows disk units that are attached to the system but have not been assigned to an ASP (nonconfigured status).

Adding disk units to an existing auxiliary storage pool
Use this information to add disk units to an existing auxiliary storage pool (ASP).

Changing the storage threshold for an auxiliary storage pool
The storage threshold for an auxiliary storage pool (ASP) determines when the system warns you that the space allocated for the ASP is almost full. The default threshold for an ASP is 90%.

Moving a disk unit to a different auxiliary storage pool
You might want to move a disk unit from one auxiliary storage pool (ASP) to another.

Removing a disk unit from an auxiliary storage pool
You need to remove a disk unit from an auxiliary storage pool (ASP) for these reasons.

Deleting an auxiliary storage pool
When you delete a user auxiliary storage pool (ASP), the status of all the disk units that are assigned to the ASP becomes nonconfigured. Any data that is in the ASP is deleted.

Stopping device parity protection on an input/output adapter
These instructions only apply to input/output adapters (IOAs) that support device parity protection.

Including a disk unit in device parity protection
When you attach a new disk unit to an existing input/output adapter (IOA) that has device parity protection, you can include the disk unit into an existing parity set with other disk units.

Excluding a disk unit from device parity protection
You can exclude a disk unit that is attached to an input/output adapter (IOA) from device parity protection.

Starting mirrored protection
Use this procedure to start mirrored protection for a specific auxiliary storage pool (ASP) on your system.

Stopping mirrored protection
When you stop mirrored protection, one unit from each mirrored pair becomes nonconfigured.

Starting disk compression
You can start disk compression from the Dedicated Service Tools (DST) menu.

Stopping disk compression
Use this procedure to stop disk compression by using the dedicated service tools (DST).

Related information

Accessing service tools using DST

Stopping dedicated service tools

To stop dedicated service tools (DST), follow these steps.

1. If you do not want to see the displays for a manual initial program load (IPL), return the system to automatic mode.

   If you want to see the displays, leave the system in manual mode.
2. Press F3 until you return to the Use Dedicated Service Tools (DST) menu.
3. From the Use Dedicated Service Tools (DST) menu, select option 1 (Perform an IPL).

   Note: Do not perform an IPL if you are performing a complete system recovery. The system might take significantly longer than normal to complete the IPL. Some functions that you perform by using DST, such as starting mirrored protection, require the system to do additional work during the IPL before the system is available for your use.

Related tasks

Displaying disk hardware configuration (service tools method)
To display your disk hardware configuration using the service tools method, follow these steps.

Displaying your disk configuration (software view)
When you display your software disk configuration, you see how disk units are assigned to auxiliary storage pools (ASPs) and how they are protected. A separate display shows disk units that are attached to the system but have not been assigned to an ASP (nonconfigured status).

Adding disk units to an existing auxiliary storage pool
Use this information to add disk units to an existing auxiliary storage pool (ASP).

Changing the storage threshold for an auxiliary storage pool
The storage threshold for an auxiliary storage pool (ASP) determines when the system warns you that the space allocated for the ASP is almost full. The default threshold for an ASP is 90%.

Moving a disk unit to a different auxiliary storage pool
You might want to move a disk unit from one auxiliary storage pool (ASP) to another.

Removing a disk unit from an auxiliary storage pool
You need to remove a disk unit from an auxiliary storage pool (ASP) for these reasons.

Deleting an auxiliary storage pool
When you delete a user auxiliary storage pool (ASP), the status of all the disk units that are assigned to the ASP becomes nonconfigured. Any data that is in the ASP is deleted.

Starting system service tools

To access system service tools (SST), follow these steps.

1. At the command line, type STRSST (the Start System Service Tools command).
2. At the Start Service Tools (STRSST) Sign On display, enter your service tools user profile and password and press Enter
3. The System Service Tools (SST) menu appears:

```
                    System Service Tools (SST)
Select one of the following:
     1. Start a service tool
     2. Work with active service tools
     3. Work with disk units
     4. Work with discette data recovery
     5. Work with system partitions
```

Related information

Accessing service tools using SST

Stopping system service tools

To stop system service tools (SST), follow these steps.

1. Press F3 (Exit) until you return to the System Service Tools (SST) menu.
2. Press F3 (Exit) again. You are shown the Exit System Service Tools display.
3. Press the Enter key to end SST.

Related tasks

Displaying your disk configuration (software view)
When you display your software disk configuration, you see how disk units are assigned to auxiliary storage pools (ASPs) and how they are protected. A separate display shows disk units that are attached to the system but have not been assigned to an ASP (nonconfigured status).

Adding disk units to an existing auxiliary storage pool

Use this information to add disk units to an existing auxiliary storage pool (ASP).

Changing the storage threshold for an auxiliary storage pool
The storage threshold for an auxiliary storage pool (ASP) determines when the system warns you that the space allocated for the ASP is almost full. The default threshold for an ASP is 90%.

Displaying your disk configuration

This topic describes how to display or print your current disk configuration. It also explains some of the fields that appear on the display or listing.

For some purposes, such as planning a mirrored configuration, you need to view your disk configuration both from a hardware perspective and from a software perspective. The hardware perspective shows how the disk units are attached by bus, input/output adapter (IOA), and controller. The software perspective shows how disk units are assigned to auxiliary storage pools (ASPs) and how they are protected.

You can use dedicated service tools (DST), system service tools (SST), or commands to display your disk configuration. When you are planning changes to your disk configuration, use SST and commands to print your current configuration before you begin to make changes. After you have made changes, you can use DST to verify the new configuration before you end DST.

Displaying your disk configuration (hardware view)

You can display your disk hardware configuration by using the dedicated service tools (DST) method and or the command method.

When you display your hardware disk configuration, you see all disk-related components that are physically attached to your system. This includes disk units whose software status is nonconfigured because they are not assigned to an auxiliary storage pool (ASP) yet.

Displaying disk hardware configuration (command method)

To display the hardware configuration of disk units on your system using the command method, follow these steps.

1. On a command line, type `WRKHDWRSC TYPE(*STG)` and press the Enter key. The Work with Storage Resources display appears. The display shows buses, input/output processors (IOPs), and controllers.

```
                         Work with Storage Resources
                                                              System:   RCHASDP4
Type options, press Enter.
  9=Work with resource


    Opt  Resource        Status               Text
       CMB01         Operational          Combined function IOP
       DC01          Operational          Disk Storage Controller
       DC02          Operational          Disk Storage Controller
       DC05          Operational          Tape Controller
```

2. If you want to see the detail about disk units that are attached to a controller, type 9 (Work with resource) in the Option column for the controller.

To print the hardware configuration of disk units on your system, follow this step:

On a command line, type `DSPHDWRSC TYPE(*STG) OUTPUT(*PRINT)` and press the Enter key.

Figure 33 on page 386 shows part of the listing that you receive:

```
                          Display Spooled File
File  . . . . . :   QSYSPRT                          Page/Line   1/1
Control . . . . .   +15                              Columns     1 - 78
Find  . . . . . .
*...+....1....+....2....+....3....+....4....+....5....+....6....+....7....+...
                                          Display Hardware Resources
5770SS1 V7R1M0  950602
                                               Storage Resources List
                                                        --------
                          Serial      Part           Frame
Resource        Type-Model  Number      Number         ID
  CMB01         9162-001   10-00000     0000086G7917   1
  DC01          6602-030   00-0193825                  1
  DD001         6602-030   00-0193825                  1
  DC02          6602-030   00-17900                    1
  DD002         6602-030   00-17900                    1
```

Figure 33. Display hardware resource listing

Displaying disk hardware configuration (service tools method)

To display your disk hardware configuration using the service tools method, follow these steps.

1. Start DST or SST.
2. From the Use Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 7 (Start a service tool).
   b) Select option 4 (Hardware service manager) on the Start a Service Tool display.

   or from the System Service Tools (SST) menu, follow these steps:
   a. Select option 1 (Start a service tool).
   b. Select option 7 (Hardware service manager) on the Start a Service Tool display.
3. To print the hardware disk configuration, press F6 (Print configuration).

   If you are using SST, the output is sent to a spooled file. If you are using DST, output goes directly to a printer because spooling is not active. If your system already has a printer defined for DST, the output is sent to that printer. If you do not have a DST printer defined, you are prompted with a list of attached printers.
4. To display the configuration, select option 2 (Logical hardware resources) from the Hardware Service Manager menu.

   From this display, you can select to show system bus, processor, or main storage resources.
5. To see additional detail, type 5 (Display detail) in the Option column next to each controller and press the Enter key.
6. If you have no other tasks to perform, end DST or SST.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Displaying your disk configuration (software view)

When you display your software disk configuration, you see how disk units are assigned to auxiliary storage pools (ASPs) and how they are protected. A separate display shows disk units that are attached to the system but have not been assigned to an ASP (nonconfigured status).

To display the software configuration of disk units on your system, follow these steps:

1. From the User Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.

c) Select option 1 (Display disk configuration) on the Work with Disk Configuration display.

or from the System Service Tools (SST) menu, follow these steps:

a. Select option 3 (Work with disk units).

b. Select option 1 (Display disk configuration) on the Work with Disk Units display.

The Display Disk Configuration menu appears.

```
                          Display Disk Configuration
  Select one of the following:
     1. Display disk configuration status
     2. Display disk configuration capacity
     3. Display disk configuration protection
     4. Display non-configured units
     5. Display device parity status
     6. Display disk hardware status
```

2. Select option 1 to see the Display Disk Configuration Status display:

```
                    Display Disk Configuration Status
                 Serial                  Resource
 ASP  Unit   Number      Type Model  Name        Status
   1                                             Unprotected
         1   00-0193825  6602  030   DD001       Configured
         2   00-0163477  6602  074   DD019       DPY/Active
         3   00-0190494  6602  070   DD036       DPY/Active
         6   00-17900    6602  030   DD002       Configured
   3                                             Unprotected
         4   00-0330477  6602  074   DD005       DPY/Active
         5   00-0323200  6602  074   DD033       DPY/Active




Press Enter to continue.

F3=Exit        F5=Refresh         F9=Display disk unit details
F11=Disk configuration capacity   F12=Cancel
```

Note: If you are performing a complete system restore, all the disk units on the system might not report in right away. Verify that the number of disk units displayed matches the number of disk units physically attached to the system. If they do not match, wait a few minutes and press F5 (Refresh) until all of the disk units report in.

3. If the lower right of the display says `More...`, you can page forward to see additional units.
4. To display the capacity of your disk units and how much capacity is used, press F11 from the Display Disk Configuration Status display or select option 2 from the Use Dedicated Service Tools (DST) menu:

```
                    Display Disk Configuration Capacity
                                                 --Protected--  --Unprotected--
ASP  Unit  Type  Model  Threshold  Overflow     Size  %Used     Size  %Used
  1                       90%     No           1805    *       2063    *
       1   6602  030                             0  0.00%     1031    *
       2   6602  074                           773    *          0  0.00%
       3   6602  070                          1031    *          0  0.00%
       6   6602  030                             0  0.00%     1031    *
  3                       90%     No           1547    *          0  0.00%
       4   6602  074                           773    *          0  0.00%
       5   6602  074                           773    *          0  0.00%
```

5. To display the disk protection that is configured for each disk unit, press F11 again:
6. To display nonconfigured disk units, press F11 from the Display Disk Configuration Protection display or select option 4 from the Display Disk Configuration menu:

```
                    Display Non-Configured Units
Serial                 Resource
Number      Type  Model  Name        Capacity  Status
00-0313374  6602   074   DD003          773    DPY/Active
```

7. To print the software disk configuration, use the print key from the displays.

   If your system already has a printer defined for DST, the output is sent to that printer. If you do not have a DST printer defined, you are prompted with a list of attached printers. When you are using DST, output goes directly to the printer because spooling is not active.
8. If you have no other tasks to perform, end DST or SST.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Stopping system service tools
To stop system service tools (SST), follow these steps.

Disk configuration and status displays

This topic explains some of the fields that appear on the displays that you use to look at your disk configuration and status. You can view online information for all the fields and their possible values.

Unit field: A unit number is assigned by the system to identify a specific disk unit. The unit number is a software function and does not appear when you display the hardware configuration. When disk units are protected by mirrored protection, both disk units in a mirrored pair are assigned the same unit number.

Resource Name field: The system resource manager assigns a resource name to every hardware device that is physically attached to the system. This resource name is the link between the hardware and the software definition of the hardware. When you add a disk unit to an auxiliary storage pool (ASP), you use the resource name to identify which disk unit to add.

Status field for the ASP: The display shows the status of an entire ASP. This status indicates the software disk protection that is in effect for the ASP. The possible values are as follows:

Unprotected	Mirrored protection is not active for the ASP. However, device parity protection might be active for some or all of the disk units in the ASP. You need to look at the individual disk units to determine the level of protection for the ASP.
Mirrored	The ASP is fully protected. Mirrored protection has been started for the ASP. All of the disk units in the ASP are protected either by mirrored protection or by device parity protection.

Status–disk unit: The display also shows the status of individual disk units. The possible values are as follows:

Operational	The disk unit is operational and ready to accept input or output operations.
Not operational	The device cannot communicate with the input/output processor (IOP). You should verify that the unit is powered on.
Not ready	The device cannot perform media-related functions, but it can still communicate with the IOP.
Busy	The device is not available for processing any commands on this connection.
Read/write protected	The device cannot process either a read or a write operation. A device might be in this state due to a cache problem, a device configuration problem, or other types of problems that might cause a data integrity exposure.

Write protected	The device cannot accept write operations. Read operations are allowed.
Performance degraded	The device is functional, but performance might be impacted due to other hardware problems (such as a problem with the IOP cache).
Redundant failure	The device is functional, but availability might be impacted due to other problems (such as a redundant power supply problem). Service is required to prevent additional failures that will stop input and output operations to the device.
DPY/Failed	This unit is part of a disk unit subsystem that has device parity protection. The disk unit failed within its device parity set, causing the loss of data protection for the device parity set.
DPY/Unprotected	This unit is part of a disk unit subsystem that has device parity protection. Data protection is no longer in effect due to a failure in another resource.
DPY/Rebuilding	This unit is part of a disk unit subsystem that has device parity protection. Data protection is being rebuilt.
DPY/Active	This unit is part of a disk unit subsystem that has device parity protection. The unit is operational and ready to accept input or output operations.
DPY/Resyncing	This unit is part of a disk unit subsystem that has device parity protection. The subsystem is in the process of recreating the redundancy data for the device parity set. All units in the set that are being synchronized will have this status.
DPY/Unknown	This unit is part of a disk unit subsystem that has device parity protection. The status of this unit is not known to the system.
Active	This unit is one of a mirrored pair. It is capable of having data written to it or read from it.
Suspended	This unit is one of a mirrored pair. It is not capable of having data written to it or read from it. The data on this unit is not current. For example, if the disk needs repair action or has been manually suspended, it will be in a Suspended state.
Resuming	This unit is one of a mirrored pair. The current data is being copied (or will be copied) to this unit from the other active unit of the mirrored pair.
Unprotected	The device is in a state that cannot be determined.

Chapter 22. Working with auxiliary storage pools

You can use system service tools (SST) and dedicated service tools (DST) to work with auxiliary storage pools (ASPs). ASPs are also called *disk pools*. You can create new encrypted or unencrypted user ASPs, or add disk units to an existing ASP.

When you are making changes to the disk configuration on your system, see the Chapter 21, “Configuring disks and disk protection,” on page 369 topic for the correct sequence of steps for your situation.

You can use IBM Navigator for i to work with all types of ASPS, including independent ASPs.

Related concepts

Recovering an overflowed user auxiliary storage pool
Reset a user auxiliary storage pool (ASP) in overflowed status as soon as possible. An overflowed ASP affects system performance. It also makes recovery more difficult and might increase the amount of data that is lost if a failure occurs.

Related tasks

Creating objects in a nonlibrary user auxiliary storage pool
You can create journals, journal receivers, and save files in nonlibrary user auxiliary storage pool (ASP).

Related information

Managing independent disk pools
Managing disk pools

Initializing and formatting disk units

Use this information to initialize and format disk units without device parity protection.

To initialize and format disk units without device parity protection, follow these steps:

1. Start Dedicated Service Tools (DST) or System Service Tools (SST).
   Enter your service tools user ID and password.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 2 (Work with disk unit recovery) on the Work with Disk Units display.
   c) Select option 6 (Disk unit problem recovery procedures) on the Work with Disk Unit Recovery display.

   or from the System Service Tools (SST) menu, follow these steps:

   a. Select option 3 (Work with disk units).
   b. Select option 3 (Work with disk unit recovery) on the Work with Disk Units display.
   c. Select option 2 (Disk unit problem recovery procedures) on the Work with Disk Unit Recovery display.
3. Select option 1 (Initialize and format disk unit) on the Disk Unit Problem Recovery Procedures display.
4. Select option 1 (Select) on the Select Disk Units for Initialize and Format display for each disk unit that you want to initialize. Then press Enter.
5. Press F10 on the Confirm Initialize and Format Disk Unit display.
6. Wait for the message "Initialize and format completed successfully".
7. Press F3 to return to the DST or SST menu.

© Copyright IBM Corp. 1997, 2025

Adding disk units to an existing auxiliary storage pool

Use this information to add disk units to an existing auxiliary storage pool (ASP).

Do things in the right sequence:

- If you want to have device parity protection for the disks that you are adding, you should start device parity protection before you add the disk units to an ASP.
- If you do not want to have device parity protection for the disks that you are adding, you should initialize and format the disk units before you add them to an ASP.
- If you have more than one ASP on your system, you should plan how you want to add the new disk units before you begin this procedure.

When you (or your service representative) physically attach a new disk unit to your system, the status of the new disk unit is nonconfigured. `Nonconfigured` status means that a disk unit has not yet been assigned to an ASP on the system. You can assign disk units to an existing ASP. You create a new ASP by assigning disk units to it.

To assign nonconfigured disk units to an ASP, follow these steps:

1. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) On the Use Dedicated Service Tools (DST) display, select **Work with disk units**.
   b) On the Work with Disk Units display, select **Work with disk configuration**.
   c) On the Work with Disk Configuration display, select **Work with ASP configuration**.
   d) On the Work with ASP Configuration display, select **Add units to ASPs and balance data**.

   or from the System Service Tools (SST) menu, follow these steps:
   a. On the System Service Tools (SST) display, select **Work with disk units**.
   b. On the Work with Disk Units display, select **Work with disk configuration**.
   c. On the Work with Disk Configuration display, select **Add units to ASPs and balance data**.
2. On the Add Units to ASP and Balance Data display, select **Add units to existing ASPs and Balance Data**.

```
                Add Units to ASPs and Balance Data

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs and Balance data
```

3. On the Specify ASPs to Add Units to display, enter the ASP number (1 through 255) for each disk unit that needs to be added to an existing ASP. This display lists all the disk units that are not configured.

```
                       Specify ASPs to Add Units to

   Specify the existing ASP to add each unit to.

Specify  Serial                                  Resource
  ASP    Number      Type  Model  Capacity  Name
  1      21-6C597    4327   050      70564  DD007
  1      50-128840F  2107   A84      70564  DD004
  1      50-128940F  2107   A84      70564  DD005
  2      50-128A40F  2107   A85      35165  DD011
         50-128B40F  2107   A85      35165  DD003
         68-0C8BA12  6717   050       8589  DD008
         68-0C9D209  6717   050       8589  DD009
         68-606E0    6718   050      17548  DD006
         68-0CDAB10  6718   050      17548  DD010
         21-05348    4327   050      70564  DD012
         21-05322    4327   050      70564  DD013


F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

Note: If you are performing a complete system restore, all the disk units on the system might not report in right away. Verify that the number of disk units displayed matches the number of disk units physically attached to the system. If they do not match, wait a few minutes and press F5 (Refresh) until all of the disk units report in.

a) If you require more than one ASP, type an ASP number next to each disk unit that you want to configure.

   Number 1 is reserved for the system ASP. You can enter a number from 2 to 255. Numbers 33 to 255 are reserved for independent ASPs.

b) After you complete all units, press Enter.

4. The Confirm Add Units and Balance Data display shows what the entire system configuration will be when you add the units. Verify this configuration against your planned configuration.

```
                    Confirm Add Units and Balance Data

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter=Confirm Add units and Balance data.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data.
Press F12=Cancel to return and change your choice.

           Serial                  Resource
ASP  Unit  Number      Type Model   Name       Protection
  1                                           Unprotected
       1   68-0CDAB35  6718  050    DD001      Unprotected
       3   21-6C597    4327  050    DD007      Unprotected
       4   50-128840F  2107  A84    DD004      Unprotected
       5   50-128940F  2107  A84    DD005      Unprotected
  2                                           Unprotected
       2   50-100240F  2107  A82    DD002      Unprotected
       6   50-128A40F  2107  A85    DD011      Unprotected

F9=Resulting Capacity              F10=Add and Balance
F11=Display Encryption Status      F12=Cancel
```

5. You can press F9 (Resulting capacity) to see how the change will affect your disk utilization.

   You are shown the Resulting Capacity display:

```
                            Resulting Capacity
The configuration change that you requested would result in the
 following ASP capacities.

Press Enter to continue.
               ----------Current----------  ----------Proposed----------
               --Protected--  -Unprotected-  --Protected--  -Unprotected-
ASP  Threshold    Size  %Used    Size  %Used   Size  %Used   Size    %used
  1      90%         0  0.00%    70564 23.98%     0  0.00%   211692  8.00%
  2      90%         0  0.00%    35165 00.00%     0  0.00%    70330  0.00%
```

   a) Press Enter or F12 (Cancel) to return to the Confirm Add Units display.

6. Press Enter or F10 to confirm the selected disk units. If you want to make changes, press F12 to return to step “3” on page 392

   - If you press Enter, and you selected the option to balance data in step “1” on page 392, the system will balance the data in the ASP and move data to the newly added disk units.
   - If you press F10, the system will balance the data in the ASP and move data to the newly added disk units.

   Adding units can take from several minutes to several hours. During that time, you are shown the Function Status display.

```
                                Function Status

You selected to add units.




                                    5 % Complete
```

The system updates the display periodically.

Note: You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished adding disk units.

The time it takes the system to add units depends on the type, model, and size of each unit being added and the ability of the system to do multiple adds at the same time.

7. If you have no other tasks to perform, end DST or SST.

Related concepts

Managing auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) or IBM Navigator for i to manage auxiliary storage pools (ASPs).

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Stopping system service tools
To stop system service tools (SST), follow these steps.

Related information

Disk management

System values

Creating an unencrypted auxiliary storage pool

Use this information to create an unencrypted basic auxiliary storage pool (ASP) and add disk units to the ASP. You can create multiple ASPs at a time. You can assign disk units to a new ASP.

Do things in the right sequence:

- If you want to have device parity protection for the disks that you are adding, you should start device parity protection before you create an ASP.
- If you do not want to have device parity protection for the disks that you are adding, you should initialize and format the disk units before you create an ASP.
- You should plan how you want to add the new disk units to each ASP before you begin this procedure.

To create an independent ASP, you must use IBM Navigator for i or the Configure Device ASP (CFGDEVASP) command. For more information, see Configuring independent disk pools.

To create an unencrypted basic ASP, follow these steps:

1. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.

d) Select option 3 (Add units to ASPs) on the Work with ASP Configuration display.

or from the System Service Tools (SST) menu, follow these steps:

a. Select option 3 (Work with disk units).

b. Select option 2 (Work with disk configuration) on the Work with Disk Units display.

c. Select option 2 (Add units to ASPs) on the Work with ASP Configuration display.

2. On the Add Units to ASP display, enter 1 (Create unencrypted ASPs) to create one or more ASPs.

```
                          Add Units to ASPs

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs
```

3. On the Specify New ASPs to Add Units to display, enter the ASP number (2 through 32) for each disk unit that needs to be added to a new user ASP. This display lists all the disk units that are not configured.

```
                    Specify New ASPs to Add Units to

Specify the new ASP to add each unit to.

Specify  Serial                                    Resource
  ASP    Number           Type  Model  Capacity  Name
    2    21-6C597         4327   050      70564  DD007
    2    50-128840F       2107   A84      70564  DD004
    3    50-128940F       2107   A84      70564  DD005
    3    50-128A40F       2107   A85      35165  DD011
         50-128B40F       2107   A85      35165  DD003
         68-0C8BA12       6717   050       8589  DD008
         68-0C9D209       6717   050       8589  DD009
         68-606E0         6718   050      17548  DD006
         68-0CDAB10       6718   050      17548  DD010
         21-05348         4327   050      70564  DD012
         21-05322         4327   050      70564  DD013

F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

Note: If you are performing a complete system restore, all the disk units on the system might not report in right away. Verify that the number of disk units displayed matches the number of disk units physically attached to the system. If they do not match, wait a few minutes and press F5 (Refresh) until all of the disk units report in.

a) If you require more than one ASP, type an ASP number next to each disk unit that you want to configure.

Number 1 is reserved for the system ASP. You can enter a number from 2 to 32. Numbers 33 to 255 are reserved for independent ASPs.

b) After you complete all units, press Enter.

4. The Confirm Add Units display shows what the entire system configuration will be when you add the units. Verify this configuration against your planned configuration.

```
                              Confirm Add Units

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data on units.
Press F12=Cancel to return and change your choice.

            Serial                  Resource
ASP  Unit   Number      Type Model  Name       Protection
  1                                            Unprotected
       1    68-0CDAB35  6718  050   DD001      Unprotected
  2                                            Unprotected
       3    21-6C597    4327  050   DD007      Unprotected
       4    50-128840F  2107  A84   DD004      Unprotected
  3                                            Unprotected
       5    50-128940F  2107  A84   DD005      Unprotected
       6    50-128A40F  2107  A85   DD011      Unprotected

F9=Resulting Capacity              F10=Add and Balance
F11=Display Encryption Status      F12=Cancel
```

5. You can press F9 (Resulting capacity) to see how the change will affect your disk utilization.

   You are shown the Resulting Capacity display:

```
                              Resulting Capacity
The configuration change that you requested would result in the
 following ASP capacities.

Press Enter to continue.
               -----------Current----------  ----------Proposed----------
               --Protected--  -Unprotected-  --Protected--  -Unprotected-
ASP  Threshold    Size  %Used    Size  %Used    Size  %Used    Size  %used
  1      90%         0  0.00%    1967 23.98%    2950  0.07%    3934 12.02%
  2      90%                                   2950  0.07%       0  0.00%
```

   a) Press Enter or F12 (Cancel) to return to the Confirm Add Units display.

6. If you are satisfied with the configuration, press the Enter key to add the disk units to the ASP.

   If you want to make changes, press F12 to return to step “3” on page 395.

   Adding units can take from several minutes to several hours. During that time, you are shown the Function Status display.

```
                               Function Status

You selected to add units.









                                    5 % Complete
```

   The system updates the display periodically.

   Note: You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished adding disk units.

   The time it takes the system to add units depends on the type, model, and size of each unit being added and the ability of the system to do multiple adds at the same time.

7. If you have no other tasks to perform, end DST or SST.

8. If you created the user ASP (encrypted or unencrypted) using SST, you must perform a normal IPL to use integrated file system objects on the ASP. If you used DST to create the user ASP, you do not need to perform this IPL.

Creating an encrypted auxiliary storage pool

Use this information to create an encrypted basic auxiliary storage pool (ASP) and add disk units to it.

Important: If you installed option 45 and the system is being recovered from the most recent SAVSYS media, you need to take one of the following actions.

- Do an initial program load (IPL) of the system.
- Reinstall option 45 before you can create an encrypted ASP.

To create an independent ASP, you must use IBM Navigator for i or the **Configure Device ASP (CFGDEVASP)** command. For more information, see Configuring independent disk pools.

- If you want to have device parity protection for the disks that you are adding, start device parity protection before you create an ASP.
- If you do not want to have device parity protection for the disks that you are adding, initialize and format the disk units before you create an ASP.
- Plan how you want to add the new disk units to each ASP before you begin this procedure.

To configure an encrypted basic ASP, complete these steps:

1. If you are creating an encrypted ASP for the first time, install IBM i Option 45 (Encrypted ASP Enablement) by using the GO LICPGM command.

   Option 45 needs to be installed only one time.
2. Start Dedicated Service Tools (DST) or System Service Tools (SST).

   Enter your service tools user ID and password.
3. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
   d) Select option 3 (Add units to ASPs) on the Work with ASP Configuration display.

   Or from the **System Service Tools (SST)** menu, follow these steps:
   a. Select option 3 (Work with disk units).
   b. Select option 2 (Work with disk configuration) on the Work with Disk Units display.
   c. Select option 2 (Add units to ASPs) on the Work with ASP Configuration display.
4. On the 'Add Units to ASP' display, enter 2 (Create encrypted ASPs) to create encrypted ASPs.

```
                          Add Units to ASPs

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs
```

5. On the 'Specify New Encrypted ASPs to Add Units to' display, enter the ASP number to which you want to add disk units.

   User ASPs 2 through 32 can be encrypted. You can create multiple encrypted ASPs and add disk units to them.

```
                      Specify New Encrypted ASPs to Add Units to

Specify the new ASP to add each unit to.
All the new ASPs will be encrypted.

Specify  Serial                                      Resource
  ASP    Number           Type  Model  Capacity     Name
    2    21-6C597         4327   050      70564     DD007
    3    50-128840F       2107   A84      70564     DD004
         50-128940F       2107   A84      70564     DD005
         50-128A40F       2107   A85      35165     DD011
         50-128B40F       2107   A85      35165     DD003
         68-0C8BA12       6717   050       8589     DD008
         68-0C9D209       6717   050       8589     DD009
         68-606E0         6718   050      17548     DD006
         68-0CDAB10       6718   050      17548     DD010
         21-05348         4327   050      70564     DD012
         21-05322         4327   050      70564     DD013

F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

a) If you require more than one ASP, type an ASP number next to each disk unit that you want to configure.

   Number 1 is reserved for the system ASP. You can enter a number from 2 to 32. Numbers 33 to 255 are reserved for independent ASPs.

b) After you complete all units, press Enter.

6. The 'Confirm Add Units' display shows what the entire system configuration after you add the units. Verify this configuration against your planned configuration. Press F11 to display the encryption status of the ASP.

```
                                                          Confirm Add Units

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data on units.
Press F12=Cancel to return and change your choice.

            Serial                   Resource
ASP  Unit   Number      Type Model   Name       Protection
  1                                             Unprotected
       1    68-0CDAB35  6718  050    DD001      Unprotected
  2                                             Unprotected
       3    21-6C597    4327  050    DD007      Unprotected
  3                                             Unprotected
       2    50-128840F  2107  A84    DD002      Unprotected
```

7. If you are satisfied with the configuration, press the Enter key to add the disk units to the encrypted ASP.

   If you want to make changes, press F12 to return to step 8.

   Adding disk units can take from several minutes to several hours. During that time, you are shown the Function Status display.

```
                               Function Status

You selected to add units.




                                    5 % Complete
```

   The system updates the display periodically.

Note: You can press F16 to return to the **Use Dedicated Service Tools (DST)** menu if you have other tasks to complete. However, you cannot attempt any disk configuration tasks or end DST until the system finishes adding disk units.

The time that it takes the system to add units depends on the following factors.

- The type, model, and size of each unit that is added.
- The ability of the system to do multiple adds at the same time.

8. End DST or SST.
9. If you created the user ASP (encrypted or unencrypted) using SST, you must complete a normal IPL to use integrated file system objects on the ASP. If you used DST to create the encrypted user ASPs, you do not need to do this IPL.

Changing the storage threshold for an auxiliary storage pool

The storage threshold for an auxiliary storage pool (ASP) determines when the system warns you that the space allocated for the ASP is almost full. The default threshold for an ASP is 90%.

To change the storage threshold for an ASP, follow these steps:

1. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 2 (Work with disk configuration).
   c) Select option 3 (Work with ASP threshold).

   or from the Use Dedicated Service Tools (DST) menu, follow these steps:
   a. Select option 4 (Work with disk units).
   b. Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c. Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
   d. Select option 5 (Change ASP threshold) on the Work with ASP Configuration display.
2. The Select ASP to Change Threshold display is shown.

```
                      Select ASP to Change Threshold
   Type option, press Enter.
1=Select
                                        --Protected--   --Unprotected--
Option  ASP  Threshold  Overflow         Size  %Used      Size  %Used
          1        90%     No           6046  0.31%      7676  6.36%
   1      2        90%     No           2950    *           0  0.00%
```

3. On the Select the ASP to Change Threshold display, select the ASP that you want to have a different threshold.

   Press the Enter key. The following display is shown.

```
                           Change Storage Threshold
                                --Protected--  --Unprotected--
ASP   Threshold   Overflow      Size   %Used     Size  %Used
  2       90%       No          2950    *          0  0.00%

This is an unprotected ASP.  The threshold represents the amount of
unprotected storage used before a warning message is sent to the system
operator.   Type choice, press Enter.
New threshold . . . . . .     88% 1-100
```

4. Type your choice for the *New threshold* prompt and press the Enter key.
5. If you have no other tasks to perform, end DST or SST.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools

To stop dedicated service tools (DST), follow these steps.

Stopping system service tools
To stop system service tools (SST), follow these steps.

Changing the storage threshold for the system auxiliary storage pool
It is important to prevent the system auxiliary storage pool (ASP) from filling to capacity. If this occurs, the system ends abnormally. You can try to avoid this by specifying a storage threshold that warns you of a potential space shortage.

Related information

System values

Changing the storage threshold for the system auxiliary storage pool

It is important to prevent the system auxiliary storage pool (ASP) from filling to capacity. If this occurs, the system ends abnormally. You can try to avoid this by specifying a storage threshold that warns you of a potential space shortage.

One way to establish this threshold is through dedicated service tools (DST) or system service tools (SST). Use the same procedures as you would when setting the storage threshold for any other ASP.

Note: Establishing the threshold through DST will not prevent the system from ending abnormally. It will only notify you when the system ASP reaches the capacity threshold.

You can also protect the system ASP from filling to capacity by using the QSTGLOWLMT and QSTGLOWACN system values. The QSTGLOWLMT system value specifies the percentage of unallocated auxiliary storage that is remaining when the critical storage lower limit is reached. If the system reaches that limit, the QSTGLOWACN system value specifies what action the system should take. Using this method allows the system to actively prevent an abnormal shutdown instead of sending a warning of the condition.

Note: Using these system values does not affect any existing storage threshold that you might have set through DST.

You can use the QSTGLOWLMT and QSTGLOWACN system values on the following commands:

CHGSYSVAL	RTVSYSVAL
DSPSYSVAL	WRKSYSVAL

The following procedure demonstrates how to use these system values. (The Work with System Value (WRKSYSVAL) command is used as an example.)

1. At a command line, type WRKSYSVAL and press Enter.

   You are shown the Work with System Values display.

```
                          Work with System Values
                                                          System: YOURSYS
Position to  . . . . .   _______  Starting characters of system value
Subset by Type . . . .   _______  F4 for list

Type options, press Enter.
  2=Change   5=Display

         System
Option   Value       Type    Description
  _      QSTGLOWACN  *STG    Auxiliary storage lower limit action
  _      QSTGLOWLMT  *STG    Auxiliary storage lower limit
```

2. Type a 2 in the **option** field to change QSTGLOWACN and press Enter.

   You must have *ALLOBJ and *SECADM authority to change QSTGLOWACN. You are shown the Change System Value display.

```
                              Change System Value

System value . . . . . :   QSTGLOWACN
Description  . . . . . :   Auxiliary storage lower limit action


Type choice, press Enter.

  Action . . . . . .   *MSG_______     *MSG
                                       *CRITMSG
                                       *REGFAC
                                       *ENDSYS
                                       *PWRDWNSYS
```

3. On the Change System Value display, type the name of the action that you want the system to perform after reaching the critical storage lower limit.

   Press the Enter key. The actual actions that are performed by the action names are as follows:

   ***MSG**
   The system sends the CPI099C message to the QSYSMSG and QSYSOPR message queues. (The system also sends this message when you select any one of the other actions.)

   ***CRITMSG**
   The system sends the CPI099B critical message to the user who is specified in the service attribute to receive critical messages.

   ***REGFAC**
   The system submits a job to call exit programs that are registered for the QIBM_QWC_QSTGLOWACN exit point.

   ***ENDSYS**
   The system ends to the restricted state.

   ***PWRDWNSYS**
   The system powers down immediately and restarts

4. From a command line, type DSPSYSVAL and press the Enter key.

   The Display System Value display is shown.

```
                              Display System Value

System value . . . . . :   QSTGLOWLMT
Description  . . . . . :   Auxiliary storage lower limit

Lower limit  . . . . . :   1.0000        0-100 percent
```

   The lower limit value is the lowest amount of unused storage that can exist in the system ASP before the system performs the QSTGLOWACN action. (You can use the WRKSYSSTS command to view the amount of storage that is currently being used in the system ASP.) The system is shipped with the QSTGLOWLMT system value set to 5.0. Any change you make to this system value takes effect immediately.

Related tasks

Changing the storage threshold for an auxiliary storage pool
The storage threshold for an auxiliary storage pool (ASP) determines when the system warns you that the space allocated for the ASP is almost full. The default threshold for an ASP is 90%.

Related information

System values

Moving a disk unit to a different auxiliary storage pool

You might want to move a disk unit from one auxiliary storage pool (ASP) to another.

You want to create a user ASP for journal receivers and move some of the disk units on your system to the new user ASP. You can accomplish this in one process. When you move a disk unit to an ASP that does not exist, the system creates the ASP.

You might also decide to move disk units because you no longer need to have user ASPs on your system and you want to move all the disk units back to the system ASP.

Note: Disk units cannot be moved to or from an independent auxiliary storage pool.

Restrictions when changing your ASP configuration: Consider these things when you are planning to move disk units from an ASP:

- The system might take a long time to move the disk unit because it must copy the data from that disk unit to other disk units in the ASP.
- You cannot move unit 1 (the load source unit) from the system ASP.
- You cannot move disk units from a user ASP that is overflowed.
- You cannot move disk units in and out of the same ASP in the same operation.
- When mirrored protection is active for an ASP, you cannot move disk units in and out of the ASP. You must remove disk units in pairs from a mirrored ASP. You can then add them to a different ASP.
- When mirrored protection is active for the ASP that contains the disk units, you must remove both disk units of a mirrored pair.
- When you remove a disk unit, it becomes nonconfigured.

To move disk units between ASPs, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
3. Select option 6 (Move units from one ASP to another) from the Work with ASP Configuration display.
   The Specify ASP to Move Disk Units display is shown.

```
                    Specify ASP to Move Disk Units
  To move units to different ASPs, specify the ASP that you want to move each
one to in the 'New ASP' field.   Specify the units to be moved, press Enter.

  New  Current        Serial                    --Protected-- --Unprotected--
  ASP    ASP    Unit  Number      Type  Model   Size  %Used    Size   %Used
         1                                        0   0.00%    4124   41.50%
                  1   00-0193825  6602   030      0   0.00%    1031   82.00%
                  2   00-0163477  6602   030      0   0.00%    1031   29.00%
   2              3   00-0190494  6602   030      0   0.00%    1031   27.00%
   2              4   00-17900    6602   030      0   0.00%    1031   28.00%
```

4. Type the number of the ASP you want to move the disk units to in the New ASP column and press the Enter key.
   If you specify an ASP that does not currently exist on your system, the system creates a new ASP. If the move operation would leave the source ASP with insufficient storage, you receive an error message.

   If you see the Confirm Move of Unit display, skip to step 6.

   The Confirm Continuation display is shown if the storage management directories are not usable:

```
                              Confirm Continuation
   In order to proceed the system must perform internal processing that might
take several minutes during which the system might appear inactive.   Press
Enter to continue. Press F12=Cancel to return and change your choice.
```

5. Determine whether you want to cancel the procedure or continue.

   If you want to continue, press the Enter key.

6. The Confirm Move of Unit display is shown:

```
                            Confirm Move of Unit
   Moving units will take several minutes.   Press Enter to confirm your
choice to move the units. Press F9=Capacity information to display the capacity.
Press F12=Cancel to return to change your choice.
                Serial                    --Protected-- --Unprotected--
   ASP    Unit  Number      Type  Model   Size  %Used    Size  %Used
   1                                        0   0.00%    2062  83.00%
            1   00-0193825  6602  030
            2   00-0163477  6602  030
   2                                        0   0.00%    2062   0.01%
            3   00-0190494  6602  030
            4   00-17900    6602  030
```

   Press F9 (Capacity information) to display the resulting capacity.

```
                              Resulting Capacity
   The configuration change that you requested would result in the following
ASP capacities.   Press Enter to continue.
                   ----------Current----------  ----------Propose----------
                   --Protected--  -Unprotected-  --Protected--  -Unprotected-
 ASP   Threshold    Size  %Used    Size  %Used    Size   %Used   Size   %Used
  1        90%        0   0.00%    4124 41.50%      0    0.00%   2062   83.00%
  2        90%                                     0    0.00%   2062    0.01%
```

7. Press the Enter key to return to the Confirm Move of Unit display.
8. Press the Enter key on the Confirm Move of Units display to move the selected disk units.

   The system will move the data off the selected disk units to the remaining disk units in the source ASP. The move can take several minutes during which the system appears inactive.

9. When the move operation is complete, you return to the Work with ASP Configuration display.
10. If you have no other tasks to perform, end DST.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Removing a disk unit from an auxiliary storage pool

You need to remove a disk unit from an auxiliary storage pool (ASP) for these reasons.

- You want to physically remove a disk unit from your system.
- You want to move a disk unit from an ASP that has mirrored protection to another ASP. You can remove pairs of disk units from a mirrored ASP without stopping mirrored protection. You then add the disk units to the target ASP.

Considerations when changing your ASP configuration: Consider these things when you are planning to remove disk units from an ASP:

- The system might take a long time to remove the unit because it must copy the data from that unit to other units in the ASP.
- To proceed, the system must perform internal processing that might take several minutes during which the system might appear inactive.
- When you remove a disk unit, it becomes nonconfigured.

Restrictions when changing your ASP configuration: Consider these restrictions when you are planning to remove disk units from an ASP:

- You cannot remove unit 1 (the load source unit) from the system ASP.
- You cannot remove disk units from a user ASP that is overflowed.
- When mirrored protection is active for the ASP that contains the disk units, you must remove both units of a mirrored pair.
- If you are going to physically remove a disk unit from an input/output adapter (IOA), you must either exclude the disk unit or stop device parity protection first.

This task describes how to remove a disk unit with dedicated service tools (DST). For other ways to remove a disk unit, see Removing a disk unit from a disk pool.

To remove a disk unit, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.
   d) Select option 7 (Work with removing units from configuration) on the Work with ASP Configuration display.
   e) Select option 3 (Remove units from configuration) on the Work with Removing Units From Configuration display.
3. You are shown the Remove Units from Configuration display.

```
                      Remove Units from Configuration

Type options, press Enter.
  4=Remove unit from configuration

                 Serial                        Resource
OPT  Unit  ASP   Number      Type   Model      Name            Status
      2     1    10-00A7529  9332   400        DD010           Configured
      3     1    10-00A4936  9332   400        DD012           Configured
      4     1    10-00A4936  9332   400        DD019           Configured
 4    5     1    10-00A7498  9332   400        DD025           Configured
 4    6     1    10-00A7498  9332   400        DD036           Configured
      7     1    10-00A7530  9332   400        DD042           Configured
      8     1    10-00A7530  9332   400        DD052           Configured
```

4. Type a 4 (Remove unit from configuration) in the OPT column for each unit that you want to remove and press the Enter key.

   If the remove operation would leave the ASP with insufficient storage, you receive an error message.

   If you see the Confirm Remove Disk Units display, skip to 6.

   The Confirm Continuation display might be shown before the Confirm Remove Disk Units display if the storage management directories are not usable.

```
                              Confirm Continuation

To proceed, the system must perform internal processing
that might take several minutes during which the system might
appear inactive.

Press Enter to continue.
Press F12=Cancel to return and change your choice.
```

5. Determine whether you want to cancel the procedure or continue.

   If you want to continue, press the Enter key.
6. The Confirm Remove Disk Units display is shown:

```
                            Confirm Remove Disk Units

Removing disk units will take several minutes.

Press Enter to confirm remove of disk units.
Press F9=Capacity information to display the capacity information.
Press F12=Cancel to return to change your choice.


                  Serial                        Resource
OPT   Unit  ASP   Number      Type   Model      Name            Status
 4      5    1   10-00A7498  9332    400       DD010           Configured
 4      6    1   10-00A7498  9332    400       DD012           Configured
```

Press F9 (Capacity information to display the resulting capacity.

```
                               Resulting Capacity

The configuration change that you requested would result in the
following ASP capacities.

Press Enter to continue.
               ----------Current----------  ----------Modified--------
               --Protected--  -Unprotected-  --Protected-- -Unprotected-
ASP   Threshold  Size  %Used   Size  %Used   Size   %Used   Size   %Used
 1        90%       0   0.00%  1600  52.70%     0    0.00%  1200   70.26%
```

7. Press the Enter key to return to the Confirm Remove Disk Units display.
8. Press the Enter key on the Confirm Remove Disk Units display to remove the selected units.

   The system moves the data off the units selected to be removed to the remaining units in the source ASP. The remove can take several minutes or several hours during which the system appears inactive.

   Notes:

   - The time it takes to remove a unit depends on the disk unit type and model.
   - If the data on the unit being removed is severely fragmented and the amount of storage used is high, the remove operation might take several hours.

When the remove operation is complete, you return to the Work with ASP Configuration display. If you have no other tasks to perform, end DST.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Deleting an auxiliary storage pool

When you delete a user auxiliary storage pool (ASP), the status of all the disk units that are assigned to the ASP becomes nonconfigured. Any data that is in the ASP is deleted.

The normal procedure is to remove all objects from the ASP before you delete the ASP. You can do this either by moving the objects to a different ASP or by deleting the objects. If you delete an ASP that has objects in it, those objects are marked by the system as damaged or destroyed.

You cannot delete ASP 1, which is the system ASP and holds the operating system.

To delete a user ASP, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.

c) Select option 3 (Work with ASP configuration) on the Work with Disk Configuration display.

3. Select option 2 (Delete user ASP) on the Work with ASP Configuration display and press the Enter key.

```
                              Delete User ASP
  Type option, press Enter    4=Delete
                                        --Protected--  --Unprotected--
Option  ASP  Threshold  Overflow      Size  %Used     Size  %Used
         1       90%    No             600  77.84%       0  0.00%
         2       90%    No               0   0.00%     200  0.53%
         3       90%    No               0   0.00%     200  0.53%
```

4. Type a 4 in the **Option** field of the ASP you want to delete and press the Enter key.

   The Confirm Delete of User ASP display is shown.

```
                         Confirm Delete Of User ASP
  Warning: Deleting a user ASP will remove all units of that ASP from the
configuration. The units will become nonconfigured. Press F10 to confirm
your choice for 4=delete Press F12=Cancel to return to change your choice
                                    --Protected--  --Unprotected--
Option  ASP  Threshold  Overflow    Size  %Used     Size  %Used
   4      2       90%    No            0   0.00%     200  0.53%
```

5. Press F10 (Confirm) to confirm that delete of the ASP.

   The delete operation might take several minutes.

6. If you have no other tasks to perform, end DST.

Related concepts

Managing auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) or IBM Navigator for i to manage auxiliary storage pools (ASPs).

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping dedicated service tools
To stop dedicated service tools (DST), follow these steps.

Calculating space requirements for an auxiliary storage pool

When you are planning to change the disk configuration or disk protection on your system, you need to calculate the space requirements for the change before you begin.

You want to ensure that your system has sufficient disk storage for the changes. See Calculating disk space requirements.

Displaying the objects in a user auxiliary storage pool

To print a list of all the objects in a user auxiliary storage pool (ASP), use the Display Object Description (DSPOBJD) command and specify the object types that you want to display.

For a nonlibrary user ASP, specify object types *FILE (savf), *JRN, and *JRNRCV. For a list of objects that are supported in an independent ASP, see Supported and unsupported object types. The object description information includes the ASP where the object is located.

To list all the documents in a user ASP, use the Query Document Library (QRYDOCLIB) command:

```
QRYDOCLIB … QRYDFN(*IF(*ASP *EQ 4))
```

To determine which ASP an object is in, use the DSPOBJD command and look at the number shown on the **Auxiliary storage pool** field. To determine which ASP a document library object (DLO) is in, use the Display Document Library Object Name (DSPDLONAM) command. Look at the number shown on the **Auxiliary storage pool** field.

If the object is an integrated file system object, use the Display Object Links (DSPLNK) command. Select option 8 (Display attributes) to determine which ASP the object is in.

Balancing an auxiliary storage pool

The ASP balancing function improves system performance by balancing disk usage across all the disk units in an auxiliary storage pool (ASP).

You can use the Start ASP Balance (STRASPBAL) command to start the type of balancing that you want to use.

- Capacity balancing
- Usage balancing
- Hierarchical Storage Management (HSM) balancing
- Media preference balancing

Note: You cannot balance journal receivers among the disk units of an ASP if the journal does not have one of the receiver size options specified. If *MAXOPT1 and *MAXOPT2 are not specified for the RCVSIZOPT parameter on the **Created Journal** (CRTJRN) or **Change Journal** (CHGJRN) commands, the journal receiver can be spread across only 10 disk units. If either *MAXOPT1 or *MAXOPT2 is specified, then the journal receiver can be spread across as many as 100 units. Using RCVSIZOPT(*RMVINTENT) also affects how the journal receiver uses the disk units in the ASP.

Before usage balancing or hierarchical storage management (HSM) balancing is used, you must run the Trace ASP Balance (TRCASPBAL) command. This command starts a trace function that collects statistics on the data in the ASPs that you want to balance. Data that is used often is referred to as high use or hot data. Data that is not often used is referred to as low use or cold data.

To end the ASP balancing function, use the End ASP Balance (ENDASPBAL) command.

Capacity balancing

When you use capacity balancing, the data on the disk units within an auxiliary storage pool (ASP) is distributed evenly across all the units.

Instead of some units containing the majority of the data, each unit has an equal percentage of used and unused space. This type of balancing is useful when you add new disk units to an ASP.

Usage balancing

Usage balancing is useful when the auxiliary storage pool (ASP) contains some disk units that are used to a higher degree than other disk units in the ASP.

The Trace Auxiliary Storage Pool Balance (TRCASPBAL) command must finish collecting statistics before usage balancing can begin. When you use usage balancing, the high use and low use data on each unit in the ASP is redistributed to balance the arm utilization of each unit within the specified ASP.

Hierarchical storage management balancing

Hierarchical storage management (HSM) balancing can only be used auxiliary storage pools (ASPs) that contain a mixture of compressed and non-compressed disk units.

The Trace Auxiliary Storage Pool Balance (TRCASPBAL) command must finish collecting statistics before hierarchical storage management balancing can begin. When you use hierarchical storage management balancing, the high use and low use data on each unit in the ASP is redistributed. The high use data is moved to high performance units and the low use data is moved to low performance units. After the balance activity has completed, the system clears the trace information.

Media preference balancing

Media Preference (MP) balancing moves data between Hard Disk Drives (HDDs) and Solid State Disks (SSDs).

Data that has the Media Preference attribute is moved from HDDs to SSDs and data without the Media Preference attribute is moved from SSDs to HDDs. Media Preference balancing can only be used with auxiliary storage pools (ASPs) that contain a mixture of HDDs and SSDs.

Transferring objects between auxiliary storage pools

You can move entire libraries or folders from one auxiliary storage pools (ASPs) to another. Special procedures are used to move a library that contains journals because a journal and the journaled objects must be in the same basic user ASP or the same independent ASP group.

The "Working with nonlibrary user auxiliary storage pools" on page 415 topic discusses the procedures for working with nonlibrary user ASPs.

You cannot directly move objects between ASPs because the Move Object (MOVOBJ) command and the Move Document (MOVDOC) command move only the pointer to the object. They do not physically copy data from one location to another. In general, follow these steps to move an object to a different ASP:

1. Sign as QSECOFR.
2. Save the object and its private authorities by specifying the PVTAUT(*YES) parameter.
3. Delete the object from the system. If you are transferring the object from one independent ASP to another independent ASP, this step is not required.
4. Restore the object to the target ASP by using the RSTASP parameter on the RST*xxx* command. If you are restoring objects to an independent ASP, use the RSTASPDEV parameter. If you need to restore the private authorities of the object, use the PVTAUT(*YES) parameter.

These restrictions apply when specifying the RSTASP parameter for a basic user ASP:

- When you attempt to restore an object to a different basic ASP from its library, the ASP must be a nonlibrary user ASP and the object must be a journal, a journal receiver, or a save file.
  - For journals, journal receivers, and save files, if a library exists on the ASP, you receive an error message and the object is not restored.
  - For other object types, the object is restored to the ASP that contains the library.
- If you try to restore an object to a basic user ASP by explicitly specifying the user ASP that you want for the RSTASP parameter and the designated user ASP does not exist, you receive a message. The object is not restored.
- If you restore an object and specify RSTASP(*SAVASP), and if the basic ASP from which the object was saved no longer exists, the object is restored to the system ASP. You receive an informational message.

If you try to restore an object to an independent ASP by explicitly specifying the independent ASP that you want for the RSTASPDEV parameter and the designated independent ASP does not exist, you receive a message. The object is not restored.

Transferring a library to a different auxiliary storage pool

Use this procedure to move a library to a different auxiliary storage pool (ASP).

This example moves the CUSTLIB library from ASP 1 to ASP 2.

1. Sign on as QSECOFR.
2. Save the library and its private authorities:

```
SAVLIB  DEV(device-name) LIB(library-name)
        ASPDEV(asp1) PVTAUT(*YES)
```

3. Delete the library: `DLTLIB LIB(CUSTLIB)`

4. Restore the library and its private authorities to the new ASP:

```
RSTLIB  DEV(device-name) SAVLIB(library-name)
        RSTASPDEV(asp2) PVTAUT(*YES)
```

As an alternative to the preceding method, you can consider using the Move Library to ASP (QHSMMOVL) application programming interface (API) to transfer the library to a different ASP. The QHSMMOVL API handles object authorities as part of the library transfer process. You cannot use the QHSMMOVL API for libraries that contain journals, journal receivers, or objects that are journaled.

Related tasks

Restoring authorities for independent auxiliary storage pool data
You can use the methods shown in these figures to restore authority for an independent auxiliary storage pool (ASP).

Restoring private authorities for selected objects
If you need to recover or migrate a few objects when the user profiles already exist on the target system, the quickest method for restoring private authorities for selected objects is to specify PVTAUT(*YES) with any of the save and restore commands.

Related information

Application programming interfaces

Transferring a folder to a different auxiliary storage pool

Use this procedure to move a folder to a different auxiliary storage pool (ASP).

This example moves the HRFLR folder from ASP 1 to ASP 2. Do not move IBM-supplied folders (those starting with Q) to a user ASP. These folders must be in the system ASP.

1. Save the private authorities for the folder: `SAVSECDTA DEV(TAP01)`.
2. Save the folder: `SAVDLO DLO(*ALL) FLR(HRFLR) DEV(TAP01)`.
   Consider saving the object twice to 2 different media volumes.
3. Delete the folder: `DLTDLO DLO(*ALL) FLR(HRFLR)`.
   Do not skip this step. If you restore a document library object (DLO) to an ASP and it already exists in another ASP, you receive an error message. The restore operation continues with the next DLO. If you issue a restore command for a large number of DLOs, you will receive a large number of error messages.
4. Place the system in a restricted state: `ENDSBS *ALL *IMMED`.
5. Restore the private authorities you saved in step 1: `RSTUSRPRF USRPRF(*ALL) DEV(TAP01)`
6. Restore the folder to the new user ASP: `RSTDLO DLO(*ALL) SAVFLR(HRFLR) RSTASP(2)`
7. Restore authority to the folder and its objects: `RSTAUT`

You can move more than one folder at a time by specifying multiple folders on the Save Document Library Object (SAVDLO) and Restore Document Library Object (RSTDLO) commands. If you save DLOs from more than one ASP, you must specify sequence numbers on the RSTDLO command.

Transferring journals and objects to a different auxiliary storage pool

If you use a library user auxiliary storage pool (ASP), both the objects you are journaling and the journal must be in the same ASP.

For the purpose of recovery as well as performance, it is suggested that the journal receiver be placed in a different user ASP. If the journal is in the system or a basic user ASP, place the journal receiver in another basic user ASP. If the journal is in an independent ASP, place the journal receiver in secondary ASP when the journal is stored in the primary ASP of the disk pool group. If a failure occurs in the ASP that contains the objects and the journal, you do not lose both the objects and the journaled changes that are in the receiver. Placing your objects and journal receiver in the same user ASP also causes contention between access to the object and access to the journal receiver.

Use the following procedure to move a journal and the associated journaled objects to a different ASP. This procedure applies to library user ASPs (where the journal and its library are in the same ASP). If the journal is in a nonlibrary user ASP, see the"Working with nonlibrary user auxiliary storage pools" on page 415 topic.

To transferring journaled objects, follow these steps:

1. Sign on QSECOFR.
2. Save the journal using the Save Object (SAV), Save Object (SAVOBJ) or the Save Library (SAVLIB) command. Specify the PVTAUT(*YES) parameter on the save command to save the private authorities.
3. Because the journal and journaled objects must be in the same ASP, the objects must also be moved to the same user ASP before you can resume journaling the objects after the move.
4. Save any objects that you journal and save any logical files that have their access paths journaled. Specify the PVTAUT(*YES) parameter on each save command to save the private authorities.

   You can use the Work with Journal Attributes command to determine which objects that you journal. Consider saving the journal and journaled objects twice to two different media volumes.
5. Delete the journaled objects using the appropriate delete command.
6. Delete the journal using the Delete Journal (DLTJRN) command.
7. Delete the library that contained the journal.
8. Create the library for the journal in the user ASP using the Create Library (CRTLIB) command:
   - For basic ASPs, specify:

     ```
     CRTLIB LIB(library-name) ASP(asp-number)
     ```
   - For independent ASPs, specify:

     ```
     CRTLIB LIB(library-name) ASP(*ASPDEV) ASPDEV(asp-device-name)
     ```

   Note: The new library must have the same name as the library in which the journal was originally located.
9. Restore the journal to the library in the user ASP using the Restore Object (RSTOBJ) command.

   Specify PVTAUT(*YES) on the restore command to restore the private authorities. For independent ASPs, specify the RSTASPDEV parameter.
10. Restore the previously journaled objects to the library or directory in the user ASP.

    If you want to restore the previously journaled objects to their original libraries or directories, you must first move those libraries or directories to the user ASP. You move libraries and directories to a different ASP by saving them, deleting them, and restoring them to the new ASP. Specify PVTAUT(*YES) on each save and restore command.

    Restoring the previously journaled objects automatically resumes journaling for the objects if the journal already exists.
11. Save the journaled objects so that the journaled changes can be applied, if necessary.

    When journaling starts, the system assigns a journal identifier (JID) to the object. Typically the JID that is assigned is the same JID that the object had when it was saved. The object must be saved after the JID is assigned.

Creating objects in a library user auxiliary storage pool

You can create an object in a specific auxiliary storage pool (ASP) by placing it in a library or folder that is in the ASP.

You can place a library or a folder in an ASP in one of the following ways:

- Specify a value for the ASP parameter on the Create Library (CRTLIB) command or the Create Folder (CTRFLR) command. The default value for the ASP parameter is 1. For independent ASPs, you must also specify the ASP device name in the ASPDEV parameter. The CTRFLR command works for user ASPs only, but not for independent ASPs.
- Restore the library or the folder to a specific ASP by using the RSTASP parameter on the restore command.

Integrated file system objects can reside in a user ASP through the use of user-defined file systems (UDFSs) by following these steps:

1. Create a user defined file system in the ASP.
2. Mount the UDFS over another directory and use it through the mounted-over path.

Related information

User-defined file systems (UDFSs)

Example: Placing a document in a basic auxiliary storage pool

Here is an example of how to place a document in a specific basic auxiliary storage pool (ASP).

1. To create a folder in a basic ASP, use the ASP parameter on the Create Folder (CRTFLR) command:

```
CRTFLR FLR(ASP3FLR) ASP(3)
```

2. To create a document in that folder, use the Create Document (CRTDOC) command or a program that creates documents.

When you create a document or another folder in ASP3FLR, the new document or folder is automatically placed in ASP 3.

When you create the first folder in a basic ASP, the system creates the corresponding library. For example, when you create the ASP3FLR folder, the system creates the QDOC0003 library if it does not already exist. You should never create a QDOCnnnn library yourself. This can cause unpredictable results.

Example: Placing an object in a user auxiliary storage pool

Here is an example of how to place a journal receiver in a specific user auxiliary storage pool (ASP).:

1. Create a library for the journal receiver:
   - For basic ASPs, specify:

```
CRTLIB LIB(ASP2LIB) ASP(2)
```

   - For independent ASPs, specify:

```
CRTLIB LIB(library name) ASP(*ASPDEV) ASPDEV(asp-dev-name)
```

2. If the journal receiver is being place in an independent ASP, use the Set ASP Group (SETASPGRP) command to associate the job with the correct independent ASP.
3. Create the journal receiver in the library that you created in the user ASP:

```
CRTJRNRCV JRNRCV(ASP2LIB/RCVINASP2)
```

Example: Creating a user-defined file system in user auxiliary storage pool

Here is an example of creating a user-defined file system （UDFS） in a user auxiliary storage pool (ASP) using the Create User-Defined File System (CRTUDFS) command and the Add Mounted File System (MOUNT) command.

1. Create a UDFS in ASP 2.

   All objects created in this new file system will also reside in ASP 2:

   ```
   CRTUDFS UDFS('/dev/qasp02/asp2dir.udfs')
   ```

2. Mount the UDFS over directory '/myLocalPath':

   ```
   MOUNT TYPE(*UDFS) MFS('/dev/qasp02/asp2dir.udfs') MNTOVRDIR('/myLocalPath')
   ```

3. Create a directory in the UDFS:

   ```
   CRTDIR DIR('/myLocalPath/newDir')
   ```

Placing journal receivers in a user auxiliary storage pool

By placing journal receivers in a user auxiliary storage pool (ASP), you might improve your system performance.

Dedicating a user ASP to the journal receivers for a single journal is the suggested approach. Placing journal receivers in a user ASP simplifies recovery.

Placing journal receivers in a user auxiliary storage pool

Use this procedure to place journal receivers in a user auxiliary storage pool (ASP).

The following procedure uses the example of the CUSTJRN journal and journal receivers that use the naming convention CUSTR*nnnn*.

1. Create a library for the journal receiver in the specified user ASP:
   - For a basic ASP, specify:

     ```
     CRTLIB LIB(CUSTJRNR) ASP(4)
     ```

   - For an independent ASP, specify:

     ```
     CRTLIB LIB(CUSTJRNR) ASP(*ASPDEV) ASPDEV(asp-dev-name)
     ```

2. Use the Work with Journal Attributes (WRKJRNA) command to find the name of the currently attached journal receiver: `WRKJRNA JRN(CUSTJRNR/CUSTJRN)`.

   For the example, assume the currently attached journal receiver is CUSTR0005.
3. Use F17 (Display attributes) to determine the attributes for the current receiver.
4. If the journal receiver is being place in an independent ASP, use the Set ASP Group (SETASPGRP) command to associate the job with the correct independent ASP.
5. Create a new journal receiver in the library that is in the user ASP.

   Use a name that continues your naming convention. Specify the attributes that you displayed. For example, if the threshold of the current receiver is 100000, you would specify the following command:

   ```
   CRTJRNRCV JRNRCV(CUSTJRNR/CUSTR0006) THRESHOLD(100000)
   ```

6. Change journal receivers so the new journal receiver is attached and actively receiving journal entries:

```
CHGJRN JRN(CUSTJRN) JRNRCV(CUSTJRNR/CUSTR0006)
```

7. You can save the detached journal receiver and delete it from the system.
8. In the future, when you change journal receivers and specify JRNRCV(*GEN), the system will create the new journal receiver in the CUSTJRNR library in ASP 4.

Moving journal receivers from an overflowed basic user auxiliary storage pool

To maintain journaling for objects, follow these steps.

1. Use the Work with Journal Attributes (WRKJRNA) command to determine the names of the journal receivers associated with the journal: WRKJRNA JRN(*library-name/journal-name*)
2. Use F17 to display the attributes of the receiver that is attached.
3. If the journal receiver to be moved is attached to a journal, create a new journal receiver on a different auxiliary storage pool (ASP) by using the Create Journal Receiver (CRTJRNRCV) command.

   Assume the currently attached receiver is CUSTR0005. Use a name for the journal receiver that continues your naming conventions.

   If the journal receiver is in a user ASP, follow these steps:

   a. Create a new library in a different ASP, such as LIBJNEW: CRTLIB LIB(LIBJNEW) ASP(4)
   b. Create a new journal receiver in the library. Specify the attributes that you displayed. For example, if the threshold of the current receiver is 100000, you would specify the following:

      ```
      CRTJRNRCV JRNRCV(LIBJNEW/CUSTR0006) THRESHOLD(100000)
      ```

   If the journal receiver is in a nonlibrary user ASP, create a new journal receiver in a different nonlibrary user ASP or in the system ASP: CRTJRNRCV JRNRCV(CUSTJRNR/CUSTR0006) ASP(5)
4. Change the journal by using the Change Journal (CHGJRN) command.

   Specify the newly created journal receiver on the JRNRCV parameter: CHGJRN JRN(CUSTJRNR/CUSTJRN) JRNRCV(*library-name*/CUSTR0006)
5. Save the journal receivers from the overflowed user ASP.

   If the journal receivers are the only objects in the library, use the Save Library (SAVLIB) command. If other objects are in the library, use the Save Object (SAVOBJ) command.
6. If you used the SAVLIB command in step 5, delete the library from the overflowed user ASP by using the Delete Library (DLTLIB) command.

   If you used the SAVOBJ command, delete the journal receivers using the Delete Journal Receiver (DLTJRNRCV) command.
7. Journal receivers can be restored only to the library from which they were saved.

   The steps required to restore these journal receivers if you need them for a recovery depend on whether they were in a library user ASP or a nonlibrary user ASP. For journal receivers in a nonlibrary user ASP, you can restore them to any ASP, as long as their library is in the system ASP. For libraries that were in a library user ASP, you must ensure that the ASP has adequate space before restoring the journal receivers to the library.

Resetting a journal with a status of overflowed

If a journal object has a status of overflowed, you must delete it and restore it to reset its status.

Note: Use the Display Object Description (DSPOBJD) command to determine whether a specific object in a library has a status of overflowed. Use the Display Object Links (DSPLNK) command and select option 8 to determine whether a specific object in a directory has a status of overflowed.

Because journals and journaled objects must be in the same auxiliary storage pool (ASP), the best method for dealing with an overflowed journal is to restore it to the same basic ASP. If you restore the journal to a different ASP, you must also move all the journaled objects to that ASP.

If you want to move the journal and journaled objects to a different ASP, follow the procedure in the "Transferring journals and objects to a different auxiliary storage pool" on page 409 topic.

Before beginning this procedure, ensure that you have freed enough space in the overflowed ASP to prevent the journal from overflowing when it is restored.

1. Use the Work with Journal Attributes (WRKJRNA) command to print information about journaled objects and the receiver directory: `WRKJRNA JRN(library-name/journal-name) OUTPUT(*PRINT)`.
2. Use the Save Object (SAVOBJ) command to save the journal that must be reset.
3. Save the journal receivers that are associated with the journal by using the SAVOBJ command.
4. End journaling for any objects being journaled as follows:
   a) Access paths:

   ```
   ENDJRNAP JRN(library-name/journal-name) FILE(*ALL)
   ```

   b) Physical database files:

   ```
   ENDJRNPF JRN(library-name/journal-name) FILE(*ALL)
   ```

   c) Integrated file system objects:

   ```
   ENDJRN OBJ(*ALL) JRN('QSYS.LIB/library-name.LIB/journal-name.JRN')
   ```

   d) Issue the following command to record the inheritance rules and journaling attributes for each journaled library:

   ```
   DSPLIBD LIB(library-name) OUTPUT(*PRINT)
   ```

   e) Journaled libraries:

   ```
   ENDJRNLIB LIB(*ALL) JRN(library-name/journal-name)
   ```

   f) All other object types:

   ```
   ENDJRNOBJ OBJ(*ALL) OBJTYPE(*ALL) JRN(library-name/journal-name)
   ```

5. Deactivate any remote journals that are associated with the journal by using the Change Journal State (QjoChangeJournalState) API or the Change Remote Journal (CHGRMTJRN) command.
6. Delete the journal: `DLTJRN JRN(library-name/journal-name)`.
7. Restore the journal to the same library and the same ASP.

   If the journal was in library user ASP, you do not need to specify the ASP parameter on the Restore Object (RSTOBJ) command. If the journal was in a nonlibrary user ASP, specify `RSTASP(*SAVASP)` on the RSTOBJ parameter.
8. Start journaling again for each object that was journaled as follows:
   a) Database physical files:

   ```
   STRJRNPF FILE(library-name/file-name) JRN(library-name/journal-name)
   ```

   b) Access paths:

```
STRJRNAP FILE(library-name/file-name) JRN(library-name/journal-name)
```

c) Integrated file system objects:

```
STRJRN OBJ('object-path-name')
       JRN('/QSYS.LIB/library-name.LIB/journal-name.JRN')
```

d) Journaled libraries:

```
STRJRNLIB LIB(library-name) JRN(library-name/journal-name)
```

Use the inheritance rules and journaling attributes that you saved for the journaled libraries in step 4d.

e) All other object types:

```
STRJRNOBJ OBJ(library-name/file-name) OBJTYPE(object-type)
          JRN(library-name/journal-name)
```

You printed a list of objects in step 1.

9. Reestablish the journal receiver chain:
   a) Type WRKJRN and press the Enter key.
   b) On the prompt display, type the name of the journal and the library.
   c) On the Work with Journals display, select option 9 (Associate receivers with journal).
   d) When you receive a message that the receivers have been associated, press F12 to return.
   e) Type the following command:

```
WRKJRNA JRN(library-name/journal-name)
```

10. Use the Display Object Description (DSPOBJD) command to display the object description for the journal.

    Ensure that the journal is no longer in overflowed status.

Working with nonlibrary user auxiliary storage pools

Here are the object types that are allowed in a nonlibrary user auxiliary storage pool (ASP)..

- Journal
- Journal receiver
- Save file

Creating objects in a nonlibrary user auxiliary storage pool

You can create journals, journal receivers, and save files in nonlibrary user auxiliary storage pool (ASP).

When you create an object in a nonlibrary user ASP, the library for the object must exist in the system ASP. For example, to create a save file in a nonlibrary user ASP, type the following command:

```
CRTSAVF FILE(SAVFLIB/DSTSAV) ASP(4)
```

In this command, 4 is the number of the user ASP where you are placing the save file. The library for the save file is in the system ASP and ASP 4 does not contain any libraries.

After the object is created, all storage for the object resides in the designated user ASP. Changes and additions to that object are also made in the user ASP. If the ASP becomes full, it overflows into the system ASP.

It is suggested that all journals and journal receivers on the system have unique names. The Reclaim Storage (RCLSTG) command renames them if duplicate names are found when objects are placed in library QRCL and the user cannot rename them to their original name.

Monitor the size of objects to prevent them from overflowing into the system ASP with the MAXRCDS parameter on the Create Save File (CRTSAVF) command, and the THRESHOLD parameter on the Create Journal Receiver (CRTJRNRCV) command.

Related concepts

Working with auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) to work with auxiliary storage pools (ASPs). ASPs are also called *disk pools*. You can create new encrypted or unencrypted user ASPs, or add disk units to an existing ASP.

Transferring an object to a nonlibrary auxiliary storage pool

Use this example procedure to move an object to a nonlibrary auxiliary storage pool (ASP).

In this example, the DSTSAVF save file is moved to ASP 4. The library (SAVFLIB) for the DSTSAVF is in the system ASP. You can use this procedure only for save files, journals, and journal receivers.

To transfer an object to a nonlibrary ASP, follow these steps:

1. Sign on as QSECOFR.
2. Save the save file and its private authorities:

```
SAVOBJ  OBJ(DSTSAVF) LIB(SAVFLIB) OBJTYPE(*FILE)
        DEV(TAP01) PVTAUT(*YES)
```

3. Delete the save file: `DLTF FILE(DSTSAVF)`
4. Restore the save file and its private authorities to ASP 4:

```
RSTOBJ  OBJ(DSTSAVF) SAVLIB(SAVFLIB) RSTASP(4)
        DEV(TAP01) PVTAUT(*YES)
```

Transferring a journal to a nonlibrary user auxiliary storage pool

Use this procedure to move a journal to a different nonlibrary basic user auxiliary storage pool (ASP) and to reassociate any previously journaled objects with that journal.

1. Sign on as QSECOFR.
2. Save the journal with either the Save Object (SAV), Save Object (SAVOBJ), or Save Library (SAVLIB) command. Specify the PVTAUT(*YES) parameter to save the private authorities for the journal object.
3. List the objects being journaled: `WRKJRNA JRN(`*`library-name/journal-name`*`) OUTPUT(*PRINT)`
   a) Print the journaling attributes of the journaled files.

   ```
   DSPOBJD OBJ(file-name) OBJTYPE(*FILE) DETAIL(*FULL) OUTPUT(*PRINT)
   ```

   b) Print the inheritance rules and journaling attributes for each journaled library.

   ```
   DSPLIBD LIB(library-name) OUTPUT(*PRINT)
   ```

   c) Print the journaling attributes of the journaled data areas.

   ```
   DSPOBJD OBJ(data-area-name) OBJTYPE(*DTAARA) DETAIL(*FULL) OUTPUT(*PRINT)
   ```

   d) Print the journaling attributes of the journaled data queues.

```
DSPOBJD OBJ(data-queue-name) OBJTYPE(*DTAQ) DETAIL(*FULL) OUTPUT(*PRINT)
```

e) Print the journaling attributes of the integrated file system objects.

```
DSPOBJD OBJ(file-system-object-name) DETAIL(*EXTENDED) OUTPUT(*PRINT)
```

4. Stop journaling access paths for the journal: `ENDJRNAP FILE(*ALL) JRN(library-name/journal-name)`
5. Stop journaling physical files for the journal: `ENDJRNPF FILE(*ALL) JRN(library-name/journal-name)`
6. Stop journaling integrated file system objects: `ENDJRN OBJ(*ALL) JRN('/QSYS.LIB/library-name.LIB/journal-name.JRN'/)`
7. Stop journaling libraries for the journal: `ENDJRNLIB LIB(*ALL) JRN(library-name/journal-name)`
8. Stop journaling all other object types for the journal: `ENDJRNOBJ OBJ(*ALL) OBJTYPE(*ALL) JRN(library-name/journal-name)`
9. Deactivate any remote journals that are associated with the journal by using the Change Journal State (QjoChangeJournalState) API or the Change Remote Journal (CHGRMTJRN) command.
10. Delete the journal with the Delete Journal (DLTJRN) command.
11. Restore the journal and its private authorities to the user ASP: `RSTOBJ OBJ(journal-name) OBJTYPE(*JRN) RSTASP(asp-number) PVTAUT(*YES)`
12. Use the Start Journal Physical File (STRJRNPF), Start Journal Access Path (STRJRNAP), Start Journal Object (STRJRNOBJ), Start Journal Library (STRJRNLIB), and Start Journal (STRJRN) commands to resume journaling for any previously journaled objects.

    Refer to the lists that you printed in step 3. When resuming journaling of the previously journaled objects, use the journaling attributes and inheritance rules information that you saved in step 3.

Placing journal receivers in a nonlibrary basic auxiliary storage pool

Use this procedure to place journal receivers in a nonlibrary basic auxiliary storage pool (ASP).

The following procedure uses the example of the CUSTJRN journal and journal receivers that use the naming convention CUSTR*nnnn*. This example assumes that the library for the journal receiver (CUSTJRNR) already exists in the system ASP.

1. Use the Work with Journal Attributes (WRKJRNA) command to find the name of the currently attached journal receiver: `WRKJRNA JRN(CUSTJRNR/CUSTJRN)`.
   For the example, assume that the currently attached journal receiver is CUSTR0005.
2. Use the Display Journal Receivers Attributes (DSPJRNRCVA) command to determine the attributes for the current receiver: `DSPJRNRCVA JRNRCV(CUSTJRNR/CUSTR0005)`
3. Create a new journal receiver in a user ASP.

   Use a name that continues your naming convention. Specify the attributes that you displayed, such as the threshold:

```
CRTJRNRCV JRNRCV(CUSTJRNR/CUSTR0006) ASP(4)
THRESHOLD(100000)
```

4. Change journal receivers so the new journal receiver is attached and actively receiving journal entries:

```
CHGJRN JRN(CUSTJRN) JRNRCV(CUSTJRNR/CUSTR0006)
```

5. You can save the detached journal receiver and delete it from the system.
6. In the future, when you change journal receivers and specify `JRNRCV(*GEN)`, the system will create the new journal receiver in the CUSTJRNR library in ASP 4.

Chapter 23. Working with device parity protection

You can use system service tools (SST) and dedicated service tools (DST) to work with device parity protection on your system.

When you are making changes to the disk configuration or disk protection on your system, you need to perform tasks in the correct sequence. See the “Choosing the right procedure for configuring disks” on page 369 topic to determine the correct tasks for your situation.

Related information

Disk protection

Starting device parity protection

Often, you start device parity protection when you attach a new disk subsystem to your system.

Ideally, you should start device parity protection before you add the disk units to the software disk configuration. When you start device parity protection for disk units that are already being used, the system must move data off of the disk units to other disk units in the auxiliary storage pool (ASP) in order to make space available for the parity data. The performance for starting device parity protection is much better if the system does not have to move data.

When you start device parity protection, the system does validity checking and moves data from the required units, if necessary. For some types of disk units, you or your service representative must perform tasks with the disk subsystem when you start device parity protection.

Note: If you plan to start device parity protection for disk units that are already part of your disk configuration, check the following before you start device parity protection.

- The configuration must be complete and no disk units can be missing in any ASPs that contain disk units that are to have device parity protection. This is because the system must move data off the disks that are to be protected to make room for parity information.
- The disk units that will become device-parity protected cannot be in an ASP that has mirrored protection active. If the disk units are in an ASP that has mirrored protection, you must stop mirrored protection before starting device parity protection.
- When you start device parity protection, you reduce the capacity of some of the disk units in the subsystem. The system must have sufficient storage in each affected ASP to make room for redundant parity data.

Starting device parity protection for an input/output adapter

You can start device parity protection for an input/output adapter (IOA).

The topic lists the rules and describes the procedure for starting protection on IOAs:

- Maximum number of parity sets allowed - 8 (some adapters might support fewer parity sets)
- Maximum number of devices per parity set - 18
- Minimum number of devices per parity set - 3 for RAID 5
- Minimum number of devices per parity set - 4 for RAID 6
- All devices in a parity set must be the same capacity

The IOA starts the fewest number of parity sets needed to protect all the devices of the same capacity. For example, to protect 10 devices, it starts one parity set of ten devices. To protect 11 devices, it starts two parity sets: one parity set of seven devices and one parity set of four devices. You can use IBM Navigator for i to optimize your parity sets for availability, capacity, performance, or balanced (a mixture of the two).

© Copyright IBM Corp. 1997, 2025

Hot spare is a function that is built into the input/output adapter. When a disk unit that is protected with device parity and hot spare fails, the IOA immediately replaces the failing disk with a standby disk and starts the device parity rebuild process. These functions take place at the time of the failure without the need for manual intervention. The hot spare feature reduces exposure to secondary failures for other devices in the parity set. It also prevents loss of data in the parity set.

Note: Displays might differ.

When you work with device parity protection, the displays you see might differ slightly from the displays in these examples because of differences in disk subsystem requirements. The best way to verify that you are running the correct step in each procedure is to refer to the titles of the displays.

1. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 5 (Work with device parity protection) on the Work with Disk Configuration display.
2. Select option 2 (Start device parity protection) on the Work with Device Parity Protection display and press the Enter key.

   You are shown the Start Device Parity Protection display where you can select to start the following functions:

   - RAID 5 or RAID 6 protection.
   - Device parity with the hot spare function.

```
                    Start Device Parity Protection
   Select the subsystems to start device parity protection.
   Type choices, press Enter.
   1=Start device parity protection
                 Parity   Serial                      Resource
         Option    Set   Number      Type   Model   Name
          1          2   3314025     6502   001     SI01
                     1   0000318     9337   225     DD013
          -      ------  ----------  -----  -----   --------
```

3. Type a 1 in the Option column for the disk unit subsystems that you want to prepare to start device parity protection. Press the Enter key.

   If you are shown the following display, press Enter to continue.

```
                            Confirm Continuation
   In order to proceed the system must perform internal processing that might
take several minutes during which the system might appear inactive.  Once you
confirm to continue, the system must perform an IPL when you leave Work with
Disk Configuration functions.

Press Enter to continue.
Press F12=Cancel to return to change your choice.
```

4. Press the Enter key to continue.

   You are shown the Confirm Starting Device Parity Protection display. The display shows all the disk unit subsystems that you selected and the individual disk units that are eligible to be started. Disk units that have an asterisk (*) in the ASP and Unit columns are not yet configured.

```
                  Confirm Starting Device Parity Protection
During the
preparation for starting device parity protection, data will be moved from
parts of some disk units.  This might take several minutes for each subsystem
selected.   Press Enter to continue. Press F12=Cancel to return and change
your choice.
                 Parity            Serial                        Resource
         Option    Set  ASP  Unit  Number      Type  Model    Name
            1        2              3314025     6502   001     SI01
            1        2   *     *   00-0051556  6603   074     DD056
            1        2   *     *   00-0020525  6603   074     DD025
            1        2   *     *   00-0024030  6603   074     DD030
            1        2   *     *   00-0013026  6603   074     DD026
            1        2   *     *   00-0024519  6603   074     DD019
            1        2   *     *   00-0046338  6603   074     DD038
```

5. Notice - Pressing the Enter key initiates the procedure for starting device parity protection.

   After begun, this procedure continues to run until it is complete. If the subsystems that you selected are correct, press the Enter key to continue. The status display shows how the operation is proceeding. When the system completes its preparation for starting device parity protection on the selected subsystems, the following display is shown.

```
                   Start Device Parity Protection Status
   The operation to start device parity protection will be done in several
phases.  The phases are listed here and the status will be indicated when
known.
Operation                                    Status
Initialize disk units  . . . . . . . . . . :  Running
Prepare to start . . . . . . . . . . . . . :  ____________________
Start device parity protection . . . . . :  ___ %
```

6. Press the Enter key to return to the Work with Device Parity Protection menu.

Related information

Disk pools

Stopping device parity protection

When preparing to stop device parity protection, the system does validity checking to make sure that stopping device parity protection does not leave the system in a configuration that is not supported.

The following restriction applies when you stop device parity protection:

- You cannot stop device parity protection on a subsystem when a unit in that subsystem is in a mirrored auxiliary storage pool (ASP). In order to stop device parity protection, mirrored protection must be stopped first.

Stopping device parity protection on an input/output adapter

These instructions only apply to input/output adapters (IOAs) that support device parity protection.

To stop device parity protection on an IOA, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 5 (Work with device parity protection) on the Work with Disk Configuration display.
3. Select option 3 (Stop device parity protection) on the Work with Device Parity Protection display and press the Enter key.

   The following display is shown.

```
                        Stop Device Parity Protection
   Select the subsystems to stop device parity protection.
   Type choices, press Enter.
   1=Stop device parity protection
                   Parity  Serial
Resource Option     Set   Number      Type   Model  Name
            1         1   10-4453031  6502    001   SI01
                      2   00-00341    9337    212   DC09
```

4. Type a 1 in the Option column for the disk unit subsystems that you want to stop device parity protection.

   Press the Enter key. The following display is shown.

```
                    Confirm Stop Device Parity Protection
   Warning: Disk units connected to these subsystems will not be protected
after you confirm your choices.
   Press Enter to continue.
   Press F12=Cancel to return and change your choices.
                  Parity            Serial
Resource Option    Set  ASP  Unit  Number      Type  Model  Name
            1        1             10-4453031  6502   001   SI01
            1        1    8     3  00-0334602  6602   050   DD012
            1        1    8     4  00-0334673  6602   050   DD011
            1        1    8     5  00-0231915  6602   050   DD010
            1        1    8     6  00-0334610  6602   050   DD009
            1        1    8     7  00-0274937  6602   050   DD008
```

5. Notice: At this point, pressing the Enter key initiates the procedure for stopping device parity protection.

   After this procedure has begun, you cannot cancel it. If the subsystems you have selected are correct, press the Enter key to continue. You will get status displays.

```
                     Stop Device Parity Protection Status
   The operation to stop device parity protection will be done in several
phases.  The phases are listed here and the status will be indicated when
known.
Operation                                      Status
Prepare to stop  . . . . . . . . . . . . . :  Completed
Stop device parity protection  . . . . . :  Completed
   WARNING: There are now unprotected disk units on this system. When a system
has unprotected, exposed, or suspended disk units attached to it, disk related
failures might affect the availability of the system and can cause loss of data.
```

   Note: If you have not yet received "Completed" status you can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished starting device parity protection.

6. When the status shows Completed, press the Enter key to return to the Work with Device Parity Protection menu.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Removing disk units that have device parity protection for System i while powered on

Here are the detailed steps for removing disk units that have device parity protection for your system while powered on.

1. **Attention:** If possible, use a wrist strap to prevent electrostatic discharge from damaging your hardware.

   If you do not have a wrist strap, touch a metal surface of the system or expansion unit before installing or replacing hardware.

If you received the 2209 Disposable Wrist Strap, follow these steps:

a. Attach the disposable wrist strap.

b. When you unroll the strap, remove the liner from the copper foil at the end.

c. Attach the adhesive part of the copper foil to an exposed, unpainted metal surface on the frame of the unit.

Note: Follow the same precautions you would use without the wrist strap. The 2209 Disposable Wrist Strap is for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.

2. Find the Disk unit locations for your system unit or your expansion unit.
3. Remove the disk unit cover. Squeeze the latch and pull out.
4. Be sure you are signed on with at least service level authority.
5. Type `strsst` on the command line and press Enter.
6. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display and press Enter.

   Note: The service tools password is case sensitive.
7. Select **Start a service tool** from the System Service Tools (SST) display. Press Enter.
8. Select **Hardware service manager** from the Start a Service Tool display. Press Enter.
9. Select **Device Concurrent Maintenance** from the Hardware Service Manager display. Press Enter.
10. Type the Physical Location (Frame ID and position) where you will be replacing the disk unit.

    Note: The physical location of the disk unit to be replaced should have been provided by customer support.

    In this example, Physical Location: Frame ID 1 Position D06, you would install a disk unit in frame 1 (base system), disk unit position 6.
11. Select option 1 (Remove device) for the action to be performed.
12. Set the time delay for 1 minute—01.

    Do not press Enter at this time.

    Note: Locate the light in the numbered arrow by the disk unit you are replacing. When you press Enter, after a one minute delay, this light comes on and begins to flash rapidly. You then have 9 seconds to slide the disk unit into place.
13. Press Enter on the console.
14. When the light starts flashing, squeeze and pull the handle of the disk unit out toward you before you remove the disk unit.

    If the handle is not all the way out, the disk unit will not slide out of the system.
15. Support the bottom of the disk unit assembly as you slide it completely out.

    Do not hold the disk unit by the handle.
16. The light will stop flashing and remain off when the disk unit is removed.
17. Return to the console and wait until the Concurrent Maintenance Results display appears. Press F12.
18. Find the package that contains the new disk unit assembly and remove it from the static-protective package.

    Attention: Disk units are fragile. Handle with care.
19. Unlock the handle by squeezing and pulling it out toward you before you install the disk unit.

    If the handle is not all the way out, the disk unit will not slide into the system unit or the expansion unit.
20. Support the bottom of the disk unit assembly as you slide it half way into the system unit or the expansion unit.

    Do not push it any further.

Note: You install the disk units horizontally in the 7123, 7136, and 7137 disk unit enclosure.

21. The physical locations you entered in step 11 of this procedure might still appear on the display.

    If not, retype the Physical Location (Frame ID and position) where you will be installing the new disk unit. In this example, Physical Location: Frame ID 1 Position D06, you will install a disk unit in frame 1 (base system), disk unit position 6.
22. Select **option 2 (Install device)** for the action to be performed.
23. Set the time delay for one minute—01.

    Do not press Enter at this time.

    Note: Locate the light in the numbered arrow by the disk unit you are replacing. When you press Enter, after a one minute delay, this light comes on and begins to flash rapidly. You then have nine seconds to slide the disk unit into place.
24. Press Enter on the console.
25. When the light starts flashing, slide the disk unit all the way in and then lock it in place by closing the latch.
26. The light stops flashing and remains on when you have installed the disk unit properly.
27. Return to the console and wait until the Concurrent Maintenance Results display appears. Press Enter.
28. When the Hardware Service Manager display appears, press F3.
29. Select **Work with Disk Units** on the System Service Tools display. Press Enter.
30. Select **Work with disk unit recovery** on the Work with Disk Units display. Press Enter.
31. Select **Rebuild disk unit data** on the Work with Disk unit recovery display. Press Enter.
32. Select **1** to rebuild the disk unit displayed (disk unit displayed is the disk unit that you removed) on the Rebuild Disk Unit Data display. Press Enter.
33. Press Enter on the Confirm Rebuild Disk Unit Data display. The rebuild process might take several minutes to complete.
34. Install the disk unit covers.

    Align the tabs on the disk unit cover with the slots in the system unit. Push the cover in until it latches.
35. Install or close the system unit or expansion unit front cover.
36. Press F5 to refresh the display until the **Percent complete** filed shows 5%.
37. When the display shows at least 5% complete, you can either continue to monitor this display to completion, or press F3 (Exit) to return to the Work with disk units display.
38. Press F3 (Exit) to return to the System service tools display.
39. Press F3 (Exit) to return to the exit SST display and press Enter.

The replacement procedure is complete.

Removing disk units that have device parity protection for System i while powered off

The topic details the steps for removing disk units that have device parity protection for your system while your system is powered off.

1. Press F3 repeatedly to exit System Service Tools and return to the main menu.
2. Power off the system unit.
3.

   Attention: If possible, use a wrist strap to prevent electrostatic discharge from damaging your hardware. If you do not have a wrist strap, touch a metal surface of the system or expansion unit before you install or replace hardware.

   If you received the 2209 Disposable Wrist Strap, follow these steps:

a. Attach the disposable wrist strap.

b. When you unroll the strap, remove the liner from the copper foil at the end.

c. Attach the copper foil adhesive to an unpainted metal surface on the frame of the unit.

Note: Follow the same precautions that you use without the wrist strap. The 2209 Disposable Wrist Strap is for static control. It does not increase or decrease your risk of receiving electric shock when you use or work on electrical equipment.

4. Remove the disk unit covers. Squeeze the latch and pull out.
5. Locate the disk unit that you want to replace.
6. Unlock the handle by squeezing it and pulling it out toward you before you remove the disk unit.

   If the handle is not all the way out, the disk unit cannot slide out of the system.
7. Support the disk unit assembly as you slide it completely out.

   Do not hold the disk unit by the handle.
8. Find the package that contains the new disk unit assembly and remove it from the static-protective package.

   Write the serial number (use the last 4 digits) of the disk unit here _____ before you install it.

   Attention: Disk units are fragile. Handle with care.
9. Squeeze and pull out the handle toward you before you install the disk unit.

   If the handle is not all the way out, the disk unit does not slide into the system unit or the expansion unit. Do not hold the disk unit by the handle.
10. Make sure that the disk unit handle is fully extended when you slide it into the system unit or the expansion unit.

    Support the disk unit assembly as you slide it into the system unit or the expansion unit until it stops.
11. Close the disk unit handle firmly, locking it into place.
12. Install the disk unit covers.

    The enclosure cover for the disk unit has tabs. Align the tabs with the slots on the system unit enclosure. Push the cover in until it latches.
13. Install or close the unit front cover.
14.

    DANGER: An electrical outlet that is not correctly wired can place hazardous voltage on metal parts of the system or the products that attach to the system. It is the customer's responsibility to verify that the outlet is correctly wired and grounded to prevent an electrical shock. (RSFTD201) Plug the following power cords into an electrical outlet.

    - System unit
    - System unit console
    - Expansion units
15. Plug in and power on all attached devices, such as printers and displays.
16. Activate your logical partition by using the HMC.

    Wait for approximately 10 seconds before the power-on light illuminates and data appears in the Function/Data display.

    Note: The system takes several minutes to power on to complete an initial program load (IPL). When the IPL is complete, the IBM i signon display appears on your system unit console.
17. Follow these steps to verify your disk unit configuration:

    a) Type `strsst` on the command line of the IBM i session and press Enter.

    b) Select **Work with disk units** from the System Service Tools (SST) display. Press Enter.

    c) Select **Work with disk unit recovery** from the Work with Disk Units display. Press Enter.

    d) Select **Rebuild disk unit data** on the Work with Disk unit recovery display. Press Enter.

e) Select **1** to rebuild the disk unit displayed (the disk unit that is displayed is the disk unit that you removed) on the Rebuild Disk Unit Data display. Press Enter.

f) Press Enter on the Confirm Rebuild Disk Unit Data display. The rebuild process might take several minutes to complete.

g) Install the disk unit covers. Align the tabs on the disk unit cover with the slots in the system unit. Push the cover in until it latches.

h) Install or close the system unit or the front cover of the expansion unit.

i) Press F5 to refresh the display until shows 5% complete.

j) When the display shows at least 5% complete, you can either monitor this display through completion or press F3 (Exit) to return to the Work with disk units display.

k) Press F3 (Exit) to return to the System service tools display.

l) Press F3 (Exit) to return to the exit SST display and press Enter.

Removing disk units that have device parity protection for a IBM i

Use these steps to remove disk units that have device parity protection for your system.

1. Prepare to remove the disk drive from the system configuration.
   a) Determine the protection status of the disk drive that you are replacing.
   b) A service provider must replace the unprotected disk drives.
      Before you replace the disk drive, the service provider must do one of the following tasks -
      - Save the data and initialize the disk drive.
      - Initialize the disk drive and copy to disk.
2. Locate the disk drive to be removed.
3. Stop the system or logical partition that controls the disk drive location.
4. If you stopped the logical partition, do not disconnect the power source.
   If you stopped the system, disconnect the power source from the system by unplugging the system.

   DANGER: This system might be equipped with a second power supply. Before you continue with this procedure, verify that all power sources to the system are disconnected .
5. Remove the system unit door or expansion unit door, or open the rack front door.
6. Remove the front cover.
7. Unlock the disk drive handle by squeezing it and pulling it out toward you.

   Attention: Attach a wrist strap to an unpainted metal surface of your hardware to prevent electrostatic discharge from damaging your hardware. When you use a wrist strap, follow all electrical safety procedures. A wrist strap is for static control. It does not increase or decrease your risk of receiving electric shock when you use electrical equipment. If a wrist strap is not available, before you remove the product from ESD packaging and install hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds.
8. Support underneath the disk drive as you slide it out of the system or expansion unit.
   Do not hold the disk drive by the handle.
9. If a disk drive filler is available, follow these steps:

   Note: IBM suggests that disk drive slots be refilled with either another disk drive or a disk drive filler. Filling the disk drive slot helps to cause proper air flow for cooling and helps maintain optimal electromagnetic interference (EMI) compliance. Order feature 6598 to have four more disk slot fillers included with the order.

   a. Attach the disk drive bezel to the filler.
   b. Press in and then tighten thumbscrew B to secure the bezel to the disk drive.

c. Install the disk drive filler in the system or expansion unit

10. Replace the front cover.
11. Install or close the system, expansion unit, or rack front door.
12. Reconnect all system power cables.
13. Reconnect the system to the power supply by plugging it in.
14. Start the system or logical partition.

    To start a system that does not use a Hardware Management Console (HMC), follow these steps:

    a. Before you press the power button on the system control panel, verify that power is connected to the system unit.

       Note: On the OpenPower 710 or model 9110-510, the power button is on the operations panel. All system power cables are connected to a power source. The power-on light F is slowly flashing. Display D contains 01 V=F.

    b. Press the power button A on the control panel.

       Note: When you press the power button the system does not start, use the Advanced System Management Interface (ASMI) to start the system. First, set up access to the ASMI.

       i) Start the system by using the ASMI.

       ii) After you press the power button, verify that the power-on light flashes visibly faster. The system cooling fans are activated after approximately 30 seconds and accelerate to operating speed. Progress indicators, also referred to as checkpoints, appear on the control panel display while the system is being started. The power-on light on the control panel stops flashing and remains on, indicating that system power is on

    Start the system by using the HMC.

    a. In the navigation area, open **System and Partition**.
    b. Select **System Management**.
    c. In the contents area, open the server on which the partition profile is located.
    d. Open **Partitions**.
    e. Open the logical partition for the partition profile.
    f. Right-click the partition profile and select **Activate**.
    g. Select the partition profile that you want to activate and click **OK**.

The replacement procedure is complete.

Including a disk unit in device parity protection

When you attach a new disk unit to an existing input/output adapter (IOA) that has device parity protection, you can include the disk unit into an existing parity set with other disk units.

Attention: Including disk units into an existing parity set with PCIe2, PCIe3, and newer generations of SAS RAID adapters can result in a large performance degradation. The degradation can occur because the adapter uses firmware in place of specialized hardware assists when drives are included into an existing parity set. For this reason, the Include disk operation is not recommended with PCIe2, PCIe3, and newer generations of SAS RAID adapters. Systems with POWER8 technology and beyond do not allow the Include disk operation. Alternatives for adding disk units to the system are to stop parity on the existing set and then start parity with the additional disks in the set, or to create a new parity set of three or more disks with the additional disk units. Adding three or more disk units requires you to create a new parity set.

If you want to include a disk unit in a parity set and the disk units of that parity set are in an independent disk pool, the independent disk pool must be varied off. For all other disk pools, you must restart your system to dedicated service tools (DST) mode before you can include them in a parity set.

For a Small Computer System Interface (SCSI) IOA, the disk units to include must be the same capacity as the other disk units in the parity set. For a Serial Attached SCSI (SAS) IOA, the disk units to include can be the same capacity or larger than the other disk units in the parity set.

Note: You cannot include a disk unit if that disk unit has already been added to an auxiliary storage pool (ASP) that has mirrored protection. You must stop mirrored protection before including the disk unit. Stopping mirrored protection must be done from the DST menu. Adding mixed protection on the same IOP requires mirroring to be stopped and restarted.

To include disk units in a device parity set, perform the following steps:

1. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 2 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 8 (Work with Device Parity Protection) on the Work with Disk Configuration display.
   d) Select option 2 (Include unit in device parity protection) on the Work with Device Parity Protection display.

   or from the Use Dedicated Service Tools (DST) menu, follow these steps:

   a. Select option 4 (Work with disk units).
   b. Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c. Select option 5 (Work with device parity protection) on the Work with Disk Configuration display.
   d. Select option 4 (Include unit in device parity protection) on the Work with Device Parity Protection display.

   The Include Disk Units in Device Parity Protection display appears:

```
              Include Disk Units in Device Parity Protection
   Select the units to be included in Device Parity Protection.
   Type choices, press Enter.
   1=Include unit in device parity protection
                 Parity              Serial
Resource Option    Set  ASP  Unit  Number      Type  Model  Name
           1         1    8     7  00-0274937  6602   050   DD008
```

2. Select option 2 for RAID 5 parity protection. Or select option 6 for RAID 6 parity protection.
3. If the disk units that you selected are to be included in device parity protection, confirm this by pressing the Enter key.

   After the include operation has completed, the following display is shown.

```
           Include Disk Units in Device Parity Protection Status
   The operation to include units in the device parity protection will be done
in several phases.  The phases are listed here and the status will be indicated
when known.
Operation                                    Status
Prepare to include units . . . . . . . . . :  Completed
Include units  . . . . . . . . . . . . . . :  Completed
```

   Note: You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished including disk units in device parity protection.
4. Press the Enter key to return to the Work with Device Parity Protection menu.

Related tasks

Starting dedicated service tools

To start dedicated service tools (DST), follow these steps.

Excluding a disk unit from device parity protection

You can exclude a disk unit that is attached to an input/output adapter (IOA) from device parity protection.

You cannot exclude a disk unit that has been assigned to an auxiliary storage pool (ASP) that has mirrored protection. To exclude a disk unit from device parity protection, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 5 (Work with device parity protection) on the Work with Disk Configuration display.
3. Select option 5 (Exclude unit from device parity protection) on the Work with Device Parity Protection display and press the Enter key.

   The following display is shown.

```
              Exclude Disk Units from Device Parity Protection
   Select the units to be excluded from Device Parity Protection.
   Type choices, press Enter.
   1=Exclude unit from device parity protection
                 Parity            Serial
Resource Option    Set  ASP  Unit  Number      Type  Model  Name
            1        1    8     7  00-0274937  6602   070   DD008
```

   This display shows only disk units that are eligible to be excluded. A disk unit is eligible to be excluded if it does not contain parity information. If the disk units that you want to remove are not eligible to be excluded, you must stop device parity protection instead. Then physically remove the disk units and restart device parity protection.
4. Type a 1 in the Option column for the disk units that you want to exclude from device parity protection and press the Enter key.

   The following display is shown.

```
                      Confirm Disk Units to be Excluded
   Press Enter to confirm your choice to have the system exclude the selected
units from device parity protection
Press F12=Cancel to return to change your choice
                 Parity            Serial
Resource Option    Set  ASP  Unit  Number      Type  Model  Name
            1        1    8     7  00-0274937  6602   050   DD008
```

5. If the disk units you selected are to be excluded from device parity protection, confirm this by pressing the Enter key.

   After the exclude operation has completed, the following display is shown.

```
           Exclude Disk Units in Device Parity Protection Status
   The operation to exclude units from the device parity protection will be done
in several phases.  The phases are listed here and the status will be indicated
when known.
Operation                                    Status
Prepare to exclude units . . . . . . . . . :  Completed
Exclude units  . . . . . . . . . . . . . . :  Completed
```

   Note: You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished excluding disk units in device parity protection.
6. Press the Enter key to return to the Work with Device Parity Protection menu.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Displaying device parity status

Use these steps to display device parity status.

1. From the System Service Tools (SST) menu, follow these steps:
   a) Select option 3 (Work with disk units).
   b) Select option 1 (Display disk configuration) on the Work with Disk Units display.

   or from the Use Dedicated Service Tools (DST) menu, follow these steps:
   a. Select option 4 (Work with disk units).
   b. Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c. Select option 1 (Display disk configuration) on the Work with Disk Configuration display.
   d. Select option 1 (Display disk configuration status) on the Display Disk Configuration display.

   The Display Disk Configuration menu appears.
2. Select option 5 (Display device parity status) on the Display Disk Configuration display and press the Enter key.
   The Display Device Parity Status display appears:

```
                    Display Device Parity Status

Parity                                    Resource                Hot Spare
  Set    ASP   Unit   Type    Model   Name       Status     Protection
     1                571E    001      DC11       RAID 5          N
        *     *      6717    078     DD0020     Active
        *     *      6717    078     DD0027     Active
        *     *      6717    078     DD0034     Active
        *     *      6717    078     DD0051     Active
     2                571E    001     DC12        RAID 6           N
        *     *      4327    099     DD0060     Active
        *     *      4327    099     DD0057     Active
        *     *      4327    099     DD0134     Active
        *     *      4327    099     DD0228     Active
```

   The display is organized by device parity set. It includes controllers that can support device parity protection and all of the disk units that have the hardware capability for device parity protection. The possible values for the Status column are as follows:

   Active
   This unit is part of a disk unit subsystem that has device parity protection. This unit is fully operational.

   Failed
   This unit is part of a disk unit subsystem that has device parity protection. This unit has failed. If another unit in the disk unit subsystem fails, data could be lost.

   % Rebuilt
   This unit is part of a disk unit subsystem that has device parity protection. The data on this unit is being rebuilt from other units in the disk unit subsystem.

   Unprotected
   This unit is part of a disk unit subsystem that has device parity protection. This unit is operational. However, another unit in the disk unit subsystem has failed or is being rebuilt. If another unit in the disk unit subsystem fails, data might be lost.

   % Resyncing
   The parity for the parity set is being built from the data within the disk unit subsystem.

Chapter 24. Working with mirrored protection

This topic describes how to start and stop mirrored protection and the rules that apply when you set up a mirrored auxiliary storage pool (ASP).

Related information

Mirrored protection

Planning cross-site mirroring

Mirrored protection: Configuration rules

Here are the configuration rules for mirrored protection.

- Mirrored protection is configured by auxiliary storage pool (ASP) number. The system configures the two units of a mirrored pair within an ASP.
- Mirrored protection requires an even number of storage units for each size of disk unit in the ASP that is being mirrored. An uneven number of storage units for any size of disk unit prevents mirrored protection from starting. This requirement does not apply to disk units that have device parity protection.
- The system attempts to assign the two storage units of a mirrored pair so that a failed unit can be repaired while the system continues to use the other mirrored unit. For a hardware configuration where this is not possible, the repair of the failing unit must be delayed until the system can be powered down. This might be true for a failing mirrored unit which is sharing the same controller or input/output processor (IOP) as its mirrored unit.
- Standard disk mirroring support requires that the mirrored units are at specific input and output addresses on the system. (For Version 3 Release 6, you can override this restriction with a patch.) Both units in the pair must be attached to the multi-function IOP because the system must be able to perform an initial program load (IPL) from either unit. Therefore, the system attempts to assign the mirrored units for unit 1 of the system ASP first. If you are mirroring the system ASP, mirrored protection does not start if valid mirrored units for unit 1 cannot be found.

  Note: Remote mirroring support removes this requirement and provides IOP-level or bus-level protection.

Related tasks

Replacing a mirrored unit

A unit selected to replace the failed mirrored unit must satisfy all of the mirrored protection configuration rules and restrictions when it is paired with the remaining unit in the mirrored pair.

Starting mirrored protection

Use this procedure to start mirrored protection for a specific auxiliary storage pool (ASP) on your system.

Before you attempt to start mirrored protection, you should ensure that the ASP meets the requirements that are described in the “Mirrored protection: Configuration rules” on page 431 topic. If you are performing several disk configuration and disk protection tasks, see the Chapter 21, “Configuring disks and disk protection,” on page 369 topic to determine the correct sequence of steps for your situation.

To start mirrored protection, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual IPL to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 4 (Work with mirrored protection) on the Work with Disk Configuration display.
3. Select option 2 (Start mirrored protection) on the Work with Mirror Protection display.

© Copyright IBM Corp. 1997, 2025

```
                  Select ASP to Start Mirrored Protection

Select the ASPs to start mirrored protection on.

Type options, press Enter
  1=Select
Option   ASP   Protection
  _        1   Unprotected
  1        2   Unprotected
```

4. Select the ASP or ASPs to be mirrored on the Select ASP to Start Mirrored Protection display and press the Enter key.

   You might see the following display:

```
                           Confirm Continuation
   To proceed, the system must perform directory recovery, which might take a
significant amount of time.  The system might appear inactive during this time.

Press Enter to continue. Press F12=Cancel to return and change your choices.
```

   Press the Enter key to continue.

5. The system shows a confirmation display of the new mirrored protection configuration, including the levels of protection.

   Notice that half of the previous unit numbers for units being mirrored in the ASP no longer exist. The storage units for those unit numbers have been paired with the storage units for the remaining unit numbers to make mirrored pairs. Also notice that device parity units in the ASP have not been affected.

```
                     Confirm Start Mirrored Protection

Press Enter to confirm your choice to start mirrored
protection.  During this process the system will be IPLed.
You will return to the DST main menu after the IPL is
complete.  The system will have the displayed protection.

Press F12 to return to change your choice.

            Serial                  Resource
ASP  Unit   Number      Type Model  Name        Protection
  1                                             Unprotected
       1    00-48519    6606  030   DD010       Unprotected
  2                                             Mirrored
       2    00-1000341  9337  211   DD012       Disk Unit
       2    00-5000341  9337  211   DD015       Disk Unit
       3    00-0186325  6602  074   DD019       Device Parity
       4    00-0162516  6602  074   DD025       Device Parity
       5    00-0238703  6602  074   DD052       Device Parity
```

6. If the configuration is what you had planned and you do not have other configuration changes to make, skip to step 7.

   If the configuration is not what you had planned, for example, the level of protection is less, you have the following options:

   - Verify that the correct ASP was selected. Verify that any new storage units have been added to the correct ASP.
   - Determine if additional hardware is required to achieve the planned level of protection.
   - Determine if the existing hardware needs to be connected differently to achieve the planned level of protection. Contact your technical support organization for assistance.
   - Consider continuing the start mirrored protection process which will provide better availability than non-mirrored protection, rather than waiting until additional hardware arrives so that you can achieve your planned level of protection. After you receive and install the additional hardware, use the Table 66 on page 369 topic to determine the procedure for configuring your disk storage correctly. Even on very large systems, the tasks to stop mirroring, add units, and start mirrored protection can be done in several hours.

7. Place the system in Normal mode and press the Enter key to accept the configuration.

The system performs the first part of the process for starting mirrored protection. During that time, you are shown the Function Status display:

```
                          Function Status
You selected to start mirrored protection.              5 % Complete
```

The system updates the display periodically.

Note: You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished starting mirrored protection.

The system continues the start mirrored protection process in the What the system does when you start mirrored protection topic without further operator intervention.

8. After the system reaches the Command Entry display, you might want to create the QSYSMSG message queue to receive messages.

   If you have this message queue, the system sends certain critical messages to it. You can monitor the message queue with a program to make sure the messages are not missed.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

What the system does when you start mirrored protection

The system performs these steps when mirrored protection is started.

1. Data is moved off half of the storage units in the selected auxiliary storage pools (ASPs). This can take from a few minutes to a few hours, depending on the amount of data that must be moved.

   Objects that are created on a preferred unit can be moved to another unit. The preferred unit number might no longer exist when mirrored protection starts.
2. New control information is written to disk, describing the new mirrored system configuration.
3. After the data is moved and the control information is written, the system performs an initial program load (IPL). However, when you start mirroring only on independent ASPs, the system does not perform an IPL.
4. When the system reaches dedicated service tools (DST), the previously selected ASPs are mirrored, although the two storage units in the mirrored pairs are not yet synchronized.

   If the keylock switch is in the Manual position, you have the option to perform other configuration changes or perform an IPL. If you do not have configuration changes to make, select the option to perform an IPL and press the Enter key.

   If the keylock switch is in the Normal position, the system automatically continues the IPL.
5. When the system continues the IPL past DST, the mirrored pairs are synchronized during storage management recovery. This can take a few hours, although this long recovery time only occurs when mirrored protection is first started, and not during every IPL on a mirrored system. The progress of the synchronization process is displayed and periodically updated on the system or logical partition control panel. The system will display the code SRC C6xx 4205 where xx indicates the completed percentage of the synchronization process. On very large systems, the entire start mirrored protection process can take up to approximately 8 to 10 hours.
6. After storage management recovery is completed, the selected ASPs have mirrored protection.

Mirrored protection configuration errors

There can be no missing, active disk units anywhere in the configuration when mirroring is being started. Units with a status of missing must be powered, repaired, or replaced.

Starting mirrored protection can fail if there is insufficient storage available in the auxiliary storage pool (ASP) to contain the current data in the ASP. The percentage that is used in the ASP must normally be

less than half of the ASP threshold. The exception to this occurs when the ASP contains device parity protected disk units that can allow starting mirrored protection with a greater percent used.

There must be sufficient storage units in the ASP for the system to create mirrored pairs. If you receive a message that indicates that the system cannot pair unit 1 or other units, see the “Mirrored protection: Configuration rules” on page 431 topic.

Stopping mirrored protection

When you stop mirrored protection, one unit from each mirrored pair becomes nonconfigured.

Before you can stop mirrored protection for an auxiliary storage pool (ASP), at least one unit in each mirrored pair in that ASP must be present and not suspended. To control which mirrored unit of each pair becomes nonconfigured, you can suspend the storage units that you want to become nonconfigured. For units that are not suspended, the selection is automatic.

To stop mirrored protection, follow these steps:

1. If you are not already using dedicated service tools (DST), perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Service Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 4 (Work with mirrored protection) on the Work with Disk Configuration display.
3. Select option 3 (Stop mirrored protection) on the Work with Mirror Protection display.

   You are shown the Select ASP to Stop Mirrored Protection display:

```
                    Select ASP to Stop Mirrored Protection

Select the ASPs to stop mirrored protection on.

Type options, press Enter
  1=Select

Option   ASP   Protection
  1       2    Mirrored
```

4. Select the ASP or ASPs for which mirrored protection is to be stopped on the Select ASP to Stop Mirrored Protection display and press the Enter key.

   The Confirm Stop Mirrored Protection display is shown:

```
                      Confirm Stop Mirrored Protection

Press Enter to confirm your choice to stop mirrored
protection.  During this process the system will be IPLed.
You will return to the DST main menu after the IPL is
complete.  The system will have the displayed protection.

Press F12 to return to change your choice.

            Serial                   Resource
ASP  Unit   Number      Type Model   Name        Protection
  1                                              Unprotected
       1    00-48519    6606  030    DD010       Unprotected
  2                                              Unprotected
       2    00-1000341  9337  211    DD012       Unprotected
       3    00-0186325  6602  074    DD019       Device Parity
       4    00-0162516  6602  074    DD025       Device Parity
       5    00-0238703  6602  074    DD052       Device Parity
```

5. Press the Enter key to confirm your choice.

   The system stops mirrored protection for the ASPs you requested and performs an IPL. However, when you stop mirroring on only independent ASPs, the system does not perform an IPL.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Chapter 25. Working with disk compression

This topic describes how to start and stop disk compression. It also discusses various considerations for using disk compression.

Introduction to disk compression

Disk compression is a technology that increases the apparent capacity of disk storage devices by encoding the data to take up less physical storage space on disk.

Disk compression is performed in the disk subsystem controller and does not affect the IBM i processor. The compression of data is performed automatically on each write command. The decompression of data is performed automatically on each read command. With the exception of a performance impact, disk compression is transparent to applications. The performance of compressed disk drives is slower than the performance of non-compressed disk drives. This is due to the overhead of compression and decompression, and the variations in the length of the data that is written to disk.

Typically, data that is found on disk units has a wide range of access requirements. You can choose to move data that is accessed infrequently, or data that does not require high performance input/output (I/O) rates, to compressed disk units. Disk compression is intended to make infrequently accessed data available online at a lower cost. This storage alternative is positioned between non-compressed disk unit storage and optical or tape storage.

Compressed disks have the same disk subsystem availability options of device parity protection and mirrored protection as non-compressed disks. Disk compression is only supported in user auxiliary storage pools (ASPs).

Restrictions and considerations

These restrictions and considerations apply to disk compression.

- A compression-capable storage controller is required for the compressed disk units.
- Disk compression is only allowed on certain types of disk units. Contact your service provider for a list of the disk units that are capable of disk compression.
- If you are using V4R4 or later on your system, you can start or stop disk compression on configured and nonconfigured disk units.

  Note: If the disk unit is configured, you can only start disk compression if the disk unit is less than or equal to 92 if there is room in the auxiliary storage pool (ASP) for moving data off the disk unit so that the amount of data left on the disk unit is 92% of its uncompressed capacity.
- Disk units that are compressed cannot be added to the system ASP. You can only add compressed units to a user ASP.
- A User ASP might contain compressed and noncompressed disk units.
- A device parity protection set might contain compressed and non-compressed disk units.
- If you are using mirrored protection, both disk units in the mirrored pair must have the same compression status (either both compressed or both non-compressed).
- You might experience a zero to 30% performance degradation on I/O operations to compressed disk units. Your results might vary based on your actual work load characteristics.

Related tasks

Starting disk compression

© Copyright IBM Corp. 1997, 2025

You can start disk compression from the Dedicated Service Tools (DST) menu.

Disk compression and capacity

Capacity gains might vary and are determined by the amount of compression that can be performed on the data.

An overall 2:1 compression ratio should be expected for a disk unit. A maximum 4:1 compression ratio for a disk unit might be achievable in certain circumstances. Additionally, a compression ratio for a set of data on that disk unit might reach a maximum of 8:1. The IBM i software enforces the overall disk unit maximum compression ratio of 4:1. Because capacity is dependent on how well the data compresses, the capacity that is displayed will change as data is written to or deleted from the disk. The capacity that is displayed for a compressed disk unit is the amount of data on the disk plus an estimate of how much additional data can fit on the disk.

The following example shows the calculation and display of capacity by the system for compressed disk units. The disk unit capacities are available on the Display Disk Configuration Capacity display under the DST menus and the SST menus. The capacities are also available on the Work with Disk Status (WRKDSKSTS) display.

Note: If you have the licensed program Performance Tools for IBM Performance Tools for i (5770-PT1) installed on the system, you might use the system report to display the compression ratio. (You can find the ratio in the Disk Compression Statistics section in the system report.)

1. Before starting compression, a nonconfigured 6602 model 050 has a capacity of 1031 megabytes.

```
                          Display Non-Configured Units

Serial                      Resource
Number       Type    Model    Name      Capacity    Status
83-0135199   6602      050    DD005         1031    Non-configured
83-0306044   6602      050    DD006         1031    Non-configured
```

2. After starting compression, the 6602 model number changes to 060, and the capacity doubles.

```
                          Display Non-Configured Units

Serial                      Resource
Number       Type    Model    Name      Capacity    Status
83-0135199   6602      060    DD005         2062    Non-configured
83-0306044   6602      060    DD006         2062    Non-configured
```

3. Two compressed disk units are added into user ASP 2.

```
                    Display Disk Configuration Capacity

                                              ----Protected----   ---Unprotected---
ASP  Unit  Type  Model  Threshold  Overflow     Size   % Used        Size   %Used
 1                         90%    No              0    0.00%       21372  17.26%
       1   6607    050                           0    0.00%        4194  29.25%
       2   6713    050                           0    0.00%        8589  14.33%
       3   6713    050                           0    0.00%        8589  14.34%

 2                         90%    No              0    0.00%        4124   0.10%
       5   6602    060                           0    0.00%        2062   0.10%
       6   6602    060                           0    0.00%        2062   0.10%
```

4. After writing data to the user ASP, capacities and percentages that are used are displayed.

```
                    Display Disk Configuration Capacity

                                                 ----Protected----   ---Unprotected---
ASP  Unit  Type  Model  Threshold  Overflow     Size   % Used        Size   %Used
 1                        90%    No              0    0.00%       21372  17.26%
       1   6607    050                          0    0.00%        4194  29.25%
       2   6713    050                          0    0.00%        8589  14.33%
       3   6713    050                          0    0.00%        8589  14.34%

 2                        90%    No              0    0.00%        5812  61.06%
       5   6602    060                          0    0.00%        2905  61.04%
       6   6602    060                          0    0.00%        2907  61.08%
```

5. The following calculations can be performed to determine how well the data is being compressed, and the estimated disk unit capacity. These calculations can be performed on a user ASP basis, as well as on individual disk units.

Amount	Calculation
Logical data written	Size * (% Used / 100)
Physical free space	(Size * (1-(% Used / 100)) / 2
Physical space used	Non-compressed size - Physical free space
Compression ratio of data written	Logical data written / Physical space used
Estimated disk capacity	Logical data written + (2 * Physical free space)

Using the values listed for Unit 5 in the preceding step with these formulas produces the following table:

Amount	Calculation
Logical data written	2905 * .6104 = 1773
Physical free space	(2905 * (1-(.6104)) / 2 = 566
Physical space used	1031 - 566 = 465
Compression ratio	1773 / 465 = 3.8
Estimated disk capacity	1773 + (2 * 566) = 2905

Disk unit full considerations

While space is being reserved or data is being written to compressed disk units, it is possible for a unit to become temporarily full.

The storage subsystem controller will detect this situation and attempt to reposition the data on the unit in order to maximize the amount of data that can be stored on the unit. This repositioning of the data increases the effective size of the disk unit. If the storage subsystem controller cannot immediately service a system request, a failure will be returned to the system. The system bases its response to this failure on the type of input/output (I/O) operation that is requested. One of the following scenarios might occur:

- The system responds to this failure by overflowing the auxiliary storage pool (ASP). Once the overflow occurs, this I/O request will be performed in the system ASP and will succeed.
- The system responds to this failure by displaying an A6xx 0277 system reference code (SRC) on the control panel in the system unit. The system displays this attention SRC until storage space becomes available on the disk unit that is denoted in the attention SRC.

When the system displays an A6xx 0277 attention SRC on the control panel, it also logs a corresponding A6xx 0277 record in the Product Activity Log. This occurs each time this disk unit full condition is detected. The system also sends message CPI116C

```
Compressed disk unit &1 is full
```

to the QSYSOPR message queue. The system will reissue the failed I/O operation and will continue to display the attention SRC on the system or logical partition control panel until the condition is corrected. When the storage subsystem controller creates sufficient space on the compressed unit to contain the system request, the I/O operation completes successfully and the system resumes normal processing.

While this attention SRC is being displayed, some I/O operations to the affected compressed disk unit might be suspended. As a result, you might observe that jobs which issue I/O operations to the affected unit appear to hang.

To reduce the likelihood of system operations hanging while the storage subsystem recovers from a disk unit full condition, it is suggested that ASPs with compressed units operate with a storage threshold of less than or equal to 90%.

As space on the disk unit continues to be used, eventually the storage subsystem controller can no longer store more data on the unit. At this point, the storage subsystem controller will return a failure on any system requests requiring storage space.

Related concepts

How the system responds to disk unit full 448 reference code A6xx 0277
The system bases its response to the disk unit full condition on the type of input/output (I/O) operation that caused the condition.

How the system responds to disk unit full 448 reference code A6xx 0277

The system bases its response to the disk unit full condition on the type of input/output (I/O) operation that caused the condition.

If the system request is reserving additional storage space in the auxiliary storage pool (ASP), the compression recovery policy for the ASP determines the system response. You set this policy by using the Change ASP Attribute (CHGASPA) command.

The system might respond to the disk unit full condition in one of the following ways:

- If the compression recovery policy is *OVERFLOW, the system responds to this condition by overflowing the ASP. Once the overflow occurs, this I/O request will be performed in the system ASP and will succeed. This is the system default compression recovery policy for all user ASPs.
- If the compression recovery policy is *RETRY, the system responds to this condition by displaying an A6xx 0277 SRC on the control panel for the system or logical partition, and repeatedly retrying the failed I/O operation. The system displays this attention SRC until storage space becomes available on the disk unit and the I/O operation succeeds. If storage space does not become available on the disk unit, the ASP overflows. The system removes the attention SRC from the control panel, and the I/O operation is successfully performed in the system ASP.
- If the compression recovery policy is *WAIT, the system responds to this condition by displaying an A6xx 0277 SRC on the control panel and repeatedly retrying the failed I/O operation. If storage space does not become available on the disk unit, the ASP does not overflow. The user must take one of the corrective actions that are discussed in “SRC code A6xx 0277” on page 441.

Related concepts

Disk unit full considerations
While space is being reserved or data is being written to compressed disk units, it is possible for a unit to become temporarily full.

Recovering an overflowed user auxiliary storage pool

Reset a user auxiliary storage pool (ASP) in overflowed status as soon as possible. An overflowed ASP affects system performance. It also makes recovery more difficult and might increase the amount of data that is lost if a failure occurs.

SRC code A6xx 0277

When you see the A6xx 0277 SRC code in the control panel on the system or logical partition, select the appropriate compression reference code word for additional information.

The compression reference code word is either word 15 or word 17.

Table 77. Word formats for SRC codes

Word for the SRC code	Word format	Description
15	CCEE 0000 CC indicates the operation in progress with the following values: • 84 is an allocate operation • 2x is a write operation where x is 1, 2, or 4 EE indicates the error code with the following values: • 00 is the storage subsystem controller cannot obtain any additional storage space on the unit • 02 is the storage subsystem controller will be able to obtain additional storage space on the unit	This word is the compression reference code that indicates two things. First, it indicates which operation was in progress. Second, it indicates whether the storage subsystem controller can obtain additional storage space on the system. This reference code word was formerly word 17 in the V4R4 release and in previous releases.
16	uuuu uuuu	This word describes the unit address of the disk unit.
17	BBBB ccbb	This word defines the bus, card, and board address of the disk unit.

Perform one of the four actions in response to SRC A6xx 0277.

User action 1

You need to wait for the storage subsystem controller to reposition the data on the disk unit.

If the error code for EE of the attention system reference code (SRC) is 02, the storage subsystem controller will eventually obtain additional storage space on the unit, such that the input/output (I/O) operation will succeed. If the system does not resume normal processing within 20 minutes, contact your next level of support.

User action 2

You need to make storage space available in the auxiliary storage pool (ASP) that contains the disk unit that is denoted in the system reference code (SRC).

Word 16 contains the unit address of the disk unit. Word 17 or 7 is BBBB ccbb. Correlate the unit address (logical address) with a resource name or serial number. The ASP that contains the disk unit can be determined by using the Display Disk Configuration Status display under the DST and SST menus.

If the error code for EE of the attention SRC is 00, the storage subsystem controller determined that the disk unit is full.

Perform one or more of the following actions:

- Delete objects that are not needed from the ASP.
- Save objects that are not needed from the ASP by specifying STG(*FREE) on the Save Object (SAVOBJ) command.
- Move one or more libraries to a different ASP.

  Note: You cannot use the Move Object (MOVOBJ) command to do this. You must save the library, delete it, and then restore it to a different ASP.
- Move one or more folders to a different ASP by saving the folder, deleting it, and restoring it to a different ASP.
- Increase the storage capacity by adding disk units to the ASP.

User action 3

You need to change the compression recovery policy to the desired system behavior.

For more information about the Change Auxiliary Storage Pool Attributes (CHGASPA) command, refer to the IBM i online help.

User action 4

A system reference code (SRC) is displayed that indicates an auxiliary storage pool (ASP) requires more storage. You need to do an initial program load (IPL) that more storage space can be made available in ASP.

Word 16 contains the unit address of the disk unit. Word 17 or 7 is BBBB ccbb. Correlate the unit address (logical address) with a resource name or serial number. The ASP that contains the disk unit can be determined by using the Display Disk Configuration Status display under the DST and SST menus.

If the error code for EE of the attention SRC is 00 and the system is holding critical resources, the system hangs. The suggested recovery procedure is to reload the initial program on the system. The system must be in Manual mode. Take the following steps:

1. Force the system to write changed data in main storage to disk storage by pressing the power button twice to stop the system.

   Wait for system activity to stop.

   Changed data that is in main storage cannot be written to the disk unit. Therefore, the power off that was described earlier, hangs.
2. Start an IPL.
   a) Place the system in manual mode.
   b) Press the Function Select switch (or buttons) to display 03 in the Function display.
   c) Press the Enter button on the control panel.
3. On the following Manual mode IPL, do one of the following tasks -
   - Increase the storage capacity by adding disk units to the ASP at dedicated service tools (DST).
   - Start the system in a restricted state. Make storage space available in the ASP that contains the disk unit that was denoted in the attention SRC. Make the space available by do one or more of the following steps:
     – Delete objects that are not needed from the ASP.
     – Delete the data for objects that are not needed from the ASP. The data can be deleted by saving the objects with the Save Object (SAVOBJ) command and specifying STG(*FREE).
     – Move one or more libraries to a different ASP.

       Note: The object cannot be moved to a different ASP by using the Move Object (MOVOBJ) command. Instead, ave the library, delete it, and restore it to a different ASP.
     – Move one or more folders to a different ASP by saving the folder, deleting it, and restoring it to a different ASP.

Examples of A6xx 0277

You can learn the examples that illustrate two situations where an A6xx 0277 SRC record will be generated, and any associated actions you might need to take.

The two situations are as follows:

17 8402 0110 (V4R4 and previous) OR 15 or 5 8402 0000 (V4R5)
In this example, an allocate operation (84) is being attempted and the error code (02) denotes the operation is being retried and will eventually succeed. You do not need to take any additional action. If the system does not resume normal processing within 20 minutes, contact your next level of support.

If you want the system to overflow the user ASP into the system ASP, take "User action 3" on page 442 that specifies the compression recovery policy *OVERFLOW.

17 2000 0110 (V4R4 and previous) OR 15 or 5 2000 0000 (V4R5)
In this example, a write operation (20) is being attempted and the error code (00) denotes that the operation is being retried indefinitely, because the storage subsystem controller has determined that there is no storage space available on the disk unit. Take either "User action 2" on page 441 or "User action 4" on page 442.

Starting disk compression

You can start disk compression from the Dedicated Service Tools (DST) menu.

To start disk compression from DST, perform the following steps:

1. If your use your 2748 storage I/O controller for compression, make sure that you set the jumper for compression mode before you continue with the following steps.
2. If you are not already using DST, perform a manual initial program load (IPL) to start DST.
3. From the Use Dedicated Services Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 6 (Work with disk compression) on the Work with Disk Configuration display.
4. Select option 2 on the Work with Disk Compression.

```
                     Work with Disk Compression

Select one of the following:

     1. Display disk compression status
     2. Start compression on disk units
     3. Stop compression on disk units



Selection

F3=Exit     F12=Cancel
```

5. Select the disk units that you want to start compression on from the Select Disk Units for Start Compression display.

   Note: You can only start compression on a configured unit if the disk unit is less than or equal to 92% full.

```
             Select Disk Units for Start Compression

Select the units for start compression.

Type choice, press Enter.
  1=Start compression
                Serial                    Resource
OPT Unit   ASP  Number     Type   Model  Name         Status
 1              68-7F0DB   6607    050   DD005        Non-configured
                68-5FB0B   6713    050   DD001        Non-configured



F3=Exit     F12=Cancel
```

6. You are shown the Confirm Disk Units for Start Compression display.

   This display shows the approximate amount of time that is required to start disk compression, and the current and proposed sizes of the disk unit.

```
             Confirm Disk Units for Start Compression

Estimated time for this operation to complete : 14-16 Minutes

Press Enter to confirm your choice to compress the disk units.
Press F9=Capacity Information to display the resulting capacity.
Press F12=Cancel to return to change your choice.

                   Serial                    Current     Proposed
OPT  ASP   Unit    Number      Type   Model  Size        Size
 1                 68-7F0DB    6607    050   4194        8388



F9=Resulting Capacity           F12=Cancel
```

7. At the Confirm Disk Units for Start Compression display, press Enter to confirm your choice of disk units that you want to start compression on.

   You are shown the Start Compression on Disk Unit Status display.

```
              Start Compression on Disk Unit Status

Estimated time for this operation to complete : 14-16 Minutes

The operation to start compression on the selected disk units
will be done in multiple phases.  The phases are listed here
and the status will be updated as the phase proceeds.

Phase                                      Status

Prepare to start compression . . . . . :   0 % Complete
Start compression  . . . . . . . . . . :
Prepare to compress data . . . . . . . :
Compress data  . . . . . . . . . . . . :



Wait for next display or press F16 for DST menu
```

8. After the start compression operation completes, you are returned to the Work with Disk Compression display, and shown a completion message.

```
                    Work with Disk Compression

Select one of the following:

     1. Display disk compression status
     2. Start compression on disk units
     3. Stop compression on disk units



Selection

F3=Exit     F12=Cancel
 The requested compression operation completed successfully.
```

Related concepts

Restrictions and considerations
These restrictions and considerations apply to disk compression.

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Stopping disk compression

Use this procedure to stop disk compression by using the dedicated service tools (DST).

1. If you are not already using DST, perform a manual initial program load (IPL) to start DST.
2. From the Use Dedicated Services Tools (DST) menu, follow these steps:
   a) Select option 4 (Work with disk units).
   b) Select option 1 (Work with disk configuration) on the Work with Disk Units display.
   c) Select option 6 (Work with disk compression) on the Work with Disk Configuration display.
3. Select option 3 on the Work with Disk Compression display.

```
                    Work with Disk Compression

Select one of the following:

     1. Display disk compression status
     2. Start compression on disk units
     3. Stop compression on disk units



Selection

F3=Exit     F12=Cancel
```

4. Select the disk units that you want to stop compression on from the Select Disk Units for Start Compression display.

   Note: Compression can only be stopped if there is room in the ASP for moving data off of the disk unit. After the data has been moved off, the amount of data that is left on the disk is less than or equal to 92% of its uncompressed capacity.

```
            Select Disk Units for Stop Compression

Select the units for stop compression.

Type choice, press Enter.
  1=Stop compression
                Serial                    Resource
OPT Unit   ASP  Number     Type   Model  Name         Status
 1             68-7F0DB   6607    060   DD005        Non-configured



F3=Exit     F12=Cancel
```

5. You are shown the Confirm Disk Units for Stop Compression display.

   This display shows the approximate amount of time that is required to stop disk compression, and the current and proposed sizes of the disk unit.

```
            Confirm Disk Units for Stop Compression

Estimated time for this operation to complete : 1-2 Minutes

Press Enter to confirm stop compression.
Press F9=Capacity Information to display the resulting capacity.
Press F12=Cancel to return to change your choice.

                   Serial                    Current     Proposed
OPT  ASP   Unit    Number      Type   Model  Size        Size
 1                 68-7F0DB    6607    060   8388        4194



F9=Resulting Capacity           F12=Cancel
```

6. At the Confirm Disk Units for Stop Compression display, press Enter to confirm your choice of disk units that you want to stop compression on.

   You are shown the Stop Compression on Disk Unit Status display.

```
             Stop Compression on Disk Unit Status

Estimated time for this operation to complete : 1-2 Minutes

The operation to stop compression on the selected disk units
will be done in multiple phases.  The phases are listed here
and the status will be updated as the phase proceeds.

Phase                                      Status

Prepare to stop compression  . . . . . :   0 % Complete
Stop compression . . . . . . . . . . . :
```

7. After the stop compression operation completes, you are returned to the Work with Disk Compression display, and shown a completion message.

```
                    Work with Disk Compression

Select one of the following:

     1. Display disk compression status
     2. Start compression on disk units
     3. Stop compression on disk units



Selection

F3=Exit     F12=Cancel
 The requested compression operation completed successfully.
```

Related tasks

Starting dedicated service tools
To start dedicated service tools (DST), follow these steps.

Procedural sequences for configuring disks and protection

Each respective configuration change requires you to perform procedures in a specific order. You can find a list, showing the order that you should use when performing the procedures.

Checklist: Adding a new I/O compression-capable storage controller

This checklist shows the sequence of tasks that you use to add a new compression-capable input/output (I/O) storage controller and disk units to your system.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you to diagnose any problems that occur.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 78. Adding a new I/O storage controller and disk units

Task	What to do	Where to read more about it
___ Task 1	Install the new storage controller in the system. This is normally done by a service representative.	
___ Task 2	Physically attach disk units to the new storage controller. This is normally done by a service representative.	
___ Task 3	Start dedicated service tools (DST).	“Starting dedicated service tools” on page 382.
___ Task 4	Print your current disk configuration.	“Displaying your disk configuration” on page 385.
___ Task 5	If you want to have device parity protection for the storage controller, start device parity protection now.	“Starting device parity protection for an input/output adapter” on page 419.
___ Task 6	Start disk compression on nonconfigured disk units.	“Starting disk compression” on page 443.

Table 78. Adding a new I/O storage controller and disk units (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 7	Add nonconfigured disk units to the correct auxiliary storage pools (ASPs). You can add the disk units to an existing user ASP or to a new ASP. If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have equal capacities.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 8	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it at this time.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 9	If you created any new ASPs and you want those ASPs to have mirrored protection, start mirrored protection now.	"Starting mirrored protection" on page 431.
___ Task 10	Verify that your disk configuration is correct.	"Displaying your disk configuration" on page 385.
___ Task 11	End DST.	"Stopping dedicated service tools" on page 383.

Checklist: Adding disk units to an existing compression-capable storage controller

This checklist shows the sequence of tasks that you use to add disk units to an existing compression-capable storage controller.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you to diagnose any problems that occur.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 79. Adding compressed disk units to an existing storage controller		
Task	**What to do**	**Where to read more about it**
___ Task 1	Physically attach disk units to an existing storage controller. This is normally done by a service representative.	
___ Task 2	Start dedicated service tools (DST) or system service tools (SST).	"Starting dedicated service tools" on page 382 or "Starting system service tools" on page 384.
___ Task 3	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 4	Include the disk units that you want to protect in device parity protection.	"Including a disk unit in device parity protection" on page 427.
___ Task 5	Start disk compression on nonconfigured disk units.	"Starting disk compression" on page 443.

Table 79. Adding compressed disk units to an existing storage controller (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 6	Add nonconfigured disk units to the correct auxiliary storage pools (ASPs). You can add the disk units to an existing user ASP or to a new ASP. If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have equal capacities.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 7	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it at this time.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 8	If you created any new ASPs and you want those ASPs to have mirrored protection, start mirrored protection now.	"Starting mirrored protection" on page 431.
___ Task 9	Verify that your disk configuration is correct.	"Displaying your disk configuration" on page 385.
___ Task 10	End DST or SST.	"Stopping dedicated service tools" on page 383 or "Stopping system service tools" on page 384.

Checklist: Moving disk units from the system ASP to a user ASP

This checklist shows the sequence of tasks that you use to move one or more disk units from the system auxiliary storage pool (ASP) to a user ASP. This checklist assumes that the disk units are attached to a compression-capable subsystem storage controller.

Before you begin, make a copy of this checklist. Complete the appropriate areas as you or the service representative perform the configuration tasks. This checklist provides an important record of your actions. It might help you diagnose any problems that occur.

Attention: When you perform the tasks in this checklist, the system moves large amounts of data. Make sure that you have saved your system completely in the event that you need to recover from an error situation.

Most tasks in the checklist include references to other topics in this document. Refer to these topics if you need more information about how to perform a particular task.

Table 80. Moving disk units from the system ASP to a user ASP		
Task	**What to do**	**Where to read more about it**
___ Task 1	Print your current disk configuration.	"Displaying your disk configuration" on page 385.
___ Task 2	Calculate the space requirements for both the source and target ASPs for the disk units.	"Calculating space requirements for an auxiliary storage pool" on page 406.
___ Task 3	Use option 21 from the Save menu to save your entire system.	
___ Task 4	Start dedicated service tools (DST).	"Starting dedicated service tools" on page 382.

Table 80. Moving disk units from the system ASP to a user ASP (continued)		
Task	**What to do**	**Where to read more about it**
___ Task 5	Remove disk units that you plan to add to a different ASP.	"Removing a disk unit from an auxiliary storage pool" on page 403.
___ Task 6	If you want to use device parity protection, start device parity protection (if necessary), and include the disk units that you want to protect. If you do not want to use device parity protection, proceed to the next step.	"Starting device parity protection for an input/output adapter" on page 419 (if necessary), and "Including a disk unit in device parity protection" on page 427.
___ Task 7	Start disk compression on nonconfigured disk units.	"Starting disk compression" on page 443.
___ Task 8	Add nonconfigured disk units to the correct ASPs. You can add the disk units to an existing user ASP or to a new ASP. If you are adding disk units to an ASP that has mirrored protection and the new disk units do not have device parity protection, you must add pairs of disk units that have equal capacities.	"Adding disk units to an existing auxiliary storage pool" on page 392.
___ Task 9	If you created a new ASP on your system when you added disk units, the system set the storage threshold of the ASP to 90%. If you want a different threshold, change it at this time.	"Changing the storage threshold for an auxiliary storage pool" on page 399.
___ Task 10	If you created any new ASPs and you want those ASPs to have mirrored protection, start mirrored protection now.	"Starting mirrored protection" on page 431.
___ Task 11	Verify that your disk configuration is correct.	"Displaying your disk configuration" on page 385.
___ Task 12	End DST.	"Stopping dedicated service tools" on page 383.
___ Task 13	If necessary, move objects between ASPs.	"Transferring objects between auxiliary storage pools" on page 408.

Recovering from error codes

You might encounter system reference code (SRC) when working with disk compression.

This section discusses some causes of SRCs, and how to respond to them. You might want to see the Chapter 22, "Working with auxiliary storage pools," on page 391 topic for information on moving disk units into and out of auxiliary storage pools (ASPs).

Recovering from SRC 6xxx 7051

You are receiving this message because the compressed device and the compression input/output adapter (IOA) are not compatible.

1. Did you remove the disk unit from another system?

   Yes

   No

↓

Go to step 6.

2. Was it removed from the auxiliary storage pool (ASP) of the other system?

Yes

No

↓

Go to step 4.

3. Stop compression on the disk unit.

This ends the procedure.

4. Do you want to save the data that is on the disk unit?

Yes

No

↓

Stop compression on the disk unit.

This ends the procedure.

5. Return the disk unit to its original system and IOA and perform the following tasks:
   a) Remove the disk unit from the ASP.
   b) Stop compression on the disk unit.

Reinstall the disk unit to this system.

This ends the procedure.

6. The input/output processor (IOP) or IOA that you are using is not compatible with the disk unit.

Do you want to save the data that is on the disk unit?

Yes

No

↓

Stop compression on the disk unit.

This ends the procedure.

7. If you come here from another procedure, return there and follow the procedure, otherwise contact your next level of support.

This ends the procedure.

Recovering from SRC 6xxx 7052

You are receiving this message because the counter that tracks the number of write operations to this drive has reached 75% of its maximum value.

This counter is used to determine if data on the drive is current with data within input/output adapter (IOA) memory. Because this counter is large, it is not necessary to stop compression for this drive immediately, but you should do so during your next scheduled maintenance. If this counter wraps, data can be lost.

To stop compression and restart compression on the disk drive to reset the write count, perform the following tasks:

1. Perform a manual mode initial program load (IPL) to dedicated service tools (DST).
2. To find the resource name of the disk drive with the problem, perform the following tasks:
   a) Select the Use dedicated service tools option.
   b) Select the Start a service tool option.
   c) Select the Hardware service manager option.
   d) Select the Work with service action log option.

e) Select the time frame of the problem.

f) Record the resource name that is associated with the 6xxx 7052 entry in the system reference code (SRC) column.

3. Remove the disk unit from the auxiliary storage pool (ASP).
4. Stop compression on the disk unit.
5. Start compression on the disk unit.
6. Add the disk drive back to the ASP from which it was removed.

This ends the procedure.

Chapter 26. Managing auxiliary storage pools

You can use system service tools (SST) and dedicated service tools (DST) or IBM Navigator for i to manage auxiliary storage pools (ASPs).

By separating libraries, documents, or other objects in an ASP, you protect them from data loss if a disk unit in a different ASP fails. The importance of protecting the data in the ASP and the capacity of the ASP are factors in determining what protection mechanisms (mirroring or device parity, for example) you choose to implement. You can make independent ASPs available or unavailable as you need the information in them. You can take low-usage data offline until it needs to be accessed, thus reducing system start time. In a clustered environment, you can configure independent ASPs to switch between systems.

Using ASPs also increases performance. You can place libraries or objects in an ASP and the ASP is used exclusively for those objects. If you do extensive journaling, a dedicated disk unit for the journal receiver can also improve journaling performance.

Note: Placing many active journal receivers in the same user ASP is not productive. The resulting contention between writing to more than one receiver in the ASP can slow system performance. For maximum performance, place each active journal receiver in a separate user ASP.

If the amount of data in a storage pool increases, you might need to increase the size of the storage pool. Conversely, if data in a storage pool decreases, you might want to decrease the size of that storage pool and use the disk space elsewhere.

Changing the size of an ASP by adding a disk unit, removing a disk unit, moving a disk unit, or deleting an ASP from the system. You must typically have QSECOFR authority to access these tasks.

Related concepts

Working with ASP trace and ASP balance
The three types of auxiliary storage Pool (ASP) balancing actions that you can choose.

Related tasks

Adding disk units to an existing auxiliary storage pool
Use this information to add disk units to an existing auxiliary storage pool (ASP).

Deleting an auxiliary storage pool
When you delete a user auxiliary storage pool (ASP), the status of all the disk units that are assigned to the ASP becomes nonconfigured. Any data that is in the ASP is deleted.

Determining adequate disk storage
If you need to know how much disk storage is available on your system, use the Work with System Status (WRKSYSSTS) command.

Related information

Disk management

Working with ASP trace and ASP balance

The three types of auxiliary storage Pool (ASP) balancing actions that you can choose.

- Capacity balance
- Hierarchical storage management (HSM) balance
- Usage balance

The balancing actions use the results of previous ASP traces to determine disk unit usage. Therefore, an ASP balance is more effective if you run an ASP trace first.

The following are restrictions and considerations for balancing actions:

© Copyright IBM Corp. 1997, 2025

- The usage balance uses the results of previous ASP traces to determine the disk unit usage. Trace an ASP before you run a usage balance.
- The HSM balance uses the results of previous ASP traces to determine the disk unit usage. Trace an ASP before you can run an HSM balance.
- The HSM balance requires a mixture of compressed and non-compressed disk units in the ASP.
- The system experiences degraded performance during the Trace activity and Balance activity.
- You can use the Trace and Balance functions in only ASPs that contain more than one disk unit.
- You can add an unconfigured disk unit to the configuration during the time a Trace is running. In this case, the system automatically includes the newly configured disk unit in the trace activity.
- You can run a single trace activity or balance activity only at a specific time to an ASP.
- The amount of improvement in system throughput that is achieved by running the balance depends on many of the following factors.
  - The amount of main storage.
  - The number of processors.
  - The level of activity on the system.
  - The amount of cache in the storage subsystem.
  - The quantity of disk units under each I/O processor in each storage subsystem.

Related concepts

Managing auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) or IBM Navigator for i to manage auxiliary storage pools (ASPs).

Auxiliary storage pool trace
The Trace Auxiliary Storage Pool Balance (TRCASPBAL) command monitors the frequency that data is accessed on the disk units in the auxiliary storage pool (ASP).

Capacity balance

The capacity balance function rearranges the data on all of the disk units within an auxiliary storage pool (ASP).

The capacity balance moves data so that each disk unit has an equal percentage of used and unused space. This is useful when you add new units to an ASP. We want to avoid the situation where several disk units contain the majority of the data and the newly added disk units contain very little data. This situation leads to poor system performance. The balance function spreads the data in the ASP evenly across all of the disk units.

Below is a display which shows the effects of using capacity balance. Before using capacity balance the recently added Unit 4 contained very little data. The system storage management allocates newly created data to the disk unit that has the lowest percentage of capacity used. Thus, the system routes all new storage allocations to the unit 4. If the system uses that newly created data frequently, a potential bottleneck is created. The system directs all of the input/output (I/O) operations to a single disk unit instead of spreading it across all units in the ASP. The capacity balance performed on the ASP allows the data to be evenly distributed across all of the disk units in the ASP. This means that the distribution of future space allocation on the disk units in the ASP is even across all of the disk units in the ASP. That ensures that the I/O to those allocates are also evenly spread out amongst the disk units instead of being concentrated on the newly added disk unit.

Unit	Before capacity balancing		After capacity balancing	
	Disk size	% used	Disk size	% used
1	1967	54.59%	1967	55.69%
2	1031	68.45%	1031	55.80%

Unit	Before capacity balancing		After capacity balancing	
	Disk size	% used	Disk size	% used
3	1031	68.41%	1031	55.76%
4	1031	0.30%	1031	55.77%

To start capacity balance by a control language (CL) command use the Start Auxiliary Storage Pool Balance (STRASPBAL) command. For example, if you want to start a capacity balance on ASP 4 and have it run for 25 minutes, enter the following command: `STRASPBAL ASP(4) TYPE(*CAPACITY) TIMLMT(25)`.

If you want to end a capacity balance before the time limit requested is reached, use the End Auxiliary Storage Pool Balance (ENDASPBAL) command. For example, if you want to end a running capacity balance on ASP 4, enter the following command: `ENDASPBAL ASP(4)`.

Hierarchical storage management (HSM) balance

The high-use and low-use data on each disk unit in the auxiliary storage pool (ASP) is redistributed. This action is done so that the high-use data resides on "fast" disk units and the low-use data resides on compression disk units, which are typically slower than standard disk units.

The ASP selected for a HSM balance must have a combination of compression disk units and non-compression disk units. You can only run a HSM balance following a trace ASP balance. The trace ASP balance function monitors the input/output (I/O) activity on each of the disk units in the ASP to determine where the high-use and low-use data resides.

Compression disk units have larger capacity, but are somewhat slower than non-compression disk units. This is due to the overhead of compression and decompression, and the variations in the length of the data that is written to disk. Typically, data that is found on disk units has a wide range of access requirements. The HSM balance function moves data that is accessed infrequently to compression disk units. Disk compression makes infrequently accessed data available online at a lower cost. System throughput improves when you move high-use data off of compression disk units. Moving the low-use data to the large compression disk units makes additional capacity available on the standard disk units so high-use data can be allocated.

The Start Auxiliary Storage Pool Balance (STRASPBAL) command is used to perform the HSM balance function. For example, if you want to run a HSM balance on ASP 4 for 25 minutes, enter the following command: `STRASPBAL ASP(4) TYPE(*HSM) TIMLMT(25)`.

If you want to end a HSM balance before the time limit requested is reached, use the End Auxiliary Storage Pool Balance (ENDASPBAL) command. For example, if you want to end a running HSM balance on ASP 4, enter the following command: `ENDASPBAL ASP(4)`.

Usage balance

The usage balance attempts to balance the usage of the disk units in an auxiliary storage pool (ASP).

The usage balance can only be done following a trace ASP balance. The trace ASP balance function monitors the input/output (I/O) activity on each of the disk units in the ASP. It then determines where the frequently used and infrequently used data resides. The usage balance function utilizes that trace information. It adjusts the data on the disk units so future system activity will be more equally balanced amongst the disk units in the ASP.

If the system determines that all of the disk units are of approximately equal utilization, the balance will end very quickly. The usage balance function uses the trace information in its calculations. If the trace data is old, or if your applications have changed to reference different data since the trace was run, the usage balance might result in very little improvement in your system throughput. It might even cause the throughput to decrease somewhat.

The Start Auxiliary Storage Pool Balance (STRASPBAL) command is used to perform the archive balance function. For example, if you want to start a usage balance on ASP 4 to run for 25 minutes, enter the following command: `STRASPBAL ASP(4) TYPE(*USAGE) TIMLMT(25)`.

If you want to end a usage balance before the time limit requested is reached, use the End Auxiliary Storage Pool Balance (ENDASPBAL) command. For example, if you want to end a running usage balance on ASP 4, enter the following command: `ENDASPBAL ASP(4)`.

Auxiliary storage pool trace

The Trace Auxiliary Storage Pool Balance (TRCASPBAL) command monitors the frequency that data is accessed on the disk units in the auxiliary storage pool (ASP).

Each input/output (I/O) to the disk units is monitored, and the results are recorded for use by the balance commands. The statistics which are collected are cumulative. For example, suppose you start one trace and it runs for 35 minutes. Then you start another trace on that ASP and it runs for 15 minutes. The second group of statistics is added to the first collection and the cumulative result is used to balance the ASP.

Select an ASP that you want the system to monitor. The system will record all of the I/O activity on the disk units in that ASP. For example to start a trace on ASP 4 which will run for 35 minutes, enter the following command: `TRCASPBAL ASP(4) SET(*ON) TIMLMT(35)`.

If you want to end a trace before the time limit requested on the start trace is reached, use the Trace Auxiliary Storage Pool Balance (TRCASPBAL) command. For example, if you want to end the trace on ASP 4, enter the following command: `TRCASPBAL ASP(4) SET(*OFF)`.

The collected statistics on each disk units I/O activity can be cleared by use of the TRCASPBAL command. You can clear old trace data if you do not want that data to be used in determining the locations of the high-use and low-use data on the disk units in the ASP. Use the TRCASPBAL command to clear the trace data. For example, if you want to clear the trace data which has been collected from ASP 4, enter the following command: `TRCASPBAL ASP(4) SET(*CLEAR)`.

Related concepts

Working with ASP trace and ASP balance
The three types of auxiliary storage Pool (ASP) balancing actions that you can choose.

Determining adequate disk storage

If you need to know how much disk storage is available on your system, use the Work with System Status (WRKSYSSTS) command.

1. Type `WRKSYSSTS` at the command line.

   The Work with System Status display appears.
2. Record the values that are shown for system auxiliary storage pool (ASP) and % system ASP used.
3. Use those values in the formula that is shown below to calculate the available free space on your system:

```
System ASP x (100 - % system ASP used) x 0.01 = Available disk storage
```

If the amount of storage is less than you need to complete your task, you must create more disk space. You can create more space by adding additional disk units or by cleaning your system of files and programs you no longer use.

Related concepts

Managing auxiliary storage pools
You can use system service tools (SST) and dedicated service tools (DST) or IBM Navigator for i to manage auxiliary storage pools (ASPs).

Part 7. Backup and recovery tools and techniques

© Copyright IBM Corp. 1997, 2025

Chapter 27. Examples: Techniques and programming for backup and recovery

You can use several techniques to assist with and manage your backup and recovery processes.

Examples: Backup and recovery

Here are several program examples for working with backup and recovery.

Using the Retrieve Journal Entry command in a program

You can use the Retrieve Journal Entry (RTVJRNE) command in a control language (CL) program to retrieve a journal entry and place it in variables in the program.

You can retrieve the following items:

- Sequence number
- Journal code
- Entry type
- Journal receiver name
- Library name for the journal receiver
- Journal entry-specific data

For example, you can use this command to automate your recovery procedures or to change the journal receivers and then save them.

In Figure 34 on page 459, the RTVJRNE command determines when job 000666/QPGMR/WORKST01 last opened file ORDENTP:

```
PGM
DCL &SEQ# TYPE(*DEC) LEN(10 0)
DCL &JRNE TYPE(*CHAR) LEN(200)
DCL &DATE TYPE(*CHAR) LEN(6)
DCL &TIME TYPE(*CHAR) LEN(6)
RTVJRNE JRN(DSTJRN/JRNLA) FILE(DSTPRODLIB/ORDENTP) +
        RCVRNG(DSTJRN/RCV30 DSTJRN/RCV27) +
        FROMENT(*LAST) TOENT(*FIRST) SEARCH(*DESCEND) +
        JRNCDE(F) ENTTYP(OP) JOB(000666/QPGMR/WORKST01) +
        RTNSEQNBR(&SEQ#) RTNJRNE(&JRNE)
CHGVAR &DATE (%SST(&JRNE 19 6))
CHGVAR &TIME (%SST(&JRNE 25 6))
ENDPGM
```

Figure 34. Program for retrieving journal entries

Control language program to handle escape conditions

You normally use the Apply Journaled Changes (APYJRNCHG) command to perform object recovery. However, usable journal receivers are needed to successfully use this command.

If usable journal receivers are not found, an escape message is sent. The following example demonstrates how this escape condition can be handled in a control language (CL) program by prompting for restoration of the required receiver. This example uses database files. You can extend this example to use the APYJRNCHG command with all object types that support journaling.

```
FILERECOV: PGM
           .
           .
           APYJRNCHG JRN(JRNLIB/JRNA) FILE((LIBA/FILEA)) +
                 RCVRNG(RCVLIB/RCV1 *CURRENT)
           MONMSG MSGID(CPF7053 CPF9801) +
```

© Copyright IBM Corp. 1997, 2025

```
              EXEC(CALL PGM(FIXLIB/RSTRCV) PARM(FILERECOV))
           .
           .
           ENDPGM
           .
           .
RSTRCV: PGM PARM(&PGMNM)
        /* Recover a nonexistent or unusable receiver        */
        /* in RCVRNG by prompting for a restore of           */
        /* receiver.
DCL *PGMNM TYPE(*CHAR) LEN(10)         /* name of program    */
                                       /* calling RSTRCV     */
                                       /* that received      */
                                       /* CPF7053 or         */
                                       /* CPF9801            */
DCL &MSGDATA TYPE(*CHAR) LEN(22) /* variable for   */
                                       /* CPF7053 or         */
                                       /* CPF9801            */
DCL &MSGDID TYPE(*CHAR) LEN(7) /* escape message  */
                                       /* ID                 */
DCL &RCVNAME TYPE(*CHAR) LEN(10) /* name of       */
                                       /* receiver to        */
                                       /* restore            */
DCL &RCVLIB TYPE(*CHAR) LEN(10) /* library name   */
                                    /* of receiver to        */
                                    /* restore               */
DCL &RCODE TYPE(*CHAR) LEN(2) VALUE(x'0001')
/* reason code 1 of CPF7053         */
RCVMSG PGMQ(*SAME &PGMNM) MSGTYPE(*EXCP) WAIT(0) +
   RMV(*NO) MSGDTA(&MSGDATA) MSGID(&MSGID)
```

```
IF COND(&MSGID *EQ 'CPF9801') THEN(DO)   /* CPF9801 occurred */
CHGVAR &RCVNAME (MSGDATA 1 10))     /* get receiver     */
                                       /* from message     */
                                       /* data             */
CHGVAR &RCVLIB (MSGDATA 11 10))    /* get library      */
                                       /* name from        */
                                       /* message data     */
? RSTOBJ OBJ(&RCVNAME) SAVLIB(&RCVLIB) OBJTYPE(*JRNRCV)
                                    /* display RSTOBJ prompt */
  ENDDO
ELSE DO
IF COND((&MSGID *EQ 'CPF7053') & (MSGDATA 1 2) +
   *EQ &RCODE)) THEN(DO)  /*CPF7053 RC(1) occurred*/
CHGVAR &RCVNAME (MSGDATA 3 10))  /* get receiver     */
                                       /* name from        */
                                       /* message data     */
CHGVAR &RCVLIB (MSGDATA 13 10))  /* get library      */
                                       /* name from        */
                                       /* message data     */
? RSTOBJ OBJ(&RCVNAME) SAVLIB(&RCVLIB) OBJTYPE(*JRNRCV)
                                /* display restore prompt */
    ENDDO
  ELSE
.
.
 ENDDO
ENDPGM
```

Writing output to save media using the Receive Journal Entry command

Here is an example of using the Report Program Generator (RPG) program as the exit program for the Receive Journal Entry (RCVJRNE) command.

Note: You might want to consider using the remote journal function to transfer your journal receiver data to another system instead of the following program.

Figure 35 on page 461 shows an RPG program that is being used as the exit program for the RCVJRNE command. This example writes output onto tape media. See "Journal entries written to an intersystem communications function file" on page 462 for a discussion of changing the sample to write output to an IBM i ICF file.

```
SEQNBR *... ... 1 ... ... 2 ... ... 3 ... ... 4 ... ... 5 ... ... 6 ... ... 7

  1.00      FTAPE    O   F     300            SEQ
  2.00      IJRNENT      DS                            300
  3.00      I                                        1   5OJOENTL
  4.00      C           *ENTRY    PLIST
  5.00      C                     PARM           JRNENT
  6.00      C                     PARM           CALLCD  1
  7.00      C             CALLCD  IFEQ '1'                       Entry rcvd
  8.00      C* Ensure journal entry is not being truncated
  9.00      C             JOENTL  CABGT300       RETURN     H1   If GT output
 10.00      C                     ADD  1         OUTRCD  70     Bump ctr
 11.00      C                     EXCPTOUTPUT                   Output
 12.00      C                     END                           Entry rcvd
 13.00      C             CALLCD  IFEQ 'O'                      Rdy to wait
 14.00      C                     EXSR FORCE                    Force out
 15.00      C                     END                           Rdy to wait
 16.00      C                     SHTDN                    31   Test shtdwn
 17.00      C   31                DO                            If shtdwn
 18.00      C                     EXSR FORCE                    Force out
 19.00      C                     MOVE '9'       CALLCD         Set to end
 20.00      C                     SETON                    LR   Set LR
 21.00      C                     END                           If shtdwn
 22.00      C             RETURN  TAG                           Return tag
 23.00      C                     RETRN                         Return
 24.00      C             FORCE   BEGSR                         Force out
 25.00      C             OUTRCD  IFNE *ZERO                    If rcds
 26.00      C                     FEOD TAPE                     FEOD
 27.00      C                     Z-ADDO         OUTRCD         Reset
 28.00      C                     END                           If rcds
 29.00      C                     ENDSR                         End subr
 30.00      OTAPE    E                OUTPUT
 31.00      0                         JRNENT
```

Figure 35. Program for writing the RCVJRNE command output to save media

Related information

Journal management

Considerations when writing to tape

A separate job must be in continuous operation and dedicated to converting entries to tape.

Before issuing the Receive Journal Entry (RCVJRNE) command, your job should issue an Override with Tape File (OVRTAPF) command, specifying fixed-length blocked records, to direct the Report Program Generator (RPG) file TAPE to a tape device.

You should not consider this approach with a streaming tape device. A user auxiliary storage pool (ASP) is a preferred solution instead of a tape. However, this approach is similar to writing journal entries to a communications line.

The RPG program is written assuming that the largest journal entry to be passed is 300 bytes. This is the size that is given to the data structure JRNENT. It allows a record size of 175 bytes plus the 125 bytes of journal entry identifier information and qualifier information. A check is made in the program to ensure that the record image is not being truncated:

- If a code of 1 is passed from the RCVJRNE command, the program ensures that the journal entry does not exceed 300 bytes. If it does, the program sets on the H1 indicator and returns. The program adds 1 to the counter and writes the record to the tape output file. Because this is an output-only file, RPG automatically blocks the records within the RPG program.

  When full, the block is passed to tape data management, where additional blocking can occur and double-buffering to the tape device is provided. This ensures that the tape performance is optimal. Because the records are not written directly to tape when the program requests the output, there can be some interval of time before the records are written to the external media.
- When a code of 0 is passed from the RCVJRNE command, no more entries exist in the journal. On the return to the RCVJRNE command, the DELAY parameter value specified on the RCVJRNE command is used to wait before checking for additional entries. To avoid keeping the records in the various buffers

while the delay occurs, the program forces the records to the tape device by using the force end-of-data operation (FEOD).

This causes all records in either the RPG or tape data management buffers to be written to the tape device, and a device completion notification to be received before proceeding to the next instruction. If there is less than a full block of records, a short block is written to tape. Tape data management correctly handles the short block if the tape is read in a subsequent program. When the return occurs to the RCVJRNE command, the delay time occurs whether any journal entries have arrived since the last time the exit program was called.

The RPG program increments a counter each time a record is written and resets it when the FEOD operation is used. The program issues the FEOD operation only if a record has been written which avoids calling tape data management when there are no records to be written. (If tape data management has no records in its buffers when the FEOD operation occurs, no empty block is written, but system overhead occurs.)

The RPG program uses the SHTDN operation code to check for requests to end the job from external functions such as an End Job (ENDJOB) or End Subsystem (ENDSBS) command with OPTION(*CNTRLD). If end-of-job is requested, the program forces the records from the buffers, sets the counter to 9 (which tells the RCVJRNE command to complete normally, and sets the LR indicator on). The RETRN operation is then issued and:

- If LR is on, the program's working storage is returned to the system.
- If LR is off, the program remains active and waits to be called again by the RCVJRNE command.

Writing to tape occurs either by the buffers being full or when the FEOD operation is used. This trade-off allows good performance when many journal entries are written and minimizes the number of times the FEOD operation is used to ensure that the entries are actually on the tape. With the sample program, the value of the DELAY parameter and the work management specifications for your job (for example, pool size and priority) are the major factors controlling the frequency with which the entries are written and the performance implications on the system for this function.

If the system ends abnormally while the job is running, so that a successful end-of-file indication is not written, the subsequent reading of the tape can produce results that cannot be predicted. Blocks that were successfully written can be correctly read. The last block and any subsequent data that is on the tape from a previous use can produce results that cannot be predicted. Copy the tape to a database file and examine the contents before using the data.

The journal sequence numbers are in ascending sequence (unless they have been reset) and can be used to determine where the logical end-of-file exists. To avoid confusion, delete the tapes used for this type of approach.

Assume, for example, that the largest record size journaled was 175 bytes and the tape record size 300 bytes, as in Figure 35 on page 461. If you need to increase the tape record size, change the value of 300 in the RPG file description specification, the input specification, and factor 2 of the CABGT operation code. If there are some significantly larger records being journaled, consider how much excess media is used. An alternative would be to examine the individual fields (**JOENTL**) and write two or more small records for each large record.

Journal entries written to an intersystem communications function file

The topic discusses the differences in programming when you use an intersystem communications function (ICF) file instead of a tape file as output for the Receive Journal Entry (RCVJRNE) command.

Refer to the program in the Figure 35 on page 461 topic.

If you use an ICF file to transmit journal entries to another system, the end-of-data operation (FEOD) operation does not apply. Instead, data description specifications (DDS) words (for example, FRCDTA) are used to force records from the buffers.

Typically the number of blocks that are transmitted to tape that uses records less than 175 bytes is a minimal performance consideration. However, on communications lines this number can be significant. To

avoid sending unnecessary trailing blanks, consider decreasing the length of the record that is transmitted that uses the variable length function (VARLEN DDS keyword).

Appendix A. Licensed Internal Code installation error displays

Use the Licensed Internal Code installation error displays to determine what actions to take. For example, the system might be missing the load-source disk.

One of the following three displays might be shown if you chose option 1 (restore) on the install selection menu and the selected disk is not currently a load-source disk unit. In this case, the restore cannot be done. If the selected disk is the correct one to install to, return to the selection menu and choose the correct install option 2, 3, 4, or 5 to initialize the disk and perform an installation. If the selected disk is not the correct one, or if an existing load-source disk unit should have been found, follow the appropriate procedures to determine why the correct disk did not report in, or was not recognized.

If there is information about a missing disk (second or third of the following three displays), it indicates what the last load-source disk unit was on this system. If that disk still exists (it has not been removed or replaced), determine why it was not found. If that disk has been removed or replaced, then this is just informational and might not indicate an error.

```
            Restore Licensed Internal Code

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The disk selected has not previously been a load source.  The
restore of the Licensed Internal Code cannot be done.


Press Enter to return to the Install Licensed Internal Code
screen.
```

```
             Restore Licensed Internal Code

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

    The load source disk could not be found (see disk information
    below).


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The disk selected has not previously been a load source.  The
restore of the Licensed Internal Code cannot be done.


Press Enter to return to the Install Licensed Internal Code
screen.
```

© Copyright IBM Corp. 1997, 2025

```
                     Restore Licensed Internal Code

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

   The load source disk and its mirrored pair could not be found
   (see disk information below).


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----
   ----------     ----    ---    ----      ----        ----

The disk selected has not previously been a load source.  The
restore of the Licensed Internal Code cannot be done.


Press Enter to return to the Install Licensed Internal Code
screen.
```

The following display can be shown if you chose option 1 (restore) on the install selection menu, but the release level of the Licensed Internal Code on the install media cannot be restored over the current release level on disk. Verify that you have the correct install media (version/release/modification level). If the level is correct, then you must do an initialize and install to get the new Licensed Internal Code installed over the existing Licensed Internal Code on the disk.

```
                      Restore Licensed Internal Code

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The release level of the Licensed Internal Code on the distribution
media cannot be restored over the existing release level on the
selected disk.


Press Enter to return to the Install Licensed Internal Code
screen.
```

The following display can be shown if you chose option 1 (restore) on the install selection menu and the selected disk is currently a load-source disk unit, but the pertinent data on the disk cannot be read and, therefore, a restore cannot be done. You must do an initialize and installation to install the new Licensed Internal Code on this disk.

```
                      Restore Licensed Internal Code

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The selected disk cannot be read.  The restore of the Licensed
Internal Code cannot be done.


Press Enter to return to the Install Licensed Internal Code
screen.
```

The following display can be shown if two load-source disk units are detected on your system. The best disk was selected for the restore or installation. The data about the other disk is informational so you can verify that the correct disk was selected. If the correct disk was not selected, follow the appropriate procedures to disable or remove the selected disk , so that the other disk is selected when you rerun the task.

```
                  Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Warning:
    Another load source disk has also been found on this system.
    If you continue the restore or install, the disk listed
    above will be used.


Additional load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


Press Enter to continue the restore or install on the selected
disk.
```

The following display is shown if mirroring is active but one disk of the load source mirrored pair cannot be found. The restore or installation can still continue on the selected disk, but will not be mirrored until the missing disk becomes again active. You might want to follow the appropriate procedures to determine why one of the disks were not found.

```
                  Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Warning:
    The mirrored unit for this load source was not found (see
    disk information below).  The restore or install can continue
    on the selected load source.  The missing mirrored unit will
    be suspended when the restore or install is complete.


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


Press Enter to continue the restore or install on the selected
disk.
```

The following two displays are shown if the disk selected for the install to is not the same disk that was previously the load source on this system. If the missing disk should still exist (it has not been removed or replaced), determine why it was not found. If the disk was removed or replaced, this data is just informational and might not indicate an error.

```
                  Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Warning:
    The load source disk could not be found (see disk information
    below).


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


Press Enter to continue the initialize and install on the
selected disk.
```

```
                  Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Warning:
    The load source disk and its mirrored pair could not be found
    (see disk information below).


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----
   ----------     ----    ---    ----      ----        ----


Press Enter to continue the initialize and install on the
selected disk.
```

The following display is shown if mirroring is active and the active load-source disk unit cannot be found. One unit of the load source mirrored pair was found but is not currently active. You can install to it, but you will not be allowed to initial program load (IPL) past dedicated service tools (DST) with it. You might want to follow the appropriate procedures to determine why the active load-source disk unit cannot be found.

```
                  Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Warning:
    A load source disk could not be found (see disk information
    below).
    The disk selected to be the load source (see above) is
    suspended.  You can install to it and perform an IPL from it to
    get to DST and perform disk diagnostics.  However, you will not
    be able to perform an IPL past DST with it.


Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

Press Enter to continue the restore or install on the selected
disk.
```

One of the following three displays is shown in the following situations:

- No valid, functional load-source disk unit exists on the system.
- No disk units can be made into load-source disk units.

If information is supplied about a missing disk (second and third of the three displays), it indicates what was the last load-source disk unit on this system. If that disk still exists (it has not been removed or replaced), then determine why it was not found. If that disk has been removed or replaced, then this data is just informational and might not be the reason for the error.

```
               Install Licensed Internal Code - Error

Error:
    A disk could not be selected to be the load source.
    You can return to the dedicated service tools display and
    run diagnostics to determine why a disk could not be
    selected.

    Correct the problem and install the Licensed Internal Code
    again.


Press Enter to return to the dedicated service tools display.
```

```
                   Install Licensed Internal Code - Error

Error:
    The load source disk could not be found (see disk information
below).

Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

    A disk could not be selected to be the load source.
    You can return to the dedicated service tools display and
    run diagnostics to determine why a disk could not be
    selected.

    Correct the problem and install the Licensed Internal Code
    again.


Press Enter to return to the dedicated service tools display.
```

```
                   Install Licensed Internal Code - Error

Error:
    The load source disk and its mirrored pair could not be found
    (see disk information below).

Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----
   ----------     ----    ---    ----      ----        ----

    A disk could not be selected to be the load source.
    You can return to the dedicated service tools display and
    run diagnostics to determine why a disk could not be
    selected.

    Correct the problem and install the Licensed Internal Code
    again.

Press Enter to return to the dedicated service tools display.
```

One of the following two displays is shown if a disk is found, but it is not at a valid address to be the load source.

If there is information about a missing disk (the second display), it indicates what the last load-source disk unit was on this system. If that disk still exists (it has not been removed or replaced), determine why it was not found. If it has been removed or replaced, this is just informational and might not be the reason for the error.

```
                  Install Licensed Internal Code - Error

Error:
    A disk was found, but it is not at a valid address to be the
    load source device.

Selected disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


The install cannot be done.  Press Enter to return to the Dedicated
Service Tools display.
```

```
                  Install Licensed Internal Code - Error

Error:
    A disk was found, but it is not at a valid address to be the
    load source device.

Selected disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The following disk was a load source previously, but could not be
found.

Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


The install cannot be done.  Press Enter to return to the Dedicated
Service Tools display.
```

The following display is shown if an existing load-source disk unit is found, but is not at a valid address to be the load source. If it was intentionally moved, determine why no other disk can be found to which to install. If this is the correct disk , determine why it is not at a valid address.

```
                  Install Licensed Internal Code - Error

Error:
    The following disk was a load source previously, but it is not
    currently at a valid address to be the load source device.

Selected disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----


The install cannot be done.  Press Enter to return to the Dedicated
Service Tools display.
```

The following display is shown if an existing load-source disk unit was found and:

- Is not at a valid address to be the load source.
- Is one unit of a mirrored pair.
- Is currently not the active load source.

Information about the previously active load source is also listed. If this load source is still available to install to, use the appropriate procedures to determine why it cannot be found.

```
                  Install Licensed Internal Code - Error

Error:
    The following disk was a load source, but it is not currently
    active, and it is not at a valid address to be the load source
    device.

Selected disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The following disk was the previously active load source, but it
could not be found.

Missing load source disk :
  Serial Number   Type   Model   I/O Bus   Controller  Device
   ----------     ----    ---    ----      ----        ----

The install cannot be done.  Press Enter to return to the Dedicated
Service Tools display.
```

Appendix B. Example: Disaster recovery plan

The objective of a disaster recovery plan is to ensure that you can respond to a disaster or other emergency that affects information systems and minimize the effect on the operation of the business. When you have prepared the information described in this topic collection, store your document in a safe, accessible location off site.

Section 1. Example: Major goals of a disaster recovery plan

Here are the major goals of a disaster recovery plan.

- To minimize interruptions to the normal operations.
- To limit the extent of disruption and damage.
- To minimize the economic impact of the interruption.
- To establish alternative means of operation in advance.
- To train personnel with emergency procedures.
- To provide for smooth and rapid restoration of service.

Section 2. Example: Personnel

You can use the tables in this topic to record your data processing personnel. You can include a copy of the organization chart with your plan.

Data processing personnel			
Name	**Position**	**Address**	**Telephone**

Data processing personnel			
Name	**Position**	**Address**	**Telephone**

Data processing personnel			
Name	**Position**	**Address**	**Telephone**

© Copyright IBM Corp. 1997, 2025

Data processing personnel			
Name	**Position**	**Address**	**Telephone**

Section 3. Example: Application profile

You can use the Display Software Resources (DSPSFWRSC) command to complete the table in this topic.

Application profile				
Application name	**Critical Yes / No**	**Fixed asset Yes / No**	**Manufacturer**	**Comments**

Comment legend:

1. Runs daily.
2. Runs weekly on _____.
3. Runs monthly on _____.

Section 4. Example: Inventory profile

You can use the Work with Hardware Products (WRKHDWPRD) command to complete the table in this topic.

Application profile					
Manufacturer	**Description**	**Model**	**Serial number**	**Own or leased**	**Cost**

Application profile					
Manufacturer	**Description**	**Model**	**Serial number**	**Own or leased**	**Cost**

Notes:

1.
This list should be audited every _____ months.

2.
This list should include the following items:

Processing units
Disk units
Models
Workstation controllers
Personal computers
Spare workstations
Telephones
Air conditioner or heater
System printer
Tape and optical devices
Controllers
I/O Processors
General data communication
Spare displays
Racks
Humidifier or dehumidifier

Miscellaneous inventory		
Description	**Quantity**	**Comments**

Note: This list should include the following items:

Tapes
PC software
File cabinet contents or documentation
Tape vault contents
Optical media
CDs and DVDs
Emulation packages
Language software (such as COBOL and RPG)
Printer supplies (such as paper and forms)

Related reference

Section 12. Disaster site rebuilding
Use this information to do disaster site rebuilding.

Section 5. Information services backup procedures

Use these procedures for information services backup.

- IBM i environment
  - Daily, journals receivers are changed at _____ and at _____.
  - Daily, a saving of changed objects in the following libraries and directories is done at _____:
    - _____
    - _____
    - _____
    - _____

- _____
- _____
- _____
- _____

The preceding procedure also saves the journals and journal receivers.

– On _____ at _____ a complete save of the system is done.
– All save media is stored off-site in a vault at _____ location.

• Personal Computer

– It is suggested that all personal computers be backed up. Copies of the personal computer files should be uploaded to the IBM i environment on _____ (date) at _____ (time), just before a complete save of the system is done. It is then saved with the normal system save procedure. This provides for a more secure backup of personal computer-related systems where a local area disaster can wipe out important personal computer systems.

Section 6. Disaster recovery procedures

For any disaster recovery plan, these three elements should be addressed.

Emergency response procedures
To document the appropriate emergency response to a fire, natural disaster, or any other activity in order to protect lives and limit damage.

Backup operations procedures
To ensure that essential data processing operational tasks can be conducted after the disruption.

Recovery actions procedures
To facilitate the rapid restoration of a data processing system following a disaster.

Disaster action checklist

This checklist provides possible initial actions that you might take following a disaster.

1. Plan initiation:
   a) Notify senior management
   b) Contact and set up disaster recovery team
   c) Determine degree of disaster
   d) Implement proper application recovery plan dependent on extent of disaster
   e) Monitor progress
   f) Contact backup site and establish schedules
   g) Contact all other necessary personnel–both user and data processing
   h) Contact vendors–both hardware and software
   i) Notify users of the disruption of service
2. Follow-up checklist:
   a) List teams and tasks of each
   b) Obtain emergency cash and set up transportation to and from backup site, if necessary
   c) Set up living quarters, if necessary
   d) Set up eating establishments, as required
   e) List all personnel and their telephone numbers
   f) Establish user participation plan
   g) Set up the delivery and the receipt of mail
   h) Establish emergency office supplies

i) Rent or purchase equipment, as needed

j) Determine applications to be run and in what sequence

k) Identify number of workstations needed

l) Check out any off-line equipment needs for each application

m) Check on forms needed for each application

n) Check all data being taken to backup site before leaving and leave inventory profile at home location

o) Set up primary vendors for assistance with problems incurred during emergency

p) Plan for transportation of any additional items needed at backup site

q) Take directions (map) to backup site

r) Check for additional magnetic tapes, or optical media if required

s) Take copies of system and operational documentation and procedural manuals.

t) Ensure that all personnel involved know their tasks

u) Notify insurance companies

Related tasks

Section 7. Recovery plan for mobile site
This topic provides information about how to plan your recovery task at a mobile site.

Recovery startup procedures for use after actual disaster

Consider these recovery startup procedures for use after actual disaster.

1. Notify _____ Disaster Recovery Services of the need to utilize service and of recovery plan selection.

   Note: Guaranteed delivery time countdown begins at the time _____ is notified of recovery plan selection.

   - Disaster notification numbers

     _____or _____

   These telephone numbers are in service from _____ am until _____ pm Monday through Friday.
2. Disaster notification number

   _____

   This telephone number is in service for disaster notification after business hours, on weekends, and during holidays. Please use this number only for the notification of the actual disaster.
3. Provide _____ with an equipment delivery site address (when applicable), a contact, and an alternate contact for coordinating service and telephone numbers at which contacts can be reached 24 hours a day.
4. Contact power and telephone service suppliers and schedule any necessary service connections.
5. Notify _____ immediately if any related plans should change.

Section 7. Recovery plan for mobile site

This topic provides information about how to plan your recovery task at a mobile site.

1. Notify _____ of the nature of the disaster and the need to select the mobile site plan.
2. Confirm in writing the substance of the telephone notification to _____ within 48 hours of the telephone notification.
3. Confirm all needed backup media are available to load the backup machine.
4. Prepare a purchase order to cover the use of backup equipment.

5. Notify _____ of plans for a trailer and its placement (on _____ side of _____).
6. Depending on communication needs, notify telephone company (_____) of possible emergency line changes.
7. Begin setting up power and communications at _____.
   a) Power and communications are prearranged to hook into when trailer arrives.
   b) At the point where telephone lines come into the building (_____), break the current linkage to the administration controllers (_____).

      These lines are rerouted to lines going to the mobile site. They are linked to modems at the mobile site.

      The lines currently going from _____ to _____ would then be linked to the mobile unit via modems.
   c) This can conceivably require _____ to redirect lines at _____ complex to a more secure area in case of disaster.
8. When the trailer arrives, plug into power and do necessary checks.
9. Plug into the communications lines and do necessary checks.
10. Begin loading system from backups.
11. Begin normal operations as soon as possible:
    a) Daily jobs
    b) Daily saves
    c) Weekly saves
12. Plan a schedule to back up the system in order to restore on a home-base computer when a site is available. (Use regular system backup procedures).
13. Secure mobile site and distribute keys as required.
14. Keep a maintenance log on mobile equipment.

Related tasks

Disaster action checklist
This checklist provides possible initial actions that you might take following a disaster.

Related reference

Section 9. Restoring the entire system
You can learn how to restore the entire system.

Mobile site setup plan

You can attach the mobile site setup plan here.

Communication disaster plan

You can attach the communication disaster plan, including the wiring diagrams here.

Electrical service

You can attach the electrical service diagram here.

Section 8. Recovery plan for hot site

An alternate hot site plan should provide for an alternative (backup) site. The alternate site has a backup system for temporary use while the home site is being reestablished.

1. Notify _____ of the nature of the disaster and of its desire for a hot site.
2. Request air shipment of modems to _____ for communications. (See _____ for communications for the hot site.)

3. Confirm in writing the telephone notification to _____ within 48 hours of the telephone notification.
4. Begin making necessary travel arrangements to the site for the operations team.
5. Confirm that you have enough save media and that it is packed for shipment to restore on the backup system.
6. Prepare a purchase order to cover the use of the backup system.
7. Review the checklist for all necessary materials before departing to the hot site.
8. Make sure that the disaster recovery team at the disaster site has the necessary information to begin restoring the site.
9. Provide for travel expenses (cash advance).
10. After arriving at the hot site, contact home base to establish communications procedures.
11. Review materials brought to the hot site for completeness.
12. Start to load the system from the save media.
13. Begin normal operations as soon as possible:
    a) Daily jobs
    b) Daily saves
    c) Weekly saves
14. Plan the schedule to back up the hot-site system in order to restore on the home-base computer.

Related reference

Section 12. Disaster site rebuilding
Use this information to do disaster site rebuilding.

Alternate-site system configuration

You can attach the alternate-site system configuration here.

Section 9. Restoring the entire system

You can learn how to restore the entire system.

To get your system back to the way it was before the disaster, use the procedures in “Checklist 20: Recovering your entire system after a complete system loss” on page 93.

Before you begin - Find the following save media, equipment, and information from the onsite tape vault or the offsite storage location -

- If you install from the alternate installation device, you need both your save media and the distribution media that contains the Licensed Internal Code.
- All save media from the most recent complete save operation
- The most recent save media from saving security data (SAVSECDTA or SAVSYS)
- The most recent save media from saving your configuration, if necessary
- All save media that contains journals and journal receivers that you saved since the most recent daily save operation
- All save media from the most recent daily save operation
- PTF list (stored with the most recent complete save media, weekly save media, or both)
- Save media list from most recent complete save operation
- Save media list from most recent weekly save operation
- Save media list from daily saves
- History log from the most recent complete save operation
- History log from the most recent weekly save operation
- History log from the daily save operations

- The Installing, upgrading, or deleting IBM i and related software PDF.
- The Recovering your system PDF.
- Telephone directory
- Modem manual
- Toolkit

Related tasks

Section 7. Recovery plan for mobile site
This topic provides information about how to plan your recovery task at a mobile site.

Section 10. Rebuilding process

The management team must assess the damage and begin the reconstruction of a new data center.

If the original site must be restored or replaced, the following questions are some of the factors to consider:

- What is the projected availability of all needed computer equipment?
- Will it be more effective and efficient to upgrade the computer systems with newer equipment?
- What is the estimated time needed for repairs or construction of the data site?
- Is there an alternative site that more readily can be upgraded for computer purposes?

After the decision to rebuild the data center has been made, go to "Section 12. Disaster site rebuilding" on page 482.

Section 11. Testing the disaster recovery plan

In successful contingency planning, it is important to test and evaluate the plan regularly.

Data processing operations are volatile in nature, resulting in frequent changes to equipment, programs, and documentation. These actions make it critical to consider the plan as a changing document.

Table 81 on page 480 should be helpful for conducting a recovery test.

Table 81. Checklist for testing the disaster recovery plan

Item	Yes	No	Applicable	Not applicable	Comments
Conducting a Recovery Test					

Table 81. Checklist for testing the disaster recovery plan (continued)

Item	Yes	No	Applicable	Not applicable	Comments
1. Select the purpose of the test. What aspects of the plan are being evaluated? 2. Describe the objectives of the test. How will you measure successful achievement of the objectives? 3. Meet with management and explain the test and objectives. Gain their agreement and support. 4. Have management announce the test and the expected completion time. 5. Collect test results at the end of the test period. 6. Evaluate results. Was recovery successful? Why or why not? 7. Determine the implications of the test results. Does successful recovery in a simple case imply successful recovery for all critical jobs in the tolerable outage period? 8. Make suggestions for changes. Call for responses by a given date. 9. Notify other areas of results. Include users and auditors. 10. Change the disaster recovery plan manual as necessary.					
Areas to be tested					

Table 81. Checklist for testing the disaster recovery plan (continued)

Item	Yes	No	Applicable	Not applicable	Comments
1. Recovery of individual application systems by using files and documentation stored off-site. 2. Reloading of system save media and performing an initial program load (IPL) by using files and documentation stored off-site. 3. Ability to process on a different computer. 4. Ability of management to determine priority of systems with limited processing. 5. Ability to recover and process successfully without key people. 6. Ability of the plan to clarify areas of responsibility and the chain of command. 7. Effectiveness of security measures and security bypass procedures during the recovery period. 8. Ability to accomplish emergency evacuation and basic first-aid responses. 9. Ability of users of real time systems to cope with a temporary loss of online information. 10. Ability of users to continue day-to-day operations without applications or jobs that are considered noncritical. 11. Ability to contact the key people or their designated alternates quickly. 12. Ability of data entry personnel to provide the input to critical systems by using alternate sites and different input media. 13. Availability of peripheral equipment and processing, such as printers and scanners. 14. Availability of support equipment, such as air conditioners and dehumidifiers. 15. Availability of support: supplies, transportation, communication. 16. Distribution of output produced at the recovery site. 17. Availability of important forms and paper stock. 18. Ability to adapt plan to lesser disasters.					

Section 12. Disaster site rebuilding

Use this information to do disaster site rebuilding.

- Floor plan of data center.
- Determine current hardware needs and possible alternatives.
- Data center square footage, power requirements and security requirements.
  - Square footage _____
  - Power requirements _____

- Security requirements: locked area, preferably with combination lock on one door.
- Floor-to-ceiling studding
- Detectors for high temperature, water, smoke, fire and motion
- Raised floor

Related tasks

Section 8. Recovery plan for hot site
An alternate hot site plan should provide for an alternative (backup) site. The alternate site has a backup system for temporary use while the home site is being reestablished.

Related reference

Section 4. Example: Inventory profile
You can use the Work with Hardware Products (WRKHDWPRD) command to complete the table in this topic.

Vendors

You can attach the vendors information here.

Floor plan

You can include a copy of the proposed floor plan here.

Section 13. Record of plan changes

Keep your plan current, and keep records of changes to your configuration, your applications, and your backup schedules and procedures.

You can get print a list of your current local hardware by typing the following command:

```
DSPLCLHDW OUTPUT(*PRINT)
```

Appendix C. Recovering your system

The 'Recovering your system' topic provides instructions to completely recover your entire system.

Use these steps if you need to recover your system to the same system (restore to the same system with the same serial number). Use these recovery instructions only if you saved your entire system by using either of the following options:

- Save menu Option 21
- Equivalent save commands as **Save** menu Option 21:

  - SAVSYS
  - SAVLIB LIB(*NONSYS) ACCPTH(*YES)
  - SAVDLO DLO(*ALL) SAVFLR(*ANY)
  - SAV DEV('/QSYS.LIB/*tape-device-name*.DEVD') OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) UPDHST(*YES)

Notes:

- Use rzarmrcvsit.dita for any of the following cases.
  - The system uses the 'Alternate Installation Device Setup' feature that you can define through dedicated service tools (DST) for D-IPL.[1]
  - Mounted user-defined file systems (UDFS) exist on the system before the save.
- Use rzarmrstiasp.dita if your recovery includes independent auxiliary storage pools (ASPs).
- If you use journaling, refer to the section rzarmrstfbj.dita for important restore considerations before you proceed with recovery.
- If you are recovering to a system with a different serial number, use rzarmsteps2.dita.

Check off each item as you complete the task.

1. If you use Operations Console with local area network (LAN) connectivity, you must disconnect the Operations Console from the system. Then resynchronize your person computer device ID and connect it to the Operations Console by using the same as the IBM i device ID.

   After you begin the step, you must continue until you complete the IPL from SAVSYS media. If you are using Operations Console with IBM i V5R4, or later, you must reset only the service tools device ID password on the system. The client automatically attempts to sign on using the reset device ID password for Operations Console.
2. If you are using Operations Console, LAN, or direct connection, disconnect from your system and close Operations Console.

   Restart Operations Console and reconnect by using 'Service Tool User ID' QSECOFR and password of QSECOFR to sign on to your system.
3. Load the initial program on the system from the first SAVSYS media.
   a) Mount the first SAVSYS media on the alternate IPL device. Wait for the READY status.
   b) At the central processing unit (CPU) control panel, place the system in MANUAL mode.
   c) Press the Function Select switch to display 02 (IPL) in the Function display.
   d) Press Enter.
   e) Press the Function Select switch to display D (IPL from tape or DVD) in the Data display.
   f) Press Enter.

[1] Refer to rzarmusaid.dita.

© Copyright IBM Corp. 1997, 2025

g) If the system is powered down, press the power button on the system to power on the system. Skip to Step “4” on page 486. Otherwise, continue with Step “3.h” on page 486.

h) If the system is powered on, press the Function Select switch to display 03 (continue the IPL) in the Function display.

i) Press Enter. IBM

4. You see one of the following two displays. This shows the language on the distribution media.

   Press Enter to select the language, and then press Enter on the confirmation display.

```
                          Select a Language Group

 The language feature shown is the service tools
 language feature installed on the system.

 Attention: To keep the same language, ensure that the
 media you use for installing the operating system
 matches the language feature shown. If the operating
 system  media does not match what is shown, the
 installation process will attempt to install the
 operating system in a different language feature than
 Licensed Internal Code.

 Type choice, press Enter.

 Language feature . . . . . .  2924
```

```
LANGUAGE FEATURE ===>   2924
```

5. On the 'Install Licensed Internal Code' display, select option 1 (Install Licensed Internal Code).

```
                            Install Licensed Internal Code


Select one of the following:

   1. Install Licensed Internal Code
   2. Work with dedicated service tools (DST)
   3. Define alternate installation device

   Selection
   1
```

6. If a valid and functional load source disk unit was not detected on the partition, you are presented with the 'Select Load Source Device' display. Select the disk unit that you want to become the load source. If you do not see the 'Select Load Source Device' display, continue with step “7” on page 487.

```
                                 Select Load Source Device

Type 1 to select, press Enter.

                                    Sys  Sys   I/O       I/O
Opt  Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
     68-0CCDDA23      6718   050     12    11       0      0     7     0
     68-0DDA4212      6718   050     12    11       0      0     6     0
     68-0C211074      6718   050     12    11       0      0     3     0
     68-0DDE1201      6718   050     12    11       0      0     1     0
     68-0CDDEA44      6718   050     55    11       0      0     4     0
```

a) Press F10 to confirm the selection of the disk that you want to become the load-source disk unit.

```
                                Confirm Load Source Device

You have chosen the following device to be the load source.
Press F10 to confirm your choice.

                                Sys  Sys   I/O       I/O
Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
68-0CDDEA44      6718   050     55    11       0      0     4     0
```

7. On the 'Install Licensed Internal Code (LIC)' display, select option 2 (Install Licensed Internal Code and Initialize System) to start a "Scratch Install" of the LIC.

```
                     Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
   Serial Number        Type        Model        I/O Bus        Controller        Device
   xx-xxxxxxx       xxxx     xxx        x           x            x

Select one of the following:

   1. Restore Licensed Internal Code
 =>2. Install Licensed Internal Code and Initialize System
   3. Install Licensed Internal Code and Recover Configuration
   4. Install Licensed Internal Code and Restore Disk Unit Data
   5. Install Licensed Internal Code and Upgrade Load Source

Selection
2
```

8. On the 'Install LIC and Initialize System - Confirmation' display, press F10 to confirm the initialization and to continue the installation.

```
                   Install LIC and Initialize System - Configuration

Warning:

   All data on this system will be destroyed and the Licensed
   Internal Code will be written to the selected disk if you
   choose to continue the initialize and install.

   Return to the install selection screen and choose one of the
   other options if you want to perform some type of recovery
   after the install of the Licensed Internal Code is complete.


   Press F10 to continue the install.
   Press F12 (Cancel) to return to the previous screen.
   Press F3 (Exit) to return to the install selection screen.
```

a. The 'Initialize the Disk - Status' display appears.

```
                              Initialize the Disk - Status

The load source disk is being initialized.

Estimated time to initialize in minutes:  xx
Elapsed time in minutes . . . . . .:  0.0
```

b. The 'Install Licensed Internal Code - Status' display appears.

```
                        Install Licensed Internal Code - Status


Install of the Licensed Internal Code in progress.

            +--------------------------------------+
Percent     |             XX %                     |
complete    +--------------------------------------+

Elapsed time in minutes . . . . . .:    x.x

Please wait.
```

9. You might see the 'Accept Console' display. If it appears, press F10 to accept the current console.

   On the 'Accept And Set New Console Type On This IPL' display, press Enter to continue.

10. On the 'Disk Configuration Attention Report' display, press F10 to accept any problems and to continue.

```
                           Disk Configuration Attention Report

Type option, press Enter
   5=Display Detailed Report


Press F10 to accept all the following problems and continue.
The system will attempt to correct them.


OPT Problem
_ New disk configuration



F3=Exit                   F10=Accept the problems and continue
```

11. On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)).

```
                            IPL or Install the System

Select one of the following:

   1. Perform an IPL
   2. Install the operating system
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the operating system
   5. Save Licensed Internal Code


Selection
3
```

12. Sign on to DST with service tools user QSECOFR, with the password for the QSECOFR service tools user ID.

```
                        Dedicated Service Tools (DST) Sign On

Type choices, press Enter.

Service tools user . . . . . . . . QSECOFR
Service tools password . . . . . . QSECOFR
```

    a. The password for the service tools profile QSECOFR is expired on a new system and must be changed at first use. If you see the 'Change Service Tools User Password' display, enter in all uppercase letters the current password QSECOFR and a new password, along with the verification password.

```
                        Change Service Tools User Password

Service tools user profile name . . . . : QSECOFR
Password last changed . . . . . : mm/dd/yy

Type choices, press Enter.

Current password  . . . . . : QSECOFR
New password  . . . . . : _______
New password (to verify). . . . . : _______
```

b. Select option 4 (Work with Disk Units).

c. Select option 1 (Work with Disk Configuration).

d. Select option 3 (Work with ASP Configuration).

e. Select option 3 (Add Units to ASPs).

13. On the 'Add Units to ASPs' display, enter 3 to add units to existing ASPs, and add disk units to the system ASP.

```
                           Add Units to ASPs

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs
```

14. On the 'Specify New ASPs to Add Units to' display, enter the ASP number 1 to add disk units to the system ASP (1).

a) Enter the ASP number on each selected unit.

```
                     Specify New ASPs to Add Units to

Specify the new ASP to add each unit to.

Specify  Serial                                    Resource
  ASP    Number          Type  Model  Capacity  Name
    1    21-6C597        4327   050      70564  DD007
    1    50-128840F      2107   A84      70564  DD004
    1    50-128940F      2107   A84      70564  DD005
         50-128A40F      2107   A85      35165  DD011
         50-128B40F      2107   A85      35165  DD003
         68-0C8BA12      6717   050       8589  DD008
         68-0C9D209      6717   050       8589  DD009
         68-606E0        6718   050      17548  DD006
         68-0CDAB10      6718   050      17548  DD010
         21-05348        4327   050      70564  DD012
         21-05322        4327   050      70564  DD013

F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

b) After you complete all units, press Enter.

c) If the list of units is correct, press Enter to start initializing the units.

15. On the 'Problem Report' display, press F10 to ignore the problems and continue.

```
                                  Problem Report

   Note: Some action for the problems listed below might need to
   be taken. Please select a problem to display more detailed
   information about the problem and to see what possible
   action can be taken to correct the problem.


Type option, press Enter.
   5=Display Detailed Report


OPT Problem
   _ Unit possibly configured for Power PC AS

F3=Exit          F10=Ignore problems and continue            F12=Cancel
```

16. On the 'Confirm Add Units' display, press Enter to confirm the selected units.

```
                             Confirm Add Units

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data on units.
Press F12=Cancel to return and change your choice.

            Serial                   Resource
ASP  Unit   Number      Type Model   Name       Protection
  1                                              Unprotected
        1   68-0CDAB35  6718  050    DD001      Unprotected
        3   21-6C597    4327  050    DD007      Unprotected
        4   50-128840F  2107  A84    DD004      Unprotected
        5   50-128940F  2107  A84    DD005      Unprotected
        2   50-100240F  2107  A82    DD002      Unprotected
        6   50-128A40F  2107  A85    DD011      Unprotected

F9=Resulting Capacity              F10=Add and Balance
F11=Display Encryption Status      F12=Cancel
```

a. The 'Function' Status display shows the percentage of completion.

b. The

   ```
   Selected units have been added successfully
   ```

   message appears when the system completes the 'Add Units' process.

   Note: If you require more than ASP 1 configured, you can return to step “13” on page 489 and create more ASPs by selecting 1 to create unencrypted ASPs or selecting 2 to create encrypted ASPs. You can create multiple user ASPs at the same time by specifying ASP numbers 2 through 32 on the selected disk units.

c. Press F12 to return to the 'Work with Disk Configuration' display.

d. If your system requires mirrored protection, continue to Step “16.e” on page 490. If your system does not require mirrored protection, press F3 until you see the 'Exit Dedicated Service Tools (DST)' display. Select option 1 to exit DST and press Enter and continue to Step “16.g” on page 491.

e. To start mirrored protection for your system, follow these instructions:

   i) On the 'Work with Disk Configuration' display, select option 4 (Work with Mirrored Protection).

   ii) On the 'Work with Mirrored Protection' display, select option 2 (Start Mirrored Protection).

   iii) Select an ASP by placing a 1 next to it. Press Enter to start mirrored protection.

   iv) On the 'Confirm Continuation' display, press Enter.

   v) On the 'Confirm Start Mirrored Protection' display, press Enter.

   vi) The 'Function Status' display shows the completion status of the 'Start Mirrored Protection' request.

vii) The

```
Start mirroring completed successfully
```

message appears on the 'Disk Configuration Information Report' display.

viii) Press Enter to continue.

f. To allow security system values to change during your recovery, follow these steps:

i) On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)). Press Enter to continue.

ii) Sign on to DST with your service tools username and password.

iii) Select option 13 (Work with system security).

iv) Type 1 in the **Allow system value security changes** field and press Enter.

v) Press F3 or F12 to get back to the 'IPL or Install the System' display.

g. If you use the Operations Console, follow these instructions to switch from the current console, or to verify the console mode value:

i) On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)). Press Enter to continue.

ii) Sign on to DST with service tools user QSEC0FR, and use your new QSECOFR password.

iii) On the 'Use Dedicated Service Tools (DST)' display, select option 5 (Work with DST environment). Press Enter to continue.

iv) On the 'Work with DST Environment' display, select option 2 (System Devices). Press Enter to continue.

v) On the 'Work with System Devices' display, select option 6 (Console Mode). Press Enter to continue.

vi) On the 'Select Console Type' display, select option 2 (Operations Console). Press Enter to continue.

vii) Press F3 or F12 to get back to the 'IPL or Install the System' display.

17. On the 'IPL or Install the System' display, select option 2 (Install the Operating System).

```
                        IPL or Install the System
                                                                    System: XXXX
Select one of the following:

   1. Perform an IPL
   2. Install the Operating System
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the Operating System
   5. Save Licensed Internal Code

Selection
2
```

a. On the 'Confirm Install of the Operating System' display, press Enter.

b. Make a selection on the 'Install Device Type Selection' display and press Enter

c. The 'Select a Language Group' display shows the primary language feature currently on your save media. Press Enter to accept this value.

```
                          Select a Language Group

Note:  The language feature shown is the language feature
installed on the system.

Type choice, press Enter.

Language feature . . . . . .  2924
```

d. The 'Confirm Language Feature Selection' display appears. Press Enter to continue.

18. On the 'Add All Disk Units to the System' display, select option 1 (Keep the current disk configuration).

```
                          Add All Disk Units to the System

Select one of the following:

1. Keep the current disk configuration
2. Perform disk configuration using DST
3. Add all units to the system auxiliary storage pool (ASP)
4. Add all units to the system ASP and balance data

Selection
1
```

Note: This display does not appear if you selected all disk units that are known to the system on Step "14" on page 489.

19. The 'Licensed Internal Code IPL in Progress' display shows the IPL progress.

```
                        Licensed Internal Code IPL in Progress

IPL:
   Type . . . . . . . . . . . . . . . :    Attended
   Start date and time  . . . . . . . :    xx/xx/xx  xx:xx:xx
   Previous system end  . . . . . . . :    Abnormal
   Current step / total . . . . . . . :       16      16
   Reference code detail  . . . . . . :    C6004065

IPL step                                   Time Elapsed   Time Remaining
  Commit Recovery                            xx:xx:xx        xx:xx:xx
  Data Base Initialization                   xx:xx:xx        xx:xx:xx
  Journal IPL Clean up                       xx:xx:xx        xx:xx:xx
  Commit Initialization                      xx:xx:xx        xx:xx:xx
 >Start the operating system                 xx:xx:xx        xx:xx:xx

Item:
  Current / Total . . . . . . . :

Sub Item:
  Identifier . . . . . . . . . . :    Searching tape for install media
  Current / Total  . . . . . . . :

Wait for next display or press F16 for DST main menu
```

20. On the 'Install the Operating System' display, select option 2 (Change install options). Select to restore your network attributes from your save media. If the date and time are not correct, type new values. Press Enter to continue.

```
                             Install the Operating System

Type options, press Enter.

Install
   option . . . . . 2      1=Take defaults (No other
                             options are displayed)
                           2=Change install options

Date
   Year . . . . . . xx     00-99
   Month. . . . . . xx     01-12
   Day. . . . . . . xx     01-31

Time
   Hour . . . . . . xx     00-23
   Minute . . . . . xx     00-59
   Second . . . . . xx     00-59
```

21. On the 'Specify Install Options' display, specify the following installation options.

```
                          Specify Install Options

 Type choices, press Enter.

    Restore option . . . . 1      1=Restore programs and language objects
                                    from current media set
                                  2=Do not restore programs or
                                    language objects
                                  3=Restore only language objects
                                    from current media set
                                  4=Restore only language objects
                                   from a different media set using the
                                    current install device

    Job and output
      queues option . . . . 2     1=Clear,  2=Keep

    Distribute IBM i on
      available disk units. 1     1=Yes, 2=No
```

If your ASPs had access path recovery times for ASPs that are not reconfigured, the access path recovery times cannot be restored. Use the **Edit Recovery Times for Access Paths (EDTRCYAP)** command to set the times after you reconfigure your ASP configuration.

22. On the 'Specify Restore Options' display, specify 1 for each of the following choices. Then press Enter.

```
                          Specify Restore Options

 Type choices, press Enter.

      Restore from the installation media:

         System information. . . 1    1=Restore,  2=Do not restore
         Edit descriptions . . . 1    1=Restore,  2=Do not restore
         Message reply list. . . 1    1=Restore,  2=Do not restore
         Job descriptions. . . . 1    1=Restore,  3=Keep customization
         Subsystem descriptions. 1    1=Restore,  3=Keep customization
```

23. The IBM i 'Installation Status' display shows the installation status of the IBM i installation profiles and libraries.

```
Message ID . . : CPI2070                                  IBM i Installation Status

        +-----------------------------------------------------------------------------+
Stage 1 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-----------------------------------------------------------------------------+
        0          20              40          60          80          100

Installation                                                                  Objects
Stage                                                          Completed      Restored

1  Creating needed profiles and libraries . . .:
2  Restoring programs to library QSYS . . . . .:
3  Restoring language objects to library QSYS .:
4  Updating program table . . . . . .:
5  Installing database files  . . . . . .:
6  Installing base directory objects. . . . . .:
```

24. The system installs the remaining IBM i objects.

```
Message ID . . : CPI2070                                    IBM i Installation Status

        +-------------------------------------------------------------------------------+
Stage 4 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-------------------------------------------------------------------------------+
        0              20              40              60              80              100

Installation                                                    Objects
Stage                                              Completed    Restored

   2  Restoring programs to library QSYS  . . . . :      x          09
   3  Restoring language objects to library QSYS. :      x          03
>> 4  Updating program table  . . . . . . . . . :      x
   5  Installing database files . . . . . . . . . :      x
   6  Installing base directory objects . . . . . :
   7  Completing installation . . . . . . . . . . :
```

25. On the 'Sign On' display, logon with user QSECOFR. You do not need to enter a password now.
26. On the 'IPL Options' display, enter the correct values system date, time, and time zone.

    Set the following options to Y:

    - Start system to restricted state
    - Set major system options
    - Define or change the system at IPL

```
                                    IPL Options

Type choices, press Enter.

   System date  . . . . . . . . . . . . . . . xx/xx/xx   MM / DD / YY
   System time  . . . . . . . . . . . . . . . xx:xx:xx   HH : MM : SS
   System time zone . . . . . . . . . . . .   Q0000UTC   F4 for list
   Clear job queues . . . . . . . . . . . . . N          Y=Yes, N=No
   Clear output queues  . . . . . . . . . . . N          Y=Yes, N=No
   Clear incomplete job logs  . . . . . . . . N          Y=Yes, N=No
   Start print writers  . . . . . . . . . . . N          Y=Yes, N=No
   Start system to restricted state . . . . . Y          Y=Yes, N=No

   Set major system options . . . . . . . . . Y          Y=Yes, N=No
   Define or change system at IPL . . . . . . Y          Y=Yes, N=No
```

    a. On the 'Set Major System Options' display, select **N** to disable automatic configuration.

```
                             Set Major System Options

Type choices, press Enter.

Enable automatic configuration . . . . N         Y=Yes, N=No
Device configuration naming  . . . . . NORMAL    *NORMAL, *S36
                                                 *DEVADR
Default special environment  . . . . . *NONE     *NONE, *S36
```

27. The 'Define or Change the System at IPL' display appears.
    a) Select option 3 (System Value Commands).
    b) On the 'Change System Value Commands' display, select 3 (Work with System Values).
    c) On the 'Work with System Values' display, select the system value that you plan to change by placing a 2 next to it.

       Press Enter only after you select all the values.

       Update the following system values. Write down the existing values so you can update them after the recovery, if necessary.

       - QALWOBJRST to *ALL
       - QFRCCVNRST to 0
       - QIPLTYPE to 2

- QINACTITV to *NONE
- QJOBMSGQFL to *PRTWRAP
- QJOBMSGQMX to a minimum of 30 with a suggestion of 64
- QLOGOUTPUT to *JOBLOGSVR
- QLMTDEVSSN to 0
- QLMTSECOFR to 0
- QMAXSIGN to *NOMAX
- QPFRADJ to 2
- QPWDEXPITV to *NOMAX
- QSCANFSCTL adds *NOPOSTRST
- QVFYOBJRST to 1

d) After the system changes the system values, press F3 twice to return to the 'Define or Change the System at IPL' display.

e) On the 'Define or Change the System at IPL' display, press F3 to exit, and continue the IPL.

28. On the 'Change Password' display, type QSECOFR as the current password.

    Enter a new password. Reenter the password to verify and press Enter. (New password cannot be QSECOFR.)

```
Change Password

Password last changed . . . . . xx/xx/xx

Type choices, press Enter.

   Current password . . . . . . . . . QSECOFR
   New password . . . . . . . . . . . _______
   New password (to verify) . . . . . _______
```

29. To configure tape units that are 3490 Model E or F use the following instructions:

a) Use the **Work with Hardware Resource (WRKHDWRSC)** command to determine the tape controller name: `WRKHDWRSC TYPE(*STG)`

b) Locate the tape controller on the 'Work with Hardware Resources' display.

c) Type a 9 (Work with resource) next to the tape controller name and press the Enter key.

Note: If the tape controller is not listed on the display, you need to select other resources, such as disk storage controllers. For some system models, tape units are now attached through combined-function input/output processors (IOPs). Look through the resources until you find the tape unit you want.

d) Locate the resource name for the tape unit (such as TAP01).

e) Enter a 5 (Work with Configuration Descriptions) in the **Opt** column next to the tape resource name and press the Enter key.

You are shown the 'Work with Configuration Descriptions' display.

f) Type a 1 (Create) in the **Opt** field and a name for the tape device description (such as TAP01) in the **Description** field.

Press the Enter key. You are shown the 'Create Device Description (Tape)' display.

g) Change any values that you want to change, then press the Enter key (twice) to create the device description.

You are shown the 'Work with Configuration Descriptions' display again. The device that you created appears on the display.

h) Type an 8 (Work with configuration status) in front of the new device description.

You are shown the 'Work with Configuration Status' display.

i) Type a 1 (Vary on or Make available) in front of the new device.

If the status does not change to Varied On or Available, wait a few minutes. Then press F5 (Refresh). If the status still does not change to Varied On or Available, follow normal problem analysis for the device.

j) Press F3 until you return to the main menu.

```
IBM i Main Menu

Select one of the following:

   1. User tasks
   2. Office tasks
   3. General system tasks
   4. Files, libraries, and folders
   5. Programming
   6. Communications
   7. Define or change the system
   8. Problem handling
   9. Display a menu
  10. Information Assistant options
  11. IBM i Access tasks

  90. Sign off

Selection or command
```

30. On the IBM i 'Main Menu' display, type the command, WRKRPYLE, and check to see whether CPA3709 is there.

    If it is not, determine an available sequence number and then press F6 to add `MSGID(CPA3709) RPY(G)` by using the available sequence number. Press F5 to Refresh and verify that you added CPA3709.

    a. Type the command `CHGJOB INQMSGRPY(*SYSRPYL)` to update the current job to use the system reply list for inquiry messages.

31. On the IBM i 'Main Menu' display, type `GO RESTORE` to access the IBM i 'Restore' display.

    a) On the 'Restore' display, select option 21 (Restore System and User Data).

    b) Press Enter to continue.

32. On the 'Specify Command Defaults' display, enter the name of the tape drive that you are using for the restore.

    a) **Set Prompt for command** to N.

    b) **Set Message queue delivery** to *NOTIFY.

    c) **Set Spooled file data** to *NEW.

```
                          Specify Command Defaults

Type choices, press Enter.

  Devices  . . . . . .   TAP01       Names
                         ----------
                         ----------
                         ----------

  Prompt for commands  . . . . .   N            Y=Yes, N=No

  Message queue delivery . . . .   *NOTIFY      *BREAK, *NOTIFY

  Restore to different system      N            Y=Yes, N=No

  Spooled file data  . . . . . .   *NEW         *NEW, *NONE

  End TCP/IP wait time . . . . .   300          *NONE, 1-9999 seconds

  Print joblog . . . . . .   Y            Y=Yes, N=No
```

    d) Press Enter to continue ending the subsystems.

    The restore process starts to run unattended. The restore process stops only if the system requires a tape mount for the restore process to continue. The restore process runs to completion.

33. If you receive message CPF3204 (`Cannot find object needed for file xxxx in xxx`), during the restore of your libraries, the necessary based-on physical files were not yet restored when the logical files were attempting to restore.

    You will need to restore logical files after the physical files are restored by using one of the following options:

    - Run this command to restore only the logical files that were not restored previously.

      ```
      RSTLIB SAVLIB(*NONSYS) DEV(media-device-name)
      OPTION(*NEW) MBROPT(*ALL) ALWOBJDIF(*ALL)
      ```

    - Run this command for each library that has logical files that need to be restored.

      ```
      RSTLIB SAVLIB(library-name) DEV(media-device-name)
      OPTION(*NEW)  MBROPT(*ALL) ALWOBJDIF(*ALL)
      ```

    - Run this command for each logical file that needs to be restored.

      ```
      RSTOBJ OBJ(logical-file-name) SAVLIB(library-name)
      DEV(media-device-name) OBJTYPE(*FILE) MBROPT(*ALL) ALWOBJDIF (*ALL)
      ```

34. Update program temporary fix (PTF) information for all PTF save files in library QGPL by typing: `UPDPTFINF`.
35. Reapply any PTFs that you applied since the last time that you saved your system data.
    a) Locate the most recent cumulative PTF tape.
    b) From a command-line, enter `GO  PTF` to access the **PTF** menu.
    c) Select option 8 (Install program temporary fix package), on the **PTF** menu.

       This installs all the PTFs in the cumulative PTF package for the licensed programs that are installed on your system. Refer to the 'IBM i PTF Shipping Information Letter' for special instructions.

       Notes:

       - If you want to restore individual PTFs, see the Basic system operations topic for more information about applying individual PTFs.
       - If you do not have the PTFs that you need, order them and apply them later.
36. If necessary, change the system values back to the original values that you wrote down in step “27.c” on page 494, by using the Work with System Value (WRKSYSVAL) command.
37. If you do not know the password for the restored QSECOFR profile, change the password before you sign off. Type the following command:`CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)`.
38. To place scheduled jobs on hold, type `WRKJOBSCDE` and select option 3 to hold any scheduled jobs.

    You need to release the jobs in Step “43” on page 498 after you complete the restore.
39. Type the command `SIGNOFF *LIST` or `DSPJOBLOG * *PRINT`.

    Check the job log to verify that the job restored all objects. To verify whether the job restored all objects, spool the job log for printing along with any remaining spooled job output. Check for any error messages. Correct the errors and restore those objects from media.
40. IPL of the system.
    a) On the CPU control panel, select the **NORMAL** setting.
    b) Type the following command: `PWRDWNSYS OPTION(*IMMED) RESTART(*YES *FULL) IPLSRC(B)`
41. If IBM Content Manager OnDemand for i is installed on your system, restart the journaling for Content Manager OnDemand by typing the following commands.

```
CALL QRDARS/QRLCSTRJ PARM('RLC')
CALL QRDARS/QRLCSTRJ PARM('RLR')
```

42. If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3 installed, run command `CALL QCCADEV/QCCAELOAD`.
43. Release the jobs that you previously held in step “38” on page 497.

Related information

Console Advanced Topics

System values

Appendix D. Recovering your system to a different system

Use this information if you need to recover your system to a different system (for example, a system with a different serial number).

Note: If you are doing a system migration, you must also see the Data migration topic.

Note: If the serial number of recovery system is different than the source system and it is to be used in production, new license keys need to be obtained for the system. License keys have a 70-day grace period before they need to be added.

Use these recovery instructions only if you saved your entire system by using either of the following options:

- **Save** menu Option 21
- Equivalent save commands as **Save** menu Option 21:

  - SAVSYS
  - SAVLIB LIB(*NONSYS) ACCPTH(*YES)
  - SAVDLO DLO(*ALL) SAVFLR(*ANY)
  - SAV DEV('/QSYS.LIB/*tape-device-name*.DEVD') OBJ(('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) UPDHST(*YES)

Important:

- Use rzarmrcvsit.dita for any of the following cases:
  - Mounted user-defined file systems (UDFS) were on the system before the save.
  - The system uses the 'Alternate Installation Device Setup' feature that you can define through dedicated service tools (DST) for D-IPL.[2]
- Use rzarmrstiasp.dita if your recovery includes independent auxiliary storage pools (ASPs).
- If you use journaling, refer to the section rzarmrstfbj.dita for important restore considerations before you proceed with recovery.
- If a tape or optical device other than the primary installation device is to be used for the recovery, refer to the rzarmusaid.dita topic.

Check off each item as you complete the task.

1. If you use Operations Console with local area network (LAN) connectivity, you must disconnect the Operations Console from the system. Then resynchronize your person computer device ID and connect it to the Operations Console by using the same as the IBM i device ID.

   After you begin this step, you must continue until you complete the initial program load (IPL) from SAVSYS media. If you are resynchronizing Operations Console in V5R3 or later, use the following instructions. If you are resynchronizing a previous release of Operations Console, see Console Advanced Topics.

   a. If Operations Console is connected to the system, select the connection name in Operations Console that you are changing; then click **Connection** > **Disconnect**.
   b. On the Operations Console, select the connection name that you are changing, then click **Connection** > **Properties**.
   c. Select the 'Device ID' page.
   d. Click **Reset**.

[2] Refer to rzarmusaid.dita.

© Copyright IBM Corp. 1997, 2025

e. Specify the correct Access password and click **OK**.

f. Click **OK**.

g. Select the connection name that you changed, then click **Connection** > **Connect**.

2. If you are using Operations Console, LAN, or direct connect, disconnect from your system and close Operations Console.

   Restart Operations Console and reconnect by using the user ID of QSECOFR and password of QSECOFR to sign on to your system.

3. IPL the system from the first SAVSYS media.

   a) Mount the first SAVSYS media on the alternate IPL device. Wait for the READY status.

   b) At the central processing unit (CPU) control panel, place the system in MANUAL mode.

   c) Press the Function Select switch to display 02 (IPL) in the Function display.

   d) Press Enter.

   e) Press the Function Select switch to display D (IPL from tape or DVD) in the Data display.

   f) Press Enter.

   g) If the system is powered down, press the power button on the system to power on the system. Skip to Step “4” on page 500. Otherwise, continue with Step “3.h” on page 500.

   h) If the system is powered on, press the Function Select switch to display 03 (continue the IPL) in the Function display.

   i) Press Enter.

4. You see one of the following two displays. This shows the language on the distribution media.

   Press Enter to select the language, and then press Enter on the confirmation display.

```
                          Select a Language Group

The language feature shown is the service tools
language feature installed on the system.

Attention: To keep the same language, ensure that the
media you use for installing the operating system
matches the language feature shown. If the operating
system  media does not match what is shown, the
installation process will attempt to install the
operating system in a different language feature than
Licensed Internal Code.

Type choice, press Enter.

Language feature . . . . . .  2924
```

```
LANGUAGE FEATURE ===>   2924
```

5. On the 'Install Licensed Internal Code' display, select option 1 (Install Licensed Internal Code).

```
                          Install Licensed Internal Code


Select one of the following:

   1. Install Licensed Internal Code
   2. Work with dedicated service tools (DST)
   3. Define alternate installation device

   Selection
   1
```

6. If a valid and functional load source disk unit was not detected on the partition, you are presented with the 'Select Load Source Device' display. You can select the disk unit that you want to become the load source. If you do not see the 'Select Load Source Device' display, continue with step “7” on page 501.

```
                                Select Load Source Device

Type 1 to select, press Enter.

                                  Sys  Sys   I/O       I/O
Opt  Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
     68-0CCDDA23      6718   050     12    11       0      0     7     0
     68-0DDA4212      6718   050     12    11       0      0     6     0
     68-0C211074      6718   050     12    11       0      0     3     0
     68-0DDE1201      6718   050     12    11       0      0     1     0
     68-0CDDEA44      6718   050     55    11       0      0     4     0
```

   a) Press F10 to confirm the selection of the disk that you want to become the load-source disk unit.

```
                                  Confirm Load Source Device

You have chosen the following device to be the load source.
Press F10 to confirm your choice.

                                Sys  Sys   I/O       I/O
Serial Number    Type  Model   Bus  Card  Adapter   Bus   Ctl   Dev
68-0CDDEA44      6718   050     55    11       0      0     4     0
```

7. On the 'Install Licensed Internal Code (LIC)' display, select option 2 (Install Licensed Internal Code and Initialize System) to start a "Scratch Install" of the Licensed Internal Code.

```
                     Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:

   Serial Number        Type        Model        I/O Bus        Controller        Device

   xx-xxxxxxx       xxxx     xxx        x           x            x

Select one of the following:

   1. Restore Licensed Internal Code
 =>2. Install Licensed Internal Code and Initialize System
   3. Install Licensed Internal Code and Recover Configuration
   4. Install Licensed Internal Code and Restore Disk Unit Data
   5. Install Licensed Internal Code and Upgrade Load Source

Selection
2
```

8. On the 'Install LIC and Initialize System - Confirmation' display, press F10 to confirm the initialization and to continue the installation.

```
                    Install LIC and Initialize System - Configuration

Warning:

   All data on this system will be destroyed and the Licensed
   Internal Code will be written to the selected disk if you
   choose to continue the initialize and install.

   Return to the install selection screen and choose one of the
   other options if you want to perform some type of recovery
   after the install of the Licensed Internal Code is complete.


   Press F10 to continue the install.
   Press F12 (Cancel) to return to the previous screen.
   Press F3 (Exit) to return to the install selection screen.
```

a. The 'Initialize the Disk - Status' display appears.

```
                            Initialize the Disk - Status

The load source disk is being initialized.

Estimated time to initialize in minutes:  xx
Elapsed time in minutes . . . . . .:  0.0
```

b. The 'Install Licensed Internal Code - Status' display appears.

```
                        Install Licensed Internal Code - Status


Install of the Licensed Internal Code in progress.

            +--------------------------------------+
Percent     |             XX %                     |
complete    +--------------------------------------+

Elapsed time in minutes . . . . . .:    xxxx.x
```

9. You can see the 'Accept Console' display. If it appears, press F10 to accept the current console.

   On the 'Accept And Set New Console Type On This IPL' display, press Enter to continue.

10. On the 'Disk Configuration Attention Report' display, press F10 to accept any problems and to continue.

```
                          Disk Configuration Attention Report

Type option, press Enter
   5=Display Detailed Report


Press F10 to accept all the following problems and continue.
The system will attempt to correct them.


OPT Problem
_ New disk configuration



F3=Exit                   F10=Accept the problems and continue
```

11. On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)).

```
                          IPL or Install the System
                                                                  System: XXXX
Select one of the following:

   1. Perform an IPL
   2. Install the operating system
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the operating system
   5. Save Licensed Internal Code


Selection
3
```

12. Sign on to DST with service tools user QSEC0FR, with the password for the QSEC0FR service tools user ID.

```
                       Dedicated Service Tools (DST) Sign On

Type choices, press Enter.

Service tools user . . . . . . . . QSECOFR
Service tools password . . . . . . QSECOFR
```

    a. The password for the service tools profile QSECOFR is expired on a new system and must be changed at first use. If you see the 'Change Service Tools User Password' display, enter in all uppercase letters the current password QSECOFR and a new password, along with the verification password.

```
                       Change Service Tools User Password

Service tools user profile name . . . . : QSECOFR
Password last changed . . . . . . . . . : mm/dd/yy

Type choices, press Enter.

Current password  . . . . . . . . . . . : QSECOFR
New password  . . . . . . . . . . . . . : _______
New password (to verify). . . . . . . . : _______
```

    b. Select option 4 (Work with Disk Units).

    c. Select option 1 (Work with Disk Configuration).

    d. Select option 3 (Work with ASP Configuration).

    e. Select option 3 (Add Units to ASPs).

13. On the 'Add Units to ASPs' display, enter 3 to add units to existing ASPs, and add disk units to the system ASP.

```
                          Add Units to ASPs

Select one of the following:

   1. Create unencrypted ASPs
   2. Create encrypted ASPs
   3. Add units to existing ASPs
```

14. On the 'Specify New ASPs to Add Units to' display, enter the ASP number 1 to add disk units to the system ASP (1).

    a) Enter the ASP number on each selected unit.

```
                    Specify New ASPs to Add Units to

Specify the new ASP to add each unit to.

Specify  Serial                                    Resource
  ASP    Number          Type  Model  Capacity    Name
    1    21-6C597        4327   050      70564   DD007
    1    50-128840F      2107   A84      70564   DD004
    1    50-128940F      2107   A84      70564   DD005
         50-128A40F      2107   A85      35165   DD011
         50-128B40F      2107   A85      35165   DD003
         68-0C8BA12      6717   050       8589   DD008
         68-0C9D209      6717   050       8589   DD009
         68-606E0        6718   050      17548   DD006
         68-0CDAB10      6718   050      17548   DD010
         21-05348        4327   050      70564   DD012
         21-05322        4327   050      70564   DD013

F3=Exit     F5=Refresh     F11=Display disk configuration capacity
F12=Cancel
```

b) After you complete all units, press Enter.

c) If the list of units is correct, press Enter to start initializing the units.

15. On the 'Problem Report' display, press F10 to ignore the problems and continue.

```
                                Problem Report

   Note: Some action for the problems listed below might need to
   be taken. Please select a problem to display more detailed
   information about the problem and to see what possible
   action can be taken to correct the problem.


Type option, press Enter.
   5=Display Detailed Report


OPT Problem
   _ Unit possibly configured for Power PC AS

F3=Exit          F10=Ignore problems and continue           F12=Cancel
```

16. On the 'Confirm Add Units' display, press Enter to confirm the selected units.

```
                            Confirm Add Units

Add will take several minutes for each unit.  The system will
have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.
Press F9=Capacity Information to display the resulting capacity.
Press F10=Confirm Add and Balance data on units.
Press F12=Cancel to return and change your choice.

            Serial                   Resource
ASP  Unit   Number      Type Model   Name       Protection
  1                                             Unprotected
       1    68-0CDAB35  6718  050    DD001      Unprotected
       3    21-6C597    4327  050    DD007      Unprotected
       4    50-128840F  2107  A84    DD004      Unprotected
       5    50-128940F  2107  A84    DD005      Unprotected
       2    50-100240F  2107  A82    DD002      Unprotected
       6    50-128A40F  2107  A85    DD011      Unprotected

F9=Resulting Capacity              F10=Add and Balance
F11=Display Encryption Status      F12=Cancel
```

a. The 'Function Status' display shows the percentage of completion.

b. The

`Selected units have been added successfully`

message appears when the system completes the 'Add Units' process.

Note: If you require more than ASP 1 configured, you can return to step “13” on page 503 and create more ASPs by selecting 1 to create unencrypted ASPs or selecting 2 to create encrypted ASPs. You can create multiple user ASPs at the same time by specifying ASP numbers 2 through 32 on the selected disk units.

c. Press F12 to return to the 'Work with Disk Configuration' display.

d. If your system requires mirrored protection, continue to Step “16.e” on page 505. If your system does not require mirrored protection, press F3 until you see the 'Exit Dedicated Service Tools (DST)' display. Select option 1 to exit DST and press Enter and continue to Step “16.g” on page 505.

e. To start mirrored protection for your system, follow these instructions:

i) On the 'Work with Disk Configuration' display, select option 4 (Work with Mirrored Protection).

ii) On the 'Work with Mirrored Protection' display, select option 2 (Start Mirrored Protection).

iii) Select an ASP by placing a 1 next to it. Press Enter to start mirrored protection.

iv) On the 'Confirm Continuation' display, press Enter.

v) On the 'Confirm Start Mirrored Protection' display, press Enter.

vi) The 'Function Status' display shows the completion status of the 'Start Mirrored Protection' request.

vii) The

```
Start mirroring completed successfully
```

message appears on the 'Disk Configuration Information Report' display.

viii) Press Enter to continue.

f. To allow security system values to change during your recovery, follow these steps:

i) On the 'IPL or Install the System' display, select option 3 (Use Dedicated Service Tools (DST)). Presses Enter to continue.

ii) Sign on to DST with your DST username and password.

iii) Select option 13 (Work with system security).

iv) Type 1 in the **Allow system value security changes** field and press Enter.

v) Press F3 or F12 to get back to the 'IPL or Install the System' display.

g. If you use Operations Console, follow these instructions to switch working from local console to operations console.

i) On the 'IPL or Install the System' display, select option 3 (Use dedicated service tools (DST)). Press Enter to continue.

ii) Sign on to DST with service tools user QSECOFR, and the new password.

iii) On the 'Use Dedicated Service Tools (DST)' display, select option 5 (Work with DST environment). Press Enter to continue.

iv) On the 'Work with DST Environment' display, select option 2 (System Devices). Press Enter to continue.

v) On the 'Work with System Devices' display, select option 6 (Console Mode). Press Enter to continue.

vi) On the 'Select Console Type' display, select option 2 (Operations Console). Press Enter to continue.

vii) Press F3 or F12 to get back to the 'IPL or Install the System' display.

17. On the 'IPL or Install the System' display, select option 2 (Install the Operating System).

```
                        IPL or Install the System

Select one of the following:
                                                                        System: XXXX
   1. Perform an IPL
   2. Install the Operating System
   3. Use Dedicated Service Tools (DST)
   4. Perform automatic installation of the Operating System
   5. Save Licensed Internal Code

Selection
2
```

a. Make a selection on the 'Install Device Type Selection' display and press Enter.

   Note: Depending on the installation device type selected, other input might be needed.

b. On the 'Confirm Install of the Operating System' display, press Enter.

c. The 'Select a Language Group' display shows the primary language feature that is on your save media. Press Enter to accept this value.

```
                           Select a Language Group

Note:  The language feature shown is the language feature
installed on the system.

Type choice, press Enter.

Language feature . . . . . .  2924
```

d. The 'Confirm Language Feature Selection' display appears. Press Enter to continue.

18. On the 'Add All Disk Units to the System' display, select option 1 (Keep the current disk configuration).

```
                          Add All Disk Units to the System

Select one of the following:

1. Keep the current disk configuration
2. Perform disk configuration using DST
3. Add all units to the system auxiliary storage pool
4. Add all units to the system ASP and balance data

Selection
1
```

Note: This display does not appear if you selected all disk units that are known to the system on step “14” on page 503.

19. The 'Licensed Internal Code IPL in Progress' display shows the IPL progress.

```
                        Licensed Internal Code IPL in Progress

IPL:
   Type . . . . . . . . . . . . . . . :    Attended
   Start date and time  . . . . . . . :    xx/xx/xx  xx:xx:xx
   Previous system end  . . . . . . . :    Abnormal
   Current step / total . . . . . . . :       16      16
   Reference code detail  . . . . . . :    C6004065

IPL step                                   Time Elapsed   Time Remaining
  Commit Recovery                            xx:xx:xx        xx:xx:xx
  Data Base Initialization                   xx:xx:xx        xx:xx:xx
  Journal IPL Clean up                       xx:xx:xx        xx:xx:xx
  Commit Initialization                      xx:xx:xx        xx:xx:xx
 >Start the operating system                 xx:xx:xx        xx:xx:xx

Item:
  Current / Total . . . . . . . :

Sub Item:
  Identifier . . . . . . . . . . :    Searching tape for install media
  Current / Total  . . . . . . . :

Wait for next display or press F16 for DST main menu
```

20. On the 'Install the Operating System' display, select option 2 (Change install options).

    This option allows network attributes to be restored from your save media. If the date and time are not correct, type new values. Press Enter to continue.

```
                             Install the Operating System

Type options, press Enter.

Install
   option . . . . . 2      1=Take defaults (No other
                             options are displayed)
                           2=Change install options

Date
   Year . . . . . . xx     00-99
   Month. . . . . . xx     01-12
   Day. . . . . . . xx     01-31

Time
   Hour . . . . . . xx     00-23
   Minute . . . . . xx     00-59
   Second . . . . . xx     00-59
```

21. On the 'Specify Install Options' display, specify the following installation options.

```
                             Specify Install Options

  Type choices, press Enter.

     Restore option . . . . _1_    1=Restore programs and language objects
                                     from current media set
                                   2=Do not restore programs or
                                     language objects
                                   3=Restore only language objects
                                     from current media set
                                   4=Restore only language objects
                                    from a different media set using the
                                     current install device

     Job and output
       queues option . . . . _2_   1=Clear,  2=Keep

     Distribute IBM i on
       available disk units. _1_   1=Yes, 2=No
```

    If your ASPs had access path recovery times for ASPs that are not reconfigured, the access path recovery times cannot be restored. Use the **Edit Recovery Times for Access Paths (EDTRCYAP)** command to set the times after you reconfigure your ASP configuration.

22. On the 'Specify Restore Options' display, specify 1 for each of the following choices. Then press Enter.

```
                           Specify Restore Options

 Type choices, press Enter.

      Restore from the installation media:

         System information. . . 1    1=Restore,  2=Do not restore
         Edit descriptions . . . 1    1=Restore,  2=Do not restore
         Message reply list. . . 1    1=Restore,  2=Do not restore
         Job descriptions. . . . 1    1=Restore,  3=Keep customization
         Subsystem descriptions. 1    1=Restore,  3=Keep customization
```

23. The IBM i 'Installation Status' display shows the installation status of the IBM i installation profiles and libraries.

```
Message ID . . : CPI2070                                  IBM i Installation Status

        +-------------------------------------------------------------------------------+
Stage 1 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-------------------------------------------------------------------------------+
        0          20              40          60          80          100

Installation                                                                    Objects
Stage                                                          Completed       Restored

1  Creating needed profiles and libraries . . .:
2  Restoring programs to library QSYS . . . . .:
3  Restoring language objects to library QSYS .:
4  Updating program table . . . . . . . . . . .:
5  Installing database files  . . . . . . . . .:
6  Installing base directory objects  . . . . .:
```

24. The system installs the remaining IBM i objects.

```
Message ID . . : CPI2070                                  IBM i Installation Status

        +-------------------------------------------------------------------------------+
Stage 4 |xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx |
        +-------------------------------------------------------------------------------+
        0              20              40              60              80              100

Installation                                                Objects
Stage                                            Completed  Restored

   2  Restoring programs to library QSYS  . . . . :      x         09
   3  Restoring language objects to library QSYS. :      x         03
>> 4  Updating program table  . . . . . . . . . :      x
   5  Installing database files . . . . . . . . . :      x
   6  Installing base directory objects . . . . . :
   7  Completing installation . . . . . . . . . . :
```

25. On the 'Sign On' display, logon with user profile QSECOFR.

    You do not need to enter a password now.

26. On the 'IPL Options' display, enter the correct values for the system date, time, and time zone. Set the following options to Y.

    - Start system to restricted state
    - Set major system options
    - Define or change the system at IPL

```
                                IPL Options

Type choices, press Enter.

   System date  . . . . . . . . . . . . . . . xx/xx/xx   MM / DD / YY
   System time  . . . . . . . . . . . . . . . xx:xx:xx   HH : MM : SS
   System time zone . . . . . . . . . . . .   Q0000UTC   F4 for list
   Clear job queues . . . . . . . . . . . . . N          Y=Yes, N=No
   Clear output queues  . . . . . . . . . . . N          Y=Yes, N=No
   Clear incomplete job logs  . . . . . . . . N          Y=Yes, N=No
   Start print writers  . . . . . . . . . . . N          Y=Yes, N=No
   Start system to restricted state . . . . . Y          Y=Yes, N=No

   Set major system options . . . . . . . . . Y          Y=Yes, N=No
   Define or change system at IPL . . . . . . Y          Y=Yes, N=No
```

a. On the 'Set Major System Options' display, select **Y** to enable automatic configuration.

   Note: If you are using these instructions during a data migration, select **N** so that automatic configuration is not enabled.

```
                          Set Major System Options

Type choices, press Enter.

Enable automatic configuration . . . . Y         Y=Yes, N=No
Device configuration naming  . . . . . *NORMAL   *NORMAL, *S36
                                                 *DEVADR
Default special environment  . . . . . *NONE     *NONE, *S36
```

27. The 'Define or Change the System at IPL' display appears.
    a) Select option 3 (System Value Commands).
    b) On the 'System Value Commands' display, select option 3 (Work with System Values).
    c) On the 'Work with System Values' display, select the system value that you plan to change by placing a 2 next to it.

       Press Enter only after you select all the values.

       Update the following system values. Write down the existing values so you can update them after the recovery, if necessary.

       - QALWOBJRST to *ALL
       - QFRCCVNRST to 0
       - QIPLTYPE to 2
       - QINACTITV to *NONE
       - QJOBMSGQFL to *PRTWRAP
       - QJOBMSGQMX to a minimum of 30 with a suggestion of 64
       - QLMTDEVSSN to 0
       - QLMTSECOFR to 0
       - QLOGOUTPUT to *JOBLOGSVR
       - QMAXSIGN to *NOMAX
       - QPFRADJ to 2
       - QPWDEXPITV to *NOMAX
       - QSCANFSCTL add *NOPOSTRST
       - QVFYOBJRST to 1

    d) After the system changes the system values, press F3 twice to return to the 'Define or Change the System at IPL' display.
    e) On the 'Define or Change the System at IPL' display, press F3 to exit, and continue the IPL.
28. On the 'Change Password' display, type QSEC0FR as the current password. Enter a new password.

Reenter the password to verify and press Enter. (New password cannot be QSECOFR.)

If you are doing a data migration and you set Enable automatic configuration to **NO**, use the following steps to create a tape device to do the restore. If you set Enable automatic configuration to **YES**, skip to step "30" on page 510.

29. To configure tape units that are 3490 Model E or F, or that are not models 34xx, use the following instructions:
    a) Use the Work with Hardware Resource (WRKHDWRSC) command to determine the tape controller name: `WRKHDWRSC TYPE(*STG)`
    b) Locate the tape controller on the 'Work with Hardware Resources' display.
    c) Type a 9 (Work with resource) next to the tape controller name and press the Enter key.

       Note: If the tape controller is not listed on the display, you need to select other resources, such as disk storage controllers. For some system models, tape units are now attached through combined-function input/output processors (IOPs). Look through the resources until you find the tape unit you want.
    d) Locate the resource name for the tape unit, such as TAP01.
    e) Enter a 5 (Work with Configuration Descriptions) in the **Opt** column next to the tape resource name and press the Enter key.

       You are shown the 'Work with Configuration Descriptions' display.
    f) Type a 1 (Create) in the **Opt** field and the name of the tape device description, such as TAP01, in the **Description** field.

       Press the Enter key. You are shown the 'Create Device Description (Tape)' display.
    g) Change any values that you want to change, then press the Enter key (twice) to create the device description.

       You are shown the 'Work with Configuration Descriptions' display again. The device that you created is on the display.
    h) Type an 8 (Work with configuration status) in front of the new device description.

       You are shown the 'Work with Configuration Status' display.
    i) Type a 1 (Vary on or Make available) in front of the new device.

       If the status does not change to Varied On or Available, wait a few minutes. Then press F5 (Refresh). If the status still does not change to Varied On or Available, follow normal problem analysis for the device.
    j) Press F3 until you return to the main menu.
30. On the IBM i 'Main Menu' display, type the command, `WRKRPYLE`, and check to see whether CPA3709 is there.

    If it is not, determine an available sequence number and then press F6 to add `MSGID(CPA3709) RPY(G)` by using the available sequence number. Press F5 to Refresh and verify that you added CPA3709.

    a. Type the command `CHGJOB INQMSGRPY(*SYSRPYL)` to update the current job to use the system reply list for inquiry messages.
31. On the IBM i 'Main Menu' display, type `GO RESTORE` to access the 'Restore' display.
    a) On the 'Restore' display, select option 21 (Restore System and User Data).
    b) Press Enter to continue.
32. On the 'Specify Command Defaults' display, enter the name of the tape drive that you are using for the restore.
    a) Set **Prompt for command** to `N`.
    b) Set **Message queue delivery** to `*NOTIFY`.
    c) Set **Restore to different system** to `Y`.
    d) Set **Spooled file data** to `*NEW`.

e) Press Enter to continue ending the subsystems.

   The restore process starts to run unattended. The restore process stops only if the system requires a tape mount for the restore process to continue. The restore process runs to completion.

33. If a controller or device no longer has a valid resource on the system, use the **Work with Hardware Product (WRKHDWPRD)** command to correct the hardware resource names.
34. If you receive message CPF3204 `Cannot find object needed for file xxxx in xxx`, during the restore of your libraries, the necessary based-on physical files were not yet restored when the logical files were attempting to restore.

    You will need to do the restore of logical files after the physical files are restored by using one of the following options:

    - Run this command to restore only the logical files that were not restored previously:

      ```
      RSTLIB SAVLIB(*NONSYS) DEV(media-device-name)
      OPTION(*NEW) MBROPT(*ALL) ALWOBJDIF(*COMPATIBLE)
      ```

    - Run this command for each library that has logical files that need to be restored:

      ```
      RSTLIB SAVLIB(library-name) DEV(media-device-name)
      OPTION(*NEW)  MBROPT(*ALL) ALWOBJDIF(*COMPATIBLE)
      ```

    - Run this command for each logical file that needs to be restored:

      ```
      RSTOBJ OBJ(logical-file-name) SAVLIB(library-name)
      DEV(media-device-name) OBJTYPE(*FILE) MBROPT(*ALL) ALWOBJDIF(*COMPATIBLE)
      ```

35. Update program temporary fix (PTF) information for all PTF save files in the QGPL library by typing: `UPDPTFINF`.
36. Reapply any PTFs that you applied since the last time that you saved your system data.
    a) Locate the most recent cumulative PTF tape or CD.
    b) From a command line, enter `GO  PTF` to access the **PTF** menu.
    c) Select option 8 (Install program temporary fix package), on the **PTF** menu.

       This installs all the PTFs in the cumulative PTF package for the licensed programs that are installed on your system. Refer to the 'IBM i PTF Shipping Information Letter' for special instructions.

       Notes:

       - If you want to restore individual PTFs, see the Basic system operations topic for more information about applying individual PTFs.
       - If you do not have the PTFs that you need, order them and apply them later.

37. If necessary, change the system values back to what they were in step “27.c” on page 509 by using the Work with System Value (WRKSYSVAL) command.
38. If you do not know the password for the restored QSECOFR profile, change the password before you sign off.

    Type the following command and press Enter:

    ```
    CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)
    ```

39. To place scheduled jobs on hold, type `WRKJOBSCDE` and select option 3 to hold any scheduled jobs.

    You need to release the jobs in step “44” on page 512 after you complete the restore.
40. Type the command `SIGNOFF *LIST` or `DSPJOBLOG * *PRINT`.

    Check the job log to verify that the job restored all objects. To verify whether the job restored all objects, spool the job log for printing along with any remaining spooled job output. Check for any error messages. Correct the errors and restore those objects from media.

41. IPL the system.
    a) On the CPU control panel, select the **NORMAL** setting.
    b) Type the following command and press Enter:

```
PWRDWNSYS OPTION(*IMMED) RESTART(*YES
*FULL) IPLSRC(B)
```

42. If IBM Content Manager OnDemand for i installed on your system, restart the journaling for Content Manager OnDemand by typing the following commands.

```
CALL QRDARS/QRLCSTRJ PARM('RLC')
CALL QRDARS/QRLCSTRJ PARM('RLR')
```

43. If you have the Cryptographic Device Manager licensed program, 5733-CY2 or 5733-CY3 installed, run command `CALL QCCADEV/QCCAELOAD`.
44. Release the jobs that you previously held in Step “39” on page 511.

You completed the recovery.

Related information

System values

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 grant 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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

© Copyright IBM Corp. 1997, 2025

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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 illustrate 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

This Recovering your system publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be

trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

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

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

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

Java and all Java-based trademarks and logos are trademarks of Oracle, Inc. in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Bibliography

Publications that provide additional information about topics that are described or referred to in this document.

- *Are You Saving the Right Stuff?*, G325-6153.

  The *Are You Saving the Right Stuff?* poster is a visual reminder of the methods available for saving and restoring information on your system. It also highlights the software, hardware, people, and service offerings that are available to help you with backup, recovery, and availability.

- Storage Solutions.

  The Storage Solutions topics helps you make choices about which storage technologies are right for your company. Examples of storage technologies include disk, tape, optical storage, and virtual storage.

- Backup, Recovery, and Media Services (BRMS).

  Backup, Recovery, and Media Services (BRMS) helps you implement a disciplined approach to managing your backups and provides you with an orderly way to retrieve lost or damaged data.

- Control language.

  The Control language (CL) topic allows system programmers and system administrators to write programs that use IBM i commands and other IBM-supplied commands.

- Communications Management .

  Communications Management provides general configuration information about network interfaces, network servers, line, controller, device, mode, NetBIOS, and class-of-service descriptions. It also provides information about configuration lists, and connection lists.

- Hierarchical Storage Management .

  Hierarchical Storage Management provides an overview of hierarchical storage management principles. It also describes considerations for planning for the use of dynamic retrieval that is supplied with BRMS. The manual also provides information on the implementation of hierarchical storage management.

- ILE Concepts.

  ILE Concepts provide concepts and terminology that are related to the Integrated Language Environment (ILE) architecture of the IBM i licensed program. Subjects that are covered include module creation, binding, message handling, the running and debugging of programs, and exception handling.

- Local Device Configuration .

  Local Device Configuration provides the system operator or system administrator with information on how to do an initial configuration and how to change the configuration. It also contains conceptual information about device configuration.

- Performance .

  Performance provides the programmer with the information that is needed to collect data about the system, job, or program performance. Other performance information includes tips for printing and analyzing performance data to identify and correct inefficiencies that might exist. Information about the manager and agent feature is included.

- Security Reference.

© Copyright IBM Corp. 1997, 2025

The Security Reference topic provides the programmer (or someone who security officer responsibilities) with information about system security concepts, planning for security, and setting up security on the system. Security Reference does not describe security for specific licensed programs, languages, and utilities.

- Simple Network Management Protocol (SNMP) Support .

  Simple Network Management Protocol (SNMP) Support provides the system operator, programmer, or system administrator with information for configuring the IBM i environment to use simple network management protocol (SNMP).
- Installing, upgrading, or deleting IBM i and related software.

  The Installing, upgrading, or deleting IBM i and related software topic provides the system operator or system administrator with step-by-step procedures for installing licensed programs, program temporary fixes (PTFs), and secondary languages from IBM.

  The Installing, upgrading, or deleting IBM i and related software topic is also for users who already have a IBM i model with an installed release and want to upgrade to a new release.
- Basic system operations.

  The Basic system operations topic provides information about working with job output, working with printer output, and working with support functions. It also provides information about message handling, device communications, and cleaning up your system.
- Networking.

  The Networking topic provides information for configuring TCP/IP support and applications. The applications that are included are Telnet, File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), line printer requester (LPR), and line printer daemon (LPD).
- Work management.

  Work management is an important building block within the IBM i operating system.

Index

Special Characters

()
  recovery procedures
    complete data loss, not overflowed 73
    complete data loss, overflowed 77, 84
    no data loss 70
    some data loss 81, 104
(ASP)
  calculating space requirements 406
  changing threshold 399, 400
  deleting 377
  moving disk unit 402
  moving folder 409
  moving library 408
  recovery procedures
    complete data loss, not overflowed 82
  removing disk unit 403
  transferring journals 409
  transferring objects 408
  user
    calculating space requirements 406
    changing threshold 399, 400
    deleting 377
    journal receivers 412
    moving disk unit 402
    removing disk unit 403
    transferring objects 408
(auxiliary storage pool)
  determining overflowed status 183
  overflowed
    deleting objects 186
  recovery procedure
    loss, overflowed 66
(CHGASPA) Change ASP Attribute command 440
(DST)
  starting 382
  stopping 383
(non-load source disk unit failure)
  recovery procedures
    some data loss 71
(SST) 384
*ALLOBJ (all-object) special authority
  restoring 215

Numerics

34xx tape units
  creating tape configuration 335, 495, 510

A

A900 2000 SRC (system reference code)
  recovery 157
abbreviated install
  definition 142
abend 53
abnormal end
  definition 159
  restarting system 159
abnormal IPL (initial program load) 159
access path
  editing rebuild during IPL 163
  recovery times
    restoring 149
  restoring 246
access path recovery time
  recovering 209
accessing
  (DST) 382
action
  mirrored protection recovery 287, 292
active 389
active disk unit status 389
adding disk units 392, 394
addressability 172
adopted authority object
  allowing restore operation 34
after 172
all-object (*ALLOBJ) special authority
  restoring 215
allow object differences (ALWOBJDIF) parameter
  authorization lists 216
  effect 23
  purpose 23
allow object restore operation (QALWOBJRST) system value 34
allow user domain objects (QALWUSRDMN) system value 29
allowing
  restore
    adopted authority objects 34
    sensitive objects 34
    system-state programs 34
alternate installation device
  how to verify and select, during a manual installation 364
  IPL-type D from devices without IOP 361
alternate IPL device 123
alternate restart device 123
ALWOBJDIF (allow object difference) parameter
  database file 241
  member 241
ALWOBJDIF (allow object differences) parameter
  authorization lists 216
  effect 23
  purpose 23
APIs
  QsrRestore 265
Apply Journaled Changes (APYJRNCHG) command
  broken receiver chain 283
  unbroken receiver chain 282
applying
  journaled changes
    broken receiver chain 283

applying *(continued)*
journaled changes *(continued)*
determining whether to 280
QAOSDIAJRN journal 284
unbroken receiver chain 282
APYJRNCHG (Apply Journaled Changes) command
broken receiver chain 283
unbroken receiver chain 282
ASP
independent 133
ASP (auxiliary storage pool)
overflowed
moving journal receivers 413
resetting journal 413
ASP group 53
ASPs 415
assigning to auxiliary storage pool (ASP) 392, 394
audit (QAUDJRN) journal
creating during restore 40
authority
private
restoring 218
restoring 217
authority holder
restoring 214
authorization list
restoring 214
restoring link 216
automatic configuration
enabling
during recovery 153
automatic IPL after power restored (QPWRRSTIPL) system value 159
auxiliary storage (ASP)
high percentage used 167
auxiliary storage configuration
checklist
adding disk to IOA 372
adding disk units without 371
adding input/output adapter 373
addinginput/output adapter 373
deleting (ASP) 377
moving disk units 375, 376
new system 370
removing disk units 378–380
auxiliary storage pool 53
auxiliary storage pool (ASP)
encrypted
restoring 309
overflowed
moving journal receivers 413
resetting journal 413

B

backup 116
BRMS synchronization 353
broken receiver chain
applying journaled changes 283
bus failure 294
busy 388
busy disk unit status 388

C

calculating
disk space 406
calendar server (QCALSRV) subsystem
ending 28
Change ASP Attribute (CHGASPA) command 440
Change Service Tools User Password display 134
changed object
restoring
by library 278
by object 278
cumulative 278
directories 279
not cumulative 278
changing
auxiliary storage pool (ASP)
storage threshold 399
storage threshold
auxiliary storage pool (ASP) 399
system auxiliary storage pool (ASP) 400
system auxiliary storage pool (ASP)
storage threshold 400
changing system
at IPL
during recovery 153
checklist
disk configuration
adding disk to IOA 372
adding disk units without 371
adding input/output adapter 373
deleting (ASP) 377
moving disk units 375, 376
new system 370
removing disk units 378–380
cleaning up
hardware configuration 228
clearing
job queue
during recovery 148
output queue
during recovery 148
command, CL
Apply Journaled Changes (APYJRNCHG)
broken receiver chain 283
unbroken receiver chain 282
APYJRNCHG (Apply Journaled Changes)
broken receiver chain 283
unbroken receiver chain 282
CHGASPA (Change ASP Attribute) 440
Display Database Relations (DSPDBR) 252
DSPDBR (Display Database Relations) 252
Edit Check Pending Constraint (EDTCPCST) 165
EDTCPCST (Edit Check Pending Constraint) 165
End Subsystem (ENDSBS)
QCALSRV (calendar server) subsystem 27
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27
using 27
ENDSBS (End Subsystem)
QCALSRV (calendar server) subsystem 27
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27

command, CL *(continued)*
ENDSBS (End Subsystem) *(continued)*
using 27
RCLDLO (Reclaim Document Library Object) 262
RCLSTG (Reclaim Storage)
duplicate names in QRCL 29
object ownership 29
procedure 32, 173
QALWUSRDMN (allow user domain objects) system value 29
recovering 173
user domain object 29
what system does 29
why to run 167
RCVJRNE (Receive Journal Entry)
writing output to save media 460
Receive Journal Entry (RCVJRNE)
writing output to save media 460
Reclaim Document Library Object (RCLDLO) 262
Reclaim Storage (RCLSTG)
duplicate names in QRCL 29
object ownership 29
procedure 32, 173
QALWUSRDMN (allow user domain objects) system value 29
recovering 173
user domain object 29
what system does 29
why to run 167
Remove Defer ID (RMVDFRID) 251
Rename Directory Entry (RNMDIRE)
restoring mail 265
Rename Document Library Object (RNMDLO)
restoring documents 265
Restore (RST)
changed objects 279
how to use 265
Restore (RST) command
restrictions 269
restrictions when restoring documents 271
Restore Authority (RSTAUT)
nonrestricted state system 219
Restore Configuration (RSTCFG) 227
Restore Deferred Object (RSTDFROBJ) 251
Restore Document Library Object (RSTDLO)
maximum number of DLOs 263
media error 42
output 261
overview 260
renaming document 264
restoring authority 264
restoring descriptive information 264
restoring ownership 264
Restore Library (RSTLIB)
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
Defer ID parameter 26
media error 41
multiple concurrent 233
OPTION parameter 232
overview 231
Restore Licensed Program (RSTLICPGM) 260
Restore Object (RSTOBJ)
command, CL *(continued)*
Restore Object (RSTOBJ) *(continued)*
Defer ID parameter 26
multiple concurrent 235
Restore User Profiles (RSTUSRPRF) 211
Retrieve Journal Entry (RTVJRNE)
using in program 459
RMVDFRID (Remove Defer ID) 26, 251
RNMDIRE (Rename Directory Entry)
restoring mail 265
RNMDLO (Rename Document Library Object)
restoring documents 265
RST (Restore)
changed objects 279
how to use 265
RST (Restore) command
restrictions 269
restrictions when restoring documents 271
RSTAUT (Restore Authority)
nonrestricted state system 219
RSTCFG (Restore Configuration) 227
RSTDFROBJ (Restore Deferred Object) 251
RSTDLO (Restore Document Library Object)
maximum number of DLOs 263
media error 42
output 261
overview 260
renaming document 264
restoring authority 264
restoring descriptive information 264
restoring ownership 264
RSTLIB (Restore Library)
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
Defer ID parameter 26
media error 41
multiple concurrent 233
OPTION parameter 232
overview 231
RSTLICPGM (Restore Licensed Program) 260
RSTOBJ (Restore Deferred Objects) 26
RSTOBJ (Restore Object)
Defer ID parameter 26
multiple concurrent 235
RSTUSRPRF (Restore User Profiles) 211
RTVJRNE (Retrieve Journal Entry)
using in program 459
Save Library (SAVLIB)
determining what command was used 304
SAVLIB (Save Library)
determining what command was used 304
common terminology 53
complete data loss, no 61
complete data loss, not overflowed 62
compression recovery policy 440
concurrent add
disk unit 369
configuration
cleaning up 228
errors for mirrored protection 433
restoring
problems with SRM information 228
rules for mirrored protection 431

configuration list
recovering 209
configuration object
restoring to a different system 228
configuring
disk
adding disk to IOA 372
adding disk units without 371
adding input/output adapter 373
deleting (ASP) 377
moving disk units 375, 376
new system 370
removing disk units 378–380
Confirm Delete ASP Data display 197
Confirm Move of Unit display 403
Confirmation Continuation display 402
considerations
mirrored protection 293
console
problem
during recovery 157
console type
changing during restore 230
constraint
pending
editing during IPL 164
conversion
program 257
CPA3388 message 261
CPF0975 message
during recovery 157
CPF7088 message 40
CPF8113 message 167
CPI0953 message 182
CPI0954 message 182
creating
tape configuration
for 34xx tape units 335, 495, 510
for non-34xx tape units 157
creating document library objects (DLOs) 411
creating objects 410, 415
creation date
database file
restoring 241
current release-to-previous release support
installing previous release compiler 313
using TGTRLS (target release) parameter 313

D

damaged
database file 29
document
restoring 263
folder
restoring into 263
IBM-supplied user profile 166
job description 166
job queue 166
journal 167
journal receiver 168
journaled object 169
object
damaged *(continued)*
object *(continued)*
without library 29, 167
operating system object 166
output queue 167
QAOSS (text index) database files 167
damaged object
recovery 165
data
restoring save file 258
data area
journaled
restoring 236
restoring
objects being journaled 236
data queue
journaled
restoring 236
restoring
objects being journaled 236
database
restore order 26, 251
restoring
referential constraints 249
trigger program 250
database file
constraint
editing during IPL 164
damaged 29, 167
deleting 252
journaled
damaged 169
not synchronized 169
member
damaged 167
multiple members
example 240
QAOSS (text index)
damaged 167
renaming
during restore 242
restoring
access paths 246
ALWOBJDIF (allow object difference) parameter 241
considerations 240
creation date 241
different member set 243
files being journaled 236
MAXMBRS (maximum members) parameter 243
MBROPT (Member option) parameter 243
member locking 241
members, list of 244
shared formats 248
database relations
displaying 252
dedicated service tools (DST)
options 381
default
owner (QDFTOWN) user profile
restoring objects 216
deferred restore
dependent objects
logical files 26

deferred restore *(continued)*
dependent objects *(continued)*
physical files 26
physical files 251
define or change system at IPL
during recovery 153
defining system
at IPL
during recovery 153
definition 53–55, 182, 388, 392, 394
deleting
(ASP) 377
journal 253
journal receiver 256
physical file 252
dependent file
restoring 249
device
recovering after restore operation 229
starting
during recovery 152
device configuration
restoring 227
device error
recovery actions 287
unrecoverable 293
device parity protection 419, 421
device parity status 430
different release 258
directory
restoring
changed objects 279
restoring objects 265
directory entry
renaming
restoring mail 265
disaster recovery
sample plan 473
disk
adding while system active 369
concurrent add 369
recovering configuration 136
understanding configuration status 385
disk compression
considerations 437
disk unit capacity 438
disk unit full considerations 439
error codes
SRC 6xxx 7051 450
SRC 6xxx 7052 451
restrictions 437
SRC A6xx 0277 441
starting 443
stopping 445
disk configuration
checklist
adding disk to IOA 372
adding disk units without 371
adding input/output adapter 373
addinginput/output adapter 373
deleting (ASP) 377
moving disk units 375, 376
new system 370
disk configuration *(continued)*
checklist *(continued)*
removing disk units 378–380
recovering 136
DISK CONFIGURATION ATTENTION REPORT display 134
disk configuration status
displaying 385
printing 385
understanding 385
disk failure
pump 57
recovery strategy 56
recovery with mirrored protection 87
disk pool 53
disk protection
configuration sequences
adding a storage controller 447
disk space
calculating 406
disk unit
adding to system 371
failure
load source unit before IPL 294
formatting 391
hot spare 419
initializing 391
moving 402
nonconfigured status
reasons 145
removing
from ASP 403
disk unit failure
recovery strategy 56
disk unit full
system response 440
disk unit in device parity protection 427, 429
disk unit number 388
disk unit status 388, 389
display 160
Display Access Path Status display 164
Display Constraint Status display 165
Display Database Relations (DSPDBR) command 252
Display Disk Configuration Capacity display 183, 184
displaying
database relations 252
displaying device parity status 430
displaying objects 406
displays
LIC error 465
licensed internal code installation error 465
distribution media
restoring Licensed Internal Code 123
restoring licensed program 141
distribution services (QAOSDIAJRN) journal
applying journaled changes 284
DLO (document library object)
maximum number on RSTDLO command 263
reclaiming 262
renaming
restoring documents 265
restoring
descriptive information 264
media error 42
overview 260

DLO (document library object) *(continued)*
restoring *(continued)*
renaming document 264
using RST (Restore) command 271
restoring authority 264
restoring ownership 264
document
restoring
damaged 263
overview 260
document library 190
document library object (DLO)
maximum number on RSTDLO command 263
reclaiming 262
renaming
restoring documents 265
restoring
descriptive information 264
media error 42
overview 260
renaming document 264
using RST (Restore) command 271
restoring authority 264
restoring ownership 264
document library objects (DLOs) 411
DPY/Active 389
DPY/Active disk unit status 389
DPY/Failed 389
DPY/Failed disk unit status 389
DPY/Rebuilding 389
DPY/Rebuilding disk unit status 389
DPY/Resyncing 389
DPY/Resyncing disk unit status 389
DPY/Unknown 389
DPY/Unknown disk unit status 389
DPY/Unprotected 389
DPY/Unprotected disk unit status 389
DSPDBR (Display Database Relations) command 252
DST ()
starting 382
stopping 383
DST (dedicated service tools)
options 381
DST Sign On display 134

E

Edit Check Pending Constraint (EDTCPCST) command 165
Edit Check Pending Constraints display 156, 164
edit description
recovering 209
restoring 150
Edit Rebuild of Access Paths display 155, 163
EDTCPCST (Edit Check Pending Constraint) command 165
enabling
automatic configuration
during recovery 153
encrypted auxiliary storage pool
restoring 309
encrypted tape
restoring 308
encryption
saving keys 308
using BRMS 307
encryption *(continued)*
using encrypting tape 308
End Subsystem (ENDSBS) command
QCALSRV (calendar server) subsystem 27
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27
using 27
ending
(DST) 383
mirrored protection 434
subsystem
QCALSRV (calendar server) subsystem 27
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27
using 27
ENDSBS (End Subsystem) command
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27
using 27
entire system
restore operation
unattended 203
restoring 203
error
mirrored protection configuration 433
permanent read 293
restore operation
not recoverable 41
recoverable 41
SRM (system resource management) information 228
unrecoverable device 293
error displays
LIC 465
licensed internal code installation 465
error handling
mirrored protection
disk 293
error message 293
escape condition
handling with program 459
example
database file
multiple members 240
handling escape conditions 459
using Retrieve Journal Entry (RTVJRNE) command in program 459
excluding 429
excluding disk unit from device parity protection 429
excluding in device parity protection 429

F

failed unit 88
failure
active mirrored load source 295
bus 294
I/O processor 294
load source unit
before IPL 294
file
constraint
editing during IPL 164
database
shared formats 248

file *(continued)*
journaled
restoring 236
restoring
logical 244
folder
restoring
damaged 263
overview 260
procedure 263
transferring
different ASP 409
formatting disk units 391

H

hard link
restoring 266
hardware
configuration
cleaning up 228
hardware resource 388
hot spare 419
human error
recovery strategy 55

I

IBM-supplied user profile
damaged 166
ICF file
writing journal entries 462
including 427
including disk unit in device parity protection 427
independent 54
independent ASP
recovering disk configuration after complete system loss 133
independent ASP (auxiliary storage pool)
recovery procedures
complete data loss 91
no data loss 89
some data loss 90
initial program load (IPL)
after abnormal end 159
disk related failure of load source unit 294
editing check pending constraints 164
editing rebuild of access paths 163
options
during recovery 153
performing normal 43
restoring operating system 142
selecting options
restoring operating system 151
initializing disk units 391
Input/OutputAdapter
adding 373
install options
selecting
restoring operating system 147
Install the Operating System display 136
installation
abbreviated
definition 142
installation device
alternate 361
installation error displays
licensed internal code 465
integrated file system object
journaled
damaged 169
not synchronized 169
restoring 236
restoring
integrated file system objects being journaled 236
interpreting 388
IOA
adding disk 372
IPL (initial program load)
after abnormal end 159
disk related failure of load source unit 294
editing check pending constraints 164
editing rebuild of access paths 163
normal 43
options
during recovery 153
performing normal 43
restoring operating system 142
selecting options
restoring operating system 151
IPL Options display 139, 152, 162, 334, 494, 508

J

job description
damaged 166
job number
resetting counter
during recovery 148
job queue
clearing during recovery 148
damaged 166
journal
damaged 167
deleting 253
overflowed
resetting 413
recovering from QRCL library 175, 176
restoring 252
transferring into user ASP 409
journal entry
RTVJRNE (Retrieve Journal Entry) command 459
transmitting to another system 462
writing to ICF file 462
journal receiver
damaged 168
deleting 256
directory
correcting 255
moving
from overflowed ASP 413
placing in 412
recovering from QRCL library 175, 176
restoring 252
journaled changes

journaled changes *(continued)*
applying
broken receiver chain 283
determining whether to 280
unbroken receiver chain 282
journaled file
restoring 236
journaled integrated file system object
restoring 236
journaled libraries 236
journaled object
damaged 169
not synchronized 169
journaling
applying changes 280

L

library
journaled
restoring 236
locking during restore procedure 38
moving
different ASP 408
replaying changes
Apply Journaled Changes (APYJRNCHG) 236
restoring
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
media error 41
OPTION parameter 232
overview 231
restoring private authorities 408
saving
determining what command was used 304
library list
changing
during recovery 153
LIC
error displays 465
installation error displays 465
licensed internal code
installation error displays 465
Licensed Internal Code
restoring
preparation 122
starting 122
steps 123
using distribution media 123
using function code
SRC (system reference) codes 125
Licensed Internal Code IPL in Progress display 135, 136
licensed program
restoring
choosing procedure 142
manual IPL 142
overview 141
preparation 141
reasons 141
selecting install options 147
steps 143
using distribution media 141
licensed programs
licensed programs *(continued)*
restoring 260
link
restoring 266
load source
failure
active mirrored 295
unknown status 297
locking
database members
during restore 241
restore processing 38
logical file
restoring 244
logical partitions
restoring 231
loss, not overflowed 62
lost object 29, 167

M

main storage dump 159
major system options
setting
during recovery 153
setting during recovery 153
manual IPL (initial program load)
restoring operating system 142
materialized query table 251
maximum members (MAXMBRS) parameter
exceeding during restore 243
MAXMBRS (maximum members) parameter
exceeding during restore 243
MBROPT (Member option) parameter 243
media error
during RSTDLO procedure 42
during RSTLIB procedure 41
member
damaged 167
locking
during restore 241
renaming
during restore 242
restoring
ALWOBJDIF (allow object difference) parameter 241
creation date 241
Member option (MBROPT) parameter 243
menu 153
message
CPA3388 261
CPF7088 40
CPF8113 167
CPI0953 182
CPI0954 182
message reply list
restoring 150
microcode
restoring 122
mirrored protection
active load source failure 295
configuration errors 433
configuration rules 431
device error

mirrored protection *(continued)*
device error *(continued)*
recovery actions 287
disk error handling 293
ending 434
I/O processor or bus failure 294
missing disk units 294
nonconfigured unit
using for replacement 291
permanent read error
recovery actions 287
recovery actions
errors and failures 287
performed by the service representative 292
recovery steps 87
replacing unit 289
resuming 293
starting 431
stopping 434
unknown load source status 297
using spare nonconfigured unit 291
mirrored unit
replacing 289
resuming 289
suspending 288
mirroring
device error
recovery actions 287
permanent read error
recovery actions 287
moving
disk unit 402
disk units 375, 376
folder
different ASP 409
journal receiver
overflowed ASP (auxiliary storage pool) 413
library
different ASP 408
object
different ASP 416
user profile
different system 215

N

network
database 249
restoring 249
network attribute
recovering 209
resetting when restoring to a different system 154
new system
configuring disk 370
no data loss 59
non-34xx tape units
creating tape configuration 157
non-load source unit
recovery procedure
complete data loss, no 72
nonconfigured disk unit
reasons 145
nonconfigured status 392
nonconfigured unit
nonconfigured unit *(continued)*
mirrored protection 291
nonload source unit
recovery procedure
complete data loss, not overflowed 73
complete data loss, overflowed 77
no data loss 70
normal initial program load (IPL) 43
not operational 388
not operational disk unit status 388
not ready 388
not ready disk unit status 388
not synchronized
journaled file 169

O

object
damaged 170
lost owner 29
ownership
restoring 216
primary group
restoring 216
restore sequence 25
restoring
RSTOBJ (Restore Object) command 235
saving
previous release system 313
transferring
between ASPs 408
different ASP 416
without library 29, 167
object in directory
restoring 265
object ownership
ALWOBJDIF (allow object differences) parameter 216
ObjectConnect 16
objects 410, 415
operating system
damaged object 166
restoring
choosing procedure 142
encryption 143
manual IPL 142
overview 141
preparation 141
reasons 141
selecting install options 147
steps 143
using distribution media 141
operational 388
operational disk unit status 388
order
restoring objects 25
output
RSTDLO (Restore Document Library Object) command 261
output queue
clearing during recovery 148
damaged 167
overflowed
deleting objects 186
determining status 183

overflowed *(continued)*
displaying 183, 184
overflowed object
deleting 186
overflowed status 182
overview 16
ownership
ALWOBJDIF (allow object differences) parameter 216
restoring 216

P

parallel restore operations 44
parent file
restoring 249
pending constraint
editing during IPL 164
performance degraded 389
performance degraded disk unit status 389
permanent error 287
permanent read error
mirrored protection
recovery actions 287, 293
physical file
deleting 252
physical files
deferred restore 251
placing in 417
placing journal receivers 417
power failure
recovery strategy 55
previous release system
saving objects 313
previous release-to-current release support
considerations when restoring configuration objects 322
restoring data to current release system 322
saving data on previous release system 322
primary 54
primary ASP 54
primary group
ALWOBJDIF (allow object differences) parameter 217
restoring 216
printer writer
starting
during recovery 152
private authority
restoring 217, 218
profile
QDFTOWN (default owner)
restoring objects 216
program
conversion 257
recreation 257
restoring 256
translation 257
validation value 256
program failure
recovery strategy 55
program temporary fix (PTF)
restoring 273
programming example
handling escape conditions 459
using Retrieve Journal Entry (RTVJRNE) command in program 459
PTF (program temporary fix)
restoring 273
pump (disk pump) 57
PVTAUT parameter 218, 408

Q

QALWOBJRST (allow object restore operation) system value 34
QALWUSRDMN (allow user domain objects) system value 29
QAOSDIAJRN (distribution services) journal
applying journaled changes 284
QAPZ files
attempting to restore 39
QASRRSTO output file
field descriptions 47
format 44
QAUDJRN (audit) journal
creating during restore 40
QCALSRV (calendar server) subsystem
ending 28
QDFTOWN (default owner) user profile
assigning ownership during reclaim storage procedure 29
restoring objects 216
QPWRRSTIPL (automatic IPL after power restored) system value 159
QRCL (recovery) library
duplicate names 29
journal 176
journal receiver 176
using for recovery 175
QRYDOCLIB (Query Document Library) 190
QRYDOCLIB (Query Document Library) command 190
QsrRestore API 265
QSYSMSG message queue
error messages 293
QSYSOPR message queue
error messages 293
QSYSWRK (subsystem monitor) subsystem
ending 28
Query Document Library (QRYDOCLIB) 190
Query Document Library (QRYDOCLIB) command 190
querying 190
QVFYOBJRST (verify object on restore) system value 34

R

RCLDLO (Reclaim Document Library Object) command 262
RCLSTG (Reclaim Storage) command
duplicate names in QRCL 29
object ownership 29
procedure 32, 173
QALWUSRDMN (allow user domain objects) system value 29
recovering 173
user domain object 29
what system does 29
why to run 167
RCVJRNE (Receive Journal Entry) command
writing output to save media 460
read error 293
read/write protected 388

read/write protected disk unit status 388
Receive Journal Entry (RCVJRNE) command
    writing output to save media 460
receiver
    restoring 252
receiver chain
    broken
        applying journaled changes 283
    definition 255
    unbroken
        applying journaled changes 282
receiver directory
    correcting 255
Reclaim Document Library Object (RCLDLO) command 262
Reclaim Storage (RCLSTG) command
    duplicate names in QRCL 29
    object ownership 29
    procedure 32, 173
    QALWUSRDMN (allow user domain objects) system value 29
    recovering 173
    user domain object 29
    what system does 29
    why to run 167
reclaiming
    document library object (DLO) 262
    storage
        duplicate names in QRCL 29
        procedure 32, 173
        QALWUSRDMN (allow user domain objects) system value 29
        recovering 173
        user domain object 29
        what the system does 29
        why to run 167
recoverable error
    restore operation 41
recovering
    access path recovery times 149, 209
    configuration lists 209
    devices that will not vary on 229
    disk configuration 136
    edit descriptions 150, 209
    Licensed Internal Code
        preparation 122
        starting 122
        steps 123
        using distribution media 123
    licensed program
        choosing procedure 142
        manual IPL 142
        overview 141
        preparation 141
        reasons 141
        selecting install options 147
        steps 143
        using distribution media 141
    message reply list 150
    network attributes 209
    object 285
    reply list entries 209
    restoring
        preparation 122
        starting 122
recovering *(continued)*
    restoring *(continued)*
        steps 123
        using distribution media 123
    service attributes 149
    system information 149, 209
    system management objects 149
    system reply list 150
    system values 149, 209
    System/36 environment 230
    tape controller 229
    unsuccessful restore operation 41
    user information
        choosing procedure 105
        using Restore menu option 21 110
        using Restore menu options 22 and 23 113
    using changed objects 277
    using journaling 277
    vital product data information 56, 131
recovery
    considerations for mirrored protection 293
    damaged objects 165
    disaster
        sample plan 473
    mirrored protection 293
    unreadable sectors 165
recovery (QRCL) library
    duplicate names 29
    journal 176
    journal receiver 176
    using for recovery 175
recovery actions
    mirrored protection 287, 292
    performed by the service representative
        mirrored protection 292
recovery checklist
    complete system loss
        including independent ASP 97
    independent ASP
        complete data loss 91
        no data loss 89
        some data loss 90
    mirrored protection 87
    non-
        complete data loss, no 72
        complete data loss, not overflowed 73
        complete data loss, overflowed 77
        no data loss 70
        some data loss 71
    user information
        using commands 107
        using Restore menu option 21 110
        using Restore menu options 22 and 23 113
recovery from unreadable sectors during disk failure 165
recovery procedure
    complete system loss
        including independent ASP 97
    independent ASP
        complete data loss 91
        no data loss 89
        some data loss 90
    mirrored protection 87
    non-
        complete data loss, no 72

recovery procedure *(continued)*
non- *(continued)*
complete data loss, not overflowed 73
complete data loss, overflowed 77
no data loss 70
some data loss 71
user information
using commands 107
using Restore menu option 21 110
using Restore menu options 22 and 23 113
recovery procedures 88
recovery steps
complete system loss
including independent ASP 97
independent ASP
complete data loss 91
no data loss 89
some data loss 90
mirrored protection 87
non-
complete data loss, no 72
complete data loss, not overflowed 73
complete data loss, overflowed 77
no data loss 70
some data loss 71
user information
using commands 107
using Restore menu option 21 110
using Restore menu options 22 and 23 113
recovery strategy
disk failure 56
human error 55
power failure 55
program failure 55
selecting 53
system failure 55
recovery with 87
redundant failure 389
redundant failure disk unit status 389
referential constraint
pending
editing during IPL 164
restoring 249
release-to-release support 313
removing
disk unit
from ASP 403
disk units 378–380
failed disk unit 196
removing failed unit 88
Rename Directory Entry (RNMDIRE) command
restoring mail 265
Rename Document Library Object (RNMDLO) command
restoring documents 265
renaming
database file
during restore 242
directory entry
restoring mail 265
document library object
restoring documents 265
reply list
restoring 150
reply list entry
reply list entry *(continued)*
recovering 209
resetting
job number counter 148
journal
overflowed status 413
resource not detected status
correcting 228
resource, hardware 388
Restore (RST) command
changed objects 279
how to use 265
restrictions 269
restrictions when restoring documents 271
Restore Authority (RSTAUT) command
nonrestricted state system 219
Restore Configuration (RSTCFG) command 227
Restore Document Library Object (RSTDLO) 190
Restore Document Library Object (RSTDLO) command
maximum number of DLOs 263
media error 42
output 261
overview 260
renaming document 264
restoring authority 264
restoring descriptive information 264
restoring ownership 264
Restore Library (RSTLIB) 188, 193
Restore Library (RSTLIB) command
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
media error 41
multiple concurrent 233
OPTION parameter 232
overview 231
Restore Licensed Program (RSTLICPGM) command 260
restore menu
option 21 (entire system) 203
option 22 (system data only) 203
option 23 (all user data) 203
Restore menu
commands run by menu options 201
using 201
Restore Object (RSTOBJ) command
multiple concurrent 235
restore operations
parallel 44
restore output file
field descriptions 47
format 44
restore strategy
disk failure 56
human error 55
power failure 55
program failure 55
selecting 53
system failure 55
Restore User Profiles (RSTUSRPRF) command 211
restoring
*ALLOBJ (all-object) special authority 215
access path recovery times 149
access paths 246
allowing sensitive programs 34

restoring *(continued)*
ALWOBJDIF (allow object differences) parameter 216, 217
authority
document library object 264
selected objects 218
authority holders 214
authorization list link 216
authorization lists 214
changed objects
by library 278
by object 278
cumulative 278
not cumulative 278
changing console type 230
configuration
problems with system resource management (SRM) information 228
correct sequence 25
damaged document 263
database file
ALWOBJDIF (allow object difference) parameter 241
creation date 241
different member set 243
how to 240
MAXMBRS (maximum members) parameter 243
MBROPT (Member option) parameter 243
member locking 241
renaming 242
deferred
dependent objects 26
physical files 251
determining tape contents 304
different system
network attributes 154
DLO (document library object)
maximum number 263
DLOs (document library objects)
overview 260
document library object
descriptive information 264
media error 42
document library object (DLO)
renaming document 264
document library objects (DLOs)
overview 260
documents
media error 42
overview 260
edit descriptions 150
encrypted auxiliary storage pools 309
encrypted data using encrypting tape 308
encrypted data using software encryption 307
entire system
unattended 203
error is not recoverable 41
error is recoverable 41
file with trigger program 250
files being journaled 236
folders
overview 260
procedure 263
hard link 266
restoring *(continued)*
integrated file system objects being journaled 236
into damaged folder 263
journal receivers 252
journals 252
libraries
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
OPTION parameter 232
overview 231
library
media error 41
Licensed Internal Code
preparation 122
starting 122
steps 123
using distribution media 123
licensed program
choosing procedure 142
manual IPL 142
overview 141
preparation 141
reasons 141
selecting install options 147
steps 143
using distribution media 141
licensed programs 260
link 266
list of members 244
locking objects 38
logical files 244
logical partitions 231
member
renaming 242
message reply list 150
mounted UDFS 239
object
how to 235
multiple names 266
replaying journaled changes 285
object authority 217
objects being journaled 236
objects in directories 265
operating system
hardware encryption 143
software encryption 143
ownership
document library object 264
parts of system 19
private authority 218
program temporary fixes (PTF) 273
programs
ability to use the latest hardware features 257
creating *MODULE, *PGM, *SRVPGM objects 257
PTF (program temporary fixes) 273
QAPZ files 39
QGPL (general purpose) library
QAPZ files 39
QUSRSYS (user system) library
QAPZ files 39
referential constraints 249
related objects 25
Restore System Information 275

restoring *(continued)*
RSTSYSINF 275
save file data 258
Save System Information 275
SAVSYSINF 275
security considerations 34
security information
object authorities 217
object ownership 216
ownership 216
primary group 216
private authorities 217, 218
sequence 210
user profile 211
service attributes 149
shared formats 248
soft link 266
symbolic link 266
system information 149
system management objects 149
system reply list 150
system values 149
unit 295
unmounted UDFS 238
unsuccessful 41
user profile
different system 215
procedure 211
using Restore menu 201
verifying success 39
restoring logical partitions 231
restoring programs 258
restricted state
definition 27
starting 27
Resulting Capacity display 403
resuming
mirror protection 293
mirrored unit 289
resuming status 389
retranslation 257, 258
Retrieve Journal Entry (RTVJRNE) command
using in program 459
RNMDIRE (Rename Directory Entry) command
restoring mail 265
RNMDLO (Rename Document Library Object) command
restoring documents 265
RST (Restore) command
changed objects 279
how to use 265
restrictions 269
restrictions when restoring documents 271
RSTAUT (Restore Authority) command
nonrestricted state system 219
RSTCFG (Restore Configuration) command 227
RSTDLO (Restore Document Library Object) 190
RSTDLO (Restore Document Library Object) command
maximum number of DLOs 263
media error 42
output 261
overview 260
renaming document 264
restoring authority 264
restoring descriptive information 264
RSTDLO (Restore Document Library Object) command *(continued)*
restoring ownership 264
RSTLIB (Restore Library) 188, 193
RSTLIB (Restore Library) command
*ALLUSR libraries 232
*IBM libraries 232
*NONSYS libraries 232
media error 41
multiple concurrent 233
OPTION parameter 232
overview 231
RSTLICPGM (Restore Licensed Program) command 260
RSTOBJ (Restore Object) command
multiple concurrent 235
RSTUSRPRF (Restore User Profiles) command 211
RTVJRNE (Retrieve Journal Entry) command
using in program 459

S

S/36 environment
recovering 230
save file data
restoring 258
Save Library (SAVLIB) command
determining what command was used 304
saving
determining tape contents 304
Encryption Key Manager 308
library
determining what command was used 304
object
previous release system 313
unit 295
SAVLIB (Save Library) command
determining what command was used 304
SAVSYS (Save System) media
definition 122
secondary 54
secondary ASP 54
sector damage 57
security information
components 210
private authority
restoring 218
restoring
sequence 210
sequence
restoring 210
user profile
restoring 211
security-relevant object
allowing restore operation 34
Select ASP to Delete Data From display 197
Select DST Console Mode display 135
Select Product to Work with PTFs display 152, 162
selecting
install options
restoring operating system 147
sensitive object
allowing restore 34
sequence
restoring objects 25
service attribute

service attribute *(continued)*
restoring 149
service representative
mirrored protection recovery action 292
set major system options
during recovery 153
Set Major System Options display 153
setting
major system options
during recovery 153
soft link
restoring 266
software encryption
restoring data 307
some data loss 60
space, disk
calculating 406
spare nonconfigured unit
using 291
special authority
*ALLOBJ (all-object)
restoring 215
Specify ASP to Move Disk Units display 402
Specify Install Options display 148
Specify Restore Options display 149
SQL materialized query table 251
SRC (system reference code)
A6xx 0277 441
A900 2000
recovery 157
Licensed Internal Code using 125
SRM (system resource management) information
correcting problems 228
SST 384
SST (system service tools)
options 381
stand-alone 186, 192
starting
(DST) 382
device
during recovery 152
mirrored protection 431
printer writer
during recovery 152
system
after abnormal end 159
starting device parity protection 419
starting system
display 160
status
disk
understanding 385
unknown load source 297
stopping
(DST) 383
mirrored protection 434
stopping device parity protection 421
storage
reclaiming
duplicate names in QRCL 29
procedure 32, 173
QALWUSRDMN (allow user domain objects) system value 29
recovering 173
storage *(continued)*
reclaiming *(continued)*
user domain object 29
what the system does 29
why to run 167
unit
not operational 294
storage unit
not operational 294
subsystem
ending
QCALSRV (calendar server) subsystem 27
QSYSWRK (subsystem monitor) subsystem 27
restricted state 27
using 27
subsystem monitor (QSYSWRK) subsystem
ending 28
suspended 389
suspended disk unit status 389
suspended status 389
suspending
mirrored units 288
symbolic link
restoring 266
synchronization
BRMS 353
recovery considerations 293
synchronizing
system
methods overview 342
planning and procedures 341
system
parts 19
system data
restoring 203
system failure
recovery strategy 55
system information
recovering 209
restoring 149
system management object
restoring 149
system reference code (SRC)
A900 2000
recovery 157
Licensed Internal Code using 125
system reply list
restoring 150
system resource management (SRM) information
correcting problems 228
system service tools 384
system service tools (SST)
options 381
system state program
allowing restore operation 34
system status display
auxiliary storage
high percentage used 167
system value
allow object restore operation (QALWOBJRST) 34
allow user domain objects (QALWUSRDMN) 29
automatic IPL after power restored (QPWRRSTIPL) 159
changing
during recovery 153

system value *(continued)*
QALWOBJRST (allow object restore operation) 34
QALWUSRDMN (allow user domain objects) 29
QPWRRSTIPL (automatic IPL after power restored) 159
QSCANFSCTL (scan control) 37
QVFYOBJRST (verify object on restore) 34
recovering 209
restoring 149
scan control (QSCANFSCTL) 37
System/36 environment
during recovery 153
recovering 230

T

tape
save
determining what command was used 304
tape configuration
creating
for 34xx tape units 335, 495, 510
for non-34xx tape units 157
tape controller
recovering after restore 229
target release (TGTRLS) parameter
valid values 313
temporary error 287
terminology 53
TGTRLS (target release) parameter
valid values 313
threshold
auxiliary storage pool (ASP)
changing 399
system auxiliary storage pool (ASP)
changing 400
time-out disk error 294
transferring
existing journals into a user ASP 409
folder
different ASP 409
library
different ASP 408
object
different ASP 416
objects between ASPs 408
user profile
different system 215
translation 258
trigger
restoring 250
trigger program
restoring 250

U

UDFS ASP 55
unattended restore operation 203
unbroken receiver chain
applying journaled changes 282
unit
mirrored
resuming 289
suspending 288
unit *(continued)*
missing mirrored disk 294
not operational storage 294
restoring 295
saving 295
spare nonconfigured 291
unit number 388
unprotected status 389
unreadable sectors
recovery 165
unrecoverable device error 293
unrecoverable error
restore operation 41
unsuccessful
restore operation 41
Use DST display 134
user 392, 394, 405, 406, 410, 411, 415
user auxiliary storage pool (ASP)
calculating space requirements 406
changing threshold 399, 400
deleting 377
journal receivers 412
moving disk unit 402
removing disk unit 403
transferring objects 408
user data
restoring 203
user domain object
reclaiming 29
user information
recovering
choosing procedure 105
using commands 107
user profile
*ALLOBJ (all-object) special authority
restoring 215
IBM-supplied
damaged 166
moving to different system 215
restoring 211
user-defined file system (UDFS) ASP 55
user-defined file systems
restore actions 239
restoring 178, 238, 239
using backup 116

V

validation value 256
verify object on restore (QVFYOBJRST) system value 34
verifying
successful restore 39
vital product data information 56, 131

W

when restoring programs 258
working with
display 160
working with device parity protection 419
working with mirrored protection 431
write protected 389
write protected disk unit status 389

writing
output using RCVJRNE (Receiver Journal Entry)
command 460

Product Number: 5770-SS1