z/OS

DFSMS Advanced Copy Services

SC35-0428-01

z/OS

DFSMS Advanced Copy Services

SC35-0428-01

Note
Before using this information and the product it supports, be sure to read the general information under “Notices” on page 315.

Second Edition, March 2002

This edition applies to Version 1 Release 3 of z/OS™ (5694-A01) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC35-0428-00.

Order publications through your IBM® representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation
RCF Processing, Department M86/050
5600 Cottle Road
San Jose, CA 95193-0001
United States of America

IBMLINK from US: STARPUBS at SJEVM5
IBMLINK from Canada: STARPUBS at TORIBM
IBM Mail Exchange: USIB3VVD at IBMMAIL
Internet: starpubs@us.ibm.com
World Wide Web: http://www.storage.ibm.com/software/sms/sdm/index.htm

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:
- Title and order number of this book
- Page number or topic related to your comment

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

© Copyright International Business Machines Corporation 1994, 2002. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Book xi
Required Product Knowledge xi
Referenced Publications xi
Accessing z/OS DFSMS Books on the Internet xi
Using LookAt to look up message explanations xii
Accessing Licensed Books on the Web xii
How to Send Your Comments xiii

Summary of Changes xv
Summary of Changes for SC35-0428-01 z/OS Version 1 Release 3 xv
New Information xv
Changed Information xv
Moved Information xv

Part 1. Advanced Copy Services Overview 1

Chapter 1. Introducing Advanced Copy Services 3
In this chapter 3
Overview of the advanced copy services functions 3
Extended remote copy (XRC) 4
Peer-to-peer remote copy (PPRC) 4
FlashCopy 5
Snapshot 5
Concurrent copy 5
Using the ESS copy services web interface 6

Part 2. Extended Remote Copy 7

Chapter 2. What is Remote Copy? 9
In this chapter 9
Overview of remote copy 9
Using remote copy for disaster recovery 10
Using remote copy for workload migration 11
Deciding which remote copy option to use 12
Choosing extended remote copy 13
Choosing peer-to-peer remote copy 14
Using channel extenders 15
Combining copy services operations 16

Chapter 3. Planning for XRC 19
In this chapter 19
XRC requirements 19
XRC software requirements 19
XRC hardware requirements 19
XRC operational considerations 20
Introducing the XRC components 21
XRC primary systems 21
System data mover 22
XRC volume terminology 24
XRC data sets 25
XRC storage control sessions 26
Establishing XRC secondary volumes 26

© Copyright IBM Corp. 1994, 2002

Evaluating remote copy with DFSMShsm™ 27
Defining storage requirements 28
Virtual storage requirements 28
Real storage requirements 29
Paging space requirements 29
Library requirements 30
Configuring extended remote copy. 30
Choosing an XRC configuration. 30
Including DASD channel extenders 32
Determining XRC resource needs 32
Evaluating how XRC affects the primary storage subsystem 32
Evaluating how XRC affects the recovery storage subsystem 33
Determining the recovery site storage capacity 33
Determining recovery system access to resources 34
Establishing storage subsystem channel connections 34
Establishing XRC ESCON connections 34
Establishing XRC data mover connections 35
Estimating XRC system resources for data mover operations 35
Allocating XRC journal data sets 35
Protecting access to XRC commands 35

Chapter 4. Setting Up the XRC Environment 37
In this chapter 37
Installing and configuring XRC 37
Virtual storage messages 39
Controlling access to XRC resources. 39
Controlling ICKDSF activity to XRC volumes 39
Controlling access to XRC commands 40
Placing TSO commands in an RACF-protected library 40
Defining resource profiles in the RACF FACILITY class 40
Including XRC TSO commands in automated procedures 41
Specifying XRC journal, control, and state data sets 42
Specifying the journal data sets. 43
Specifying the control data set 46
Specifying the state data set 46
Copying the catalog and control data sets 47
Specifying XRC utility volumes 48
Identifying volume pairs. 48
Optimizing XRC performance 49
Configuring for high performance 49
Balancing storage control configurations 52
Distributing work loads 52
Maximizing access to journal data sets 52

Chapter 5. Administering Your XRC Environment 55
In this chapter 55
Issuing extended remote copy TSO commands 55
XADDPAIR–adding volume pairs or utility volumes. 56
XADVANCE—updating secondary volumes 61
XCOUPLE—coupling XRC sessions 63
XDELPAIR–deleting volume pairs 66
XEND–ending a session 70
XQUERY–querying a session 75
XRECOVER–recovering data on the recovery system 82
XSET–changing session parameters 84
XSTART–starting a session 89

XSUSPEND–suspending volumes or sessions 91
ANTRQST application programming interface overview 99

Chapter 6. Managing Extended Remote Copy Operations 101
In this chapter 101
Starting or restarting an XRC session 101
Using XRC utility devices 102
Addressing the utility device selection 102
Adding an XRC volume pair 104
Initializing volume pairs 105
Adding volumes in a volume group 105
Adding a volume with the FULLCOPY or QUICKCOPY option 105
Adding a volume with the NOCOPY option 106
Adding a volume with the SUSPENDED option 107
Adding a volume with the DONOTBLOCK option 108
Understanding the XADDPAIR command 108
Using the ATTIME, DRAIN, and IMMEDIATE keywords 109
Deleting an XRC volume pair 110
Using tertiary XRC volumes 111
Example: using XRC tertiary volumes 111
Using the XSUSPEND command 112
Suspending XRC volume pairs 112
Suspending an XRC session 113
Adding back suspended volumes to the XRC session 114
Using the XSET TIMEOUT parameter 115
Running the ICKDSF program 116
Using XRC XQUERY reports 116
Examples of XQUERY reports in an uncoupled environment 117
Examples of XQUERY reports in a coupled environment 120
Examples of XQUERY ADVANCE and XQUERY RECOVER reports 123
Accessing secondary volumes while XRC is active 124
Ending an XRC session 125
Applying XRC operational scenarios 125
Example: applying maintenance 125
Example: preparing for disaster 126
Example: migrating data 127
Example: switching system modes 128
Example: changing the session error recovery level 129
Example: changing a volume pair's error recovery level 130
Example: suspending and resynchronizing a volume pair 131
Example: suspending and resynchronizing a session 132
Example: changing volume channel connection addresses 133
Examples: changing TIMEOUT intervals for suspended sessions 133
Identifying XRC system interactions 136
XRC varies volumes offline and online 136
Using the DEVSERV and IDCAMS functions 136

Chapter 7. Managing Coupled XRC Sessions 137
In this chapter 137
Overview of CXRC 137
Configuring a coupled XRC environment 138
Identifying sessions in a coupled environment 138
Choosing a coupled XRC configuration 139
Defining data sets in a coupled environment 142
Managing coupled XRC operations 145
Coupling XRC sessions 145

Querying coupled status for XRC sessions 146
Creating a backup copy of secondary volumes (XADVANCE) 147
Restarting an inactive XRC coupled session 147
What to do if the master session is in coupled HOLD status 150
Using the XCOUPLE RELEASE command 155
What to do if a session is in COUPLE_FAILED status 155
Removing coupled sessions 156
Generating dumps of address spaces 156
Applying CXRC operational scenarios 158
Example: recovering from a disaster using the XRECOVER command 158
Example: applying software maintenance in a coupled environment 159

Chapter 8. XRC Data Recovery Operations 161
In this chapter 161
Understanding the XRC volume synchronization process 161
Overview of consistency group time 161
Terms used for consistency group time 162
Changes to consistency group time 163
Example: illustration of consistency group time 164
Understanding the importance of timestamped writes 165
Recovering data with XRC—uncoupled and coupled sessions 166
Recovering data with XRC 166
Creating a recovery volume report 168
Advancing consistency time using the XADVANCE command 169

Chapter 9. Migrating Data with XRC 171
In this chapter 171
Steps for migrating data with XRC 171
Migration operation considerations 172

Chapter 10. Recovering from Error Conditions Using XRC 175
In this chapter 175
Specifying the ERRORLEVEL parameter in coupled and uncoupled sessions 175
Comparing the ERRORLEVEL options 175
ERRORLEVEL(VOLUME) 176
ERRORLEVEL(SESSION) 176
ERRORLEVEL(group_name) 177
Steps for recovering from environmental errors 177
Steps for recovering from processor, address space, and XRC data set failures 178

Part 3. Peer-to-Peer Remote Copy 179

Chapter 11. Planning for PPRC 181
In this chapter 181
PPRC requirements 181
PPRC software requirements 181
PPRC hardware requirements 181
PPRC operational considerations 182
Examining PPRC configuration options 183
Establishing the PPRC configuration 184
Basic PPRC configurations 185
PPRC and DASD channel extenders 185
Determining PPRC resource needs 185
Steps for evaluating how PPRC affects the primary storage subsystem 186
Evaluating how PPRC affects the recovery storage subsystem 186
Determining the recovery site storage capacity 186

Determining recovery system access to resources 187
Establishing storage subsystem channel connections 187
Establishing PPRC ESCON connections 187
Copying the catalog and control data sets 188
Controlling access to PPRC commands 188

Chapter 12. Setting Up the PPRC Environment 189
In this chapter 189
Installing peer-to-peer remote copy 189
Controlling access to PPRC resources 190
Protecting access to PPRC commands 190
Controlling ICKDSF activity to PPRC volumes 191
Identifying volume pairs 192
Steps for including PPRC TSO commands in automated procedures 193
Configuring storage subsystem resources for PPRC 194
Configuring IBM storage controls 194
Accessing PPRC secondary volumes 195
Optimizing PPRC performance 195

Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions 199
In this chapter 199
PPRC command overview 199
CDELPAIR – deleting pairs 200
CDELPAIR command syntax 200
CDELPAIR command examples 201
CDELPATH – deleting paths 201
CDELPATH command syntax 201
CDELPATH command examples 202
CESTPAIR – establishing pairs 202
CESTPAIR command syntax 202
CESTPAIR command example 205
CESTPATH – establishing paths 205
CESTPATH command syntax 206
CESTPATH command examples 207
CGROUP – controlling volume groups 208
CGROUP command syntax 208
CGROUP command examples 208
CQUERY – querying status 209
CQUERY command syntax 209
CQUERY command examples 210
CRECOVER – recovering data on the recovery system 210
CRECOVER command syntax 210
CRECOVER command example 211
CSUSPEND – suspending pairs 211
CSUSPEND command syntax 213
CSUSPEND command example 214

Chapter 14. Managing Peer-to-Peer Remote Copy Operations 215
In this chapter 215
Managing the PPRC system 215
System IPL volumes 215
ESCON manager 215
Resource Measurement Facility™ (RMF™) data 216
DEVSERV and IDCAMS 216
Identifying peer-to-peer remote copy volume states 216
Establishing paths 218

Managing volumes and paths 221
Adding a PPRC volume pair 221
Querying PPRC volumes 222
Querying PPRC paths 227
Suspending PPRC volume pair operations 230
Deleting a PPRC volume pair 231
Deleting PPRC paths 232

Chapter 15. PPRC Data Recovery Operations 233
In this chapter 233
Examining disaster recovery from a system viewpoint 233
Resuming operations at the recovery site 234
Managing errors during recovery 236
Errors encountered when writing to PPRC volumes established with CRIT(YES) 236
Errors that occur in the absence of write activity 236
Errors that cause a PPRC pair to become suspended 236

Chapter 16. Moving and Migrating Data with PPRC 239
In this chapter 239
Steps for moving the work load with PPRC 239
Steps for migrating data with PPRC 240
Steps for copying the secondary volumes with PPRC 241

Chapter 17. Recovering from PPRC Error Conditions 243
In this chapter 243
Understanding host error recovery procedures 243
Preparing for PPRC error recovery 244
Putting PPRC error recovery procedures into effect 244
RAMAC device sparing with PPRC 249
Steps to spare a RAMAC drawer 250

Chapter 18. PPRC Dynamic Address Switching (P/DAS) 251
In this chapter 251
How does P/DAS work? 251
Preparing for P/DAS operations 252
Hardware and software requirements 252
Environmental conditions 252
Additional P/DAS considerations 253
P/DAS operations 253
Typical PPRC configuration before P/DAS operations 253
Steps for using P/DAS in a nonsysplex, shared-disk environment 254
Typical PPRC configuration after a P/DAS function has completed 256
Steps for using P/DAS in a sysplex environment 256
Replying to P/DAS-related messages 259
System actions for switch pair, and swap 260
P/DAS error conditions 262
Storage-related failures 262
I/O-related failures 262
Validation failures 263
P/DAS operations example 263

Part 4. FlashCopy 265

Chapter 19. What Is FlashCopy? 267
In this chapter 267

FlashCopy requirements 267
Determining FlashCopy software requirements 267
Determining FlashCopy hardware requirements 267
Using FlashCopy 268
Steps for initializing sessions 268
Copying data sets with FlashCopy 268
FlashCopy TSO commands 269
Controlling access to FlashCopy commands 269
FlashCopy establish (FCESTABL) command 270
FlashCopy query (FCQUERY) command 274
FlashCopy withdraw (FCWITHDR) command 275

Part 5. Snapshot Copy 277

Chapter 20. What Is Snapshot Copy? 279
In this chapter 280
Snapshot copy requirements 280
Determining Snapshot copy software requirements 280
Determining Snapshot copy hardware requirements 280
Using Snapshot copy 280
Steps for initializing sessions 280
Steps for providing a Snapshot working-space data space 281
Copying data with Snapshot copy 282
Moving data sets with Snapshot copy 282
Using Snapshot copy, PPRC, and RVA together 283
Snapshot SMF information 284

Part 6. Concurrent Copy 285

Chapter 21. What Is Concurrent Copy? 287
In this chapter 287
Understanding the concurrent copy solution 287
Concurrent copy requirements 288
Determining concurrent copy software requirements 288
Determining concurrent copy hardware requirements 288
Determining concurrent copy resource needs 288
Central and expanded storage 288
Storage subsystem cache 289
Channel paths 290
Using concurrent copy 290
Steps for initializing sessions 290
Copying the data 291
SMF information 291
Scheduling concurrent copy operations 291

Part 7. Appendixes 293

Appendix A. Advanced Copy Services Diagnostic Aids 295
Restarting data mover address spaces (SDM, XRC, CC, Snapshot) 295
Diagnosing system data mover functions with the MVS MODIFY command 295
AUTO_READD operation (XRC) 296
DUMP operation (XRC, CC, Snapshot) 296
DVCDATA operation (XRC, CC) 297
LISTDVCS operation (XRC) 298
LISTSESS operation (XRC, CC) 298

RESTART operation (XRC) 299
RCVRSESS operation (SDM, XRC, CC) 300
SCDATA operation (XRC, CC) 300
SCTRAP operation (XRC) 301
TERMDVC operation (XRC) 302
TERMSESS operation (XRC, CC) 302
XENDDUMP operation (XRC) 304
XRCTRAP/SUPRDUMP operation (XRC) 304
Peer-to-peer remote copy diagnostic aids. 304

Appendix B. Configuring the Data Mover and PPRC Options on IBM 3990 and 9390 Storage Controls 305

Appendix C. SMF Type 42 Records 307
XRC information in SMF type 42 records 307
Concurrent copy information in SMF type 42 records 309
Concurrent copy SMF type 42 subtype 4 records 309

Appendix D. Advanced Copy Services Messages 311

Appendix E. Accessibility 313
Using assistive technologies 313
Keyboard navigation of the user interface. 313

Notices 315
Programming Interface Information 316
Trademarks. 316

Glossary 317

Index 325

About This Book

The purpose of this publication is to help you understand and use IBM® Advanced Copy Services functions. It describes three “dynamic” copy functions and several “point-in-time” copy functions. These functions provide backup and recovery of data should a disaster occur to your data center. The dynamic functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy, and are known collectively as remote copy. FlashCopy, SnapShot, and Concurrent Copy are the point-in-time copy functions.

Required Product Knowledge

To use this book effectively, you should be familiar with the following information:

- Programming, especially programming with TSO commands
- Current disaster recovery and workload migration procedures at your location

Referenced Publications

The following publications are referenced in this book:

Publication Title	Order Number
z/OS TSO/E Command Reference	SA22-7782
z/OS DFSMS Introduction	SC26-7397
z/OS DFSMSdfp Storage Administration Reference	SC26-7402
z/OS DFSMSdfp Advanced Services	SC26-7400
z/OS DFSMSdss Storage Administration Guide	SC35-0423
z/OS DFSMSdss Storage Administration Reference	SC35-0424
z/OS MVS System Messages, Volume 1 (ABA-AOM)	SA22-7631
z/OS Security Server RACF Command Language Reference	SA22-7687
z/OS Security Server RACF Security Administrator's Guide	SA22-7683
IBM Enterprise Storage Server Web Users Interface Guide	SC26–7346
IBM Enterprise Storage Server Introduction and Planning Guide	GC26–7294
Implementing the Enterprise Storage Server in Your Environment	SG24–5420
Using Enterprise Systems Connection Directors	GA23–0354

Accessing z/OS DFSMS Books on the Internet

In addition to making softcopy books available on CD-ROM, IBM provides access to unlicensed z/OS softcopy books on the Internet. To find z/OS books on the Internet, first go to the z/OS home page: http://www.ibm.com/servers/eserver/zseries/zos

From this Web site, you can link directly to the z/OS softcopy books by selecting the Library icon. You can also link to IBM Direct to order hardcopy books.

© Copyright IBM Corp. 1994, 2002

Using LookAt to look up message explanations

LookAt is an online facility that allows you to look up explanations for z/OS messages, system abends, and some codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can access LookAt from the Internet at:

`http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html`

or from anywhere in z/OS where you can access a TSO command line (for example, TSO prompt, ISPF, z/OS UNIX System Services running OMVS).

To find a message explanation on the Internet, go to the LookAt Web site and simply enter the message identifier (for example, `IAT1836` or `IAT*`). You can select a specific release to narrow your search. You can also download code from the *z/OS Collection*, SK3T-4269 and the LookAt Web site so you can access LookAt from a PalmPilot (Palm VIIx suggested).

To use LookAt as a TSO command, you must have LookAt installed on your host system. You can obtain the LookAt code for TSO from a disk on your *z/OS Collection*, SK3T-4269 or from the LookAt Web site. To obtain the code from the LookAt Web site, do the following:

1. Go to http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html.
2. Click the **News** button.
3. Scroll to **Download LookAt Code for TSO and VM.**
4. Click the ftp link, which will take you to a list of operating systems. Select the appropriate operating system. Then select the appropriate release.
5. Find the **lookat.me** file and follow its detailed instructions.

To find a message explanation from a TSO command line, simply enter: **lookat** *message-id*. LookAt will display the message explanation for the message requested.

Note: Some messages have information in more than one book. For example, IEC192I has routing and descriptor codes listed in *z/OS MVS Routing and Descriptor Codes*. For such messages, LookAt prompts you to choose which book to open.

Accessing Licensed Books on the Web

z/OS licensed documentation in PDF format is available on the Internet at the IBM Resource Link Web site at:

`http://www.ibm.com/servers/resourcelink`

Licensed books are available only to customers with a z/OS license. Access to these books requires an IBM Resource Link Web userid and password, and a key code. With your z/OS order you received a memo that includes this key code.

To obtain your IBM Resource Link Web userid and password log on to:

`http://www.ibm.com/servers/resourcelink`

To register for access to the z/OS licensed books:

1. Log on to Resource Link using your Resource Link user ID and password.

2. Click on **User Profiles** located on the left-hand navigation bar.
3. Click on **Access Profile.**
4. Click on **Request Access to Licensed books.**
5. Supply your key code where requested and click on the **Submit** button.

If you supplied the correct key code you will receive confirmation that your request is being processed. After your request is processed, you will receive an e-mail confirmation.

Note: You cannot access the z/OS licensed books unless you have registered for access to them and received an e-mail confirmation informing you that your request has been processed.

To access the licensed books:

1. Log on to Resource Link using your Resource Link userid and password.
2. Click on **Library**.
3. Click on **zSeries**.
4. Click on **Software**.
5. Click on **z/OS**.
6. Access the licensed book by selecting the appropriate element.

How to Send Your Comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other DFSMS documentation:

- Visit our Web site: http://www.storage.ibm.com/software/sms/sdm/index.htm
- Send your comments by e-mail to:
  - IBMLink™ from US: starpubs@us.ibm.com
  - IBMLink from Canada: STARPUBS at TORIBM
  - IBM Mail Exchange: USIB3VVD at IBMMAIL
  - Internet: starpubs@us.ibm.com

  Be sure to include the name of the book, the part number of the book, version and product name, and if applicable, the specific location of the text you are commenting on (for example, a page number or a table number).
- Fill out one of the forms at the back of this book and return it by mail or by giving it to an IBM representative. If the form has been removed, address your comments to IBM Corporation, RCF Processing Department 61C, 9000 S. Rita Road, Tucson, Arizona 85744-0001, U.S.A.

Summary of Changes

This book contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

You may notice changes in the style and structure of some content in this book, for example, procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our books.

Summary of Changes for SC35-0428-01 z/OS Version 1 Release 3

This book contains information previously presented in *z/OS Version 1 Release 1 DFSMS Advanced Copy Services* (SC35-0428-00).

The following sections summarize the changes to that information.

New Information

This edition includes the following new information:

- Support for fixed utility devices. The XSET and XADDPAIR commands have been enhanced to allow you to specify a dedicated (fixed) device to be used for reading data from a storage subsystem. The XQUERY command has also been enhanced to show the type of utility devices used.
- Support of XRC recovery of tertiary volumes. The XRECOVER and XADVANCE commands have been enhanced to allow you to recover to copies of secondary volumes (other than the original secondary volumes).

Changed Information

The following information was changed in this edition:

- Virtual storage requirements enhancements. Guidelines are provided to help you manage virtual storage used by the ANTAS*nnn* address spaces.

Moved Information

Coupled Extended Remote Copy (CXRC) information that was previously described across several chapters is now included in its own chapter called “Chapter 7. Managing Coupled XRC Sessions” on page 137.

© Copyright IBM Corp. 1994, 2002

Part 1. Advanced Copy Services Overview

© Copyright IBM Corp. 1994, 2002

Chapter 1. Introducing Advanced Copy Services

As businesses become more dependent on information technology to conduct their operations, the requirements for protecting their data, backing it up, and making it available, 24 hours a day and seven days a week are forever increasing. In the event of disaster or system failures, without a plan for recovery, loss of data can lead to significant financial cost or maybe total failure for businesses.

IBM Advanced Copy Services provides a combined hardware and software solution to help address these issues. Advanced Copy Services is a collection of functions that provides best-of-breed solutions to the complex challenges of disaster recovery, data migration, and data duplication. Many of these functions run on IBM Enterprise Storage Server™ (ESS). The ESS, designed for mid-range and high-end environments, gives you a high performance, continuous availability, high capacity storage subsystem that can be configured according to your requirements.

This chapter provides overviews to help you understand the functions of Advanced Copy Services.

In this chapter

This chapter introduces you to the functions of Advanced Copy Services. It includes the following sections:

Topic . . .	See page . . .
Overview of the advanced copy services functions	3
Using the ESS copy services web interface	6

Overview of the advanced copy services functions

Advanced Copy Services comprises the following functions and enhancements:

- Extended Remote Copy (XRC), which includes the following enhancements:
  - Coupled Extended Remote Copy (CXRC)
  - Multiple Extended Remote Copy (MXRC)
- Peer-to-Peer Remote Copy (PPRC)
- FlashCopy
- Snapshot
- Concurrent Copy (CC)

Table 1 shows how each of the Advanced Copy Services copy functions falls into one of two groups. Dynamic copy functions constantly update the secondary copy as applications make changes to the primary data source. Point-in-time copy functions provide an instantaneous copy, or view, of what the original data looked like at a specific point in time.

Table 1. Dynamic copy and point-in-time copy functions

Type of copy	Copy function names
Dynamic copy of data	XRC, CXRC, MXRC, PPRC
Point-in-time copy of data	FlashCopy, Snapshot, and Concurrent Copy

© Copyright IBM Corp. 1994, 2002

Note: Advanced Copy Services largely relies on a data movement engine, the system data mover (SDM), to efficiently and reliably move large amounts of data between storage devices. For more information, see “System data mover” on page 22.

Extended remote copy (XRC)

Remote copy offers two options for your disaster recovery and workload migration needs: extended remote copy (XRC) and peer-to-peer remote copy (PPRC). XRC addresses the problem of unrecoverable data that occurs between the last, safe backup of a primary system to a recovery system and the time when the primary system fails.

XRC provides an asynchronous copy operation, over distance, with minimal performance impact to primary system DASD I/O write operations. For more information about XRC, see “Chapter 2. What is Remote Copy?” on page 9.

Table 2 provides related XRC topics and their locations.

Table 2. XRC-related sections

For information about . . .	See . . .
Planning for XRC	“Chapter 3. Planning for XRC” on page 19
Setting up your XRC environment	“Chapter 4. Setting Up the XRC Environment” on page 37
Issuing XRC commands	“Chapter 5. Administering Your XRC Environment” on page 55
Managing your XRC environment	“Chapter 6. Managing Extended Remote Copy Operations” on page 101
Managing Coupled XRC sessions	“Chapter 7. Managing Coupled XRC Sessions” on page 137
Recovering your data with XRC	“Chapter 8. XRC Data Recovery Operations” on page 161
Migrating data with XRC	“Chapter 9. Migrating Data with XRC” on page 171
Recovering from error conditions using XRC	“Chapter 10. Recovering from Error Conditions Using XRC” on page 175

Peer-to-peer remote copy (PPRC)

Like XRC, Peer-to-Peer Remote Copy (PPRC) addresses the problem of unrecoverable data that occurs between the last, safe backup of a primary system to a recovery system and the time when the primary system fails.

A PPRC data copy to the recovery storage subsystem is synchronous with the primary volume’s I/O operation. For more information about PPRC, see “Overview of remote copy” on page 9.

For other related PPRC topics, see Table 3 provides related PPRC topics and their locations.

Table 3. PPRC-related topics

For information about . . .	See . . .
Planning for PPRC	“Chapter 3. Planning for XRC” on page 19
Setting up your PPRC environment	“Chapter 4. Setting Up the XRC Environment” on page 37
Issuing PPRC commands	“Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199
Managing your PPRC environment	“Chapter 14. Managing Peer-to-Peer Remote Copy Operations” on page 215
Recovering your data with PPRC	“Chapter 15. PPRC Data Recovery Operations” on page 233

FlashCopy

FlashCopy provides a point-in-time copy of data for backup and recovery operations. Both source and target volumes must reside on the same logical subsystem. FlashCopy creates a time zero (T0) copy of a volume. The T0 copy of a volume is current with the volume at the time the action was initiated.

DFSMSdss™ automatically invokes FlashCopy when you issue the COPY FULL command and the copy operation is for data on a subsystem that supports FlashCopy functions.

For more information on FlashCopy, see “Chapter 19. What Is FlashCopy?” on page 267.

Snapshot

Snapshot is a function of RAMAC® Virtual Array (RVA) Storage Subsystems that allows you to make a very quick copy of a set of tracks (an entire volume, a data set, or just a random set of tracks). The copy operation is completed with only a few I/Os to the device. For more information on Snapshot, see “Chapter 20. What Is Snapshot Copy?” on page 279.

Table 4. Snapshot-related topics

For information about . . .	See . . .
Snapshot requirements	“Snapshot copy requirements” on page 280
Using Snapshot, PPRC, and RVA together	“Using Snapshot copy, PPRC, and RVA together” on page 283

Concurrent copy

Concurrent copy is an extended function that enables data center operations staff to generate a copy or a dump of data while applications are updating that data. Concurrent copy delivers a copy of the data, in a consistent form, as it existed before the updates took place. For more information on Concurrent Copy, see “Chapter 21. What Is Concurrent Copy?” on page 287.

Using the ESS copy services web interface

IBM includes a Web-browser interface called StorWatch™ Enterprise Storage Server (ESS) Copy Services. The interface is part of the ESS storage subsystem and can be used to perform FlashCopy and PPRC functions.

The default support for the Web Copy Services is for open-systems environments. If you plan to use Copy Services functions using the StorWatch ESS Copy Services interface, ensure that your IBM service representative enables the proper support to use System/390® volumes. The *IBM Enterprise Storage Server Introduction and Planning Guide* contains a worksheet for ESS Copy Services that helps the service representative determine which option to enable. For more information, see *IBM Enterprise Storage Server Web Users Interface Guide*.

Part 2. Extended Remote Copy

© Copyright IBM Corp. 1994, 2002

Chapter 2. What is Remote Copy?

Remote copy is a storage-based disaster recovery and workload migration solution that allows you to copy data to a remote location in real time. “Remote copy” refers to two related Advanced Copy Services functions: Extended Remote Copy (XRC) and Peer-to-Peer Remote Copy (PPRC).

Note: Remote copy is an extended function on most IBM storage controls, along with the appropriate levels of DFSMSdfp™, and either DFSMS/MVS™ or z/OS.

Enhancements to XRC include Coupled Extended Remote Copy (CXRC) and Multiple Extended Remote Copy (MXRC). For more information on coupled XRC, see “Chapter 7. Managing Coupled XRC Sessions” on page 137.

In this chapter

This chapter to provides overviews on these remote copy (XRC and PPRC) solutions, and how to choose which one to use in your environment. It includes the following sections:

Topic . . .	See page . . .
Overview of remote copy	9
Using remote copy for disaster recovery	10
Using remote copy for workload migration	11
Deciding which remote copy option to use	12
Using channel extenders	15
Combining copy services operations	16

Overview of remote copy

Remote copy operates with two systems: A primary system at one location and a recovery system at another location. You can locate both systems in the same building or at remote locations. The recovery system only needs to be in place for the time when a recovery is required, and can be a stand-alone system. Each system has specific DASD that processes data that you have identified as remote copy-managed. In case of a disaster at your primary location, your recovery system accesses data from the recovery system DASD.

Once established, remote copy provides the following benefits:

- Provides a real-time copy of all SMS-managed and non-SMS-managed data.
- Copies changes of your primary system data to your recovery system as you make those changes to your primary system data.
- Is application independent. You do not need separate copy functions for multiple applications or databases.
- Supports all DASD data needed for application recovery.
- Provides an improved recovery time on the recovery system, especially in the case of a disaster.

Note: To help switch applications from the primary site to the recovery location, you may want to automate operations, with a program such as NetView™ or

© Copyright IBM Corp. 1994, 2002

Geographically Dispersed Parallel Sysplex® (GDPS™). Your installation can make this switch with a minimal impact on end users.

Using remote copy for disaster recovery

Disasters occur in many forms. Some disasters happen suddenly, when they do, they stop all processing at a single point in time. A more likely scenario happens when a disaster interrupts operations in stages that occur over several seconds or even minutes. This is often referred to as a rolling disaster. You must, therefore, plan for recovery from a potential disaster that causes system failures that are *immediate, intermittent, or gradual*. IBM's remote copy functions address this real-world situation.

Recovering from a disaster without remote copy

Figure 1 shows how a data processing facility attempts disaster recovery without remote copy. The existing system cannot keep data at the recovery site current; therefore, any updates made since the last periodic backup may be lost. Without remote copy, you must determine what updates have occurred since the last backup, and then try to manually recreate the missing data.

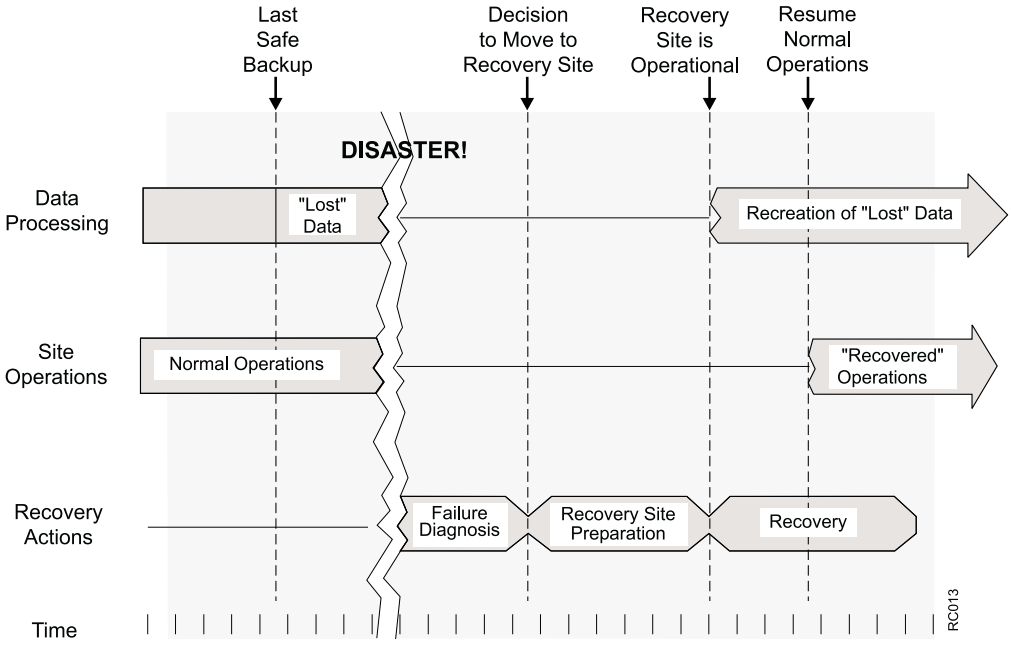

Figure 1. Current disaster recovery scenario (without remote copy). These actions portray a typical sequence of events.

Recovering with remote copy

Figure 2 on page 11 shows how disaster recovery proceeds with remote copy solutions.

Comparing Figure 1 with Figure 2 on page 11, you find that remote copy improves the current disaster recovery process by allowing a much shorter recovery time with little or no data loss. Remote copy promotes faster recovery because it constantly updates the records at the recovery site to match the primary application records. In the event of a disaster, remote copy minimizes or even eliminates data loss.

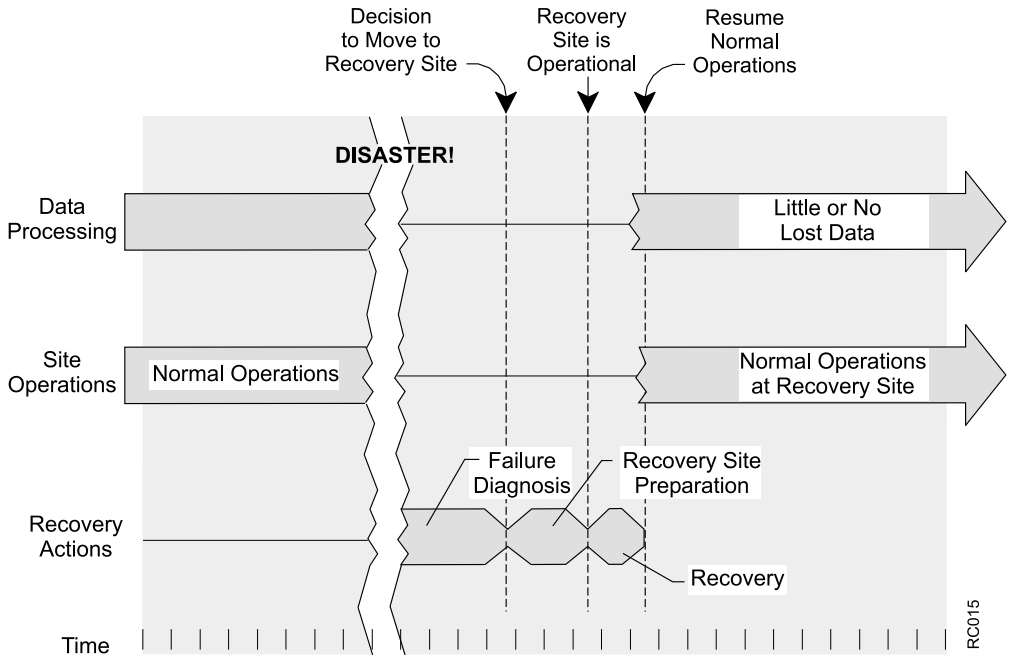

Figure 2. Disaster recovery scenario with remote copy. Remote copy allows recovery to take place faster and more completely.

The remote copy XRC and PPRC options maintain current copies of your vital data as follows:

- With PPRC, no DASD data is lost between the last update at the primary system and the recovery at the recovery system site.
- With XRC, the only data that is lost is data that is in transit between the time that the primary system fails and the recovery at the recovery system site occurs.

Using remote copy for workload migration

In addition to disaster recovery, you can use remote copy for migrating the workload from one location to another or for moving (or migrating) data from one set of DASD volumes to another, with minimal effect on continuous availability applications. For example, reasons for migrating data from one volume to another might include the following:

- Removing volumes from use
- Performing maintenance on a storage control
- Moving applications to a new data center in another location

The following list describes information for using remote copy to migrate data.

For information about . . .	See . . .
Migrating a remote-copy DASD volume to any other DASD type.	“Chapter 9. Migrating Data with XRC” on page 171
Migrating work from the primary site to a secondary site.	“Chapter 16. Moving and Migrating Data with PPRC” on page 239
Redirecting application I/O from one set of DASD volumes to another set of the same type.	“Chapter 18. PPRC Dynamic Address Switching (P/DAS)” on page 251

Restrictions for migrating data

Remote copy can copy data from one volume to another if the number of bytes per track and the number of tracks per cylinder are the same between the two devices, and the target device has an equal or greater number of cylinders. The only way to copy data between unlike devices or from larger devices to smaller capacity devices is by using DFSMSdss™ (or DFDSS, or a comparable product) logical dump, restore, or copy functions. Although the 3380 DASD volumes do not have the same number of bytes per track as the 3390 DASD volumes, you can use DFSMSdss to copy data between the two unlike devices. For more information on DFSMSdss copy functions, see *z/OS DFSMSdss Storage Administration Guide*.

Advantages of using remote copy to migrate data

Using remote copy over other data migration methods provide the following advantages:

- Minimal application down time — Remote copy updates the remote migration copy at the same time that you update the application. When the migration copy is complete, you simply quiesce application programs, ensure that all updates are on the migration volumes, and start up the application using the new volumes.
- Minimal application impact — Application performance is not affected during the migration. You can continue to update data normally, with a pause to switch to the new volumes.
- Migration automation — There is less operator intervention that is required to perform the migration because there is no need for a dump and restore process.

Deciding which remote copy option to use

Based on your specific needs, you can choose between one of the two remote copy disaster recovery options: Extended remote copy (XRC) or peer-to-peer remote copy (PPRC).

The following sections describe these two options in more detail. Table 5 compares both the extended remote copy and peer-to-peer remote copy options.

Table 5. Comparison of extended remote copy and peer-to-peer remote copy

Topic . . .	Extended remote copy	Peer-to-peer remote copy
Application independent?	Yes	Yes
Type of solution	Hardware and software	Hardware and software
Design priority	Minimize impact on application performance	Maintain current data at recovery site
Recovery	Entire session	Volume-by-volume basis
Recovery site data is current?	Except the bytes in transit	Secondary DASD kept fully current
Support software for system with data mover	All supported OS/390® and z/OS releases	All supported OS/390 and z/OS releases
Primary system support software	Must support timestamping function	All supported OS/390 and z/OS releases
Support for large environments	Yes	No
Channel type	ESCON®®, fiber, or parallel	ESCON only for PPRC link (ESCON, fiber, or parallel for host links)
Supported copy distance	Any distance supported (1)	Maximum ESCON distance between storage controls
Copy operation	Asynchronous	Synchronous (2)

Table 5. Comparison of extended remote copy and peer-to-peer remote copy (continued)

Topic . . .	Extended remote copy	Peer-to-peer remote copy
Primary site storage control	Cached storage subsystem	Cached storage subsystem
Recovery site storage control	Cached storage subsystem	Cached storage subsystem

Notes:
1. In theory, if the data mover were located in a third location, midway between both storage subsystems, the maximum total copy distance would be twice the maximum ESCON distance. Adding DASD channel extenders can further extend the distance by sending the data across telecommunication lines. See “Using channel extenders” on page 15 for more information.
2. The initial volume copy to the secondary device is asynchronous.

Choosing extended remote copy

Extended remote copy is a combined hardware and software solution to the problem of accurate and rapid disaster recovery. XRC also provides a DASD and workload migration solution.

XRC is designed for sites that match the following criteria:

- Must maintain the highest levels of performance on their primary system.
- Support extended distances between volume copies.
- Can accept a gap of a few seconds between writes on the primary system and the subsequent write updates on the recovery system.
- Support a large number of primary storage controls and DASD volumes using CXRC, an XRC enhancement. See “Chapter 7. Managing Coupled XRC Sessions” on page 137 for more information on CXRC.

Protecting your enterprise’s data requires that related updates are applied to the secondary volumes in the same order as they were applied on the primary volumes. Maintaining data integrity becomes especially critical when a volume is updated by multiple applications, or when a data set exists on multiple volumes spread across multiple storage controls. XRC’s design strategy ensures that secondary updates are applied on a consistent basis across multiple storage controls. This update sequencing is necessary in order to avoid data integrity problems and potential data loss.

XRC supports channel extenders, ESCON, FICON™, and parallel channel operation. With ESCON and FICON channels, the XRC function provides for the recovery system to reside at an extended distance from the primary system.

Figure 3 on page 14 presents an overview of the XRC option of remote copy. The workload (or DASD) migration configuration is basically the same, except that the “recovery site” would be considered the “migration target”.

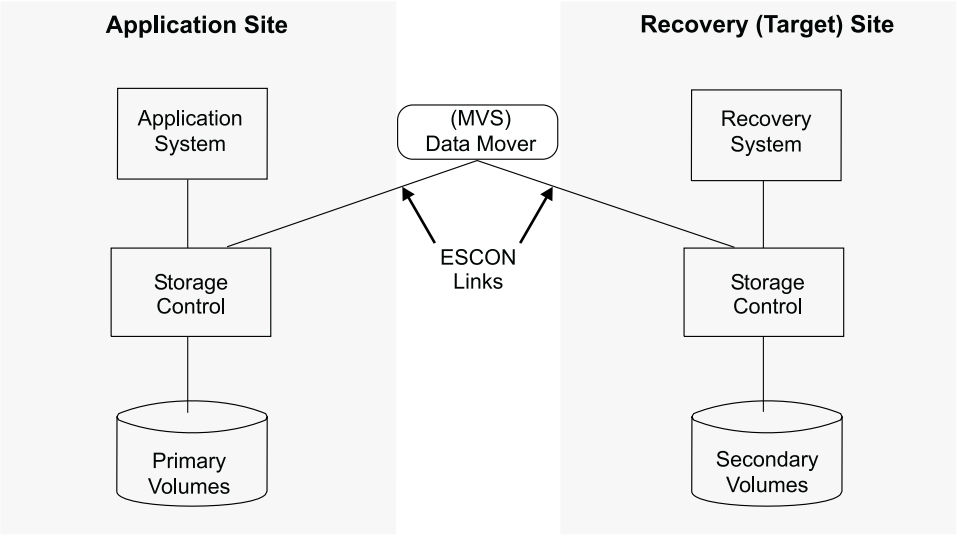

Figure 3. Extended remote copy overview

How extended remote copy (XRC) works

XRC is implemented in a cooperative way between cached storage subsystems and DFSMS/MVS host system software. With XRC, copies of updated data are automatically sent to the recovery system. This is done asynchronously to data updates on the primary system, and with a minimal increase to DASD write response time at the application.

Choosing peer-to-peer remote copy

Peer-to-peer remote copy, available on all supported OS/390 and z/OS releases, is a hardware solution to the problem of accurate and rapid disaster recovery, and also provides a workload migration and DASD migration solution.

Peer-to-peer remote copy is designed for those sites that conform to the following:

- Need the recovery system to always be fully current with the primary system.
- Can accept some performance impact to application write I/O operations at the primary location.

Figure 4 on page 15 presents an overview of the PPRC option of remote copy, based on a disaster recovery need. The workload (or DASD) migration configuration is basically the same, except that the “recovery site” would be considered the “migration target.”

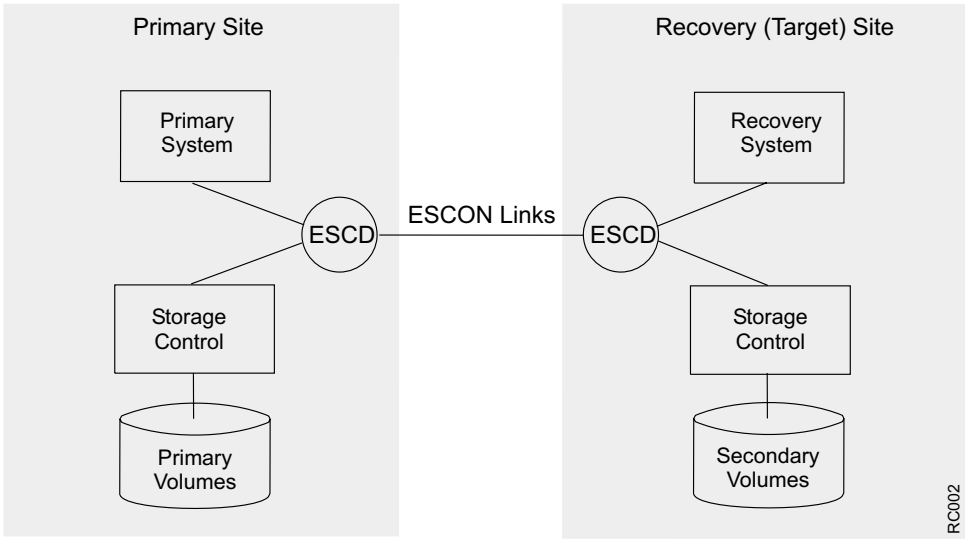

Figure 4. PPRC overview

How peer-to-peer remote copy works (PPRC)

The following describes the PPRC copy functions:

- A PPRC data copy to the recovery storage subsystem is synchronous with the primary volume's I/O operation. This means that the primary system writes data to a primary volume's storage subsystem DASD, the data is transferred to cache and nonvolatile storage (NVS), and the storage control sends channel end status to the host.
- The primary site storage control then initiates an I/O channel program to the recovery site storage control to transfer the updated data.
- The primary site storage control returns device end status to the primary system when the transfer to the recovery site storage control cache and NVS is complete.
- The primary system notifies the application program that the operation is complete.

Note: The PPRC copy function does not consider the primary system DASD write operation complete until the data that is sent to the recovery storage subsystem has received channel end and device end status from the secondary volume storage control. Each primary system write to the recovery subsystem causes an increase to the primary system response time.

Using channel extenders

Remote copy supports channel extenders for instances where a direct ESCON connection is not possible. DASD channel extenders perform the following functions:

- Allow you to configure remote copy DASD between locations that are either beyond the maximum ESCON channel distance, or where ESCON fiber optic cables cannot physically be placed.
- Operate as a pair of devices. One device attaches to a host system channel and the second unit connects to the storage control, which considers the attached

extender to be a system channel. Both devices are connected through a network of fiber-optic cables or telecommunication lines.

An advantage of XRC is that it minimizes the effect on applications when data moves across long distances. DASD channel extenders provide an important part of this solution. There may be a small delay between the time that the application makes the update and when the data mover receives the update. This potential delay is a result of the time that is required for signals to travel greater distances and the channel-to-network protocol translation necessary to achieve transmission at the extended distance.

For more information on channel extenders, see “PPRC and DASD channel extenders” on page 185 and “Including DASD channel extenders” on page 32.

Combining copy services operations

Remote copy-capable IBM storage controls provide support for various copy services. These services include dual copy, concurrent copy, snapshot, FlashCopy, XRC, and PPRC. Various hardware platforms support different combinations of these capabilities as shown in Table 6 and Table 7 on page 17.

Note: The RAMAC array DASD dynamic sparing feature is disabled when the RAMAC volume is a member of a dual copy or PPRC volume pair.

To determine whether a combination works, select one volume type from the top row and one from the leftmost column. The top row represents the copy service function in which a device is currently active, and the leftmost column represents the copy service function in which a device may become active. Notes for both tables can be found at the bottom of Table 7 on page 17.

Table 6. Allowable combinations of copy operations on the same device. This table continues with Table 7 on page 17.

If device is → can it also become ↓?	FlashCopy source	FlashCopy target	Concurrent copy source	Snapshot copy source	Snapshot copy target
XRC source	Yes	Yes	Yes	No	No
XRC target	Yes	Yes (3)	Yes	Yes	No
PPRC source	Yes	Yes	Yes	Yes	Yes
PPRC target	Yes	Yes (3)	No	Yes	No
Concurrent copy source (2)	Yes	Yes	Yes	No	No
FlashCopy source	No	No	Yes	No	No
FlashCopy target	No	No	No	No	No
Snapshot copy source	No	No	No	No	Yes
Snapshot copy target	No	No	No	Yes	No

Table 7. Allowable combinations of copy operations on the same device (continued). This is a continuation of Table 6 on page 16.

If device is → can it also become ↓?	XRC source	XRC target	PPRC source	PPRC target
XRC source	No	Yes	Yes	No
XRC target	Yes	No (4)	Yes	No
PPRC source	Yes	Yes	No	No
PPRC target	No	No	No	No
Concurrent copy source (2)	Yes	Yes	Yes	No
FlashCopy source	Yes	Yes	Yes	Yes (1)
FlashCopy target	No	No (4)	No	No
Snapshot copy source	No	Yes	No	Yes
Snapshot copy target	No	No	No	No

Notes:

1. The primary control unit could see the long busy condition while the PPRC secondary is destaging all cache modified data in order to become a FlashCopy pair.
2. These columns and rows should be interpreted on an extent basis.
3. The operation is allowed. Updates of the affected extents will result in implicit removal of the FlashCopy relationship.
4. Host software disallows this combination.

Chapter 3. Planning for XRC

Before installing and using XRC, you need to plan for it. Understanding the system requirements helps you prepare for a successful installation of XRC. If you are considering a coupled XRC environment, see also “Chapter 7. Managing Coupled XRC Sessions” on page 137.

In this chapter

This chapter is intended to help you plan for extended remote copy. It includes the following sections:

Topic . . .	See page . . .
XRC requirements	19
XRC operational considerations	20
Introducing the XRC components	21
Establishing XRC secondary volumes	26
Evaluating remote copy with DFSMShsm™	27
Defining storage requirements	28
Configuring extended remote copy	30
Determining XRC resource needs	32
Protecting access to XRC commands	35

XRC requirements

You must meet the software and hardware requirements below in order to plan for and successfully install XRC. Table 5 on page 12 includes information about XRC hardware and software requirements.

XRC software requirements

The data mover function requires all supported OS/390 and z/OS releases. The minimum MVS™ level that supports the IBM 9037 sysplex timer is MVS/ESA™ 4.1.0 with MVS/DFP™ 3.3.0.

XRC hardware requirements

The following topics describe hardware requirements for XRC:

- ESA/390™ hardware.
- Your ES/9000™ environment must be compatible. The XRC primary system builds on the existing sysplex concept to ensure sequence consistency. You must have a sysplex timer if multiple processing units update XRC primary volumes. The system clock is sufficient in a single processor or virtual server environment.
- The system data mover must have access to the control, state, and journal data sets, and to the primary and secondary copy volumes. If you are going to configure a coupled environment, all system data mover hosts must have access to a shared volume that contains the master data set.
- All host primary systems must have a common time reference. This is necessary so that XRC can provide data consistency for all volume updates across all attached host systems. Host system MVS/DFP software performs timestamping

© Copyright IBM Corp. 1994, 2002

of all application I/O write operations to active XRC volumes. Therefore, all channel programs issued to XRC-managed volumes are timestamped with a common time reference.

Examples: The following are examples of common time references:

- The system time-of-day clock provides the common time reference for environments with either a single processor, or with multiple virtual servers defined. In this case, the minimal MVS level of the primary systems is MVS/SP™ 3.1.0 with MVS/DFP 3.3.0. (ESCON attachment requires MVS 3.1.0e or above.)
- In an environment with multiple processing units, the sysplex timer, or equivalent, provides the common time reference for application programs. XRC and application programs both require a common time reference. The minimum MVS level that supports the IBM 9037 sysplex timer is MVS/ESA 4.1.0 with MVS/DFP 3.3.0.

- A compatible secondary volume must be available for each primary volume you wish to copy. The secondary volume must have the identical track capacity and number of tracks per cylinder and either the same or larger volume capacity.

XRC supported devices

For the **primary disk subsystem**, XRC supports all storage controls that have XRC-capable LIC.

Recommendation: The **secondary disk subsystem** should be equivalent to the primary system storage control, with similar cache and NVS sizes, for the best XRC system performance and recoverability.

XRC operational considerations

This section addresses physical and operational limitations for XRC, including limits on the number of active volumes and sessions, and volume format restrictions.

The following limitations apply to XRC operations:

- Each volume may be an XRC primary in only one session. Each secondary volume can belong to only one XRC volume pair. However, a secondary volume in one XRC session may be a primary volume in another session. A volume cannot be both primary and secondary within the same session.
- Each logical storage subsystem (LSS) in an Enterprise Storage Server (ESS) storage subsystem can support up to 64 XRC sessions. This session limit is a combination of XRC and concurrent copy sessions. The 3990 Storage Control supports up to four XRC sessions per SSID. Each connected MVS system can manage a unique set of volumes behind the storage subsystem.
- You can define a maximum of 256 volumes for each LSS in an ESS storage subsystem. The 3990 Storage Control supports up to 128 volumes.
- Each XRC session can support a maximum of 80 storage control sessions.
- With MXRC, up to five XRC sessions may be running per LPAR.
- All volumes that are part of XRC copy operations must conform to the following **volume format**, **track**, and **access method** restrictions, to ensure data integrity:
  - Volumes must have a standard format for record zero (R0). Volumes with R0 data lengths longer than eight bytes can cause a track format error to remain undetected when the storage control formats the track in cache.
  - You cannot assign alternate tracks within the user area. User data can be overlaid if you assign a user track on the primary address as an alternate track for a secondary address.

  - All storage control Define Extent commands must specify normal access authorization mode. Data written to an XRC primary device while the storage control is in diagnostic or device support mode is *not* copied to the XRC secondary volume. It is therefore important to remove volumes from the XRC session before running a utility program like ICKDSF.
- Both primary and secondary volumes must be varied online to the XRC system when you issue the XADDPAIR command. The system applications must have a common time reference.
- Ensure that only the data mover can update XRC secondary volumes. Do not allow non-XRC applications to update XRC secondary volumes. Non-XRC applications may read from secondary volumes, but to avoid potential data integrity exposures, you must prevent operations from writing to XRC secondary volumes. See “Establishing XRC secondary volumes” on page 26 for more information.
- Assign the secondary volume a different volume serial number than the primary system volume. This will allow both volumes to be online to the system that contains the data mover. The XRECOVER command changes the secondary volume serial number to that of the primary system volume as part of the recovery operation. If you plan to read from secondary volumes during normal copy operations, either use a different catalog or refer to the data sets by explicit volume serial numbers.
- The DFSMS™ address space must already be active by the time the XRC address space becomes active. To obtain accurate System Management Facilities (SMF) accounting, the DFSMS address space must remain active for the duration of the XRC session, as XRC and DFSMS both write SMF records. The system data mover writes SMF type 42 subtype 11 records. See “XRC information in SMF type 42 records” on page 307 for record details.
- You may need to build a channel command word (CCW) chain to correct the lack of timestamping in stand-alone programs and in operating systems that do not support timestamping (bypass start I/O drivers).
- **For VM Systems:** When you operate XRC as a guest under VM, define all XRC volumes to VM as unsupported disks. You cannot define them as minidisks.

  Note: For a disk device that is defined as unsupported under VM, the SET PATH GROUP command is unable to reestablish the path group once a path has been removed.

Introducing the XRC components

This section describes the following major system components in the XRC environment:

- Primary systems
- System data mover
- XRC volumes
- Storage subsystems

XRC primary systems

Primary systems are the host systems where an application runs. XRC can run on the primary system or on any other system that is connected to the XRC-managed primary volumes that the application is updating. If there are multiple primary systems, they must have a common clock reference (a system clock) to maintain data consistency across multiple storage subsystems.

System data mover

The system data mover (SDM) is a DFSMS/MVS® component that interacts with data storage subsystems and with various advanced copy services functions to efficiently move large amounts of data. As updates occur to primary volumes, the SDM manages the process of copying those updates to secondary volumes.

The SDM ensures that updates to secondary volumes are made in the same order in which they were made to the primary volumes, maintaining sequence consistency.

The following topics describe XRC-specific SDM functions, the SDM address spaces, and the XRC sessions descriptions.

Functions that are performed by the system data mover

Table 8 describes the XRC operations and the functions that are performed by the system data mover.

Table 8. Functions that are performed by the system data mover

During this operation . . .	The system data mover . . .
Data shadowing	• Monitors all updates made to the XRC primary volumes. • Transfers the updated records to the journal data sets (only when SESSIONTYPE is XRC). • Maintains the XRC session status on the state data set. • Applies the updates to the XRC-managed secondary volumes while recording the activity in the control data set.
Disaster recovery	• Collects information about XRC session status and unapplied updates from the control, state, and journal data sets, and from the master data set, if the session is coupled. • Applies any remaining consistency groups to the secondary volumes. • Relabels the secondary volumes, and makes them ready for use at the recovery site.
Migration	• Monitors all updates that are made to the XRC primary volumes. • Maintains record status on the state data set. • Applies the updates to the XRC-managed secondary volumes.
Migration recovery	Relabels the secondary volumes and makes them ready for use at the recovery site.

Address spaces for XRC

The system data mover for XRC operates in specific system address spaces that start with “ANT.” The functions that are provided by these address spaces are described in Table 9 on page 23.

Note: The system data mover generates messages that begin with “ANT.” *z/OS MVS System Messages, Vol 1 (ABA-AOM)* describes functions that the address spaces provide.

Table 9. Address space descriptions

Address space	Description
ANTAS000	• This address space handles TSO commands and API requests that control XRC. • ANTAS000 and the point-in-time address space, ANTMAIN, are automatically started during IPL. • The data mover automatically reinitializes the address space if it is canceled.
ANTAS*nnn* refers to: ANTAS001 ANTAS002 ANTAS003 ANTAS004 ANTAS005	An ANTAS*nnn* address space will start for each XSTART command issued.

Notes:

1. Because they process a large number of I/O requests, ANTAS000 and ANTAS*nnn* are unswappable. The ANTAS*nnn* address space runs at the MASTER dispatching priority.
2. If you are running CXRC or MXRC, you may run up to five XRC sessions with address spaces of ANTAS001–ANTAS005 in a single LPAR, with or without coupling them.
3. If the ANTAS000 address space is canceled, XRC automatically restarts it. Except for TSO commands and API requests, there is no impact to data mover functions that run in the ANTAS*nnn* address space. If an event cancels the ANTAS*nnn* address space, XRC ends all active storage control sessions with storage controls that do not support hardware bitmapping.

XRC session state descriptions

An XRC session may be in one of several states of operation. Table 10 describes the different XRC session states.

Table 10. XRC session state descriptions

State of session	Description	For information, see . . .
Active	The data mover has been initiated with an XSTART command. XRC accepts the XADDPAIR, XDELPAIR, XQUERY, XSUSPEND, XSET, and XEND commands.	“Chapter 5. Administering Your XRC Environment” on page 55
Inactive	The data mover has been stopped by an XEND command. The only commands accepted by XRC are the XSTART and XRECOVER commands.	“Chapter 5. Administering Your XRC Environment” on page 55
Recovery	An XRECOVER command has initiated the data mover, which is available to update secondary volumes with unapplied journal records. XRC only accepts the XQUERY command after you have issued the XRECOVER command. The XRC session returns to the inactive mode when the recovery function is complete.	“XRECOVER–recovering data on the recovery system” on page 82

Table 10. XRC session state descriptions (continued)

State of session	Description	For information, see . . .
Suspended	The XSUSPEND(*session_id*) command with the TIMEOUT option has suspended the XRC session. The ANTAS*nnn* address space is no longer active. You can issue an XSTART command to restart the suspended XRC session or an XRECOVER command to recover the session at a recovery site.	"XSUSPEND–suspending volumes or sessions" on page 91

XRC volume terminology

The following terms occur throughout this publication to describe volumes and volume pairs that belong to an XRC session.

XRC volume descriptions

Table 11 describes the volume types.

Table 11. XRC volume descriptions

Volume type . . .	Description . . .
Primary	The source volume that XRC copies to a remote location.
Secondary	The target volume that receives the updates that are made to the primary volume. See "Establishing XRC secondary volumes" on page 26 for more information.
Source	Another name for the primary volume.
Target	Another name for the secondary volume.
Active	A primary volume in an XRC session that receives data that is also copied to a remote location. The status of an active volume can be DUPLEX, SEQCHECK, COPY, or PENDING.
Suspended	An established primary volume in an XRC session whose data XRC is no longer copying to a remote location. The volume remains in the XRC session. See "XSUSPEND–suspending volumes or sessions" on page 91 for more information.
Utility	A primary XRC volume that is available to the system data mover for reading data from a storage control. For more information about utility volumes, see "Using XRC utility devices" on page 102.
XRC	An established primary and secondary volume that is identified by the XADDPAIR command. The status of a volume pair can be COPY, DUPLEX, PENDING SEQCHECK, or SUSPENDED. You have the option to copy an XRC volume pair to a remote location.

XRC volume pair status descriptions

Table 12 describes the status of a volume.

Table 12. XRC volume pair status descriptions

Volume pair status . . .	Description . . .
COPY	A volume pair that XRC is currently synchronizing. Volume pairs in COPY status are not part of the XRC session and, are therefore, not eligible for recovery. XRC does not consider copy volume pairs during ERRORLEVEL processing, nor does it include them during XRECOVER and XADVANCE processing.
DUPLEX	A volume pair for which volume synchronization or resynchronization has completed and that has been successfully recovered, or that was established with the NOCOPY option. The pair is active, and is neither SUSPENDED, PENDING, nor in COPY status. The volume pair is consistent up to the session-specified timestamp reported in message ANTQ8231I.
PENDING	A volume pair that XRC has not yet synchronized or resynchronized. Volume pairs in PENDING status are not part of the XRC session and, are therefore, not eligible for recovery. XRC does not consider a volume pair in PENDING status during ERRORLEVEL processing, nor does it include the volume pair during XRECOVER and XADVANCE command processing.
SEQCHECK	A volume pair that contains nontimestamped writes. This temporary condition indicates that data consistency is uncertain, for this volume pair, in the event of a disaster recovery. The pair automatically returns to the DUPLEX status when XRC successfully writes a timestamped update to the secondary volume.
SUSPENDED	An established volume pair in an XRC session where XRC is no longer copying data from the primary volume to the secondary volume. The volume pair remains in the XRC session.

XRC data sets

This section defines terms for the XRC data sets. An XRC system requires one of each of the state, control and journal data sets. A Coupled XRC system requires a single master data set in addition to each of the above sets per XRC subsystem.

XRC data set descriptions

The following terms apply to the data sets in an XRC session.

Data set type	Description
State data set	The state data set contains status of the XRC session and of associated volumes that XRC is managing. The state data set is updated if an XADDPAIR, XDELPAIR, XSET, XSUSPEND, XRECOVER, or XEND command is issued, or whenever a volume state changes. If a session is coupled, the state data set identifies the master session it is coupled to, and is updated when an XCOUPLE command is issued.

Data set type	Description
Control data set	The control data set contains consistent group information on the secondary volumes and the journal data set. It contains information necessary for recovery operations. The control data set acts as the table of contents for the session. **Example:** The control data set keeps track of data written to secondary volumes, the location of unwritten data in the journal set, and which group to start recovery with.
Journal data set	The journal data set contains temporary user data that is created when records in the primary volume are changed or updated.
Master data set	The master data set ensures recoverable consistency among all XRC subsystems contained within the Coupled XRC system.

XRC storage control sessions

This section defines terms for XRC storage control sessions that are associated with primary volumes in an XRC volume pair. A storage control session is a logical entity that is created for processing updates to the XRC primary volumes. It is used to group sets of primary XRC volumes that are being processed by an XRC session within the storage control.

The number of storage control sessions that may be in effect for a single storage control depends on the capability of the storage control. For example, IBM 3990 and 9390 Storage Controls can each manage a maximum of four XRC storage control sessions. Each logical storage subsystem (LSS) in an Enterprise Storage Server (ESS) storage subsystem can manage a maximum of 64 storage control sessions (XRC or CC) in any combination.

Storage control session states

The following states apply to the XRC storage control sessions for a primary volume in an XRC session.

Session state . . .	Description . . .
Quiesced	A non-ESS storage control session is in quiesced state after an XSUSPEND command has completed. The session remains quiesced until either you issue an XSTART command and the storage control session becomes active, or the quiesce timeout condition expires and the storage control session ends.
Suspended storage control session	An ESS storage control session is in a suspended state after either receiving a system reset or when the XRC session becomes suspended. When XRC stops reading from the ESS, the ESS suspends the storage control session. The session remains suspended until you issue an XADDPAIR command for a volume that is associated with the storage control session.

Establishing XRC secondary volumes

All XRC secondary volumes must stay dedicated exclusively to XRC use. It is recommended that you make the secondary volumes offline and unavailable to other host systems, but not to the data mover, thereby ensuring that *no* write I/O operations occur outside the XRC data mover function. You can do this by making the volumes part of a system-managed, disabled storage group, and by keeping the volumes offline to all but the data mover system. Ensure the data integrity of the

secondary volumes by not allowing other functions to make updates to them. For more information about storage groups, refer to *z/OS DFSMSdfp Storage Administration Reference*.

Secondary volumes may have data read from them while they are in an XRC session. However, if primary system programs change the VTOC pointers, data can appear to have been copied improperly when read from another system. To avoid this, refresh the in-storage system pointers by varying the secondary volumes off and then on to the system that is reading from the volume. Allow the volumes to reach duplex state before doing this refresh operation.

Evaluating remote copy with DFSMShsm™

If any remote copy-eligible volumes are also eligible for either primary space migration or interval migration, you must make backup copies of these migration-eligible data sets available to the disaster recovery site. The need for these backup copies can be better understood when we compare the DFSMShsm™ migration process to the remote copy process.

Note: Although backups are always useful, remote copy (especially XRC) reduces the need for them by keeping the offsite copy far enough away from the original to avoid the loss of both copies. The only need for backups, therefore, is to allow retrieval of older versions of existing data sets.

DFSMShsm migration from a primary volume consists of the following steps:

1. A move or delete from the primary volume

   _____

2. A change in catalog pointers to a DFSMShsm-managed control data set

   _____

3. The reallocation of the data on another DASD device or tape volume

   _____

If you are also copying the primary volume with remote copy, the above DFSMShsm actions appear as a delete of the primary as far as remote copy is concerned. Since the secondary volume is a track-for-track image of the primary volume, the delete performed against the primary volume during the migration process is then reflected to the remote copy secondary volume, which is then deleted. You must take steps to ensure that you are not left without current copies of these data sets.

The best solution is to have policies in place that control how volumes can be removed from or added to remote copy control. One way is to take advantage of DFSMShsm management-class and storage-group SMS constructs, both of which have options to prevent automatic migration of these data sets.

Sometimes data is written to DASD for a very short length of time, as with a temporary work file or when the NOBACKUP keyword is specified. This might also be the case if the volume using remote copy services is being managed under Tape Mount Management (TMM). With TMM migrating the data every hour without backup, you must ensure that a valid backup is available to the remote recovery site.

Recommendation: An excellent way to achieve this backup is to use a fast replication copy function (such as concurrent copy, FlashCopy, or snapshot copy) to copy the secondary volumes before migration.

Ensure that you have adequate cache capacity to accommodate both functions, as the cache resources must now serve both remote copy and concurrent copy. Without adequate policies in place, data sets available for migration or deletion by DFSMShsm cannot be guaranteed to be covered by remote copy.

Defining storage requirements

This section defines the storage requirements that are needed to support XRC.

Virtual storage requirements

XRC operates in up to six address spaces on a single LPAR. ANTAS000 is always active and processes XRC commands. Up to five XRC sessions may be started at a time. Each unique XRC session that is started on a system occupies its own address space, ANTAS001 through ANTAS005. See “Address spaces for XRC” on page 22 for specific information on XRC session address spaces.

The ANTAS000 address space uses 65 MB of virtual storage, which includes a 20 MB trace data space.

Recommendation: Do not restrict the virtual storage that is used by the ANTAS*nnn* address spaces. However, if you need to limit virtual storage usage, the following provides a guideline as to the amount of virtual storage that is needed by each ANTAS*nnn* address space.

An ANTAS*nnn* address space requires a maximum of 100 MB of virtual storage that is used for volume initialization, trace data space, and other non-storage control session specific virtual storage. For each storage control session added to an XRC session using an XADDPAIR command, add an additional 35 MB of virtual storage, up to a maximum of 1.7 GB. The exact amount depends on the number of storage control sessions and the amount of data being moved. If you find that XRC commands are failing due to a virtual storage shortage, you can reduce virtual storage usage by issuing XDELPAIR commands for all volumes in one or more storage control sessions.

Note: The numbers that are provided above for the ANTAS*nnn* address space are relative to running XRC sessions with standard IBM 3390-3 size devices. Larger amounts of virtual storage will be used by XRC when devices with larger numbers of cylinders are added to a session. Also, the update activity of storage control sessions may cause an ANTAS*nnn* address space to obtain virtual storage in a fragmented manner, and therefore cause XRC to use additional virtual storage.

Ensure that local paging space on the data mover system is sufficient to support your planned remote copy configuration.

XRC needs a minimal amount of space (0.2 KB) in the extended common service area (ECSA). All other storage used is within the XRC address spaces.

XRC supports a maximum of 80 storage control sessions for each XRC session.

If your system indicates that a pageable storage shortage exists, the following are ways to manage a pageable storage shortage if it is caused by XRC:

- Increase the real storage available to the system.
- Decrease the fixed storage requirement for XRC using the XSET command with the PAGEFIX option, and a value smaller than the current value. It might be necessary to suspend and restart the session using the XSUSPEND, XSTART, XSET, XADD SUSPENDED commands to receive a timely response to these commands. (XRC might not honor a lower XSET PAGEFIX command very quickly because the system is likely experiencing severe paging overhead.)
- Reduce the number of XRC sessions running on a system by issuing the XSUSPEND command for one or more sessions and then issuing the XSTART command on a different system.
- Reduce the number of storage control sessions XRC is using by issuing XDELPAIR commands for all volumes in a storage control session and then issuing the XADDPAIR command for those volumes to a different XRC session (on a different system) or using the same storage control session name as used for other volumes in the same subsystem id.
- Separate XRC from other applications that use a lot of fixed pages by running it on a different system.

Real storage requirements

XRC uses storage buffers for I/O operations. The real storage consumption is for fixed pages for the duration of input or output. XRC operations require real storage as a function of the number of copied volumes, and the activity level against each volume.

Provide enough real storage on the data mover system to support the projected peak I/O load. The actual amount of real storage that is required at a given point depends on how much concurrent activity is occurring. For example, volume initialization requires up to 1.2 MB of real storage for four concurrent volume copies. Volume update activity can use up to a maximum of 10% more real storage than is specified with the XSET command PAGEFIX limit, up to a limit of 35 MB per primary storage control. All storage beyond the PAGEFIX limit is fixed and freed as required.

Paging space requirements

XRC requires additional paging space to support the trace data space and the two address spaces.

Table 13 shows the data and address space requirements for the two address spaces.

Table 13. Address space requirements

Address space	Data space requirements	Address space requirements
ANTAS000	15 MB	2 MB
ANTAS*nnn*	15 MB	up to 2 GB

Recommendation: Data mover paging should use expanded storage.

See “Address spaces for XRC” on page 22 for more information on address spaces.

Library requirements

The XRC function resides in the following libraries:

- SYS1.NUCLEUS (increases by 0.1 KB)
- SYS1.LINKLIB (increases by 2 MB)
- SYS1.LPALIB (increases by 0.2 KB)
- SYS1.CMDLIB (increases by 100 KB)

Configuring extended remote copy

A critical aspect of recovery system performance (or “target” system performance, for data migration tasks) is the configuration of the full XRC system. This section describes basic XRC configurations, and those that include DASD channel extenders. If you are considering a coupled XRC configuration see“Choosing a coupled XRC configuration” on page 139 for more information before continuing to plan your system.

Choosing an XRC configuration

In XRC configurations, different host systems can share common primary disk volumes. The system data mover collects all application updates and sends them to the secondary volumes.

Figure 5 on page 31 outlines three XRC configuration variations. The major difference between these configurations is the location of the system data mover.

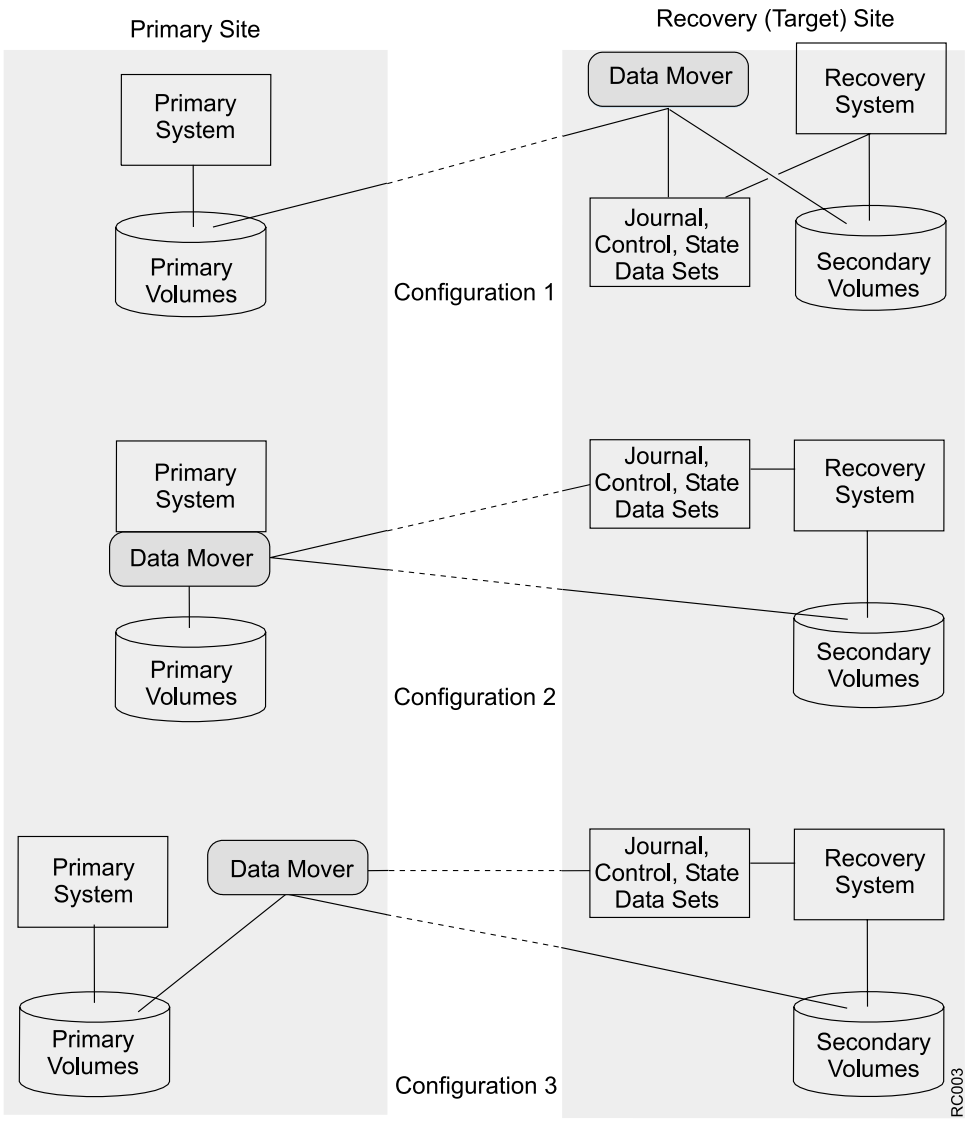

Figure 5. Basic XRC configurations

Each of the configurations above places the data mover at a different location, as explained in Table 14.

Table 14. Data mover locations in basic XRC configurations

Configuration . . .	Location of data mover . . .
Configuration 1	The data mover resides at the recovery site.
Configuration 2	The data mover is on one of the application site host systems.
Configuration 3	The data mover is at the primary site, but not on the primary host system itself.

Recommendation: It is a good idea to keep all primary volumes online to the system on which the data mover is running.

Rule: All secondary volumes must be online when the XADDPAIR command processes. You can then vary the secondary volumes offline to protect them from

updates that originate from non-XRC address spaces, if needed. You can use any other non-XRC volumes on the systems at either site for local applications, as they are independent of XRC activity.

Including DASD channel extenders

DASD channel extenders work in an XRC configuration when the data mover is at either the primary site or the remote site. An XRC configuration with a DASD channel extender operates most efficiently when the data mover software runs on a system at the secondary location. The data mover then accesses the primary disks through the DASD channel extender connection.

There may be situations where you cannot or do not want to have a continuously running host at the secondary site. If so, you can run the data mover on a system at the primary location with the DASD channel extender that is connected to disk at the secondary location. This setup, however, is much less efficient than the previously mentioned configuration.

High quality DASD channel extenders are transparent to the I/O software (except for the performance cost). XRC software can access primary or secondary disks that are configured through extenders. See “Using channel extenders” on page 15 for general information about DASD channel extenders.

Determining XRC resource needs

This section contains information and recommendations about the resources needed to establish an XRC configuration that meets disaster recovery and workload migration needs.

Before installing Extended Remote Copy, review Table 15.

Table 15. Considerations for determining XRC resource needs

Consider . . .	On page . . .
Evaluating how XRC affects the primary storage subsystem	32
Evaluating how XRC affects the recovery storage subsystem	33
Determining the recovery site storage capacity	33
Determining recovery system access to resources	34
Establishing storage subsystem channel connections	34
Establishing XRC ESCON connections	34
Establishing XRC data mover connections	35
Estimating XRC system resources for data mover operations	35
Allocating XRC journal data sets	35

Evaluating how XRC affects the primary storage subsystem

Storage subsystem use can affect overall system performance, and it is therefore necessary to determine how adding XRC activity might affect the primary storage subsystem.

Perform the following steps to determine the effect of adding XRC activity to your primary storage subsystem.

1. Make note of the write rate for the volumes that you want to copy on each primary site storage control.

   _____

2. Add this rate, as if it was all fast write, to the total storage subsystem use.

   _____

3. Decide if this impact on the storage subsystem allows applications to achieve acceptable disk I/O performance. If it does not, you may need to redistribute volumes among storage controls or add an additional primary site storage control.

   _____

Evaluating how XRC affects the recovery storage subsystem

XRC lets you allocate recovery system volumes on any number of recovery storage subsystems. The storage controls at the recovery site do not have to have a one-to-one correlation with their application storage subsystem counterparts. Project the performance impact on recovery storage subsystems to be sure that it is within acceptable limits.

For XRC, recovery site storage subsystems must be able to accommodate the XRC volume write rate. Include all other recovery storage subsystem activity in your assessment, such as access by other applications and systems to non-XRC volumes on the same storage subsystem.

With XRC, you also need to plan for the bandwidth capacity of the journal data sets. The journal data sets, which are either striped or unstriped sequential access method (SAM) data sets, must be able to accept the write rate of copied disk from the primary system. See “Specifying XRC journal, control, and state data sets” on page 42 for more information on XRC journal data sets.

Determining the recovery site storage capacity

As part of planning for XRC installation, you must consider the number and type of disk volumes, according to the steps outlined in Table 16.

Table 16. Steps for disk volume determination

Stage	Description
Determine the number of necessary storage subsystem disk volumes.	Assume that you must copy all volumes from the primary subsystem to the recovery system. You can then consider each volume individually to determine if you will need it for recovery and subsequent application operations on the recovery system. **Example:** You do not have to copy disk volumes that contain only page data sets to the recovery site storage subsystem.

Table 16. Steps for disk volume determination (continued)

Stage	Description
Determine the number of disk volumes to be copied to the recovery site storage subsystem.	Consider which application and support data sets will be needed to start applications on the recovery subsystem at recovery time. **Example:** User catalogs are a good example of such data sets. If the recovery subsystem cannot find data volumes associated with user catalogs, jobs will not start successfully after recovery.
Consider disk volumes.	Consider what disk volumes, such as DFSMShsm migration level 1 data sets, may be needed for recovery system operations.

Determining recovery system access to resources

Rules: To meet disaster recovery needs, recovery systems must obey the following:

- Be able to run the applications that use the copied data.
- Be able to process commands to the secondary volumes to end the secondary copy state and to become available for primary system allocation.
- Have access to the data mover software, TSO commands, and to the state, control, and journal data sets.
- Have access to the CXRC master data set.See “Allocating the CXRC master data set” on page 143 for additional information on host access to the master coupled data set.

Establishing storage subsystem channel connections

For XRC, disk channel connections include those to primary and recovery systems. Connect both primary and recovery system disk devices to the system that contains the data mover. For some storage controls, you can use ESCON, FICON, or parallel channels between hosts and storage controls.

Tip: It may be useful when planning an initial XRC **test** installation (with a limited number of volumes) to have both the primary and secondary volumes on the same storage subsystem, thereby conserving resources.

Establishing XRC ESCON connections

You can configure IBM storage controls to support ESCON ports (connected by multimode fiber-optic cables). ESCON Directors, if your configuration includes them, can have Extended Distance Feature (XDF) ports with singlemode fiber-optic cables for longer link lengths.

You can only use 9036 ESCON Remote Channel Extenders to convert from multimode fiber links to XDF and back, and as XDF extender links. Use these or other comparable repeaters when the distance between systems is greater than 3 km and your configuration does not include ESCON Directors.

You can only dynamically switch one ESCON Director in a path to an I/O unit. The other ESCON Director, if present, provides a static connection. Refer to the appropriate configuration planning documentation for information about ESCON configurations with ESCON Directors.

Rule: Recovery site storage subsystems must have access to the host systems that will run primary applications when the primary system is out of service. A host system at the recovery site must be able to issue recovery commands as part of recovery system takeover procedures.

Establishing XRC data mover connections

All primary and secondary disk subsystems require connection to the system that runs the data mover software.

Note: Some storage controls allow you to use parallel channels instead of ESCON channels for the XRC disk devices that run the asynchronous software data mover. Parallel channel connections, however, may not provide the bandwidth under load or sufficient distance for disaster protection.

Estimating XRC system resources for data mover operations

The XRC data mover must have sufficient I/O transfer bandwidth and have access to both primary and recovery system disk devices. A good initial estimate includes enough system resources to process approximately three times the application I/O write rate to volumes that will be XRC primary volumes. Actual resources should be less since the data mover performs much of its I/O in more efficient multiple-I/O-transfer channel programs.

You can specify an I/O dispatching priority of either IOQ=PRTY or IOQ=FIFO in the IEAIPSxx member of the SYS1.PARMLIB. Specifying IOQ=PRTY ensures that the data mover operates at the highest priority.

Allocating XRC journal data sets

Include disk allocation for XRC journal data sets in your resource planning. When allocating journal data sets to disk volumes, consider the total, concurrent data transfer demand of all XRC writes. For very high total write rate systems, you may have to assign journal data set volumes to multiple storage controls.

Allocating a striped SAM data set also can be worthwhile. Verify that there is sufficient recovery storage control capacity when both journal data sets and secondary copy volumes are present on any single storage control. For more information, see “Specifying the journal data sets” on page 43 and “Maximizing access to journal data sets” on page 52.

Protecting access to XRC commands

Each installation is responsible for preventing the unauthorized use of XRC TSO commands and for controlling ICKDSF activity to XRC-managed volumes. The following are two ways to protect access to XRC commands:

- Put the commands in a library that is protected by RACF® (a component of the SecureWay® Security Server for z/OS) (See “Placing TSO commands in an RACF-protected library” on page 40 for more instructions on placing commands in a protected library.)**or**
- Define resource profiles in the RACF FACILITY class and restrict access to those profiles. (See “Defining resource profiles in the RACF FACILITY class” on page 40 for instructions on defining resource profiles.)

Chapter 4. Setting Up the XRC Environment

This section describes how to set up the extended remote copy environment. It explains the system components that interact with XRC, and describes how to do the following:

Note: If you are considering a CXRC environment, see “Chapter 7. Managing Coupled XRC Sessions” on page 137 before continuing with this chapter.

In this chapter

This chapter is intended to help you setup and install extended remote copy. It includes the following sections:

Topic . . .	See page . . .
Installing and configuring XRC	37
Virtual storage messages	39
Controlling access to XRC resources	39
Controlling access to XRC commands	40
Including XRC TSO commands in automated procedures	41
Specifying XRC journal, control, and state data sets	42
Copying the catalog and control data sets	47
Specifying XRC utility volumes	48
Identifying volume pairs	48
Optimizing XRC performance	49

Installing and configuring XRC

To install and configure the extended remote copy support, perform the following steps:

1. Install the appropriate z/OS release and the maintenance that is specified in the Copy Services web site at:http://www.storage ibm.com/software/sms/sdm/index.html

   You can also find related information under Technical Updates at the same web site.

2. Authorize XRC TSO commands by adding the command names to the AUTHCMD PARM parameter of the IKJTSO*xx* member of SYS1.PARMLIB. See “Chapter 5. Administering Your XRC Environment” on page 55 for specific XRC TSO commands. After adding the XRC command names to the IKJTSO*xx* member, issue the TSO command PARMLIB UPDATE(*xx*) to activate the new IKJTSO*xx* member.

   To restrict XRC commands to certain users, follow the instructions in “Controlling access to XRC commands” on page 40.

3. Based on the application or applications that you plan to copy, determine all primary (source) and secondary (target) devices.

© Copyright IBM Corp. 1994, 2002

4. Allocate the appropriate journal, control, and state data sets on disk devices that have connectivity to the data mover and either the recovery or the target migration system. You can change the default SYS1 high-level qualifier (HLQ) in any of the examples in Table 17 to a name that corresponds with the HLQ issued on the XSTART command:

Table 17. Data Set Allocation Command Examples

Command . . .	**Notes . . .**
SYS1.XCOPY.*session_id*.CONTROL	This is only required with SESSIONTYPE(XRC).
SYS1.XCOPY.*session_id*.STATE	
SYS1.XCOPY.*session_id*.JRNL01	This is only required with SESSIONTYPE(XRC).
SYS1.XCOPY.*session_id*.JRNL*xx*	This is only required with SESSIONTYPE(XRC). A minimum of two journals are required (JRNL01 and JRNL02), up to JRNL16.

Note: Ensure that the ANTAS*nnn* address space has update authority to these data sets. See “Address spaces for XRC” on page 22 for more information on address spaces.

5. Issue the XSTART TSO command on the system that contains the data mover. Specify the session ID that is associated with the journal data set names. Select the proper level of recovery for your environment. See “XSTART–starting a session” on page 89.

6. Issue the XSET commands on the data mover system to tune the environment as you desire. See “XSET–changing session parameters” on page 84.

7. Issue the XADDPAIR TSO commands on the system data mover to perform the copy of primary to secondary volumes. See “XADDPAIR–adding volume pairs or utility volumes” on page 56.

8. Issue the XQUERY TSO command on the data mover’s system to verify the XRC installation options currently in effect. See “XQUERY–querying a session” on page 75.

9. If you are using XRC for disk devices migration, follow the disk migration scenario that is described in “Chapter 9. Migrating Data with XRC” on page 171.

10. If you are using XRC for disaster recovery, XRC continues to copy updates from the primary volumes to the secondary volumes for as long as the system is running, or until you issue a command to end the volume copies.

Virtual storage messages

If you receive pageable-storage shortage messages, the data mover system is constrained by lack of real or virtual storage. In this case, either reduce the number of primary storage controls being managed by the data mover or increase the amount of real storage or paging space for the ANTAS*nnn* address space. See “Address spaces for XRC” on page 22 for more information on address spaces.

Controlling access to XRC resources

This section describes how to protect XRC commands and control volume repair activities.

Controlling ICKDSF activity to XRC volumes

The storage administrator must control ICKDSF activity to primary and secondary XRC-managed volumes.

Note: XRC does not prevent users from varying either the primary or secondary volumes offline after the volume initialization is complete. The data mover continues to read record updates, and then applies them to the secondary volumes, regardless of whether the volumes are online to the data mover system.

XRC does not always copy ICKDSF operations from the primary to secondary volumes. If you use ICKDSF to change the VTOC or volume serial number, or repair either a primary or secondary volume in an XRC session, you must completely recopy the volume pair with the XADDPAIR command.

XRC only copies changes that are made with normal access authorization. XRC cannot manage ICKDSF operations that are not performed within normal access authorization mode.

Proceed as follows to perform an ICKDSF change or repair on a remote copy-managed volume:

1. Issue the XDELPAIR command to delete the volume pair from the remote copy session.

   _____

2. Run ICKDSF to change or repair the volume.

   _____

3. Issue the XADDPAIR command, with the default copy option, to add the repaired volume pair to the remote copy session. The default copy option ensures that XRC correctly initializes the volume.

   _____

Attention: Failure to follow the above procedure when repairing volumes with ICKDSF can result in potential data integrity exposures to offline volumes.

Do not reconfigure a device via the MVS Hardware Configuration Definition™ (HCD) while that device is a member (primary or secondary) of an XRC pair. If you must use ICKDSF on either volume of an XRC pair, you must first remove the pair from the session with the XDELPAIR command.

Controlling access to XRC commands

The following are two ways to protect access to XRC commands:

- Put the commands in a library that is protected by RACF® (a component of the SecureWay® Security Server for z/OS) **or**
- Define resource profiles in the RACF FACILITY class and restrict access to those profiles.

Placing TSO commands in an RACF-protected library

Place the TSO commands in a RACF-protected library to restrict XRC TSO commands to authorized storage administrators.

To RACF-protect XRC commands, perform the following steps:

1. Issue the following RDEFINE command for each XRC command, and for each command abbreviation that you want defined to RACF:

   ```
   RDEFINE PROGRAM cmdname ADDMEM('SYS1.CMDLIB')/volser/NOPADCHK) UACC(NONE)
   ```

   The following terms apply to the above example:

   cmdname Defines the XRC TSO command name or an abbreviation of a command. Issue a separate RDEFINE command for each of the XRC commands, and any command abbreviations you plan to use. Examples of XRC command abbreviations are XADD, XDEL, and XRCV. RACF can only perform checking on commands and abbreviations that are defined to it.

   volser Defines the name of the volume that contain the SYS1.CMDLIB data set.

2. Issue the PERMIT command for all commands and authorized XRC TSO command users as follows:

   ```
   PERMIT cmdname CLASS(PROGRAM) ID(name) ACCESS(READ)
   ```

   The following terms apply to the above example:

   cmdname Defines the XRC TSO command name, or an abbreviation of a command.

   name Defines the user ID receiving RACF access authority for that command name.

3. Issue the SETROPTS command from a user ID that has the appropriate authority:

   ```
   SETROPTS CLASSACT(PROGRAM) WHEN(PROGRAM) REFRESH
   ```

Defining resource profiles in the RACF FACILITY class

You can limit the use of XRC commands by defining resource profiles in the RACF FACILITY class and restricting access to those profiles. To use a protected command, you need read-access authority to the applicable profile.

Table 18 lists the XRC commands and the FACILITY class profiles that can restrict them. See *z/OS Security Server RACF Security Administrator's Guide* for details on activating the RACF FACILITY class and defining and authorizing users to the XRC command profiles.

Table 18. XRC FACILITY class profile names

Command	Profile name
FCESTABL	STGADMIN.ANT.XRC.COMMANDS
FCWITHDR	
XADDPAIR	
XADVANCE	
XCOUPLE	
XDELPAIR	
XEND	
XQUERY	
XRECOVER	
XSET	
XSTART	
XSUSPEND	
FCQUERY	STGADMIN.ANT.XRC.XQUERY
XQUERY	
Note: Authorize FCQUERY and XQUERY command use with the STGADMIN.ANT.XRC.COMMANDS profile or the STGADMIN.ANT.XRC.XQUERY profile. XRC first checks STGADMIN.ANT.XRC.COMMANDS for authorization. If authorization is not permitted with the STGADMIN.ANT.XRC.COMMANDS profile, XRC checks the STGADMIN.ANT.XRC.XQUERY profile for authorization to issue the FCQUERY and XQUERY commands.	

Examples: The following examples activate the RACF FACILITY class, define the profile for the XRC commands, and give user STGADMIN authority to use this profile.

- Activate the RACF FACILITY class:

```
SETROPTS CLASSACT(FACILITY)
```

- Define the profile for XRC commands, and authorize user STGADMIN to use this profile:

```
RDEFINE FACILITY STGADMIN.ANT.XRC.COMMANDS UACC(NONE)

PERMIT STGADMIN.ANT.XRC.COMMANDS CLASS(FACILITY) -
       ID(STGADMIN) ACCESS(READ)
```

Including XRC TSO commands in automated procedures

You can include XRC TSO commands in automated procedures, such as REXX execs. Within the automated procedure, use the MSGROUTEID parameter to specify the TSO user ID to which you want XRC to send messages.

XRC routes operational messages to the user ID that issued the command. In order to have these messages appear in the SYSLOG, ensure that the SYS1.PARMLIB CONSOLxx member permits “Write to Programmer” on the ROUTCODE parameter.

Specifying XRC journal, control, and state data sets

This section describes the XRC journal, control, and state data sets. These data sets can reside on any device, and can be attached to any storage control that supports the sequential access method (SAM) and the partitioned data set extended (PDSE) formats.

If these data sets are Resource Access Control Facility (RACF) protected, ensure that the ANTAS*nnn* address space, which updates them, has ALTER access authority on both the data mover and recovery systems. See “Address spaces for XRC” on page 22 for more information on address spaces. RACF protects these vital data sets and ensures that no changes are made to them whenever XRC is not active.

Before activating XRC, allocate and catalog the necessary support data sets on volumes that both the system data mover and the recovery system host can access. Table 19 shows which data sets should be allocated with each command.

Table 19. Data set allocation guide

For . . .	Allocate . . .
a SESSIONTYPE of MIGRATE	the state data set.
a SESSIONTYPE of XRC	the journal, control, and state data sets.
CXRC	the master data set.

Note: Because the control and state data sets must be SMS-managed, use the IDCAMS recatalog function when you activate the recovery system if these data sets do not already reside there.

Guidelines: Follow these guidelines when defining journal, control, and state data sets.

- The state and control data sets can reside on the same volume, but neither should be on any volume that also contains journal data sets. Because the performance of the journal data set and control data set is so important, be sure to follow the guidelines that are described in “Maximizing access to journal data sets” on page 52.
- To avoid a potential deadlock situation, place the XRC state, control, and journal data sets in a user catalog that contains only entries for these data sets. If you plan to copy the volume that contains these data sets, ensure that the volume does not also contain data sets that application programs are updating. Either of these conditions can also cause a deadlock condition when XRC sessions are suspended or restarted.

XRC cannot dynamically accept changes to the journal, state, or control data sets. If you wish to change these data sets, you must follow the instructions listed in Table 20 on page 43.

Table 20. Instructions for changing journal, control, and state data sets

To change . . .	Choose . . .	Follow these directions . . .
Journal or control data sets		1. Suspend the session with the XSUSPEND command using the TIMEOUT option. 2. Reallocate the data sets. 3. Copy the data set members to the new data sets. 4. Delete the old data sets and rename the new data sets to the old data set names. 5. Reissue the XSTART command and subsequent commands.
State data sets	to end the session and reallocate the data set. **or**	1. Issue the XEND command. 2. Reallocate the state data set. 3. Reissue the XSTART and subsequent commands.
	to suspend the session and allocate a new data set.	1. Issue the XSUSPEND command. 2. Allocate a new state data set. 3. Copy the suspended state data set members into the new state data set. 4. Delete the original state data set. 5. Rename the new state data set to the original state data set name. 6. Restart the session by issuing XSTART and subsequent commands.

Specifying the journal data sets

Journal data sets contain checkpoint records of the changes that are made, and are read only when you issue an XRECOVER command.

XRC uses a journal data set name in the form:

SYS1.XCOPY.*session_id***.JRNL***nn* or *hlq*.**XCOPY.***session_id***.JRNL***nn*

Example: The following is an example of a journal data set name:

```
SYS1.XCOPY.DALLAS.JRNL01
```

Required specifications

The following requirements are necessary when you specify journal data sets:

- Allocate journal data sets on disk volumes. Each journal data set must be a fixed-block SAM data set with the following attributes:

  ```
  DCB=(RECFM=FB,LRECL=7680,BLKSIZE=7680,DSORG=PS),SPACE=(TRKS,(1200,0))
  ```

  Note: If the LRECL and BLKSIZE are not set to 7680, the data mover will attempt to force the allocation to this size.
- Specify at least two journal data sets. If you allocate less than the maximum number of journal data sets, XRC issues the following message:

  ```
  DATA SET SYS1.XCOPY.session_id.JRNLnn NOT IN CATALOG OR CATALOG
   CANNOT BE ACCESSED
  ```

  A high-level qualifier name can replace SYS1 in the example. XRC issues the above message for the first journal data set that is not found, even if there are two or more journal data sets allocated.
- Do not use compression algorithms on journal data sets.
- Observe the following when naming journal data sets:
  - Specify up to eight alphanumeric characters for the *session_id* portion of the journal data set name. Use the same name that you will specify with the *session_id* parameter of the XSTART command.
  - Name the journal data sets contiguously with a number, *nn*, beginning at 01 up to a maximum of 16.

    Example: If you define journal data sets 01, 02, and 04, XRC ignores data set 04 because you did not define data set 03.

Recommended specifications

Use the following steps as guidelines to help to improve the efficiency of your XRC operations:

- Allocate journal data sets as single-extent data sets. If secondary extents are used, the data set must be filled prior to starting the XRC session if you want XRC to use the full amount of space that you have allocated for journal use.
- Define an even number of journal data sets, preferably four to eight for smaller XRC configurations (up to about 750 volumes). For large XRC configurations, define from eight to 16 journal data sets.
- Define the journal data sets as striped SAM data sets.

  When you define the journal data sets as striped SAM data sets on different volumes, you enable the system data mover to spread its I/O operations across multiple volumes and storage controls. Spreading the work load reduces contention for the same physical device. If you use nonstriped journal data sets, each data set must reside on a single volume.
- Allocate a minimum of 100 cylinders for each journal data set.
- Make journal data sets large enough to allow the system data mover to offload data to secondary volumes during peak data load periods. See Note.
- Do not put the control data set on any of the journal data set volumes.
- Minimize non-XRC activity to the journal data set volumes.

Note: The maximum total journal data set size is 2 GB. Allocating more than 1.5 GB of space for all journal data sets, however, does not provide any additional performance benefit.

Optimal performance specifications

The following steps are recommended to achieve optimal journal performance:

- Allocate the journal data sets on high-performance volumes, and place them behind storage controls that use DASD fast write, cache functions, and have large NVSs.
- Attach the journal data sets with high bandwidth ESCON or FICON channels, preferably capable of very high transfer rates.
- Allocate the same amount of space for all journal data sets.
- Put the journal data set volumes behind storage controls that do not also process primary or secondary volumes.
- Spread the striped journal data sets across as many disk volumes as possible, preferably with each journal data set on a separate volume. Processing performance is improved because XRC assigns a separate internal subtask for each pair of allocated journal data sets.

Table 21 lists journal data set patterns that optimize journal performance.

Table 21. Journal data set patterns for optimal performance

Striped data sets	Nonstriped data sets
For two journal data sets:	
VOL1 VOL2 JRNL01 JRNL01 JRNL02 JRNL02	VOL1 JRNL01 JRNL02
For four journal data sets:	
VOL1 VOL2 VOL3 VOL4 JRNL01 JRNL01 JRNL03 JRNL03 JRNL02 JRNL02 JRNL04 JRNL04	VOL1 VOL2 JRNL01 JRNL03 JRNL02 JRNL04
For eight journal data sets:	
VOL1 VOL2 VOL3 VOL4 JRNL01 JRNL01 JRNL03 JRNL03 JRNL02 JRNL02 JRNL04 JRNL04 VOL5 VOL6 VOL7 VOL8 JRNL05 JRNL05 JRNL07 JRNL07 JRNL06 JRNL06 JRNL08 JRNL08	VOL1 VOL2 JRNL01 JRNL03 JRNL02 JRNL04 VOL3 VOL4 JRNL05 JRNL07 JRNL06 JRNL08
For 16 journal data sets:	
VOL1 VOL2 VOL3 VOL4 JRNL01 JRNL01 JRNL03 JRNL03 JRNL02 JRNL02 JRNL04 JRNL04 VOL5 VOL6 VOL7 VOL8 JRNL05 JRNL05 JRNL07 JRNL07 JRNL06 JRNL06 JRNL08 JRNL08 VOL9 VOL10 VOL11 VOL12 JRNL09 JRNL09 JRNL11 JRNL11 JRNL10 JRNL10 JRNL12 JRNL12 VOL13 VOL14 VOL15 VOL16 JRNL13 JRNL13 JRNL15 JRNL15 JRNL14 JRNL14 JRNL16 JRNL16	VOL1 VOL2 JRNL01 JRNL03 JRNL02 JRNL04 VOL3 VOL4 JRNL05 JRNL07 JRNL06 JRNL08 VOL5 VOL6 JRNL09 JRNL11 JRNL10 JRNL12 VOL7 VOL8 JRNL13 JRNL15 JRNL14 JRNL16

See “Maximizing access to journal data sets” on page 52 for more information about optimizing XRC performance.

Specifying the control data set

The recovery operation uses the control data set to determine what data XRC must still write to the secondary volumes. The control data set contains control records that indicate:

- The last set of data that is written to the secondary volumes
- The amount of unwritten data that exists in the journal
- The location of this unwritten data

XRC uses a control data set name in one of the following forms:

SYS1.XCOPY.*session_id***.CONTROL** or *hlq***.XCOPY.***session_id***.CONTROL**

The XRC session ID can be up to eight characters long and must be the same name that you will specify with the *session_id* parameter of the XSTART command.

Example: The following is an example of a control data set name:

```
SYS1.XCOPY.DALLAS.CONTROL
```

Guidelines: Use the following guidelines to allocate the control data set:

- Place the control data set on a different volume from the journal data sets.
- Allocate the control data set without defining secondary extents.
- Allocate the control data set in one of two ways:
  - Choose Sequential Allocation — Sequential allocation provides much higher performance than PDSE allocation, and is the recommended choice. The data mover function detects which type of control data set you have allocated and supports the specified data set type.

    You can allocate the control data set on a disk device to be a high-performance sequential data set. As such, it will be a standard PS data set with the following attributes:

    ```
    DCB=(RECFM=FB,LRECL=15360,BLKSIZE=15360,DSORG=PS),SPACE=(TRKS,(2,0))
    ```

    If you define the control data set as a PS data set, you should allocate two tracks.
  - Choose PDSE Allocation — You can allocate the control data set on disk as an SMS-managed, partitioned data set extended (PDSE) data set. PDSE allocation is not recommended, however, due to the higher performance with sequential allocation. If you choose PDSE, define the data set with the following attributes:

    ```
    DCB=(RECFM=FB,LRECL=4096,BLKSIZE=4096,DSORG=PO),DSNTYPE=LIBRARY
    ```

    If you define the control data set as a PDSE, you should allocate approximately 480 KB of disk space. This is sufficient space regardless of the size of the environment.
- Some environments can benefit from specifying the STORCLAS parameter to ensure that SMS allocates the PDSE control data set to the proper SMS-managed volume.

Specifying the state data set

The state data set defines and contains the status of the XRC session and of associated volumes pairs that XRC is currently managing. The state data set is updated whenever an XADDPAIR, XDELPAIR, XSET, XSUSPEND, XRECOVER, or XEND command is issued, or whenever a volume state changes.

XRC uses a state data set name in one of the following forms:

SYS1.XCOPY.*session_id***.STATE** or *hlq***.XCOPY.***session_id***.STATE**

Rule: The XRC session ID can be up to eight characters long and must be the same name that you will specify with the *session_id* parameter of the XSTART command.

Example: The following is an example of a state data set name:

```
SYS1.XCOPY.DALLAS.STATE
```

Guidelines: The following are guidelines for the state data set:

- Allocate the state data set on disk as an SMS-managed partitioned data set extended (PDSE) data set with the following attributes:

  ```
  DCB=(RECFM=FB,LRECL=4096,BLKSIZE=4096,DSORG=PO),DSNTYPE=LIBRARY
  ```

- Allocate ten tracks per storage control session. Try to plan for expected future growth when you initially allocate the state data set, as it may be inconvenient to reallocate it later.
- Allocate one track for each volume pair in the storage control session, again planning for expected future growth. For volumes that have a larger capacity than a 3390 Model 3 (3339 cylinders), allocate one additional track for every 3000 additional cylinders. Because XRC posts an error if the state data set runs out of space, ensure that the allocation size is sufficient to prevent this situation. You cannot add additional storage control sessions or volume pairs if there is insufficient space in the state data set, nor will you be able to couple a session to a master session.

  Perform the following steps to allocate a new, larger data set:

  1. Issue XSUSPEND *session_id* TIMEOUT(*hh.mm.ss*)

     _____

  2. Copy all members from existing state data set to the new, larger data set.

     _____

  3. Rename the existing state data set to an unused name.

     _____

  4. Rename the new, larger data set to the pre-existing state data set name.

     _____

  5. Issue XSTART and resume XRC operations.

     _____

  See “Chapter 5. Administering Your XRC Environment” on page 55 for the syntax for specific commands.

Note: Some environments can benefit from specifying a STORCLAS parameter to ensure that SMS allocates the state data set to the proper SMS-managed volume.

Copying the catalog and control data sets

You can ensure faster disaster recovery if you use remote copy functions to copy volumes that contain the master catalog, key user catalogs, and system control data sets to the recovery system.

Remote copy automatically copies all data sets on the managed volumes. Managed volumes can include catalog data sets and control data sets. XRC can copy *any* data set type, which includes JES spool, JES checkpoint, system volumes, and paging volumes.

Specifying XRC utility volumes

XRC utility volumes are specified according to the configuration of the storage subsystem. See Table 22 when specifying XRC utility volumes. For more information, see “Using XRC utility devices” on page 102.

Table 22. Decision process for utility volume specification

If . . .	Then . . .	Notes . . .
Operating with non-IBM ESS storage subsystems:	Specify at least one (ideally more than one) low-activity volume for each application (primary) storage control session as a utility volume (or volumes) for the data mover.	• If multiple volumes exist in the storage control session, they need not be dedicated, as the data mover switches to the least-busy volume at different times. XRC data mover performance improves when you specify several candidate volumes that have little or no activity. This, in turn, improves the application I/O performance to volumes in XRC sessions. • See “Using XRC utility devices” on page 102 for more information.
Operating with the ESS storage subsystem:	You have the option to define a single-cylinder volume for specific use as a utility volume.	See “Using XRC utility devices” on page 102 for more information.

Identifying volume pairs

The storage administrator must identify the application volumes to copy, and create extended remote copy volume pairs. XRC provides an image copy of a volume on a record-for-record basis. There is a one-to-one correspondence between the record on the primary volume track and the record on the secondary volume track. Because the data written on the primary disk is written to the same tracks on the secondary disk, the secondary disk must have the same track size and number of tracks per cylinder, and the same or larger volume capacity, as the primary. (If these criteria are not met, the XADDPAIR command will fail.) Ensure that secondary volumes are part of a storage group that does not allow allocations, and dedicate these volumes to XRC use only. For more information on the XADDPAIR command, see “XADDPAIR–adding volume pairs or utility volumes” on page 56.

Remote copy supports volumes, not data sets. As a result, all data sets on the volumes that are copied are part of remote copy activity, and are therefore copied to the target volume. This support is application (IBM or non-IBM) independent, and supports all data set types.

Because applications deal with data sets and not volumes, multivolume data sets require special attention. Multivolume data set types include data sets that reside on multiple volumes, striped data sets, and VSAM spheres. Unless you copy all volumes of a multivolume data set, XRC will copy only part of the data set. Other data sets on the copied volume may be usable, but the multivolume data sets will not be. If multivolume data sets are critical for recovery, you must copy all volumes on which these data sets reside. In addition, if these data sets are critical to the system, start the XRC session by issuing the XSTART or XADDPAIR command with the ERRORLEVEL(SESSION) or (*group_name*) parameter specified.

In the event of an error, specifying ERRORLEVEL(SESSION) ensures data consistency across all volume pairs in the entire XRC session. Specifying ERRORLEVEL(*group_name*) ensures that data is consistent across all volume pairs that belong to the specified group. Errors within the specified group do not affect the update activity of volume pairs that belong to other groups.

If the data sets require consistency as a group, issue the XADDPAIR command for these related data sets. Specify the ERRORLEVEL parameter with a group name that is meaningful for the related volumes. Specified as such, an error for one volume pair would suspend the entire set of volumes that are associated with the named group. After correcting the error, you can add the volume group back to the XRC session.

For more ERRORLEVEL parameter information, see “Chapter 10. Recovering from Error Conditions Using XRC” on page 175.

Optimizing XRC performance

The information in this section adds to the basic guidelines that are presented in “Specifying XRC journal, control, and state data sets” on page 42. See your IBM marketing representative for the latest remote copy performance and configuration information and recommendations for your specific environment.

Ideally, for the best XRC system performance and recoverability, the recovery system would be equivalent to the primary system. In reality, you may only need to back up a subset of your primary system as a result of a disaster. The following sections outline additional ways to optimize overall XRC system performance.

Configuring for high performance

Extended remote copy works together with cached storage controls to efficiently manage existing system resources for disaster recovery. The system data mover issues I/Os to drain record updates from the cache that are designated to be copied by XRC. Each logical XRC session may have one or more storage control sessions for each primary storage control. A logical XRC session manages all volumes that are associated with those storage control readers.

High performance can be achieved by ensuring that the storage subsystems have adequate resources, including cache, NVS, and path connections to handle the peak workload that is associated with remote copy sessions (and coexisting concurrent copy operations). The system data mover software requires adequate processor MIPs, multiple processors, and real storage to accomplish its function in an efficient and timely fashion. The data mover, without adequate resources, may not be able to drain the storage control cache rapidly enough under peak workload

conditions. Based on the workload, if this continues for an extended period of time, the cache could become overcommitted, and ultimately this could affect the performance of your primary systems.

Recommendations for controlling the storage control cache

Use the recommendations listed in this section as guidelines to help reduce the likelihood of overcommitting the storage control cache:

Define one or more separate, low-activity volumes: Defining one or more separate, low-activity volumes specifically for data mover I/O activity is helpful for storage subsystems that do not have the functional capabilities of the ESS storage subsystem. By having a dedicated reader volume, the data mover can avoid conflicts with application programs that share a utility volume. Sharing a utility volume can negatively impact the performance of both the data mover and the application program.

The XRC data mover software issues its I/O to the XRC utility volume that is associated with each primary storage control. This is part of the data transfer process. You may assign one or more XRC utility volumes to each primary storage control by specifying the SCSESSION keyword on the XADDPAIR command. Each specified unique SCSESSION will start a separate reader. All volumes in an LSS (or SSID) with the same SCSESSION share the same reader. See “Using XRC utility devices” on page 102 for more information.

Assign multiple storage control sessions: Assign multiple storage control sessions with the SCSESSION keyword on the XADDPAIR command. Each storage control session has an independent data mover task that is committed to it. The data mover is thereby capable of reading concurrently from each of the storage control sessions. Although it is possible to combine independent applications into a single session, this increases the data mover’s workload. Providing separate storage control sessions for each independent application achieves the best performance.

Table 23 provides a summary of XRC session configurations.

Table 23. Quick reference for XRC session limits

You can configure up to . . .	**For each . . .**
4 XRC sessions, and up to 64 combined XRC and CC sessions	Primary site 3990 or 9390 Storage Control
64 combined XRC and CC sessions	Primary site ESS logical storage subsystem
256 volumes	ESS logical storage subsystem
128 volumes	3990 or 9390 Storage Control
80 storage control sessions	XRC session
1 XRC session, and up to 16 combined XRC and CC sessions	volume
5 XRC sessions (coupled or uncoupled)	LPAR
14 XRC sessions	master CXRC session

XRC maintains update sequence integrity among the volumes that you assign as part of an MVS XRC storage control session. The data mover, however, cannot maintain update sequence integrity between different XRC sessions unless the different sessions are coupled to the same master session. Therefore, if you are not

running CXRC, keep all volumes associated with an application that requires update sequence integrity within the same MVS XRC session.

Use caution when combining both XRC and Concurrent copy operations on the same storage control: Use caution when combining both XRC and concurrent copy operations on the same storage control for storage controls that do not support hardware bitmapping.

Concurrent copy and XRC both use storage control cache. If you run both simultaneously, you could deplete the cache resources on storage controls that do not support hardware bitmapping. In this case, the storage subsystem may cancel concurrent copy activity. You may instead want to run concurrent copy during low XRC update activity periods, or else run concurrent copy against the secondary volumes of the XRC pair.

Increase the storage control cache size: You can help manage the peak write activity load for XRC-managed volumes by increasing the size of the storage control cache. Due to the large volume of data being moved, storage controls that support XRC may require significantly more cache memory than do those without this ability. The peak load for the storage control dictates the amount of cache that is required. Storage controls require more cache to prevent canceled XRC sessions as the peak load becomes heavier.

Depending on your application write I/O rate, the primary storage subsystem may require additional cache. Sophisticated applications with high write content can require as much as 6 GB or more of subsystem cache. The actual requirements depend on the application, and on the capability of the storage subsystem.

Increase the data mover's transfer performance: Increase the transfer performance by connecting it to the primary storage control with the highest bandwidth channels. If you use channel extenders, ensure that the compression feature is installed, since this will typically double the bandwidth.

Optimize the recovery system setup: Allocate the XRC control, state, and journal data sets, and the secondary volumes behind storage controls that have large cache and NVS storages. This configuration, in addition to speeding the XRC copy process, also provides a powerful platform for the recovery operation. For the best setup, configure the secondary site storage controls to have at least the same cache and NVS capacity as the primary site storage controls. In addition, if you have the journal data sets on an ESS, you may want to place them on volumes that utilize parallel access volumes (PAVs), because the journal data sets will take advantage of this capability.

Note: If the data mover abruptly ends with an ANTX5118E message, the most likely reason is that the I/O rate of the primary subsystem has exceeded the capacity of the subsystem at the receiving end.

Provide a stand-alone data mover system or LPAR: If possible, place the data mover on its own system with a large amount of real storage. This ideal setup allows the data mover to allocate a large number of track-size buffers and to optimize its I/O performance with minimal impact to primary application I/O. A separate system also minimizes the impact to the data mover, eliminating a potential bottleneck due to capacity problems during peak I/O update times.

Assign data mover fixed pages: To take advantage of the available real storage, issue the XSET PAGEFIX command to assign the data mover 35 MB per defined

reader. In addition, keep the pages fixed by setting the tuning table at offset X'b2' to X'80'. You may accomplish this by issuing the following console command after an XSTART command has completed:

```
F ANTASnnn,P .XMTUN+B2  x'00'  x'80'
```

Provide adequate host resources for the data mover: XRC copy operations rely on host processing resources. Therefore, analyze the host resources that are necessary for your applications. Ensure that the data mover system has the necessary processing capacity required to maintain optimal performance. This includes defining processors with the appropriate MIP capability (about 4-6 MIPs per 100 write updates) and multiprocessor capability (2-4 processors per data mover).

Balancing storage control configurations

Configure your system so that primary system activity does not exceed the capacity of the recovery site storage control.

XRC supports configurations where many primary site storage controls can funnel their updated data to a single storage control on the recovery system. However, you must consider the number of copied primary volumes and the rate of write activity to these volumes in your overall work load evaluation for the configuration. An inadequate configuration can cause XRC to cancel the copy operation to one or more volumes, or even end the storage control session.

Distributing work loads

Avoid directing all update activity to a small set of common volumes on a single recovery site storage control. Typically, only a small number of devices receive most of the activity on a particular storage control. Usually two to four devices account for 50% or more of the total update activity and 8-10 volumes account for 80% of the update activity during any stress timeframe. If, for example, two primary site storage controls channel their most active volumes into a single recovery site storage control, the recovery subsystem could become rapidly overcommitted with copy activity.

Maximizing access to journal data sets

Where and how you allocate journal data sets affects XRC performance, except in migration mode. XRC journals all copied data, and the capacity of the journal data sets must be enough to support the primary system I/O rate. You do not need journal and control data sets when the XRC session is running in SESSIONTYPE(MIGRATE) mode.

XRC can occasionally fill all of the journal data sets, in which case it must wait for secondary update processing to release journal space before continuing. If the full-journal condition is not relieved within a reasonable amount of time, XRC suspends the session, and then restarts it. If this happens, you may need to enlarge the journal data sets. Follow the procedure that is documented in “Specifying XRC journal, control, and state data sets” on page 42 to enlarge the journal data sets.

Perform the following steps to enlarge the journal data sets.

1. Suspend the session if it is still active.

   _____

2. Enlarge the size of the journal data sets **or** allocate additional ones.

   _____

3. Restart the XRC session with the XSTART command (See “XSTART–starting a session” on page 89).

   _____

Note: The state and control data sets can reside on the same volume. The control data set, however, should not reside on the same volumes as the journal data sets.

See “Specifying the journal data sets” on page 43 for information on determining the correct journal data set size, and for more information about optimizing performance.

Chapter 5. Administering Your XRC Environment

This chapter describes extended remote copy (XRC) Time Sharing Option (TSO) commands used for starting and managing your XRC sessions. XRC can be managed through TSO commands and the ANTRQST API described in “ANTRQST application programming interface overview” on page 99.

XRC processes commands in the order that you issue them, except for the XQUERY and XSET commands, which are processed immediately. All XRC TSO commands can be put into a SYS1.PROCLIB member, included within standard job control language (JCL), or put into CLISTs.

In this chapter

This chapter provides the commands used for administering your XRC environment. The tasks associated with using these commands are described in “Chapter 6. Managing Extended Remote Copy Operations” on page 101.

Table 24. XRC TSO commands

Topic . . .	See page . . .
Issuing extended remote copy TSO commands	55
XADDPAIR–adding volume pairs or utility volumes	56
XADVANCE—updating secondary volumes	61
XCOUPLE—coupling XRC sessions	63
XDELPAIR–deleting volume pairs	66
XEND–ending a session	70
XQUERY–querying a session	75
XRECOVER–recovering data on the recovery system	82
XSET–changing session parameters	84
XSTART–starting a session	89
XSUSPEND–suspending volumes or sessions	91
ANTRQST application programming interface overview	99

Issuing extended remote copy TSO commands

Extended remote copy TSO commands have the same restrictions as other TSO commands. The TSO abbreviation convention requires that you specify as much of the command name or parameter as is necessary to distinguish it from other command names or parameters. The command syntax diagrams show all unique abbreviations. See *z/OS TSO/E Command Reference* for details on TSO commands.

You can issue XRC commands to control an XRC session as soon as TSO becomes operational. Table 25 on page 56 lists the XRC TSO commands and indicates from which sessions you can issue them. The mode of the system data mover determines which command you can issue. See “XRC session state descriptions” on page 23 for session descriptions.

© Copyright IBM Corp. 1994, 2002

Table 25. Extended remote copy time sharing option (TSO) commands

TSO command	Can this command be issued on: (1)						
	An active session?	An inactive session?	A suspended session?	A recovery session?	A coupled, active session?	A coupled, inactive session?	A master session?
XADDPAIR	Yes	No	No	No	Yes	No	No
XADVANCE	Yes	Yes	Yes	No	Yes	Yes	No
XCOUPLE ADD	Yes	No	No	No	No	No	No
XCOUPLE DELETE	No	No	No	No	Yes	No	No
XCOUPLE PURGE	No	No	No	No	No	Yes	No
XCOUPLE RELEASE	No	No	No	No	No	No	Yes
XDELPAIR	Yes	No	No	No	Yes	No	No
XEND	Yes	No	No	No	Yes (2)	No	Yes (3)
XQUERY	Yes	No	No	Yes	Yes	No	Yes
XQUERY MASTER	No	No	No	No	No	No	Yes
XRECOVER	No	Yes	Yes	Yes	No	Yes	No
XSET	Yes	No	No	No	Yes	No	No
XSTART	No	Yes	Yes	No	No	Yes	No
XSUSPEND (Volume)	Yes	No	No	No	Yes	No	Yes (3)
XSUSPEND (Session)	Yes	No	No	No	Yes (2)	No	Yes (3)

Notes:

1. There are several exceptions to the guidelines listed above. Reference the individual commands in this chapter for the exceptions for each command.
2. You may issue the XEND and the XSUSPEND commands to a coupled, active session after all volumes have been suspended (AVS) or there are no volumes in the session (NOV).
3. The XEND, XSUSPEND (VOLUME) and XSUSPEND (SESSION) commands issued to a master session will cause the specified command to be coordinated on the member coupled sessions.

XADDPAIR–adding volume pairs or utility volumes

Use the XADDPAIR command to:

- Define primary and secondary volumes that you want to add to an XRC session. When the command processes, XRC copies all changes from the specified primary volume to the specified secondary volume, unless you specify NOCOPY.
- Add volume pairs that were previously suspended. XRC automatically resynchronizes suspended pairs and processes them ahead of new volume pairs (which require a full-volume copy).
- Add a utility volume to a storage control session.

- Specify fix or float utility device support. See “Using XRC utility devices” on page 102 for more information.

You can specify a maximum of 50 volume pairs with a single XADDPAIR command. The primary and secondary volumes must be compatible. For example, they must meet the following conditions:
- The same number of bytes on each track
- The same number of tracks on each cylinder
- The same or larger number of cylinders on the secondary volume

The XADDPAIR command by default performs an image copy of the primary volume to the secondary volume, unless the volume is a utility volume. Applications may continue to update the primary volumes while the XADDPAIR copy operation is processing. “Adding an XRC volume pair” on page 104 contains more information about the XADDPAIR command.

Note: Multiple-volume data sets require special handling. To ensure that the entire data set is recoverable, add *all volumes* of a multivolume data set to the session.

XADDPAIR command syntax

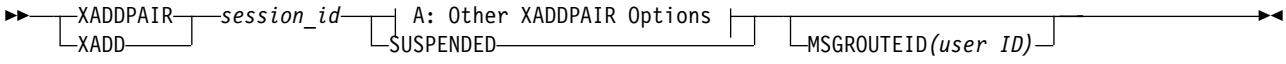

A: other XADDPAIR options: You can issue any of the following parameters in place of the SUSPENDED parameter.

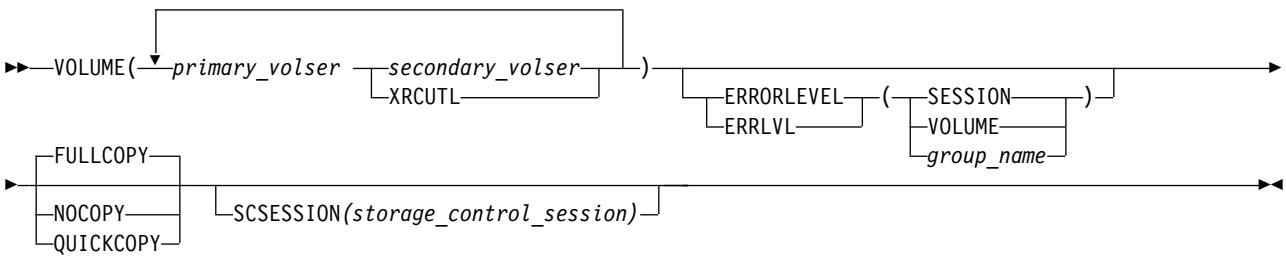

Required parameters:

session_id Specifies the name of an XRC session. The name you specify can be any name acceptable to TSO that matches the name of a session previously specified on the XSTART command, except for the word ALL, which is reserved. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

VOLUME Specifies the volume serial numbers of the primary and secondary volume pairs you wish to add to an XRC session. Use either a blank or a comma to separate each primary and secondary volume serial number within each pair and to separate the pairs. The VOLUME parameter is mutually exclusive with the SUSPENDED parameter.

XRCUTL Specifies that the data mover can use the primary volume as a utility volume for a storage control session. This means that a utility volume is *not*

shadowed to a secondary volume when the secondary volser specified is XRCUTL.

Specifying XRCUTL in conjunction with the UTILITY parameter on the XSET command determines the type of device (fixed or floating) that XRC uses to read data from a storage control. For more information, see “Using XRC utility devices” on page 102.

When you specify XRCUTL, XRC ignores the NOCOPY, DONOTBLOCK, ERRORLEVEL, FULLCOPY, and QUICKCOPY options.

SUSPENDED Specifies that XRC adds all suspended volumes back into the session and schedules them for resynchronization. This parameter is mutually exclusive with the VOLUME, DONOTBLOCK, NOCOPY, SCSESSION, ERRORLEVEL, FULLCOPY, and QUICKCOPY parameters. Suspended volumes keep the same characteristics that existed when the suspension occurred.

Note: Do not use this option to add volumes back into the session if you plan to change any of the volume pair characteristics (such as error level, session assignment, or blocking assignment)

Optional parameters:

DONOTBLOCK

Specifies that the application I/O to the specified primary volumes is *not* blocked , even if the data mover is unable to offload updates from the volumes rapidly enough. If you do not specify this parameter, the default is to block the primary volumes, on storage controls that support this function. This parameter is ignored if you have specified SUSPENDED or XRCUTL.

Note: If you do not explicitly specify DONOTBLOCK, the parameter will not be kept across suspended volumes when they are re-added to the session.

ERRORLEVEL

Specifies how the XRC session responds when an error occurs that causes the session to become unable to process a volume. You can abbreviate this parameter as **ERRLVL.**

Note: Suspended volumes keep the same characteristics that existed when the suspension occurred.

The error level you specify with the XADDPAIR command overrides the value that was specified in the XSTART command. If you do not specify the ERRORLEVEL parameter, the error level defaults to the value that was specified on the XSTART command.

For a volume resynchronization, you may change the error level for a volume pair from what it was prior to the suspension. XRC ignores this parameter for utility volumes or if you also specify SUSPENDED.

To specify how XRC should process certain error conditions, select one of the values:

SESSION Specifies that, if a permanent error that is associated with a duplex primary or secondary volume occurs, XRC suspends *all volume pairs* in the session, regardless of the volume status. If the session is coupled to other XRC sessions, all volumes in the coupled XRC sessions are also suspended.

If you specify the SESSION parameter, volume pairs that are suspended from the session will negatively affect the usability of the remaining secondary volumes. Specify the SESSION parameter to ensure that all secondary volumes necessary for recovery are consistent up to the time of failure.

Note: To couple an XRC session to a master session, use the ERRORLEVEL(SESSION) parameter to ensure a recoverable consistency across all coupled sessions.

VOLUME Specifies that, if a permanent error occurs, the XRC session suspends only the volume pair or pairs that are associated with the error. All other volumes continue to process.

group_name Specifies that, if an error to a duplex primary or secondary volume occurs, XRC suspends all volume pairs that are associated with that group name. The *group_name* can be any name acceptable to TSO. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks. Do not use the reserved names VOLUME and SESSION.

Note: The *group_name* parameter is not supported across coupled XRC sessions.

NOCOPY Specifies that XRC add a volume pair to an XRC session without initially copying the full primary volume. XRC *only* copies updates to the volume after it adds the volume. This assumes that the storage administrator had previously copied the volume. XRC puts the volumes directly into DUPLEX status.

See “Adding a volume with the NOCOPY option” on page 106 for information about using the NOCOPY option.

If you do not specify NOCOPY:

- XRC copies the primary volume onto the secondary volume at the same time that additional application updates are made to both copies.
- Volumes are initially in PENDING status and go to COPY status when XRC selects them for synchronization.
- When synchronization is complete, the volumes enter DUPLEX status.

Note: XRC ignores NOCOPY when establishing utility volumes and when resynchronizing a volume. XRC copies data to the secondary volumes during resynchronization.

FULLCOPY Specifies that XRC make an image copy of *each* track of the primary volume onto the secondary volume at the same time that additional application updates are made to both copies. The volumes are initially in pending status and enter duplex state when the initial volume copy completes. FULLCOPY is the default.

Note: If you do not specify a COPY option using the XADDPAIR command, XRC uses the option that is specified on the XSET COPY command. XRC ignores the FULLCOPY option when establishing utility volumes and when resynchronizing a volume. If you do not specify an XSET COPY option, XRC uses the default of FULLCOPY.

QUICKCOPY Specifies that XRC make an image copy of *allocated* tracks on the primary volume onto the secondary volume at the same time that additional application updates are made to both copies. XRC reads the VTOC of the primary volume to determine the allocated tracks on the volume. The volumes are initially in PENDING status and enter DUPLEX status when the quick volume copy completes.

XRC ignores the QUICKCOPY option when establishing utility volumes and when resynchronizing a volume. For more information, see “Adding a volume with the FULLCOPY or QUICKCOPY option” on page 105.

Note: To ensure data integrity, the initial processing for QUICKCOPY must issue a reserve and then a release for the primary volume during the initial phase of the synchronization process. If access to the primary volume is through a channel extender and the connection fails while XRC has the volume reserved, applications at the primary site will not be able to access the primary volume.

SCSESSION Specifies that XRC assign the volumes that are associated with this command to the value that you specify with *storage_control_session*. You can assign either one or two alphabetic characters for the *storage_control_session* value. If you do not specify this parameter, the default value is “--”.

XRC ignores any change to the SCSESSION designation if this XADDPAIR command is resynchronizing volume pairs or if you also specify the SUSPENDED parameter. If you want to change the volume assignment, you must first issue an XDELPAIR command, then reissue the XADDPAIR command for the volume pair. This will require that you perform a full synchronization copy for the volume pair.

The SCSESSION parameter is valid for utility volumes as well as for regular volume pairs. See “Using XRC utility devices” on page 102 for more information.

MSGROUTEID

Specifies the user ID to which XRC messages associated with this

command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command syntax-checking routine are sent to the user ID that issues the TSO command.

XADDPAIR command examples

The following are examples of the XADDPAIR command:

```
XADDPAIR MAIN1 VOLUME(PRIM02 SECD02) MSGROUTEID(OPER1)

XADDPAIR MAIN1 VOLUME(PRIM02 SECD02) ERRLVL(VOLUME)

XADDPAIR MAIN1 VOLUME(PRIM02 SECD02) SCSESSION(AB)

XADDPAIR MAIN1 SUSPENDED

XADDPAIR MAIN1 VOLUME(PRIM01 SECD01) QUICKCOPY

XADDPAIR MAIN1 VOLUME(PRIM02,XRCUTL)
```

For examples of the XADDPAIR command to add utility volumes, see “Using XRC utility devices” on page 102.

XADVANCE—updating secondary volumes

You can issue the XADVANCE TSO command to both coupled and noncoupled sessions.

The XADVANCE command differs from the XRECOVER command in that it obtains a consistent set of secondary volumes without updating the secondary volume VOLSERs. This makes it easier for you to create a time-consistent backup of the secondary volumes before restarting XRC and resynchronizing volumes.

The XADVANCE command requires all volumes in the session to be suspended. It may be used with an active or an inactive session. You cannot use this command when the session is active on another system, or if the session has any nonsuspended volumes. The XADVANCE command does not apply any data to volumes that were suspended because of a command, or that were suspended because of errors that occurred before the last consistency group.

If an error occurs while writing to a secondary volume that results in all volumes being suspended, that volume will be suspended at a different consistency time than the other volumes in the session. To resolve the time discrepancy, correct the error on the secondary volume, and issue the XADVANCE command to update that volume to the same consistency time as other volumes in the session.

Use the XADVANCE TSO command to update secondary volumes of a session to a time that is consistent with the volumes of other sessions that are coupled to the master session. The XADVANCE command applies secondary volume updates to the minimum master journal time or the pending master command time. This is the same time to which an XRECOVER command would recover the session.

If RACF® facility class protection is being used for XRC, the XADVANCE command should be added to the resource profile using the procedure documented for existing XRC commands. The XADVANCE command must also be defined as an authorized TSO command.

The XADVANCE command produces an XQUERY report that displays *YES* in the ADV column to indicate each volume in the session that was eligible to have updates applied. A *NO* in the column indicates that the volume was not eligible to have updates applied. The report also displays *NO* for volumes that have not reached duplex, that were suspended through a command, or that were suspended at a time before the consistency time that exists as of the beginning of the XADVANCE operation.

Issuing an XADVANCE command to a master session automatically generates an XQUERY MASTER report for the coupled sessions.

See “Example: resuming operations after a sympathetic suspension” on page 152 for a scenario that uses XADVANCE to resume XRC operations after a sympathetic suspension caused by an error.

XADVANCE command syntax

The syntax of the XADVANCE command is:

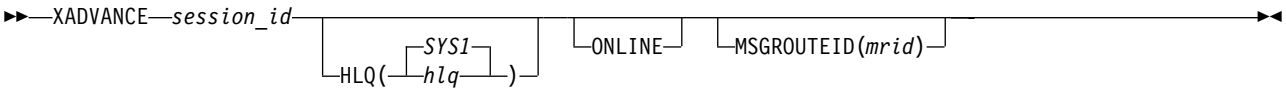

Required parameters:

session_id Specifies the name of an XRC session that is to have its secondary volumes updated to a consistent time. This can be any name acceptable to TSO that matches the name of a session that was previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

Optional parameters:

HLQ*(hlq)* Specifies the name of the high-level qualifier for the state, control, and journal data sets which was previously specified on the XSTART command. The HLQ name must be from one to eight characters, and can be any name acceptable to TSO. SYS1 is the default value.

ONLINE Specifies that any online volume that matches the volser of a secondary volume will be used as a secondary volume when this command processes. See “Using tertiary XRC volumes” on page 111 for more information.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XADVANCE command example

The following is an example of the XADVANCE command:

```
XADVANCE UMC HLQ(SYS1) MSGROUTEID(OPER3)
```

XCOUPLE—coupling XRC sessions

Use the XCOUPLE TSO command to add XRC sessions to, or delete XRC sessions from, a master session. You can also use this command to manage the status of the master session. Specifically:

- The XCOUPLE(ADD) command couples the specified XRC session into a master session, which is then used to coordinate the consistency time and data recovery of multiple XRC sessions. When two or more XRC sessions are coupled, the journal and consistency times of the sessions are maintained such that the entire group of sessions can always be processed by the XRECOVER or XADVANCE commands to a consistent point in time. The XCOUPLE(ADD) command is not allowed for sessions where you specified SESSIONTYPE(MIGRATE) on the XSTART command.
- The XCOUPLE(DELETE) command uncouples the specified active XRC session from the master session. In a coupled environment, when the last volume is deleted, the system will automatically perform the XCOUPLE DELETE command to remove the session from the master session.
- The XCOUPLE(PURGE) command is issued to an inactive coupled session. If the master session is active, this command clears information about the coupled XRC session from the master data set.
- The XCOUPLE(RELEASE) command removes the HOLD status of the master session. Use the RELEASE option only if you no longer need to protect the recoverable state of the coupled XRC sessions in the master session and are ready to resume normal remote copy operations. For more information, see “Managing coupled XRC operations” on page 145.

Note: If you issue the XSTART(*session_id*) command to an XRC session using SESSIONTYPE(MIGRATE), you cannot issue the XCOUPLE command to couple the session to a master session. If the XRC session has been migrated and you wish to couple it, follow these steps:

1. Issue the XEND or XSUSPEND command to the session.
2. Restart the session by using SESSIONTYPE(XRC).
3. Issue the XCOUPLE command to the XRC session

If you restart a coupled session and specify SESSIONTYPE(MIGRATE), the session type is changed to XRC. To restart a coupled session using MIGRATE, you must first issue the XCOUPLE PURGE command to uncouple the session.

Notes:

1. If RACF facility class protection is being used for XRC, the XCOUPLE command should be added to the resource profile by using the procedure documented for existing XRC commands. The XCOUPLE command must also be defined as an authorized TSO command.
2. The maximum number of coupled sessions that are supported by this release of CXRC is 14.

If you issue an XSTART command to restart a previously coupled session, and the automatic XCOUPLE ADD processing fails, subsequent XADDPAIR commands also fail with a return code indicating that the XCOUPLE processing is incomplete. To resolve this condition and to permit XADDPAIR commands to complete, you must either:

- Correct the condition that caused the XCOUPLE ADD command to fail and reissue the XCOUPLE ADD command to the XRC session that encountered the problem.

- Issue the XCOUPLE DELETE command to uncouple the XRC session.

See “XCOUPLE—coupling XRC sessions” on page 63 for more information on adding or deleting sessions from a master session.

XCOUPLE command syntax

The XCOUPLE command includes two syntax diagrams:

- The first syntax diagram allows you to:
  - Add (*couple*) the specified XRC session to the group of sessions that is associated with the master session.
  - Delete (*uncouple*) the specified XRC session from the group of sessions that is associated with the master session.
  - Purge information about the specified coupled XRC session from the master data set when the XRC session is inactive.
- The second syntax diagram allows you to remove the HOLD status of the master session.

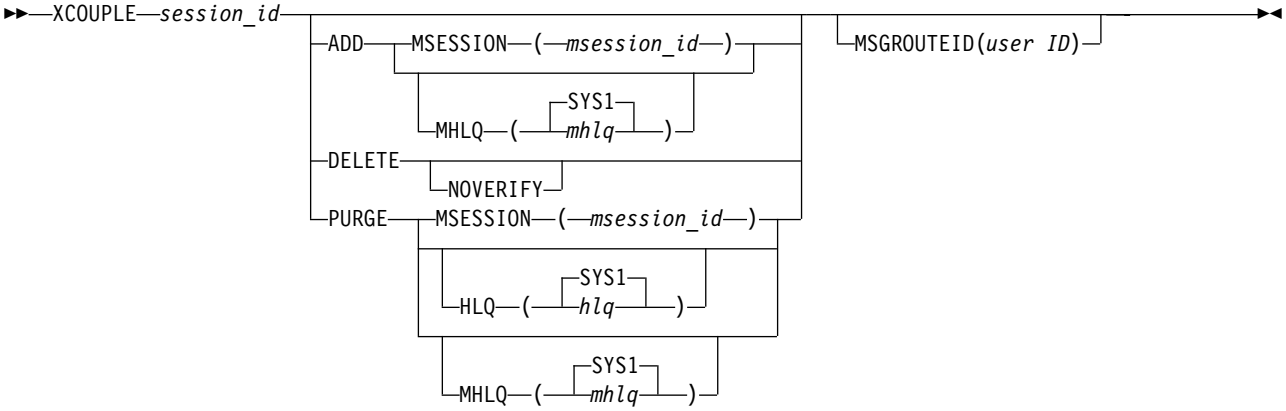

Required parameters:

session_id
: Specifies the name of an XRC session that is to have its coupling status changed. It can refer to an uncoupled session or a coupled XRC session. This name can be any name acceptable to TSO that matches the name of a session that was previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores the leading blanks and the trailing blanks.

(ADD | DELETE | PURGE)
: **(ADD)** indicates that the specified XRC session *session_id* is to be added or coupled to the group of sessions that is associated with the master session. **(DELETE)** indicates that the specified active XRC session *session_id* is to be removed or uncoupled from the group of sessions associated with the master session. **(PURGE)** removes information about the XRC session from the master data set when the XRC session is inactive.

MSESSION*(msession_id)*
: Specifies the name of the master session. The master session is a logical entity that is used to coordinate session commands and data

consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores the leading blanks and the trailing blanks. This parameter is not used on the XCOUPLE DELETE command.

Optional parameters:

HLQ*(hlq)* Specifies the name of the high-level qualifier for the state, control, and journal data sets to which the XCOUPLE command refers. The XCOUPLE HLQ name must be from one to eight characters, and can be any name acceptable to TSO. SYS1 is the default value.

MHLQ*(mhlq)* Specifies the high-level qualifier of the master session control data set. The default is SYS1. The MHLQ is used only with the XCOUPLE ADD and XCOUPLE PURGE commands. It is not a valid option for the XCOUPLE DELETE command.

MSGROUTEID
Specifies the user ID to which XRC messages that are associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

NOVERIFY Specifies that XRC not issue message ANTT0025A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

The syntax of the XCOUPLE command to remove the HOLD status of the master session is:

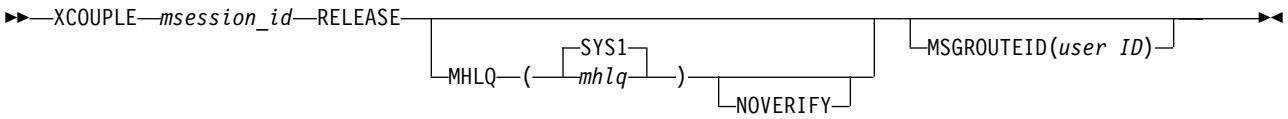

Required parameters:

msession_id Specifies the name of the master session. The master session is a logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores the leading blanks and the trailing blanks.

RELEASE Indicates that all manual, operator-initiated actions have been performed to protect the recoverable state of coupled XRC sessions that are associated with the specified master session. The HOLD status of the master session is to be removed.

Optional parameters:

MHLQ*(mhlq)* Specifies the high-level qualifier of the master session control data set. The default is SYS1.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

NOVERIFY Specifies that XRC not issue message ANTT0025A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

XCOUPLE command examples

The following are examples of the XCOUPLE command:

```
XCOUPLE STJOSEPH ADD MSESSION(HOSPITAL) MHLQ(MSTRHOSP)

XCOUPLE GENERAL DELETE MSGROUTEID(OPER2)

XCOUPLE STJOSEPH PURGE MSESSION(HOSPITAL) MHLQ(MSTRHOSP) HLQ(STJOS1)

XCOUPLE HOSPITAL RELEASE MHLQ(MSTRHOSP)
```

XDELPAIR–deleting volume pairs

Use the XDELPAIR TSO command to specify that XRC delete volume pairs from an XRC session. When the XDELPAIR command processes, XRC stops processing either a primary volume and its associated secondary volume, or a utility volume. You may delete up to 100 volume pairs by listing the 100 primary volume serial numbers in a single XDELPAIR command.

For each volume, the XDELPAIR command returns a timestamp indicating that the data on the secondary volume is consistent with the data that was on the primary volume at the specified time.

If you issue an XDELPAIR command for one or more volume pairs in a session, the command operates the same in either a coupled or noncoupled environment. Deleting the last volume in a coupled session will cause the session to be uncoupled from the master data set.

For more information about deleting utility volumes, see “Using XRC utility devices” on page 102.

Notes:

1. XRC does not process any volumes in a list if there is a command syntax error, or if any of the volumes in the list are not part of the XRC session.
2. You can delete XRC volume pair with the XDELPAIR command, regardless of the status of the pair. Volume pairs in pending, duplex, copy, seqcheck, and suspend status all become inactive when deleted, and are no longer eligible for recovery.

3. An XRC session remains active even when you have deleted all of its pairs. Issue an XEND command to end the session or an XSUSPEND command to suspend the session.

XDELPAIR command syntax

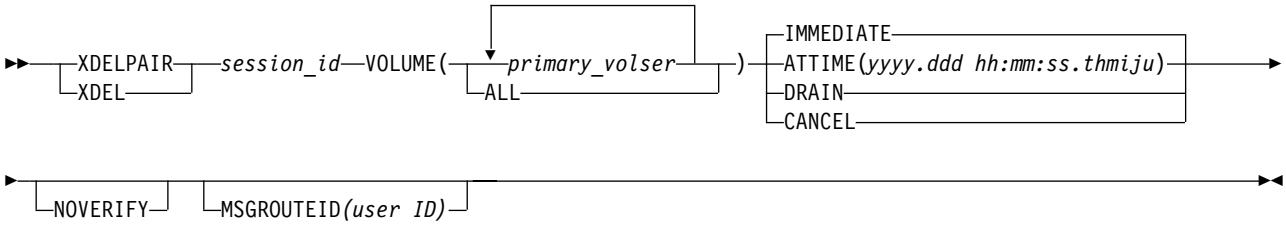

Required parameters:

session_id Specifies the name of an XRC session. This can be any name acceptable to TSO that matches the name of a session that was previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

VOLUME Specifies the volume serial numbers of the primary volumes to be deleted from an XRC session. Do not specify the volume serial numbers of the secondary volumes. The XDELPAIR command automatically removes the associated secondary volumes from the session.

See the IMMEDIATE, ATTIME, and DRAIN parameters for information that describes when volumes are deleted based on the value specified with the VOLUME parameter. See also “Addressing the utility device selection” on page 102 for more information.

You can specify one of the following values:

primary_volser Specifies the volume serial numbers of the primary volumes to delete from an XRC session. To specify multiple volume serial numbers, separate the volume serial numbers with a blank or a comma.

ALL Specifies that XRC delete all eligible volumes in the XRC session.

Optional parameters:

ATTIME Specifies that XRC delete the specified volume pairs after applying all updates, up to the specified time, to the secondary volumes. Table 26 on page 68 describes how volumes are deleted when you specify the ATTIME parameter.

XRC will not apply any updates beyond the specified Universal Time Coordinated (UTC) timestamp to the secondary volumes. All affected secondary volumes are consistent up to the UTC timestamp that is reported on successful completion of the command.

You can issue another XDELPAIR command to change the ATTIME value to either earlier than or later than a previously specified time. The command fails if the specified ATTIME value is earlier than the

consistency group time that XRC is currently processing. If you issue multiple commands to a single volume pair, XRC always uses the last valid ATTIME value that you specified with this command.

Note: You cannot issue an XDELPAIR command with the ATTIME parameter to delete a volume that is in the suspended state. You can only delete suspended volumes with an XDELPAIR IMMEDIATE command or an XEND command.

If all volume pairs are deleted or suspended (because of an error or by command) before the time specified on the ATTIME expires, XRC cancels the XDELPAIR command and clears the ATTIME timestamp.

The system data mover does not process record updates for the deleted volume that are received in the storage control after the specified time. Record updates in the journal data set that are associated with the deleted volume are discarded and are not available for a subsequent recovery. The storage control continues to record updates to other primary volumes. The data mover then processes these updates.

Table 26. Deleting volumes with the ATTIME parameter specified

Specifying the ATTIME parameter and . . .	Results in . . .
VOLUME(ALL)	Volumes that are *not* fixed utility devices will be deleted based on the time specified on the ATTIME parameter. Fixed utility devices will not be deleted.
VOLUME(*primary_volser*)	• If the specified volser is a fixed utility device, the command will be rejected, with return code 4082. • If the specified volser is not a fixed utility device, it will be deleted based on the time specified on the ATTIME parameter. If more than one volser is specified, and one of the volsers is for a fixed utility device, the command will be rejected. If the list of volsers does not include a fixed utility device, XRC will delete each volume in the list based on the time specified on the ATTIME parameter. (The list of volsers can be across multiple storage controls sessions.)

CANCEL Specifies that XRC cancel the previous XDELPAIR command with an ATTIME or DRAIN request. The CANCEL option is only valid when there is a pending XDELPAIR command.

DRAIN Specifies that XRC:

1. Reads the record updates from the cache on all storage controls and obtains the timestamp of the last update for each storage control session.
2. Determines the most recent timestamp.
3. Processes the XDELPAIR command when all records, up to and including the most recent timestamp, have been applied to the secondary volumes that are being deleted from the session.

Table 27 describes how volumes are deleted when you specify the DRAIN parameter.

Note: You cannot issue an XDELPAIR command with the DRAIN parameter to delete a volume that is in the suspended state. You can only delete suspended volumes with an XDELPAIR IMMEDIATE command or an XEND command.

If all volume pairs are deleted or suspended (because of an error or by command) before XRC can drain all updates from the primary storage control, XRC cancels the XDELPAIR(DRAIN) command.

Table 27. Deleting volumes with the DRAIN parameter specified

Specifying the DRAIN parameter and . . .	Results in . . .
VOLUME(ALL)	Volumes that are *not* fixed utility devices will be deleted after all updates from the storage control cache have been applied to the secondary volumes. Fixed utility devices will not be deleted.
VOLUME(*primary_volser*)	• If the specified volser is a fixed utility device, the command will be rejected, with return code 4082. • If the specified volser is not a fixed utility device, it will be deleted after all updates from the storage control cache have been applied to the secondary volumes. If more than one volser is specified, and one of the volsers is for a fixed utility device, the command will be rejected. If the list of volsers does not include a fixed utility device, XRC will delete each volume in the list after all updates from the storage control cache have been applied to the secondary volumes. (The list of volsers can be across multiple storage controls sessions.)

IMMEDIATE Specifies that XRC immediately process the delete request. IMMEDIATE is the default.

Table 28 on page 70 describes how volumes are deleted when you specify the IMMEDIATE parameter.

XRC deletes the volume pairs when the current consistency group has been applied. The storage control continues to record updates to other primary volumes, and the data mover processes these updates. The secondary volumes are consistent up to the UTC timestamp that is reported on the successful completion of the command.

Table 28. Deleting volumes with the IMMEDIATE parameter specified

Specifying the IMMEDIATE parameter and . . .	Results in . . .
VOLUME(ALL)	If a storage control session is in fixed utility mode when this command processes: 1. The storage control session is changed to floating utility mode 2. All volumes in that storage control session are deleted.
VOLUME(*primary_volser*)	• If the specified volser is not a fixed utility device, it is deleted. • If the specified volser is the only device in the storage control session and it is a fixed utility device, the volume is deleted. Otherwise, the command is rejected, with return code 4082. If more than one volser is specified, and one of the volsers is for a fixed utility device and it is not the only device in the storage control session, the command will be rejected. If the list of volsers does not include a fixed utility device, XRC will delete each volume in the list. (The list of volsers can be across multiple storage controls sessions.)

NOVERIFY Specifies that XRC not issue message ANTT0025A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XDELPAIR command examples

The following are examples of the XDELPAIR command:

```
XDELPAIR MAIN1 VOLUME(PRIM01 PRIM04 PRIM05) ATTIME(1997.234 12:14:20.002001)

XDELPAIR MAIN1 VOLUME(PRIM02 PRIM03) DRAIN
```

XEND–ending a session

Use the XEND command to stop all XRC activity to active volumes and to end the XRC session. Issuing the XEND command stops all updates to the secondary volumes at the time specified.

In a coupled environment, you can issue the XEND command to a master session to have all sessions end with their volumes consistent to the same time. See “XEND in a coupled environment” on page 73 for information on ending coupled sessions.

The XEND command returns a timestamp indicating that the data on the secondary volumes of all active pairs is consistent with the corresponding data that was on the primary volumes at the specified time.

After you issue the XEND command, XRC prompts you (unless you have specified NOVERIFY) to confirm the request to end the session. You must specify either `YES` or `NO`. If the XEND command is invoked through a CLIST written in REXX, you should either specify NOVERIFY or place your response to the prompt in the data stack before issuing the command. In a batch environment, you can stack the response.

An XEND command that specifies an ATTIME value overrides a time that was specified with a previous XEND command. Also, an XEND command that specifies an ATTIME value overrides a pending XSUSPEND(*session_id*) command function.

All secondary volumes are consistent up to the UTC timestamp that is reported by successful completion of the command. The storage control session ends after the time that is determined by the XEND command. Record updates in the journal data set are discarded and not available for a subsequent recovery. If you issue an XRECOVER command, the secondary volume serial numbers of volume pairs in the duplex or seqcheck state are changed to match (“clipped to”) the primary volume serial number, but no journal updates are applied to the volumes. The secondary volume serial numbers of suspended volume pairs are also clipped, provided that the pair had completed initial volume synchronization prior to being suspended.

XEND command syntax

The syntax of the XEND command is:

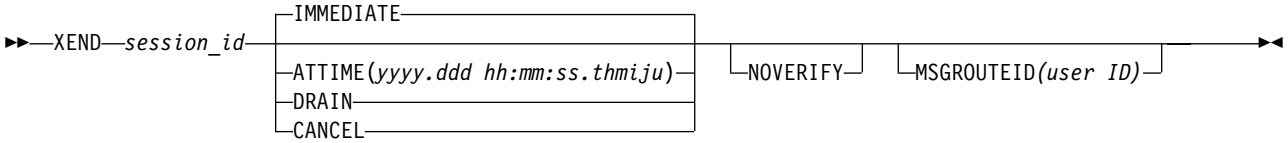

Required parameters:

session_id Specifies the name of an XRC session. This can be any name acceptable to TSO that matches the name of a session previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

Optional parameters:

IMMEDIATE Specifies that XRC immediately process the request to end the session. XRC stops updates to the secondary volumes of active pairs once the current consistency group has been applied. XRC ends the active session. Secondary volumes are consistent up to the UTC timestamp that is reported on the successful completion of the command. IMMEDIATE is the default.

ATTIME Specifies that XRC end the specified XRC session after applying all updates, up to the time that is specified, to the secondary volumes. XRC will not apply updates to the secondary volumes beyond the specified universal time coordinated (UTC) timestamp.

When XRC encounters a record update that has a timestamp later than the time specified, XRC ends the session. All affected secondary volumes are consistent up to the UTC timestamp that is reported on successful completion of the command.

You can issue another XEND command to change the ATTIME value to either earlier than or later than a previously specified time. XRC issues a message that indicates when it has accepted a new ATTIME value. The command fails if the specified ATTIME value is earlier than the consistency group time that XRC is currently processed. If you issue multiple XEND ATTIME commands, XRC always uses the last valid ATTIME value that is specified.

If all volume pairs become deleted or suspended (either due to an error, or by command) prior to the specified ATTIME, XRC cancels the XEND command and clears the ATTIME timestamp.

DRAIN Specifies that XRC:

1. Read the record updates from all storage control caches and obtain the timestamp of the last update for each storage control session.
2. Determine the most recent timestamp.
3. Perform the XEND command when XRC has applied all records, up to and including the most recent timestamp, to the secondary volumes.

The DRAIN parameter supports operations such as migration, where applications are quiesced prior to issuing the DRAIN request. This parameter allows you to quiesce application program updates to the primary volumes and ensures that all updates have been applied to the secondary volumes before the session is ended.

If non-timestamped record updates exist in a storage control, the DRAIN function ensures that all updates in the storage control at the time the command was issued have been successfully applied to the secondary volume or volumes before the command is performed.

If all volume pairs become deleted or suspended (either due to an error, or by command) before XRC can drain all updates from the primary storage control, XRC cancels the XEND(DRAIN) command.

CANCEL Specifies that XRC cancel the previous XEND command that was issued with an ATTIME or DRAIN request specified. If a pending XEND command was originally issued using the name of the master session, you cannot issue the CANCEL option for that command using the name for the local session.

NOVERIFY Specifies that XRC not issue message ANTT0022A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is logged off, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XEND command examples

The following are examples of the XEND command:

```
XEND MAIN1 MSGROUTEID(OPER1) DRAIN

XEND MAIN1
```

XEND in a coupled environment

In a coupled environment, you can issue the XEND command to a master session to have all sessions end with their active duplex volumes at a consistent time across the coupled sessions. You can issue the XEND command to an individual XRC session only:

- After the session has been uncoupled
- If the session is in a coupled state with all volumes suspended
- After there are no volumes in the session

If you uncouple an XRC session using the XCOUPLE command with the DELETE option before you issue the XEND command, the session will no longer be consistent with the other sessions in the master session. Therefore, you will not be able to recover all volumes to a consistent state with the other volumes in the master session.

You may end an XRC session that is coupled in one of three ways:

- Issue an XEND to the master session.
- Suspend or delete all volumes in the XRC session and then issue an XEND to the XRC session.
- Issue an XCOUPLE DELETE to the XRC session and then issue an XEND to the XRC session.

XEND processing for a coupled session will automatically generate an XQUERY MASTER report. See “XQUERY MASTER in a coupled environment” on page 80 for further information on this report.

XEND in a coupled environment command syntax: The syntax of the XEND in a coupled environment command is:

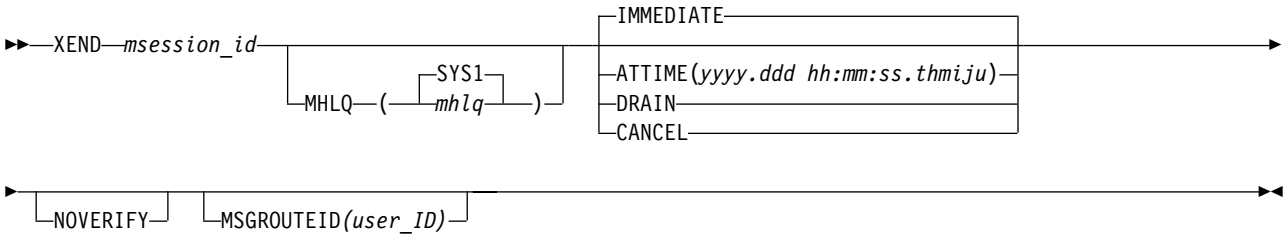

Required parameters:

msession_id Specifies the name of the master session. The master session is a logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores the leading blanks and the trailing blanks.

Optional parameters:

MHLQ*(mhlq)* Specifies the high-level qualifier of the master session control data set. The default is SYS1.

IMMEDIATE When you issue the XEND command by using a master session name, the command will be processed on all XRC sessions that are coupled to the master session. In this instance, the IMMEDIATE parameter is processed as ATTIME equal to a time that will permit all coupled sessions to be ended at the same consistency time.

ATTIME When the XEND command is issued using a master session name, the ATTIME parameter is processed as global ATTIME if no component session has received an application update timestamp past the requested timestamp. If any session has passed the time, the command is rejected.

DRAIN Specifies that all sessions in a coupled master session determine a 'drain time'. When you specify a XEND command by using a master session name, the DRAIN parameter is processed as ATTIME by using the latest component drain time.

CANCEL Specifies that XRC cancel the previous master XEND command that has not yet processed. The CANCEL option cancels a pending XEND command that was issued to the master session but not yet processed. If the XEND command has started processing on any session, the CANCEL request fails. (The CANCEL option is valid only when there is a pending XEND command.)

NOVERIFY Specifies that XRC not issue message ANTT0022A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command to end the master session without waiting for verification.

MSGROUTEID

Specifies the user ID to which XRC messages that are associated with the processing of this command are routed. If the specified user ID is logged off, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XEND in a coupled environment command examples: The following are examples of the XEND in a coupled environment command:

```
XEND MSTR1 MSGROUTEID(OPER1) DRAIN

XEND MSTR3 MHLQ(MSTRMHLQ) IMMEDIATE
```

XQUERY–querying a session

Use the XQUERY command to request status information that is related to an active XRC session. The XQUERY command allows you to customize your information request. An enhancement to the XQUERY function allows you to issue an XQUERY MASTER command to a master session. See “XQUERY MASTER in a coupled environment” on page 80 for additional information on master session reports.

You can generate various summary reports by selecting XQUERY command parameter options. For more information on the types of reports that can be generated, see “Using XRC XQUERY reports” on page 116.

The session consistency group time reported in the XQUERY report is the timestamp of the last application update that was applied to all active secondary volumes in the XRC session. You can use the reported consistency group time to monitor application update activity. The consistency group time does not change when there is no application update activity to any of the primary volumes.

The time reported with IDLE indicates how long the session has been idle. In the IDLE state, XRC has applied all updates to the secondary volumes, and the primary storage controls do not have any updates pending. The time reported with DELAY indicates how far the secondary system lags behind the primary system before all current updates are applied.

The data exposure time reported in the XQUERY report provides an approximation of the time difference between data written to the primary volumes and data secured on the journal data set. If the data exposure time is zero, all data written to primary volumes for the session has been secured on the journal data set. If the data exposure time is greater than zero, then the data that has not yet been secured on the journal data set is potentially nonrecoverable. You may use the data exposure time to identify bottlenecks that might prevent the maximum amount of data from being recovered in the event of a disaster.

See “Using XRC XQUERY reports” on page 116 for XQUERY command output examples.

XQUERY command syntax

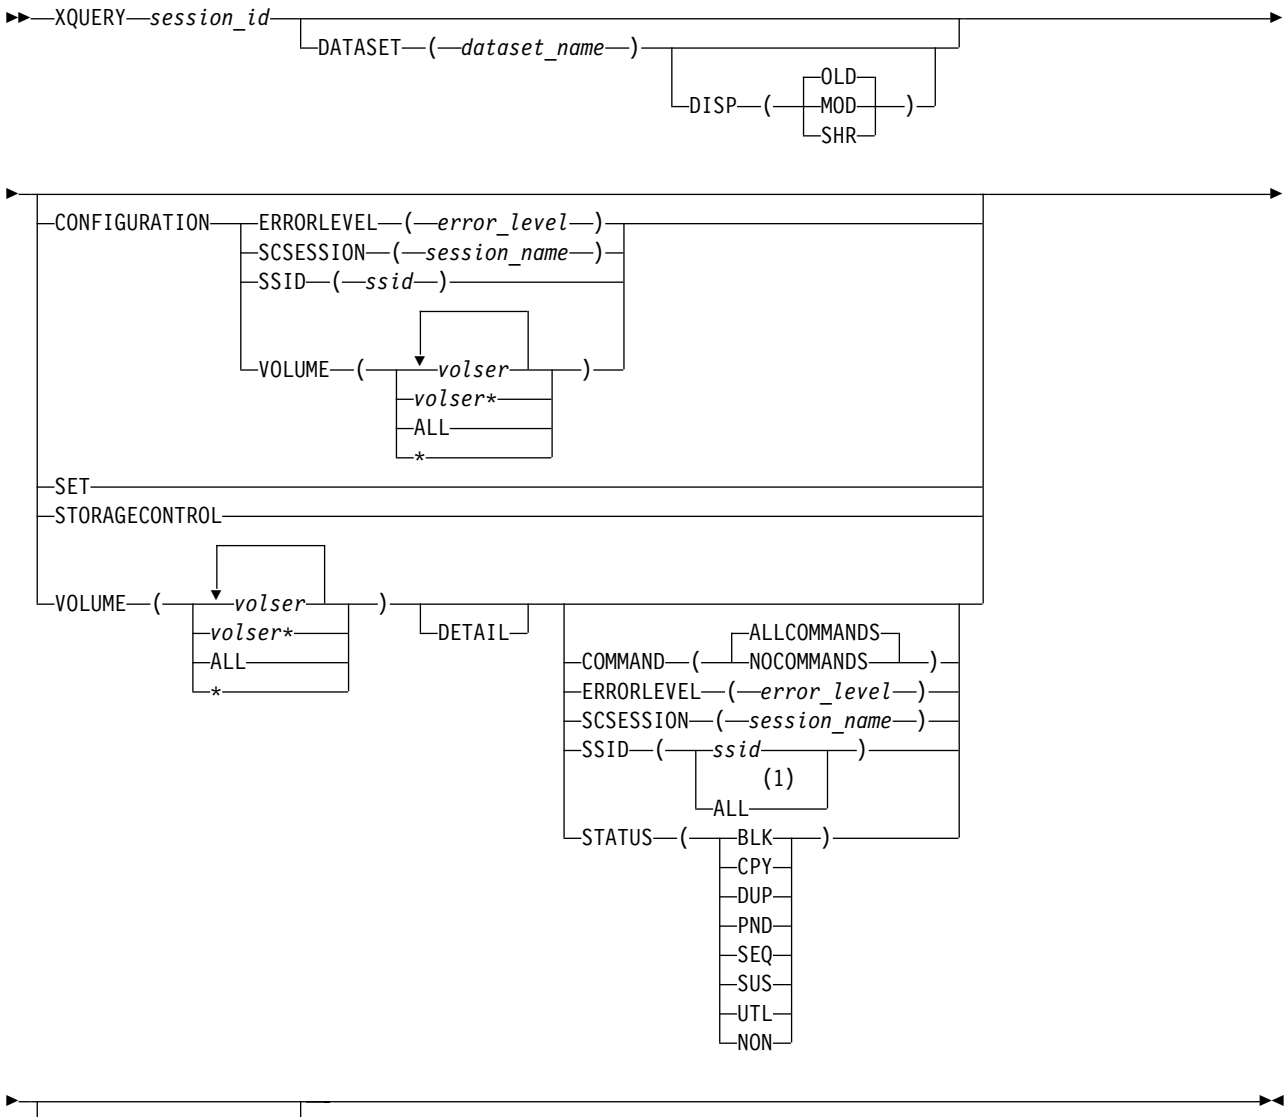

Notes:

1 SSID(ALL) is valid only if you specify the VOLUME and DETAIL parameters. Normally, the VOLUME DETAIL output is sorted by the residual count (RES CNT) in descending order, and then by primary volumes (PRIM VOL) in ascending order. If SSID(ALL) is specified, the output is sorted by SSID, then RES CNT, and then PRIM VOL.

Required parameters:

session_id Specifies the name of an XRC session. This can be any name acceptable to TSO that matches the name of a session previously specified on the XSTART command, or you may specify ALL, which returns information for any active XRC session on this MVS host. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

If you specify the session ID only, without the SET, CONFIGURATION, STORAGECONTROL, or VOLUME parameters, XRC generates a summary report of the session status.

Optional parameters:

DATASET Specifies that XRC direct the XQUERY output to the sequential data set that is identified by *dataset_name*. Standard TSO naming conventions apply for *dataset_name*. The target data set must be a sequential data set and must be large enough to accommodate the XQUERY reports you want to collect. You can specify DATASET with or without DISP; the default is DISP(OLD).

Allocate a data set for the XQUERY function with the following attributes:

```
DCB=(RECFM=FB,LRECL=80,BLKSIZE=6400,DSORG=PS)
```

DISP Specifies how XRC writes the XQUERY output to the data set, as follows:

MOD Specifies that XRC append the output to the data set.

OLD Specifies that XRC clear the data set before it receives output.

SHR Specifies that XRC clear the data set be cleared before it receives output. SHR also allows multiple allocations of the same data set.

If you specify DISP, you must also specify DATASET.

CONFIGURATION

Specifies that XRC generate a configuration report to provide the details of the XRC session configuration. CONFIGURATION is mutually exclusive with the SET, STORAGECONTROL, and VOLUME parameters (VOLUME, however, is a valid subparameter of CONFIGURATION). For an example of the XQUERY configuration report, see “Using XRC XQUERY reports” on page 116. If you specify only the CONFIGURATION parameter, XRC generates a report on the entire configuration.

You can use combinations of the ERRORLEVEL, SCSESSION, SSID, and VOLUME parameters on a single configuration report. The report will include the volumes that meet *all* of the specified filters.

ERRORLEVEL

Specifies that the requested report displays only volumes that are defined with the specified error level (session, volume, or *group_name*). If you do not specify ERRORLEVEL, the requested report will contain a list of all volumes which meet the specified report criteria.

SCSESSION Specifies that the requested report displays volumes that are associated with the storage control session that is specified by *session_name*. This parameter is only valid for the VOLUME or CONFIGURATION reports.

SSID Specifies that the requested report displays only

volume pairs that are associated with the storage control that matches the specified SSID. If you do not specify SSID, the requested report contains a list of all volumes that meet the specified report criteria. You can request a report that is based on the primary or secondary SSID.

VOLUME Specifies that the requested report displays only volumes that meet the specified volume serial number criteria, as well as any other specified report criteria. The VOLUME parameter has the following options:

- VOLUME(*volser*) provides a report of the primary or secondary volume or list of volumes that meet the specified criteria.
- VOLUME(*volser**) provides a report of the primary volumes that match a single volume serial number prefix. Specify the prefix with one to five characters that are followed by an asterisk (wildcard character).
- VOLUME(ALL) provides a report of all volumes in the session that meet the specified criteria.
- VOLUME(*) provides a report of all volumes in the session that meet the specified criteria and that also match the character and asterisk (wildcard) pattern specified with (*).

You can specify a maximum of 100 primary or secondary volume serial numbers on a single command. If any volumes in the list are not part of the XRC session, XRC includes the volumes in the report with an indication of this status.

SET Specifies a report on the current operational values for the XRC session, which are stored in the state data set. SET is mutually exclusive with the STORAGECONTROL, VOLUME, and CONFIGURATION parameters. See “Using XRC XQUERY reports” on page 116 for an example of an XQUERY SET report.

STORAGECONTROL

Specifies a report for all primary storage controls that have volumes in the XRC session. STORAGECONTROL is mutually exclusive with the SET, VOLUME, and CONFIGURATION parameters. For an example of an XQUERY storage control report, see “Using XRC XQUERY reports” on page 116.

Tip: The STORAGECONTROL report includes residual counts and session delay time as part of the status of each of the storage controls. Periodically, issuing this report provides a good way to monitor the XRC system.

VOLUME Specifies that the requested report displays only volumes that meet the specified volume serial number criteria, as well as any other specified report criteria. For example, specifying:

- VOLUME(*volser*) displays the primary or secondary volume or list of volumes that meet the specified criteria.

- VOLUME(*volser**) displays the primary volumes that match a single volume serial number prefix. Specify the prefix with one to five characters that are followed by an asterisk (wildcard character).
- VOLUME(ALL) displays all volumes in the session that meet the specified criteria.
- VOLUME(*) displays all volumes in the session that meet the specified criteria and that also match the character and asterisk (wildcard) pattern specified with (*).

You can specify a maximum of 100 primary or secondary volume serial numbers on a single command. If any volumes in the list are not part of the XRC session, XRC includes the volumes in the report with an indication of this status. VOLUME is mutually exclusive with the SET, STORAGECONTROL, and CONFIGURATION parameters.

You can use combinations of the COMMAND, ERRORLEVEL, SCSESSION, SSID, and STATUS parameters on a single volume report. The report will include the volumes that meet *all* of the specified filters.

DETAIL Specifies that a volume report provides specific detailed information relative to device activity level. This parameter is only valid for the VOLUME report.

COMMAND Specifies that XRC report volumes based on the last pending command issued for the volume, as follows:

ALLCOMMANDS
Specifies that XRC report all volumes that have a pending XDELPAIR or XSUSPEND command request.

NOCOMMANDS
Specifies that XRC report all volumes that do *not* have a pending XDELPAIR or XSUSPEND command request.

ERRORLEVEL
Specifies that the requested report displays only volumes that are defined with the specified error level (session, volume, or *group_name*). If you do not specify ERRORLEVEL, the requested report contains a list of all volumes that meet the specified report criteria.

SCSESSION Specifies that the requested report displays volumes that are associated with the storage control session that is specified by *session_name*. This parameter is only valid for the VOLUME or CONFIGURATION reports.

SSID Specifies that the requested report displays only volume pairs that are associated with the storage control that matches the specified SSID. If you do not specify SSID, the requested report contains a list of all volumes that meet the specified report

criteria. You can request a report that is based on the primary or secondary SSID.

STATUS Specifies that the volume report displays only the listed volumes whose current status matches the status qualifier. You may select only one option per report. The options are as follows:

BLK Specifies that XRC display device-blocked volumes.

CPY Specifies that XRC display copy volume pairs.

DUP Specifies that XRC display duplex volume pairs.

NON Specifies that XRC display volumes that are not duplex volume pairs. That is, copy, pending, and suspended volume pairs are displayed.

PND Specifies that XRC display pending volume pairs.

SUS Specifies that XRC display suspended volume pairs.

SEQ Specifies that XRC display seqcheck volume pairs.

UTL Specifies that XRC display utility volumes that you have added with the XRCUTL parameter.

MSGROUTEID

Specifies the user ID to which XRC messages that are associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command. If you specify both DATASET and MSGROUTEID, XQUERY output is directed to the data set that is specified with DATASET and not to the MSGROUTEID user ID. Error messages are directed to the MSGROUTEID user ID.

XQUERY command examples

The following are examples of the XQUERY command. If you do not use single quotation marks, TSO by default prefixes the user ID to the data set name.

```
XQUERY MAINZ VOLUME(ALL) DETAIL STATUS(DUP)

XQUERY MAINZ VOLUME(flr*) DATASET('XQUERY.OUTPUT1') DISP(MOD)

XQUERY TUCRUN CONFIGURATION ERRORLEVEL(SESSION) STATUS(CPY) COMMAND(NOCOMMANDS)
```

XQUERY MASTER in a coupled environment

Issue the XQUERY MASTER command to a master session to display information about coupled XRC sessions associated with the master session, including the XRC

session name, session status, command pending status, journal delta time, RCV/ADV delta time, master session status, master session recoverable time, and master pending command.

XQUERY MASTER in a coupled environment command syntax: The syntax of the XQUERY MASTER command in a coupled environment is:

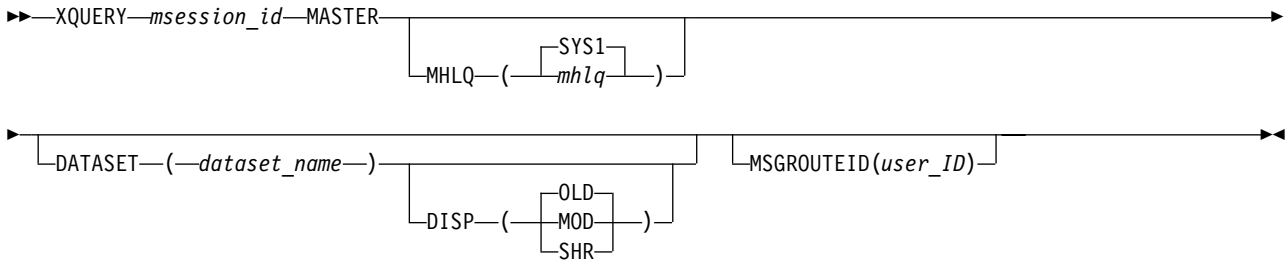

Required parameters:

msession_id Specifies the name of the master session. The master session is a logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

Optional parameters:

DATASET Specifies that XRC direct the XQUERY output to the sequential data set that is identified by *dataset_name*. Standard TSO naming conventions apply for *dataset_name*. The target data set must be a sequential data set and must be large enough to accommodate the XQUERY reports you want to collect. You can specify DATASET with or without DISP; the default is DISP(OLD).

DISP Specifies how XRC writes the XQUERY output to the data set, as follows.

MOD Specifies that XRC append the output to the data set.

OLD Specifies that XRC clear the data set before it receives output.

SHR Specifies that XRC clear the data set before it receives output. SHR also allows multiple allocations of the same data set.

If you specify DISP, you must also specify DATASET.

MASTER Specifies that XQUERY produce a new master coupling report. It may be specified in conjunction with the DATASET and MSGROUTEID parameters.

MHLQ*(mhlq)* Specifies the high-level qualifier of the master session control data set. The default is SYS1.

MSGROUTEID

Specifies the user ID to which XRC messages that are associated with the processing of this command are routed. If the specified

user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command. If you specify both DATASET and MSGROUTEID, XQUERY output is directed to the data set that is specified with DATASET and not to the MSGROUTEID user ID. Error messages are directed to the MSGROUTEID user ID.

XQUERY MASTER in a coupled environment command examples: The following are examples of the XQUERY MASTER in a coupled environment command:

```
XQUERY HOSPITAL MASTER MHLQ(SYS2) DATASET (XQUERY.OUTPUT2) DISP(OLD)

XQUERY SCHOOLS MASTER MHLQ(SYS1) MSGROUTEID(QADAMS)
```

XRECOVER–recovering data on the recovery system

Use the XRECOVER command to bring data on the XRC recovery system to a consistent, recoverable state in the event of a disaster to the primary system.

The first command at the recovery site must be the XRECOVER command. You will receive a “session not active” message if you issue another XRC command before an XSTART or XRECOVER command has completed.

When you issue the XRECOVER command to each individual XRC session, it determines the common consistency time to which all coupled sessions can be recovered. All appropriate updates are then applied to the target XRC session to reach that consistency time.

Issue the XRECOVER command on the recovery system to complete the updating of all journal data to the secondary (target) volumes. The secondary volume serial numbers of volume pairs in the duplex or seqcheck state are changed to match (“clipped to”) the primary volume serial number. The secondary volume serial numbers of suspended volume pairs are also clipped, provided that the pairs have completed initial volume synchronization prior to being suspended. For SESSIONTYPE(XRC), the XRC recovery function must continue to have access to the appropriate journal, control, and state data sets that were in use at the time of the failure, or when the XRC session ended. A migration session requires only the state data set.

If the XRC session has either been suspended or ended normally by a command, the XRECOVER command processing applies the journaled data to update secondary volumes that were suspended while XRC was writing the last consistency group when the XRC session was active. XRC attempts to apply all outstanding data to the secondary volumes so that these volumes are consistent with the rest of the session.

If the suspended volumes can be made consistent, XRC then changes the volume status from SUS to RCV and clips the secondary. If the XRC session has continued to apply consistency groups, then the volume will stay in SUS status. XRC relabels the secondary volumes to the primary volume serial numbers, provided that the volume pair had completed initial volume synchronization prior to the recovery action.

You can also issue the XRECOVER command to recover a group of interlocked coupled sessions to the same consistency time.

In order to provide master session consistency, CXRC applies updates to volumes on the sessions that are behind to allow them to advance to the forward session. CXRC will forward-recover volumes that are suspended because of an error and that have the necessary journal data available. It will *not* forward-recover volumes that are suspended because of an XSUSPEND command.

Finally, the XRECOVER command automatically generates a recovery report. The report includes the last applied timestamp that defines the recovered, consistent state for all volumes. The XRECOVER command also automatically generates an XQUERY MASTER report for coupled sessions.

Note: Do not issue the XRECOVER command to an active XRC session. If a session is active when you issue the XRECOVER command, XRC rejects the command. You must first issue an XSUSPEND or XEND command for the session, and then make all primary volumes offline to the data mover system. The XRECOVER command that follows the XSUSPEND or XEND command relabels each eligible secondary volume with the primary volume serial number. See “Chapter 8. XRC Data Recovery Operations” on page 161 for more information.

See “Example: recovering from a disaster using the XRECOVER command” on page 158 for a scenario of using XRECOVER in a coupled environment.

XRECOVER command syntax

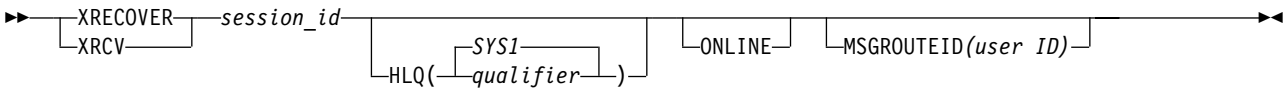

Required parameters:

session_id Specifies the name of an XRC session. This can be any name acceptable to TSO that matches the name of a session that was previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

Optional parameters:

HLQ Specifies the name of the high-level-qualifier that the XRECOVER command will refer to when you issue it. The XRECOVER HLQ name must be the same name that was specified on the XSTART command that initiated the XRC session, if an HLQ was specified. If you started the session without specifying an HLQ, the XRECOVER HLQ name must be from one to eight characters, and can be any name acceptable to TSO. SYS1 is the default value.

ONLINE Specifies any online volume that matches the volser of a secondary volume can be used for recovering data. See “Using tertiary XRC volumes” on page 111 for more information.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is

not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XRECOVER command example

The following is an example of the XRECOVER command:

```
XRECOVER MAIN1 MSGROUTEID(OPER9)
```

XSET–changing session parameters

Use the XSET command to change the characteristics of an active session without requiring that you suspend and restart the session. The XSET command allows you to dynamically change XRC session control values.

When the XSET command processes, characteristics of the XRC session are updated. The changes remain in effect for the duration of the session, even through session suspension and restart, unless changed with another XSET command.

XSET command syntax

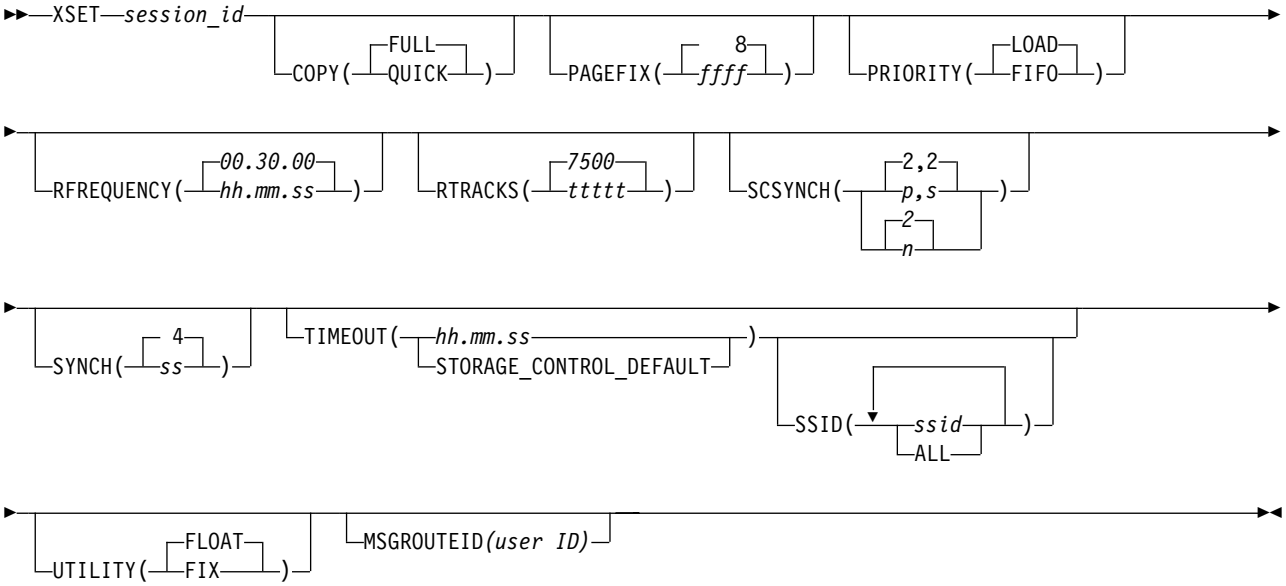

Required parameters:

session_id

Specifies the name of an XRC session. This can be any name acceptable to TSO that matches the name of a session that was previously specified on the XSTART command. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

Optional parameters:

COPY Specifies the extent to which the primary volume will be copied to the secondary volume. Possible values for this option are:

FULL Specifies that XRC make a copy of each track of

the primary volume onto the secondary volume at the same time that additional application updates are made to both copies. The volumes are initially in pending status and enter duplex state when the initial volume copy completes.

QUICK Specifies that XRC make a copy of only allocated tracks on the primary volume onto the secondary volume at the same time that additional application updates are made to both copies. XRC reads the VTOC of the primary volume to determine the allocated tracks on the volume. The volumes are initially in pending status and enter duplex state when the quick volume copy completes.

Note: To ensure data integrity, the initial processing for QUICKCOPY must issue a reserve and then a release for the primary volume during the initial phase of the synchronization process. If access to the primary volumes is through a channel extender and the connection fails while XRC has the volume reserved, applications at the primary site will not be able to access the primary volume.

PAGEFIX Specifies the maximum amount of real storage that XRC will keep page fixed to process I/O operations. This storage is freed back to the system when the data mover no longer needs the storage. The data mover, however, can fix and free any amount of additional real storage, for short-term durations, as required, up to 35 MB per storage control session.

The default value is 8 MB. You can specify the value from zero (no permanent page-fixed storage) to 9999. Changes specified with the PAGEFIX parameter take place when the next set of storage control buffers are processed.

Example: To minimize the processor load for two storage controls, you would set the PAGEFIX value to 70. The two storage controls divide the 70 MB to get 35 MB each, which is the maximum per storage control session. However, you should allocate 35 MB of storage for each unique SCSESSION level that you initiate. If you have two primary storage controls and start two storage control sessions on each, you would PAGEFIX 140 MB of storage.

Attention: Unpredictable results can occur if you specify an amount of permanently page-fixed storage that exceeds or is near the limits of available processor storage. An IPL may be necessary.

PRIORITY Specifies the method to use when choosing the next volume pair to be synchronized or resynchronized. Changes specified with the PRIORITY parameter take place when the next volume is processed. The default value is LOAD. Possible values are:

FIFO Specifies that the system data mover select volumes in the order that is specified on the XADDPAIR command.

LOAD Specifies that the system data mover select volumes whose primary and secondary storage controls have the least load. XRC considers a primary storage control to be overloaded if

it has more than X'500' pending record updates. XRC considers a secondary storage control to be overloaded if any of its primary storage controls have more than X'500' pending record updates.

XRC bypasses a volume pair if the primary or secondary storage control is overloaded. If bypassed, the volume remains the next eligible candidate for XRC to select when the load decreases.

RFREQUENCY

Specifies how long the system data mover will wait before it resets the resynchronization bitmaps. You can specify the frequency in hours, minutes, and seconds. The time can range from 00.00.30 to 18.00.00. The default time is 30 minutes. If you set the frequency value to zero, XRC does not use elapsed time to determine whether to reset the bitmaps.

Volume resynchronization can take longer if you specify a long frequency time; it is likely that more data will have changed on the primary volumes. A shorter time can put a greater demand on data mover host MIPs. If you set the frequency value to zero, XRC does not use elapsed time to determine whether to reset the storage control session bitmap.

RTRACKS Specifies the number of tracks that must change before the data mover resets the resynchronization bitmaps. You can specify a value from 0 to 99999. The default value is 7500 tracks. If you set the tracks value to zero, or to a value that is greater than the number of tracks on a volume, XRC does not use the number of changed tracks to determine whether to reset the bitmaps.

SCSYNCH Specifies the maximum number of volume pairs that the data mover can synchronize or resynchronize concurrently per storage control.

You can specify the XSET SCSYNCH keyword as SCSYNCH(*p,s*) or as SCSYNCH(*n*).

Examples: These are possible examples for setting the XSET SCSYNCH keyword:

- `XSET SCSYNCH(2)`
- `XSET SCSYNCH(2,2)`
- `XSET SCSYNCH(4,2)`

The default setting is SCSYNCH(2,2). A single value, *n*, specifies the maximum number of volume pairs that the data mover can concurrently synchronize or resynchronize per storage control. The limit specified with *n* applies to both primary and secondary storage controls. The parameter value can be from zero to 45.

Multiple values (*p,s*) allow you to separately specify the maximum number of volume pairs that the data mover can concurrently synchronize or resynchronize on both the primary and the secondary storage controls. Both values apply to storage controls that have XRC volume pairs in that XRC session. The limit for the primary storage controls is *p*, and *s* is the limit for the secondary storage controls. The parameter values can be from one to 45. The

parameter value applies to a logical storage subsystem (LSS) regardless of the number of storage control sessions that are active on the storage subsystem.

Set the SCSYNCH value in conjunction with the SYNCH parameter, which specifies the *session* limits for concurrent volume synchronization tasks. Changes specified with the SCSYNCH parameter take place when the next volume initialization is processed.

SYNCH Specifies the maximum number of volume synchronization or resynchronization tasks that XRC can concurrently start in the XRC session. The range of *ss* is from zero to 45 tasks; the default value is four. Specify a value of zero to ensure that XRC does not start any new volume initializations. XRC will continue with existing volume synchronization or resynchronization tasks. Changes specified with the SYNCH parameter take place when XRC processes the next volume.

SYNCH specifies an XRC session-level value. For each synchronization task that it starts, XRC fixes real page storage that is based on the following formula:

```
{ SYNCH * (number of volume pairs in CPY status) * 360K }
```

This storage is in addition to the real storage that is used by normal update processing. The real storage remains page-fixed for the total time the volumes are being synchronized or resynchronized. Set the SYNCH value in conjunction with the SCSYNCH value to control concurrent storage control activity.

TIMEOUT Specifies the primary storage control timeout value for application impact. This value specifies the maximum time that applications are unable to update volumes before the storage control suspends, or ends, the storage control session. Storage controls that do *not* support hardware bitmapping end their affected storage control sessions at the end of the specified TIMEOUT duration. XRC must then resynchronize all volumes in these sessions with a full-volume copy. Storage controls that support hardware bitmapping suspend their affected storage control sessions and perform hardware bitmapping of changes to primary volumes. XRC will not need to perform a full-volume copy to resynchronize these volumes.

You should specify this parameter in hours, minutes, and seconds. Values can range from a minimum of one second (00.00.01) to a maximum of 18 hours (18.00.00). Specify TIMEOUT(**STORAGE_CONTROL_DEFAULT**) to request that XRC use the default set in the storage control, which is normally five minutes (00.05.00). Changes specified with the TIMEOUT parameter take place when you add a new storage control to the XRC session. You can also specify which storage subsystems are immediately affected by the TIMEOUT parameter, using the SSID parameter.

You can override the timeout value by issuing a value with the XSUSPEND(TIMEOUT) command. The new timeout value remains in effect until you restart the XRC session. It is recommended that

the TIMEOUT value never be set less than the missing-interrupt handler (MIH) timeout interval or the channel extender timeout interval for disk volumes.

Note: The XSET TIMEOUT value specifies the acceptable application impact for normal data mover operations. For comparison, the XSUSPEND TIMEOUT value specifies the acceptable application impact (expected duration) for a planned outage.

SSID Specifies the SSIDs that the value specified on the TIMEOUT is applied to, thereby qualifying the scope of the TIMEOUT keyword.

If you do not specify the SSID parameter, then only storage controls that are added in the future will get the new timeout value. If you specify SSID(ALL), all current and future storage controls immediately get the new timeout value.

If you specify a specific set of SSIDs, only the listed storage controls immediately use the new timeout value. Storage controls that you add in the future will get the original timeout value, not the one specified in the TIMEOUT parameter.

You can specify up to nine SSIDs on a single XSET command. Specify each SSID with up to four character values. Separate each SSID by a comma or a blank.

UTILITY Specifies how the utility device is selected for reading data from the primary storage control. The values include:

FLOAT Specifies the utility device for a storage control session will be selected by the storage control. FLOAT is the default.

FIX Specifies the utility device for a storage control session is to be a specified (fixed) device.

Note: Be aware that issuing the XSET command with the UTILITY parameter specified does not activate or deactivate fixed utility device support. This support is enabled when you issue an XADDPAIR command with a secondary volser of XRCUTL. For more information, see “Using XRC utility devices” on page 102.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is logged off, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command syntax-checking routine are sent to the user ID that issues the TSO command.

XSET command examples

The following are examples of the XSET command:

```
    XSET  OLYMPIC  PAGEFIX(70)  PRIORITY(FIFO)  SCSYNCH(3)
          SYNCH(8)  TIMEOUT(00.40.00) MSGROUTEID(OPER1)

    XSET  PANDA  PAGEFIX(140)  PRIORITY(LOAD)  SYNCH(8) TIMEOUT(01.30.00)
RTRACKS(18000) FREQUENCY(04.00.00)  MSGROUTEID(KING9) SSID(ALL)

XSET PANDA UTILITY(FIX) PAGEFIX(12)
```

XSTART–starting a session

Use the XSTART command to:

- Start a new XRC session
- Restart a suspended XRC session

When the suspended XRC session becomes active, the suspended volumes are not automatically resynchronized. You have to issue an XADDPAIR command to resynchronize all volume pairs within the session. You can set up the XSTART command to automatically start the session when you initialize the system.

See “Example: coupling new or existing XRC sessions” on page 145 for a scenario that shows you how to use XSTART to couple an XRC session to a master session.

XSTART command syntax

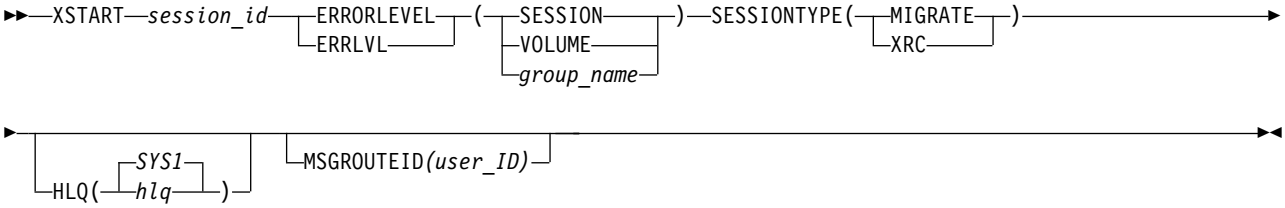

Required parameters:

session_id Specifies the name of an XRC session. This can be any name acceptable to TSO except for the word ALL, which is reserved. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

ERRORLEVEL

Specifies how the XRC session responds when an error occurs that causes the session to become unable to process a volume. You can abbreviate this parameter as **ERRLVL**.

The error level you specify with the XSTART command is the default for all volumes that do not have an error level specified with an XADDPAIR command. The error level specified with the XADDPAIR command overrides the error level specified with the XSTART command. For example, if you specify ERRORLEVEL(SESSION) for a volume pair on an XADDPAIR command, an error on that duplex volume pair results in the suspension of all volumes in the XRC session, even if ERRORLEVEL(VOLUME) was specified on the XSTART command.

XRC supports multiple levels of error processing within a single session. Critical volumes can have ERRORLEVEL(SESSION), while all other volumes can use ERRORLEVEL(VOLUME) or ERRORLEVEL(*group_name*).

SESSION Specifies that, if a permanent error to a duplex primary or secondary volume occurs, XRC suspends *all volume pairs* in the XRC session, regardless of the volume status. Suspending any volume pairs from the session can negatively affect the usability of the remaining secondary volumes. Specify SESSION to ensure that all secondary volumes in the entire XRC session that are necessary for recovery are consistent up to the time of failure.

For sessions that will be coupled, use ERRORLEVEL(SESSION) to ensure a recoverable consistency across all coupled sessions.

VOLUME Specifies that, if a permanent error occurs on a duplex volume, the XRC session suspends the volume pair or pairs that are associated with the error. All other volumes continue to process.

group_name Specifies that, if an error to a duplex primary or secondary volume pair occurs, XRC suspends all volume pairs that are associated with that group name.

The *group_name* can be any name acceptable to TSO. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks. Do not use the reserved names VOLUME and SESSION.

Note: For coupled sessions, error processing for volumes that are contained in a group will only pertain to the volumes within that session.

SESSIONTYPE

Specifies the operating mode (MIGRATE or XRC) for the XRC session. To change the mode of operation, suspend the session with an XSUSPEND command, followed by an XSTART command with a new SESSIONTYPE parameter.

MIGRATE Specifies that the data mover apply updates to the secondary volumes on a consistency group basis, but not write updates to the journal data sets. You must already have allocated the XRC state data set. The journal and control data sets are not used.

XRC commands maintain data consistency within a migration session in the same way as in a recovery session. If a disaster occurs while in SESSIONTYPE(MIGRATE) mode, XRC cannot guarantee data consistency across the secondary volumes, as there is no journaling activity. You can

issue an XRECOVER command after ending a session in MIGRATE mode to clip the secondary volume serial numbers.

XRC Specifies that XRC operate in recovery mode. A SESSIONTYPE(XRC) session makes use of journal volumes to guarantee secondary volume consistency group time. If XRC removes a volume from the session due to an I/O error, XRC can no longer update the secondary volume. Subsequent updates to the primary volume are lost if you must recover the applications on the recovery system.

Optional parameters:

HLQ Specifies the name of the high-level-qualifier for the state, control, and journal data sets to which the XSTART command refers. The XSTART HLQ name must be from one to eight characters, and can be any name acceptable to TSO. SYS1 is the default value.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XSTART command example

The following is an example of using the XSTART command:

```
XSTART MAIN1 SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) MSGROUTEID(OPER1)
```

This command:

- Starts an XRC session called MAIN1
- Suspends an XRC volume pair if an error occurs when specifying the ERRORLEVEL(VOLUME)
- Routes XRC messages that result from processing this command to userID OPER1

XSUSPEND–suspending volumes or sessions

Use the XSUSPEND TSO command to either suspend an XRC session or to suspend one or more volume pairs from the XRC session. In a coupled environment, you may either suspend all XRC sessions coupled to the master session or suspend all volume pairs from sessions that are coupled to the master session. See “XSUSPEND in a coupled environment” on page 96 for information on suspending coupled sessions.

When an XRC session is suspended, the data mover is no longer active and the volume pairs are in suspended state. Session suspension should occur as part of a planned outage when you stop applications or intend the outage to be brief. Suspending an XRC session cancels all pending volume commands; you must reissue pending commands when you restart the session.

The data mover remains active when a volume pair is in suspended state. If a storage control supports a hardware bitmap, the data mover does not read record updates from it. If a storage control does not support a hardware bitmap, the data mover continues to read record updates from the primary storage control that owns the suspended volume. A volume-level suspension does not affect application activity.

After you issue the XSUSPEND command, XRC prompts you (unless you have specified NOVERIFY) to confirm the request to suspend the session or set of specified volumes. If the XSUSPEND command is invoked through a CLIST written in REXX, you should either specify NOVERIFY or place your response to the prompt in the data stack prior to issuing the command. In a batch environment, you can stack the response.

XSUSPEND command syntax

The syntax of the XSUSPEND command is:

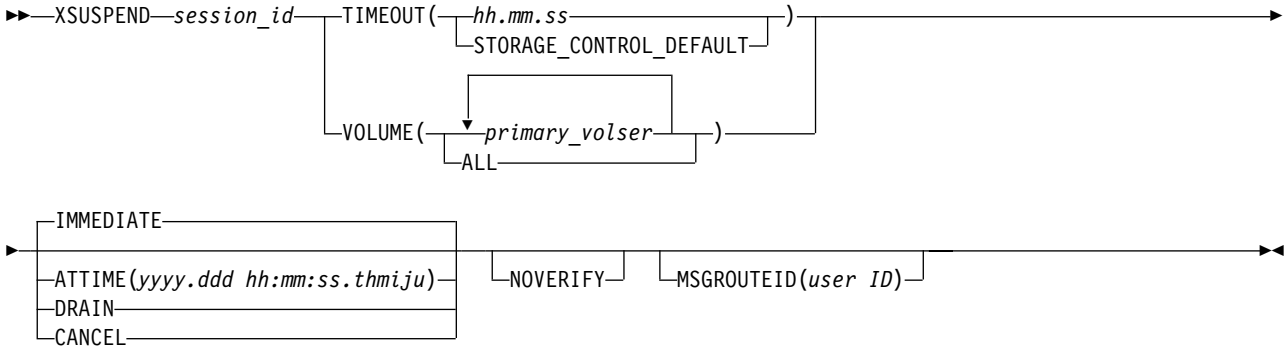

Required parameters:

session_id Specifies the name of an XRC session. This can be any name that was previously specified on an XSTART command, except for the word ALL. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores leading blanks and trailing blanks.

TIMEOUT Specifies the minimum time that a storage control can wait for the XRC session to restart and begin to offload records from the primary storage controls. If you do not restart the XRC session within the time that you specified with the TIMEOUT variable, the primary storage controls take action to recover cache resources that are associated with the storage control session. Storage controls that do not support hardware bitmapping *end* their XRC sessions at the end of the specified TIMEOUT duration. Storage controls that support hardware bitmapping *suspend* their XRC sessions and perform hardware bitmapping of changes to primary volumes.

The TIMEOUT parameter allows you to control the impact that a suspension operation has on the application program. The minimum TIMEOUT value is one second (00.00.01); the maximum TIMEOUT value is 18 hours (18.00.00). See “Suspending an XRC session” on page 113 for more information about the potential impact on the application program.

The TIMEOUT timeframe begins when update activity causes storage control cache to exceed a threshold value. If an IEA482I message is received, application updates remain stopped until the “busy-control-unit” condition is relieved. If the storage control ends the storage control session, you must restart the storage control session and fully synchronize all volumes in the session.

When you specify TIMEOUT, XRC suspends the entire XRC session, and ends the ANTAS*nnn* data mover address space. The TIMEOUT parameter is mutually exclusive with the VOLUME parameter.

When you issue the XSUSPEND command with the TIMEOUT(STORAGE_CONTROL_DEFAULT) parameter, XRC uses the default that is set in the storage control. This value is normally five minutes (00.05.00).

Notes:

1. If all of the storage subsystems in your environment support the hardware bitmap function, you can set a low timeout value; the hardware bitmap ensures a rapid resynchronization time. If some or none of the storage subsystems in your environment support the hardware bitmap function, set the timeout limit to minimize the effect on applications. Take into consideration how long a full-volume resynchronization will take for storage subsystems without hardware bitmapping.
2. ESS storage subsystems suspend their sessions as a part of XSUSPEND processing. Therefore, the TIMEOUT parameter does not affect these subsystems.
3. The XSUSPEND TIMEOUT parameter specifies the acceptable application impact (expected duration) for a planned outage. For comparison, the XSET TIMEOUT parameter specifies the acceptable application impact for normal data mover operations.

VOLUME Specifies the volume serial numbers of the primary volumes that XRC is to suspend from the session. When XRC suspends a volume, the data mover continues to read record updates from the primary storage control. The data mover does not apply updates to the secondary volume while the volume remains in a suspended state. You may specify up to 100 primary volume serial numbers on a command. Use either a blank or a comma to separate each primary volume serial number.

The XSUSPEND command will fail if you include a utility volume in the volume list. You can add or delete utility volumes, but you cannot suspend them.

Note: Although you cannot explicitly suspend a utility volume, XRC can internally suspend a utility volume.

When you specify XSUSPEND VOLUME(ALL), XRC suspends all volume pairs in the session, but does *not* suspend the utility volumes. The ANTAS*nnn* address space remains active. The data mover continues to read record updates for storage controls that do not support a hardware bitmap.

The VOLUME parameter is mutually exclusive with the TIMEOUT parameter.

Optional parameters:

ATTIME Specifies that XRC suspend the specified volume pairs or the XRC session after all updates, up to the time specified with the timestamp, and apply them to the secondary volumes. XRC does not apply any updates to the secondary volume beyond the specified UTC timestamp.

When XRC encounters a record update that has a timestamp later than the time specified, XRC suspends the volume pair. All affected secondary volumes are consistent up to the UTC timestamp that is reported on successful completion of the command.

If you specify VOLUME(ALL) with the ATTIME parameter, XRC processes volumes that are in duplex, pending, seqcheck, and copy status. The XSUSPEND command does *not* process utility volumes when you specify VOLUME(ALL).

You may issue another XSUSPEND command to change the ATTIME value to be either earlier or later than a previously specified time. XRC issues a message that indicates that it has accepted a new ATTIME value. The command fails if the specified ATTIME value is earlier than the consistency group time that XRC is currently processing. If you issue multiple commands for either a single volume pair or an entire session, XRC always uses the last valid ATTIME value that was specified.

If all volume pairs become deleted or suspended (either due to an error or by command) prior to the specified ATTIME, XRC cancels the XSUSPEND command and clears the ATTIME timestamp.

For *volume* suspension, record updates that are received in the storage control after the specified time are maintained in the software bitmap for a subsequent resynchronization operation. For *session* suspension, storage controls that support hardware bitmapping suspend their storage control sessions and perform hardware bitmapping of changes to primary volumes. The data mover stops reading record updates from the storage control. If a storage control does not support a hardware bitmap, the storage control accumulates record updates in cache until either you restart the XRC session, or the storage control reaches the timeout value.

If you issue an XRECOVER command to a session that contains suspended volumes pairs, XRC clips those secondary volumes with the primary volume serial numbers, provided that the volume pair had reached duplex prior to being suspended. XRC does not apply journal updates to the volumes. The data mover ensures that the proper data is recovered from the journal data set so that volumes are consistent to the reported timestamp.

DRAIN Specifies that XRC reads the record updates from all storage control caches and obtains the timestamp of the last update for each storage control session. Specifying this parameter ensures that XRC applies all updates to the secondary volumes before it suspends either the volume pairs, or the XRC session.

If a storage control has never received a timestamped record, the DRAIN function fails. XRC cancels the command if it suspends the volume pair because of an error prior to completion of the function.

If you specify VOLUME(ALL) with the DRAIN parameter, XRC processes volumes that are in duplex, pending, seqcheck, and copy status. The XSUSPEND command does *not* process utility volumes when you specify VOLUME(ALL).

All secondary volumes are consistent up to the UTC timestamp that is reported by successful completion of the command. The storage control continues to receive record updates after the time that is determined by the XSUSPEND DRAIN function. These updates are maintained in the software or hardware bitmap if this is a volume suspension, or maintained in the primary storage control if this is a session suspension.

If an XRECOVER command is issued, the secondary volumes of volume pairs in the duplex state or the seqcheck state are clipped to the primary volume serial number. XRC does not apply journal updates to the volumes. Although XRC does not apply journal updates to suspended volumes, XRC does clip volume serial numbers to pairs that were in duplex state prior to the suspension. The data mover ensures that the proper data is recovered from the journal data set so that volumes are consistent to the reported timestamp.

If all volume pairs become deleted or suspended (either due to an error, or by command) before XRC can drain all updates from the primary storage control, XRC cancels the XSUSPEND command.

CANCEL Specifies that XRC cancel the previous XSUSPEND command with an ATTIME or DRAIN request. The CANCEL option is only valid when there is a pending XSUSPEND command, and cannot be specified with the ATTIME, DRAIN, or IMMEDIATE parameters.

IMMEDIATE Specifies that XRC immediately process the suspend request. Updates to the specified secondary volumes stop when the current consistency group has been applied. The storage control continues to record updates to primary volumes.

For *volume* suspension, XRC suspends the volume pair or pairs from the active session. For *session* suspension, XRC suspends all volume pairs from the active session. In either case, the secondary volumes are consistent up to the UTC timestamp that is reported on the successful completion of the command. IMMEDIATE is the default.

If you want to process all volumes in the session, except for utility volumes and volumes that are already suspended, specify both the VOLUME(ALL) and IMMEDIATE parameters. The XSUSPEND command does *not* process suspended volumes or utility volumes.

NOVERIFY Specifies that XRC not issue message ANTT0028A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XSUSPEND command examples

The following are examples of the XSUSPEND command:

```
XSUSPEND TUCSON CANCEL

XSUSPEND TUCSON TIMEOUT(S) MSGROUTEID(OPER1)

XSUSPEND TUCSON V(prod14) ATTIME(1997.234 12:14:20.002001)

XSUSPEND TUCSON V(prod15,prod16) CANCEL
```

XSUSPEND in a coupled environment

The following section provides information about:

- Suspending all coupled sessions in a master session
- Suspending volumes in all coupled sessions associated with the master session

XSUSPEND of all coupled XRC sessions: In a coupled environment, you can issue the XSUSPEND command to a master session. All sessions are suspended with their active duplex volumes at a consistent time across the coupled sessions.

You can issue the XSUSPEND command to an individual XRC session only:

- After the XRC session has been uncoupled
- If the XRC session is in a coupled state with all volumes suspended
- If there are no volumes in the XRC session

If you issue an XSUSPEND command with TIMEOUT(*hh.mm.ss*) to an XRC session that is part of a coupled session, the command is rejected. A return code is returned that indicates that the session is coupled.

To uncouple and suspend an XRC session that is part of a coupled session, you can use one of the following methods:

- Issue an XSUSPEND command to the master session. The session will be consistent to the same time as the other sessions in the master session.
- Suspend or delete all volumes in the XRC session and then issue an XSUSPEND command to the XRC session. Or issue an XCOUPLE command with the DELETE option and then issue the XSUSPEND command to the XRC session. In either case, the session will no longer be consistent to the same time with the other sessions in the master session. As a result, you will not be able to recover all volumes to a consistent time with the other volumes in the master session.

If you issue an XSUSPEND command for one or more volume pairs in a session, the command operates in the same manner in a coupled or noncoupled environment. The system suspends the specified volume or set of volumes and manages the remainder normally. If the last volume in a coupled session is suspended, the status of the session changes to NON-INTERLOCKED and the session is not kept consistent to the same time as the other coupled sessions.

The IMMEDIATE, ATTIME, and DRAIN parameters allow you to control when the suspension of the master session or an uncoupled XRC session occurs.

After you issue the XSUSPEND command to the master session, XRC prompts you (unless you have specified NOVERIFY) to confirm the request to suspend the

master session or set of specified volumes. You must specify either YES or NO. If the XSUSPEND command is invoked through a CLIST written in REXX, you should either specify NOVERIFY or place your response to the prompt in the data stack prior to issuing the command. In a batch environment, you can stack the response.

XSUSPEND of all volumes in coupled XRC sessions: In a coupled environment, you can suspend the volumes for all the sessions coupled to the master session by issuing the XSUSPEND command to a master session and specifying the VOLUME(ALL) option.

Issuing the XSUSPEND *msession_id* VOLUME(ALL) command causes individual XSUSPEND commands to be issued to each XRC coupled session in the master session, which results in all volumes for the session to be suspended. You cannot specify the CANCEL option with the XSUSPEND command if issued to a master session name.

XSUSPEND in a coupled environment command syntax: The syntax of the XSUSPEND in a coupled environment command is:

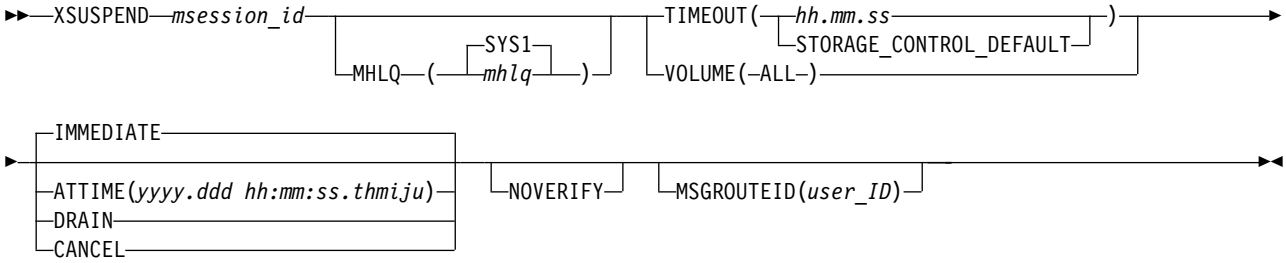

Required Parameters:

msession_id Specifies the name of the master session. The master session is a logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. The maximum length of the name is eight characters. Do not include embedded blanks. TSO ignores the leading blanks and the trailing blanks.

TIMEOUT The TIMEOUT(*hh.mm.ss*) parameter may be applied with the XSUSPEND(*msession_id*) command only after a coupled XRC session is uncoupled.

VOLUME Uncouple XRC sessions from the master session by issuing the VOLUME(ALL) parameter to the master session. You may also use an abbreviated syntax of V(ALL).

Optional parameters:

MHLQ*(mhlq)* Specifies the high-level qualifier of the master session control data set. The default is SYS1. If you specify the MHLQ parameter on the XSUSPEND command, and the *session_id* specified is an XRC session name, the MHLQ parameter will be ignored.

ATTIME When XSUSPEND is issued using a master session name, the ATTIME parameter is processed as global ATTIME if no component session consistency time has passed the requested timestamp. If any session has passed the time, the command is rejected.

DRAIN All sessions in a coupled master session determine a ‘drain time’. When XSUSPEND is issued using a master session name, the DRAIN parameter is processed as ATTIME by using the latest component drain time.

CANCEL Specifies that XRC cancel the previous master XSUSPEND command that has not yet processed. The CANCEL option cancels a pending XSUSPEND command that was issued to the master session, but not yet processed. If the XSUSPEND command has started processing on any session, the CANCEL request fails.

A pending XSUSPEND VOLUME(ALL) to a master session cannot be canceled from the master session. Instead, you must issue an XSUSPEND CANCEL VOL(*volser_list*) from each individual coupled session to specify all volumes in that session by primary volume serial number.

IMMEDIATE When you issue the XSUSPEND command by using a master session name, the command will be processed on all XRC sessions that are coupled to the master session. In this instance, the IMMEDIATE parameter is processed as ATTIME equal to a time that permits all coupled sessions to be ended at the same consistency time.

NOVERIFY Specifies that XRC not issue message ANTT0028A, which prompts the user to confirm that XRC should continue to process the command. If you select this option, XRC immediately processes the command without waiting for verification.

MSGROUTEID

Specifies the user ID to which XRC messages associated with the processing of this command are routed. If the specified user ID is not logged on, TSO saves messages in the TSO BRODCAST data set and displays them after the next logon.

Note: Messages issued by the initial command parser routine are sent to the user ID that issues the TSO command.

XSUSPEND in a coupled environment command examples: The following are examples of the XSUSPEND in a coupled environment command:

```
XSUSPEND MSTR3 TIMEOUT(S) MSGROUTEID(OPER1)

XSUSPEND MSTR2 V(ALL) ATTIME(2000.234 12:14:20.002001)
```

ANTRQST application programming interface overview

The ANTRQST macro provides an application program call to the z/OS system data mover's application programming interface (API). This macro allows you to call XRC, PPRC, FlashCopy, and Snapshot Copy functions. You may need to ensure that necessary address spaces, such as ANTAS000, ANTAS001, or ANTMAIN, are active when you execute the macro.

The ANTRQST macro has been restructured to allow the following new and enhanced functions:

- PPRC request CESTPATH RESETHP(YES | NO) parameter
- CESTPAIR ONLINSEC(YES | NO) parameter
- New FlashCopy commands FCESTABL, FCQUERY, and FCWITHDR
- All XRC and CXRC commands

The ANTRQST macro is described in detail in *z/OS DFSMSdfp Advanced Services*.

Chapter 6. Managing Extended Remote Copy Operations

This chapter describes how to manage XRC copy operations. The commands associated with performing these tasks are described in “Chapter 5. Administering Your XRC Environment” on page 55.

In this chapter

This chapter includes the following sections:

Topic . . .	See page . . .
Starting or restarting an XRC session	101
Using XRC utility devices	102
Adding an XRC volume pair	104
Using the ATTIME, DRAIN, and IMMEDIATE keywords	109
Deleting an XRC volume pair	110
Using tertiary XRC volumes	111
Using the XSUSPEND command	112
Using the XSET TIMEOUT parameter	115
Running the ICKDSF program	116
Using XRC XQUERY reports	116
Accessing secondary volumes while XRC is active	124
Ending an XRC session	125
Applying XRC operational scenarios	125
Identifying XRC system interactions	136

Starting or restarting an XRC session

Use the XSTART command to start an XRC session or to restart a suspended session. See “XSTART–starting a session” on page 89 for additional information about the XSTART command.

1. You can start the XRC session automatically at system initialization time by putting the JCL for a TSO job, that includes the XSTART command, in a SYS1.PROCLIB member.
2. The operating system then starts the SYS1.PROCLIB member from a SYS1.PARMLIB COMMAND*xx* member.

Note: The storage administrator must issue the XSTART command to start the XRC session if the SYS1.PROCLIB member does not include the XSTART command.

XRC can either be restarted or experience a new start, depending on the steps taken after issuing an XSUSPEND command. XRC will undergo a new start if an XRECOVER command is issued following the XSUSPEND command. If a XRECOVER command is not issued, XRC can restart. A new start will always occur after you issue an XEND command to end the session; however, a restart is not possible at that point.

© Copyright IBM Corp. 1994, 2002

When an error occurs and you have specified ERRORLEVEL(SESSION), XRC suspends all volumes in the session but leaves the XRC session active. XRC leaves the secondary volumes in a consistent state within the suspended session. You can correct the error and add the volumes to the active session without having to restart the session. There is minimal impact to the primary application. XRC can rapidly resynchronize the suspended volumes when you add them to the session. This is possible because either hardware bitmaps or software bitmaps maintain the changes to data on the primary volumes.

The following command starts the extended remote copy session with error level set to SESSION. All of the volumes in the session contain associated data. If an error to an XRC duplex state volume occurs, XRC suspends all volumes in the session. If an error occurs with a coupled, interlocked session, the volumes in all coupled sessions are suspended.

```
XSTART DALLAS SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
```

In the preceding example, the minimum syntax supported is:

```
XSTART DALLAS S(XRC) E(SESSION)
```

The following command starts the XRC session with error level set to VOLUME. It directs XRC to suspend any active volume pair that encounters an error.

```
XSTART DALLAS SESSIONTYPE(XRC) ERRORLEVEL(VOLUME)
```

Note: When you restart a suspended XRC session with the XSTART command, the data mover determines which tracks require resynchronization on all of the suspended volumes. The resynchronization process does not occur, however, until you issue an XADDPAIR command for the volume pair.

Using XRC utility devices

This section discusses utility devices that XRC uses to read from a primary storage control. How XRC selects utility devices is described in “Addressing the utility device selection”.

Addressing the utility device selection

Enhancements have been made to the XSET command that allow you to specify the method that is used to select utility devices that XRC uses to read data from a storage control. You can choose that a specific (fixed) device be used or allow the storage control to determine which utility device (floating) to use. For more information, see “Fixed utility devices” on page 103 and “Floating utility devices” on page 103.

For fixed utility devices, XRC always reads data from the storage control using this device, thus eliminating application contention for devices. The following recommendations apply to these devices:

- Select dedicated fixed utility devices that are not used for any other purpose.
- Ensure that the device is not reserved by other applications.
- Ensure that the first device added to a session for a particular storage control be a fixed utility device.
- Define the fixed utility device as a single cylinder device if a storage control is an IBM 2105 ESS. Only a single fixed utility device may be specified for a particular storage control subsystem identifier (SSID).

For floating utility devices, XRC allows the storage control to indicate which device to use when reading data from the storage control.

The utility device method is controlled through the XSET command, and the utility device to be used is specified using the XADDPAIR command with the secondary volser of XRCUTL. The UTILITY parameter of the XSET command, when used in conjunction with the XADDPAIR command, allows you to:

- Set up a storage control session to use a particular device as a fixed utility device.
- Change to a different fixed utility device for a storage control session.
- Disable the fixed utility support for a particular storage control session and return to the default floating utility device support.

Note: Be aware that issuing the XSET command with the UTILITY parameter specified does not activate or deactivate fixed utility device support. This support is enabled when you issue an XADDPAIR command with a secondary volser of XRCUTL.

Fixed utility devices

Table 29 describes device utility support when the XSET command is specified with UTILITY(FIX) and the XADDPAIR command with a secondary volser of XRCUTL is specified.

Table 29. Specifying UTILITY(FIX) and the XADDPAIR command with a secondary volser of XRCUTL

If . . .	**Then . . .**
The primary volser is *not* in the session.	It is added and it becomes the fixed utility device.
The primary volser is already in the session, its secondary volser is XRCUTL, and there is no outstanding XDELPAIR command issued against it.	The primary volser becomes the fixed utility device.
The primary volser is already in the session, its secondary volser is XRCUTL and there is an outstanding XDELPAIR command issued against it.	The XADDPAIR command is rejected, with return code 4027.
The primary volser is already in the session and its secondary volser is *not* XRCUTL.	The XADDPAIR command is rejected, with return code 489.

Floating utility devices

Table 30 describes utility support when the XSET command is specified with UTILITY(FLOAT) and the XADDPAIR command with a secondary volser of XRCUTL is specified.

Table 30. Specifying UTILITY(FLOAT) and the XADDPAIR command with a secondary volser of XRCUTL

If . . .	**Then . . .**
The primary volser is *not* in the session.	It is added to the session. The fixed utility mode is deactivated if it is active and support returns to floating utility mode.
The primary volser is already in the session and its secondary volser is XRCUTL.	The fixed utility mode is deactivated if it is active and support returns to floating utility mode.

Table 30. Specifying UTILITY(FLOAT) and the XADDPAIR command with a secondary volser of XRCUTL (continued)

If . . .	Then . . .
The primary volser is already in the session and its secondary volser is *not* XRCUTL.	The XADDPAIR command is rejected, with return code 489.

More about utility devices

Table 31 describes utility support when the UTILITY parameter is specified with either (FIX or FLOAT) and the XADDPAIR command with a secondary volser is not XRCUTL:

Table 31. Specifying UTILITY(FIX or FLOAT) and the XADDPAIR command with a secondary volser that is not XRCUTL

If . . .	Then . . .
The primary volser is *not* in the session.	The primary volser is added to the session.
The primary volser is already in the session and its secondary volser is XRCUTL, but it is not a fixed utility device.	The XADDPAIR command processes normal volume initialization.
The primary volser is already in the session, its secondary volser is XRCUTL and it is a fixed utility device.	The XADDPAIR command is rejected, with return code 4082.
The primary volser is already in the session and its secondary volser is not XRCUTL.	The XADDPAIR command is rejected with either: • Return code 490, if the new secondary volser matches an existing secondary volser. • Return code 489, if the new secondary volser is different that an already existing secondary volser.

Adding an XRC volume pair

The XADDPAIR TSO command specifies both the primary (source) and secondary (target) volumes of a volume pair that you want to add to an XRC session. You can add up to 50 volume pairs with each XADDPAIR command. See "XADDPAIR–adding volume pairs or utility volumes" on page 56 for the XADDPAIR command syntax.

Notes:

1. When you add a suspended volume pair back to the session, the data on the recovery (target) volume is incomplete (not ready for recovery) until the volume pair reaches duplex state. Choose the resynchronization time carefully to allow the volume pairs to quickly reach duplex state.
2. After XRC receives an XADDPAIR command, it slightly delays the start of volume initialization to allow you time to issue additional XADDPAIR commands.
3. A high activity level on the primary volume's storage control can delay the initial volume copy process. XRC will not start, or continue with, an initial volume copy or resynchronization process on a storage control that has a storage control session that exceeds the update threshold value (X'500'). Also, XRC will stop the initial volume copy or resynchronization of any other volume pair that shares the same secondary storage control with a volume pair whose primary storage control has a storage control session that exceeds the update threshold value. If the volume pair's copy-completed percentage is not changing as expected, you

may want to observe the activity level for the associated primary and secondary storage controls. When the activity level declines, the volume copy continues.

Initializing volume pairs

The XRC volume pair initialization function selects the next volume to synchronize or resynchronize, and prioritizes the volumes from the command's list of volumes. The first priority is to resynchronize existing volume pairs, then to add new volume pairs to the XRC session. XRC selects and adds new pairs as needed according to the PRIORITY subparameter (FIFO or LOAD), which you can specify with the XSET command. See Table 32 for more information.

Table 32. PRIORITY specification with the XADDPAIR command

If you specify . . .	Then . . .	Notes . . .
PRIORITY(FIFO)	XRC processes the volumes that were listed by an XADDPAIR command in the order that they are listed.	1. First, XRC processes all volumes that are eligible for resynchronization. 2. Next, XRC processes all volumes that are eligible for synchronization. XRC processes the volumes within the limits that you specified in the SYNCH and SCSYNCH parameters of the XSET command.
PRIORITY(LOAD)	XRC processes the volumes that were listed by an XADDPAIR command based on the primary and secondary storage control load in this order:	1. First, XRC processes all volumes that are eligible for resynchronization. XRC can defer the initialization if the associated storage control has too much activity at the time that the volume is eligible. The volume continues to be eligible for processing when the activity level has decreased. 2. Next, XRC processes all volumes that are eligible for synchronization. Again, XRC can defer the initialization if the primary storage control activity is too high. XRC performs both resynchronization and synchronization processing within the limits that the SYNCH and SCSYNCH parameters of the XEND command specify.

Adding volumes in a volume group

The ERRORLEVEL keyword of the XADDPAIR command provides an opportunity to assign a common group name to volumes that are related. XRC suspends all of the volumes in a common group when an error occurs to any one of the volumes that belong to that group.

See "XADDPAIR command syntax" on page 57 for information about adding related volumes as groups. "Chapter 10. Recovering from Error Conditions Using XRC" on page 175 describes additional ERRORLEVEL parameter details.

Adding a volume with the FULLCOPY or QUICKCOPY option

Unless you specify NOCOPY, XRC makes an initial copy of the primary volume. Use Table 33 on page 106 to determine which copy option to choose.

Table 33. Copy option determination

If you use . . .	Then . . .	Notes . . .
the default FULLCOPY option	XRC creates this initial copy of the secondary volume by copying all of the tracks from the primary volume.	You may change the default by using the XSET COPY option. See “XSET command syntax” on page 84.
the QUICKCOPY option	XRC copies *only* allocated tracks on the primary volume to the secondary volume during this initialization of the secondary volume. Initialization of the secondary volume is performed while copying primary volume updates to the secondary volume.	During this processing, XRC does not copy the volume serial number from the primary volume to the secondary volume. Therefore, during normal copy operations, a display of the secondary volume serial number shows its original serial number. During XRECOVER processing, XRC sets the secondary volume serial number equal to the serial number of the primary volume; this is not performed for XADVANCE processing.

Therefore, if you specify either the FULLCOPY or QUICKCOPY option when you add a volume pair to an XRC session, XRC ensures the initial data integrity for the secondary volume for an XRC volume pair. XRC ensures that it has successfully copied all tracks (with FULLCOPY) or all allocated tracks (with QUICKCOPY) on the primary volume, and all updates that are made to the volume during initialization, to the secondary volume. As a result, when volume initialization is complete, the secondary volume is in a known, consistent state.

Notes:

1. FULLCOPY is the default option unless you change the default by using the XSET COPY option.
2. As part of the resynchronization function, XRC automatically copies all changed and any not previously copied tracks whenever suspended volume pairs are added again to a session. For this resynchronization function, XRC ignores any copy option specified.

Use the QUICKCOPY option with caution

As a part of QUICKCOPY initial processing, XRC must issue a reserve for the primary volume. After the tracks that contain the allocated space for the volume have been identified, XRC issues a release during the initial phase synchronization phase. If access to the primary volume is through a channel extender and the connection to the primary volume is lost while XRC has the reserve, XRC will not be able to release the volume. Applications then will not be able to access the primary volume.

Adding a volume with the NOCOPY option

The NOCOPY option bypasses the initial primary to secondary volume copy and eliminates the initial data transfer between volumes. This lets you quickly initialize a large number of XRC-managed volumes.

Use the NOCOPY option with caution

Because of potential data integrity exposures, use the NOCOPY option of the XADDPAIR command only in TSO command mode, not in automated procedures. Specify NOCOPY only when installation procedures can guarantee that the primary and secondary volumes are equivalent at the time that you issue the command.

The NOCOPY option specifies that the primary and secondary volumes are identical and that it is not necessary to copy the primary volume to the recovery system. Be careful when you specify the NOCOPY option when any of the following are true:

- You issue the XRC TSO XADDPAIR command as part of automatically initiated XRC subsystem activity, because the NOCOPY option can introduce potential data integrity exposures during a system failure.

  Example: An application may have completed a successful write operation to a primary volume by placing the data in the storage control's nonvolatile storage. Before the data mover can copy the data to the recovery system, however, the data mover system fails. If you specify NOCOPY when you restart the data mover system, XRC does *not* copy the primary volume to the secondary volume. XRC will not synchronize them even though the two volumes may no longer contain exactly the same data. There is a data integrity exposure here because latent updates exist , the data mover has not received the updates, and there is no indication of missing data.
- The XADDPAIR command follows a DFSMSdss full-volume restore that included the COPYVOLID parameter. In this case, you must manually set the secondary volume serial number to a unique ID, and vary it online.

Note: The XADDPAIR TSO command with the NOCOPY option requires a few seconds to complete across the sysplex. If application hosts initiate disk write updates to the volumes before the volumes are under XRC control, those updates will not be timestamped. (Only write I/O operations to volumes that are part of an XRC session are timestamped.) This temporary inconsistency has passed when an XQUERY command returns the information that there are no volumes in seqcheck status.

The XADDPAIR command with NOCOPY is equivalent to the dual copy function with the NOCOPY option. **You must use it with care**. When you specify NOCOPY, both the primary and secondary volumes must contain exactly the same data at the point in time when the pair is established. To ensure that the volume contents match, perform the following steps:

1. Stop all application activity to the primary volumes.
2. Use DFSMSdss to dump the full primary volumes. Specify the ALLDATA and ALLEXCP options to dump both allocated used and allocated unused space on the volumes. (Some applications refer to the allocated unused space of a data set.)
3. Restore the full secondary volumes.
4. Ensure that no applications have made updates to the primary volumes during the volume dump and restore process.
5. Issue the XADDPAIR command with the NOCOPY option.
6. Once the XRC pair is established, restart the applications to the primary volumes. If this is part of an automated procedure, you can check for the ANTA8101I or ANTA8004I message to detect when the pair is under XRC control.

Adding a volume with the SUSPENDED option

Use the XADDPAIR command SUSPENDED keyword to restore the volume configuration to the state that existed prior to a suspension. The keyword processes all pairs that are suspended at the time the XADDPAIR command is issued. XRC then places each primary volume in an active state so that it is eligible to start resynchronization, and so that each secondary volume can receive updates.

Example: Say that there are 1000 volumes in the XRC session. 200 are in suspended state and 800 are in duplex state. An XADDPAIR SUSPENDED command would start the volume initialization process for the 200 suspended volumes. XRC initializes the volumes subject to the SYNCH and SCSYNCH values that are currently in effect.

Adding a volume with the DONOTBLOCK option

Specify the DONOTBLOCK option to exclude a volume from device-level blocking. The system data mover uses the device-level blocking function when the primary storage control supports this function. An application may generate significant updates to a primary volume in a single-channel program chain, such as a defragmentation or full-volume restore operation. The device-level blocking algorithm paces the operation to allow the data mover to off-load updates in an efficient manner. Consider using the DONOTBLOCK option only for high-performance, sensitive volumes that generate very short updates, such as IMS WADs volumes.

Understanding the XADDPAIR command

The following are examples of XADDPAIR commands:

- To add utility volumes and regular volume pairs to two storage controls, use the following commands:

```
XADDPAIR DALLAS VOLUME(PRIMA1 XRCUTL,PRIMB1 XRCUTL)
XADDPAIR DALLAS VOLUME(PRIMA2 SEC001,PRIMB2 SEC002)
```

- To add utility volumes and volume pairs to two storage controls, use the following command:

```
XADDPAIR DALLAS VOLUME(PRIMA1 XRCUTL,PRIMB1 XRCUTL,PRIMB2 SEC001,
         PRIMB2 SEC002)
```

- To add two volume pairs to the XRC session, use the following command. In this example, XRC copies the primary volumes PRIM01 and PRIM02 to the secondary volumes SECO01 and SECO02, respectively. A blank or comma is used to separate each primary and secondary volume serial number within each pair; a comma is used to separate the pairs.

```
XADDPAIR DALLAS VOLUME(PRIM01 SECO01,PRIM02 SECO02)
```

- To add a volume pair with the MSROUTEID and NOCOPY options, use the following command. In this example, the command adds a volume pair to the XRC session, routes XRC messages to user ID E299000, and specifies that XRC only copy updates to the primary volume. XRC does not copy the initial volume image.

```
XADDPAIR DALLAS VOLUME(PRIM01 SECO01,PRIM02 SECO02)
         MSGROUTEID(E299000) NOCOPY
```

  Note: XRC ignores the NOCOPY keyword, however, if you issue this XADDPAIR command to resynchronize the volume pair.

- To add volume pairs with the ERRORLEVEL option, use the following command. In this example, the command adds two volume pairs to the XRC session so that both pairs belong to a group that is called DB2GROUP:

```
XADDPAIR DALLAS VOLUME(PRIM01 SECO01,PRIM02 SECO02)
         ERRORLEVEL(DB2GROUP)
```

- To add volume pairs with an SCSESSION keyword, use the following command. In this example, the command adds two volume pairs to the storage control session called GA. The XADDPAIR SCSESSION keyword provides a way to manage volumes for a single storage control. If there are multiple storage control sessions within that storage control, then all of the sessions are affected.

```
XADDPAIR DALLAS VOLUME(PRIM01 SECO01,PRIM02 SECO02)
         SCSESSION(GA)
```

Guideline: For the best system data mover performance, assign heavy use volumes to separate storage control sessions.

Example: use JCL to start an XRC session

You can run the following JCL example as a batch job, or process it from a SYS1.PROCLIB member to automate the procedures that are shown here. This example shows how to use JCL to start an XRC session, add a single volume pair to the session, and route message output to a specific user ID.

```
//IKJEFT01  JOB,
//          MSGLEVEL=(1,1),MSGCLASS=A,NOTIFY=user ID
//STARTXRC  EXEC PGM=IKJEFT01
//STEPLIB DD DSN=authorized.cmdlib,DISP=SHR
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
  PROFILE MSGID
  XSTART DALLAS SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) MSGROUTEID(user ID)
  XADDPAIR DALLAS VOLUME(PRIM01,SECO01) MSGROUTEID(user ID)
/*
```

Example: use a CLIST to start an XRC session

You can process the following TSO CLIST in a standard TSO environment. This example shows how to use a CLIST to start an XRC session and add two volume pairs to the session.

```
PROFILE MSGID
XSTART DALLAS SESSIONTYPE(XRC) ERRORLEVEL(VOLUME)
XADDPAIR DALLAS VOLUME(PRIM01,SECO01,PRIM02,SEC02)
```

Using the ATTIME, DRAIN, and IMMEDIATE keywords

The ATTIME, DRAIN, and IMMEDIATE keywords let you control the consistency group time when you delete, end, or suspend volume pairs. Consider the following guidelines when using these functions.

- Use the ATTIME keyword in environments where applications continue to update the primary volumes even after the time specified on the ATTIME keyword has passed. XRC will not apply primary updates to the secondary volumes beyond the time you have specified on the ATTIME keyword.
- XRC processes timestamps with respect to the timestamp that is generated by updates on the primary system. Depending on the data mover system load, the XRC session consistency group time for the primary system may lag behind the data mover system time. As a result, ATTIME events can sometimes lag behind the data mover system time.
- The ATTIME and DRAIN keywords do not apply to volumes that are in the suspended state. You can only delete these volumes with an XDELPAIR IMMEDIATE command or an XEND command.
- Specify the VOLUME(ALL) and IMMEDIATE keywords if you want to process **all volumes** in the session, regardless of their volume states. The exception is that XSUSPEND ALL IMMEDIATE will *not* include utility volumes or volumes that are already suspended.
- It is possible that all volume pairs can become deleted or suspended, either due to an error or by command, either before the specified ATTIME occurs or before the data mover can DRAIN the stored updates from the primary storage control. If this happens, XRC cancels the command that contains the affected ATTIME or DRAIN parameter, and clears any timestamp that is related to that command.

- Under certain circumstances, the data mover must approximate the time when the specified ATTIME occurs on the primary system. This can happen, for example, if you specify an ATTIME value with a command, but the application activity to the primary volumes has already been quiesced. The current timestamp of the primary system is therefore unavailable to the data mover, so the data mover must approximate the time. XRC conservatively approximates the time to ensure that enough time has passed on the primary system. XRC does not apply data to the secondary volumes beyond the specified ATTIME. The ATTIME event may therefore happen later than expected. XRC guarantees that the data on all secondary volumes is consistent as of the reported ATTIME.
- When you issue XSUSPEND or XEND commands to a master session, the corresponding action is processed on all XRC sessions that are coupled to the master session. The IMMEDIATE parameter is processed as ATTIME equal to a time that permits all coupled sessions to be ended at the same consistency time. The ATTIME parameter is processed as global ATTIME if no component session consistency time has passed the requested timestamp. If any session has passed the time, the command is rejected. The DRAIN parameter is processed as ATTIME by using the latest component drain time.

Deleting an XRC volume pair

Use the XDELPAIR command to delete volumes or utility devices from an XRC session. For more information, see “XDELPAIR–deleting volume pairs” on page 66 and “Using XRC utility devices” on page 102.

The XDELPAIR command returns a timestamp that identifies the application time that the volume was last consistent with the XRC session. After you issue the XDELPAIR command, XRC prompts you (unless you specified NOVERIFY) to confirm the request to delete the volumes. You can start the XDELPAIR command through a CLIST in either a foreground or a batch environment. If you do, you should either specify NOVERIFY or place your response to the prompt in the data stack. See the CLIST examples that follow.

Examples: The following are examples of commands to delete one or more volume pairs:

- The following command deletes a volume pair from the XRC session:

```
XDELPAIR DALLAS VOLUME(PRIM01)
```

- The following command deletes three volume pairs from the XRC session and routes messages to user ID E299000:

```
XDELPAIR DALLAS VOLUME(PRIM01 PRIM02 PRIM03) MSGROUTEID(E299000)
```

- You can process the following TSO CLIST in a standard TSO environment for commands that require a response. The following example shows how to use a CLIST to delete two volume pairs:

```
PROFILE MSGID
CONTROL PROMPT
XDELPAIR DALLAS VOLUME(PRIM01,PRIM02)
DATA PROMPT
YES
ENDDATA
```

- The following REXX CLIST example adds a new volume pair to the XRC session and deletes another pair from the session:

```
/* REXX */
'PROFILE MSGID'
'PROFILE PROMPT'
SETPROMPT=PROMPT('ON')
```

```
'XADD   A V(XRCF40,SECF51)'
PUSH 'YES'
'XDEL   A V(XRCF26)'
EXIT
```

- You can submit the following batch job to the system:

```
//XRCJOB   JOB   MSGLEVEL=(1,1),MSGCLASS=H,REGION=2048K,CLASS=A
//XSTART   EXEC  PGM=IKJEFT01
//STEPLIB  DD    DSN=IBMUSER.CLIST.CLIST,DISP=SHR
//SYSTSPRT DD    SYSOUT=*
//SYSTSIN  DD    *
  PROFILE MSGID
  EXEC 'TSOUSER.CLIST.CLIST(JOB)' C
/*
```

- The above JCL starts the following CLIST. The CLIST is in the data set TSOUSER.CLIST.CLIST.

```
CONTROL PROMPT
PROFILE MSGID
XDEL NEVADA V(PRIM00)
DATA PROMPT
YES
ENDDATA
```

Note: You do not need to include the SETPROMPT and PUSH statements if you specify the NOVERIFY option.

Using tertiary XRC volumes

The XRECOVER and XADVANCE commands have been enhanced to allow you to recover to copies of secondary volumes (other than the original secondary volumes, therefore the name tertiary).

A tertiary copy may be made either to create volumes that the recovery system will use when it takes over the primary site's application workload or as a backup of the secondary volumes to ensure that there is a consistent copy of all recoverable volumes in place while XRC secondary volumes are resynchronized. Rather than have the recovery system use the XRC secondary volumes when taking over the primary site's workload, the restart will use the tertiary volumes.

If the keyword ONLINE is specified on either the XRECOVER or XADVANCE command, any online volume that matches the volser of a secondary volume can be used for recovering data.

Example: using XRC tertiary volumes

Suppose you want to set up a test environment using copies of XRC secondary volumes to allow your application testing to use real data. The following steps show you how to copy secondary volumes to tertiary volumes using FlashCopy commands against secondary volumes that exist on an IBM 2105 ESS storage control:

1. Suspend all volume pairs in the XRC session using the XSUSPEND VOLUME(ALL) command.

   _____

2. Issue FlashCopy commands to copy secondary volumes to tertiary volumes.

   _____

3. Issue FlashCopy commands to copy journal, state, control, and master data sets to tertiary volumes.

4. Issue the XADDPAIR command with the SUSPENDED parameter specified to add suspended volumes back into the session. Wait until all suspended pairs are in duplex state.

5. On the tertiary volumes, rename the journal, state, control, and master data sets to a different session name.

6. Vary the FlashCopy secondary volumes offline to the system on which the XRECOVER or XADVANCE commands will be issued.

7. Vary the tertiary volumes online to the system on which the XRECOVER or XADVANCE commands will be issued.

8. Issue the XRECOVER or XADVANCE command with the ONLINE parameter that specifies the renamed session. Remember that the XRECOVER command will clip (match) the volsers to secondary volsers, but the XADVANCE command will not.

Using the XSUSPEND command

The following section includes information suspending XRC volume pairs and XRC sessions.

Suspending XRC volume pairs

Whenever volume pairs are in a suspended state, XRC does not apply primary volume updates to the suspended secondary volumes. A track-level bitmap is maintained of changed tracks that are associated with the suspended primary volumes. Suspension messages report the consistency group time for the suspended volumes. See “XSUSPEND–suspending volumes or sessions” on page 91 for the XSUSPEND command syntax.

The suspension of a volume pair, as the result of an error, may cause XRC to take action on other volume pairs as a result. The ERRORLEVEL that you specified for the suspended volume pair, and the scope of the error, both determine the actions that are taken by XRC. See “Chapter 10. Recovering from Error Conditions Using XRC” on page 175 for more information about each ERRORLEVEL specification.

Table 34. Hardware bitmap support considerations

If . . .	Then . . .
XRC suspends volume pairs on storage controls that support hardware bitmaps	the hardware maintains a track-level bitmap of primary system updates.
XRC suspends volume pairs on storage controls that do *not* support hardware bitmaps	the data mover continues to record the track-level bitmaps for suspended volumes as long as the ANTAS*nnn* data mover address space remains active.

If errors occur while a volume pair is in suspended state, the pair remains in a suspended state and the ERRORLEVEL parameter has no effect on suspended volume pairs. Certain severe errors such as loss of the software or hardware bitmap records may require that you do a full-volume resynchronization on the volume pair when you add it back to the XRC session.

Suspending an XRC session

For storage controls that support hardware bitmaps, suspension of an XRC session has no application performance impact. For storage controls that do not support hardware bitmaps, a suspended XRC session can affect the system environment. The actual impact depends on the configuration, the activity level of the volume pairs in the suspended session, and the primary storage control's timeout value.

The capability of the storage control largely determines what happens when an XRC session becomes suspended. Storage controls that are able to do so manage record updates with a hardware bitmap. For these types of volumes, the outage (planned or unplanned) can be of any duration. The primary impact is that the resynchronization time can be longer for an extended outage.

See "XSUSPEND–suspending volumes or sessions" on page 91 for the XSUSPEND command syntax.

Managing cache resources on storage subsystems that support hardware bitmaps

When XRC suspends a storage control session on a storage subsystem, such as an ESS, that supports hardware bitmaps, the storage control does not need to store changes in its cache storage. The hardware maintains a bitmap of primary system updates.

Managing cache resources on storage subsystems that do not support hardware bitmaps

When XRC suspends a storage control session on a storage control that does *not* support hardware bitmaps, the storage control continues to accumulate changes in its cache storage. At the end of the TIMEOUT duration, the storage control ends the storage control session. When you restart the XRC session with an XSTART command, the data mover drains the storage control cache and accumulates the updates in a software bitmap. When you issue an XADDPAIR command, the software uses the software bitmap to resynchronize the volumes in the XRC session.

A suspended storage control session that experiences a lot of update activity places a greater demand on existing cache resources. This suspended session has an increased potential to significantly affect primary system performance. The storage control may cancel storage control sessions as a result.

If you do not issue the XSTART command before the timeout expires, the storage control *ends* the XRC session. You can monitor the current timeout for a storage control by issuing an "F ANTAS000,SCDATA" command. (See "SCDATA operation (XRC, CC)" on page 300 for command specifics.) If the timeout is close to expiration, follow the procedure that is described in "Examples: changing TIMEOUT intervals for suspended sessions" on page 133 to avoid having to recopy the volumes on the storage control.

If a suspended XRC session requires significant cache resources, other applications that are accessing the same storage control can encounter IEA482I, IEF196I, or ADR348E messages. You may choose to cancel the jobs that generate the messages, then rerun them after you have restarted the XRC session. Be aware that DFSMSdss concurrent copy sessions and suspended XRC sessions both use cache resources. The potential exists to overrun these resources, resulting in the messages that are described above.

During any planned outage, primary systems that issue RESERVE commands to the utility volume will enter into a long-busy condition. Either a storage control report or the ″F ANTAS000,SCDATA″ command can identify the utility volumes in a suspended storage control session.

Avoiding deadlock conditions

XRC allocates the state, control, and journal data sets when you issue an XSTART command. The catalog that is accessed by the HLQ keyword on the XSTART command controls access to these data sets. Ensure that application programs have not reserved this catalog data set, as this can cause a long busy condition (with resultant IEA482I or IEF196I messages). If this occurs, the XSTART function fails with a return code of 435 or 486 when attempting to allocate these data sets. You will need to cancel the jobs causing the deadlock situation.

You may avoid this deadlock condition by having the XRC state, control, and journal data sets in a user catalog that contains only entries for these data sets. Also ensure that either XRC is not copying the volume these data sets are on, or the volume does not contain data sets that application programs are updating. Either of these conditions can cause a deadlock condition.

Adding back suspended volumes to the XRC session

Issue the XADDPAIR command to add individual volume pairs, all suspended volume pairs, or suspended utility volumes back into an XRC session. See “XADDPAIR–adding volume pairs or utility volumes” on page 56 for the XADDPAIR command syntax.

Adding back specific suspended volume pairs

Issue an XADDPAIR command to add specific suspended volume pairs back to the XRC session. An XADDPAIR command that is issued to an active volume pair generates an error message. The XADDPAIR command resynchronizes each volume pair by copying all tracks of data that have changed since the time of suspension. If the volume pair was in pending state, the initialization copy continues from the point where the pair became suspended. It also copies all data that has changed since the pair became suspended.

You can change the keyword options when you reissue an XADDPAIR command to a suspended volume pair, with the following constraints:

- You cannot change the primary and secondary volumes in the pair. They must be the same volumes that you initially specified.
- You may change the ERRORLEVEL value. Suspended volumes will keep the same characteristics that existed when the suspension occurred.
- You may change the DONOTBLOCK value. If you do not explicitly specify DONOTBLOCK, the parameter will not be kept across suspended volumes when they are re-added to the session.
- You may add the MSGROUTEID keyword, or specify a new routing.
- You cannot change the SCSESSION value.
- XRC will ignore the NOCOPY, FULLCOPY, or QUICKCOPY keyword.

Adding back all suspended volume pairs

Resynchronize **all** of the suspended volumes that are in an XRC session with the SUSPENDED parameter of the XADDPAIR command. The volume pair keeps the same characteristics that you specified with the XADDPAIR command that created the volume pairs.

The SUSPENDED parameter provides a rapid way to restore the environment as it existed prior to either a session-level suspension or the suspension of all volumes

in the session. However, you cannot use this option to change the characteristics (such as error level and session assignment) of a volume pair.

Adding a suspended utility volume back to the storage control session

If an error causes the primary storage control session to end, the system data mover places the utility volume or volumes in a suspended state. These suspended utility volumes are then not accessible for use.

After you fix the problem that caused the error, reissue an XADDPAIR command to reactivate suspended utility volumes and add them back into the storage control session.

Using the XSET TIMEOUT parameter

The XSET TIMEOUT keyword specifies the maximum application impact that is allowed by an XRC session. As a result of the TIMEOUT value, a storage control that supports hardware bitmapping automatically *suspends* an XRC session if the data mover does not read from the storage control sidefile within the timeout period. This action frees up storage control resources. The hardware bitmap and application programs continue with no performance impact. XRC resynchronizes the volumes, based on the hardware bitmap, when it receives a subsequent XADDPAIR command for the volumes on that storage control. The volumes are not consistent with the session consistency group time until the resynchronization is complete.

A storage control that does not support hardware bitmapping *ends* the storage control session if the timeout interval expires. This frees all storage control resources so that there is no application performance impact. The fact that the storage control session has ended, however, means that you must completely synchronize the volumes in the XRC session from scratch. The volumes are not consistent with the XRC session consistency group time until the resynchronization is complete.

Recommendation: Set the XSET TIMEOUT command to a value that is acceptable for your primary application production environment. It is recommended that you *do not* set this value lower than the MIH interval for the volumes. The default value is 5 minutes. The TIMEOUT interval does not begin until the residual in the primary storage control exceeds a high threshold value. You can minimize impact to application programs with the TIMEOUT value, application update rate, and the storage control cache size.

Running the ICKDSF program

Table 35 contains information about running the ICKDSF REFORMAT REFVTOC program.

Table 35. Determining when to run the ICKDSF REFORMAT REFVTOC program

If . . .	Then . . .	Comments . . .
An XRC volume pair is established or reestablished using the NOCOPY option	Run the ICKDSF REFORMAT REFVTOC program on the primary volume after the pair is established.	The volume table of contents (VTOC) and indexed VTOC information then reflect the primary volume that is being accessed.
The primary and secondary of this pair is switched at any time	Run the ICKDSF REFORMAT REFVTOC program on the new primary volume	The VTOC and indexed VTOC information then reflect the new primary volume.
The secondary volume is removed from the pair for use as a simplex volume	Run the ICKDSF REFORMAT REFVTOC program on the volume	The VTOC and indexed VTOC information then reflect the correct information for the device.

Using XRC XQUERY reports

Use the XQUERY command to query the status of an active XRC session. The XQUERY command can generate the following reports:

- Session summary
- Volume summary
- Settings for session defaults
- Storage control activity
- Detail of storage control activity
- Detail of volume activity
- Configuration activity
- Master session summary (in coupled environments only)

"XQUERY–querying a session" on page 75 explains the command parameters that produce these reports.

CXRC enhances XQUERY reports which return a consistency time to report a recoverable delay. For a coupled session, the recovery delay is the difference between the XRC session's latest controller time and the master session's journal time. For an uncoupled session, the recovery delay is the difference between the latest time reported by a controller session under that XRC session and the journal time of the session. For the duration of this delay, XRECOVER cannot recover applications. The recovery delay can be thought of as the degree of exposure of data to a disaster.

Notes:

1. For a configuration report or volume report, you can use combinations of the filtering keywords. The report will include the configurations (or volumes) that meet *all* of the specified filters. See the XQUERY CONFIGURATION and VOLUME keywords for details.
2. The order of the volume pairs in an XQUERY report may differ from the order in which the volume pairs were added to the XRC session.

3. “Recovering data with XRC—uncoupled and coupled sessions” on page 166 describes a special XQUERY report that the XRECOVER command automatically generates as part of a recovery session.
4. XRC automatically generates an XQUERY storage control report when an XSUSPEND command suspends an XRC session.
5. If you direct XQUERY command output to a data set, XRC does not also provide the output to the system’s log. Remember that when you specify the data set name, standard TSO conventions apply. Unless you enclose the data set name in quotation marks, TSO appends the user ID to the front of the data set name.
6. XQUERY command output in the SYSLOG can contain a connection code at the end of the ANTQ8200I message and at the beginning of each of the lines associated with that connection code. The system generates this code for any multiline WTO.
7. XRC automatically generates an XQUERY MASTER report during an XEND, XSUSPEND TIMEOUT(), or XADVANCE of a coupled session.

Examples of XQUERY reports in an uncoupled environment

The reports in this section are examples of XQUERY reports that are generated in an uncoupled environment.

XQUERY summary report

The following is an example of an XQUERY summary report that summarizes the operating parameters for the XRC session. This report was generated when the XQUERY SESSION1 command processed.*z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY SUMMARY REPORT - 002
ANTQ8203I ------------------------------------------------------------
ANTQ8238I TOTAL=37  DUP=32  CPY=1  PND=1  SUS=2  SEQ=0  UTL=1
ANTQ8231I DATA CONSISTENT(1998.266 19:41:27.998450)  DELAY(00:00:01.83)
ANTQ8240I DATA EXPOSURE(00:00:00.68)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(VOLUME)  HLQ(SYS1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8234I PND SESSION CMD = XSUSPEND ATTIME(1998.267 12:00:00.000000)
ANTQ8236I PND XSUSPEND TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8201I XQUERY SUMMARY REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY volume report

The following is an example of an XQUERY volume report. This report provides a complete volume summary and is similar to the report that the initial remote copy release provided. This report was generated when the XQUERY SESSION1 VOLUME(ALL) command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY VOLUME REPORT - 003
ANTQ8211I PRIM   SEC     ERROR  SYNCH
ANTQ8212I VOL    VOL     LEVEL    % STA CMD OP -------TIMESTAMP--------
ANTQ8203I ------------------------------------------------------------
ANTQ8213I PRIM01 SECD01 SESSION     DUP
ANTQ8213I PRIM02 SECD02 VOLUME   21 CPY
ANTQ8213I PRIM03 SECD03 SESSION     SEQ
ANTQ8213I PRIM04 SECD04 VOLUME   95 SUS        1997.135 11:59:43.583528
ANTQ8213I PRIM05 SECD05 SESSION  90 SUS        1997.135 12:48:48.071334
ANTQ8213I PRIM06 SECD06 VOLUME      DUP DEL DR 1997.135 12:56:58.276463
ANTQ8213I PRIM07 SECD07 VOLUME      DUP
ANTQ8213I PRIM08 SECD08 VOLUME      DUP SUS AT 1997.135 12:49:25.984596
ANTQ8213I PRIM09 SECD09 VOLUME      DUP SUS DR 1997.135 12:59:28.048573
ANTQ8213I PRIM10 SECD10 VOLUME      DUP
ANTQ8213I PRIM11 SECD11 IMSGROUP 95 CPY DEL AT 1997.135 12:38:18.771624
ANTQ8213I PRIM12 SECD12 IMSGROUP 00 PND DEL DR 1997.135 12:58:21.771624
ANTQ8213I PRIM13 SECD13 VOLUME   57 CPY
ANTQ8213I PRIM14 SECD14 DB2GROUP 00 PND SUS AT 1997.135 12:58:30.894533
ANTQ8213I PRIM15 SECD15 DB2GROUP 83 CPY SUS DR 1997.135 12:58:38.781519
ANTQ8213I PRIM16 SECD16 VOLUME   00 PND
ANTQ8213I PRIM17 SECD17 VOLUME      SEQ SUS AT 1997.135 12:58:30.894533
ANTQ8213I PRIM18 SECD18 GROUP1      SEQ SUS DR 1997.135 12:58:38.781519
ANTQ8213I PRIM19 SECD19 VOLUME      SEQ
ANTQ8213I PRIM20 SECD20 VOLUME   ** SUS        1997.135 12:00:45.936853
ANTQ8213I PRIM21 SECD21 VOLUME   42 SUS        NO_TIME_AVAILABLE
ANTQ8213I PRIM22 XRCUTL          -- UTL
ANTQ8213I PRIM23                 -- NIS
ANTQ8213I PRIM24 XRCUTL          -- UTL
ANTQ8238I TOTAL=23  DUP=6   CPY=4  PND=3  SUS=4  SEQ=4  UTL=2
ANTQ8231I DATA CONSISTENT(1997.135 12:48:27.998450)  DELAY(00:00:00.82)
ANTQ8240I DATA EXPOSURE(00:00:00.56)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(VOLUME)  HLQ(SYS1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8201I XQUERY VOLUME REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY volume detail report

The following is an example of an XQUERY volume detail report that was generated when the XQUERY SESSION1 V(ALL) DETAIL command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY VOLUME_DETAIL REPORT - 002
ANTQ8216I PRIM   SEC     ERROR  SYNCH          RES  THD       SC SC
ANTQ8217I VOL    VOL     LEVEL    % STA CMD OP CNT  CNT  SSID SN ID
ANTQ8203I ------------------------------------------------------------
ANTQ8218I PRIM01 SECD01 SESSION     DUP        0089 0500 8700 -- 01
ANTQ8218I PRIM02 SECD02 SESSION     DUP        0499 0500 8702 -- 01
ANTQ8218I PRIM03 SECD03 SESSION     DUP*       0620 0500 1010 AA 02
ANTQ8218I PRIM04 SECD04 SESSION  82 CPY        0056 0500 1011 -- 02
ANTQ8218I PRIM05 SECD05 SESSION  24 CPY*       05DB 0500 8702 BB 03
ANTQ8218I PRIM06 SECD06 SESSION     DUP        01AB 0500 8702 CC 04
ANTQ8218I PRIM07 SECD07 SESSION     SUS        0000 0500 8702 CC 04
ANTQ8238I TOTAL=7  DUP=4  CPY=2  PND=0  SUS=1  SEQ=0  UTL=0
ANTQ8231I DATA CONSISTENT(2000.325 23:20:49.504484)  DELAY(00:00:02.35)
ANTQ8240I DATA EXPOSURE(00:00:00.92)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(HLQ1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8201I XQUERY VOLUME_DETAIL REPORT COMPLETE FOR SESSION(SESSION1)
```

z/OS MVS System Messages, Vol 1 (ABA-AOM) contains details about the XRC messages that are listed within the preceding report.

XQUERY set report

The following is an example of an XQUERY set report. The values shown are the defaults from the state data set or are the values that you have set with the XSET

command. This report may be generated in either a coupled or an uncoupled environment. This report was generated when the XQUERY SESSION1 SET command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY SET REPORT - 003
ANTQ8203I ------------------------------------------------------------
ANTQ8261I HLQ(HLQ1)  SCTRAP(OFF)  TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8262I SYNCH(4)  SCSYNCH(2,2)  PRIORITY(LOAD)  SUPRDUMP(OFF)
ANTQ8263I SCBUFFERS(576)  PAGEFIX(8)  PACE(5)  COPY(FULL)
ANTQ8264I RFREQUENCY(00.30.00)  RTRACKS(7500)
ANTQ8201I XQUERY SET REPORT COMPLETE FOR SESSION(SESSION1)
```

z/OS MVS System Messages, Vol 1 (ABA-AOM) contains details about the XRC messages that are listed within the preceding report.

XQUERY storage control report

The following is an example of an XQUERY storage control report that was generated when the XQUERY SESSION1 STORAGECONTROL command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY STORAGECONTROL REPORT - 003
ANTQ8241I      SC SC S RES  UTIL                  CURRENT
ANTQ8242I SSID SN ID T CNT  VOL    TIMEOUT  --STORAGE CONTROL TIME--
ANTQ8203I ------------------------------------------------------------
ANTQ8243I 0240 B  02   0000 PRIM02 00.05.00 2000.135 11:58:37.459125
ANTQ8243I 8080 B  03   0000 PRIM31 00.05.00 2000.135 11:59:43.583528
ANTQ8243I 8081 A  01   0000 PRIM01 00.05.00 2000.135 11:59:49.534643
ANTQ8243I 1010 A  01   0000 PRIM22 00.05.00 2000.135 11:59:50.526463
ANTQ8243I 1010 B  02   0000 PRIM34 00.05.00 2000.135 11:59:51.474577
ANTQ8243I 1012 A  01   0000 PRIM15 00.05.00 2000.135 11:59:53.563784
ANTQ8238I TOTAL=26 DUP=18  CPY=1  PND=0  SUS=1  SEQ=0  UTL=6
ANTQ8231I DATA CONSISTENT(2000.135 11:59:53.563784)   IDLE(00:00:05.9)
ANTQ8240I DATA EXPOSURE(00:00:00.00)
ANTQ8201I XQUERY STORAGECONTROL REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY configuration report

The following is an example of an XQUERY configuration report that was generated when the XQUERY SESSION1 CONFIGURATION command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY CONFIGURATION REPORT - 003
ANTQ8281I ---------PRIMARY----------  -----SECONDARY------
ANTQ8284I      SC SC
ANTQ8282I SSID SN ID DVCN CCA VOLSER  SSID DVCN CCA VOLSER
ANTQ8203I ------------------------------------------------------------
ANTQ8283I A760 AA 04 0F40 04  PRIM01  A764 0F50 03  SECD01
ANTQ8283I A760 AA 04 0F41 08  PRIM02  A764 0F51 0E  SECD02
ANTQ8283I 1010 AA 01 0F60 02  PRIM03  A764 0F52 01  SECD03
ANTQ8238I TOTAL=3  DUP=3  CPY=0  PND=0  SUS=0  SEQ=0  UTL=0
ANTQ8231I DATA CONSISTENT(2001.018 16:41:15.390719)   IDLE(00:40:06.3)
ANTQ8240I DATA EXPOSURE(00:00:00.00)
ANTQ8201I XQUERY CONFIGURATION REPORT COMPLETE FOR SESSION(SESSION1)
```

The following is an example of an XQUERY configuration report with volume filtering. This report was generated when the XQUERY SESSION1

CONFIGURATION V(PRIM0*) command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY CONFIGURATION REPORT - 003
ANTQ8281I ---------PRIMARY----------  -----SECONDARY------
ANTQ8284I      SC SC
ANTQ8282I SSID SN ID DVCN CCA VOLSER  SSID DVCN CCA VOLSER
ANTQ8203I ------------------------------------------------------------
ANTQ8283I A760 AA 04 0F40 04  PRIM31  A764 0F50 03  SECD01
ANTQ8283I A760 AA 04 0F41 08  PRIM20  A764 0F51 0E  SECD02
ANTQ8283I 1010 AA 01 0F60 02  PRIM01  A764 0F52 01  SECD03
ANTQ8238I TOTAL=3  DUP=3  CPY=0  PND=0  SUS=0  SEQ=0  UTL=0
ANTQ8231I DATA CONSISTENT(2001.018 16:41:15.390719)   IDLE(00:40:06.3)
ANTQ8240I DATA EXPOSURE(00:00:00.00)
ANTQ8201I XQUERY CONFIGURATION REPORT COMPLETE FOR SESSION(SESSION1)
```

Examples of XQUERY reports in a coupled environment

The reports in this section are examples of XQUERY reports that are generated in a coupled environment. Use the name of the master session when you enter the command to generate the report.

XQUERY summary report

The following is an example of an XQUERY summary report that summarizes the operating parameters for the XRC master session. This report was generated when the XQUERY SESSION1 command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS002)
ANTQ8202I XQUERY SUMMARY REPORT - 002
ANTQ8203I ------------------------------------------------------------
ANTQ8238I TOTAL=37  DUP=32  CPY=1  PND=1  SUS=2  SEQ=0  UTL=1
ANTQ8231I DATA CONSISTENT(1998.266 19:41:27.998450)  DELAY(00:00:01.83)
ANTQ8240I DATA EXPOSURE(00:00:00.68)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(VOLUME)  HLQ(SYS1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8246I MSESSION(MSESS1) MHLQ(SYS1) COUPLED(INTERLOCKED)
ANTQ8234I PND MASTER CMD = XSUSPEND ATTIME(1998.267 12:00:00.000000)
ANTQ8236I PND XSUSPEND TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8201I XQUERY SUMMARY REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY volume report

The following is an example of an XQUERY volume report. This report provides a complete volume summary of the status of all volumes in the master session. This report was generated when the XQUERY SESSION1 VOL(ALL) command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY VOLUME REPORT - 003
ANTQ8211I PRIM   SEC     ERROR  SYNCH
ANTQ8212I VOL    VOL     LEVEL    % STA CMD OP -------TIMESTAMP--------
ANTQ8203I ------------------------------------------------------------
ANTQ8213I PRIM01 SECD01 SESSION     DUP
ANTQ8213I PRIM02 SECD02 VOLUME   21 CPY
ANTQ8213I PRIM03 SECD03 SESSION     SEQ
ANTQ8213I PRIM04 SECD04 VOLUME   95 SUS        1997.135 11:59:43.583528
ANTQ8213I PRIM05 SECD05 SESSION  90 SUS        1997.135 12:48:48.071334
ANTQ8213I PRIM06 SECD06 VOLUME      DUP DEL DR 1997.135 12:56:58.276463
ANTQ8213I PRIM07 SECD07 VOLUME      DUP
ANTQ8213I PRIM08 SECD08 VOLUME      DUP SUS AT 1997.135 12:49:25.984596
ANTQ8213I PRIM09 SECD09 VOLUME      DUP SUS DR 1997.135 12:59:28.048573
ANTQ8213I PRIM10 SECD10 VOLUME      DUP
ANTQ8213I PRIM11 SECD11 IMSGROUP 95 CPY DEL AT 1997.135 12:38:18.771624
ANTQ8213I PRIM12 SECD12 IMSGROUP 00 PND DEL DR 1997.135 12:58:21.771624
ANTQ8213I PRIM13 SECD13 VOLUME   57 CPY
ANTQ8213I PRIM14 SECD14 DB2GROUP 00 PND SUS AT 1997.135 12:58:30.894533
ANTQ8213I PRIM15 SECD15 DB2GROUP 83 CPY SUS DR 1997.135 12:58:38.781519
ANTQ8213I PRIM16 SECD16 VOLUME   00 PND
ANTQ8213I PRIM17 SECD17 VOLUME      SEQ SUS AT 1997.135 12:58:30.894533
ANTQ8213I PRIM18 SECD18 GROUP1      SEQ SUS DR 1997.135 12:58:38.781519
ANTQ8213I PRIM19 SECD19 VOLUME      SEQ
ANTQ8213I PRIM20 SECD20 VOLUME   ** SUS        1997.135 12:00:45.936853
ANTQ8213I PRIM21 SECD21 VOLUME   42 SUS        NO_TIME_AVAILABLE
ANTQ8213I PRIM22 XRCUTL          -- UTL
ANTQ8213I PRIM23                 -- NIS
ANTQ8213I PRIM24 XRCUTL          -- UTL
ANTQ8238I TOTAL=23  DUP=6   CPY=4  PND=3  SUS=4  SEQ=4  UTL=2
ANTQ8231I DATA CONSISTENT(1997.135 12:48:27.998450)  DELAY(00:00:00.82)
ANTQ8240I DATA EXPOSURE(00:00:00.56)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(VOLUME)  HLQ(SYS1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8246I MSESSION(MSESS1) MHLQ(SYS1) COUPLED(INTERLOCKED)
ANTQ8201I XQUERY VOLUME REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY set report

The following is an example of an XQUERY set report that was generated when the XQUERY SESSION1 SET command processed. The values shown are the defaults from the state data set or are the values that you have set with the XSET command. This report may be generated in either a coupled or an uncoupled environment.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY SET REPORT - 003
ANTQ8203I ------------------------------------------------------------
ANTQ8261I HLQ(HLQ1)  SCTRAP(OFF)  TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8262I SYNCH(4)  SCSYNCH(2,2)  PRIORITY(LOAD)  SUPRDUMP(OFF)
ANTQ8263I SCBUFFERS(576)  PAGEFIX(8)  PACE(5)  COPY(FULL)
ANTQ8264I RFREQUENCY(00.30.00)  RTRACKS(7500)
ANTQ8201I XQUERY SET REPORT COMPLETE FOR SESSION(SESSION1)
```

z/OS MVS System Messages, Vol 1 (ABA-AOM) contains details about the XRC messages that are listed within the preceding report.

XQUERY storage control report

The following is an example of an XQUERY storage control report for a master session. This report was generated when the XQUERY SESSION1 STORAGECONTROL command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY STORAGECONTROL REPORT - 004
ANTQ8241I      SC SC S RES  UTIL                 CURRENT
ANTQ8242I SSID SN ID T CNT  VOL    TIMEOUT  --STORAGE CONTROL TIME--
ANTQ8203I ------------------------------------------------------------
ANTQ8243I ssid scsn scid status residual_count utility_volume timeout storage_control_time
ANTQ8238I TOTAL=16  DUP=12  CPY=3  PND=0  SUS=0  SEQ=0  UTL=1
ANTQ8231I DATA CONSISTENT(2000.135 11:59:53.563784)  IDLE(00:01:21.4)
ANTQ8240I DATA EXPOSURE(00:01:12.28)
ANTQ8246I MSESSION(MSESS1) MHLQ(SYS1) COUPLED(INTERLOCKED)
ANTQ8201I XQUERY STORAGECONTROL REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY volume detail report

The following is an example of an XQUERY volume detail report for a master session that was generated when the XQUERY SESSION1 VOL(ALL) DETAIL command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY VOLUME_DETAIL REPORT - 002
ANTQ8216I PRIM   SEC     ERROR  SYNCH          RES  THD       SC SC
ANTQ8217I VOL    VOL     LEVEL    % STA CMD OP CNT  CNT  SSID SN ID
ANTQ8203I ------------------------------------------------------------
ANTQ8218I PRIM01 SECD01 SESSION     DUP        0089 0500 8700 -- 01
ANTQ8218I PRIM02 SECD02 SESSION     DUP        0499 0500 8702 -- 01
ANTQ8218I PRIM03 SECD03 SESSION     DUP*       0620 0500 1010 AA 02
ANTQ8218I PRIM04 SECD04 SESSION  76 CPY        0056 0500 1011 -- 02
ANTQ8218I PRIM05 SECD05 SESSION  24 CPY*       05DB 0500 8702 BB 03
ANTQ8218I PRIM06 SECD06 SESSION     DUP        01AB 0500 A760 CC 04
ANTQ8218I PRIM07 SECD07 SESSION     SUS        0000 0500 A760 CC 04
ANTQ8238I TOTAL=7  DUP=4  CPY=2  PND=0  SUS=1  SEQ=0  UTL=0
ANTQ8231I DATA CONSISTENT(2000.325 23:20:49.504484)  DELAY(00:00:02.35)
ANTQ8240I DATA EXPOSURE(00:00:00.92)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(HLQ1)
ANTQ8233I DEFAULT TIMEOUT(STORAGE_CONTROL_DEFAULT)
ANTQ8246I MSESSION(MSESS1) MHLQ(SYS1) COUPLED(NON-INTERLOCKED)
ANTQ8201I XQUERY VOLUME_DETAIL REPORT COMPLETE FOR SESSION(SESSION1)
```

z/OS MVS System Messages, Vol 1 (ABA-AOM) contains details about the XRC messages that are listed within the preceding report.

XQUERY configuration report

The following is an example of an XQUERY configuration report for a master session that was generated when the XQUERY SESSION1 CONFIGURATION command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY CONFIGURATION REPORT - 003
ANTQ8281I ---------PRIMARY----------  -----SECONDARY------
ANTQ8284I      SC SC
ANTQ8282I SSID SN ID DVCN CCA VOLSER  SSID DVCN CCA VOLSER
ANTQ8203I ------------------------------------------------------------
ANTQ8283I A760 AA 04 0F40 04  PRIM01  A764 0F50 03  SECD01
ANTQ8283I A760 AA 04 0F41 08  PRIM02  A764 0F51 0E  SECD02
ANTQ8283I 1010 AA 01 0F60 02  PRIM03  A764 0F52 01  SECD03
ANTQ8238I TOTAL=8  DUP=6  CPY=0  PND=0  SUS=0  SEQ=0  UTL=2
ANTQ8231I DATA CONSISTENT(2001.018 16:41:15.390719)   IDLE(00:40:06.3)
ANTQ8240I DATA EXPOSURE(00:00:00.00)
ANTQ8246I MSESSION(MSESS1) MHLQ(SYS1) COUPLED(INTERLOCKED)
ANTQ8201I XQUERY CONFIGURATION REPORT COMPLETE FOR SESSION(SESSION1)
```

XQUERY MASTER report

The XQUERY MASTER report is produced by ANTAS000 from the current contents of the master coupling data set. It may be issued on any system with access to the data set, and does not require an XRC session to be active on the system where it is issued.

The following is an example of an XQUERY MASTER report that was generated when the XQUERY MSESS1 MASTER MHLQ(SYS1) command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8300I XQUERY STARTED FOR MSESSION(MSESS1) MHLQ(SYS1)
ANTQ8202I XQUERY MASTER REPORT - 002
ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
ANTQ8303I ------------------------------------------------------------
ANTQ8304I SESSION1  ACT      Y       =00:00:00.000000  +00:00:01.429395
ANTQ8304I SESSION1  ACT      Y       =00:00:00.235324  +00:00:01.456778
ANTQ8304I SESSION1  ACT      Y       =00:00:00.567735  +00:00:01.568245
ANTQ8304I SESSION1  ACT      Y       =00:00:00.837235  +00:00:01.569867
ANTQ8304I SESSION1  ACT      Y       =00:00:01.252353  +00:00:00.429764
ANTQ8304I SESSION2  ACT      N       -00:00:17.527111
ANTQ8305I TOTAL=6  ACT=6  SUS=0  END=0  ARV=0  RCV=0  UNK=0
ANTQ8308I MASTER SESSION RECOVERABLE TIME(2000.325 23:20:50.933879)
ANTQ8309I INTERLOCKED=5  NON-INTERLOCKED=1
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(MSESS1)
```

The following is another example of an XQUERY MASTER report that was generated when the XQUERY MSESS1 MASTER MHLQ(SYS1) command processed. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains details about the XRC messages that are listed within the following report.

```
ANTQ8300I XQUERY STARTED FOR MSESSION(MSESS1) MHLQ(SYS1)
ANTQ8202I XQUERY MASTER REPORT - 002
ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
ANTQ8303I ------------------------------------------------------------
ANTQ8304I SESSION1  ACT ERR  Y       =00:00:00.000000  +00:00:02.327186
ANTQ8304I SESSION2  ACT SYM  Y       +00:00:01.331259  =00:00:00.000000
ANTQ8304I SESSION3  ACT SYM  N       -00:00:02.348653
ANTQ8304I SESSION4  SUS AVS  N
ANTQ8305I TOTAL=4  ACT=3  SUS=1  END=0  ARV=0  RCV=0  UNK=0
ANTQ8306I MASTER SESSION STATUS = HOLD
ANTQ8308I MASTER SESSION RECOVERABLE TIME(2000.321 20:48:05.994409)
ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=2
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(MSESS1)
```

Examples of XQUERY ADVANCE and XQUERY RECOVER reports

The reports shown below are examples of XADVANCE and XRECOVER reports. See “Creating a recovery volume report” on page 168 for more information on generating recovery volume reports.

The following is an example of an XQUERY ADVANCE report that was generated when the XADVANCE SESSION1 HLQ(SYS1) command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY ADVANCE REPORT - 001
ANTQ8272I PRIMARY  SECONDARY STA CON  ------TIMESTAMP---------
ANTQ8203I ---------------------------------------------------------
ANTQ8273I PAS001   SES001    SUS YES  2000.322 20:55:00.356210
ANTQ8273I PAS002   SES002    SUS YES  2000.322 20:55:00.356210
ANTQ8273I PAS002   SES002    SUS NO   2000.322 19:53:00.324610
ANTQ8237I TOTAL=3   DUP=0   PND=0   SUS=3
ANTQ8231I DATA CONSISTENT(2000.322 20:55:00.356210)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(SYS1)
ANTQ8201I XQUERY ADVANCE REPORT COMPLETE FOR SESSION(SESSION1)
```

The following is an example of an XQUERY RECOVER report that was generated when the XRECOVER SESSION1 HLQ(SYS1) command processed.

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY RECOVER REPORT - 002
ANTQ8271I ------ORIGINAL-------
ANTQ8274I PRIMARY  SECONDARY STA CON CLP  ------TIMESTAMP---------
ANTQ8203I ---------------------------------------------------------
ANTQ8275I PRIM01   SECD01    DUP YES YES
ANTQ8275I PRIM02   SECD02    DUP YES YES
ANTQ8275I PRIM03   SECD03    SUS NO  YES  2000.322 19:54:00.562356
ANTQ8275I PRIM04   SECD04    PND NO  NO
ANTQ8237I TOTAL=4   DUP=2   PND=1   SUS=1
ANTQ8231I DATA CONSISTENT(2000.322 20:55:00.356210)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(SYS1)
ANTQ8201I XQUERY RECOVER REPORT COMPLETE FOR SESSION(SESSION1)
```

Accessing secondary volumes while XRC is active

Applications can read secondary volumes while XRC continues to copy data to these volumes. However, you should consider the following when accessing secondary volumes from another system:

- The secondary volume serial number is different than the primary volume serial number. Access to the secondary volume must be by a different catalog or by explicit unit, volume, and serial number reference.
- Whenever the VTOC location on a volume changes, the new location must be updated in an in-storage pointer to the VTOC. XRC performs this refresh on the data mover system whenever it copies the primary volume VTOC pointer to the secondary volume. To refresh the pointer on other systems attached to the secondary volume, issue the console commands to vary the secondary volume offline and then online. Failure to do this can result in I/O errors when attempting to access the old VTOC.
- XRC continuously copies application updates to a secondary volume that remains in the XRC session. A program that reads from the secondary volume must be able to manage the fact that the data continues to change.

  Note: VSAM volume data set (VVDS) information will reflect the primary volume, not the secondary volume.

Error conditions can occur during reads, however, due to the way that XRC handles secondary I/O functions. An example of this is an I/O error that results when the VTOC is updated before XRC has completely copied the data to the secondary volume. Applications that access secondary XRC volumes must be able to handle these temporary error situations. You can often resend the read I/O to clear these errors.

Ending an XRC session

Issue the XEND command to end the extended remote copy session prior to recovering a failed system. XRC generates a message that indicates the time of the last known consistent state of all volume pairs. See “XEND–ending a session” on page 70 for the XEND command syntax.

After you issue the XEND command, XRC prompts you (unless you have specified NOVERIFY) to confirm the request to end the session. You must specify either YES or NO. You can include the XEND command that is written in REXX. If you do, you should either specify NOVERIFY or place your response to the prompt in the data stack before you issue the command. In a batch environment, you can stack the response.

The following command ends the XRC session:

```
XEND DALLAS
```

The following command ends the XRC session and routes XRC messages to user ID E299000:

```
XEND DALLAS MSGROUTEID(E299000)
```

The XEND command stops all updates to secondary volumes after the completion of the target consistency group to that secondary volume. XRC sends messages that indicate that the XRC session has ended, and includes the consistency group timestamp for the last applied updates to the secondary volumes.

Applying XRC operational scenarios

The scenarios in this section describe possible ways to manage your XRC sessions to perform or accommodate various planned data management situations. XRC scenarios include preparing for maintenance, moving your data to another site, and changing session values. There are also scenarios for coupled XRC sessions, including coupling new or existing XRC sessions, restarting an inactive XRC coupled session after an error occurs, applying software maintenance in a coupled environment, resuming operations after a sympathetic suspension, and recovering from a disaster.

Before you begin: You need to be familiar with the commands used in this scenario. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

Example: applying maintenance

The following scenario uses XRC commands to suspend an XRC session in preparation for applying maintenance to the data mover system.

1. Start a session named TUCSON for disaster recovery with the state, control, and journal data sets under a high-level qualifier of NEW. Use the command:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

   Note: An error on any volume suspends the volume; other volumes continue to be copied.

2. Allocate 70 MB of fixed real storage to the data mover.

   ```
   XSET TUCSON PAGEFIX(70)
   ```

3. Add a volume pair to XRC session TUCSON.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SEC001) SCSESSION(AB)
   ```

   The error response level defaults to what is specified on the XSTART command. The volume is assigned to storage control session AB.

4. Add another volume pair to the XRC session, specifying the error response level as SESSION. The volume is assigned to the default storage control session "--".

   ```
   XADDPAIR TUCSON VOLUME(PRIM02,SEC002) ERRORLEVEL(SESSION)
   ```

   (Normal data processing continues.)

5. In time, maintenance needs to be applied to the data mover system. The maintenance, which may include IPLing the data mover system or refreshing SYS1.LINKLIB, is expected to take a maximum of six hours. Suspend the XRC session (XSUSPEND) and apply maintenance to the system. If the maintenance is to be applied to the data mover modules in SYS1.LINKLIB, you must cancel the ANTAS000 address space after the refresh of SYS1.LINKLIB to have the maintenance take effect.

   ```
   XSUSPEND TUCSON TIMEOUT(06.00.00)
   ```

   During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

6. Restart the TUCSON session for disaster recovery.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

7. Allocate 70 MB of fixed real storage to the data mover.

   ```
   XSET TUCSON PAGEFIX(70)
   ```

8. Add all suspended volume pairs back to the XRC session.

   ```
   XADDPAIR TUCSON SUSPENDED
   ```

   All ERRORLEVEL and SCSESSION values default to what they were prior to the suspension. XRC resynchronizes all volume pairs.

Example: preparing for disaster

The following scenario uses XRC commands to support a disaster recovery. The scenario objective is to initialize a set of volumes, copy changes on a continuous basis, and recover to a known point in time after the disaster.

1. Start a session named TUCSON for disaster recovery with the state, control, and journal data sets under a high-level qualifier of NEW. Use the command:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

Note: An error on any volume suspends the volume; other volumes continue to be copied.

---

2. Add a volume pair to the XRC session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

   The error response level defaults to what is specified on the XSTART command. XRC assigns the volume pair to the default storage control session "--".

---

3. Add another volume pair to the XRC session, specifying the error response level as SESSION.

   ```
   XADDPAIR TUCSON VOLUME(PRIM02,SECO02) ERRORLEVEL(SESSION)
   ```

   (Normal data processing continues.)

---

4. When a disaster occurs, issue these commands at the recovery site.

   ```
   XEND TUCSON XRECOVER TUCSON HLQ(NEW)
   ```

   The XRECOVER command applies in-process updates to the secondary volumes and clips the secondary volumes to the primary volume serial numbers. The recovery is then complete, with the secondary volumes consistent up to the time reported by the XRECOVER command.

Example: migrating data

The following scenario uses XRC commands to support migration mode. The scenario objective is to migrate a set of volumes from one site to another site, and ensure data consistency to a known point in time after the migration. The migration mode cannot be used for coupled sessions.

1. Start a session named TUCSON for data migration with the state data set under a high-level qualifier of NEW. Use the command:

   ```
   XSTART TUCSON SESSIONTYPE(MIGRATE) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

   Note: An error on any volume suspends the volume; other volumes continue to be copied.

---

2. Add volume pairs to the session for migration.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01 PRIM02,SECO02)
   ```

   The error response level defaults to what is specified on the XSTART command, which is VOLUME. (Normal data processing continues.)

---

3. The application programs must be quiesced prior to issuing this command. After they have been quiesced, this command ensures that all updates in the storage controls have been applied to the secondary volumes.

   ```
   XSUSPEND TUCSON DRAIN TIMEOUT(03.30.00)
   ```

The timeout duration of three hours and 30 minutes means that you plan to restart this session by an XSTART command before this timeout duration has passed.

---

4. Issue the RECOVER command to clip the secondary volumes.

   ```
   XRECOVER TUCSON HLQ(NEW)
   ```

   The migration is now complete with the secondary volumes consistent up to the time reported by the XRECOVER command.

Example: switching system modes

The following scenario uses XRC to support migration of a data center from one site to another site. After the migration, the scenario shows how the XRC session mode can be changed to activate the data mover in a disaster recovery mode.

1. Start a session named TUCSON for data migration with the state data set under a high-level qualifier of NEW.

   ```
   XSTART TUCSON SESSIONTYPE(MIGRATE) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

   Note: An error on any volume suspends the volume; other volumes continue to be copied.

---

2. Add a volume pair to the session for migration.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

   The error response level defaults to what is specified on the XSTART command.

---

3. Add another volume pair to the XRC session, specifying the error response level as VOLUME.

   ```
   XADDPAIR TUCSON VOLUME(PRIM02,SECO02) ERRORLEVEL(VOLUME)
   ```

   (Normal data processing continues.)

---

4. Quiesce the application programs prior to issuing the XSUSPEND command.

   ```
   XSUSPEND TUCSON DRAIN TIMEOUT(03.30.00)
   ```

   In this example, the DRAIN keyword ensures that all updates in the storage controls are applied to the secondary volumes. The timeout duration of three hours and 30 minutes means that you plan to restart this XRC session with an XSTART command before this timeout duration has passed.

   During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

---

5. The application programs can be restarted and the session is restarted as a disaster recovery session. The XSTART command must complete before the storage control timeout expires, or else the XRC session is ended.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME) HLQ(NEW)
   ```

   An error on any volume causes that volume to be suspended; other volumes continue to be copied. The state, control, and journal data sets are under a high-level qualifier of NEW.

---

6. Resynchronize the volume pair in the session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

   The error response level remains what was initially specified prior to the suspension.

---

7. Resynchronize the volume pair in the session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM02,SECO02) ERRORLEVEL(SESSION)
   ```

   The error response level for this volume pair is changed to SESSION. (Normal data processing continues.)

---

8. When a disaster occurs, this command is issued at the recovery site to apply in-progress updates to the secondary volumes and clip the secondary volumes to the primary volume serial numbers.

   ```
   XRECOVER TUCSON HLQ(NEW)
   ```

   The recovery is complete with the secondary volumes consistent up to the time reported by the XRECOVER command.

Example: changing the session error recovery level

The following scenario uses XRC commands to change the XRC session error recovery level.

1. Start a session named TUCSON for disaster recovery.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   An error on any duplex volume causes the XRC session and all volumes in the session to be suspended. If an error occurs, all volumes in the session must be resynchronized. The state, control, and journal data sets are under a high-level qualifier of SYS1.

---

2. Add a volume pair to the XRC session, which inherits ERRORLEVEL(SESSION) from the XSTART command.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

---

3. Add another volume pair to the XRC session, specifying the error recovery level as VOLUME.

   ```
   XADDPAIR TUCSON VOLUME(PRIM02,SECO02) ERRORLEVEL(VOLUME)
   ```

(Normal data processing continues.)

---

4. Suspend the XRC session to the specific timestamp after all updates to the specified UTC timestamp have been applied to the secondary volumes.

   ```
   XSUSPEND TUCSON ATTIME(1998.266 19:41:20.231440) TIMEOUT(03.30.00)
   ```

   The timeout duration of three hours and 30 minutes means that you plan to restart this session by an XSTART command before this timeout duration has passed.

   During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

---

5. Start a session named TUCSON for disaster recovery.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME)
   ```

   The default error recovery level was SESSION. Specifying ERRORLEVEL(VOLUME) changes the default error recovery level to VOLUME for all new volumes that are added to the session. The state, control, and journal data sets are under a high-level qualifier of SYS1.

---

6. Resynchronize these suspended volume pairs.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SEC001,PRIM02,SEC002)
   ```

   Since this is a volume resynchronization, the volume pair keeps its original error recovery level of SESSION.

---

7. Add a volume pair to the session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM03,SEC003)
   ```

   The volume pair inherits the VOLUME error recovery level from the XSTART command's ERRORLEVEL(VOLUME) parameter.

---

8. Add another volume pair to the session with an error recovery level of SESSION.

   ```
   XADDPAIR TUCSON VOLUME(PRIM04,SEC004) ERRORLEVEL(SESSION)
   ```

   If an error occurs on this volume, the entire XRC session is suspended.

Example: changing a volume pair's error recovery level

The following scenario uses XRC commands to change the error recovery level for a volume pair, possibly because the volume is no longer critical to the XRC session.

1. Start a session named TUCSON for disaster recovery, using this command:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

An error on any volume causes the XRC session and all volumes in the session to be suspended. All volumes must be resynchronized.

---

2. Add a volume pair to the XRC session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

   The volume inherits the SESSION error recovery level from the XSTART command. (Normal data processing continues.)

---

3. Suspend the volume.

   ```
   XSUSPEND TUCSON VOLUME(PRIM01)
   ```

   Storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping record application updates in software bitmaps.

---

4. Resynchronize this suspended volume and change the error recovery level from SESSION to VOLUME.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01) ERRORLEVEL(VOLUME)
   ```

Example: suspending and resynchronizing a volume pair

The following scenario describes the XRC commands that are used to suspend and resynchronize a volume pair. You might choose to do perform this scenario if a primary volume will receive many overlapping changes in a short time and, to optimize performance, and only want the final results to be copied.

1. Start a session named TUCSON for disaster recovery, using the command:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   Note: An error on any duplex volume will cause the XRC session and all duplex pairs in the session to be suspended. All volumes must be resynchronized. The state, control, and journal data sets are under a high-level qualifier of SYS1.

---

2. Add a volume pair to the XRC session, with session-level recovery from the XSTART command.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECO01)
   ```

   (Normal data processing continues.)

---

3. Suspend the volume.

   ```
   XSUSPEND TUCSON VOLUME(PRIM01)
   ```

   Storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping record application updates in software bitmaps.

4. Resynchronize the suspended volume.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SEC001)
   ```

   The state, control, and journal data sets are under a high-level qualifier of SYS1.

Example: suspending and resynchronizing a session

The following scenario describes XRC commands that are used for suspending and resynchronizing the XRC session to allow you to change to another virtual server or processor.

1. Start a session named TUCSON for disaster recovery, using the following command:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   Note: An error on any volume will cause the XRC session and all volumes in the session to be suspended. All volumes must be resynchronized. The state, control, and journal data sets are under a high-level qualifier of SYS1.

2. Add the volume pair to the XRC session.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SEC001)
   ```

   The volume inherits session-level recovery from the XSTART command. (Normal data processing continues.)

3. Suspend the XRC session, and make the planned system changes.

   ```
   XSUSPEND TUCSON TIMEOUT(03.30.00)
   ```

   The timeout duration of three hours and 30 minutes means that you plan to restart this session by an XSTART command before this timeout duration has passed.

   During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

4. Restart the XRC session on the new system.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   The state, control, and journal data sets are under a high-level qualifier of SYS1.

5. This command adds back to the session all pairs that were in suspended state at the time that you issued the XADDPAIR command.

   ```
   XADDPAIR TUCSON SUSPENDED
   ```

XRC will start volume synchronization processing, subject to the SYNCH and SCSYNCH values.

Example: changing volume channel connection addresses

When a volume is added to a session by the XADDPAIR command, the data mover uses the volume serial number to determine which storage control the volume belongs to and which channel connection address to use to address the volume. If the channel connection address for an XRC volume pair is changed while the pair is still in an active XRC session, the data mover is unable to perform the copy operation. The following scenario describes how to register a changed channel connection address with the data mover.

Alerting the data mover to a changed volume CCA

1. Volume PRIM01 is added to the XRC session on storage control A with a channel connection address of 123, with the command:

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECD01)
   ```

   (Normal data processing continues.)

   ---

2. As part of a planned configuration change, a service representative will move volume PRIM01 to storage control B with a new channel connection address of 456. Volume PRIM01 must be removed from the XRC session before the channel connection address is changed.

   ```
   XDELPAIR TUCSON VOLUME(PRIM01)
   ```

   ---

3. Volume PRIM01 is added back into the XRC session and is now on storage control B with the channel connection address of 456. You can add the volume to the session using the NOCOPY option if no application update activity has occurred during the switch-over.

   ```
   XADDPAIR TUCSON VOLUME(PRIM01,SECD01)
   ```

Examples: changing TIMEOUT intervals for suspended sessions

When an XRC session is suspended, you must provide a TIMEOUT interval. The data mover communicates this timeout value to all primary storage controls, and the session remains suspended for at least the specified timeout interval.

Changing XSUSPEND TIMEOUT intervals

The following scenario shows how to change a previously specified timeout interval. It may be that the planned outage is taking longer than you thought it would, or that it is causing unanticipated production impact. You have decided to suspend the XRC sessions.

1. Suspend the session named TUCSON. The planned outage is expected to be completed within 12 hours.

   ```
   XSUSPEND TUCSON TIMEOUT(12.00.00)
   ```

   During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do

not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

(Normal data processing continues.)

---

2. After about ten hours you realize that the timeout interval is not long enough. You restart the session named TUCSON, causing the data mover to start recording changes to volumes in the session.

```
XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(VOLUME)
```

---

3. This command starts a new timeout period of five hours:

```
XSUSPEND TUCSON TIMEOUT(05.00.00)
```

Now the XRC session will remain suspended for up to this time value, which is plenty of time to complete the planned work.

Note: During the suspension, storage controls that support hardware bitmapping record application updates in their hardware bitmaps. Storage controls that do not support hardware bitmapping do not record application updates. Instead, updates accumulate in cache until the XSTART command function is complete.

Changing TIMEOUT for all SC sessions

You can change the TIMEOUT value for all volumes in an XRC session by issuing the XSET command. Specifying the TIMEOUT option, with the SSID(ALL) subparameter, changes the time for all storage controls. When a session is suspended, the current timeout value (the last value specified) remains in effect for the restarted sessions.

An XRC session called TUCSON is currently running with a timeout value set to the default of five minutes. You wish to change the timeout value for all of the storage controls to be one minute.

1. After this command has processed, all storage controls will have a timeout value of one minute.

```
XSET TUCSON TIMEOUT(00.01.00)
```

Some of the hardware that is related to the session requires planned maintenance that will take two hours. You will need to suspend the session during this time. You have planned the suspension to happen during a time when application updates are minimal.

---

2. Issue the command:

```
XSUSPEND TUCSON TIMEOUT(02.00.00)
```

For 3990 or 9390 Storage Controls, if there are sufficient updates during this time it is possible that the storage control session may be canceled. For ESS storage subsystems, the hardware bitmap function keeps the storage control session active regardless of the timeout value that you have specified.

(Normal data processing continues.)

3. Because the hardware updates have been completed, you restart the session named TUCSON.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   The data mover automatically sets the timeout value in the storage controls to the value specified by the last XSET command, which was TIMEOUT(00.01.00).

4. XRC adds all volumes back to the session, and resynchronizes them using the hardware or software bitmaps.

   ```
   XADDPAIR TUCSON SUSPENDED
   ```

Changing TIMEOUT for specific SC sessions

You can change the timeout interval for all volumes in an XRC session that belong to a specific SSID by issuing the XSET TIMEOUT command with the SSID subparameter. The timeout interval that you specify will remain in effect across session suspensions.

An XRC session called TUCSON is currently running with a timeout value set to the default of five minutes for all storage controls. You wish to change the timeout value for storage controls 10FC and 10E2 to be one minute.

1. After this command has processed, the 10FC and 10E2 storage controls will have a timeout value of one minute.

   ```
   XSET TUCSON TIMEOUT(00.01.00) SSID(10FC, 10E2)
   ```

   All other storage controls will continue to have a timeout value of five minutes. Some of the hardware that is related to the session requires planned maintenance that will take two hours. You will need to suspend the session during this time. You have planned the suspension to happen during a time when application updates are minimal.

2. Issue the command:

   ```
   XSUSPEND TUCSON TIMEOUT(02.00.00)
   ```

   For 3990 or 9390 Storage Controls, if there are sufficient updates during this time it is possible that the storage control session may be canceled. For ESS storage subsystems, the hardware bitmap function keeps the storage control session active regardless of the timeout value that you have specified.

   (Normal data processing continues.)

3. Because the hardware updates have been completed, you restart the session named TUCSON.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   The data mover automatically sets the timeout value in the storage controls. For SSIDs 10FC and 10E2, it is set to TIMEOUT(00.01.00), the value specified by the last XSET command. The data mover sets all other storage controls to TIMEOUT(00.05.00).

4. XRC adds all volumes back to the session, and resynchronizes them using the hardware or software bitmaps.

```
XADDPAIR TUCSON SUSPENDED
```

Identifying XRC system interactions

The following topics address operational considerations when extended remote copy operations interact with other MVS system components.

XRC varies volumes offline and online

As part of normal operations, XRC ensures that the changes on these secondary volumes are identified to the data mover system. To accomplish this, XRC issues commands on the data mover system to vary secondary volumes offline and online. These actions occur at the following times:

- During initial volume copy, when XRC copies cylinder 0, track 0 to the target volume, the secondary volume is varied offline.
- After a volume pair becomes duplex, XRC varies the secondary volume online.
- During the XRECOVER function, when the secondary volume's serial number is clipped to the primary volume serial number, the secondary volume is varied online to ensure that the volume's in-storage copy of its VTOC and VTOC index is updated.

Other systems that share these secondary volumes will only have their in-storage VTOC pointers updated when these commands are issued.

If I/O errors occur when accessing the VTOC from an attached system, vary the secondary volumes offline and online to correct the condition.

Using the DEVSERV and IDCAMS functions

Determine the status of the storage control's dual copy pair, cache, and NVS by using the DEVSERV and IDCAMS functions. These functions do not report XRC copy volume pairs or states. Issuing the XQUERY TSO command from the data mover system is the only way to obtain information about XRC volumes.

Chapter 7. Managing Coupled XRC Sessions

An enhancement to XRC called Coupled Extended Remote Copy (CXRC) allows XRC sessions to be coupled together to guarantee that all volumes are consistent to the same time across all coupled XRC sessions. With CXRC, the size of configurations XRC can handle increases. CXRC can support installations that have thousands of volumes, whereas XRC is designed to support installations that have hundreds of volumes.

In this chapter

This chapter provides information about CXRC and includes the following sections:

Topic . . .	See page . . .
Overview of CXRC	137
Configuring a coupled XRC environment	138
Managing coupled XRC operations	145
Applying CXRC operational scenarios	158

Overview of CXRC

Coupled Extended Remote Copy (CXRC) is an enhancement to XRC that provides the following benefits:

- Supports large environments that have an expanded number of primary storage controls and DASD volumes, in excess of those supported by a single data mover configuration.

  Note: A large environment can be one that has more than 20 primary storage controls or more than 1000 primary volumes, or both. Installations may have configurations that consist of thousands of volumes in multiple XRC sessions.
- Provides consistent coordination between XRC sessions to help ensure that all of the volumes can be recovered to a consistent point in time.
- Minimally impacts performance to application programs and facilitates rapid recovery in the event of a disaster.
- Allows up to 14 XRC sessions to be coupled to a single master session.

Multiple XRC sessions are possible for those sites that run up to five XRC sessions in a single LPAR; the sessions may or may not be coupled. Multiple XRC sessions on an LPAR can be combined using CXRC if consistency across sessions is required.

Some XRC commands operate only on the coupled group of sessions instead of affecting individual sessions. See “Chapter 5. Administering Your XRC Environment” on page 55 for more information on CXRC TSO commands.

© Copyright IBM Corp. 1994, 2002

Configuring a coupled XRC environment

This section includes information about the following topics:

- The types of sessions configured in a coupled XRC environment
- The CXRC configurations allowed with the data mover installed at the recovery site or at the primary site
- The considerations for defining data sets in a coupled XRC environment

Identifying sessions in a coupled environment

Table 36 describes session types that are used in a coupled environment. For additional information on XRC sessions as they apply in a uncoupled environment, see “XRC session state descriptions” on page 23.

Table 36. Coupled session descriptions

Session	Description
Active Session	The system data mover is running an XRC session. Volumes can be added or deleted. You can issue the XADVANCE command and the XCOUPLE command with ADD and DELETE parameters to individual active XRC sessions.
Coupled Session	An active individual XRC session has been added to a master session.
Inactive Session	The system data mover is not active. You can issue the XADVANCE, XSTART, XRECOVER, and the XCOUPLE commands with the PURGE parameter to individual inactive XRC sessions.
Interlocked Session	This is a coupled XRC session in a master session that has a consistency time before the master journal time and a journal time equal to or after the master recoverable time. The master recoverable time is the minimum of the journal times reported by all XRC sessions. This is the latest time to which all sessions could recover their data in the event of a disaster, and still provide coupled consistency. XRC will recover all interlocked sessions such that all secondary volumes across the coupled interlocked sessions will be consistent. Because coupled consistency requires all volumes to be at the same time, no session is allowed to apply data which is after the master recoverable time.
Master Session	This is a logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session. You can issue the XQUERY, XEND, XSUSPEND, and XCOUPLE RELEASE commands to a master session.
Suspended Session	You can issue the XADVANCE command, XSTART command, XRECOVER command, and the XCOUPLE command with the PURGE parameter to individual suspended XRC sessions.

Table 36. Coupled session descriptions (continued)

Session	Description
XRC Session	You can issue the XCOUPLE ADD command to an XRC session to couple it to the master session. To uncouple the XRC session from the master session, issue the XCOUPLE DELETE command. When you issue the XCOUPLE ADD command to an individual XRC session, the session should have the following characteristics: • A session to be coupled must have a SESSIONTYPE of XRC already specified. A SESSIONTYPE of MIGRATE is not allowed for coupled sessions; the XCOUPLE ADD command will fail. • An XRC session to be coupled must be started with an ERRORLEVEL of SESSION **Example:** The command: `XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)` started a session with the correct SESSIONTYPE and ERRORLEVEL parameters for the session to be able to be coupled.

Choosing a coupled XRC configuration

With CXRC, you may now choose to use the data mover to manage your coupled data sets at the recovery site, at the primary site, or on a remote host apart from the application host.

Figure 6 on page 140 outlines a coupled XRC configuration with the data mover at the recovery location. Figure 7 on page 141 depicts a coupled XRC configuration with the data mover at the primary site. Figure 8 on page 142 shows a coupled XRC configuration with the data mover located at the primary site, but on a remote host apart from the primary host system itself.

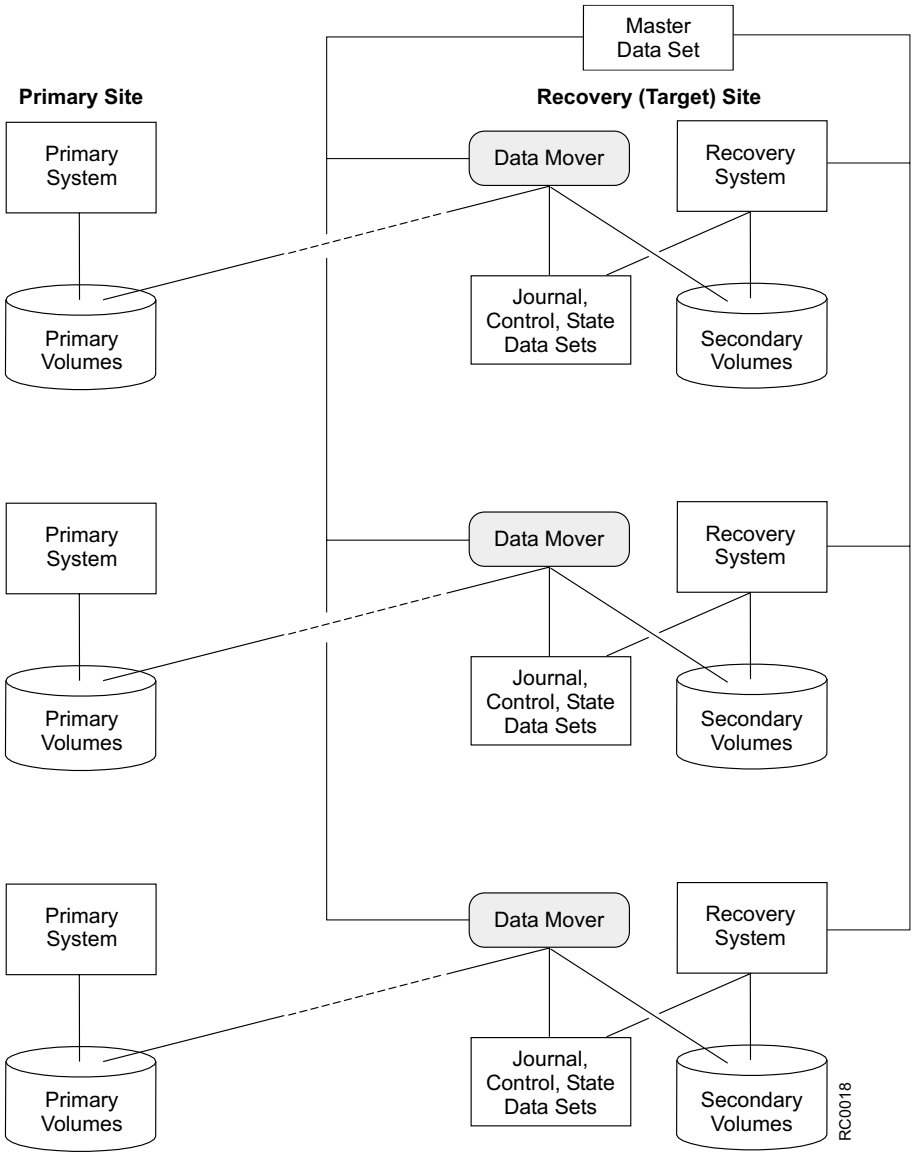

Figure 6. CXRC configuration with data mover at recovery site

Locating the data mover at the recovery site is the most viable scenario for effective disaster recovery, as shown in Figure 6. Because the data mover is not at the primary site, it is protected from any disasters that might negatively impact the primary site, such as floods, storms, fires, earthquakes, bombs, or other accidental or intentional catastrophes. To determine the location of the recovery site, consider the types of disasters that you want to protect your installation's data against. For example, if hurricanes are common in your area, you might want to establish the recovery site at a sufficiently remote location, perhaps 50–100 miles away from the primary site. The flooding and power outages that accompany a hurricane, which might adversely affect your primary site, would be unlikely to affect your remote recovery site.

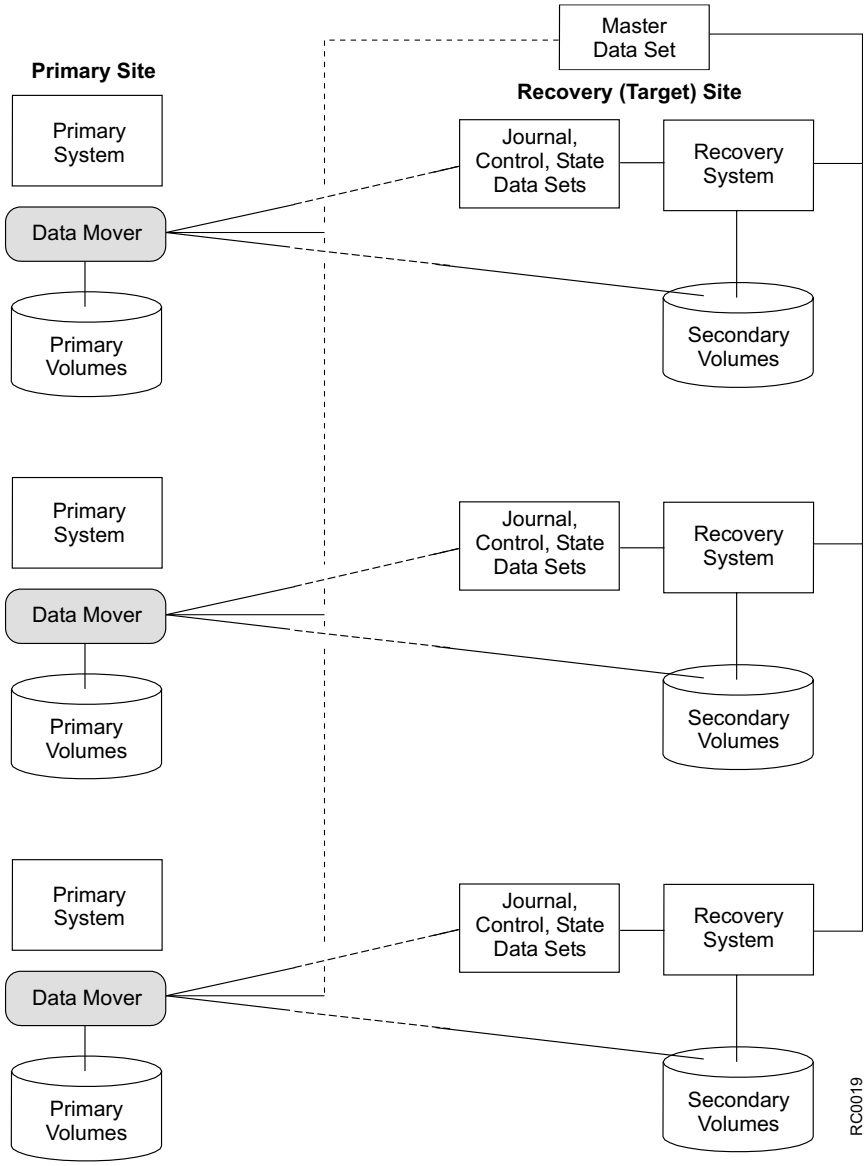

Figure 7. CXRC configuration with data mover at primary site

You may choose to locate the data mover on the primary host system at the primary site, with the recovery host system located remotely from the primary system, as depicted in Figure 7. However, this is the least effective and most vulnerable scenario for disaster recovery. If your application host goes down, the data mover will be unavailable for copy operations to the recovery site. This scenario might be acceptable for short-term testing purposes only.

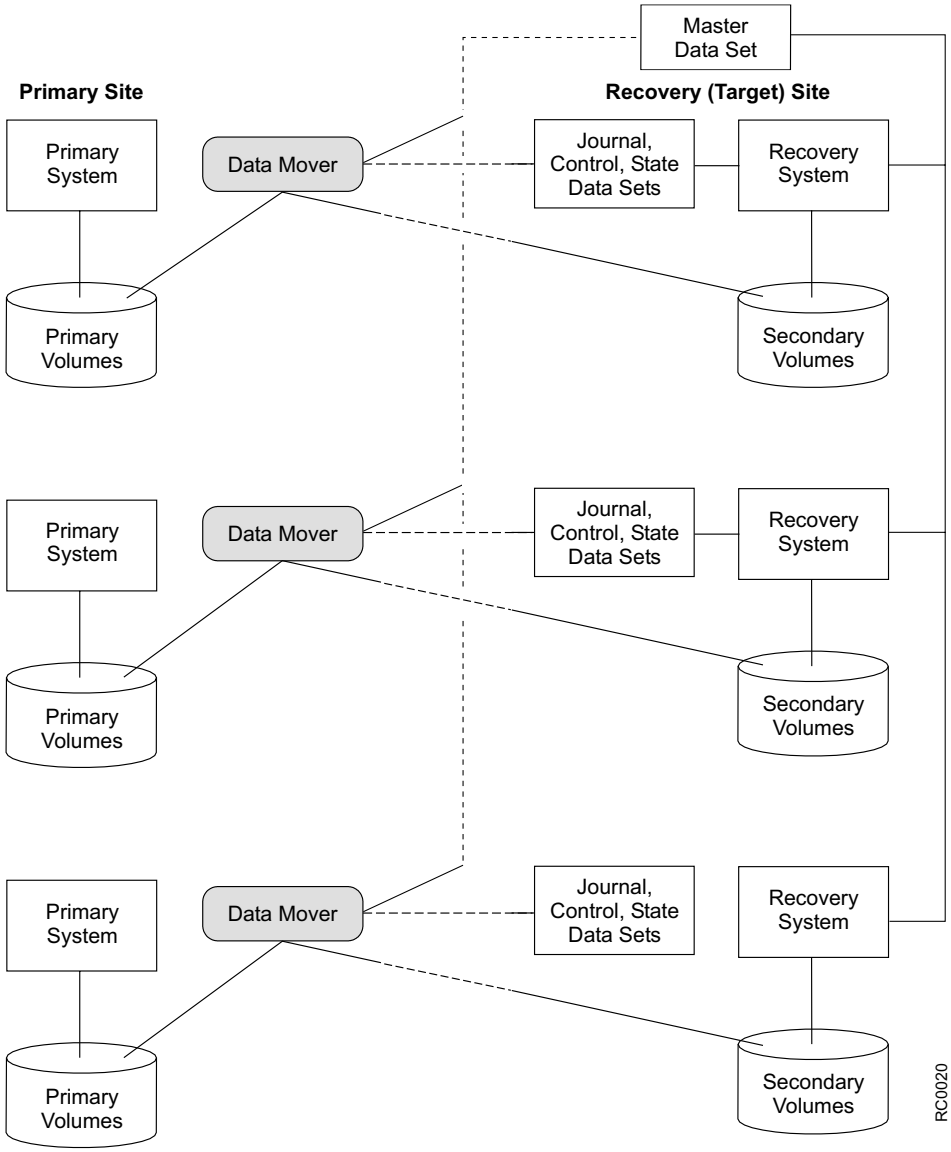

Figure 8. CXRC configuration with data mover at primary site on remote host

Locating the data mover on a remote, separate host at the primary site is an option for configuring your CXRC environment, as shown in Figure 8. There is a certain measure of safety built into having your data mover located on a separate host from your primary system, even if the hosts are at the same site. You run the risk, however, of the same disasters impacting the remote recovery host that affect the primary application host. This option is valid for early CXRC configuration testing or for special testing at any time.

Defining data sets in a coupled environment

Before activating XRC in a coupled environment, you need to allocate the master data set. There are also special considerations for journal, state, and control data sets when you work with master data sets. This section also provides instructions for changing the characteristics of a master data set once they have been established.

Allocating the CXRC master data set

All of the XRC sessions coupled in the master session continuously write to and read from the master data set. This access allows communication between hosts as data is copied from the primary site to the recovery site.

When evaluating the resources you need to successfully configure a CXRC environment, you must plan so that all primary hosts and recovery hosts have access to the master data set.

Rules: You must follow these conditions in order to create the necessary access:

- Make the master data set cataloged and accessible to each host system that processes a coupled session, as well as to the system that processes the XRECOVER command. The master data set name is *mhlq.XCOPY.msession.MASTER*.
- Note that the master data set can reside on the same volume as the state and control data sets, but not on the same volume that contains the journal data sets.
- Allocate the master data set without defining secondary extents.
- Allocate the master data set as physical sequential and not striped.
- Allocate the master data set on a single disk device.
- Note that the required size of the master data set is one cylinder, which allows 14 XRC sessions to be coupled. The required size of the master data set is fixed, which also fixes the maximum number of sessions that can be coupled into one master session.
- Preallocate the master data set size before use of the XCOUPLE command. Only the space that is allocated at the time the XCOUPLE ADD command is issued will be available for XRC use.
- Place the master data set in a user catalog that contains only entries for the master data set.
- Note that the master data set must be in a shared catalog to which all data mover hosts have access, or it must be cataloged on all data mover hosts.
- Protect the master data set from unauthorized use. See “Controlling access to XRC commands” on page 40 for more information on protecting the master data set using RACF.

Allocate the master data set with one cylinder primary space and zero cylinders secondary space as follows:

```
DCB=(RECFM=FB,LRECL=15360,BLKSIZE=15360,DSORG=PS),SPACE=(CYLS,(1,0))
```

Note: For more information about configuring the master data set, see “Specifying the master data set” on page 144.

Changing the characteristics of a master data set

XRC cannot dynamically accept changes to the master data set. If you wish to change the characteristics of the master data set once they have been established, follow this procedure:

1. Remove all coupled sessions from the master session (XCOUPLE DELETE command for active sessions or XCOUPLE PURGE command for inactive sessions).
2. Reallocate the master data set.
3. Add sessions back to the master session (XCOUPLE ADD command).

Specifying the master data set

XRC uses the master data set to keep track of multiple XRC sessions that are coupled to the same master session that is associated with the master data set. The master data set must be accessible from, and cataloged on, all system data mover hosts and all recovery hosts.

XRC uses the master data set name in one of the following forms:

- SYS1.XCOPY.*msession_id*.MASTER
- *mhlq*.XCOPY.*msession_id*.MASTER

Where:

msession_id Specifies the same name that you issued with the MSESSION parameter of the XCOUPLE ADD command. The name can be up to eight characters long. The name you specify for the *msession_id* cannot be the same as any other *session_id* that you are going to couple with the master session.

mhlq Specifies the name of the master data set high-level qualifier that is specified on the MHLQ parameter of the XCOUPLE ADD command. The *mhlq* must be one to eight characters, and can be any name acceptable to TSO. “SYS1” is the default value.

You can change the default *mhlq* to a name that corresponds with the MHLQ that is issued on the XCOUPLE command. When you issue the XSUSPEND, XEND, or XQUERY commands with the MASTER option, use the MHLQ parameter.

The following is an example of a master data set name:

```
SYS1.XCOPY.DALLAS.MASTER
```

Using journal, state, and control data sets

Guidelines: Consider the following guidelines when using journal, state, and control data sets in a coupled environment:

- Allocate the journal, control, and state data sets for the individual sessions in the same manner as the existing support.
- Add a new member to the state data set to indicate that the session is coupled and to allow the XSTART command to process the XCOUPLE ADD command automatically during restart. On a new start, this member is used to uncouple the session.

  Note: The size of this new state data set member is small and does not affect the size requirements for the state data set.
- Define from eight to 16 journal data sets to improve the efficiency of XRC in a large coupled environment.
- Allocate the control data set using physical sequential allocation.
- Note that the state data set is updated whenever you issue XCOUPLE ADD, XCOUPLE DELETE, and XCOUPLE PURGE commands.

If there is insufficient space in the state data set, you will not be able to couple an XRC session to a master session. To allocate a new, larger state data set, perform the following steps:

1. Issue XSUSPEND *msession_id* TIMEOUT(*hh.mm.ss*).
2. Copy all members from the existing state data set to the new, larger data set.
3. Rename the existing state data set to an unused name.

4. Rename the new, larger data set to the name of the pre-existing state data set.
5. Issue the XSTART command and resume XRC operations.

Managing coupled XRC operations

This section contains the following information about managing coupled XRC operations:

Topic . . .	See page . . .
Coupling XRC sessions	145
Querying coupled status for XRC sessions	146
Creating a backup copy of secondary volumes (XADVANCE)	147
Restarting an inactive XRC coupled session	147
What to do if the master session is in coupled HOLD status	150
Using the XCOUPLE RELEASE command	155
What to do if a session is in COUPLE_FAILED status	155
Removing coupled sessions	156
Generating dumps of address spaces	156

Before you begin: You need to be familiar with the commands that are used in the scenarios in this section. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

Coupling XRC sessions

Use the XCOUPLE ADD command to couple an XRC session to a master session. The first time you issue this command for a session, XRC creates a new member in the session’s state data set to indicate that the session is coupled to the specified master session. The member is also used during XSTART processing to determine if the session is to be recoupled to a master session. The XCOUPLE ADD command may only be invoked against XRC sessions that have been started with the SESSIONTYPE(XRC) parameter.

Recommendation: You should specify the ERRORLEVEL(SESSION) parameter for all coupled sessions.

Example: coupling new or existing XRC sessions

The following scenario describes the XRC commands that are used to start a new master session by coupling two XRC sessions. One of the XRC sessions is an existing session and the other is a new session. You can issue commands to the same LPAR or to different LPARs. The master data set for master session CXRC must be accessible from each LPAR.

Perform the following steps to couple new or existing XRC sessions:

1. Start a session named TUCSON, setting the error level for the session using the ERRORLEVEL parameter by entering:

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   This ensures a recoverable consistency across all coupled sessions. See “Specifying the ERRORLEVEL parameter in coupled and uncoupled sessions” on page 175 for information about what happens when an error occurs during master session processing for a coupled session.

2. Add the XRC session named TUCSON to the master session named CXRC, with a default master high-level qualifier of SYS1.

```
XCOUPLE TUCSON ADD MSESSION(CXRC) SYS1
```

3. Start a previously suspended session named PHOENIX.

```
XSTART PHOENIX SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
```

4. Add the XRC session named PHOENIX to the master session named CXRC, with a default master high-level qualifier of SYS1.

```
XCOUPLE PHOENIX ADD MSESSION(CXRC)SYS1
```

5. Add an XRC utility device and two volume pairs to session TUCSON.

```
XADDPAIR TUCSON VOLUME(UTIL01,XRCUTL,PRIM01,SEC001,PRIM02,SEC002)
```

6. Add the previously suspended volumes again to the session PHOENIX.

```
XADDPAIR PHOENIX SUSPENDED
```

7. Issue the following command to generate an XQUERY MASTER report:

```
XQUERY CXRC MASTER
```

Result: The XQUERY MASTER report command generates the following output:

```
ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
ANTQ8202I XQUERY MASTER REPORT - 002
ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
ANTQ8303I ------------------------------------------------------------
ANTQ8304I TUCSON    ACT      Y       =00:00:01.830727  =00:00:00.000000
ANTQ8304I PHOENIX   ACT      Y       +00:00:00.000000  +00:00:02.352970
ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:33:20.664559)
ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
```

Querying coupled status for XRC sessions

To determine the status of each coupled session that is associated with the specified master session, issue the XQUERY command with the MASTER parameter. The XQUERY output indicates the status of the master session, the status of each session, and the overall volume status for each session. In addition, the XQUERY command that is specified with the VOLUME or SUMMARY parameter provides the master session status.

See “Examples of XQUERY reports in a coupled environment” on page 120 for examples of XQUERY MASTER, VOLUME and SUMMARY reports.

Creating a backup copy of secondary volumes (XADVANCE)

The XADVANCE TSO command applies updates to the secondary volumes of an XRC session to a time that is consistent with the volumes of other sessions that are coupled to the master session. This is the same time to which an XRECOVER command updates secondary volumes of a session, except that the XADVANCE command does not cause the secondary volumes to be named with (clipped to) the primary volume volser. The XADVANCE command is used to advance all secondary volumes for all coupled sessions to a consistent time. Once the XADVANCE function has completed, a copy of the secondary volumes may be taken with data that is consistent to this time.

Using the XADVANCE command

Follow the directions in Table 37 for using the XADVANCE command depending on the status of the session.

Table 37. Using the XADVANCE command

If the session is . . .	Then . . .	Action . . .
Active	You may issue the XADVANCE command.	Suspend all volumes in the session and issue the command to the LPAR that the session is active in.
Inactive	You may issue the XADVANCE command.	You can issue the XADVANCE command from any LPAR.
In COUPLE_FAILED status	The XADVANCE command does not function against a session that is in COUPLE_FAILED status.	You must fix the error that occurred before you can issue the command. See “What to do if a session is in COUPLE_FAILED status” on page 155 for instructions on how to remove the COUPLE_FAILED status for a session.

Restarting an inactive XRC coupled session

You can use the following procedure to restart inactive XRC sessions. Issue the XQUERY MASTER command to determine if an XRC session is inactive. The session status of **END**, **UNK**, or **SUS** on the ANTQ8304I message may indicate that a coupled XRC session is inactive. You can issue the XQUERY SUMMARY command on the system from which the XRC session is running to determine if the session is active.

To restart inactive XRC sessions, perform the following tasks:

1. Issue the XADVANCE command for all sessions that indicate that the RCV/ADV DELTA time is positive (that is, the session consistency time is consistent with the master journal time). You must issue the command on the system on which the session was running.

   Note: After the XADVANCE command successfully completes, the data on the secondary devices is consistent with the master journal time. You can now create backup copies of your secondary devices as indicated by your installation backup procedures.

2. Issue the XSTART command for each session.

3. Issue the XQUERY MASTER command.

   Note: You must perform steps 1 through 3 for each inactive session before issuing the XRELEASE command.

---

4. Issue the XCOUPLE command with the RELEASE option if the ANTQ8306I message indicates that the master session is in HOLD status. (See “What to do if a session is in COUPLE_FAILED status” on page 155 for additional information regarding the HOLD status for a master session.)

---

5. Issue the XADDPAIR command for each pair in the session. If volume pairs are in the SUS state (suspended), you can use the SUSPENDED option of this command to add all of these volume pairs again.

Example: restarting an inactive XRC coupled session after an error occurs

The following scenario describes how to restart an inactive coupled XRC session after an error has occurred on the session. Issue the XQUERY MASTER command to determine if an XRC session is inactive. The session status of UNK may indicate that a coupled XRC session is inactive. You can issue an XQUERY SUMMARY command on the system from which the XRC session is running to determine if the session is still active.

First, an error occurs on session PHOENIX; session TUCSON reacts and suspends all volumes to the msession recovered time.

Result: This is an example of the message that was generated:

```
ANTC8402W XRC SESSION(TUCSON) COUPLED TO MSESSION(CXRC) DETECTED
COUPLING_TIMEOUT CONDITION IN XRC_SESSION(PHOENIX),
ANTC8402W (CONT) SUSPENDING ALL VOLUMES

ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM01,SEC001) FOR SESSION(TUCSON)
   AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:35:21.643364)
ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM02,SEC002) FOR SESSION(TUCSON)
   AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:35:21.643364)
```

Perform the following steps to restart an inactive XRC coupled session after an error occurs:

1. Issue the following command to generate an XQUERY MASTER report:

   ```
   XQUERY CXRC MASTER
   ```

   Result: The following output is generated from an XQUERY MASTER report for the master session named CXRC, with a default master high-level qualifier of SYS1:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT SYM  Y       +00:00:03.824936  =00:00:00.000000
   ANTQ8304I PHOENIX   UNK      Y       =00:00:00.000000  +00:00:03.470722
   ANTQ8305I TOTAL=2  ACT=1  SUS=0  END=0  ARV=0  RCV=0  UNK=1
   ANTQ8306I MSESSION STATUS=HOLD
   ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:35:21.643364)
   ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

2. You must determine why the XRC session PHOENIX did not respond so you can resolve the error. The XRC session PHOENIX may have ended for one of the following reasons:
   - The LPAR on which the XRC session PHOENIX was running had to be IPLed (Initial Program Load).
   - The XRC address space was canceled for this XRC session.
   - The LPAR where this XRC session was running lost connectivity to the master data set.

3. Correct the problem for XRC session PHOENIX.

4. The XRC session PHOENIX has a positive RCV/ADV DELTA time. If you wish to apply this journal data to this session, issue the XADVANCE command on the system from which PHOENIX is running.

   ```
   XADVANCE PHOENIX
   ```

   The XADVANCE command brings the consistency time of the session up to the msession recoverable time.

   Result: The processing of the XADVANCE command generates the following XQUERY output:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT SYM  Y       +00:00:03.824936  =00:00:00.000000
   ANTQ8304I PHOENIX   UNK      Y       =00:00:00.000000  +00:00:00.000000
   ANTQ8305I TOTAL=2  ACT=1  SUS=0  END=0  ARV=0  RCV=0  UNK=1
   ANTQ8306I MSESSION STATUS=HOLD
   ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:35:21.643364)
   ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

5. You can perform your installation backup procedures to back up the data on your secondary devices for both sessions. Both sessions are consistent to the msession recoverable time.

6. Issue the XSTART command to restart the XRC session PHOENIX.

   ```
   XSTART PHOENIX SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

7. Issue the following command to generate an XQUERY MASTER report:

   ```
   XQUERY CXRC MASTER
   ```

   Result: The following output is generated from an XQUERY MASTER report for the master session named CXRC, with a default master high-level qualifier of SYS1:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT SYM  Y       +00:00:03.824936  =00:00:00.000000
   ```

```
ANTQ8304I PHOENIX   ACT ERR  Y       =00:00:00.000000  +00:00:00.000000
ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
ANTQ8306I MSESSION STATUS=HOLD
ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:35:21.643364)
ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
```

8. Issue the XCOUPLE command with the RELEASE option to release the master session named CXRC, with a default master high-level qualifier of SYS1, from HOLD status.

   ```
   XCOUPLE CXRC RELEASE SYS1
   ```

   See “What to do if the master session is in coupled HOLD status” for additional information regarding the HOLD status for a master session.

9. Issue the XQUERY command after you issue the XCOUPLE command to ensure that the status of the master session named CXRC, with a default master high-level qualifier of SYS1, is no longer in HOLD status.

   ```
   XQUERY CXRC MASTER
   ```

   Result: The command will generate the following output:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT AVS  N
   ANTQ8304I PHOENIX   ACT AVS  N
   ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
   ANTQ8308I MSESSION RECOVERABLE TIME(NO_TIME_AVAILABLE)
   ANTQ8309I INTERLOCKED=0  NON-INTERLOCKED=2
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

10. You can now add the volume pairs to the XRC session PHOENIX with the command:

    ```
    XADDPAIR PHOENIX(VOL1,VOL2)
    ```

What to do if the master session is in coupled HOLD status

A CXRC master session is placed into the coupled HOLD status when XRC determines that the master recoverable time must remain unchanged to preserve the data consistency across all of the interlocked XRC sessions coupled to the master session.

XRC provides this data consistency function to allow manual backup of data for these sessions. The master recoverable time is the consistency time to which these coupled sessions can be recovered. This master recoverable time is protected and will not be changed when the master session is in the coupled HOLD status. XRC also disallows all actions that cause the master recoverable time to change.

Determining coupled HOLD status

You can issue the XQUERY command with the MASTER parameter to determine if a master session is in the coupled HOLD status.

Entering coupled HOLD status

The master session enters the coupled HOLD status when one of the events in Table 38 occurs.

Table 38. Triggering events for coupled hold status

If the event is . . .	Then . . .
XRECOVER command	For the first XRECOVER command you issue for any of the coupled sessions associated with the master session, XRC places the master session in the coupled HOLD state to ensure a consistent recovery time for all sessions. The master session remains in HOLD status until the last coupled session is recovered.
Error in Session	An interlocked coupled session that detects an error suspends all duplex volumes in the session and places the master session in HOLD status.
Abnormal End or Cancel of a Session	A coupled session that abnormally ends (ABENDs) or is ended by a cancel command can cause the master session to be placed in HOLD status when the coupled session is restarted.
Sympathetic Suspend of Session	An interlocked coupled session suspends all volumes in the session if one of the following occurs: • The interlocked coupled session detects that all volumes have been suspended for another interlocked coupled session, because of an error. • The interlocked coupled session detects that another interlocked coupled session has not updated the master session for some time. XRC changes the master session state to coupled HOLD state when this occurs. The following are situations that will cause a sympathetic suspend: • An error occurs to a volume of a duplex volume pair that was added with the ERRORLEVEL(SESSION) parameter. • An error occurs to a session that forces the session's address space to be restarted, such as a session level error. • A condition occurs where a session is not updating the master data set to indicate a current session consistency time. See "What to do if a session is in COUPLE_FAILED status" on page 155 for a description of the COUPLE_FAILED status.

Restrictions when the master session Is in coupled HOLD status

The following commands are not allowed for any of the coupled sessions while the master session is in the coupled HOLD status because their functions would cause the master recoverable time to change:

- XCOUPLE with the ADD parameter for a new session
- XADDPAIR

Clearing a master session from coupled HOLD status

Before you begin: Before clearing a master session from the coupled HOLD status, you must follow the procedure described in "Restarting an inactive XRC coupled session" on page 147.

Attention: After you have issued the XCOUPLE command with either the RELEASE, DELETE, or PURGE parameter, data for the specified session may not be consistent with the other sessions in the master session.

Clear the coupled HOLD status of a master session with one of the methods listed in Table 39.

Table 39. Clearing coupled HOLD status

When . . .	Then . . .	Notes . . .
All XRECOVER commands complete processing	The HOLD status clears automatically.	none
An error was detected for a session	After you have corrected the error and performed the necessary backup procedures, issue the XCOUPLE RELEASE command.	See “Using the XCOUPLE RELEASE command” on page 155.
An error was detected for a session, and the error cannot be fixed or determined	Uncouple all sessions by issuing the XCOUPLE DELETE command for active sessions and the XCOUPLE PURGE for all inactive sessions.	See “Using the XCOUPLE DELETE and XCOUPLE PURGE commands” on page 156.

Example: resuming operations after a sympathetic suspension

An error has occurred in the coupled session named TUCSON. Because you specified the error level as SESSION on the ERRORLEVEL parameter when you issued the XSTART command, XRC suspends all volume pairs in TUCSON. The session named PHOENIX issues an ANTC8402W message to indicate that a sympathetic suspend must be done and why. The session named PHOENIX suspends all volumes in the session.

The following scenario describes how to resume operations after a sympathetic suspension that is caused by an error.

An error occurs on a secondary device for session TUCSON. Session PHOENIX suspends all volumes to the msession recoverable time.

Result: This is an example of the message that was generated:

```
ANTX5104E XRC SESSION(TUCSON) ENCOUNTERED AN ERROR PROCESSING
VOLUME(SEC001) IN VOLUME PAIR(PRIM01,SEC001), RC=613 REAS=0

ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM01,SEC001) FOR SESSION(TUCSON)
  AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:33:21.680331)

ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM02,SEC002) FOR SESSION(TUCSON)
  AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:33:21.680331)

ANTX8120I ALL VOLUMES IN SESSION(TUCSON) ARE NOW SUSPENDED

ANTC8402W XRC SESSION(PHOENIX) COUPLED TO MSESSION(CXRC) DETECTED
  HOLD_STATUS CONDITION IN MSESSION(CXRC),
ANTC8402W (CONT) SUSPENDING ALL VOLUMES

ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM03,SEC003) FOR SESSION(PHOENIX)
 AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:33:24.429661)

ANTV8107I SUSPEND COMPLETE FOR VOLUME PAIR(PRIM04,SEC004) FOR SESSION(PHOENIX)
```

```
  AT CONSISTENCY_GROUP TIME(2001.080 ANTV8107I (CONT) 22:33:24.429661)

ANTX8120I ALL VOLUMES IN SESSION(PHOENIX) ARE NOW SUSPENDED
```

Perform the following steps to resume operations after a sympathetic suspension:

1. Generate an XQUERY MASTER report for the master session named CXRC, with a default master high-level qualifier of SYS1 by entering:

   ```
   XQUERY CXRC MASTER
   ```

   The report indicates the status of the XRC sessions coupled to the master session.

   Result: The output from this report allows you to determine whether an error has occurred in the coupled session named TUCSON. It also shows the HOLD status of the master session.

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT ERR  Y       =00:00:00.000000  +00:00:02.749330
   ANTQ8304I PHOENIX   ACT SYM  Y       +00:00:01.253732  =00:00:00.000000
   ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
   ANTQ8306I MSESSION STATUS=HOLD
   ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:33:24.429661)
   ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

   ---

2. Correct the problem for XRC session TUCSON.

   In this case, the RC=613 on the ANTX5104E message indicates that a hardware error was detected writing to a secondary device. You will need to determine the cause of the problem so you can resolve it.

   ---

3. Update the secondary volumes of the session named TUCSON to msession recoverable time, that is, to a time that is consistent with the volumes of the session named PHOENIX.

   ```
   XADVANCE TUCSON
   ```

   ---

4. Back up secondary devices for both sessions.

   You can perform your installation backup procedures to back up the data on your secondary devices for both sessions. Both sessions are consistent to the msession recoverable time.

   ---

5. Generate an XQUERY MASTER report for the master session named CXRC, with a default master high-level qualifier of SYS1.

   ```
   XQUERY CXRC MASTER
   ```

   Result: The following output is generated:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT ERR  Y       =00:00:00.000000  +00:00:00.000000
   ANTQ8304I PHOENIX   ACT SYM  Y       +00:00:01.253732  =00:00:00.000000
   ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
   ```

```
ANTQ8306I MSESSION STATUS=HOLD
ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:33:24.429661)
ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
```

6. Issue the XCOUPLE command with the RELEASE option to release the master session named CXRC, with a default master high-level qualifier of SYS1, from the HOLD status.

   ```
   XCOUPLE CXRC RELEASE SYS1
   ```

   See “What to do if the master session is in coupled HOLD status” on page 150 for additional information regarding the HOLD status for a master session.

7. Issue the command:

   ```
   XQUERY CXRC MASTER
   ```

   After the RELEASE function has completed, the recoverability status of each session is changed to **N** (meaning coupled NON-INTERLOCKED status).

   Result: The XQUERY MASTER command generates the following output for the master session CXRC, with the change in status displayed:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT AVS  N
   ANTQ8304I PHOENIX   ACT AVS  N
   ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
   ANTQ8308I MSESSION RECOVERABLE TIME(NO_TIME_AVAILABLE)
   ANTQ8309I INTERLOCKED=0  NON-INTERLOCKED=2
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

8. Add all suspended volumes back into the coupled session named TUCSON.

   ```
   XADDPAIR TUCSON SUSPENDED
   ```

9. Add all suspended volumes back into the coupled session named PHOENIX.

   ```
   XADDPAIR PHOENIX SUSPENDED
   ```

10. Issue the command:

    ```
    XQUERY CXRC MASTER
    ```

    Result: After you add the volume pairs again, the recoverability status is changed to **Y** (sessions are now interlocked), as displayed by the XQUERY MASTER output for master session CXRC with a default master high-level qualifier of SYS1.

    ```
    ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
    ANTQ8202I XQUERY MASTER REPORT - 002
    ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
    ANTQ8303I ------------------------------------------------------------
    ANTQ8304I TUCSON    ACT      Y       =00:00:01.830727  =00:00:00.000000
    ANTQ8304I PHOENIX   ACT      Y       +00:00:00.000000  +00:00:02.352970
    ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
    ```

```
ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:43:21.526671)
ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC) MHLQ(SYS1)
```

Using the XCOUPLE RELEASE command

You can use the XCOUPLE RELEASE command to remove the coupled HOLD status for a master session. In addition, the coupled status of each session is changed to NON-INTERLOCKED. Use caution when you issue this command. See "XCOUPLE RELEASE command considerations" for more information.

The following is a list of cases where you would use the XCOUPLE RELEASE command:

- After you issue the XQUERY command with the MASTER option.
- After you issue the XADVANCE command for all sessions that do not indicate an RCV/ADV DELTA of zero (0) on the XQUERY MASTER report. The XADVANCE command recovers the data for these sessions to the master recoverable time.
- After you perform manual backups of secondary volumes as required by your system administrator.

In addition, if you issue the XCOUPLE command with either the DELETE or PURGE parameter for the last coupled session in a master session, XRC removes the coupled HOLD status for the master session. See "Removing coupled sessions" on page 156 for usage of the XCOUPLE PURGE and XCOUPLE DELETE commands.

XCOUPLE RELEASE command considerations

When you issue the XCOUPLE RELEASE command, XRC removes the coupled HOLD status of the master session. As a result, all coupled sessions in the master session remain coupled to this master session, and the coupled status for each session changes to NON-INTERLOCKED. You can then issue the XADDPAIR command (to add volume pairs to coupled sessions) and the XCOUPLE ADD command (to add new sessions to the master session). In addition, any update writes to primary volumes of coupled sessions changes the master recoverable time.

Attention: After you have issued the XCOUPLE command with either the RELEASE, DELETE, or PURGE parameter, data for the specified session may not be consistent with the other sessions in the master session.

What to do if a session is in COUPLE_FAILED status

The XCOUPLE_FAILED status for a session occurs when a coupled XRC session is not able to read or write to the master data set because of an error. XRC issues a message to indicate the particular error that occurred. All volumes in the session are suspended. The session remains coupled to the master session.

Removing the XCOUPLE_FAILED status for a session

The master data set must be fixed in order to remove this XCOUPLE_FAILED status. Issue the XCOUPLE ADD command for the session after the master data set has been fixed. The XCOUPLE_FAILED status is then removed for the session.

When you issue the XCOUPLE DELETE command, XRC removes the XCOUPLE_FAILED status for the session. This command uncouples the session

from the master session. You may need to use the XCOUPLE DELETE command if you cannot fix the master data set for the session.

Removing coupled sessions

To remove coupled sessions from a master session, use the XCOUPLE DELETE command and the XCOUPLE PURGE command.

Using the XCOUPLE DELETE and XCOUPLE PURGE commands

The XCOUPLE DELETE command removes an active coupled session from a master session independent of the status of the coupled session. You must issue the XCOUPLE DELETE command from the LPAR in which the session is currently active. If an XQUERY MASTER command still indicates that the session is coupled to the master session after you have issued the XCOUPLE DELETE command, issue the XCOUPLE PURGE command.

The XCOUPLE PURGE command removes an inactive coupled session from a master session. You can issue the XCOUPLE PURGE command from any LPAR that has access to the associated master data set. XCOUPLE PURGE processing ensures that the session is not active in another LPAR in the same GRS ring.

The XCOUPLE DELETE and XCOUPLE PURGE commands may fail to delete the COUPLE member of the session's state data set because of an error. In this case, you will have to delete the COUPLE member using a data set utility program if you intend to add the session as a coupled session again.

Attention: After you have issued the XCOUPLE command with either the RELEASE, DELETE, or PURGE parameter, data for the specified session may not be consistent with the other sessions in the master session.

Generating dumps of address spaces

Use these procedures to help you generate the following address space dumps:

- Multiple dumps in a single LPAR
- Dumps across multiple LPARs

Steps for generating multiple dumps in a single LPAR

With the advent of CXRC, you may need to generate dumps of multiple ANTAS*nnn* address spaces at the same time. This cannot be done using a sequence of modify dump commands to the ANTAS*nnn* address space, because the subsequent dumps are not generated until the first one is finished, and they would not be done on a timely basis.

Use the following procedure to generate multiple dumps in a single LPAR:

1. From the console, enter:

   ```
   DUMP COMM=(MULTI EXAMPLE)
   ```

   where *MULTI EXAMPLE* can be any identifying text.

   ---

2. In response to the following prompt:

   ```
   xx IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
   ```

Enter the following command:

```
R xx,JOBNAME=(ANTAS*),DSPNAME=('ANTAS*'.SYSANT*),CONT
```

---

3. In response to the following prompt:

   ```
   yy IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
   ```

   Enter the following command:

   ```
   R yy,SDATA=(ALLNUC,PSA,SQA,CSA,LPA,TRT,SUM,LSQA,RGN),END
   ```

   If you are able to fit the dump operand response on one line at the console, in response to the following prompt:

   ```
   xx IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
   ```

   Enter the following command:

   ```
   R xx,JOBNAME=(ANTAS*),DSPNAME=('ANTAS*'.SYSANT*),
        SDATA=(ALLNUC,PSA,SQA,CSA,LPA,TRT,SUM,LSQA,RGN),END
   ```

If you are generating several dumps, and do not want to specify the parameters repeatedly, use the CHNGDUMP command (abbreviated as *CD*). Use *CD* to change the values for the SDUMP dump type (which is used by XRC) to enter the following command:

```
CD SET,SDUMP=(ALLNUC,PSA,SQA,CSA,LPA,TRT,SUM, LSQA, RGN)
R n1,JOBNAME=(ANTAS*),CONT
R n2,DSPNAME=('ANTAS*'.SYSANT*),END
```

To put commands in a testcase to automatically generate a dump of multiple address spaces at a certain interval, use the following EXPANDER command sequence:

```
CONSOLE CD SET,SDUMP=(ALLNUC,PSA,SQA,LPA,TRT,SUM,LSQA,RGN)
REPLY IEE094D,*,1,* JOBNAME=(ANTAS*),CONT
REPLY IEE094D,*,2,* DSPNAME=('ANTAS*'.SYSANT*)END
WAITFOR ANTC8402W 1 (or whatever message signals the condition, omit if not needed)

CONSOLE DUMP COMM=(ANTC8402W)
```

Generating dumps across multiple LPARs

Note: The above DUMP command options can be specified within the REMOTE parameter of the DUMP command if you wish to generate a dump of all XRC address spaces in a sysplex. Use the SYSLIST subparameter to list the names of the different systems to which you want the dump command sent.

The syntax for the REMOTE parameter is:

```
REMOTE(SYSLIST=(sysname1('ANTAS*'),sysname2
 ('ANTAS*')),DSPNAME=('ANTAS*'.SYSANT*))
```

You must specify the PROBDESC=SYSDLOCL parameter when you specify the REMOTE parameter (as indicated in the following example). If you specify the name of the local system in the SYSLIST list, two dumps of the local system will be generated. Therefore, specify the dump options (as indicated above) for the local system, and then specify the dump options for the remote systems.

The following example lists all responses that are needed to generate a dump of the local system, as well as the two other systems:

```
DUMP COMM=(DUMP LOCAL AND 2 OTHER SYSTEMS)
R n1,JOBNAME=(ANTAS*), DSPNAME=('ANAS*'.SYSANT*),CONT
R n2,SDATA=(ALLNUC,PSA,SQA,CSA,LPA,TRT,SUM,LSQA,RGN),CONT
R n3,PROBDESC=SYSDLOCL,CONT
R n4,REMOTE=(SYSLIST=(sysname2('ANTAS*'),sysname3
 ('ANTAS*')),DSPNAME,SDATA),END
```

Applying CXRC operational scenarios

The scenarios in this section describe possible ways to manage your CXRC sessions to perform or accommodate various planned data management situations. The following CXRC scenarios are included:

Topic . . .	See page . . .
Example: recovering from a disaster using the XRECOVER command	158
Example: applying software maintenance in a coupled environment	159

Before you begin: You need to be familiar with the commands that are used in these scenarios. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

Example: recovering from a disaster using the XRECOVER command

The following scenario describes how to use the XRECOVER command to recover from a disaster. All volumes in sessions coupled to the master session have been suspended as a result of the disaster.

Perform the following steps to suspend coupled sessions and to bring data on the sessions to a consistent, recovered state using the XRECOVER command:

1. Generate an XQUERY MASTER report for the master session named CXRC, with a default master high-level qualifier of SYS1 by entering:

   ```
   XQUERY CXRC MASTER DATASET(RPT1) DISP(OLD)
   ```

   The report indicates the status of the XRC sessions that are coupled to the master session. This allows you to determine that the XRC sessions TUCSON and PHOENIX are still active, but have an error status due to the disaster that has occurred. It also shows the status of the master session, which is HOLD.

   Result: An excerpt of the report is shown below. For an example of the entire report, see “XQUERY MASTER report” on page 123.

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    ACT SYM  Y       +00:00:03.824936  =00:00:00.000000
   ANTQ8304I PHOENIX   ACT ERR  Y       =00:00:00.000000  +00:00:02.347483
   ANTQ8305I TOTAL=2  ACT=2  SUS=0  END=0  ARV=0  RCV=0  UNK=0
   ANTQ8306I MSESSION STATUS=HOLD
   ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:35:21.643364)
   ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

   After the report has been generated, XRC directs the XQUERY output to the data set named RPT1 after first clearing the data set.

2. Suspend the XRC session named TUCSON.

   ```
   XSUSPEND TUCSON TIMEOUT(STORAGE_CONTROL_DEFAULT)
   ```

3. Suspend the XRC session named PHOENIX.

   ```
   XSUSPEND PHOENIX TIMEOUT(STORAGE_CONTROL_DEFAULT)
   ```

   Result: The following message is generated:

   ```
   ANTQ8300I XQUERY STARTED FOR MSESSION(CXRC) MHLQ(SYS1)
   ANTQ8202I XQUERY MASTER REPORT - 002
   ANTQ8302I SESSION   STA VOL INT CMD   JOURNAL DELTA      RCV/ADV DELTA
   ANTQ8303I ------------------------------------------------------------
   ANTQ8304I TUCSON    SUS SYM  Y                          =00:00:00.000000
   ANTQ8304I PHOENIX   SUS ERR  Y                          +00:00:02.347483
   ANTQ8305I TOTAL=2  ACT=0  SUS=2  END=0  ARV=0  RCV=0  UNK=0
   ANTQ8306I MSESSION STATUS=HOLD
   ANTQ8308I MSESSION RECOVERABLE TIME(2001.080 22:35:21.643364)
   ANTQ8309I INTERLOCKED=2  NON-INTERLOCKED=0
   ANTQ8301I XQUERY MASTER REPORT COMPLETE FOR MSESSION(CXRC)
   ```

Example: applying software maintenance in a coupled environment

The commands that are listed allow you to suspend coupled sessions so that you can perform maintenance on individual sessions, and then add the individual sessions back to the master session after maintenance has been completed.

The following scenario describes how to perform software maintenance operations in a coupled environment:

1. Suspend all sessions coupled to the master session, which is named CXRC and has a default master high-level qualifier of SYS1, on the day and time indicated by the ATTIME parameter by entering:

   ```
   XSUSPEND CXRC ATTIME(2001.145 18:00:00) TIMEOUT(04.00.00)
   ```

   All sessions will be suspended with their active duplex volumes at a consistent time across the coupled sessions. You may now apply whatever maintenance processes you wish. It is assumed that maintenance will take no more than four hours to complete.

2. Restart the XRC session named TUCSON. After software maintenance operations have completed, take copies of secondary volumes, if needed, to maintain consistency during the resynchronization process.

   ```
   XSTART TUCSON SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

3. Issue the command:

   ```
   XSTART PHOENIX SESSIONTYPE(XRC) ERRORLEVEL(SESSION)
   ```

   This restarts the XRC session named PHOENIX.

4. After software maintenance operations have completed, take copies of secondary volumes, if needed, to maintain consistency during the resynchronization process.

Note: Doing things in the above order allows XRC to start tracking changes to volumes that are not on a 2105 controller.

Chapter 8. XRC Data Recovery Operations

Extended remote copy (XRC) provides the ability to recover data on the recovery system in the event of a disaster to the primary system. For information about using XRC for data migration, see “Chapter 9. Migrating Data with XRC” on page 171.

In this chapter

This chapter provides information about data recovery with XRC and includes the following sections:

Topic . . .	See page . . .
Understanding the XRC volume synchronization process	161
Recovering data with XRC—uncoupled and coupled sessions	166

Understanding the XRC volume synchronization process

Understanding how XRC handles record timestamping and volume synchronization will make your recovery procedures easier to establish and accomplish. This section contains the following information about consistency group time and timestamped writes:

Topic . . .	See page . . .
Overview of consistency group time	161
Terms used for consistency group time	162
Changes to consistency group time	163
Example: illustration of consistency group time	164
Understanding the importance of timestamped writes	165

Overview of consistency group time

Keeping active volumes consistent up to a given time is an important function of XRC support. Data integrity depends on ensuring that updates occur on the secondary volumes in the same order as they occurred on the primary volumes. This update sequencing is necessary in order to avoid data integrity problems and data loss. Maintaining data integrity becomes especially critical when multiple applications update a volume, or when a data set exists on multiple volumes that are spread across multiple storage controls.

The timestamp reported by the consistency group time is the primary system timestamp up to which the data on XRC active volumes is consistent. The following commands generate messages that report the consistency group time:

- XQUERY
- XDELPAIR
- XSUSPEND
- XRECOVER
- XEND
- XADVANCE

© Copyright IBM Corp. 1994, 2002

Terms used for consistency group time

The following defined terms are related to consistency group time and are useful to help understand consistency group time. The definitions for these terms refer to the following banking transaction example:

Example: A customer opens an account with a deposit of $400, and then later removes $300. On completion of both transactions, the customer's account shows a balance of $100.

Dependent Write

A write that should not occur until a previous write operation has completed.

Example: In the banking example, the withdrawal is dependent on the deposit since the withdrawal would not be possible if the deposit had not occurred. When XRC copies this set of transactions to a secondary set of volumes, XRC must maintain the transaction order since this is a dependent write operation.

Consistent

XRC guarantees that volumes are consistent when all dependent writes have been secured on the secondary volumes.

Example: In the banking example, XRC ensures that if the withdrawal transaction is secure on the secondary volume, then XRC has also secured the deposit transaction. There is no condition where XRC would secure the withdrawal without securing the deposit.

Consistency Group

A collection of application-updated records across all storage controls in an XRC session. XRC forms consistency groups such that records that are dependent on other records (dependent writes) are secured onto the secondary volumes in a consistent way.

Example: In the banking example, XRC ensures that when a consistency group is generated for the bank transaction, the withdrawal transaction is either in the same group as the deposit transaction or a later group. It will never be in an earlier consistency group. This is true whether XRC secures the data on the same volume, different volumes behind a common storage control, or different volumes behind different storage controls. *This requirement is critical for data base applications which have data spread across numerous volumes on numerous primary storage controls in the configuration*.

Consistency Group Time

The timestamp to which XRC ensures that XRC volumes are consistent. XRC derives the consistency group time from the system-provided timestamp for all application updates. When applying a consistency group to a set of secondary volumes, XRC reports the latest application timestamp secured on the volumes.

Example: In the banking example, assume that the deposit and the withdrawal transactions were the only transactions that took place. XRC reports the consistency group time of the withdrawal, as this transaction occurred after the deposit transaction.

Data Exposure

The approximate time difference (delta time) between data written to the primary volumes and data secured on the journal data set. The data that is exposed during this time is potentially nonrecoverable. This delta time is zero if all data that has been written to primary volumes for an XRC session has been secured on the journal data set.

Changes to consistency group time

If you periodically issue XQUERY commands to the system, the consistency group time increases as long as the primary volumes continue to have update activity.

If update activity ceases for all volumes in the XRC session, the consistency group time does not change for uncoupled sessions. For coupled sessions, the consistency group time increases. In both cases, an XQUERY report displays an IDLE time that indicates how long application updates have been idle.

The consistency group time reported with XQUERY, XDELPAIR, XSUSPEND, XRECOVER, XEND, and XADVANCE messages is dependent on the status of the XRC volume pair at the time the message is issued. For active XRC volume pairs, the consistency group time reported is the consistency group time for the entire XRC session. For suspended XRC volume pairs, the consistency group time is the time that the volume pair was last consistent with the XRC session.

XRC may suspend a volume pair because of an error, or because you have suspended the volume with an XSUSPEND VOLUME or XSUSPEND TIMEOUT command. XRC reports a consistency group time for the volume pair that reflects the last consistency group that XRC secured on the secondary volume.

Other volume pairs in the XRC session continue to have update activity, and their consistency group time continues to advance. At any given moment, the latest timestamp indicates the last consistency group that XRC has successfully secured on the secondary volumes.

An error can occur while XRC attempts to write records in a consistency group to a secondary volume. XRC may successfully write some of the records in this consistency group before the error occurs. In this case, the consistency group time for that volume pair correlates to the last consistency group that XRC successfully wrote to that secondary volume. If XRC has written all of the records in the consistency group to secondary volumes, then the reported consistency group time reflects the latest timestamp of all the records in the consistency group.

Notes:

1. XRC does not assign a consistency group time to a volume in the pending state. The volume has not achieved any consistent state until it is in duplex state. The volume pair is not consistent until XRC has copied all primary volume tracks with their updates to the secondary volume.

2. If none of the volume pairs have been updated since the session was started, XRC reports “NO_TIME_AVAILABLE” for the consistency group time of the volume pair, the session, or both.

Example: illustration of consistency group time

The example in Figure 9 further illustrates the concept of consistency group time. The frozen point-in-time in the example shows how the data mover processes successive groups of volume updates. The data mover reads data updates from the primary volumes, groups them into time intervals, and then writes them to the journal and secondary volumes. The word `Done` in a column means that the action has completed.

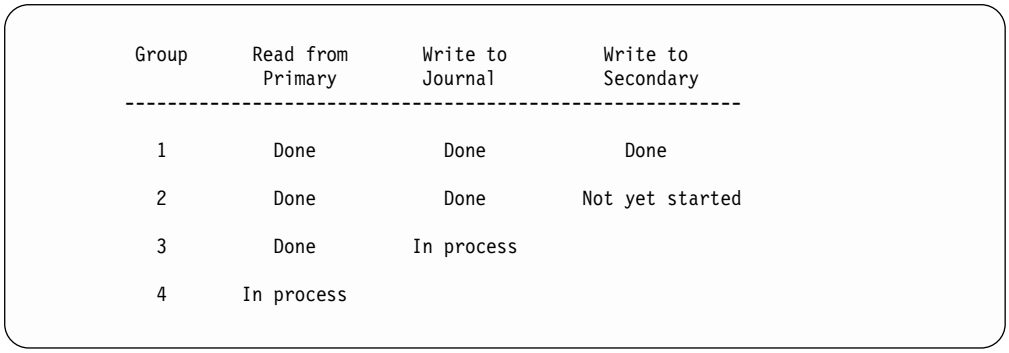

Group	Read from Primary	Write to Journal	Write to Secondary
1	Done	Done	Done
2	Done	Done	Not yet started
3	Done	In process	
4	In process		

Figure 9. A point-in-time look at the XRC volume synchronization process

The following are definitions for the column headings in Figure 9:

Group This collective information set represents all of the write updates that XRC has received from all of the primary volume’s storage controls during a given time interval.

Read from Primary A “Done” status indicates that the data mover has read the group into its storage.

Write to Journal A “Done” status indicates that the data mover has successfully written the group onto the journal data sets, which are accessible from the recovery site along with the control data set.

Write to Secondary A “Done” status indicates that the data mover has completed all update write operations on the secondary disk device for that group’s time interval.

Examples: These examples show the consistency group time that is reported for each of the following commands issued at the point in time that is shown in the example in Figure 9.:

- If you issue an XQUERY, XSUSPEND, or XEND command, XRC would report the session’s consistency group time as the Group 1 timestamp. Group 1 is the last group of updates that XRC successfully wrote to the secondary volumes.
- If you issue an XDELPAIR command to a volume pair that is in the session, XRC reports the consistency group time for the last group of updates that XRC wrote to the secondary volume. If the volume pair is still active in the session, the returned consistency group time is the same as the XRC session consistency group time.

- If you were to issue an XRECOVER command at this point in time, the secondary disk device would be consistent up to the Group 2 timestamp. That is the last completed group of interval data that is secured on the journal. Write updates from Groups 3 and 4 would be lost, and considered “bytes in transit”. In this case, the data exposure would be the time delta between Group 4 (data written to the primary) and Group 2 (data secured on the journal). Updates will only be applied to secondary volumes if the consistency group timestamp is less than or equal to the earliest interlocked session’s journal time.

Note: For a coupled session that was not ended or suspended by an XRC command, the starting *consistency_group_time* reported is the timestamp for the last known update for the session when the following conditions exist:

- A session status is noninterlocked
- Updates were not occurring when the session ended (the session was idle).

If the above conditions exist, the consistency time reported may be earlier than the master recovery time or the last session consistency time indicated in the output of an XQUERY command.

Understanding the importance of timestamped writes

You can maximize XRC data consistency when you ensure that all writes to primary volumes are timestamped. MVS device support code automatically timestamps all writes to XRC volumes.

Nontimestamped writes can exist for primary systems that are not at the required software system level. Nontimestamped writes can also exist for applications on non-MVS systems that bypass the device support code by performing direct I/O.

When XRC encounters nontimestamped write information, the following occurs:

- XRC issues an ANTX8030W message after encountering 5000 nontimestamped tracks on active volume pairs.
- The data returned with the XQUERY command shows a seqcheck condition in the status field, which warns you that the last recorded update is a nontimestamped write. The volume is no longer in seqcheck status when the last write of a subsequent consistency group is a timestamped write.

XRC copies nontimestamped writes to the secondary volume in the same way that it copies timestamped writes. A specific nontimestamped write that is dependent on other writes on other primary volumes may be applied to the secondary volume in an out-of-order sequence.

Example: If a bank account record is updated with a deposit of $400.00 and then updated with a withdrawal of $300.00, the account record would show an increase of $100.00. Assume that the processing system does *not* timestamp these data updates. The system applies the deposit update to the account record milliseconds *after* it applies the withdrawal update from a different volume. Under normal conditions, the volumes at both the primary site and the recovery site receive both updates. It is possible, though, that a disaster at the primary site could destroy the volume that contains the deposit information before the system could add this deposit data to the secondary volume. As a result, the deposit update would not be present on the recovery subsystem. The secondary volume record at the time of recovery would show a *decrease* of $300.00, with no indication that there is missing data.

Timestamped writes allow XRC to guarantee data consistency. To take advantage of the timestamping feature, ensure that all participating host systems operate with the latest device support code.

If the nontimestamped write has no dependencies on other writes, then it does not jeopardize data integrity. In either case, this condition is temporary and corrects itself during normal operations as XRC applies each set of updates to the secondary volume.

Recovering data with XRC—uncoupled and coupled sessions

This section contains the following information about recovering data with XRC:

"Recovering data with XRC"

"Creating a recovery volume report" on page 168

"Advancing consistency time using the XADVANCE command" on page 169

Recovering data with XRC

Before you begin: If the XRC data mover is running at the recovery site at the time of the failure, first issue an XEND command, then go to step 3.

Perform the following steps to have the recovery system take over for the primary system:

1. Reconfigure the path connections, if necessary, to connect local systems to the recovery system disk.

2. Start the DFSMS address space, if it is not already active on the recovery system. Run the IDCAMS recatalog function before activating the recovery system. The state, control, journal, and master (if applicable) data sets need to be cataloged on the recovery system.

   Note: If you are running multiple XRC sessions, the XRECOVER command must be issued for each individual session. If the sessions are coupled, the XRECOVER command coordinates the recovery of the sessions so that each session is recovered to the same consistency time. A master XRECOVER command does not exist.

3. Issue the XRC TSO command XRECOVER for each XRC session on the recovery system to add all valid, nonapplied journal data to the secondary (target) volumes. This action also sets the secondary volume serial numbers equal to the serial numbers of the primary volumes for those volumes which had previously reached a duplex state.

   Rule: The XRC recovery function must have access to the appropriate journal, control, and state data sets that were in use on the recovery system at the time of the failure. When the session is coupled, the recovery function must have access to the master data set.

   The XRECOVER command does not process volumes that are in the pending state, or that were suspended before reaching duplex state, as they have never reached duplex state. XRC recovery will not apply data to secondary volumes if the session was ended or suspended by command. The exception to this is when an error occurs during an update to a secondary volume while there is

data from the journal waiting to be written to the secondary volume. This allows the volume to be recovered to the same consistency time as the rest of the secondary volumes in the session.

Note: For a coupled session that was not ended or suspended by an XRC command, the starting *consistency_group_time* reported is the timestamp for the last known update for the session when the following conditions exist:

- A session status is noninterlocked
- Updates were not occurring when the session ended (the session was idle).

If the above conditions exist, the consistency time reported may be earlier than the master recovery time or the last session consistency time indicated in the output of an XQUERY command.

If XRC locates valid updates to suspended volumes, XRC applies those updates to the secondary volumes and increments their suspension times accordingly. If the volume suspension time is different than the recovery time of the XRC session, the data on the volume may not be consistent with data on other volumes.

The XRC recovery function uses the appropriate journal, control, and state data sets to put all secondary volumes in a known consistent state. The XRECOVER command, as part of XRC's recovery function, automatically creates an XQUERY volume report to assist you with recovery (see "Creating a recovery volume report" on page 168 for more information). XRC also initiates a request sequence to vary the secondary volumes offline, and then online to bring them to a ready state.

Note: In an actual disaster recovery, the primary volume would be offline when you issue the XRECOVER command. You must manually vary the primary volume offline during a disaster recovery verification test.

---

4. Some recovery system disk devices may still be in an error state after you have completed the preceding steps. If so, change the volume serial numbers to prevent the application from accessing the volume until recovery procedures have restored or have updated their contents.

   Note: It is not enough to vary these devices offline, as other applications can bring the devices online to another system.

   If the secondary has more cylinders than the primary, run ICKDSF Release 16 or above to the secondary volume to enlarge the VTOC, if required.

   You may occasionally find that you cannot access data on a successfully recovered volume because the volume's indexed VTOC has been disabled. In this case, run the BUILDIX function of ICKDSF to enable the indexed VTOC.

---

5. Restart primary systems and perform the same systems and application startup procedures that are performed on the primary system when applications start up following a system failure.

You may want to include catalog volumes as part of the data that you copy to the recovery system. Recovery system catalog entries will then be consistent with those on the primary system. Use the procedures that are listed in “Copying the catalog and control data sets” on page 47 to manage catalog updates that are not made to the recovery system.

Creating a recovery volume report

The XRECOVER command, as part of the XRC recovery function, automatically generates an XQUERY volume report to assist you with recovery.

Example: The following is an example of an XQUERY volume report that was issued by an XRECOVER command. See “XRECOVER–recovering data on the recovery system” on page 82 for further information on the XRECOVER command.

When the XRECOVER SESSION1 HLQ(SYS1) command processes, this report is generated:

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY RECOVER REPORT - 002
ANTQ8271I ------ORIGINAL-------
ANTQ8274I PRIMARY  SECONDARY STA CON CLP  ------TIMESTAMP---------
ANTQ8203I ---------------------------------------------------------
ANTQ8275I PRIM01   SECD01    DUP YES YES
ANTQ8275I PRIM02   SECD02    DUP YES YES
ANTQ8275I PRIM03   SECD03    SUS NO  YES  2000.322 19:54:00.562356
ANTQ8275I PRIM04   SECD04    PND NO  NO
ANTQ8237I TOTAL=4   DUP=2   PND=1   SUS=1
ANTQ8231I DATA CONSISTENT(2000.322 20:55:00.356210)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(SYS1)
ANTQ8201I XQUERY RECOVER REPORT COMPLETE FOR SESSION(SESSION1)
```

The XADVANCE command is also part of the XRC recovery function. The XADVANCE command differs from the XRECOVER command in that it applies updates as the XRECOVER command does, however, it does not *clip* the volumes.

Example: The following is an example of an XQUERY ADVANCE report. Issuing an XADVANCE command to a master session automatically generates an XQUERY MASTER report for the coupled sessions. See “XADVANCE—updating secondary volumes” on page 61 for further information on the XADVANCE command.

When the XADVANCE SESSION1 HLQ(SYS1) command processes, this report is generated:

```
ANTQ8200I XQUERY STARTED FOR SESSION(SESSION1) ASNAME(ANTAS001)
ANTQ8202I XQUERY ADVANCE REPORT - 001
ANTQ8272I PRIMARY  SECONDARY STA CON  ------TIMESTAMP---------
ANTQ8203I ---------------------------------------------------------
ANTQ8273I PAS001   SES001    SUS YES  2000.322 20:55:00.356210
ANTQ8273I PAS002   SES002    SUS YES  2000.322 20:55:00.356210
ANTQ8273I PAS002   SES002    SUS NO   2000.322 19:53:00.324610
ANTQ8237I TOTAL=3   DUP=0   PND=0   SUS=3
ANTQ8231I DATA CONSISTENT(2000.322 20:55:00.356210)
ANTQ8232I SESSIONTYPE(XRC)  ERRORLEVEL(SESSION)  HLQ(SYS1)
ANTQ8201I XQUERY ADVANCE REPORT COMPLETE FOR SESSION(SESSION1)
```

The XRECOVER and XADVANCE reports are similar to a regular XQUERY volume report. Volume pairs that were in duplex state at the time of recovery have had all journal data applied to their secondary volumes. The report shows active volumes

that are in pending state and suspended volumes that are in pending state. *z/OS MVS System Messages, Vol 1 (ABA-AOM)* contains additional details about the XRC messages that are listed within these reports.

Advancing consistency time using the XADVANCE command

The XADVANCE command can be used to recover secondary volumes to a consistent time. The difference between the XRECOVER command and the XADVANCE command is that the XADVANCE command does not change the volume serial labels on the secondary volumes. The XADVANCE command can be used if you wish to bring all volumes in the XRC session to a consistent state in order to make a copy of secondary volumes before restarting a session. However, for the XRECOVER command, XRC recovery will not apply data to secondary volumes if the session was ended or if the volume was suspended by command.

The exception to this is when an error occurs during an update to a secondary volume while there is data from the journal waiting to be written to the secondary volume. This allows the volume to be recovered to the same consistency time as the rest of the secondary volumes in the session.

Note: For a coupled session that was not ended or suspended by an XRC command, the starting *consistency_group_time* reported is the timestamp for the last known update for the session when the following conditions exist:

- A session status is noninterlocked
- Updates were not occurring when the session ended (the session was idle).

If the above conditions exist, the consistency time reported may be earlier than the master recovery time or the last session consistency time indicated in the output of an XQUERY command.

Table 40 outlines how updates occur when you issue the XADVANCE command for active and inactive sessions.

Table 40. Using the XADVANCE command

If the session is . . .	Then . . .
Active	All volumes must be suspended by error or by command. If volumes are suspended because of an error, issue the XADVANCE command to bring those volumes to a consistent state. The command performs the following updates depending on whether the session is coupled or uncoupled: • If the session is not coupled, this command applies all updates that have been written to the journal and have not been written to the secondary volumes. • If the session is coupled, this command applies all updates that have been written to the journal and have not been written to the secondary volumes that have timestamps less than or equal to the master session's master recoverable time.

Table 40. Using the XADVANCE command (continued)

If the session is . . .	Then . . .
Inactive	The session has been suspended or has ended. If volumes are suspended because of an error, issue the XADVANCE command to bring those volumes to a consistent state. The command performs the following updates depending on whether the session is coupled or uncoupled: • If the session is not coupled, this command applies all updates that have been written to the journal and have not been written to the secondary volumes. • If the session is coupled, this command applies all updates that have been written to the journal and have not been written to the secondary volumes that have timestamps less than or equal to the master session's master recoverable time.

Chapter 9. Migrating Data with XRC

The purpose of an XRC migration session is to migrate volumes from one location to another. The migration can be between local devices or between remote sites. The XRC migration session ensures that data for all volumes is consistent when the migration is complete. It is not important to be able to recover the data at all times. A typical migration is of relatively short duration.

Note: The XCOUPLE ADD command fails if you specify SESSIONTYPE(MIGRATE) on the XSTART command. If you restart a coupled session and you specify SESSIONTYPE(MIGRATE), the MIGRATE option will be ignored.

XRC may be the best choice for moving your data to other disk devices, depending on the device types of the primary and secondary volumes.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Steps for migrating data with XRC	171
Migration operation considerations	172

Steps for migrating data with XRC

Before you begin: You need to be familiar with the commands used in the following steps. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

The steps for data migration with XRC are as follows:

1. Initiate a migration session by specifying SESSIONTYPE(MIGRATE) on the XSTART command. The XSTART command directs the data mover to start a migration session. (XRC uses the data mover for disaster recovery purposes and for migration.)

   In migration mode, data consistency across all volumes is only important when migration is complete. Therefore, it is best to specify ERRORLEVEL(VOLUME) on the XSTART command for a migration session. If an error occurs, XRC suspends the volume pair that incurred the error. After you correct the error, issue an XADDPAIR SUSPENDED command to add back the volume pair into the XRC session. Meanwhile, XRC continues to copy other volume pairs.

   To start a session called, for example, MIGRATE, you must already have allocated the state data set with the session ID of MIGRATE. See “Specifying XRC journal, control, and state data sets” on page 42 for more information.

   Example: The following example starts a session named MIGRATE:

   ```
   XSTART MIGRATE SESSIONTYPE(MIGRATE) ERRORLEVEL(VOLUME)
   ```

2. Establish an XRC pair with an XADDPAIR command. Use the default ERRORLEVEL(VOLUME) option, and do *not* specify NOCOPY. XRC issues an ANTI8024I message when the volume or volumes are in the duplex state. This

© Copyright IBM Corp. 1994, 2002

message means that the initial copy is complete and that XRC has updated the secondary volumes with all of the changes that have occurred during the copy operation.

---

3. Monitor the initial copy process by issuing XQUERY TSO commands. The XQUERY command returns the percentage of the initial copy that has completed. During the initial copy, XRC also copies all primary volume updates to the secondary (target) volume.

---

4. Verify that the XRC volumes are in duplex state with an XQUERY command. When the volume pair attains duplex state, XRC has copied all of the primary volume contents and all subsequent updates to the XRC secondary volume. XRC continues to update the secondary volume as changes occur to the primary.

---

5. Switch to the secondary (target) volume when the copy operation is complete. To switch your application from the primary volume to the secondary volume, do the following:
   a. Stop all applications that are using the primary volume.
   b. To complete the migration at a known point in time, issue the XEND command with the ATTIME keyword or the DRAIN keyword. Specify ATTIME to ensure that XRC does not copy any updates to secondary volumes beyond the specified timestamp. Specify DRAIN when you have quiesced primary application activity and want to ensure that all application I/O has been applied to the secondary volumes before ending the session. Issue the XEND(DRAIN) command to secure all updates onto the XRC secondary volume. When all application updates have stopped, no further updates occur on the primary volumes.
   c. Issue an XRECOVER command to change the secondary volume serial number to be the same as the primary volume serial number.
   d. Vary the original primary device offline. If the secondary volume contains more cylinders than the primary volume, you must run ICKDSF Release 16 (or above) with the REFORMAT command against the secondary volume. This action expands the VTOC to reflect the additional space. If you do not run ICKDSF with REFORMAT, the system cannot recognize or use the extra space on the secondary volume.
   e. Restart all applications, which will now use the new disk volumes.

Migration operation considerations

The following considerations can affect migration mode operations:

- A migration session does not require that you allocate control data sets or journal data sets. However, it does require the state data set. Because there are no control or journal data sets, all data is consistent on the secondary volumes up to the timestamp reported when the session is ended.

  Note: Migration mode uses fewer system resources than disaster recovery mode because XRC does not write data to the journal data sets or control data sets. The primary application update rate supported in migration mode is about the same as in disaster recover mode.

- To complete the migration faster, you may want to increase the number of concurrent synchronization tasks. To do this, use the XSET command to:
  - Change the SYNCH value to increase the total number of concurrent tasks that perform the migration copy.
  - Change the SCSYNCH value to increase the maximum number of tasks that can be active concurrently on a single storage control.

  You may need to allocate additional fixed storage before starting additional tasks. See “XSET–changing session parameters” on page 84 for the XSET command syntax.
- You can change the session type whenever you issue the XSTART command. Initially, you may start a migration session to copy all volumes from one location to another. When the copy has completed, you may then choose to suspend the XRC session and restart the session as a disaster recovery session. Remember that you must define journal data sets and control data sets for a disaster recovery session.

  Note: The SESSIONTYPE of MIGRATE cannot be used for coupled XRC sessions.

For an example of a migration scenario, see “Example: migrating data” on page 127.

Chapter 10. Recovering from Error Conditions Using XRC

This chapter describes various error conditions and how to recover from them.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Specifying the ERRORLEVEL parameter in coupled and uncoupled sessions	175
Steps for recovering from environmental errors	177
Steps for recovering from processor, address space, and XRC data set failures	178

Specifying the ERRORLEVEL parameter in coupled and uncoupled sessions

The actual XRC error recovery process depends on both of the following conditions:

- The specific failure condition **and**
- The error level that is specified for that volume pair.

You can specify the default error level for the entire session when you start the XRC session by issuing the XSTART command and specifying the ERRORLEVEL parameter. You can also change the error level for specific volume pairs when you issue the XADDPAIR command with the ERRORLEVEL parameter specified. The ERRORLEVEL parameter aids in the error recovery process. It does not cause session errors.

To specify how XRC processes certain error conditions, use one of the following ERRORLEVEL subparameters:

- VOLUME
- SESSION
- *group_name*

See “XSTART–starting a session” on page 89 and “XADDPAIR–adding volume pairs or utility volumes” on page 56 for more information on these commands and their parameters.

Comparing the ERRORLEVEL options

Table 41 on page 176 compares the different options of the XADDPAIR ERRORLEVEL parameter and describes how different system failures affect the suspension of XRC volume pairs. “Active volumes” refers to volumes that are not in suspended state.

© Copyright IBM Corp. 1994, 2002

Table 41. XRC system failure volume pair suspensions, determined by ERRORLEVEL options.

ERRORLEVEL option	Single volume failure	Primary storage control failure	Data mover failure
(Volume)	One active volume pair is suspended.	All active volumes associated with the storage control session are suspended.	All active volume pairs in the session are suspended.
(Session)	All active volume pairs in the storage control session are suspended. If the session is coupled, all volumes in all other coupled sessions associated with the same master session are also suspended.	All active volumes in the storage control session are suspended. If the session is coupled, all volumes in all other coupled sessions associated with the same master session are also suspended.	All active volume pairs in the XRC session are suspended. If the session is coupled, all volumes in all other coupled sessions associated with the same master session are also suspended.
(*group_name***)**	All active volume pairs in the group are suspended.	All active volumes on the storage control are suspended, and all active volumes in this group on any storage control are suspended. See Note.	All active volume pairs in the XRC session are suspended.

Note: A primary storage control error suspends all XRC volumes that are attached to it. If any volumes on that storage control belong to an XRC named group, those volume groups become suspended on other storage controls.

ERRORLEVEL(VOLUME)

The VOLUME option suspends an XRC volume pair that encounters a device error. XRC continues to process all other volume pairs in the XRC session. You must correct the volume's error condition and then resynchronize it so that it is consistent with the other volumes in the XRC session.

If you select VOLUME and the error is associated with the XRC session, then XRC suspends all volume pairs for that storage control. Volume processing continues for other storage controls.

Note: If a volume that is suspended due to an error contains data that is dependent on data on another volume, data consistency cannot be assured in the event of a disaster recovery.

ERRORLEVEL(SESSION)

The SESSION option suspends all volume pairs in the XRC session when an error is associated with any duplex volume that has ERRORLEVEL(SESSION) specified. The data mover then continues to read the updates from the primary storage controls. It does not copy data to secondary volumes until you have added the volume pairs back to the session. You need to copy only the changed data when you add the volumes. You must correct the error condition and resynchronize the suspended volume pairs. If SESSION is selected and a failure occurs on a single storage control that attaches to any duplex volume pair with ERRORLEVEL(SESSION) specified, then all XRC volume pairs are put into

suspended state. XRC suspends volume pairs that are in pending state, and the pairs remain where they are in synchronization. Software bitmapping continues for all volumes that are part of the XRC session. In addition, if the affected session is a coupled session, all volumes in all other coupled sessions that are associated with the same master session will also be suspended. The master session is placed into HOLD status to protect the recoverable time for each XRC session.

ERRORLEVEL(group_name)

The *group_name* option allows you to associate any valid TSO name with one or more volume pairs. An error that occurs on a single storage control that attaches to any duplex volume pair that has ERRORLEVEL(*group_name*) specified causes XRC to suspend all other volume pairs that have been assigned the same group name specification. If the same *group_name* is used on other sessions that are coupled in the same master session, the volume pairs in the other sessions are not suspended.

Note: Volumes must be in duplex state before an error that is related to the associated volumes can cause other duplexed volumes to become suspended. This prevents the following error conditions for volume pairs that are still in pending state from suspending all volume pairs in the session:
- Volume serial numbers entered incorrectly
- Track copy errors during volume synchronization
- Spurious track errors during volume synchronization

You can diagnose and correct these XADDPAIR command-related errors independently of other XRC session activity. When you have corrected the error, you can make the volume pair operational by reissuing the XADDPAIR command.

Steps for recovering from environmental errors

Take action that is based on the scope and the location of the error. Automation functions can take action based on the presence of ANTX5xxxE–type errors, which report the error, the cause of the error, and the volumes that are affected.

Before you begin: You need to be familiar with the commands used in these steps. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

Use the following procedure to recover from environmental errors:

1. Correct the error.

   _____

2. Restart the primary system, if necessary.

   _____

3. Issue an XADDPAIR SUSPENDED command to add the suspended volume pairs back to the XRC session. XRC automatically copies all pending updates when the application programs become active.

   _____

Steps for recovering from processor, address space, and XRC data set failures

XRC error conditions may occur as a result of various failures on an application processor. These errors can cancel critical address spaces that reside within the affected system, and may in turn cause the application to end. The data mover may also be on the system with the failure.

Other error situations may affect only the data mover system. The data mover may be affected by an environmental problem, an internal data mover error, errors on the XRC state, control, or journal data sets, or due to the operator canceling the ANTAS*nnn* address space.

Before you begin: You need to be familiar with the commands used in these steps. For more information, refer to “Chapter 5. Administering Your XRC Environment” on page 55.

Perform the following steps to recover from processor, address space, and XRC data set failures:

1. Correct the error.

   _____

2. Restart the primary system, if necessary.

   _____

3. Restart the data mover system, if necessary. For many errors, the XRC session will restart itself.

   _____

4. Issue an XSTART command to start the XRC session, if necessary.

   _____

5. Issue an XADDPAIR SUSPENDED command to add the suspended volume pairs back to the XRC session. XRC automatically copies all pending updates when the application programs become active.

   _____

6. If the XRC session is coupled, issue the XCOUPLE *msessionid* RELEASE command.

   _____

Part 3. Peer-to-Peer Remote Copy

© Copyright IBM Corp. 1994, 2002

Chapter 11. Planning for PPRC

Before installing and using PPRC you will need to plan for it. Understanding the software and hardware requirements will help you prepare to use PPRC in your installation.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
PPRC requirements	181
PPRC operational considerations	182
Establishing the PPRC configuration	184
Determining PPRC resource needs	185
Copying the catalog and control data sets	188
Controlling access to PPRC commands	188

PPRC requirements

You must meet the software and hardware requirements below in order to plan for and successfully install PPRC. Table 5 on page 12 includes additional information about PPRC hardware and software requirements.

PPRC software requirements

MVS, VM, VSE, and ICKDSF stand-alone systems support PPRC.

MVS software requirements

PPRC TSO command functions require MVS/DFP™ 3.3.0 and above software support.

VM software requirements

VM/ESA® 2.1.0 (and above) supports PPRC operations. See the items identified as “For VM Systems” under “PPRC operational considerations” on page 182.

ICKDSF support

ICKDSF release 16 and above support PPRC operations for MVS, VM, VSE, and stand-alone platforms. The ICKDSF PPRCOPY command provides functions that correspond to the PPRC MVS TSO command set. Refer to *Device Support Facilities User’s Guide and Reference* for specific information.

PPRC hardware requirements

The following are PPRC hardware requirements:

- Both primary and recovery site storage controls must be of the same type (peer to peer).

  Example: If the primary site storage control is an IBM Enterprise Storage Server (ESS), the recovery site storage control must also be an ESS.
- Physical ESCON connections must exist between PPRC storage subsystems. These connections can be direct fiber optic connections or through ESCON Directors.

© Copyright IBM Corp. 1994, 2002

- PPRC-capable Licensed Internal Code (LIC) must be installed on both the primary and secondary storage subsystems.
- A compatible secondary volume must be available for each primary volume you wish to copy. The secondary volume must have the identical track capacity and number of tracks per cylinder and either the same or larger volume capacity.

PPRC supported devices

PPRC supports all **primary disk subsystems** and **secondary disk subsystems** that have PPRC-capable LIC. The IBM 3990 Model 6 Storage Control supports PPRC in either Basic or Enhanced mode.

PPRC operational considerations

This section addresses the physical and operational considerations for PPRC. Topics include limits on the number of active volumes and sessions, conditions for establishing PPRC paths, and volume format restrictions.

The following considerations apply to PPRC operations:

- Primary and secondary volumes may either be a member of a PPRC session or a member of a dual copy pair, but not both. See “Combining copy services operations” on page 16 for the allowable copy combinations for a device.
- To ensure data integrity, all volumes that are part of PPRC copy operations must conform to the following **volume format**, **track**, and **access method** restrictions:
  - Volumes must have a standard format for record zero (R0). Volumes with R0 data lengths longer than 8 bytes can cause a track format error to remain undetected when the storage control formats the track in cache.
  - You cannot assign alternate tracks in the user area. User data can be overlaid if you assign a user track on the primary address as an alternate track for a secondary address.
  - All storage control Define Extent commands must specify normal access authorization mode. Data written to a PPRC primary device while the storage control is in diagnostic or device support mode is not copied to the PPRC secondary volume. It is therefore important to remove volumes from the session before running a utility program like ICKDSF.
- You can only copy a primary volume to one secondary volume.
- A PPRC volume can be only primary or secondary; you cannot assign a volume to both functions at the same time.
- You must establish separate paths to transfer data in the opposite direction between subsystems. This is especially important for P/DAS volume swapping operations as described in “Chapter 18. PPRC Dynamic Address Switching (P/DAS)” on page 251.

  Example: ESCON paths established by PPRC provide connections to primary and secondary volumes, and for host data transfer. Specifically, a single path can be established *in one direction* to transfer a volume or volumes from, for example, storage subsystem A to storage subsystem B. The same logical and physical path established from A to B cannot also be used to transfer data from B to A. You must establish separate paths to transfer data from subsystem B to subsystem A (with B being the primary subsystem for those volumes). The host system, however, can transfer data on any path.
- PPRC secondary volumes can only accept a subset of I/O operations while they are secondary devices (much like dual copy volumes). IBM recommends that you

vary the attached secondary volumes offline to all attached systems. The PPRC secondary volume serial numbers are the same as those on the primary volumes.

- **For VM Systems:** Native VM supports PPRC with VM/ESA 2.1.0 and above. Observe the following when you operate PPRC as a guest under VM/ESA 2.1.0 and above:
  - PPRC volumes must either be dedicated volumes, or you must define them as fullpack minidisks. This includes DEVNO-defined minidisks.
  - The VM guest directory must include an entry that states “STDEVOPT DATAMOVER=YES”.

  For VM/ESA levels prior to 2.1.0, you must define PPRC volumes to VM as unsupported disk, and you cannot define them as minidisks.

Examining PPRC configuration options

Peer-to-peer remote copy requires ESCON paths between primary and recovery storage controls. You may define from one to eight ESCON paths, depending on the capability of the storage subsystem, from a single primary site storage control to a specific recovery site storage control. You may also attach up to four recovery site storage controls to each primary site storage control. Therefore, a total of 32 paths can be defined between a single primary site storage control and four recovery site storage controls. Figure 10 shows examples of possible ESCON path configurations between PPRC primary and recovery systems.

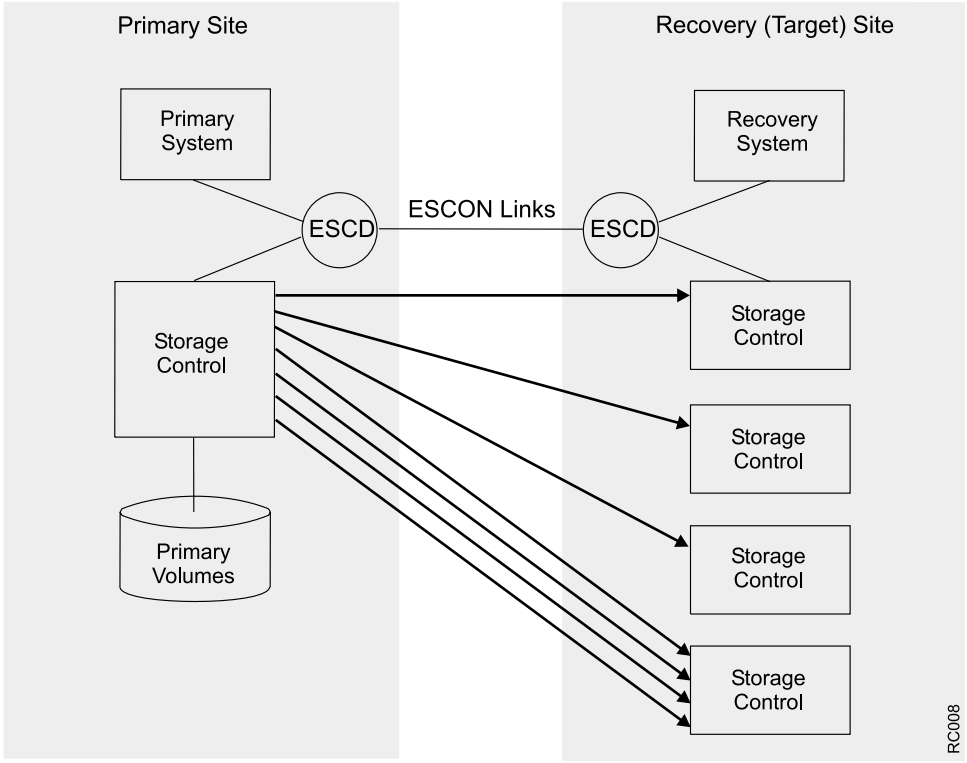

Figure 10. PPRC ESCON path options

A single storage subsystem at the recovery site can be linked to as many primary site storage controls as there are logical channels and devices available. A large cache and NVS help the recovery site storage control to accept the copy workloads from multiple primary site storage subsystems.

You can use existing ESCON paths, or add other dedicated paths. ESCON paths may include an ESCON Director (ESCD) or Directors, or 9036 ESCON Remote Channel Extenders to provide configuration flexibility.

Considering the PPRC solution

Because of the haphazard nature of disasters, and the probability that during a disaster all of your data processing equipment will not fail at the same instant, IBM offers the following guidelines if you are considering PPRC as your disaster recovery solution:

- Choose PPRC and implement these system preparations:
  - Establish the links necessary to allow you access to primary system IEA49*xx* ERP messages from the recovery site. The IEA49*xx* message indicates that a PPRC volume pair has been suspended at the time indicated. Updates are written to the suspended primary volume after this suspension time but are not copied to the suspended secondary volume.

    Having a procedure in place to relay ERP messages to the recovery site is not, by itself, a 100% disaster recovery solution. In order to ensure data integrity for all PPRC secondary volumes, you must specify CRIT(YES) *on all CESTPAIR commands*. See “Preparing for PPRC error recovery” on page 244 for additional information.
  - Specify CRIT(YES) on CESTPAIR commands. The CRIT(YES) option ensures that writes are inhibited when the PPRC pair is suspended as a result of a failure between the primary and secondary storage controls. The PPRC copy operation remains suspended until the problem is corrected and either a CESTPAIR RESYNC or CDELPAIR command is issued.

    If the default setting is changed to CRIT(YES-ALL), then specifying CRIT(YES) on CESTPAIR commands means that future write operations are stopped as a result of any failure to update a secondary device. Failures here include primary and secondary device errors. Path errors and secondary storage control failures cause PPRC to suspend the pair and stop all writes to the primary volume.

    Notes:

    1. The CRIT(YES) parameter has two modes of operation: CRIT(YES-PATHS) and CRIT(YES-ALL). The default setting, CRIT(YES-PATHS), inhibits all write operations only when a failure causes all paths to the secondary volume to be lost.
    2. Depending on the type of storage subsystem, you can change the operating mode either with a PPRC command, or from the storage control’s maintenance panel as described in “Appendix B. Configuring the Data Mover and PPRC Options on IBM 3990 and 9390 Storage Controls” on page 305. See your IBM service representative if you need to change the maintenance panel settings.
- Consider XRC instead of PPRC. The asynchronous solution, XRC, employs the MVS™ system data mover, which manages the placement of data on the secondary volumes as well as recovery of those volumes.

Establishing the PPRC configuration

The configuration of the PPRC system is vital to its recovery and performance capabilities. The following section describes basic PPRC configurations and those that include DASD channel extenders.

Basic PPRC configurations

Plan your configuration to provide for capacity requirements and redundancy requirements. Peer-to-peer remote copy allows from one to eight ESCON paths (IBM recommends at least two paths) from any primary site storage control to a single recovery site storage control. A specific primary site storage control can connect with up to four different recovery site storage controls. However, you can link a single storage subsystem at the recovery site to a maximum of 64 primary site storage controls. A large cache and NVS will help the recovery site storage control to accept the copy work loads from multiple primary site storage subsystems.

You can establish PPRC paths between the storage controls after you have installed the necessary Licensed Internal Code (LIC) in each storage control. You can either install dedicated ESCON channels between the storage controls or use existing ESCON channels connected by common ESCON Directors. When you are establishing paths, distribute them between different storage clusters and Directors to gain flexibility and availability.

Depending on the capability of your storage subsystems, you can establish up to eight paths with each CESTPATH TSO command. You need the following information before you can issue a CESTPATH command:

- Primary device number
- Primary device's storage subsystem ID (SSID)
- Primary device's storage control serial number
- Recovery storage subsystem ID (SSID)
- Recovery storage control serial number

The CESTPATH command, with its syntax and parameters, is described in "CESTPATH – establishing paths" on page 205. For additional information and examples on establishing PPRC paths, see "Establishing paths" on page 218.

PPRC and DASD channel extenders

Adding channel extenders to a PPRC configuration to support longer distances can seriously degrade the application performance. This is because of the additional time that is added to each copy write I/O. IBM therefore does not generally recommend using channel extenders with PPRC for any applications that are sensitive to response time.

See "Using channel extenders" on page 15 for general information about channel extenders.

Determining PPRC resource needs

This section contains information and recommendations about the resources needed to establish a PPRC configuration that meets disaster recovery and workload migration needs. Evaluate the resource issues listed in Table 42 before installing PPRC.

Table 42. PPRC resource considerations

Consider . . .	On page . . .
Steps for evaluating how PPRC affects the primary storage subsystem	186
Evaluating how PPRC affects the recovery storage subsystem	186
Determining the recovery site storage capacity	186

Table 42. PPRC resource considerations (continued)

Determining recovery system access to resources	187
Establishing storage subsystem channel connections	187
Establishing PPRC ESCON connections	187

Steps for evaluating how PPRC affects the primary storage subsystem

Storage subsystem use can affect overall system performance, and it is therefore necessary to determine how adding PPRC might affect the primary storage subsystem.

Perform the following steps when evaluating how PPRC affects the primary storage subsystem:

1. Determine the write rate, for each primary site storage control, for the volumes that you plan to copy.

   _____

2. Add this rate, as if it was all DASD fast write, to the total storage subsystem use.

   _____

3. Decide if this impact on the storage subsystem allows applications to achieve acceptable disk I/O performance. If it does not, you may need to redistribute volumes among storage controls or add an additional primary site storage control.

   _____

Also consider how initial volume copy and resynchronization activity affect storage control use and application performance. You may notice this when you start a copy operation, restart a system, perform some copy recovery events, or when you add volumes to the PPRC system.

Evaluating how PPRC affects the recovery storage subsystem

PPRC lets you allocate recovery system volumes on any number of recovery storage subsystems. The storage controls at the recovery site do not have to have a one-to-one correlation with their primary storage subsystem counterparts. Project the performance impact on recovery storage subsystems to be sure that it is within acceptable limits.

Higher recovery storage subsystem activity increases the write time between storage controls, thereby affecting primary disk I/O performance. Include all other recovery storage subsystem activity in your assessment, such as access by other applications and systems to non-PPRC volumes on the same storage subsystem.

Determining the recovery site storage capacity

As part of planning for PPRC installation, you must consider the number and type of disk volumes. Use Table 43 on page 187 to determine your necessary disk capacity.

Table 43. Steps for disk volume determination

Stage . . .	Description . . .
Determine the number of necessary recovery storage subsystem disk volumes.	Assume that all volumes from the primary subsystem must be copied to the recovery system. You can then consider each volume individually to determine if you need it for recovery and subsequent application operations on the recovery system. **Example:** You do not have to copy disk volumes that contain only page data sets to the recovery site storage subsystem.
Determine the number of disk volumes to be copied to the recovery site storage subsystem.	Consider which application and support data sets will be needed to start applications on the recovery subsystem at recovery time. **Example:** User catalogs are a good example of such data sets. If the recovery subsystem cannot find data volumes associated with user catalogs, jobs will not start successfully after recovery.
Consider disk volumes	Consider what disk volumes, such as DFSMShsm migration level 1 data sets, may be needed for recovery operations.

Determining recovery system access to resources

Rules: To meet disaster recovery needs, PPRC recovery systems must:

- Be able to run the applications that use the copied data.
- Be able to process commands to the secondary disk to end the secondary copy state and to become available for application system allocation.
- Have access to the TSO command module.

Establishing storage subsystem channel connections

Disk channel connections include those to primary and recovery systems. PPRC connections between storage controls must be ESCON connections.

Tip: It may be useful when planning an initial PPRC test installation (with a limited number of volumes) to have both the primary and secondary volumes on the same storage subsystem, thereby conserving resources.

Establishing PPRC ESCON connections

ESS storage subsystems support PPRC at distances up to 103 km. You can configure other IBM storage controls to support ESCON ports (connected by multimode fiber optic cables) at distances of up to three kilometers. ESCON Directors, if your configuration includes them, can have Extended Distance Feature (XDF) ports with singlemode fiber-optic cables for longer link lengths.

You can use 9036 ESCON Remote Channel Extenders to convert from multimode fiber links to XDF and back, and as XDF extender links. Use these or other comparable repeaters when the distance between systems is greater than 3 km and your configuration does not include ESCON Directors.

You can only dynamically switch one ESCON Director within a path to an I/O unit. The other ESCON Director, if present, provides a static connection. Refer to the appropriate storage subsystem planning guide for information about ESCON configurations with ESCON Directors.

Recovery site storage subsystems must have access to the host systems that will run primary applications when the primary system is out of service. A host system at the recovery site must be able to issue recovery commands as part of recovery system takeover procedures.

Copying the catalog and control data sets

You can reduce disaster recovery time if you use remote copy for the volumes that contain the master catalog, key user catalogs, and system control data sets. This may be a better choice than maintaining a current copy of the catalog and various control data sets on the recovery system.

Remote copy automatically copies all data sets on the managed volumes. These include catalog data sets and control data sets.

Controlling access to PPRC commands

The following are two ways to protect access to PPRC commands. You can:

- Put the commands in a RACF-protected library. (See “Steps for placing TSO commands in a RACF-protected library” on page 190 for further instruction.) **or**
- Define resource profiles in the RACF FACILITY class and restrict access to those profiles. (See “Defining resource profiles in the RACF FACILITY class” on page 191 for further instruction.)

Chapter 12. Setting Up the PPRC Environment

This chapter describes how to set up the peer-to-peer remote copy (PPRC) environment.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Installing peer-to-peer remote copy	189
Controlling access to PPRC resources	190
Identifying volume pairs	192
Steps for including PPRC TSO commands in automated procedures	193
Configuring storage subsystem resources for PPRC	194

Installing peer-to-peer remote copy

To install peer-to-peer remote copy, perform the following steps:

1. Install the appropriate z/OS release, and the maintenance that is specified under Technical Updates at the copy services website at: **http://www.storage.ibm.com/software/sms/sdm/sdmhome.htm**

2. Authorize PPRC TSO commands by adding the command names to the AUTHCMD PARM parameter of the IKJTSOxx member of SYS1.PARMLIB. After you have added the PPRC command names to the IKJTSOxx member, issue the TSO command PARMLIB UPDATE(xx) to activate the new IKJTSOxx member.

   To restrict PPRC commands to certain users, follow the instructions in “Controlling access to PPRC resources” on page 190.

3. Based on the applications that you plan to copy, determine which devices you want to establish as primary and secondary devices.

4. Based on the primary and secondary devices, determine the number of paths that are required to establish connections between each primary site and recovery site storage control. Use the DEVSERV system console command to determine the channel connection addresses for the devices. The serial number is on the lower left-hand part of the operator’s panel on an IBM storage control.

5. Issue the appropriate CESTPATH TSO commands to establish the PPRC paths.

6. Issue the appropriate CESTPAIR TSO commands to define the PPRC pairs.

7. Issue the CQUERY TSO command to verify the PPRC installation.

© Copyright IBM Corp. 1994, 2002

Controlling access to PPRC resources

Each installation is responsible for preventing the unauthorized use of PPRC TSO commands and for controlling ICKDSF activity to PPRC-managed volumes. This section describes how to protect PPRC commands and control volume repair activities.

Protecting access to PPRC commands

If possible, restrict PPRC TSO commands to authorized storage administrators only. You can place the commands in a library that is protected by IBM Resource Access Control Facility (RACF), and define access to resource profiles.

Steps for placing TSO commands in a RACF-protected library

Place the TSO commands in a RACF-protected library to restrict PPRC TSO commands to authorized storage administrators.

Before you begin: You need to be familiar with the commands used in these steps. For more information, refer to “Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199.

Perform the following steps to RACF-protect PPRC commands:

1. Issue the following RDEFINE command for each PPRC command, and for each command abbreviation that you want defined to RACF:

   ```
   RDEFINE PROGRAM cmdname ADDMEM('SYS1.CMDLIB')/volser/NOPADCHK) UACC(NONE)
   ```

   The following terms apply to the above example:

   cmdname Defines the PPRC TSO command name or an abbreviation of a command. Issue a separate RDEFINE command for each command and any command abbreviations you plan to use. RACF can only perform checking on commands and abbreviations that are defined to it.

   volser Defines the name of the volume that contains the SYS1.CMDLIB data set.

---

2. Issue the PERMIT command for all commands and authorized PPRC TSO command users as follows:

   ```
   PERMIT cmdname CLASS(PROGRAM) ID(name) ACCESS(READ)
   ```

   The following terms apply to the above example:

   cmdname Defines the PPRC TSO command name, or an abbreviation of a command.

   name Defines the user ID that will receive RACF access authority for that command name.

---

3. Issue the SETROPTS command from a user ID that has the appropriate authority:

   ```
   SETROPTS CLASSACT(PROGRAM) WHEN(PROGRAM) REFRESH
   ```

Defining resource profiles in the RACF FACILITY class

You can limit the use of PPRC commands by defining resource profiles in the RACF FACILITY class and restricting access to those profiles. To use a protected command, you need read-access authority to the applicable profile.

Table 44 lists the PPRC commands and the FACILITY class profiles that can restrict them. See the *z/OS Security Server RACF Security Administrator's Guide* for details on activating the RACF FACILITY class, and defining and authorizing users to the PPRC command profiles.

Table 44. PPRC FACILITY class profile names

Command	Profile Name
CDELPAIR	STGADMIN.ANT.PPRC.COMMANDS
CDELPATH	
CESTPAIR	
CESTPATH	
CGROUP	
CQUERY	
CRECOVER	
CSUSPEND	
CQUERY	STGADMIN.ANT.PPRC.CQUERY

Note: Authorize CQUERY command use with the STGADMIN.ANT.PPRC.COMMANDS profile or the STGADMIN.ANT.PPRC.CQUERY profile. PPRC first checks STGADMIN.ANT.PPRC.COMMANDS for authorization. If authorization is not permitted with the STGADMIN.ANT.PPRC.COMMANDS profile, PPRC checks the STGADMIN.ANT.PPRC.CQUERY profile for authorization to issue the CQUERY command.

Examples: The following examples activate the RACF FACILITY class, define the profile for the PPRC commands, and give user STGADMIN authority to use this profile:

- This example activates the RACF FACILITY class:

```
SETROPTS CLASSACT(FACILITY)
```

- This example defines the profile for PPRC commands, and authorizes user STGADMIN to use this profile:

```
RDEFINE FACILITY STGADMIN.ANT.PPRC.COMMANDS UACC(NONE)

PERMIT STGADMIN.ANT.PPRC.COMMANDS CLASS(FACILITY) -
       ID(STGADMIN) ACCESS(READ)
```

Controlling ICKDSF activity to PPRC volumes

The storage administrator must control ICKDSF activity to primary and secondary PPRC-managed volumes. PPRC cannot copy some ICKDSF operations from the primary to the secondary volumes. If you plan to use ICKDSF to repair either a primary or secondary volume, first remove the volume pair from the session. After the ICKDSF function is complete, reestablish the session with a CESTPAIR command.

Steps to perform an ICKDSF repair on a PPRC-managed primary volume:

Use the following steps to perform an ICKDSF repair on a PPRC-managed *primary* volume:

1. Suspend the volume pair with a CSUSPEND PRIMARY command that is directed to the primary volume. This assumes that the volume pairs are not already in a suspended state as a result of the error that necessitated the ICKDSF repair action. See “Putting PPRC error recovery procedures into effect” on page 244 for information about specific error conditions.

   _____

2. Run the latest level of ICKDSF to correct the error once the pair is in the suspended state.

   _____

3. Remove the volumes from the suspended state by using the CESTPAIR command with the RESYNC option to resume operations.

   _____

Steps to perform an ICKDSF repair on a PPRC-managed secondary volume:

Use the following steps to perform an ICKDSF repair on a PPRC-managed *secondary* volume:

1. Issue a CSUSPEND command to suspend the volume pair, thereby allowing change recording to continue. This assumes that the volume pairs are not already in a suspended state as a result of the error that necessitated the ICKDSF repair action. See “Putting PPRC error recovery procedures into effect” on page 244 for information about specific error conditions.

   _____

2. Issue the CRECOVER command to the secondary volume so that you can make repairs. (PPRC puts the volume into simplex state to permit I/O access.)

   _____

3. Run ICKDSF release 16 or above to repair the volume.

   _____

4. If data on the secondary volume is unchanged, issue the CESTPAIR command with the RESYNC option to add the repaired volume pair to the PPRC session. The RESYNC option specifies that PPRC copy all changed cylinders to the secondary volume. If the repair action has caused the secondary volume’s data to change, issue the CESTPAIR command with the COPY option.

   _____

Attention: Failure to follow the above procedures may result in potential data integrity exposures to offline volumes.

If you are planning to repair or replace a primary volume, you may want to use the P/DAS function (see “Chapter 18. PPRC Dynamic Address Switching (P/DAS)” on page 251) to minimize impact to the application I/O while the error on the primary volume is being corrected.

Identifying volume pairs

The storage administrator must identify application volumes to copy, and create peer-to-peer remote copy volume pairs. PPRC provides an image copy of a volume on a record-for-record basis. There is a one-to-one correspondence between the record on the primary volume track and the record on the secondary volume track.

PPRC writes data on the same tracks on the secondary disk as it does on the primary disk. Therefore, the secondary disk must have the same track sizes and number of tracks per cylinder, and either the same or larger volume capacities, as the primary. If the disk does not meet these criteria, the CESTPAIR command will fail. For more information on the CESTPAIR command, see “CESTPAIR – establishing pairs” on page 202.

You should make secondary volumes part of a storage group that does not allow allocations, and dedicate these volumes to PPRC use only. PPRC secondary volumes are similar to 3390 dual copy secondary volumes. All host I/O operations directed to these offline secondary volumes are command rejected.

Remote copy supports volumes, not data sets. As a result, all data sets on the volumes that are copied are part of remote copy activity, and are therefore copied to the recovery system. This support is application (IBM or non-IBM) independent, and supports all data set types.

Because applications deal with data sets and not volumes, multivolume data sets require special attention. Unless you copy all volumes of a multivolume data set, PPRC will only copy part of the data set. Other data sets on the copied volume may be usable, but the multivolume data sets will not be. If multivolume data sets are critical for recovery, you must also copy the other volumes on which these data sets reside.

Note: Multivolume data set types include data sets that reside on multiple volumes, striped data sets, and VSAM spheres.

Carefully choose the volumes to copy so as not to unnecessarily affect overall PPRC performance. Copying page data sets can affect performance.

Example: It would be wasteful to have PPRC copy a volume that contains page or other data sets that are specific to the application host system.

Steps for including PPRC TSO commands in automated procedures

You can use PPRC TSO commands in automated procedures to do certain PPRC functions. The description below, for example, outlines an automated procedure to check PPRC volume status at system IPL time and resynchronize volumes that are in suspended or duplex state.

Before you begin: You need to be familiar with the commands used in these steps. For more information, refer to “Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199.

Perform the following steps to include PPRC TSO commands in automated procedures:

1. Put the CQUERY PPRC TSO command in either a batch JCL job or PROCLIB member to query the status of all primary volumes.

   _____

2. Automatically start the procedure at IPL time by putting the procedure name into SYS1.PARMLIB(COMMNDxx). PPRC sends the results from the query to the system’s log.

   _____

3. Use NetView®, or other console automation product, to capture the output from the CQUERY command and examine specific path and volume status information.

---

4. Use NetView, or another console automation product, to initiate jobs for any paths and volume pairs that you need to reestablish. These jobs can be other procedures and can use JCL substitution.

Notes:

1. Consider attaching a host to the secondary subsystem. If regular production operations are going to take place at the secondary subsystem, install an ECAM device there to monitor real space usage at that secondary LSS. An ECAM device cannot be a secondary volume of a PPRC pair, however, because the secondary device cannot receive I/O.
2. Some storage subsystems only send the SIM down the path that issues the next SIO. Many subsystems, however, send the message down every active path group. As a result, any host that is attached to the subsystem can receive the ERP SIM messages. This is an important consideration as you plan for disaster recovery.

Configuring storage subsystem resources for PPRC

All synchronous copy operations, including PPRC, have an inherent performance penalty: Write operations to the primary subsystem must wait until the data is secure in the secondary subsystem before sending completion status. This necessary step adds a minimum of an additional fast write operation plus the data transmission times.

You must properly configure your PPRC resources to do the following:

- Minimize additional impact on system performance.
- Avoid unintended interruptions of the remote copy function.
- Maximize the data integrity of the secondary volumes.

The following subtopics address these configuration requirements.

Configuring IBM storage controls

Ensure that the service information message (SIM) alert reporting is correctly set at both PPRC storage subsystems.

The following are additional concerns for the storage controls at the recovery site:

- Attachment to at least one host for SIM reporting
- Use of remote power sequence control cables

Storage subsystem settings for SIM alerts

Set up all storage controls with PPRC operations so that PPRC routes SIM alert messages to a host console. Do not select “No Alerts”, as this option prohibits the storage control from sending any SIM alert messages to attached host systems. See “Appendix B. Configuring the Data Mover and PPRC Options on IBM 3990 and 9390 Storage Controls” on page 305 for related PPRC switch information. In the event of a subsystem failure that affects PPRC operation (such as a permanent error on a secondary volume), the subsystem must be able to alert an operator to the problem. This is especially critical at an unattended recovery site where a subsystem problem, if left unresolved, can jeopardize the entire PPRC copy operation.

Attaching a recovery site storage control to a host system

It is recommended that you keep at least one host system attached to each recovery site storage control. You can locate the host at either site, as long as it can relay SIM alert messages to a console. This is necessary so that the storage controls for the secondary volumes can offload important SIMs in the event of a problem on the recovery storage subsystem.

Using remote power sequence control cables

Make every effort to minimize power interruptions to PPRC recovery site storage subsystems.

Recommendation: It is recommended that you do not use power sequence control cables between host processors and attached storage controls that contain PPRC secondary volumes. These cables, when installed, allow an attached host processor to remotely remove power from the storage controls. If power sequence control cables are present, ensure that the power-select switches for both clusters are in the LOCAL position. The power-select switches are on an IBM Storage Control's power sequence control board. This setting removes the possibility that a processor at the recovery site can inadvertently power down a PPRC recovery storage subsystem.

Accessing PPRC secondary volumes

Dedicate all PPRC secondary volumes exclusively to PPRC use. The system automatically rejects all host read and write access as long as the PPRC pair remains active in the session, regardless of whether the secondary volume is online or offline.

Optimizing PPRC performance

This section describes how synchronous operations (like peer-to-peer remote copy) can affect various work loads, and contains guidelines for getting the best PPRC performance. See your IBM marketing representative for the latest remote copy performance and configuration information and recommendations for your specific environment.

Analyzing work load characteristics

For synchronous operations there is additional overhead for every write. Several factors can affect the overall performance of write operations, especially for synchronous copy operations. The following work load characteristics can affect synchronous write performance:

- Write content — There is an inherent performance penalty on synchronous operations that are associated with write operations. As a result, work loads with high read-to-write ratios experience much less impact than those that are write intensive. As more writes occur to the primary site storage control, there is a corresponding effect on performance.

  Some examples of extremely write-intensive work loads are SPOOL volumes and database logging volumes, which can affect I/O response times dramatically in any synchronous copy operation. A possible solution to this high write-intensive work load is to spread the SPOOL volumes across multiple volumes to minimize queueing.
- Blocksize — Work loads with small blocksizes experience better performance than do work loads with large blocksizes. As the write blocksize increases, the affect on work load performance also increases. One example of a work load that has this impact is DB2® deferred writes. A possible solution is to spread the work load across several volumes.

- Overall channel demand — As the demand increases on the host channel that services the primary application's storage control, the impact to overall performance also increases. Spread the work load across several storage controls if possible.
- Sequential write operations — Batch processing operations typically exhibit high sequential write operations. Synchronous copy operations can result in longer batch windows.

  Evaluate dump restore activities to PPRC-managed volumes. PPRC uses a data flow algorithm to copy data from primary volumes to secondary volumes during initialization. To best optimize this algorithm, avoid restoring data sets or volumes to PPRC-managed volumes. Although PPRC supports restore operations, doing so increases the amount of data that PPRC must copy from the primary to the secondary volumes. You may choose to allocate a scratch volume on the primary system, and then restore data to that volume. Establish the PPRC session, and allow normal primary-to-secondary volume processing to continue.

Balancing storage control configurations

Configure your system so that primary system activity does not exceed the capacity of the recovery site storage control. You might require performance at the secondary site that is equivalent to your primary application's performance. This could be performance when migrating work loads, or performance during a disaster recovery. If so, it is recommended that the secondary site storage controls have at least the same cache and NVS capacity as the primary site storage controls.

PPRC supports configurations where several primary site storage controls can funnel their updated data to a single storage control on the recovery system. The number of copied primary volumes and the rate of write activity to these volumes can affect performance.

Ensure that you have a sufficient number of ESCON paths established between the primary and secondary site storage controls. This is especially important in configurations where the same storage controls manage both primary and secondary volumes.

Distributing work loads

Avoid directing all update activity to a small set of common volumes on a single recovery site storage control. The performance impact at the recovery site storage control adversely affects the primary site performance.

Typically, only a few devices on a particular storage control receive most of the activity. Usually two to four devices account for 50% or more of the total update activity. If, for example, two primary storage controls channel their most active volumes into a single recovery site storage control, the recovery subsystem could become overcommitted with remote copy activity.

Table 45 provides a useful summary of the most fascinating facts for PPRC session configurations.

Table 45. Quick reference for PPRC session limits

You can configure up to . . .		
4 PPRC recovery site storage subsystems	for each	PPRC primary site storage subsystem
8 PPRC paths	between	one primary site ESS storage subsystem and one recovery site ESS storage subsystem

Table 45. Quick reference for PPRC session limits (continued)

You can configure up to . . .		
32 PPRC paths	between	one PPRC primary site ESS storage subsystem and four PPRC recovery site ESS storage subsystems
4 PPRC paths	between	one primary site 3990 or 9390 Storage Control and one recovery site 3990 or 9390 Storage Control
16 PPRC paths	between	one PPRC primary site 3990 or 9390 Storage Control and four PPRC recovery site 3990 or 9390 Storage Controls
256 volumes	for each	ESS logical storage subsystem (LSS)
128 volumes	for each	3990 or 9390 Storage Control

Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions

This section describes peer-to-peer remote copy TSO commands with their required and optional parameters, and includes an example for each command. Chapter 14 contains other examples.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
PPRC command overview	199
CDELPAIR – deleting pairs	200
CDELPATH – deleting paths	201
CESTPAIR – establishing pairs	202
CESTPATH – establishing paths	205
CGROUP – controlling volume groups	208
CQUERY – querying status	209
CRECOVER – recovering data on the recovery system	210
CSUSPEND – suspending pairs	211

PPRC command overview

Peer-to-peer remote copy TSO commands have the same restrictions as other TSO commands. The TSO abbreviation convention requires that you specify as much of the command parameter as is necessary to distinguish it from other parameters. Do not abbreviate PPRC command names. Refer to *z/OS TSO/E Command Reference* for details on TSO commands. Enter hexadecimal parameters in the form `X'nn...'`, as indicated in the examples in Chapter 14.

PPRC TSO commands, like most TSO commands, are deferred requests. You receive an initial message that reports whether the command syntax is valid, and whether TSO has accepted the command for processing. Command function processing is a separate operation that can take minutes or hours to complete. The TSO terminal, meanwhile, is free to do other processing.

Once TSO is operational, you can issue the PPRC commands to automatically initiate PPRC subsystem activity. Table 46 shows which PPRC TSO commands you can issue to primary and secondary devices. TSO processes commands in the order that you issue them.

Table 46. Peer-to-peer remote copy TSO commands

TSO command	Can command be issued to a:	
	Primary device?	Secondary device?
CDELPAIR	Yes	No
CDELPATH	Yes	No
CESTPAIR	Yes	No

© Copyright IBM Corp. 1994, 2002

Table 46. Peer-to-peer remote copy TSO commands (continued)

TSO command	Can command be issued to a:	
	Primary device?	**Secondary device?**
CESTPATH	Yes	No
CGROUP	Yes	No
CQUERY	Yes	Yes
CRECOVER	No	Yes
CSUSPEND	Yes	Yes

Note: Do not use the NOCOPY option of the CESTPAIR command if you issue this command as part of automatically initiated PPRC subsystem activity, as this may introduce potential data integrity exposures. For more information on the CESTPAIR NOCOPY option, see "CESTPAIR command syntax" on page 202.

You can put all PPRC TSO commands in a SYS1.PROCLIB member, in CLISTs, and include them within standard job JCL.

PPRC uses the Write-to-Operator (WTO) route code to route all messages to the console.

CDELPAIR – deleting pairs

Use the CDELPAIR command to specify the primary and secondary volumes to remove from PPRC. Direct the command to the primary device.

See "Deleting a PPRC volume pair" on page 231 for a CLIST example, and more information about deleting pairs.

CDELPAIR command syntax

The syntax of the CDELPAIR command is:

```
►►—CDELPAIR—DEVN(device_number)—PRIM(ssid serialno cca lss)—SEC(ssid serialno cca lss)———————►◄
```

Required parameters

DEVN Specifies the device number of the primary volume for the volume pair that PPRC is to delete. The device number is the 4-digit hexadecimal address of the device to which the I/O operation is directed.

PRIM Specifies the primary site storage subsystem ID (*ssid*), and serial number (*serialno*), the primary volume channel connection address (*cca*), and the logical storage subsystem (*lss*).

Where:

ssid Specifies a 4-digit hexadecimal address.

serialno Specifies the storage control serial number, which can include up to 12 digits, depending on the type of storage control.

cca Specifies a 2-digit hexadecimal address, displayed with the CQUERY command

lss Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *lss* value for devices in storage subsystems that support logical subsystems, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the recovery site storage subsystem ID and serial number, and the secondary volume channel connection address. The descriptions for the values are the same as for the PRIM parameter.

CDELPAIR command examples

The following examples show the CDELPAIR command.

```
  CDELPAIR DEVN(X'0C04') PRIM(X'6060' 62019 X'04') SEC(X'6061' 68006 X'09')

For devices in logical subsystems, such as ESS:

  CDELPAIR DEVN(X'0D20') PRIM(X'6114' 66143 X'08' X'06') SEC(X'6125' 68247 X'0B' X'03')
```

CDELPATH – deleting paths

Use the CDELPATH command to delete all established ESCON paths between a primary site (source) storage control and a recovery site (target) storage control. Only active paths to the specified recovery site storage control are affected; all other paths to other storage controls are unaffected.

See “Deleting PPRC paths” on page 232 for a CLIST example and for more information about deleting paths.

Note: You can use the CESTPATH command to remove only selected paths. See the note under “CESTPATH – establishing paths” on page 205 for this information.

CDELPATH command syntax

The syntax of the CDELPATH command is:

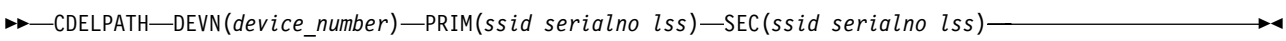

Required parameters

DEVN Specifies the device number of any disk volume that is attached to the primary site storage control without directing any operation to the volume. The device number is the 4-digit hexadecimal address of any disk volume behind the primary site storage control.

PRIM Specifies the primary storage subsystem ID (*ssid*), serial number (*serialno*), and logical storage subsystem (*lss*).

Where:

ssid Specifies a 4-digit hexadecimal address.

serialno Specifies the storage control serial number that can include up to 12 digits, depending on the type of storage control.

lss Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *lss* value for devices in storage subsystems that support logical subsystems, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the recovery site storage subsystem ID and serial number. The descriptions for the values are the same as for the PRIM parameter.

CDELPATH command examples

The following are examples of the CDELPATH command:

```
CDELPATH DEVN(X'0C04') PRIM(X'6060' 62019) SEC(X'6061' 68006)

CDELPATH DEVN(X'0D20') PRIM(X'6114' 66143 X'06') SEC(X'6125' 68247 X'12')
```

CESTPAIR – establishing pairs

Use the CESTPAIR command to specify PPRC primary and secondary volumes.

The primary and secondary volumes must have the same number of tracks on each cylinder and the same number of bytes on each track. The secondary device must have the same number or a greater number of cylinders as compared to the primary device.

Notes:

1. Peer-to-peer remote copy allows you to copy data from a smaller primary device (fewer cylinders) to a larger secondary device (more cylinders). The track size and number of tracks per cylinder, however, must be the same. The additional space on the secondary volume may not be immediately usable. End the PPRC volume pair and run ICKDSF Release 16 or above against the secondary volume to rebuild the VTOC and make the additional space usable.
2. Because multivolume data sets span multiple volumes, they require special handling. To ensure that the data set is recoverable, copy all volumes of multivolume and striped data sets.
3. For more information and examples, see “Adding a PPRC volume pair” on page 221.

CESTPAIR command syntax

The syntax of the CESTPAIR command is:

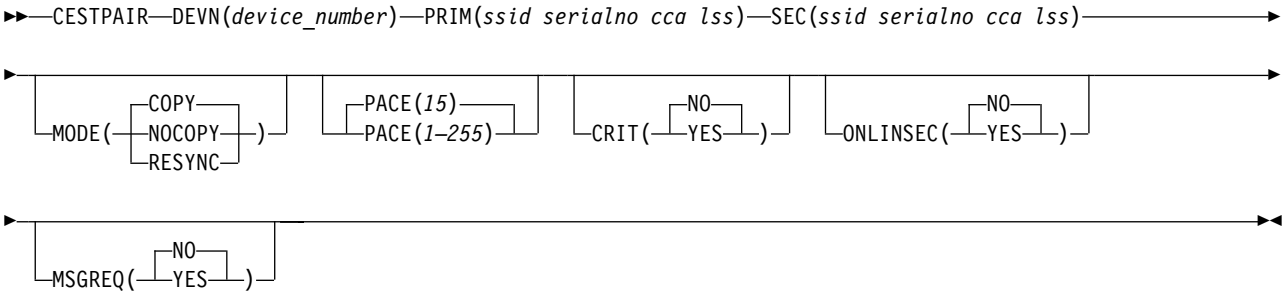

Required parameters

DEVN Specifies the device number of the primary volume for the volume pair to be established. The device number is the 4-digit hexadecimal address of the device to which the I/O operation is directed.

PRIM Specifies the primary volume's storage subsystem ID (*SSID*) and storage control serial number (*Serialno)*, primary volume channel connection address (*CCA*), and logical storage subsystem (*LSS*).

Where:

SSID Specifies a 4-digit hexadecimal address.

Serialno Specifies the storage control serial number that can include up to 12 digits, depending on the type of storage control.

CCA Specifies a 2-digit hexadecimal address, displayed with the CQUERY command

LSS Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *lss* value for devices in storage subsystems that support logical subsystems, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the secondary volume's storage subsystem ID and serial number, and the secondary volume channel connection address. The descriptions for the values are the same as for the PRIM parameter.

Optional parameters

MODE Specifies one of the following PPRC modes:

COPY Specifies that PPRC copy all tracks on the primary volume to the secondary volume. COPY is the default mode.

NOCOPY Specifies that only those tracks on the primary volume that are updated after this command is issued are to be copied to the secondary volume. Specify this mode only if the volumes are *exact* copies of each other.

The NOCOPY option assumes that the specified volumes are fully synchronized at the time the copy is started. Ensure that no I/O operations occur to either volume prior to starting the copy operation. This is necessary in order to protect the data integrity of the two volumes.

For PPRC Volume Pairs of Different Device

Types:
The following notes apply only when the volumes of a PPRC volume pair are not of the same device type (volumes do not have the same number of alternate tracks):

1. If you have established or reestablished this PPRC volume pair without specifying COPY, run ICKDSF REFORMAT REFVTOC on the primary volume after you establish the pair. Doing so

allows the VTOC and indexed VTOC information to accurately represent the primary volume.

2. If you switch the primary and secondary of this pair at any time, run ICKDSF REFORMAT REFVTOC on the new primary volume. Doing so allows the VTOC and indexed VTOC information to accurately represent the new primary volume.
3. If you remove the secondary volume from the volume pair for use as a simplex volume, you should run ICKDSF REFORMAT REFVTOC on the volume. Doing so allows the VTOC and Indexed VTOC information to accurately represent the secondary volume.

RESYNC Specifies that PPRC reestablish a suspended copy. Before you issue this option, be certain that either the primary and secondary volumes are both in a suspended state, or the secondary is in a simplex state and the primary is suspended. PPRC only copies from the primary volume to the secondary volume the cylinders that were updated during the period of suspension.

ONLINSEC Specifies whether the secondary volume attached to an ESS is online to a host system. The values are:

NO Specifies that, if the secondary volume that is attached to an ESS subsystem appears to be online to any host system, the CESTPAIR command will fail with an error. The default is NO.

YES Specifies that the CESTPAIR command be processed even if the ESS secondary volume appears to be online to a host system.

PACE Specifies the minimum number of tracks that PPRC is to copy at one time before allowing another host interrupt. The allowable range is from 1 to 255 tracks. If you omit this parameter, PPRC uses the storage control default value. PPRC ignores the PACE parameter if you also specify the NOCOPY parameter. PPRC ignores the PACE parameter with storage subsystems that support logical storage subsystems (LSS).

Table 47 shows the result of different PACE parameter values.

Table 47. PACE parameter values

Value	Effect on the CESTPAIR copy operation
1	Copy a maximum of three tracks at a time
2 to 255	Copy a maximum of 15 tracks at a time

CRIT Specifies whether PPRC allows subsequent write requests to the primary volume.

NO Specifies that, following an I/O completion error, PPRC

allows subsequent write requests to the PPRC primary volume. The primary storage control will perform change recording. The default is NO.

YES Specifies that if an I/O error occurs, PPRC either allows or does not allow subsequent writes, depending on how the storage subsystem is configured. The PPRC pair then remains in a suspended state until you correct the problem and issue either a CESTPAIR RESYNC or CDELPAIR command. See “Considering the PPRC solution” on page 184 and “Examining disaster recovery from a system viewpoint” on page 233 for related information.

MSGREQ The MSGREQ option is valid only with the COPY option.

NO Specifies that the command is complete as soon as you initiate the request to the storage control. PPRC does not wait for the copy operation to complete before issuing completion messages. The default is NO.

YES Specifies that PPRC wait until the initial full-volume copy operation is complete before issuing completion message ANTP0001I.

CESTPAIR command example

The following is an example of the CESTPAIR command:

```
CESTPAIR DEVN(X'0C02') PRIM(X'6060' 62019 X'02') SEC(X'6061' 68006 X'09')
         PACE(1) CRIT(N)
```

CESTPATH – establishing paths

Use the CESTPATH command to establish ESCON paths between a primary site (source) storage control and a recovery site (target) storage control.

Each CESTPATH command can establish up to eight paths (if the primary site storage control supports eight paths) from one primary site storage control to a single recovery site storage control. You can link up to four recovery site storage controls to one primary site storage control. These links require that you issue a separate CESTPATH command for each storage control.

Attention: The CESTPATH command is a “replace” function. Specified paths that establish a link between a primary site storage control and a recovery site storage control replace any previously established PPRC paths. Each time you issue a CESTPATH command, the paths specified on it replace the paths that are established by the last CESTPATH command that was issued. You can use the CESTPATH command to selectively add and remove paths. Use caution, however, to ensure that you do not remove paths that you want to keep.

CESTPATH command syntax

The syntax of the CESTPATH command is:

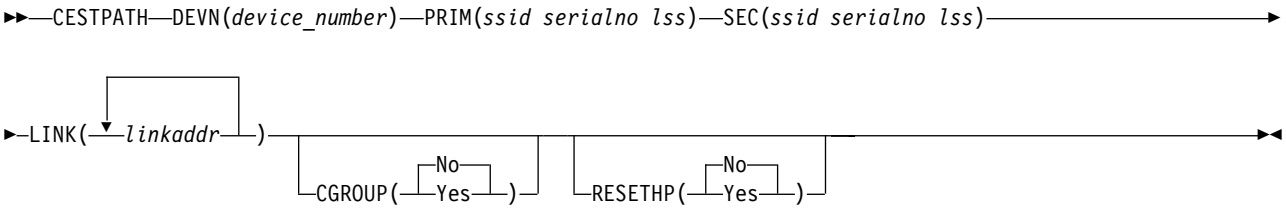

Required parameters

DEVN Specifies the device number of any simplex disk volume that is attached to the primary site storage control without directing any operation to the volume. The device number is the 4-digit hexadecimal address of any disk volume behind the primary site storage control.

PRIM Specifies the primary site storage subsystem ID (*SSID*), serial number (*serialno*), and logical subsystem (*LSS*).

Where:

SSID Specifies a 4-digit hexadecimal address.

Serialno Specifies the serial number, which can include up to 12 digits, depending on the type of storage control.

LSS Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *LSS* value for the ESS, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the recovery site storage subsystem ID and serial number. The descriptions for the values are the same as for the PRIM parameter.

LINK Specifies the addressing paths from the primary site storage control to the recovery site storage control. See “Establishing paths” on page 218 for more information on the LINK parameter.

You can specify up to eight path addresses (if the primary site storage control supports eight paths), separated with a blank, for each LINK parameter.

Each path address is an 8-digit hexadecimal address in the form: *aaaabbcc* where:

aaaa Specifies the System Adapter ID (SAID) interface identifier on the primary system that uniquely identifies the physical location of the associated storage subsystem interface. The first two values are zeros (00*aa*), and the last two values are hexadecimal SAID byte values that correspond to ESCON channel interfaces, as follows:

```
bit 0 = storage cluster (0 or 1)
bit 1 = 0
bit 2 = I/O board (0 or 1)
bit 3 = 0
bits 4-5 = slot (00, 01, 10, or 11)
bit 6 = 0
bit 7 = port number (0 or 1 for each slot)
```

bb Specifies the recovery site storage control ESCON link destination address. If an ESCON director is installed, it assigns the recovery

site storage control ESCON link destination address (bb). For direct fiber-optic cable connections between storage controls or static connections through an ESCON Director, the link address value is zero.

cc This value is always *00*, except for storage subsystems that have logical control units (LCUs). For storage subsystems that support LCUs, see the documentation for that storage subsystem.

Note: The LSS concept does ***not*** apply to the LINK parameter.

Optional parameters

CGROUP Specifies how an error that is related to a PPRC device pair is reported by a storage subsystem that contains LSSs. PPRC ignores the CGROUP parameter if the specified device is not in a storage subsystem that supports LSSs.

NO Specifies that, in subsystems that support LSSs, PPRC device pair error handling and reporting is performed the same way as a 3990 Storage Control with the automated recovery bit in the vital product data (VPD) set to Off. NO is the default.

YES Specifies that the subsystem respond to PPRC pair errors with the same status that is returned from a 3990 Model 6 Storage Control with the automated recovery bit set to On in the VPD. This bit enables the extended long busy condition for a consistency group. Automation products will be able to detect PPRC volume pair outages on these devices and take appropriate action.

RESETHP At this time, only the ESS subsystem supports the RESETHP option.

NO Specifies that PPRC *not* reset existing host system paths to the primary subsystem before establishing links to the secondary subsystem. The default is NO.

YES Specifies that PPRC reset existing host system paths to the primary subsystem before establishing links to the secondary subsystem.

CESTPATH command examples

The following are examples of the CESTPATH command:

```
CESTPATH DEVN(X'0142') PRIM(X'6060' 62019) SEC(X'6061' 68006)
          LINK(X'0004C400')

CESTPATH DEVN(X'0D20') PRIM(X'6114' 66143 X'06') SEC(X'6125' 68247 X'12')
          LINK(X'0004C400') RESETHP(YES)

CESTPATH DEVN(X'0D20') PRIM(X'6114' 66143 X'06') SEC(X'6125' 68247 X'12')
          LINK(X'0004C400') CGROUP(YES)
```

CGROUP – controlling volume groups

Issue the CGROUP command to control operations for multiple PPRC volume pairs on a single storage control. This command lets you suspend or resume all operations for all PPRC volumes on a single storage control. You must issue a separate CGROUP command to suspend or resume operations on each storage control.

CGROUP command syntax

The syntax of the CGROUP command is:

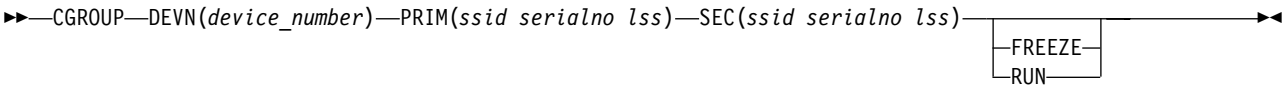

Required parameters

DEVN Specifies the device number of any disk volume that is attached to the primary site storage control without directing any operation to the volume. The device number is the 4-digit hexadecimal address of any disk volume behind the primary site storage control.

PRIM Specifies the primary site storage subsystem ID (*SSID*), serial number (*serialno*), and logical storage subsystem (*LSS*).

Where:

SSID Specifies a 4-digit hexadecimal address.

Serialno Specifies the serial number, which can include up to 12 digits, depending on the type of storage control.

LSS Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *LSS* value for the ESS, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the recovery site storage subsystem ID and serial number. The description for the values is the same as for the PRIM parameter.

Optional parameters

FREEZE Specifies that PPRC stop all updates to secondary storage controls that are related to this volume pair. PPRC does not allow any activity to the secondary volumes. You cannot specify the FREEZE parameter with the RUN parameter.

RUN Specifies that PPRC resume all operations for all PPRC volumes for the storage control. I/O to the primary volumes resumes, but the pairs remain suspended. You cannot specify the RUN parameter with the FREEZE parameter.

CGROUP command examples

The following are examples of the CGROUP command:

```
CGROUP DEVN(X'0F40') PRIM(X'6060' 62019) SEC(X'6061' 68006) FREEZE

CGROUP DEVN(X'0F40') PRIM(X'6060' 62019) SEC(X'6061' 68006) RUN
```

CQUERY – querying status

Use the CQUERY command to query the status of one volume of a PPRC volume pair, or all paths that are associated with the storage control for the device number that you specify. You can issue the CQUERY command to either a primary or secondary PPRC volume. A host system that is attached only to a primary volume cannot obtain the status of the secondary volume for that pair. In the same way, a host attached only to the secondary volume cannot obtain the status of the primary volume.

See “Querying PPRC volumes” on page 222 and “Querying PPRC paths” on page 227 for CQUERY command output examples for a volume, and for paths.

CQUERY command syntax

The syntax of the CQUERY command is:

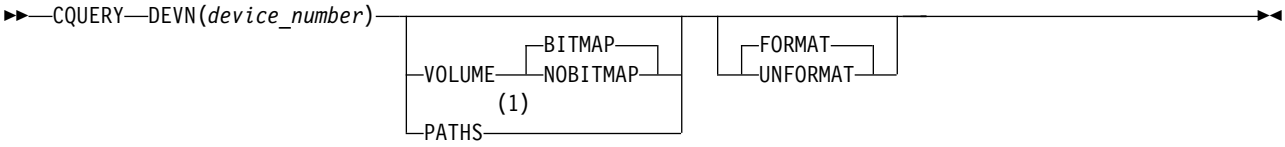

Notes:

1 BITMAP and NOBITMAP are ignored if you specify the PATHS parameter.

Required parameters

DEVN Specifies the device number of the volume that PPRC is to query. The device number is the 4-digit hexadecimal address of the device to which the I/O operation is directed.

Optional parameters

VOLUME Directs PPRC to display the status of the volume that is specified with the DEVN parameter. The default value is VOLUME.

PATHS Directs PPRC to display all of the paths that are associated with the primary site storage control, and the status of each path. If you direct this command to the secondary device, PPRC cannot make path or storage control information available. To obtain this information, issue the command to a volume on that storage control that is either in simplex state or is a primary of a duplex pair.

BITMAP Specifies that PPRC report the percentage of tracks that it has copied for a pending or suspended PPRC primary volume. This option is valid when you query a PPRC primary volume and you either specify the VOLUME keyword, or you do not specify any optional parameters. BITMAP is ignored if you specify the PATHS keyword. If you specify the VOLUME keyword, the default value is BITMAP.

NOBITMAP Specifies that PPRC not report the percentage of tracks that it has copied for a pending or suspended PPRC primary volume. This option is valid when you query a PPRC primary volume and you either specify the VOLUME keyword, or you do not specify any optional parameters.

You can use the NOBITMAP option to collect some data from a storage control even when the device is unable to respond to the CQUERY command.

Example: PPRC may not be able to collect bitmap information because the device is disabled.

FORMAT Specifies that PPRC present the CQUERY command output in a block format. FORMAT may be specified with either the VOLUME option (as in the example on page 223) or with the PATHS option (as in the example on page 227). The default value is FORMAT.

PPRC issues message ANTP0090I as part of a block report when you specify the CQUERY VOLUME with FORMAT option. PPRC issues message ANTP0095I as part of a block report when you specify the CQUERY PATH with FORMAT option.

UNFORMAT Specifies that PPRC present the CQUERY command output as in the examples on page 224 for a volume query and page 227 for a paths query. UNFORMAT is valid for both the VOLUME option and the PATHS option.

PPRC issues message ANTP0091I as part of a block report when you specify the CQUERY VOLUME with UNFORMAT option.

PPRC issues message ANTP0096I as part of a block report when you specify the CQUERY PATHS with UNFORMAT option.

CQUERY command examples

The following are examples of the CQUERY command:

```
CQUERY DEVN(X'0C04') PATHS UNFORMAT

CQUERY DEVN(X'0C04') VOLUME

CQUERY DEVN(X'0C04') NOBITMAP

CQUERY DEVN(X'0C04') VOLUME BITMAP
```

CRECOVER – recovering data on the recovery system

Use the CRECOVER command to allow the recovery system to gain control of a disk volume on its storage control. This command is issued from the recovery system. It signals the recovery site storage control to force the secondary volume into simplex state to establish recovery system control. During this process, verify the volume serial number, and relabel the volume if you desire. You can vary the volume online after this command function is complete.

CRECOVER command syntax

The syntax of the CRECOVER command is:

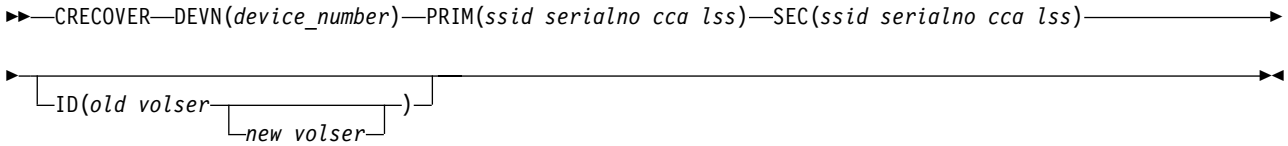

Required parameters

DEVN Specifies the secondary volume that is attached to the recovery site host. The device number is the 4-digit hexadecimal address of the device to which the I/O operation is directed.

PRIM Specifies the primary volume's storage subsystem ID (*SSID*), serial number (*serialno*), the primary volume channel connection address (*CCA*), and the logical storage subsystem (*LSS*).

Where:

SSID Specifies a 4-digit hexadecimal address.

Serialno Specifies the serial number, which can include up to 12 digits, depending on the type of storage control.

CCA Specifies a 2-digit hexadecimal address.

LSS Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *LSS* value for the ESS, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the secondary volume's storage subsystem ID and serial number, and the secondary volume channel connection address. The descriptions for the values are the same as for the PRIM parameter.

Optional parameters

ID Specifies the six-character volume serial number. If you do not specify the ID parameter, PPRC resets the volume to simplex state.

old volser Specifies the volume serial number.

If you specify *old volser* without specifying *new volser*, PPRC verifies the serial number and resets the volume to simplex state.

new volser Specifies the new volume serial number.

When you specify *new volser* with *old volser*, PPRC verifies the old serial number and resets the volume to the simplex state. PPRC then replaces the old serial number with the new serial number.

CRECOVER command example

The following is an example of the CRECOVER command :

```
CRECOVER DEVN(X'0C04') PRIM(X'6060' 62019 X'04') SEC(X'6061' 68006 X'09')
         ID(OLD408 NEW508)
```

CSUSPEND – suspending pairs

Use the CSUSPEND command to suspend PPRC operations between a primary and secondary volume pair. PPRC stops transferring data to the secondary volume. The primary storage subsystem, however, records all cylinders that change on the primary volume (except when you specify the PRIMARY option, as described below).

You can direct the CSUSPEND command to either the primary or secondary volume of a PPRC volume pair as a result of the DEVN parameter.

Notes:

1. The ESS storage subsystem does not accept the CSUSPEND command with either the PRIMARY or QUIESCE parameter, and fails the command with a command reject if you specify either parameter.
2. Avoid issuing a CDELPAIR command following a CSUSPEND command, as this sequence can cause unpredictable results.

You can issue the command with one of three combinations of the PRIMARY and QUIESCE optional parameters. The following are the three options and the results of these combinations:

When you specify neither the PRIMARY nor the QUIESCE option —
PPRC suspends the pair. Applications continue to update the primary volume. These updates are not copied to the secondary volume, but the storage control records which cylinders are changed while the pair is suspended. When you reestablish the pair, specify the RESYNC parameter of the CESTPAIR command to direct PPRC to copy only the changed cylinders to the secondary volume. PPRC then returns the secondary volume to duplex state.

This is the default option, and you can address this option to the primary or the secondary volume.

When you specify the PRIMARY option —
PPRC suspends the pair and unit checks all write data I/O that is directed to the device, except for ICKDSF media maintenance channel programs. When ICKDSF media maintenance operations have completed, you can reestablish the pair with the CESTPAIR command with the RESYNC option. (See “Controlling ICKDSF activity to PPRC volumes” on page 191 for a PPRC volume repair procedure.) You should only use the PRIMARY option as part of an ICKDSF media maintenance procedure, and the DEVN parameter must specify the primary volume.

When you specify the QUIESCE option —
PPRC defers all read data I/O and write data I/O that are directed to the device until the QUIESCE condition ends. You can end the QUIESCE condition by issuing any of the following commands:

- Issue a CSUSPEND command to the primary DEVN (without the QUIESCE or PRIMARY options)
- Issue a CSUSPEND command to the primary DEVN (with the PRIMARY option but without the QUIESCE option)
- Issue a CSUSPEND command to the secondary DEVN

The DEVN parameter must specify the primary volume when you specify the CSUSPEND option.

QUIESCE option notes:

1. APAR OW15247, APAR OW15248, or APAR OW40367 disables the QUIESCE option of the CSUSPEND command. Refer to the description of the QUIESCE option later in this section for important information about enabling this option.
2. Do not issue a CQUERY command to either member of a volume pair that you have suspended with a CSUSPEND QUIESCE command. Doing so forces the storage control to return a long busy condition to all attached host systems.

Issue the CESTPAIR command with the RESYNC option to resume PPRC operations for the volume pair.

CSUSPEND command syntax

The syntax of the CSUSPEND command is:

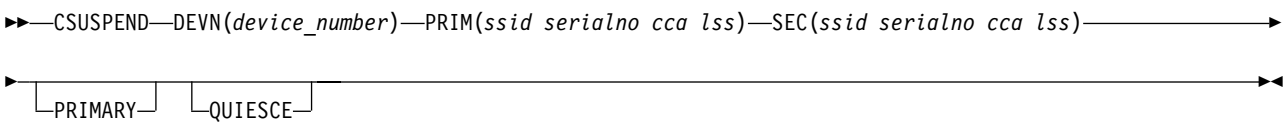

Required parameters

DEVN Specifies the device number of the primary or secondary volume of the volume pair that PPRC is to suspend. The device number is the 4-digit hexadecimal address of the device to which the I/O operation is directed.

PRIM Specifies the primary site storage subsystem ID (*SSID*), serial number (*serialno*), channel connection address (CCA), and logical subsystem (*LSS*).

where:

SSID Specifies a 4-digit hexadecimal address.

Serialno Specifies the serial number, which can include up to 12 digits, depending on the type of storage control.

CCA Specifies a 2-digit hexadecimal address.

LSS Specifies a 2-digit hexadecimal value that specifies the logical subsystem for the device.

Note: PPRC requires the *LSS* value for the ESS, but *fails* the command for devices in all other storage subsystems.

SEC Specifies the secondary volume's storage subsystem ID and serial number, and the secondary volume channel connection address. The descriptions for the values are the same as for the PRIM parameter.

Optional parameters

PRIMARY Directs PPRC to suspend the PPRC volume pair and to reject all write data I/O to the primary volume except for ICKDSF media maintenance channel programs. PRIMARY is valid only when directed to a primary volume with the DEVN parameter. The ESS storage subsystem does not accept the CSUSPEND PRIMARY parameter, and fails the command.

Use the PRIMARY option only as part of an ICKDSF media maintenance procedure because it causes the storage control to unit check all application write I/O.

QUIESCE Directs PPRC to suspend the PPRC volume pair and to send a "state change pending" indication to MVS to prevent issuance of any primary system I/O. QUIESCE is valid only when directed to a primary volume with the DEVN parameter. The ESS storage subsystem does not accept the CSUSPEND QUIESCE parameter, and fails the command with a command reject. PPRC performs only the QUIESCE option when you specify both the QUIESCE and PRIMARY options.

Attention:
APAR OW15247, APAR OW15248, or APAR OW40367 disables the QUIESCE option of the CSUSPEND command. Refer to either of these APARs for instructions about enabling the QUIESCE option of the CSUSPEND command.

If the QUIESCE option is enabled, do not issue a CSUSPEND command with the QUIESCE option to a volume with active SPOOL, PAGE, or CATALOG data sets, or to an active SYSRES volume. Doing so can cause any attached system to enter a deadlock condition, and may require a storage control IML to correct the condition.

You normally issue the CSUSPEND command with QUIESCE to all PPRC-copied volumes and follow it immediately by CSUSPEND without QUIESCE to the same set of volumes. This sequence allows you to suspend the remote copy operation with all secondary volumes in a consistent state.

CSUSPEND command example

The following is an example of the CSUSPEND command:

```
CSUSPEND DEVN(X'0C04') PRIM(X'6060' 62019 X'04') SEC(X'6061' 68006 X'09')
         PRIMARY
```

Chapter 14. Managing Peer-to-Peer Remote Copy Operations

This section explains the major PPRC system dependencies, and describes how to identify PPRC volume states, establish PPRC paths, and manage PPRC volumes and paths.

See “Chapter 15. PPRC Data Recovery Operations” on page 233 for disaster recovery procedures, and “Chapter 17. Recovering from PPRC Error Conditions” on page 243 for information related to error conditions within the PPRC environment.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Managing the PPRC system	215
Identifying peer-to-peer remote copy volume states	216
Establishing paths	218
Managing volumes and paths	221
Querying PPRC paths	227
Suspending PPRC volume pair operations	230
Deleting a PPRC volume pair	231
Deleting PPRC paths	232

Managing the PPRC system

The topics that follow describe how peer-to-peer remote copy interacts with various system components, and how these components affect PPRC.

System IPL volumes

The operating system views PPRC secondary volumes much like dual copy volumes, and permits only a limited set of I/O operations to these volumes. Therefore, in order to issue the CRECOVER command during recovery, you must IPL a recovery system that is separate from all PPRC volumes. The CRECOVER command removes all PPRC volumes from secondary PPRC volume status.

Attention: Do *not* allow the data sets that you will require to initialize the recovery system to be PPRC secondary volumes. These volumes include SYSRES, PAGE, and SPOOL volumes, and volumes that are used to startup JES and TSO.

Alternatively, you can have a second set of SYSRES, PAGE, and SPOOL volumes that PPRC keeps current. You would then use a separate set of these volumes for recovery. When recovery is complete, you can use the separate set of SYSRES, PAGE, and SPOOL volumes to initialize the recovered system. Another option is to use ICKDSF Stand-Alone to restore the necessary system data sets.

ESCON manager

PPRC paths and operations are not known to the ESCON Manager. As a result, you must consider all of the existing physical path connections when you make PPRC changes to the ESCON Director. This includes configuration changes that

© Copyright IBM Corp. 1994, 2002

are made at the ESCON Director console or through the integrated systems management feature of the ESCON Manager, because no protection (using ESCON Manager) is available.

Resource Measurement Facility™ (RMF™) data

Each primary host system write I/O operation passes through the PPRC connection to the recovery site storage control, and then to the recovery system disk devices. RMF statistics do not currently report PPRC connection activity.

RMF reports list reasons for I/O operation delays, including storage control, and ESCON Director port busy states. Even though PPRC activity may be responsible for these busy conditions, RMF does not explicitly track PPRC activity. PPRC activity may delay primary system I/O operations. You should consider this when you do your capacity planning as well as RMF data interpretation.

DEVSERV and IDCAMS

Use the DEVSERV and IDCAMS functions to determine the status of the storage control's dual copy, cache, and NVS functions. These functions do not report PPRC copy volume pairs or states. The CQUERY TSO command is the only way to obtain information about PPRC volumes.

Identifying peer-to-peer remote copy volume states

In order to manage peer-to-peer remote copy (PPRC) operations, you need to know the state of PPRC volumes. Figure 11 on page 217 shows the different volume states possible with PPRC.

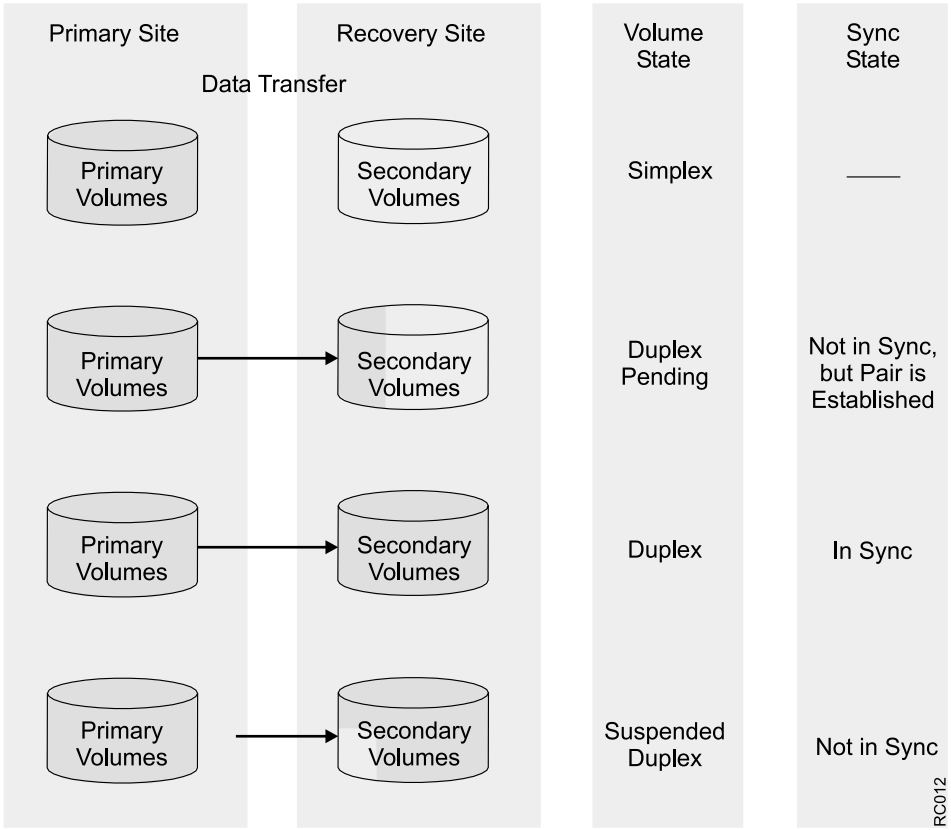

Figure 11. PPRC volume states

To determine the state of a volume, issue the CQUERY command to that volume. “Querying PPRC volumes” on page 222 describes an example of the command and the returned result.

As Figure 11 explains, at any given time a volume can be in one of the following states:

- Simplex
- Pending
- Duplex
- Suspended

The following definitions apply to volume states, as follows:

Simplex The initial state of a volume.

Pending The initial state of a defined volume pair. This state can also occur after a storage subsystem failure when PPRC reestablishes the volume pair. PPRC is in the process of copying data from the primary volume to the secondary volume.

Duplex The state of a volume pair after PPRC has completed the copy operation, and the volume pair is in synchronization.

Suspended The state of a volume pair when the storage subsystems cannot keep the primary and secondary volumes synchronized, or when either the primary host or recovery site host issues a CSUSPEND command. A PPRC volume pair will go into suspended state, for instance, when the primary system storage control fails to complete a write operation to the recovery system storage control. During this

suspended state, the primary volume's storage control records the cylinders that applications update. When you issue a CESTPAIR command with the RESYNC parameter, PPRC must only recopy data in the cylinders that have changes. Both volumes return to a synchronized state.

Establishing paths

Use the CESTPATH command to establish paths between primary site and recovery site storage controls. The storage control for a primary volume can have up to eight paths established between it and the storage control for the pair's secondary volume. Each CESTPATH command can establish up to eight of these paths.

The storage control for a specific primary (source) volume may, at any one time, link to a maximum of four storage controls for secondary (target) volumes. You must issue a separate CESTPATH command to establish each path between storage controls. See Figure 10 on page 183 for examples of PPRC path options.

Establish paths before issuing CESTPAIR commands. If you intend to establish paths in both directions between the storage control and host system (as is needed for P/DAS operations), establish each path before you establish the volume pairs.

The primary volume's storage control keeps information about the paths that connect it to the secondary volume's storage control. The primary volume's storage control automatically attempts to restart a failed path, and then resumes copy operations. If all paths to the secondary volume fail, PPRC suspends the pair. You must reestablish these paths with the CESTPATH command after you have corrected the cause of the failed paths.

Figure 12 on page 219 shows an example of two separate storage subsystems. In this example, device number 142 is the source volume (the volume that you plan to copy) and becomes the primary volume in the PPRC duplex volume pair. With respect to primary volume 142, the primary system consists of the host processor with ESCON channels, the ESCON Director (ESCD), and Storage Control A. Target volume 262 (the volume to contain the copy) resides on Storage Control B, and becomes the secondary volume in the duplex volume pair. Storage Control B, which may be in the same building or may be many kilometers away, is therefore part of the recovery system for this volume pair.

This installation has established a connection from the host to the primary volume through ESCD ports A6 and A5, and a host-to-secondary connection through ESCD ports A6 and B0.

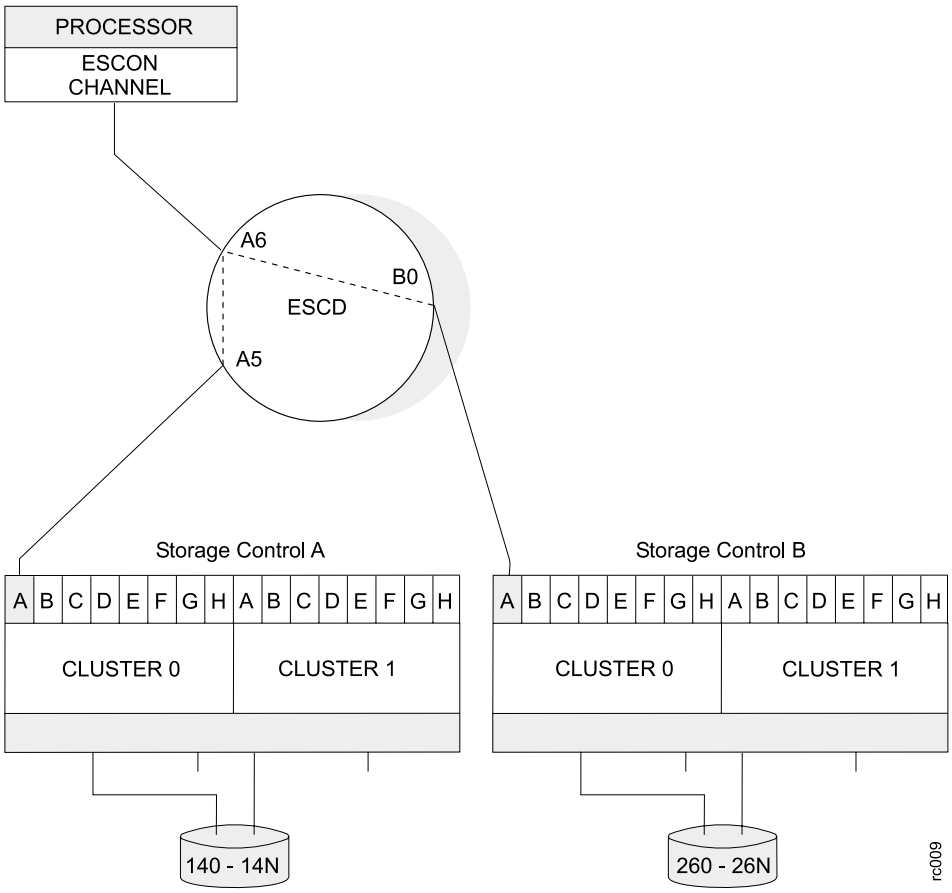

Figure 12. Storage subsystems for primary and secondary volumes

Figure 13 on page 220 shows a PPRC link established between (primary) storage control A and (recovery) storage control B. The CESTPATH command 'LINK(aaaabbcc)' parameter to establish this link would be LINK(0000B000), where the following terms apply:

```
aaaa = primary volume's storage control cluster 0, interface A
  bb = the link destination address, ESCD port B0
  cc = this value is always 00, except for storage subsystems
       that have logical control units (LCUs).
```

Notes:

1. For storage subsystems that support LCUs, see the documentation for that storage subsystem.
2. Before you issue the CESTPATH command, you may want to issue a CQUERY command to a device within each storage subsystem to determine the SSID and the serial number. The CESTPATH command requires this information. You can address the CQUERY command to either a simplex device or a duplex device.
3. You can collect path information with ICKDSF (release 16 or above) using the ANALYZE command, specifying the NOSCAN and NODRIVE keywords. Refer to *Device Support Facilities User's Guide and Reference* for more information.

The CESTPATH command for the example in Figure 13 on page 220, using primary volume 0142, is:

```
CESTPATH DEVN(X'0142') PRIM(X'6060' 62019) SEC(X'6061' 68006)
         LINK(X'0000B000')
```

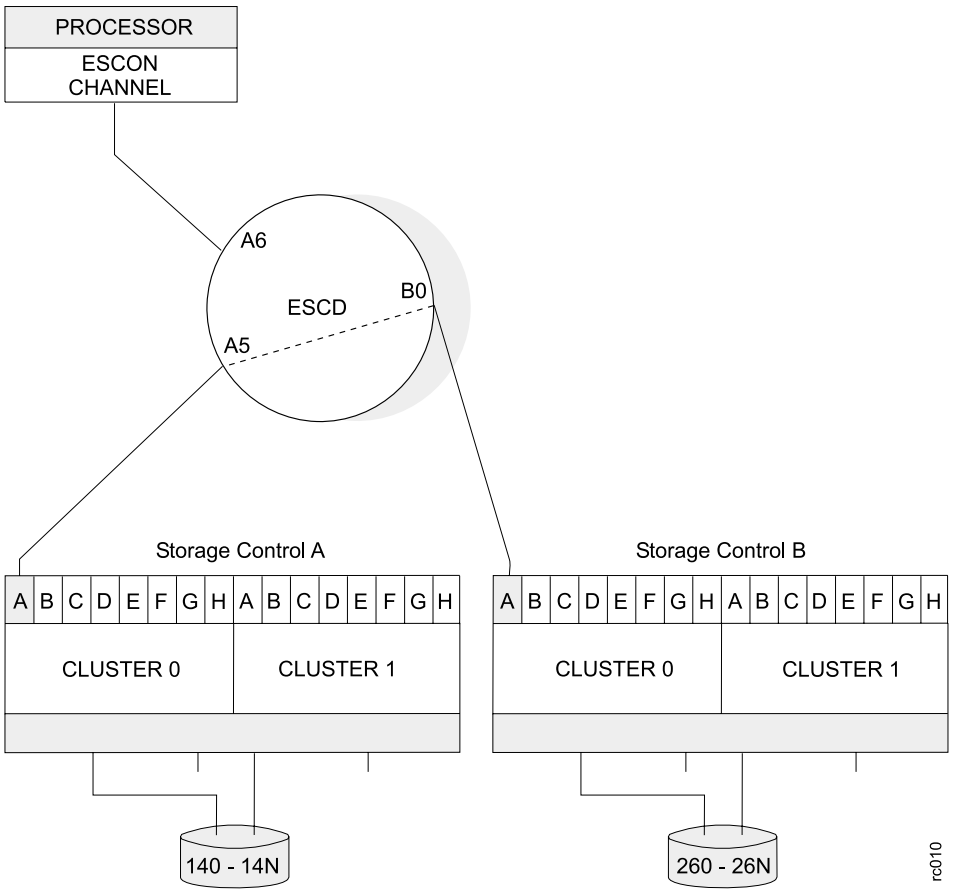

Figure 13. Establishing a PPRC path through an ESCD

Note: The CESTPATH establishes the PPRC link, only. You must also previously configure the ESCD to establish the logical system connection between ports A5 and B0.

You do not need to go through an ESCON Director to establish a PPRC link. Figure 14 on page 221 shows an example of a direct link between storage subsystem A and storage subsystem B. The installation has installed an ESCON fiber optic cable between them.

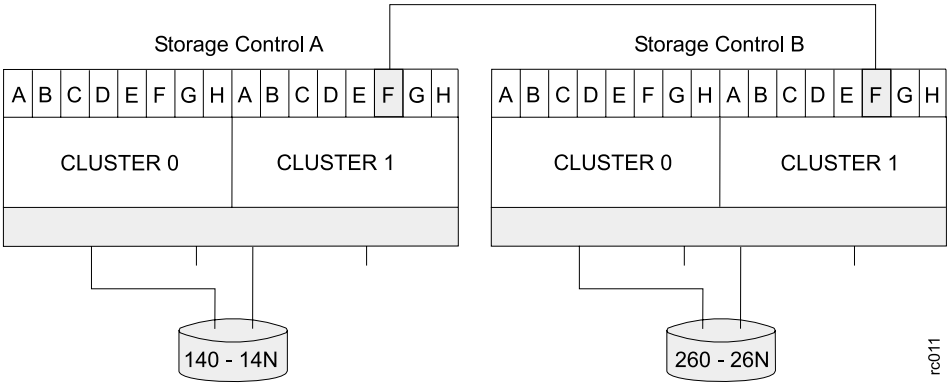

Figure 14. Establishing a direct PPRC path

The 'LINK(aaaabbcc)' parameter of the CESTPATH command for this example would be LINK(00150000), where the following terms apply:

```
aaaa = primary volume's storage control cluster 1, interface F
  bb = 00, as this is a direct connection
  cc = this value is always 00, except for storage subsystems
       that have logical control units (LCUs)
```

Notes:

1. For storage subsystems that support LCUs, see the documentation for that storage subsystem.
2. More information on the CESTPATH command LINK parameter can be found "CESTPATH – establishing paths" on page 205.

The CESTPATH command for the example in Figure 14, using primary volume 0142, is:

```
CESTPATH DEVN(X'0142') PRIM(X'6060' 62019) SEC(X'6061' 68006)
         LINK(X'00150000')
```

Example: The following example adds four ESCON paths between the primary site storage control (SSID 6060, serial number ending with 62019) and the recovery site storage control (SSID 6061, serial number ending with 68006). Issue the CESTPATH command to any device on the primary site storage control.

```
CESTPATH DEVN(X'0142') PRIM(X'6060' 62019) SEC(X'6061' 68006)
         LINK(X'0000B000' X'0001B100' X'0016B200' X'0017B300')
```

Note: The above example assumes that all of the physical and logical paths already exist.

Managing volumes and paths

This section describes how to add, delete, and suspend volume pairs, query a peer-to-peer remote copy volume, and delete a PPRC path.

Adding a PPRC volume pair

Use the CESTPAIR command to specify the PPRC primary and secondary volumes. Issue the CESTPAIR command to the primary (source) device address.

Before you issue the CESTPAIR command, you may want to first issue a CQUERY command to both simplex devices that will become the PPRC duplex pair. The

returned information includes the storage control SSID, serial number and CCA of the device. The CESTPAIR command requires this information.

Notes:

1. To change the value of a CESTPAIR command parameter for an established PPRC volume pair (change CRIT(YES) to CRIT(NO), for example), first issue a CSUSPEND command. Next, reestablish the pair with a CESTPAIR command using the RESYNC parameter, specifying all of the desired parameters.
2. If you issue a CESTPAIR command with the RESYNC option to a secondary volume that has been set to simplex with the CRECOVER command, ensure that applications have not modified that volume. If you suspect that the volume data may have changed, reestablish the pair with the CESTPAIR command COPY option. This will protect the data integrity of the two volumes.

The following command adds a PPRC volume pair:

```
CESTPAIR DEVN(X'0C40') PRIM(X'6060' 62019 X'00')
         SEC(X'6061' 68006 X'09') MODE(COPY) PACE(15) CRIT(N)
```

You can run the following JCL as a batch job, or process it from a PROCLIB to automate the procedures that are shown here. This example shows how to add four PPRC paths and a volume pair to PPRC.

```
//IKJEFT01  JOB,
//          MSGLEVEL=(1,1),MSGCLASS=A,NOTIFY=user ID
//STARTRC   EXEC PGM=IKJEFT01
//STEPLIB DD DSN=authorized.cmdlib,DISP=SHR
//SYSTSPRT DD SYSOUT=H
//SYSTSIN  DD *
  CESTPATH                                +
     DEVN(X'0C44')                        +
     PRIM(X'6060' 62019)                  +
     SEC(X'6061' 68006)                   +
     LINK(X'00008000' X'00108100' X'00018200' X'00118300')
  CESTPAIR                                +
     DEVN(X'0C40')                        +
     PRIM(X'6060' 62019 X'00')            +
     SEC(X'6061' 68006 X'09')             +
     MODE(COPY)                           +
     PACE(2)                              +
     CRIT(Y)
/*
```

Querying PPRC volumes

The CQUERY command can either query the status of one volume from a PPRC volume pair, or collect information about a volume in the simplex state. To do this, either include the VOLUME parameter, or do not specify any optional parameters. The following command queries volume 0F40:

```
CQUERY DEVN(X'0F40')
```

When you issue the CQUERY command for a volume, PPRC returns status in one of two formats, depending on whether you choose the FORMAT or UNFORMAT option. The FORMAT option, which is the default, returns status as a block message to your TSO session or batch JOBLOG, and to the SYSLOG. The UNFORMAT option allows you to save screen space by returning unformatted status information.

The path information that is available for the storage control depends on the current status of the volume to which the CQUERY command is directed, as follows:

- If the device is in **simplex** state, PPRC may display path information for the storage control for the device. The path information does not relate to the specified device.
- If the device is a **primary** of a PPRC volume pair, PPRC displays path information for the storage control. The path information does relate to the specified device. The devices that make up the volume pair are using one or more paths.
- If the device is a **secondary** of a PPRC volume pair, PPRC provides a report, but valid path information is not available. When you issue a CQUERY PATHS command to an active PPRC pair secondary volume, the formatted CQUERY report will not have serial number, SSID, or path link information displayed.

If you need path information for a storage control that has a secondary PPRC volume, issue the CQUERY PATH command. You can specify another device on that storage control that is in simplex state. You can also specify a device on that storage control that is a primary device in a PPRC pair.

Note: CQUERY command output in the SYSLOG can contain a connection code at the end of the header message and at the beginning of each of the lines associated with that connection code. The system generates this code for any multiline WTO.

The following examples show both formatted and unformatted output.

CQUERY formatted output for a primary volume

The following is an example for a CQUERY FORMAT command for a primary volume.

```
15.12.00 TSU00189 ANTP0090I CQUERY FORMATTED LVL 2           C
VOLUME REPORT
********** PPRC REMOTE COPY CQUERY - VOLUME *****************
*                                    (PRIMARY)   (SECONDARY) *
*                                     SSID CCA    SSID CCA   *
*DEVICE   LEVEL   STATE    PATH STATUS  SERIAL#   SERIAL#    *
*------ -------  --------  ----------- ---------  ---------  *
* 0F40  PRIMARY.. DUPLEX..  ACTIVE..    0240 03    8080 01   *
*       CRIT(NO).......    CGRPLB(NO). 0000090383 00000090461*
* PATHS SAID/DEST STATUS: DESCRIPTION                        *
* ----- --------- ------  -------------------                *
*   1   0011 0000   01    PATH ESTABLISHED...                *
*   2   0012 0000   01    PATH ESTABLISHED...                *
*   3   0013 0000   01    PATH ESTABLISHED...                *
*   4   0014 0000   01    PATH ESTABLISHED...                *
*   5   0015 0000   01    PATH ESTABLISHED...                *
*   6   0016 0000   01    PATH ESTABLISHED...                *
*   7   ---- ----   00    NO PATH.............               *
*   8   ---- ----   00    NO PATH.............               *
************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 0F40.
 COMPLETION CODE: 00
```

CQUERY formatted output for a secondary volume

The following is an example of output for a CQUERY FORMAT command for a suspended secondary volume.

```
15.12.00 TSU00189 ANTP0090I CQUERY FORMATTED LVL 2                C
VOLUME REPORT
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                           SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 0F42  SECONDARY  SUSPEND(5)   ACTIVE..    6060 08      6061 10   *
*       .........                          ..... 000000090383*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH.....                             *
*       ---- ----   00    NO PATH.....                             *
*       ---- ----   00    NO PATH.....                             *
*       ---- ----   00    NO PATH.....                             *
* SECONDARY WAS SUSPENDED (YMD/UTC) 1999-07-20 19.49.12.242451     *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 0F42. COMPLETION CODE: 00
```

CQUERY unformatted output for a primary volume

The following is an example of output for a CQUERY UNFORMAT command for a primary volume. See message ANTP0091I for details about the reported fields.

```
15.12.00 TSU00189 ANTP0091I CQUERY UNFORMATTED LVL 2        C
VOLUME REPORT
0F40,PRIMARY,DUPLEX,ACTIVE,
0240,03,000000090383,8080,01,000000090461,N,N,
6,00110000,01,00120000,01,00130000,01,00140000,01,
00150000,01,00160000,01,,,,,
,,,
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 0F40.
 COMPLETION CODE: 00
```

CQUERY unformatted output for a suspended secondary volume

The following is an example for a CQUERY UNFORMAT command for a suspended secondary volume as displayed on the master console. See message ANTP0091I for details about the reported fields.

```
15.12.00 TSU00189 ANTP0091I CQUERY UNFORMATTED LVL 2     C
VOLUME REPORT
 0F60,SECONDARY,SUSPEND(5),ACTIVE,
 6061,12,,6060,12,000000090462,,
 ,,,,,,,,,
 ,,,1997-04-04 20.31.05.791920
```

The preceding examples of formatted output contain the following terms:

DEVICE Displays the 4-digit hexadecimal device address.

LEVEL Displays whether the volume is a primary or secondary volume. For a primary volume, it also displays the status of the volume as defined by the CRIT parameter of the CESTPAIR command that established the pair.

If neither PRIMARY nor SECONDARY is displayed, the volume is not part of a PPRC pair.

CRIT(NO) indicates that you either issued the CESTPAIR command without the CRIT parameter, or specified CRIT(NO).

CRIT(YES-PATHS) indicates that the volume pair only be write-inhibited when there are no PPRC paths available. CRIT(YES-ALL) indicates that the primary storage control has been configured to suspend the pair and prevent I/O to the primary volume as a result of any error, including an error on the primary or secondary volume. CRIT(YES-PATHS) and CRIT(YES-ALL) both indicate that the CESTPAIR command included the CRIT(YES) parameter.

STATE Indicates whether the PPRC volume is in simplex, pending, duplex, or suspended (suspend(*n*)) state. "Identifying peer-to-peer remote copy volume states" on page 216 describes these states. The value *n* for SUSPEND is one of the following:

- (3) — The host has suspended the pair, and does not permit writes.
- (4) — The host has suspended the pair, and does permit writes. The primary site storage control will keep these updates.
- (5) — A command issued by the primary site storage control to the recovery site storage control has suspended the pair. Only a recovery site storage control returns value 5.
- (6) — Internal conditions within either storage control have suspended the pair. Either storage control returns value 6.
- (7) — The transition of the recovery volume to the simplex state has suspended the pair. Only a primary site storage control returns value 7.
- (8) — The primary site storage control has suspended the pair as a result of abnormal conditions within the recovery site storage subsystem. These conditions may involve the storage control, its attached disk devices, and the ESCON paths between the two sites. The primary site storage control detects and reports this condition.
- (9) — Either an IPL, or a power interruption to either storage control, have suspended the pair. The storage control that received the interruption returns value 9.
- (A) — A CGROUP command with the FREEZE parameter has suspended the pair. You must query a primary device to receive value A.

PATH STATUS
Displays whether the status is active or inactive.

SSID CCA Serial # PRIMARY
Displays the four-character subsystem ID (SSID) of the primary site storage control, the two-character CCA of the primary volume, and, when available, the storage control serial number (up to 12 characters).

SSID CCA Serial # SECONDARY
Displays the four-character subsystem ID (SSID) of the recovery site storage control, the two-character CCA of the primary volume, and the storage control serial number (up to 12 characters).

CRIT Displays the status of the CRIT keyword of the CESTPAIR

command. See “CESTPAIR – establishing pairs” on page 202 for an explanation of the CRIT parameter.

CGRPLB Indicates whether or not the storage subsystem is configured to support PPRC consistency grouping.

PATHS Displays the number of valid links, if there are any, between primary site and recovery site storage controls.

SAID Displays the four-character system adapter ID (SAID) of the primary site storage control. SAID and DEST (described below) make up the link address.

DEST Specifies up to four 4-digit, hexadecimal-character link addresses. The link addresses each consist of a 2-digit, hexadecimal-character ESCON port address followed by 00.

STATUS Displays a 2-digit number that indicates the detailed status of the path. See status descriptions under “Querying PPRC paths” on page 227.

DESCRIPTION

Provides a short description of the status indicator:

- 00 — No path
- 01 — Established path
- 02 — Either initialization or reinitialization has failed, or the path is inoperative
- 03 — Time out
- 04 — No resources available at primary site storage control
- 05 — No resources available at recovery site storage control
- 06 — Serial number mismatch
- 10 — Configuration error
- nn — Status is undetermined

FIRST CYL OUT OF SYNCH

This information, when available, is presented only for volumes in pending and suspended states.

If the volume is a primary volume, this is the lowest cylinder number that is out of synchronization on the primary volume. For a secondary volume, this is the lower cylinder number that was received from the primary. This update is in cache, as PPRC has not yet written this update to the secondary volume.

LAST CYL OUT OF SYNCH

This information, when available, is presented only for volumes in pending and suspended states.

If the volume is a primary volume, this is the highest cylinder number out of synchronization on the primary volume. For a secondary volume, this is the higher cylinder number that was received from the primary. This update is in cache, as PPRC has not yet written this update to the secondary volume.

PERCENT OF COPY COMPLETED

Displays the percent completed for the copy operation. This is an optional line of information that PPRC presents only if the volume is a primary volume and the information is available.

SECONDARY WAS SUSPENDED

When it is available, displays the Universal Time Coordinated (UTC) to indicate when this secondary volume became suspended. The

UTC, in the ISO format “YYYY-MM-DD HH.MM.SS.NNNNNN”, is the time-of-day clock value from the host system from which the CSUSPEND command was issued.

Querying PPRC paths

Issue the CQUERY command with the PATHS option to query the status of all PPRC paths established for a volume’s storage control.

When you issue the CQUERY command with the PATHS option, PPRC returns status in one of two formats, depending on whether you choose the FORMAT or UNFORMAT option.

- The FORMAT option, which is the default, returns status as a block message to your TSO session or batch JOBLOG, and to the SYSLOG.
- The UNFORMAT option allows you to save screen space by returning unformatted status information.

Note: CQUERY command output in the SYSLOG can contain a connection code at the end of the header message and at the beginning of each of the lines associated with that connection code. The system generates this code for any multiline WTO.

The following command queries paths established for volume 0F40:

```
CQUERY DEVN(X'0F40') PATHS
```

CQUERY unformatted output with the PATHS option

The following is an example of printed output for a CQUERY PATHS UNFORMAT command for a volume on a primary storage control. The volume has at least one path established to four secondary storage controls. Refer to message ANTP0096I for details about the reported fields.

```
15.12.00 TSU00178 ANTP0096I CQUERY UNFORMATTED LVL 2       C
PATHS REPORT
000000090383,0240,0000,
000000090461,8081,5,00110000,01,001200,01,001300,01,001400,01,
00150000,01,,,,,
000000090462,8082,1,00120100,01,,,,,,,
000000090463,8083,1,00130200,01,,,,,,,
000000090464,8084,1,00140300,01,,,,,,,
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 0F40.
COMPLETION CODE: 00
```

CQUERY formatted output with the PATHS option

The following example of CQUERY PATHS printed output displays data for a storage control that supports eight paths. Path information for paths 5-8 would not be displayed if fewer than five paths were active, or if the storage control supported only four paths.

```
15.12.00 TSU00178 ANTP0095I CQUERY FORMATTED LVL 2         C
PATHS REPORT
*************** PPRC REMOTE COPY CQUERY - PATHS **************
*PRIMARY UNIT: SERIAL#=000000090383 SSID=0240 SS=0000        *
*             FIRST      SECOND       THIRD         FOURTH   *
*           SECONDARY   SECONDARY   SECONDARY     SECONDARY  *
*SERIAL NO: 0000090461  0000090462  0000090463  000000090464 *
*     SSID:   8081        8082        8083          8084     *
*    PATHS:    6           1           1             1       *
*     SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*     --------- --  --------- --  --------- --  --------- -- *
*  1: 0011 0000 01  0012 0100 01  0013 0200 01  0014 0300 01 *
*  2: 0012 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  3: 0013 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  4: 0014 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  5: 0015 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  6: 0016 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  7: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*  8: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*                                                            *
* S* = PATH STATUS:                                         *
* 00=NO PATH          01=ESTABLISHED         02=INIT FAILED  *
* 03=TIME OUT   04=NO RESOURCES AT PRI 05=NO RESOURCES AT SEC*
* 06=SERIAL# MISMATCH 07=RESERVED      08=RESERVED           *
* 09=RESERVED         10=CONFIGURATION ERROR                *
************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 0F40.
COMPLETION CODE: 00
```

The preceding display contains the following information for the primary site storage control:

- The primary site storage control serial number (up to 12 characters)
- The four-character storage subsystem ID (SSID)

The display contains information for up to four recovery site (SECONDARY) storage controls. Under each storage control it lists:

- The serial number (up to 12 characters)
- The four-character storage subsystem ID (SSID)
- The number of paths established for that storage control
- Specific information for each path, with the SAID, followed by the link address (DEST), followed by the status of that path (S*)

If the CQUERY PATHS output is for a PPRC-capable RVA, the printed output includes the following path status information at the bottom of the printout:

```
*
* S* = PATH STATUS:                                         *
* 00=NO PATH    01=ESTABLISHED         02=INIT FAILED       *
* 03=TIME OUT   04=NO RESOURCES AT PRI 05=NO RESOURCES AT SEC*
* 06=SERIAL# MISMTCH 07=SEC SSID MISMTCH 08=ESCON LINK OFFLN *
* 09=ESTAB FAILED 0A=PATH ACTV TO HOST 0B=PATH TO SAME CLUSTR*
* 10=CONFIGURATION ERROR                                    *
************************************************************
```

The path status codes are defined as follows. **Path codes noted as "LSS only" are valid only for storage subsystems that support logical storage subsystems.**

00 — NO PATH The request to establish this path is in process, or you have not requested this entry. (The SAID/DEST

appears as "----" until the CESTPATH command operation has completed. Reissue the CQUERY command with the PATHS option.)

01 — ESTABLISHED Path is operational.

02 — INIT FAILED An attempt to establish this path has failed. Verify that the physical link, including ESCD ports, is correct for the connection you are attempting. Verify the integrity of the fiber-optic links between the requested source and the destination.

03 — TIME OUT An attempt to establish this path has failed. Verify the physical link at the primary storage control. Determine if non-PPRC system activity has degraded link response times during the establish time. Retry the CESTPATH command.

04 — NO RESOURCES AT PRI
An attempt to establish this path has failed. Verify that the primary storage control does not already have the maximum number of logical paths already established. Verify the configuration and remove any unnecessary channel paths. When the problem is resolved, issue a CQUERY command with the PATHS option to verify that the path has been automatically established.

05 — NO RESOURCES AT SEC
An attempt to establish this path has failed. Verify that the recovery site storage control does not already have the maximum number of logical paths already established. Verify the configuration and remove any unnecessary channel paths. When the problem is resolved, issue a CQUERY command with the PATHS option to verify that the path has been automatically established.

06 — SERIAL # MISMATCH The recovery site storage control specified by the SAID/DEST does not have the same serial number specified in the CESTPATH command. Verify the SAID/DEST port and the serial number parameters. A CQUERY command with the PATHS option from the host to a secondary volume can provide the serial number information.

07 — SEC SSID MISMATCH (LSS only) An attempt to establish this path has failed. Verify that the path is available and is properly defined, and then retry the CESTPATH command.

08 — ESCON LINK OFFLN (LSS only) An attempt to establish this path has failed. Retry the CESTPATH command after you have made a path available.

09 — ESTABLISH FAILED (LSS only) An attempt to establish this path has failed. If the automatic attempts to reconnect are not successful, reissue the CESTPATH command.

0A — PATH ACTIVE TO HOST
(LSS only) An attempt to establish this path has

failed. The system adapter has already established a link to the system server on this path. Retry the CESTPATH command, specifying another path.

0B — PATH TO SAME CLUSTR

(LSS only) An attempt to establish this path has failed because there is already an established path from this primary cluster to this secondary cluster.

10 — CONFIGURATION ERROR

The SAID is not valid for the primary site storage control ESCON SA card, or a PPRC path has already been established on this logical link. Verify the SAID is appropriate for the number of physical ESCON ports installed on the primary storage control SA card. Also, determine if this is already a PPRC logical path, possibly established in the other direction. To do this, issue a CQUERY command with the PATHS option to the indicated storage controls at both sites.

Example: CLIST to add a PPRC path, add a volume pair, and query PPRC

You can process the following TSO CLIST in a standard TSO environment. This example shows how to use a CLIST to add a PPRC path, add a volume pair to PPRC, and then query the status of PPRC.

```
CONTROL PROMPT
CESTPATH                              +
   DEVN(X'0143')                      +
   PRIM(X'6060' 62019)                +
   SEC(X'6061' 68006)                 +
   LINK(X'00008200')
CESTPAIR                              +
   DEVN(X'0143')                      +
   PRIM(X'6060' 62019 X'03')          +
   SEC(X'6061' 68006 X'02')           +
   MODE(COPY)                         +
   PACE(2)                            +
   CRIT(N)
CQUERY DEVN(X'0143')
```

Suspending PPRC volume pair operations

Issue the CSUSPEND command to suspend PPRC operations between the primary and secondary of a volume pair. See “CSUSPEND – suspending pairs” on page 211 for command specifics.

Attention:

APAR OW15247, APAR OW15248, or APAR OW40367 has disabled the QUIESCE option of the CSUSPEND command. Refer to either of these APARs for instructions about enabling the QUIESCE option of the CSUSPEND command.

If the QUIESCE option is enabled, do not issue a CSUSPEND command with the QUIESCE option to a volume with active SPOOL, PAGE, or CATALOG data sets, or to an active SYSRES volume. Doing so can cause any attached system to enter a deadlock condition, and may require a storage control IML to correct the condition.

The following QUIESCE option description is only relevant if you have enabled this feature.

Use the QUIESCE option, when enabled, to stop all updates to secondary volumes at a consistent point in time without having to stop the applications.

Example: If your applications use 100 volumes, issue the CSUSPEND command with the QUIESCE option to all pairs. This will bring all 100 secondary volumes to a consistent state.

The QUIESCE option causes MVS to queue all application read and write I/O operations. After these CSUSPEND commands have completed, follow them with another sequence of CSUSPEND commands, without the QUIESCE option, to the primary volumes. Doing so allows the primary applications to resume their read and write I/O to these volumes. It also directs the primary storage subsystem to keep a record of cylinders that change on the primary volumes.

You can later reestablish the PPRC volume pairs. Issue CESTPAIR commands with the RESYNC option to copy changed cylinders to the secondary volumes. Alternatively, you could break up the suspended pairs with the CDELPAIR command if you wish to use the secondary volumes for another purpose.

Example: The following command suspends PPRC operations between a volume pair. You must issue the CSUSPEND command to a valid primary or secondary PPRC device.

```
CSUSPEND DEVN(X'0F40') PRIM(X'6060' 62019 X'00')
         SEC(X'6061' 68006 X'09')
```

Issue a CESTPAIR command with the RESYNC option to resume PPRC operations for the volume pair.

Note: When you reestablish a suspended pair that was originally established with CRIT(YES) specified, you must again specify CRIT(YES) with the CESTPAIR command if you want this option to remain in effect.

Deleting a PPRC volume pair

Issue the CDELPAIR command to delete PPRC primary and secondary volumes. Direct the command to the primary device.

Before issuing a CDELPAIR command, verify that the necessary paths are active between the primary and secondary volume pair. If all of the paths that were established between both of the storage controls have been disabled, the CDELPAIR command will fail. PPRC may issue the following two messages:

```
ANTP0122I PRIMARY CU TIMED OUT ATTEMPTING TO COMMUNICATE
          WITH SECONDARY
ANTP0001I CDELPAIR COMMAND UNSUCCESSFUL FOR DEVICE nnnn.
          COMPLETION CODE: 08
```

After the CDELPAIR command times out, the state of the primary volume is simplex. The secondary volume remains in its previous state (duplex, pending, or suspended). You then need to issue a CRECOVER command to return the secondary volume to the simplex state.

The following command deletes a PPRC volume pair:

```
CDELPAIR DEVN(X'0F40') PRIM(X'6060' 62019 X'00')
         SEC(X'6061' 68006 X'09')
```

You can process the following TSO CLIST in a standard TSO environment.

Example: This example shows how to use a CLIST to delete a volume pair.

```
CONTROL PROMPT
CDELPAIR                              +
   DEVN(X'04C6')                      +
   PRIM(X'6060' 62019 X'06')          +
   SEC(X'6061' 68006 X'09')
DATA PROMPT
YES
ENDDATA
```

Deleting PPRC paths

Use the CDELPATH command to delete all established ESCON paths between primary site storage controls and recovery site storage controls. The command only affects active paths to the recovery site storage control, but not other paths to other storage controls. Issue the CDELPATH command to a valid PPRC primary device, only.

Before issuing a CDELPATH command, issue a CDELPAIR command to all active PPRC volume pairs. The CDELPATH command may cause an ANTP0121I message if you do not follow this sequence.

The following command deletes all established ESCON paths between primary site storage control *6060* and recovery site storage control *6061*.

```
CDELPATH DEVN(X'0F40') PRIM(X'6060' 62019) SEC(X'6061' 68006)
```

You can process the following TSO CLIST in a standard TSO environment.

Example: This example shows how to use a CLIST to delete all PPRC paths between two storage controls.

```
CONTROL PROMPT
CDELPATH                      +
   DEVN(X'0F40')              +
   PRIM(X'6060' 62019)        +
   SEC(X'6061' 68006)
DATA PROMPT
YES
ENDDATA
```

You can use the CESTPATH command to remove only selected paths. See the note under “CESTPATH – establishing paths” on page 205 for this information.

Chapter 15. PPRC Data Recovery Operations

A major purpose of peer-to-peer remote copy is to enable you to switch to a recovery system in the event of a disaster to the primary system. This section presents an overview of disaster recovery from a system viewpoint and describes the procedure to direct the recovery system to take over for the primary system.

PPRC is also an excellent solution for moving data between disk devices. You might want to have a data migration procedure for planned activity, rather than for a disaster. A data migration procedure is a subset of the recovery procedure that is described below. For information about using PPRC for data migration, see "Chapter 16. Moving and Migrating Data with PPRC" on page 239.

In this chapter

This chapter provides information about data recovery with PPRC and includes the following sections:

Topic . . .	See page . . .
Examining disaster recovery from a system viewpoint	233
Resuming operations at the recovery site	234
Managing errors during recovery	236

Examining disaster recovery from a system viewpoint

Disasters occur in many ways. In a sudden disaster, all volumes may be in normal full duplex mode when you begin recovery procedures. In such a case it is easy to determine how current the data is.

A disaster can also occur over a period of time. This rolling disaster is the most common situation, and necessitates that you check each volume to determine how current the data is. SYSLOG messages written by the ERP give the times when specific volumes are suspended.

It is important to understand how a rolling disaster can affect your primary system. This knowledge enables you to avoid a potential data integrity exposure when using PPRC in a disaster recovery situation. A disaster or series of disasters can result in intermittent failures to the primary system. In this situation, you may be unable to determine which volumes are recoverable because of one of the following scenarios:

- A disaster could leave many or all of your secondary volumes in a suspended state. Without the IEA49xx message to indicate when these volumes became suspended, each volume would require extensive analysis to determine how current the data is.
- If the attempt to update the secondary volume fails due to a link failure, PPRC may be unable to mark the secondary volume as suspended.
- The operating systems process all error notification actions asynchronously. Positive identification and recording of an event, such as broken PPRC ESCON links, happen after the event itself.

During a disaster recovery operation, the secondary volume would *appear* to be recoverable. There would, however, be no way to be certain that all of the data is

© Copyright IBM Corp. 1994, 2002

current. The way to minimize (or even avoid) this potential data integrity exposure is by applying the disaster recovery preparations that are described in “Considering the PPRC solution” on page 184 and “Preparing for PPRC error recovery” on page 244.

Resuming operations at the recovery site

Perform the following steps to have the recovery site assume control of operations:

1. Reconfigure ESCON Directors, if necessary, to connect local systems to recovery system Disk.

2. Query each secondary disk volume to verify that it was a secondary remote copy volume and that it is in duplex state. Continue to the steps below for volume states other than duplex state.

3. If the volume is in duplex or suspended state, check the console for SYSLOG error messages that indicate that the ERP has attempted to put the volume in suspended state. If there is an error message, recover the volume as if it was in the suspended state at the time that the IEA49xx message indicates.

4. Issue the PPRC CRECOVER command to change the secondary volume from its current state (duplex, suspended, or pending) to simplex state in order to let that volume come online. Table 48 shows suggested peer-to-peer remote copy volume recovery procedures that are based on the state of the volume at the time of system failure.

   Rules: The PPRC recovery function must have access to each secondary volume on all recovery site storage controls at the time of the recovery.

   Some recovery system volumes may not initially be usable for recovery. Change these volume serial numbers to prevent any applications from accessing the volumes until you can either restore them, or otherwise validate their contents by recovery procedures.

Table 48. Volume failure states and recovery actions

If volume is in . . .	Then do this . . .

Duplex state	Issue the CRECOVER command, as in the following example:
	`CRECOVER DEVN(X'0F40') PRIM(X'6060' 62019 X'00')`
	`SEC(X'6061' 68006 X'09') ID(DATA01)`
	This returns the volume to simplex state, fully recovered and ready to use. (Issue the CQUERY command again if you wish to ensure that the volume is in simplex state.) You can now vary the volume (DATA01 in the above example) online for system use. Review the steps above.

Table 48. Volume failure states and recovery actions (continued)

If volume is in . . .	Then do this . . .

Pending state	If the volume is in pending state, then it never reached duplex state. Issue the CRECOVER command, as in the following example:
	`CRECOVER DEVN(X'0F40') PRIM(X'6060' 62019 X'00')`
	`SEC(X'6061' 68006 X'09') ID(DATA01 DATA02)`
	This returns the volume to simplex state and changes the volume serial number from its old value to a new value. (Issue the CQUERY command again if you wish to ensure that the volume is in simplex state.) You can now vary the volume (DATA02 in the above example) online to do normal recovery from backup tapes. Once the volume state is fully recovered, the volume can be changed back to its original volume serial number so that applications can allocate and use the volume.

Simplex state	Determine why the volume is in simplex state before you accept the secondary copy as current and usable. The simplex volume may either not be a remote copy volume, or the volume may be in simplex state following a previous action that ended the volume copy pair. When you are satisfied with the volume status, vary the volume online for system use.

Suspended state	Determine what caused this state to occur. If the volume was suspended by the host or primary site storage control, data on the volume is valid. How current the data is depends on the timestamp information provided in the SYSLOG IEA49xx message. If the volume was suspended by the recovery site storage control, data on the secondary volume may be downlevel or incorrect for various reasons (i.e., an equipment check). In this case, recover from the backup data volumes.
	If the data on the volume is valid, issue the CRECOVER command, as in the following example:
	`CRECOVER DEVN(X'0F40') PRIM(X'6060' 62019 X'00')`
	`SEC(X'6061' 68006 X'09') ID(DATA01 DATA02)`
	This returns the volume to simplex state and changes the volume serial number from its old value to a new value. (Issue the CQUERY command again if you wish to ensure that the volume is in simplex state.)
	With the volume (DATA02 in the above example) now in simplex state, vary the volume online and use the system log, database logs, and backup data volumes (when appropriate) to complete the recovery of data from the time the volume went into suspended state until the time of failure. Once the volume state is fully recovered, the volume can be changed back to its original volume serial number so that applications can allocate and use the volume.

5. Restart the application systems and perform the same system and application startup procedures that you perform on the primary system when you restart there following a system failure.

   Recovery system catalog entries will be consistent with those on the primary system if catalog volumes are part of the data that was copied to the recovery

system. See “Copying the catalog and control data sets” on page 47 for guidelines on managing catalog updates that are not copied to the recovery system.

If the catalog change rate is very low, it may be more effective to copy the catalog with remote copy. Consider what other data sets are on the volumes with the catalogs, as PPRC will also copy these data sets to the recovery system.

Managing errors during recovery

This section describes errors that may occur when recovering data on the PPRC system, and the actions to take to resolve them.

Errors encountered when writing to PPRC volumes established with CRIT(YES)

An error while writing to a primary or secondary PPRC volume that was established with CRIT(YES) causes the I/O to be unit checked if there are no I/O paths available. When this occurs, there are certain error conditions where the primary is updated, but the secondary is not. You must take whatever actions are necessary to resynchronize the volumes, or recover from backup volumes.

Errors that occur in the absence of write activity

An error can occur on a primary volume that does not have any write activity at the time of the error. In this case, the volume remains in duplex state until a write occurs. To get out of this situation, send some I/O activity to the error volume, which puts it in suspended state. Then issue a CESTPAIR MODE(RESYNC) command to recover the volume.

Errors that cause a PPRC pair to become suspended

A PPRC pair may become suspended for any of the following reasons:

- Primary device write failure
- Secondary device write failure
- Secondary subsystem failure
- Communications to secondary failure
- Secondary not ready, intervention required

Either the I/O currently in process or the next I/O to the primary device of that PPRC pair is unit checked. The I/O that triggers the error may come from any attached system; therefore, the IEA49xx message can appear on any attached system’s console. An operating system issues an IEA49xx console message, indicating one of the previously listed suspension reasons, based on how the system detected this failure. It does not necessarily relate to the true cause of the failure.

If you originally established the suspended PPRC pair with CRIT(YES), the system unit checks the device when the Disk ERP reissues the I/O order. You then receive an IOS message similar to the following:

```
IOS000I 164,5D,CMD,47,0E00,00160000,VOLXXX,JOBNAME,
800000005400000F57 .....

  Sense Byte 0 '80' - command reject
  Sense Bytes 7-8 '0F57' - Device 164 in Critical State
```

An application-related error message may follow the above message.

Chapter 16. Moving and Migrating Data with PPRC

In addition to providing data recovery capability in the event of a disaster (see “Chapter 15. PPRC Data Recovery Operations” on page 233), PPRC is an efficient tool for transferring data between various disk device types. The following are specific data transfer tasks:

- Temporarily moving data to a new location, to accommodate a repair, for instance
- Migrating data to new devices
- Making a copy of the data at any point in time

This section includes procedures for all three tasks.

When you move, migrate, or copy data with PPRC, you can put the secondary (target) disk volume on the same or on a different subsystem as the primary. Ensure that each volume is attached to a storage system with the PPRC LIC installed, and neither device is part of a dual copy pair.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Steps for moving the work load with PPRC	239
Steps for migrating data with PPRC	240
Steps for copying the secondary volumes with PPRC	241

Steps for moving the work load with PPRC

Before you begin: You need to be familiar with the commands used in this procedure. See “Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199 for PPRC command information.

Perform the following steps to temporarily move data with PPRC:

1. Vary the secondary device offline to all MVS systems.

   _____

2. Issue the CESTPATH command and the CESTPAIR command with the COPY option to establish the PPRC paths and pairs.

   _____

3. Issue the CQUERY TSO command to monitor the initial copy process. (The CQUERY command returns information about the state of the PPRC volume, which changes from simplex to pending, to duplex.) During the initial copy operation, PPRC also copies all primary volume updates to the secondary (target) volume.

   _____

4. Verify that the two PPRC volumes are in the duplex state by issuing the CQUERY command.

   _____

5. Switch to the secondary (target) volume when the copy operation is complete. To switch your application from the primary volume to the secondary volume, do the following:

© Copyright IBM Corp. 1994, 2002

a. Stop all applications that are using the primary volume. PPRC is a synchronous copy; therefore the data on the primary volume is the same as the data on the secondary volume.
b. Issue the CDELPAIR command to return the PPRC primary and secondary volumes to simplex state. Optionally, you can issue a CDELPATH command if you are finished with the PPRC paths between the two storage controls.
c. If you are moving the work loads with the intention of later copying them back, ensure that both devices have the same capacity.
d. If the primary and secondary volumes are under the same MVS system, vary the original primary volume offline to MVS and the original secondary volume online to MVS. Both volumes have the same volume serial number, and only one can be online to MVS at a time.
e. Establish the volumes in the reverse direction by issuing a CESTPATH command, then a CESTPAIR command with the NOCOPY option. Issue the CSUSPEND command to suspend the pair and keep a record of changed tracks.
f. Start all applications with the new disk volumes.
g. If you are performing this procedure to enable a device repair, make necessary repairs at this time, and verify that all data is intact on the device.
h. When the secondary disk is again available and the data is intact, issue the CESTPAIR command with the RESYNC option to copy the changed data. If the data is suspect, establish the pair with the COPY option.
i. Repeat these procedures when you wish to move the work load back to the original location.

Steps for migrating data with PPRC

Before you begin: You need to be familiar with the commands used in this procedure. See “Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199 for PPRC command information.

Perform the following steps to migrate data with PPRC:

1. Vary the secondary device offline to all MVS systems.

   ---

2. Issue the CESTPATH and CESTPAIR commands with the COPY options to establish the PPRC paths and pairs.

   ---

3. Issue the CQUERY TSO command to monitor the initial copy process. (The CQUERY command returns information about the state of the PPRC volume, which changes from simplex to pending, to duplex.) During the initial copy operation, PPRC also copies all primary volume updates to the secondary (target) volume.

   ---

4. Verify that the two PPRC volumes are in the duplex state by issuing the CQUERY command.

   ---

5. Switch to the secondary (target) volume when the copy operation is complete. To switch your application from the primary volume to the secondary volume, do the following:
   a. Stop all applications that are using the primary volume. PPRC is a synchronous copy; therefore the data on the primary volume is the same as the data on the secondary volume.

b. Issue the CDELPAIR command to return the PPRC primary and secondary volumes to simplex state. Optionally, you can issue a CDELPATH command if you are finished with the PPRC paths between the two storage controls.
c. If the PPRC secondary volume is a disk type containing more cylinders than the primary PPRC volume had, run ICKDSF Release 16 or above against the original PPRC secondary volume to expand the VTOC to reflect the additional space.
d. If the primary and secondary volumes are under the same MVS system, vary the original primary volume offline to MVS and the original secondary volume online. Both volumes have the same volume serial number, and only one can be online to MVS at any time.
e. Restart all applications, which now use the new disk volume.

Steps for copying the secondary volumes with PPRC

Before you begin: You need to be familiar with the commands used in this procedure. See "Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions" on page 199 for PPRC command information.

Perform the following steps to make a copy of the secondary volumes (like volumes only) with PPRC:

1. Vary the secondary devices offline to all MVS systems.

---

2. Issue the CESTPATH and CESTPAIR commands with the COPY options to establish the PPRC paths and pairs.

---

3. Issue the CQUERY TSO command to monitor the initial copy process. (The CQUERY command returns information about the state of the PPRC volumes, which change from simplex to pending, to duplex.) During the initial copy operation, PPRC also copies all primary volume updates to the secondary (target) volumes.

---

4. Verify that the PPRC volumes are in the duplex state by issuing CQUERY commands.

---

5. When the copy operation is complete, generate a dump of the secondary (target) volumes as follows:
   a. Temporarily stop all of the applications that PPRC is copying.

      Note: CSUSPEND QUIESCE Option

      If you have reenabled the QUIESCE option (which APAR OW15247, APAR OW15248, or APAR OW40367 has disabled), you have an alternative to stopping the applications. Issue the CSUSPEND command with the QUIESCE option to the primary devices to synchronize all secondary volumes.

      Attention: If you choose to use the QUIESCE parameter, you must first enable it. For instructions and cautions about enabling QUIESCE, refer to "CSUSPEND – suspending pairs" on page 211.
   b. Issue the CSUSPEND command to the secondary volumes. PPRC now records changes to the primary volumes.

c. Issue the CRECOVER command from the secondary host to return the secondary volumes to simplex state. Specify new volume serial numbers for the secondary volumes if the primary and secondary volumes are under the same MVS system.
d. Vary the secondary volumes online. After the CRECOVER command completes, the secondary volumes will be in simplex state. It will be possible for applications to make updates to these volumes when they are varied online. *It is essential that no other applications write to the secondary volumes after the CRECOVER command completes.*
e. Use DFSMSdss physical full volume dump to dump the secondary volumes.
f. When the dumps are complete, vary the secondary volumes offline.
g. Change the secondary volume serial numbers back to their original numbers if you changed them in 5c.
h. Issue the CESTPAIR command with the RESYNC option to reestablish the PPRC pairs and resynchronize the volumes.

Note: It is essential that no other applications write to the secondary volumes that DFSMSdss is dumping. If updates are made to the secondary volumes while they are in simplex state (after they are recovered), you must reestablish the pairs and resynchronize the volumes. To do this, first issue the CDELPAIR command to the primary volumes, and then issue the CESTPAIR command with the COPY option.

If you changed the volume serial number of the secondary volume when you issued the CRECOVER command, you must restore the DFSMSdss physical full volume dump of that secondary volume. First, issue the DFSMSdss RESTORE command with the FULL and COPYVOLID keywords. Next, change the volume serial number back to the original volume serial number that the secondary volume had before you issued the CRECOVER command.

Chapter 17. Recovering from PPRC Error Conditions

This chapter describes various peer-to-peer remote copy error conditions and how to recover from them. You can usually diagnose and correct these errors independently from other PPRC session activity.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Understanding host error recovery procedures	243
Putting PPRC error recovery procedures into effect	244
RAMAC device sparing with PPRC	249

Understanding host error recovery procedures

There can be times when PPRC cannot replicate an update to a primary device onto the secondary device. In these cases, the primary device sends channel end, device end, and unit check status to the host that encountered the error. With the PPRC error recovery procedures (ERP) PTF that is installed, the host system performs the following actions:

1. The system stops all application I/O, *for that host*, to the primary volume in error. The host system also prevents other system images from accessing this volume.

2. The system writes an IEA49*xx* message to the log to indicate that the PPRC volume is now in suspended state or duplex state. (Ensure that the systems' log is common to both the primary system and the recovery system.)

3. The system puts information that is related to the specific failure into the SYS1.LOGREC data set for service personnel reference. SYS1.LOGREC can be common to both the primary system and the recovery system.

4. The system waits for enough time to let the IEA49*xx* message reach the recovery system. You can introduce an automation routine here that would receive control before the I/O operation can complete.

5. The system resumes all host application I/O to the primary volume. If you specified the CRIT(NO) parameter when you established the primary volume, the storage control allows subsequent I/O operations to continue without being unit checked. If you specified the CRIT(YES) parameter on the CESTPAIR command for the primary volume, then the storage control does the following:
   - Checks all subsequent write I/O operations.
   - Allows the host application to follow its own recovery actions.

   Note: An "A" in the unformatted CQUERY message ANTP0091I indicates that the volume pair is designated as CRIT(YES-ALL). This means that the

© Copyright IBM Corp. 1994, 2002

pair was established with the CRIT(YES) option. The storage control maintenance panel is set to inhibit writes on any failure, which includes secondary device failures.

Recommendation: Specify CRIT(YES-ALL) for logging-type volumes.

Preparing for PPRC error recovery

"Understanding host error recovery procedures" on page 243 outlines the actions that the host ERP starts when a PPRC duplex pair fails. It is possible that PPRC links fail such that the secondary device's storage control does not receive notices that the duplex volume is in suspended state. In this case, the primary host ERP issues an IEA49*xx* message. The host ERP action waits until the message is issued before allowing the primary application to continue. This delay gives an automation product such as NetView time to locate the IEA49*xx* message and automatically take appropriate actions.

In the event of a disaster, these actions provide data that is necessary to analyze the PPRC secondary volume to determine how current the data is. Refer to "Putting PPRC error recovery procedures into effect" to decide which recovery actions to take. These recovery actions can include the following:

- Locate the IEA49*xx* message via NetView and copy the message to the recovery system, thereby informing the recovery system that the PPRC duplex pair is in suspended state. The message contains the timestamp that indicates the time that the pair became suspended.
- Locate the IEA49*xx* message via NetView. Next, issue a CSUSPEND command from either host to the secondary device's subsystem to notify the storage control that the PPRC pair is in suspended state. The recovery subsystem records the suspension time.

Putting PPRC error recovery procedures into effect

Table 49 identifies how the system processes an error to a volume in a volume pair depending on how the volume and path were established. The options for the CESTPATH CGROUP parameter and the CESTPAIR CRIT parameter both affect the way the recovery system processes the volume pair when a device error occurs.

Table 49. Choosing the level of volume recovery for PPRC

What level of error recovery do you want?	Specify . . .	In the event of a PPRC volume error, the ESS storage subsystem does the following . . .	Some considerations . . .
I must have **total data consistency** between primary and secondary volumes.	CRIT(YES) on CESTPAIR command (see Note).	• Suspends the volume pair that has the error. • Unit checks all write I/O to the primary device. • Issues an IEA494I message for the suspended condition.	PPRC stops all I/O to the volume pair, and suspends the pair, on any volume error.

Table 49. Choosing the level of volume recovery for PPRC (continued)

What level of error recovery do you want?	Specify . . .	In the event of a PPRC volume error, the ESS storage subsystem does the following . . .	Some considerations . . .
I want to include **automated error recovery procedures**.	CRIT(YES) on CESTPAIR command, and CGROUP(YES) on CESTPATH command. **Note:** The CGROUP parameter directs PPRC to remove the CRIT(YES) status from a volume pair that experiences an error.	• Issues an IEA494 message for the Suspending and Extended-Long-Busy state conditions. • Suspends the volume pair that has the error. • Queues application I/O, for volume pairs added with CESTPATH CGROUP(YES), in cache for two minutes. • Resets the CRIT(YES) parameter to CRIT(NO), if applicable. • Issues an IEA494I message for the suspended condition.	After two minutes, queued I/O resumes to the storage subsystem, if possible. Each installation is responsible for the viability of the automated recovery procedures.
I prefer a **manual error recovery** process.	The default for the CESTPAIR command is CRIT(NO), and the default for the CESTPATH command is CGROUP(NO).	• Suspends the volume pair that has the error. • Does **not** unit check the primary device. • Issues an IEA494I message for the suspended condition.	Each installation is responsible for the local error recovery procedures.

Note: The CRIT(YES) parameter has two modes of operation, depending on how the storage subsystem is configured. CRIT(YES-PATHS) inhibits all write operations when all paths to a secondary volume are unavailable. CRIT(YES-ALL) inhibits writes on any failure, including secondary device failures.

Typically, a PPRC volume pair enters suspended state when the primary volume's storage control fails to complete a write operation to the secondary volume's storage control. The failed I/O operation receives channel end, device end, and unit check (with 'FB' indicated in sense byte 8), and the MVS host disk ERP issues an IEA49*xx* message. The NVS in the primary volume's storage control records cylinders that have changed while the volume pair is in suspended or duplex state. After you have corrected the conditions that caused the suspension, recopy these cylinders in order to return both volumes to a synchronized state.

Table 50 on page 246, Table 51 on page 246, and Table 52 on page 247 list general failures, failures detected by the host, and failures detected by the storage subsystem, and how to recover from them.

Table 50. General failure recovery procedures for PPRC

Error condition	Recovery action
Room power failure	If power is lost to the: • **Primary storage subsystem**, all volumes are placed in suspended or duplex state when power is restored. Use the CQUERY command to determine the exact state of all PPRC volume pairs. • **Secondary storage subsystem**, all volumes are placed in suspended or duplex state, and the primary volume's storage control records changes to the volume in its NVS. Issue the CESTPATH command to reestablish paths, and the CESTPAIR command with the RESYNC option to resynchronize a suspended or duplex PPRC volume pair, and to return the volumes to full duplex state. Since room power failures do not power off all devices at the same exact time, the RESYNC option is necessary in order to copy all updates that may have been made to the primary volumes, but not yet secured on the recovery site subsystem. The primary volume's subsystem maintains an NVS bitmap of all changed cylinders. This is important where power is restored to the primary hosts and primary volume's subsystem before it is restored to the secondary volume's subsystem. The CESTPAIR with the RESYNC option resynchronizes the volume pairs when power is restored to the secondary volume's subsystem.

Table 51. Host failure recovery procedures for PPRC

Error condition	Recovery action
Application host failure	1. Re-IPL the host. If power has not been lost to the subsystems, the peer-to-peer remote copy pairs will still be operational and in duplex state. 2. Issue the CQUERY command to verify the PPRC volume state for each PPRC volume pair. 3. If the volume pair is in duplex state, no further action is required. If the volume pair is in suspended or duplex state, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.
Last channel path failure	Host connectivity to the subsystem is lost. PPRC pair operation continues, assuming that the channel path error does not affect the PPRC link. 1. When host connectivity is restored, issue the CQUERY command to verify the PPRC volume state for each PPRC volume pair. 2. If the volume pair is in duplex state, no further action is required. If the volume pair is in suspended or duplex state, issue the CESTPATH command, and the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.
Last PPRC path failure	When the last path to the recovery site storage control is lost, the PPRC pair is placed into suspended or duplex state. When the paths are repaired or reconfigured, the primary volume's storage control automatically attempts to reestablish the paths. Issue the CESTPAIR command with the RESYNC option to return the pair to duplex state.
Device varied offline	Depending on how the device is varied offline, host connectivity to the subsystem may or may not be lost. Application connectivity to the subsystem is lost, but assuming that the device has not been varied offline in "force" or "boxed" modes, the PPRC pairs continue in operation. 1. When the device is varied back online, verify the PPRC volume status with the CQUERY command. 2. If the volume pair is in duplex state, no further action is required. If the volume pair is in suspended or duplex state, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.

Table 51. Host failure recovery procedures for PPRC (continued)

Error condition	Recovery action
Missing interrupt	A missing interrupt generally occurs when a host issues an I/O operation to a subsystem, and the operation does not complete within a specific amount of time. 1. Analyze and correct the missing interrupt condition, using information contained in the message that the host sent to the operator console. Depending on the cause of the missing interrupt, the PPRC pair operation may or may not continue. 2. When you have corrected the missing interrupt condition, issue the CQUERY command to verify the state of the affected PPRC volume pair. 3. If the volume pair is in duplex state, no further action is required. If the volume pair is in suspended or duplex state, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.

Table 52. Storage subsystem failure recovery procedures for PPRC

Error condition	Recovery action
Equipment check	A number of errors fall into this category. The error may or may not be detected on the actual I/O operation that caused the failure. Generally, an equipment check that suspends a PPRC volume pair is reported as follows: • The subsequent I/O command is unit checked (with 'FB' sense) with indicators set for equipment check, permanent error, and environmental data present. • The disk ERP issues the appropriate IEA49*xx* error message. • The PPRC volume pair is placed into suspended or duplex state. For the **primary** PPRC volume's storage control: Change recording *will not continue* for the primary volume. When you have corrected the error condition on the primary volume's storage control, issue the CESTPAIR command with the COPY or RESYNC option to recover the PPRC volume pair. See the note at the end of this table for more information on CESTPAIR command options. For the **secondary** PPRC volume's storage control: Change recording *will continue* for the primary volume. When you have corrected the error condition on the secondary volume's storage control, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair, or recopy the entire volume if its data has been corrupted by the recovery process.
Storage control failure	A number of errors fall into this category. The error may or may not be detected on the actual I/O operation that has caused the failure. Most storage control failures are reported as follows: • The I/O command that encountered the error is unit checked (with 'FB' sense) with indicators set for equipment check, permanent error, and environmental data present. • The disk ERP issues an IEA49*xx* error message. • The PPRC volume pair is placed into suspended or duplex state. For the **primary** PPRC volume's storage control: Change recording *will not continue* for the primary volume. When you have corrected the error condition on the primary volume's storage control, issue the CESTPAIR command with the COPY or RESYNC option to recover the PPRC volume pair. See the note at the end of this table for more information on CESTPAIR command options. For the **secondary** PPRC volume's storage control: Change recording *will continue* for the primary volume. When you have corrected the error condition on the secondary volume's storage control, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.

Table 52. Storage subsystem failure recovery procedures for PPRC (continued)

Error condition	Recovery action
Invalid track format	An invalid track format error is reported as follows: • Returned sense indicates an equipment check, permanent error, with environmental data present. • The disk ERP issues a “Volume Suspended” error message (the same action that occurs with dual copy). For the **primary** PPRC volume’s storage control: Change recording *will not continue* for the primary volume. When you have corrected the error condition on the primary volume’s storage control, issue the CESTPAIR command with the COPY or RESYNC option to recover the PPRC volume pair. See the note at the end of this table for more information on CESTPAIR command options. For the **secondary** PPRC volume’s storage control: Change recording *will continue* for the primary volume. When you have corrected the error condition on the secondary volume’s storage control, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.
Permanent data check	For a permanent data check on a **primary PPRC volume**, the PPRC volume pair status is not changed as a result of this specific error. The following actions occur: • A message is sent to the operator console. • The application program receives channel end, device end, and unit check in response to its I/O operation. • The disk ERP takes the appropriate recovery actions, including issuing the operator message and logging the error. • The application program performs its recovery actions. During a destage operation, a permanent data check on a track in the primary volume of a PPRC volume pair causes the primary volume’s storage control to pin that track in NVS. A subsequent I/O command is unit checked with ‘FB’ sense. The disk ERP issues the appropriate IEA49*xx* error message, and the PPRC volume pair is placed into suspended or duplex state. Change recording continues for the primary volume. When the pinned data condition is corrected on the primary volume’s storage control, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair. A permanent data check on a track in a **secondary volume** of a PPRC volume pair causes the secondary volume’s storage control to pin that track in NVS. A subsequent I/O command is unit checked with ‘FB’ sense. The disk ERP issues the appropriate IEA49*xx* error message, and the PPRC volume pair is placed into suspended or duplex state. Change recording continues for the secondary volume. When the pinned data condition is corrected on the secondary volume’s storage control, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.

Table 52. Storage subsystem failure recovery procedures for PPRC (continued)

Error condition	Recovery action
Intervention required	For the **primary PPRC volume**, this normal system operating condition is handled the same as it is for simplex devices. The host issues an “intervention required” message for the device, then takes the appropriate actions based on the operator’s reply. No data is changed on the primary volume, and the PPRC volume pair state remains unchanged. When the intervention condition is cleared, normal operations resume. For the **secondary PPRC volume**: • The subsequent I/O command is unit checked (with ‘FB’ sense) with equipment check, permanent error, and environmental data present indicators set. • The disk ERP issues an IEA49*xx* error message indicating “Intervention Required” for the secondary device. • The PPRC volume pair is placed into suspended or duplex state. Change recording *will continue* for the primary volume. When you have corrected the error condition on the secondary volume, issue the CESTPAIR command with the RESYNC option to recover the PPRC volume pair.
Cache or NVS failure, or cache reinitialization	PPRC implements write I/O operations to both primary and secondary volumes. The cache and NVS in each respective storage control hold the write updates until the updates are destaged. An NVS or cache error, or the reinitialization of cache, puts the PPRC pair into suspended or duplex state. Change recording continues. For a cache or NVS failure or cache reinitialization that affects the storage control at either end, issue the CESTPAIR command with the RESYNC option to the suspended volume after the cache and NVS have been made available. This action reestablishes the volume pair and copies any cylinders that were modified while the pair was suspended. The storage control automatically copies all cylinders from the primary volume to the secondary volume whenever an NVS failure affects the primary volume’s storage control. This is necessary so that the primary volume’s storage control can maintain the changed-cylinder map within its NVS.
Note: Generally, repairs to volumes are done using ICKDSF. The volume pair must be reinitialized if corrective actions change the data on a track so that the primary and secondary volume tracks are no longer identical. That track can be restored from a backup volume if the error causes data on a specific track to be overwritten. Select the appropriate CESTPAIR copy mode based on the completed repair action and the status of the data on the primary and secondary volumes.	

RAMAC device sparing with PPRC

The procedure described here explains how to spare (replace) a PPRC-managed RAMAC Disk drawer without requiring that you interrupt your application. Another option is to use the P/DAS function to redirect the application I/O while you correct the error on the RAMAC volume. “Chapter 18. PPRC Dynamic Address Switching (P/DAS)” on page 251 describes P/DAS operations.

You temporarily end the PPRC volume pair as part of the RAMAC device sparing procedure. During this time, there is a potential exposure if there is not a current copy of the volumes that are in the drawer. You may therefore want to set up a procedure to recover the volumes in case of a failure during this device sparing process.

Steps to spare a RAMAC drawer

Before you begin: You should be familiar with the commands used in this procedure. See “Chapter 13. Peer-to-Peer Remote Copy TSO Command Descriptions” on page 199 for PPRC command information.

Perform the following steps to spare a RAMAC drawer. Your specific recovery procedures may vary from these according to the type of data, or data sets, that are involved.

1. Issue the CDELPAIR command to end the PPRC pair.

   ---

2. Have your service representative initiate the drawer replacement procedure from the maintenance panel.

   ---

3. Reenable the RAMAC disk when the sparing process has completed.

   ---

4. When the primary and secondary volumes are available, establish the PPRC pair with a full-volume copy. When you have completed the copy and the pair is in duplex mode, the disaster recovery exposure is resolved.

   Note: You may want to consider establishing the pair to a different volume than the original secondary volume. This would preserve the original secondary volume for recovery in case an error occurs before the copy operation has completed.

Chapter 18. PPRC Dynamic Address Switching (P/DAS)

PPRC dynamic address switching (P/DAS) is a software function that provides the ability to redirect all application I/O from one PPRC volume to another PPRC volume with minimal application impact. P/DAS works with **PPRC volumes only**, and cannot redirect activity among XRC volumes.

P/DAS allows application-transparent switching of I/O to support the following tasks:
- Planned outages (device or subsystem)
- Device migration
- Workload movement

You can use P/DAS in shared-disk environments, including both sysplex environments and nonsysplex environments.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
How does P/DAS work?	251
Preparing for P/DAS operations	252
P/DAS operations	253
Replying to P/DAS-related messages	259
P/DAS error conditions	262
P/DAS operations example	263

How does P/DAS work?

P/DAS commands allow the system operator to redirect application I/Os that are currently sent to the PPRC primary volume to instead go to the secondary volume. PPRC volume pairs must be in duplex state for the P/DAS function to perform.

P/DAS manages all I/O redirection at the I/O supervisor level. P/DAS operations are therefore nearly transparent to any application program that reads from or writes to the volume. Based on the options you select, the redirection of the I/O can continue across system IPLs. The options can apply to all systems in the sysplex that are attached to the volumes.

The P/DAS function performs device switching and copying functions, as does dual copy. If you are familiar with dual copy operation, a comparison of both functions may be helpful. Table 53 summarizes some of the operating differences between the dual copy and P/DAS functions.

Table 53. Comparing P/DAS and dual copy

Function	P/DAS	Dual copy
How is switching accomplished?	Nontransparently	Transparently
Where does the operation occur?	Between storage controls	Within the same storage control
How is the function controlled?	By the operator	By the subsystem

© Copyright IBM Corp. 1994, 2002

Table 53. Comparing P/DAS and dual copy (continued)

Function	P/DAS	Dual copy
Which system component manages the swap?	MVS system	Storage subsystem
How is I/O addressed?	Operator specifies target address	Operator specifies source address

P/DAS operates similarly to DDR Swap. After a swap operation, the “To” device comes online and the “From” device goes offline.

Preparing for P/DAS operations

This section describes the hardware and software requirements, environmental conditions, and installation-specific considerations that must be satisfied before initiating P/DAS functions.

Hardware and software requirements

The P/DAS function is available on MVS/ESA™ 5.1 and above (with PTFs) in conjunction with levels of all supported OS/390 and z/OS releases (with PTFs). P/DAS operations rely on the PPRC functions of the storage control.

Environmental conditions

The following conditions must be met before a P/DAS operation is initiated:

- The source volume for the P/DAS operation must meet the following conditions:
  - **Must be the primary volume of an active PPRC pair**.
  - Must be online to the system.
  - **Must not be part of an active XRC session**.
  - Must not have any active paging data sets in use.
  - Must not have outstanding reserves (cancel any jobs that have outstanding reserves prior to issuing P/DAS commands).
- The secondary volume for the P/DAS operation must meet the following conditions:
  - Must also be the secondary volume of the PPRC pair.
  - Must not have any outstanding target device allocations to it, such as ICKDSF or application programs.
- The PPRC pair must meet the following conditions:
  - The pair must be in duplex mode, which ensures that both the source and target devices are identical in data content.
  - The pair must be of the same device type, but not necessarily the same model.
  - Both volumes must either have the same capacity, or the capacity of the target volume must be greater than that of the source volume. (See Note below.)
  - The pair must not be in a CGROUP FREEZE state when you initiate P/DAS.

  Note: Following the P/DAS switch operation, the additional capacity of the target volume will not be realized until you run a subsequent ICKDSF job against the target volume to rebuild the VTOC.
- Neither the SVC dump nor the ICKDSF processes can be active during P/DAS operations.

Additional P/DAS considerations

Depending on your installation and operations, the following items may need your attention as you plan for P/DAS operations:

- Ensure that a JES3 restart or reinitialization uses the new device addresses that result from a P/DAS operation. The JES3 INISH deck must reflect the new device addresses for devices involved in a P/DAS operation. This includes devices that are part of a JES3 spool volume.
- Ensure that automatic procedures do not involve a device that P/DAS is switching, as I/O cannot occur during a switch operation.
- You can swap the SYSRES volume, provided that the volume meets the other P/DAS requirements.
- Do not use either the PPRC CSUSPEND QUIESCE option or the CGROUP FREEZE option in conjunction with P/DAS functions.
- Consider assigning a device the same device address on all systems in a sysplex environment. Operations such as P/DAS are easier when affected devices have the same device addresses within all systems.
- P/DAS supports operations between 3390, RAMAC, and ESS devices. However, you may want to run the latest level of ICKDSF to ensure that the VTOC reflects the proper device characteristics.
- The system operator must be aware of when a P/DAS function has taken place prior to a system IPL, and ensure that the proper volume is placed online after the IPL.

  During IPL, the system operator is prompted to select which volume to place online in situations where both the primary and secondary volumes are operational and have the same volume serial number.
- Be aware that, when P/DAS reestablishes a pair in the opposite direction, it establishes the new pair with the CRIT(NO) option.

  P/DAS operates with volumes that were established with either the CRIT(YES) or CRIT(NO) keywords. P/DAS must reestablish the pair with the CRIT(NO) option because the volume pair will subsequently become suspended. To reestablish the pair with the CRIT(YES) option, issue a CSUSPEND command, then a CESTPAIR command with RESYNC and CRIT(YES) parameters.
- Do not configure P/DAS to swap between installation-static and dynamic UCBs, as the results are unpredictable.
- Do not involve volumes that contain sysplex control data sets in P/DAS volume pair operations.

P/DAS operations

The following section describes the operator commands and available options to control P/DAS functions. P/DAS operations work within sysplex, nonsysplex, and shared-disk environments. First, we examine what a PPRC configuration looks like before P/DAS functions make changes to it.

Typical PPRC configuration before P/DAS operations

Figure 15 on page 254 shows a typical PPRC configuration as it exists before issuing any P/DAS commands. You can establish PPRC paths in both directions, if desired. You do not need to use an ESCON Director. With ESCON Directors, however, PPRC pairs can be up to 43 km apart.

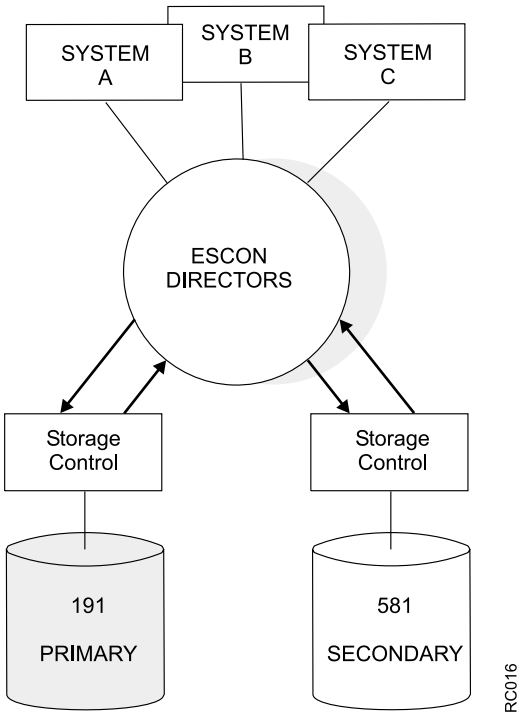

Figure 15. A typical PPRC configuration before a P/DAS operation. Disk device 191 is the PPRC primary volume, receiving updates from three attached systems through a network of ESCON Directors. PPRC copies all updates to secondary volume 581.

Steps for using P/DAS in a nonsysplex, shared-disk environment

This section describes the sequence of commands for performing a P/DAS operation in a nonsysplex, shared-disk environment. In all cases, device *ssss* refers to the primary source volume. Device *tttt* refers to the secondary target volume.

Before you begin: Identify all systems that have access to the devices that you want to switch. *You must issue the commands that follow on all systems that are connected to the devices that you plan to switch*.

Perform the following steps to use P/DAS in a nonsysplex, shared-disk environment:

1. Issue **IOACTION STOP,DEV=***ssss*.

   Initiate a P/DAS operation by issuing an IOACTION STOP command on *all* systems that are attached to the primary volume that is specified by DEV=*ssss*. This command requests that P/DAS stop I/O that is directed to device number *ssss*. DEV=*ssss* specifies a primary volume of a duplex PPRC volume pair.

   All application I/O that is issued to the primary volume remains queued in the MVS system until you issue an IOACTION RESUME command. P/DAS ends, with a simulated I/O error, any active I/O that does not complete within five seconds after the MVS system has accepted the IOACTION STOP. This action forces the disk error recovery procedures to retry the I/O operation.

   The MVS system issues the following messages to the operator console (or consoles) of all systems when the I/O for the primary device has been stopped on that system:

```
IOS600I IOACTION - THE FOLLOWING DEVICE(S) HAVE BEEN STOPPED:
                dev dev1-dev2...

IOS601I IOACTION - DEVICES REMAIN IN THE STOPPED STATE.
        USE THE 'D IOS,STOP' COMMAND TO DISPLAY THE DEVICES
```

When all attached systems have received the above messages, issue the swap command to switch volumes. MVS issues an IOS610I message in response to your D IOS,STOP command.

---

2. Issue **SWAP** *ssss*,*tttt*.

   Issue the SWAP command for the volume pairs involved in the P/DAS operation. The SWAP command directs the system to switch the source device *ssss* with the target device *tttt* and prepares the system to redirect all I/Os issued to device *ssss* to the target device *tttt*. Issue this command on all systems that are attached to the volumes that you plan to switch.

---

3. P/DAS performs validation.

   Before the SWAP operation is complete, P/DAS performs validation to ensure that it can complete the swap operation. If conditions exist that can cause a data integrity exposure, P/DAS ends the operation and generates an error message. See “P/DAS error conditions” on page 262 for error message descriptions.

   P/DAS does not proceed when it cannot determine the status of both volumes, or when the volumes are not in duplex state. Ensure that the environment meets the conditions necessary to support P/DAS. “Preparing for P/DAS operations” on page 252 describes the conditions.

   If the validation completes successfully, P/DAS issues the following message to prompt the operator as to how to proceed:

   ```
   IGF520A VERIFICATION COMPLETE: REPLY 1 TERMINATE PAIR, AND SWAP
    | 2 SWITCH PAIR, AND SWAP | 3 CONTINUE SWAP | 4 TERMINATE SWAP
   ```

---

4. Reply to the **IGF520A** message.

   “Replying to P/DAS-related messages” on page 259 describes message reply options. The system takes action according to two things:
   - The option that is selected by the operator
   - The environment that exists at the time the command is issued.

   The system issues the following completion message after the swap function is complete:

   ```
   IGF505I SWAP FROM ssss TO tttt COMPLETE
   ```

   Variable *ssss* is the source device number, and *tttt* is the target device number.

---

5. Issue **IOACTION RESUME,DEV=***tttt*.

   In a shared-disk environment, *ensure that all attached systems have received the IGF505I message before you issue the IOACTION RESUME command*. After you have received the IGF505I message, issue an IOACTION RESUME command from all systems that are attached to device *tttt*. You must issue the command to device number *tttt* because device *tttt* is now the primary volume.

The IOACTION RESUME command resumes all application I/O to device *tttt*. All application I/O that used to go to the source device *ssss* now goes to the target device *tttt* instead. This includes all I/O that is queued by MVS since the STOP command and all subsequent I/O. You must issue this command on all systems that share this volume.

The following message received at the console indicates that all application I/O has resumed and that the P/DAS operation has completed:

```
IOS607I IOACTION - THE FOLLOWING DEVICE(S) HAVE BEEN RESUMED:
                   dev,dev1-dev2
```

Typical PPRC configuration after a P/DAS function has completed

Figure 16 shows what the configuration looks like after the P/DAS function has completed. I/O has been redirected to the volume that was originally the secondary of the PPRC pair. Application I/O continues without interruption. In the example shown in Figure 16, the operator has chosen to keep both of the PPRC pair devices in operation. You can reverse PPRC operations if you have established paths in both directions.

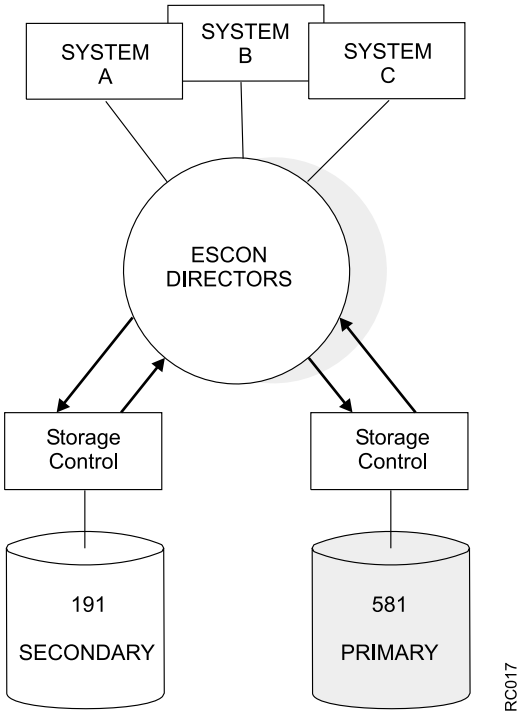

Figure 16. P/DAS swap operation complete. Disk device 581 is now the PPRC primary volume, receiving updates from attached systems through a network of ESCON Directors. PPRC copies all updates to device 191, which has now become the secondary volume.

Steps for using P/DAS in a sysplex environment

In the sysplex environment, P/DAS operations are similar to those of the nonsysplex, shared-disk environment. See “P/DAS operations example” on page 263 for an example of a sysplex scenario.

Note: The ROUTE *ALL option, which is part of the following procedures, is available on all MVS/SP 5.2 and above systems. MVS/SP 5.1 JES2 also supports the option , but MVS/SP 5.1 JES3 does not. On systems that do

not support the ROUTE *ALL option, issue the commands, without specifying ROUTE *ALL, from each of the systems.

Before you begin: You must choose one system to be the “main system.” When multiple systems are involved, you must issue the following commands on each system in the specified order.

Perform the following steps to use P/DAS in a sysplex environment:

1. Issue **ROUTE *ALL,IOACTION STOP,DEV=***ssss*.

   Initiate a P/DAS operation by issuing this command on the “main system.” This command requests that P/DAS stop I/O that is directed to device number *ssss*. ROUTE *ALL directs this action to all systems in the sysplex. DEV=*ssss* specifies a primary volume of a duplex PPRC volume pair. The MVS system stores all application I/O that is issued to the primary volume until you issue a ROUTE *ALL,IOACTION RESUME command. P/DAS ends, with a simulated I/O error, any active I/O that does not complete within five seconds after MVS has accepted the IOACTION STOP. This action forces the disk error recovery procedures to retry the I/O operation.

   The MVS system issues the following messages to the operator console for all systems in the sysplex when the I/O for the primary device has been stopped on that system:

   ```
   IOS600I IOACTION - THE FOLLOWING DEVICE(S) HAVE BEEN STOPPED:
                   dev dev1-dev2...

   IOS601I IOACTION - DEVICES REMAIN IN THE STOPPED STATE.
           USE THE 'D IOS,STOP' COMMAND TO DISPLAY THE DEVICES
   ```

   When all attached systems have received these messages, continue with the SWAP command. MVS issues an IOS610I message in response to your D IOS,STOP command.

---

2. Issue **ROUTE *ALL,SWAP** *ssss*,*tttt*.

   Issue the ROUTE *ALL,SWAP command for the volume pairs involved in the P/DAS operation. ROUTE *ALL automatically directs this action to all systems in the sysplex. This command directs the systems to switch the source device *ssss* with the target device *tttt*, and prepares the systems to redirect all I/Os issued to device *ssss* to the target device *tttt*.

---

3. P/DAS performs validation.

   Before the SWAP operation is complete, P/DAS performs validation to ensure that it can complete the swap operation. If conditions exist that can cause a data integrity exposure, P/DAS ends the operation and generates an error message. See “P/DAS error conditions” on page 262 for error message descriptions.

   P/DAS does not proceed when it cannot determine the status of both volumes, or when the volumes are not in duplex state. Ensure that the environment meets the conditions that are necessary to support P/DAS. “Preparing for P/DAS operations” on page 252 describes the conditions.

   If the validation completes successfully, P/DAS issues the following message to prompt the operator as to how to proceed:

   ```
   IGF520A VERIFICATION COMPLETE: REPLY 1 TERMINATE PAIR, AND SWAP
    | 2 SWITCH PAIR, AND SWAP | 3 CONTINUE SWAP | 4 TERMINATE SWAP
   ```

4. Select the P/DAS action.

   Attention: Do not reply to any of the system messages until all of the systems in the sysplex have issued IGF520A messages.

   Select one of the following replies from the main system console:
   - Reply with “1” to end the pair and redirect I/O to the secondary volume.
   - Reply with “2” to swap the primary and secondary volumes and redirect all I/O to the new primary volume.

   Attention: If more than one system issues the above reply, the P/DAS swap operation may fail on one or more systems. If such an error does occur, reestablish the PPRC pair on all systems and retry the P/DAS operation.

   If you choose Reply 2 (switch pair, and swap), P/DAS directs the system to take action according to the PPRC environment that is in effect at the time the command is issued, as described in Table 55 on page 260.

   The main system issues the following completion message when the swap function is complete:

   ```
   IGF505I SWAP FROM ssss TO tttt COMPLETE
   ```

   Variable *ssss* is the source device number, and *tttt* is the target device number.

5. Reply from other sysplex systems.

   Ensure that the main system has issued an IGF505I message before replying to outstanding IGF520A messages on other sysplex systems. At this point, reply to outstanding IGF520A messages on all other SYSPLEX systems by entering “3” (Continue Swap). Table 54 on page 259 describes the resulting system action.

6. Issue **ROUTE *ALL,IOACTION RESUME,DEV=***tttt*.

   Each of the attached sysplex systems issues an IGF505I message in response to your “Continue Swap” reply. Ensure that all systems in the sysplex have received the IGF505I message, and then issue the IOACTION RESUME command to device number *tttt*. You must issue the command to device number *tttt* because device *tttt* is now the primary volume. The IOACTION RESUME command resumes all application I/O to device *tttt*. All application I/O that used to go to the source device *ssss* now goes to the target device *tttt* instead.

   This command is equivalent to the IOACTION RESUME command in the nonsysplex, shared-disk environment. The difference is that the specification of ROUTE *ALL directs this action to all systems in the sysplex. Every system in the sysplex should acknowledge this command with an IOS607I message.

   Note: P/DAS does not actually redirect and resume the application I/O to the target volume until this step has completed. Each attached system redirects the I/O as it receives and processes the IOACTION RESUME command.

Replying to P/DAS-related messages

The operator's reply to the IGF52xA messages directs the P/DAS action, as described in Table 54. In some cases, the activities performed by P/DAS are based on a combination of the operator reply and the environment at the time that the P/DAS swap request is made.

Table 54. Reply options for IGF52xA messages

Your reply (to message) . . .	**Results in this system action . . .**	**Notes . . .**
Terminate Pair, and Swap (IGF520A)	This reply directs the system to perform the following functions: 1. End the PPRC pair and stop copy operations. Other PPRC operations for other pairs continue unchanged. 2. Redirect all application I/O from the source device *ssss* to the target device *tttt*. Application I/O is not affected by the device switch. The I/Os only update the secondary volume from this point on, and the primary and secondary volumes will no longer contain the same information.	In a **sysplex**, issue this reply from the main system only. **Application:** For device migration scenarios where the source volume is no longer going to be used (such as when the device lease will expire) and you want to change from one set of primary volumes to a second set of primary volumes.
Switch Pair, and Swap (IGF520A)	This reply redirects application I/O to the secondary volume. The system takes action according to the PPRC environment that is in effect at the time of the swap request. The action performed depends on the following two conditions: • How the PPRC paths are currently established • Whether the target device is the same size as the source device The resulting system actions are summarized in Table 55 on page 260.	In a **sysplex**, issue this reply from the main system only. **Application:** For device or subsystem maintenance, and workload movement.
Continue Swap (IGF520A, IGF521A, or IGF522A)	This reply prompts the system to redirect application I/Os from the primary volume to the secondary volume. Thus, all I/Os that were directed to the source device *ssss* are now directed to the target device *tttt*. After completion of this option, device *ssss* is no longer involved in any operation. Like dual copy, the additional space on the target device *tttt* cannot be used until an ICKDSF operation has been run against the target device to rebuild the VTOC so the additional volume capacity is recognized.	This reply is valid within sysplex and nonsysplex environments. In a **sysplex**, issue this reply from any system. **Application:** To remove a primary volume from the configuration.

Table 54. Reply options for IGF52xA messages (continued)

Your reply (to message) . . .	**Results in this system action . . .**	**Notes . . .**
Try Again (IGF521A or IGF522A)	This reply directs the system to perform the following functions: 1. End the pair, and return both devices to simplex state. It is the equivalent of issuing a CDELPAIR command for the PPRC pair. 2. Establish a new PPRC pair in the reverse direction. The new path is from the new primary (original secondary volume, device number *tttt*) to the new secondary (original source volume, device number *ssss*). This is the equivalent of issuing a CESTPATH and CESTPAIR command for the PPRC pair. 3. Immediately suspend the new PPRC operation. This forces all changes made to the secondary to be recorded in the bit maps for the operation. The changes are not copied to the original primary, device number *ssss*. This is the equivalent of issuing a CSUSPEND command for the PPRC pair just established. 4. Redirect application I/Os from the primary volume to the secondary volume. Thus, all I/Os that were directed to the source device (*ssss*) are now directed to the target device (*tttt*). After completion of this option, device *ssss* is no longer involved in any operation.	Before this option is used, paths have been established in the opposite direction from the target device's storage control to the source device's storage control. In a **sysplex**, issue this reply from any system.
Terminate Swap (IGF520A, IGF521A, or IGF522A)	This is a request to end the swap operation. Message IGF512I message is issued to the system operator log as follows: `IGF512I SWAP FROM ssss TERMINATED - SWAP TERMINATED BY OPERATOR` P/DAS returns error code 16 to the system.	In a **sysplex**, issue this reply from any system.

System actions for switch pair, and swap

When you reply to message IGF520A with "switch pair, and swap," P/DAS takes action that is based on the existing environment. Refer to Table 55 and to the actions labeled A through D to determine how P/DAS will respond to this reply.

Table 55. System actions for switch pair, and swap. Descriptions for the individual actions follow this table.

PPRC path direction	**System action with equivalent device geometry**	**System action with different device geometry**
PPRC paths exist in both directions	Action A	Action B
PPRC paths only exist from primary to secondary	Action C	Action D

Explanation of terms:

- **Paths exist in both directions** means that a path exists from the primary volume's storage control to the secondary volume's storage control, and another path exists from the secondary to the primary. The PPRC pair is able to operate in either direction, and the environment is ready to allow a switch of the PPRC pair from secondary to primary. When paths exist only from the primary to the secondary, the environment is not set up to allow data flow from the target volume to the source volume.
- **Equivalent device geometry** means that the primary and secondary volumes have equivalent track sizes and volume capacity.
- **Different device geometry** means that the primary and secondary volumes have the same track size but the secondary has a larger volume capacity than the primary.

The following specific actions relate to Table 55 on page 260. The table refers to these actions according to your existing environment.

Action A:

There are PPRC paths in place from the primary volume to the secondary and also from the secondary to the primary. The source and target volumes have equivalent device geometry. In this situation, selecting “switch pair, and swap” directs P/DAS to perform the following actions:

1. End the current PPRC pair.

   _____

2. Establish a new PPRC pair (with NOCOPY) to activate copying in the reverse direction (from the target device *tttt* to the source device *ssss*).

   _____

3. Immediately suspend the PPRC operation. All changes made to target device *tttt* are recorded in the bit maps for the operation, but are not to be copied to the original source device *ssss*.

   _____

4. Redirect all application I/O to target device number *tttt*. The change is transparent to the application programs. P/DAS only sends I/Os to the secondary volume from this point on, and the primary and secondary volumes will no longer contain the same information.

   _____

The changes that occur between the swap operation and the subsequent resynchronization of the volume pairs are maintained on the old secondary volume. P/DAS will copy the changed updated tracks back to the original primary volume, and will make the volumes duplexed PPRC pairs again.

Action B:

There are PPRC paths in place from the primary volume to the secondary and also from the secondary to the primary. The target device *tttt* has greater volume capacity than the source device *ssss*. In this situation, your “switch pair, and swap” choice directs P/DAS to issue the following additional message to the system operator console:

```
IGF522A UNABLE TO SWITCH, FROM DEVICE IS SMALLER THAN TO DEVICE:
 REPLY 1 TO CONTINUE SWAP | 2 TO TERMINATE SWAP
```

Select one of the message options, as described in Table 54 on page 259.

If you choose to continue with the swap, the changes that occur between the swap and the subsequent resynchronization of the volume pairs will not be maintained on the old secondary volume.

Action C:

A PPRC path exists from the primary to the secondary, but not in the reverse direction. The primary and secondary volumes have the same physical capacity. In this situation, your “switch pair, and swap” choice directs P/DAS to issue the following additional message to the system operator:

```
IGF521A NO PATH IN OPPOSITE DIRECTION: REPLY 1 TO CONTINUE SWAP |
2 TRY AGAIN | 3 TERMINATE SWAP
```

Select one of the message options, as described in Table 54 on page 259.

Action D:

A PPRC path exists from the primary to the secondary, but not in the reverse direction. The primary and secondary volumes do not have the same physical capacity. P/DAS takes no special action to use the additional capacity on the target volume. In this situation, your “switch pair, and swap” choice directs P/DAS to issue the following additional message to the system operator:

```
IGF521A NO PATH IN OPPOSITE DIRECTION: REPLY 1 TO CONTINUE SWAP |
2 TRY AGAIN | 3 TERMINATE SWAP
```

Select one of the message options, as described in Table 54 on page 259. If you reply with “2” (Try Again), P/DAS issues the following message, which requires an operator reply:

```
IGF522A UNABLE TO SWITCH, FROM DEVICE IS SMALLER THAN TO DEVICE:
REPLY 1 TO CONTINUE SWAP | 2 TO TERMINATE SWAP
```

Select one of the message options, as described in Table 54 on page 259.

P/DAS error conditions

This section describes error conditions that can occur while attempting to perform P/DAS operations or during the resulting system actions.

Storage-related failures

The system obtains storage to perform P/DAS operations. If the storage requests fail, the system issues the following error message:

```
IGF512I SWAP FROM ssss TERMINATED - GETMAIN FAILED - SP245
```

P/DAS writes this message to the system operator log and returns error code 16 to the system.

I/O-related failures

The system issues I/O to the storage subsystems when PPRC pairs are establishing and ending as part of P/DAS functions. If a storage subsystem I/O fails during these functions, MVS issues the following error message:

```
IGF512I SWAP FROM ssss TERMINATED - SWAP FAILED DUE TO I/O ERROR
```

P/DAS writes this message to the system operator log and returns error code 16 to the system.

Validation failures

One of the following error messages are generated when a validation process fails. P/DAS writes all messages to the system operator log and returns error code 16 to the system.

- The following message means that the volumes are not part of a PPRC pair or that the volumes are not in duplexed mode:

  ```
  IGF512I SWAP FROM ssss TERMINATED - NOT A VALID PPRC PAIR
  ```

- The following message means that the volume is an active paging volume:

  ```
  IOS602I IOACTION - STOP NOT ALLOWED FOR PAGING DEVICES.
          THE FOLLOWING DEVICE(S) ARE IGNORED:  dev dev1-dev2...
  ```

- The following message means that a system in the sysplex has reserved the volume:

  ```
  IOS604I IOACTION - STOP NOT ALLOWED FOR RESERVED OR RESERVED PENDING
          DEVICES.
          THE FOLLOWING DEVICE(S) ARE IGNORED:  dev dev1-dev2...
  ```

- The following message means that the volume is currently allocated:

  ```
  IOS605I IOACTION - STOP NOT ALLOWED FOR DEVICES IN USE BY THE SYSTEM.
          THE FOLLOWING DEVICE(S) ARE IGNORED:  dev dev1-dev2...
  ```

If devices remain stopped

MVS issues the following message if devices remain stopped for an extended period of time:

```
IOS601I IOACTION - DEVICES REMAIN IN THE STOPPED STATE.
        USE THE 'D IOS,STOP' COMMAND TO DISPLAY THE DEVICES
```

P/DAS operations example

Table 56 on page 264 shows the sequence of events that make up a P/DAS operation in a sysplex environment. As in all P/DAS operations, the correct command sequence is necessary to perform a successful P/DAS exchange of devices.

In this example, there are three systems in the sysplex that share volumes 191 and 581. System A is the main system and has the master console. Initialize the environment by establishing a single-direction PPRC path from the storage control for device 191 to the storage control for device 581. Device 191 is the primary volume of the PPRC pair, and device 581 is the secondary volume.

In preparation for P/DAS Operation, the storage administrator issues two CESTPAIR commands. The commands are issued:

- From any of the attached systems to establish a path from the storage control for device 191 to the storage control for device 581.
- From any of the attached systems for device 191 (primary) and device 581 (secondary). PPRC copies all of the tracks on device 191 onto device 581.

Table 56 shows what happens when the P/DAS operation begins.

Table 56. Example of a P/DAS swap function in a sysplex environment

System A (main system)	System B	System C
Operator issues: ROUTE *ALL,IOACTION STOP,DEV=191. System A stops I/O on 191, then issues IOS600I and IOS601I messages.	System B stops I/O on 191, then issues IOS600I and IOS601I messages.	System C stops I/O on 191, then issues IOS600I and IOS601I messages.
Operator issues: ROUTE *ALL,SWAP 191,581. System A performs validation, then issues IGF520A message.	System B performs validation, then issues IGF520A message.	System C performs validation, then issues IGF520A message.
Operator issues: R nn,1. System A ends PPRC pair, swaps pair, then issues IGF505I message.	After IGF505I message received, operator issues: R nn,3. System B swaps PPRC pair, then issues IGF505I message.	After IGF505I message received, operator issues: R nn,3. System C swaps PPRC pair, then issues IGF505I message.
Operator issues: ROUTE *ALL,IOACTION RESUME, DEV=581. System A directs I/O to 581, then issues IOS607I message.	System B directs I/O to 581, then issues IOS607I message.	System C directs I/O to 581, then issues IOS607I message.

Part 4. FlashCopy

© Copyright IBM Corp. 1994, 2002

Chapter 19. What Is FlashCopy?

This section explains what the FlashCopy copy function is and how it works, and describes the various hardware and software products that are necessary. It also details the new TSO commands that are available for using the FlashCopy capabilities of the Enterprise Storage Server (ESS) D/T2105 storage subsystem.

FlashCopy is an Advanced Copy Services function that can quickly copy data from source volumes to target volumes that are within the same logical storage subsystem (LSS). DFSMSdss calls the FlashCopy function when the source and target volumes are both located within the same LSS on an ESS.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Using FlashCopy	268
FlashCopy requirements	267
FlashCopy TSO commands	269

FlashCopy requirements

There are both software and hardware prerequisites for using FlashCopy.

Determining FlashCopy software requirements

DFSMS/MVS Version 1 Release 3 and subsequent releases provide FlashCopy support.

Determining FlashCopy hardware requirements

MVS/ESA SP™ 4.3 runs on the following processors:

- All ES/9000 models that support ESA/390
- All ES/3090™-9000T models that support ESA/370™ or ESA/390
- All ES/3090 E, S, J, and JH models at the SEC level for ESA/370

FlashCopy is a feature on ESS storage subsystems.

FlashCopy operates with 3390 devices, and includes 3390 devices that are in 3380 track-compatibility mode. The source and target volumes must have the same track format.

Making cache available

To use FlashCopy with an ESS storage subsystem, both the subsystem cache and NVS with battery backup must be active. The statuses of cache and DASD fast write have no effect on whether or not you can use FlashCopy. Their statuses can, however, affect the overall performance of volumes with active FlashCopy sessions.

© Copyright IBM Corp. 1994, 2002

Using FlashCopy

This section describes FlashCopy copy operation and explains how the hardware and software components work together. You can also use the ESS Specialist to perform FlashCopy functions. The ESS Specialist is a Web-enabled interface that is part of the ESS storage subsystem. Refer to *IBM Enterprise Storage Server Web Users Interface Guide* for information about configuring the copy services options of the ESS. The Redbook, *Implementing the Enterprise Storage Server in Your Environment*, is another good ESS Specialist information source.

Steps for initializing sessions

Functionally, the data that FlashCopy is copying is unavailable only long enough for DFSMSdss to initialize a FlashCopy session. Initialization normally takes only a few seconds to complete successfully. The session initialization process is a series of distinct steps, as follows:

1. Determine an appropriate time to start a backup of databases, and stop update activity to those databases.

   ---

2. Use DFSMSdss to perform a copy of the volume. You can use DFSMSdss directly from a batch job or from an application through the DFSMSdss application interface.

Normal application activity can resume as soon as the FlashCopy session initialization is complete.

Copying data sets with FlashCopy

FlashCopy is much faster than traditional methods, especially when you want to copy large amounts of data. DFSMSdss can use FlashCopy to perform a full-volume copy of the data from the source volumes to the target volumes when the following conditions are true:

- The requested DFSMSdss operation is COPY FULL.
- Both the source and target devices are in the same ESS storage subsystem and are both within the same LSS.
- The source and target devices must have the same track format.
- The designated target volume cannot also be a primary or secondary volume in an XRC or PPRC volume pair.

FlashCopy cannot accommodate relocation of tracks. The CCHH location must be the same on the target volume as it is on the source volume.

DFSMSdss automatically performs FlashCopy if possible, whether or not you specify the CONCURRENT keyword.

FlashCopy TSO commands

Enterprise Storage Subsystems FlashCopy functions can be invoked by the TSO commands in Table 57:

Table 57. FlashCopy TSO commands

Command . . .	Description . . .	See . . .
FCESTABL	Establishes a FlashCopy relationship.	"FlashCopy establish (FCESTABL) command" on page 270
FCQUERY	Queries an ESS device.	"FlashCopy query (FCQUERY) command" on page 274
FCWITHDR	Withdraws a FlashCopy relationship.	"FlashCopy withdraw (FCWITHDR) command" on page 275

The commands listed above are authorized TSO commands. They should be protected by RACF FACILITY class protection or other authorization methods. These commands should be protected in a way similar to the PPRC commands.

You must exercise caution when entering these commands. Data loss is possible if the target for a FlashCopy is not carefully selected, especially when using the NOCOPY option. For more information, see "Destaging target data when using the NOCOPY option" on page 271.

Data on the target device will be overlaid within the track extents specified on the FCESTABL command if you specify the COPY option. If you specify the FCESTABL command with the NOCOPY option specified, any source volume updates to data within the track extents causes an overlay of the corresponding track on the target device. There is no host system data protection for the track extent areas on the target device like the host system protection for the source device data.

Controlling access to FlashCopy commands

The following are two ways to protect access to FlashCopy commands:

- Put the commands in a RACF-protected library. **or**
- Define resource profiles in the RACF FACILITY class and restrict access to those profiles.

Steps for placing TSO commands in a RACF-protected library

To restrict FlashCopy TSO commands to authorized storage administrators only, place the TSO commands in a RACF-protected library.

Perform the following steps to RACF-protect FlashCopy commands:

1. Issue the following RDEFINE command for each FlashCopy command that you want defined to RACF:

   ```
   RDEFINE PROGRAM cmdname ADDMEM('SYS1.CMDLIB') UACC(NONE)
   ```

   where:

   cmdname Defines the FlashCopy command name. Issue a separate RDEFINE command for each of the FlashCopy commands. RACF can only perform checking on commands that are defined to it.

2. Issue the PERMIT command for all commands and authorized FlashCopy TSO command users as follows:

   ```
   PERMIT cmdname CLASS(PROGRAM) ID(name) ACCESS(READ)
   ```

   where:

   cmdname Defines the FlashCopy TSO command name.

   name Defines the user ID receiving RACF access authority for that command name.

3. Issue the SETROPTS command from a user ID having the appropriate authority:

   ```
   SETROPTS CLASSACT(PROGRAM) WHEN(PROGRAM) REFRESH
   ```

Defining resource profiles in the RACF FACILITY class

You can limit the use of FlashCopy commands by defining resource profiles in the RACF FACILITY class and restricting access to these profiles. To use a protected command, you need read-access authority to the applicable profile. For FlashCopy commands, there are two RACF FACILITY class resource profile names:

- STGADMIN.ANT.ESFC.COMMANDS
- STGADMIN.ANT.ESFC.FCQUERY

Note: Authorize FCQUERY command use with the STGADMIN.ANT.ESFC.COMMANDS profile or the STGADMIN.ANT.ESFC.FCQUERY profile. FlashCopy first checks STGADMIN.ANT.ESFC.COMMANDS for authorization. If authorization is not permitted with STGADMIN.ANT.ESFC.COMMANDS, FlashCopy checks the STGADMIN.ANT.ESFC.FCQUERY profile for authorization to issue the FCQUERY command.

Examples: The following RACF command examples activate the RACF FACILITY class, define the profile for the FlashCopy commands, and give user STGADMIN authority to use this profile:

1. Activate the RACF FACILITY class:
   - SETROPTS CLASSACT(FACILITY)
2. Define the profile for FlashCopy commands, and authorize user STGADMIN to use this profile:
   - RDEFINE FACILITY STGADMIN.ANT.ESFC.COMMANDS UACC(NONE)
   - PERMIT STGADMIN.ANT.ESFC.COMMANDS CLASS(FACILITY) ID(STGADMIN) ACCESS(READ)

See the *z/OS Security Server RACF Security Administrator's Guide* for details on activating the RACF FACILITY class, and on defining and authorizing users to the FlashCopy command profiles.

FlashCopy establish (FCESTABL) command

Use the FCESTABL TSO command to establish a FlashCopy relationship between the user-specified source and target devices. This relationship has two possible operations.

- With a background copy—MODE(COPY). This is the default operation.

  The COPY mode establishes the FlashCopy relationship between the source and target devices and begins the background copy of all data on the tracks within the specified extents. As soon as the relationship is established, user programs have access to two logical copies of the source data. Updates to the source volume after the FlashCopy relationship is established are *not* part of the logical copy on the target device.

  You can modify the target device data while the background copy is in progress. The target device data that you modify remains in the state of the last user update, regardless of the updates to the source device at the corresponding track location. After the background copy is complete, the data on the target device is the same as the source device when the FlashCopy relationship was established unless you modify the target tracks within the track extents. When the background copy operation completes, the FlashCopy relationship ends.
- Without the background copy—MODE(NOCOPY).

  The NOCOPY mode establishes the FlashCopy relationship between the source and target devices. As soon as the FlashCopy relationship is established, user programs have access to two logical copies of the source data. Updates to source device data within the track extents causes the original source data to be copied to its related target device track. Only the original source device data will be copied to the target track; additional updates to source data will not be copied.

  You can modify target device data within the track extents. The target device data that is modified within the track extents remains in the state of the last user update, regardless of updates to the source device data. Updates to the source device after the FlashCopy relationship is established will *not* be available for reading from the target device. A FlashCopy relationship established with NOCOPY continues until a FlashCopy Withdraw (FCWITHDR) command is issued or until all tracks within the track extents have been updated on the target device.

  The ESS will support only one relationship between the source and target devices. The number of extents for each relationship can be from 1 to 110. The maximum number of extents can be used through the ANTRQST API services for FlashCopy. Because of the TSO command length restrictions, the FlashCopy FCESTABL TSO command can have from one to five extents.

Destaging target data when using the NOCOPY option

Issue the IDCAMS SETCACHE command with the DESTAGE option whenever you use the NOCOPY option on either the TSO FlashCopy FCESTABL command or the API FlashCopy FCESTABL command request. The SETCACHE command ensures that any updates made to the target device (prior to the FCESTABL command request) have been physically written to the device.

Example: The following JCL example shows the use of the SETCACHE command:

```
//JOB1     JOB
//STEP1    EXEC PGM=IDCAMS
//FILE1    DD   UNIT=3390,VOL=SER=AAAAAA,DISP=SHR
//SYSPRINT DD   SYSOUT=*
//AMSDUMP  DD   SYSOUT=*
//SYSIN    DD *
   SETCACHE DESTAGE  FILE(FILE1)
//
```

FCESTABL command syntax

The syntax of the FCESTABL TSO command is:

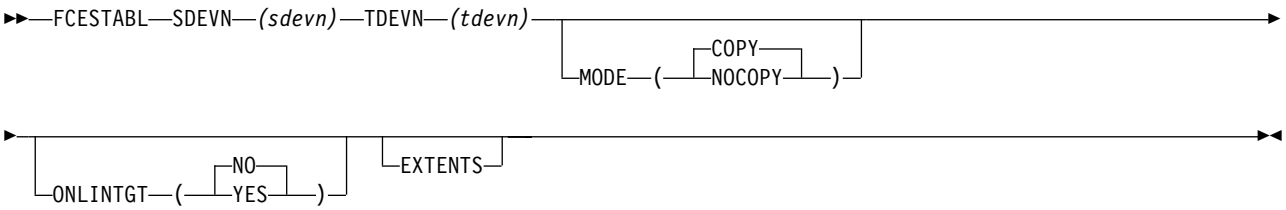

Required parameters

SDEVN Specifies the four-digit hexadecimal device number of the source device. The source device can be online or offline.

TDEVN Specifies the four-digit hexadecimal device number of the target device. The target device can be online or offline. If the target device is online, the completion of the FCESTABL command depends on the setting of the ONLINTGT parameter.

Optional parameters

MODE Specifies the type of FlashCopy operation.

- **COPY** Establishes the FlashCopy relationship between the source and target devices for the specified extent ranges. Perform a background copy of all source tracks within the specified extents to the target device. The FlashCopy relationship ends when the background copy is complete or the FCWITHDR command is issued. The default is COPY.
- **NOCOPY** Establishes the FlashCopy relationship between the source and target devices for the specified extent ranges. No background copy operation is performed. The FlashCopy relationship continues until a FCWITHDR command is received or all tracks within the track extents have been updated on the target device.

ONLINTGT Specifies whether an online device can be used as the target of a FlashCopy relationship.

- **YES** Specifies that the FlashCopy relationship can be established even if the target device appears to be online.
- **NO** Specifies that the FlashCopy relationship cannot be established if the target device appears to be online. The default is NO.

EXTENTS Specifies the first and last track of a set of contiguous tracks (an extent pair) to be part of this FlashCopy relationship. A minimum of one extent pair must be specified, up to a maximum of five extent pairs. The first track must be lower or equal to the last track value in the extent pair. All tracks within the extent pair range become part of the FlashCopy relationship.

If you do not specify this parameter, the FlashCopy FCESTABL command defaults to all tracks of the source and target volumes (full volume).

This is the content of each extent pair:

- X'ccccchhhh' X'ccccchhhh' where 'cccc' is the hexadecimal cylinder number and 'hhhh' is the hexadecimal head number. The first value is the beginning location and the second value is the ending location.

  Example: A full 3390-3 volume extent pair is specified as X'00000000' X'0D0A000E'.

  Note: The initial release of FlashCopy allows one source device and one target device in a FlashCopy relationship. The source device extents and target device extents must be identical.

  In the FCESTABL TSO command, the EXTENTS option specifies source volume begin and end extent pairs that are copied to the target volume begin and end extent pairs. This allows for a greater number of extents to be specified in the TSO command. Additionally, the source and target extent pairs may specify different locations on the source and target volumes. The total number of tracks in the relationship must be the same.

  The FCESTABL TSO command changes to require a source begin and end extent pair followed by a target begin and end extent pair, so that the maximum number of extents possible on one FCESTABL command is two. Additional FCESTABL commands to the same source or target are needed if additional extents in a FlashCopy relationship are desired.

FCESTABL command examples

The following are examples of the FCESTABL command:

```
FCESTABL SDEVN(X'ssss') TDEVN(X'tttt') MODE(NOCOPY) EXTENTS(X'00000000' X'0100000E')

FCESTABL SDEVN(X'ssss') TDEVN(X'tttt') MODE(COPY) ONLINTGT(YES)
```

1. The first command establishes a FlashCopy relationship between the source device, X'ssss', and the offline target device, X'tttt'. No background copy is performed. The FlashCopy relationship pertains to track X'00000000' (the first track) through track X'0100000E'.

   Note: The target device is offline because the ONLINTGT parameter has not been specified. The relationship will not be established if the target device is online and the ONLINTGT parameter has not been defined.
2. The second command establishes a FlashCopy relationship between the source device, X'ssss', and the target device, X'tttt'. The system will perform a background copy. This relationship can be established if the target device is online, due to the ONLINTGT(YES) parameter setting.

FlashCopy query (FCQUERY) command

Use the FCQUERY TSO command to determine the copy status of any DASD device available to the host system. If the command is issued to a device not in an ESS, certain data will not necessarily be available. The control unit information reflects the storage control type, and whatever information is available about the device will be returned.

When you issue a FCQUERY command to an ESS device that is part of a FlashCopy relationship, the specified device information includes whether the device is the source or target of a FlashCopy relationship. As much information about the related device as is available is returned. Because status information is extremely volatile in a shared device environment, these conditions may not exist when the next FlashCopy command is issued.

FCQUERY command syntax

The syntax of the FCQUERY command is:

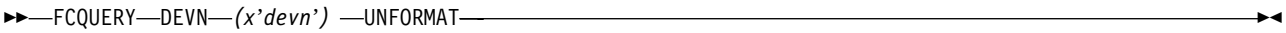

Result: The output format shown assumes that the DEVN is in a FlashCopy relationship, the DEVN identifies the source volume of a FlashCopy pair, and the UNFORMAT keyword was omitted.

```
ANTF0090I FCQUERY Formatted
DEVN SSID LSS CCA CU   SERIAL STATUS
0A4D 2830  03 0D 2105 0000325476 FC . . . . . 88%
---- 2830  03 07 2105 0000325476 FC . . . . . .
```

The device information in the DEVN field for a FlashCopy pair indicates whether the FCQUERY DEVN is the source or target device. The first output line displays device information about the source device, and the second output line displays device information about the target device. If the FCQUERY DEVN device is the source device, the DEVN displays on the first output line. If the FCQUERY DEVN device is the target device, the DEVN displays on the second output line.

The status area could have any of the following information:

SIMPLEX The volume is not in any copy mode.

CC The volume is in Concurrent Copy mode.

XRC The volume is an XRC primary volume.

PPRC The volume is in a PPRC pair.

FC If this status is displayed on the first output line, the FlashCopy was established with the NOCOPY option. The second output line always has this value.

FC....xxx% This format indicates the FlashCopy was established with the background copy option. This will be returned in the first output line only. The value is the percent of tracks copied as compared to the total tracks on the volume.

Result: The FCQUERY output for the same device situation shown above, but with the UNFORMAT keyword, would be similar to the following:

```
ANTF0095I  FCQUERY Unformatted
0A4D,2830,03,0D,2105,0000325476,FC.....88%,
----,2830,03,07,2105,0000325476,FC........
```

Result: If the FCQUERY device is not in a FlashCopy relationship, and the UNFORMAT parameter is specified, the output from the command would be similar to the following:

```
ANTF0095I  FCQUERY Unformatted
0A4D,2830,03,0D,2105,0000325476,SIMPLEX... (or some other status)
```

Required parameters

DEVN Specifies the four-digit hexadecimal device number of the volume. The query output has device and storage subsystem hardware configuration information, and the current copy services status of the specified device. If the specified device is part of an active FlashCopy relationship, then as much hardware configuration information about the related device as is available is displayed.

Optional parameters

UNFORMAT Specifies the information returned from the FCQUERY is to be formatted as a string of values separated by commas, similar to the information string returned to the ANTRQST API caller of FCQUERY. The string format is detailed in the example below. The default is formatted data.

FCQUERY command examples

The following is an example of the FCQUERY command:

```
FCQUERY DEVN(X'nnnn') UNFORMAT
```

This command queries device number X'nnnn' and generates unformatted output.

Note: If device X'nnnn' is in a FlashCopy relationship, it will be the source volume of that relationship. If it is not in a FlashCopy relationship, the generated output will state that the volume is in SIMPLEX status.

For more information on the volume states, see “XRC volume terminology” on page 24.

FlashCopy withdraw (FCWITHDR) command

The FCWITHDR TSO command is used to end a FlashCopy relationship. This command stops a FlashCopy operation with background copy active, or a FlashCopy with no background copy active. The FCWITHDR command causes all extents to be withdrawn. The source and target devices must correctly identify the devices in the original FCESTABL command, or the FCWITHDR command will fail.

FCWITHDR command syntax

The syntax of the FCWITHDR command is:

Required parameter

TDEVN Specifies a four-digit hexadecimal device number of the target device.

Optional parameter

SDEVN Specifies a four-digit hexadecimal device number of a source device in an active FlashCopy relationship.

FCWITHDR command examples

The following is an example of the FCWITHDR command:

```
FCWITHDR SDEVN(X'ssss') TDEVN(X'tttt')
```

This command ends the FlashCopy relationship between the source device, X'ssss', and the target device, X'tttt'. These must be the same device numbers as those used in the FCESTABL command, or the FlashCopy relationship will continue to be active.

Part 5. Snapshot Copy

© Copyright IBM Corp. 1994, 2002

Chapter 20. What Is Snapshot Copy?

This section explains what the Snapshot copy function is and how it works, and describes the various hardware and software products that Snapshot copy requires.

Snapshot copy is an Advanced Copy Services function that can quickly move data from the source device to the target device. DFSMSdss calls Snapshot copy when you request the function for data on hardware that supports the Snapshot capability. The two types of Snapshot copy includes:

Native Snapshot
Data is "snapped" (quickly copied) directly from the source location to the target location. This function occurs when you issue a DFSMSdss COPY command to copy volumes, tracks, or data sets from one DASD volume to another. DFSMSdss uses this method whenever the source and target data are located on like devices in the same partition on the same RVA subsystem, and no reblocking is required. DFSMSdss can use this method whether or not the CONCURRENT keyword is specified. With "native" snapshot, the copy of the data is logically and physically complete as soon as the snap is complete.

Virtual-concurrent-copy Snapshot
Data is snapped from the source location to an intermediate location and then gradually copied to the target location using normal I/O methods. This function occurs when you issue either the DFSMSdss COPY or DUMP command and you specify the CONCURRENT keyword. As the name implies, this method operates almost identically to existing concurrent copy (CC) support.

All DFSMSdss users and callers of the DFSMSdss API (such as DFSMShsm, DB2®, and IMS®) can continue to use the CONCURRENT keyword and receive functionally identical CC support. Using virtual-concurrent-copy snapshot, the copy or dump of the data is logically complete after the source data is snapped to the intermediate location, and then physically complete after the data has been moved to the target media. It is also possible to perform concurrent copy on VM minivolumes using virtual-concurrent-copy snapshot. VM minivolumes are supported if you are using RVA devices to the extent that these devices are supported by IBM Extended Facilities Product (IXFP) device reporting.

© Copyright IBM Corp. 1994, 2002

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Using Snapshot copy	280
Snapshot copy requirements	280
Using Snapshot copy, PPRC, and RVA together	283
Snapshot SMF information	284

Snapshot copy requirements

There are both software and hardware prerequisites for snapshot copy operation.

Determining Snapshot copy software requirements

DFSMS/MVS Version 1 Release 3 and subsequent releases provide snapshot copy support. You must also have the Snapshot 1.2 and IXFP 2.1 program products installed. The SIBFVP and SIBLLAPI load modules must be in a library in the LINKLST concatenation.

Determining Snapshot copy hardware requirements

MVS/ESA SP 4.3 runs on the following processors:

- All ES/9000 models that support ESA/390
- All ES/3090-9000T models that support ESA/370 or ESA/390
- All ES/3090 E, S, J, and JH models at the SEC level for ESA/370

All RVA storage subsystems support snapshot copy.

Snapshot copy operates with both 3380 and 3390 devices, but does not operate with 3390 devices in 3380 track-compatibility mode.

Making cache available

To use snapshot copy with a particular cached storage subsystem, the subsystem cache must be active. The statuses of volume-level caching, nonvolatile storage (NVS), DASD fast write, and dual copy have no effect on whether or not you can use snapshot copy. Their statuses, however, can affect the overall performance of volumes that have active snapshot sessions.

Using Snapshot copy

This section describes the Snapshot copy operation and explains how the hardware and software components work together.

Steps for initializing sessions

Functionally, the data that you are dumping or copying is unavailable only long enough for DFSMSdss to initialize a snapshot copy session. Initialization takes only a matter of seconds to complete successfully. The session initialization process is as follows:

1. Determine an appropriate time to start a backup of databases, and stop update activity to those databases.

2. Use DFSMSdss to perform a copy or dump of the data with snapshot copy. You can use DFSMSdss directly from JCL or from an application through the DFSMSdss application interface.

Normal application activity resumes as soon as the snapshot copy session initialization is complete.

Steps for providing a Snapshot working-space data space

Perform the following steps to provide a snapshot working-space data space:

1. The virtual-concurrent-copy snapshot function requires some dedicated space on one or more volumes in the same RVA subsystem as the source data set. Ensure that adequate space is available before using the CONCURRENT keyword on DFSMSdss commands that refer to data sets on an RVA subsystem. Provide the necessary space by allocating dedicated data sets, called working-space data sets, with the following naming convention:

   `SYS1.ANTMAIN.S`*`sysname`*`.SNAP`*`nnnn`*

   Variable *sysname* is the system identifier, and *nnnn* is a four-digit decimal number in the value range 0001–9999. If the system identifier is eight characters, 'S' replaces the first character.

---

2. Ensure that you catalog the working-space data sets. DFSMSdss performs a numerically-sequential catalog search for each data set, starting with SYS1.ANTMAIN.S*sysname*.SNAP0001, until it encounters a catalog locate error. A catalog locate error indicates that DFSMSdss could not find the data set. DFSMSdss cannot use data sets as working space that are beyond the data set that DFSMSdss was unable to find.

---

3. Allocate the working-space data sets as physical sequential, nonextended-format data sets. The data sets may be SMS-managed or non-SMS-managed. If the catalog search for a working-space data set indicates that the data set is multivolume, DFSMSdss will not use it as a working-space data set.

---

4. The System Data Mover (SDM) does not extend a working-space data set. A user must extend the data set by filling it with data prior to starting the DFSMSdss job if they want to allocate secondary space. SDM holds an enqueue for the data set while the working space is being used during a snapshot copy operation. It then releases the enqueue after all usage of the data set is completed. You can reallocate or extend a working-space data set when SDM does not have the data set enqueued. SDM will use the new size of the data set. You may add additional working-space data sets after ANTMAIN has completed the initialization process.

---

5. SDM uses these data sets the first time an out-of-working-space condition is encountered during a snapshot copy operation. When this condition occurs, SDM refreshes the list of working-space data sets by performing a catalog search that starts with SYS1.ANTMAIN.S*sysname*.SNAP0001.

   The LRECL and block size may be any valid combination. The system uses the tracks within the data set as the target of snapshot copy operations. Therefore, do not try to access these tracks with normal data access methods.

Recommendation: It is recommendd that you protect the working-space data sets by security, such as RACF, to ensure that unauthorized users do not have access to sensitive data.

---

6. Allocate data sets on a volume in each RVA subsystem that you have set aside for virtual-concurrent-copy snapshot. If the RVA subsystem has more than one defined device type, you must allocate a working-space data set on each device type that contains a data set that DFSMSdss snapshot copy will process. Device types are, for example, 3390 and 3380. Different models within each device type do not require separate working-space data sets.

   More than one system can simultaneously access data on the same RVA subsystem. If each system has DFSMSdss jobs that specify CONCURRENT, then each system requires at least one working-space data set.

   Example: DFSMSdss must allocate three working-space data sets to process data on an RVA subsystem from three MVS systems. The working-space data sets must include at least one of each device type that contains data that snapshot will process.

   If there is insufficient working space, the snapshot initialization for one or more data sets in the job fails. To avoid this, determine the largest total amount of data that you plan to process in a single DFSMSdss COPY or DUMP command on that RVA. The total size of all working-space data sets that you have allocated on each RVA subsystem should equal or exceed the determined amount of space.

   Note: Do not place any working-space data set on a PPRC primary volume.

Copying data with Snapshot copy

Once the initialization is complete, DFSMSdss automatically employs the snapshot copy function for any data movement operation where the CONCURRENT option applies.

Moving data sets with Snapshot copy

DFSMSdss prefers to use snapshot copy to quickly move the data from the source device to the target device. DFSMSdss can do this when the source and target devices are both in the RVA, and DFSMSdss does not need to manipulate data. This process, called native snapshot support, is much faster than traditional methods, especially when moving large amounts of data.

Native snapshot has the following requirements:

- Source and target device types (3390 or 3380, for example) must be the same.
- Source and target devices must be in the same RVA.
- Data must not require manipulation, such as reblocking.
- The data movement must not require utilities.

If the source data is in an RVA, DFSMSdss attempts to allocate the target data set on the same device type in the same RVA. This increases the probability that native snapshot can copy the data. DFSMSdss cannot allocate the target to use native snapshot if the source data set is multivolume, but is not entirely within one partition of an RVA subsystem. DFSMSdss allocates these data sets to whatever volumes are available, irrespective of their snapshot capability.

DFSMSdss automatically performs native snapshot if possible, even when you specify the CONCURRENT keyword.

Using Snapshot copy, PPRC, and RVA together

The RAMAC Virtual Array (RVA) storage subsystem provides the unique ability to combine snapshot copy functions with PPRC functions. For instance, you can use PPRC to migrate data to a different RVA subsystem; you can end the pairs and use snapshot copy at the secondary subsystem. This ability can be very beneficial for Year 2000 testing and data mining applications. However, there are some operational considerations regarding the interaction of snapshot copy with volumes that are part of a PPRC pair.

- You cannot snap (copy) data onto any volume that is part of a PPRC pair. A primary volume cannot be the target of a snapshot copy command even if it is in a suspended state.
- The primary device can be the source for a snap without having to change the state of the primary volume.
- Snapshot copy requires that space be available on the same RVA subsystem where the source of a snap resides. If you plan to use snapshot copy on an RVA subsystem that already has PPRC pairs, then provide volumes on that subsystem that are not part of a PPRC pair. You cannot, therefore, assign all of the volumes in a subsystem to be part of PPRC pairs and still use snapshot copy on that subsystem.
- If you plan to use snapshot copy at the secondary RVA subsystem, ensure that space is available at that subsystem. This real space must be independent of PPRC pairs. You cannot, therefore, have all 256 volumes as part of PPRC pairs.
- Do not put PPRC volumes into a snapshot target storage group. Also, do not include a volume that is part of a PPRC pair within a snapshot volume list. Snapshot does not check PPRC status. The snapshot copy operation will fail if you attempt to allocate any part of a data set to a volume that is part of a PPRC pair. Snapshot copy uses all volumes available to it on a space-available basis when it dynamically allocates space for the target copy.
- You cannot perform any read or write I/O to a secondary volume of a PPRC pair. However, you can snap copies of data from the secondary device and bring the PPRC pairs back to active status. The following procedure outlines a quick way to do this:

  1. Suspend the pair.

     _____

  2. Recover the secondary, returning it to simplex.

     _____

  3. Snap the secondary.

     _____

  4. Issue a CESTPAIR command with the RESYNC parameter to copy only the changes.

     _____

  You can use a similar procedure in conjunction with the CGROUP command (MVS TSO command only) to snap data from secondary devices that contain related information.

  1. Issue a CGROUP command with FREEZE to suspend all pairs on the path between a primary logical storage subsystem and a secondary logical storage subsystem.

2. Recover the secondary volumes, returning them to simplex.

3. Snap the secondary volumes.

4. Issue a CESTPAIR command with the RESYNC parameter to copy only the changed tracks of each volume.

Snapshot SMF information

The SDM writes an SMF type 42 subtype 4 record that contains session statistics. Among other things, you can use the information in this record to determine the following:

- Session initialization time
- Maximum size of host and storage subsystem sidefiles
- Number of intercepted writes.

The SMF type 42 subtype 4 record is described in detail in “Concurrent copy information in SMF type 42 records” on page 309.

Part 6. Concurrent Copy

© Copyright IBM Corp. 1994, 2002

Chapter 21. What Is Concurrent Copy?

This section explains what concurrent copy is and how it works, and describes the hardware and software requirements necessary to use concurrent copy.

Concurrent copy is a storage subsystem extended function that allows you to generate a copy or a dump of data while applications are updating that data. Other tools are available that allow you to dump databases while applications continue to update the databases. Concurrent copy, however, provides a copy of the data, in a consistent form, as it existed before the updates took place. Concurrent copy relies on cached storage subsystem Licensed Internal Code (LIC) and enhanced components of z/OS.

In this chapter

The following sections are included in this chapter:

Topic . . .	See page . . .
Understanding the concurrent copy solution	287
Using concurrent copy	290
Concurrent copy requirements	288
Determining concurrent copy resource needs	288
SMF information	291
Scheduling concurrent copy operations	291

Understanding the concurrent copy solution

Concurrent copy provides point-in-time data consistency, which is not possible with other online data dump techniques. The system serializes access to the data that you are dumping or copying just long enough for the concurrent copy session to initialize. This serialization takes a matter of seconds. Compare this to the quiesce and backup technique, which makes data unavailable for the entire duration of the dump (possibly hours).

The copy is *logically* complete as soon as you have initialized the concurrent copy environment. At that point, concurrent copy protects the original state of the data. After logical completion, the data is once again available for unrestricted application access. The copy is physically complete once the concurrent copy process finishes copying the data to the output device.

You can use concurrent copy to back up any data that can be backed up using DFSMSdss (or DFDSS) because DFSMSdss is the external interface to concurrent copy.

For example, you can back up IMS databases by using DFSMSdss. During the recovery process, IMS database recovery control (DBRC) coordinates recovery of the DFSMSdss dump and the application of updates from the IMS log. Conversely, DB2 (before Version 3) does not support those activities. Consequently, do not use concurrent copy with DB2 releases earlier than Version 3.

© Copyright IBM Corp. 1994, 2002

Concurrent copy requirements

This section describes the necessary software and hardware prerequisites for using concurrent copy.

Determining concurrent copy software requirements

All supported DFSMS/MVS and z/OS releases provide concurrent copy support. Concurrent copy requires the following components:

- DFSMSdfp™ for the SDM, new management and storage class attributes, and ISMF support of those attributes
- DFSMSdss for the external interface to concurrent copy
- DFSMShsm to call concurrent copy during backup operations.

SMS is not necessary in order to use concurrent copy. You can use concurrent copy directly from DFSMSdss without SMS. However, DFSMShsm supports concurrent copy during automatic and aggregate backup only when SMS is active.

Concurrent copy with shared storage subsystems

In configurations where a single cached storage subsystem is attached to multiple host systems (running either in separate processors or separate logical partitions), any DFSMS/MVS or z/OS system can exploit concurrent copy. Although systems that are not running DFSMS/MVS or z/OS cannot initiate concurrent copy operations, such systems can share a cached storage subsystem (including volumes with active concurrent copy sessions) with DFSMS/MVS or z/OS systems provided that the non-DFSMS/MVS or z/OS systems support 32-byte sense information.

Determining concurrent copy hardware requirements

All cached storage subsystems provide support for concurrent copy.

Concurrent copy operates with both 3380 and 3390 devices, but does not operate with 3390 devices in 3380 track compatibility mode.

To use concurrent copy with a particular cached storage subsystem, the subsystem cache must be active. The statuses of cache, nonvolatile storage (NVS), DASD fast write, and dual copy have no effect on your use of concurrent copy. Their statuses, however, can affect the overall performance of volumes with active concurrent copy sessions.

Note: Be aware that if you disable cache on a storage subsystem that has active concurrent copy sessions, the cached storage subsystem will end the concurrent copy sessions. The operations that were using the sessions will also fail.

Determining concurrent copy resource needs

A concurrent copy operation requires the system resources that are described in this section.

Central and expanded storage

The SDM uses central storage for the host sidefiles that it maintains for concurrent copy sessions. A sidefile is a storage area that is used to maintain copies of tracks within a concurrent copy domain. A concurrent copy operation maintains a sidefile in

storage subsystem cache and another in processor storage. The MVS real storage manager may page sections of the sidefile into expanded storage in response to varying system loads.

The size of the host sidefile is a secondary consideration when using concurrent copy. The SDM uses the sidefile only to store temporary copies of tracks as part of the intercepting write process. The size and use of the sidefile relate directly to the number of intercepted writes. The size of the host sidefile is unlikely to be a significant factor unless you use a concurrent copy operations during a period of very high update activity or in a very storage-constrained environment.

The SDM creates one host sidefile for each concurrent copy session. The host sidefiles for different concurrent copy sessions share a single data space, which can grow to a maximum size of 2 GB. The actual size of the data space is, however, dynamic and changes in response to the SDM requirements. The size depends on the following considerations:

- The number of active concurrent copy sessions
- The level of update activity
- The rate at which DFSMSdss reads data and writes it to the output media

In the unlikely event that the sidefiles will fill the data space, the SDM creates additional data spaces as necessary.

Another factor that affects the host sidefile size is the distribution of application writes across a device. I/O operations on a device typically concentrate on a set of tracks rather than uniformly distributing updates over the device. An application that updates a track causes an intercepted write. Subsequent updates to the same track do not result in an intercepted write. Similarly, as soon as DFSMSdss has processed a particular track, updates to that track do not result in an intercepted write. Concurrent copy only uses the host sidefile when there are intercepted writes. Therefore, the grouping of I/Os by applications significantly reduces the potential maximum size of the host sidefile.

Storage subsystem cache

Concurrent copy, by design, minimizes the amount of cache storage that is used for sidefiles. Under normal circumstances, a concurrent copy session uses less than 0.5 MB of cache. It is very unlikely that using a few simultaneous concurrent copy sessions will significantly affect the hit ratios that you achieve with a cached storage subsystem. It is, however, possible that the storage used by a concurrent copy session could have an effect on hit ratios. This is especially likely in cases where the amount of cache is only barely adequate for normal requirements. In those cases, you should increase the size of the cache regardless of whether or not you use concurrent copy.

The following are several tactics that concurrent copy uses to minimize the amount of utilized cache:

- When reading data directly from DASD, the SDM and DFSMSdss inhibit the loading of tracks into the cache.
- When the cache sidefile holds no tracks, it occupies no cache space. A sidefile occupies cache storage only when it contains copies of tracks that are saved by the cached storage subsystem during processing of intercepted writes.
- The SDM reads and deletes tracks from a cache sidefile as soon as it detects that they are present.

The cached storage subsystem monitors the size of the cache sidefiles that an active concurrent copy sessions uses. In the highly unusual case of a system failure or software problem preventing the SDM from reading tracks, the storage subsystem cancels active concurrent copy sessions when the sidefiles occupy more than half of the cache. The storage subsystem cancels sessions by starting with the session with the largest cache sidefile. This action preserves cache storage for use by sharing systems, but does not affect data integrity. Data is safe because the sidefiles contain only images of tracks before an update was applied. When concurrent copy is active, the subsystem maintains the same level of integrity for update operations that it maintains when concurrent copy is not active.

Channel paths

Because concurrent copy operations generate additional I/Os to the volumes involved, the use of concurrent copy always increases channel utilization. The impact of using concurrent copy can vary, and depends on the following conditions:

- The number of simultaneous concurrent copy sessions.
- The DFSMSdss OPTIMIZE keyword.
- The number of paths available.

Using concurrent copy

This section describes concurrent copy operation and explains how the hardware and software components work together.

Steps for initializing sessions

Functionally, data being dumped or copied is unavailable only for the time it takes for DFSMSdss to start a concurrent copy session. Initialization takes only a matter of seconds to complete successfully. The session initialization process is a series of distinct steps, as follows:

1. Determine an appropriate time to start a backup of databases, and stop update activity to those databases.

2. DFSMSdss uses concurrent copy to perform a copy or dump of the data. You can use DFSMSdss directly from JCL or from an application through the DFSMSdss application interface.

3. DFSMSdss locates the data sets to process and breaks them down into a hierarchy of volumes and extents.

4. DFSMSdss interfaces with the system data mover (SDM) to initialize a concurrent copy session.

5. The SDM determines which storage controls are attached to the affected volumes and defines a concurrent copy session to each storage control. The SDM also tells each storage control the volume and the extents on each volume that are included in this concurrent copy session. Once the SDM has completely initialized the session, it returns control to DFSMSdss.

6. DFSMSdss checks all SDM return codes. DFSMSdss releases serialization on those data sets or volumes that were successfully included in the concurrent copy session. DFSMSdss keeps serialization on data sets or volumes that are not included in the session.

---

7. DFSMSdss issues messages that indicate the status of the session. DFSMSdss uses the user interaction module (UIM) to update the calling program with the status of the session, if the application interface called DFSMSdss.

---

8. The application can take appropriate action, which may include restarting applications.

Copying the data

Once the initialization is complete, DFSMSdss releases serialization and calls the UIM, as described in the previous section. DFSMSdss then starts to read the data to perform the dump or copy operation. Applications can resume processing the data while concurrent copy is making the copy. Meanwhile, the cached storage subsystems monitor the I/Os for updates to any tracks that are included in the concurrent copy session.

DFSMSdss uses the SDM to read data from the devices within the concurrent copy session. For data that was not included in the concurrent copy session, DFSMSdss reads the data itself in the same way as for a nonconcurrent copy operation. As DFSMSdss receives tracks from the SDM or reads them itself, it writes the data to the output device in the same way as for any DFSMSdss operation.

SMF information

The SDM writes an SMF type 42 subtype 4 record that contains session statistics. Among other things, you can use the information in this record to determine the following:

- Session initialization time
- Maximum size of host and cached storage subsystem sidefiles
- Number of intercepted writes.

The SMF type 42 subtype 4 record is described in detail in “Concurrent copy information in SMF type 42 records” on page 309.

Scheduling concurrent copy operations

Many factors influence when you can schedule concurrent copy operations.

Example: The structure of your overnight batch processing determines at what stages in the processing you can use concurrent copy. Similarly, availability of tape drives may restrict the intervals during which running a concurrent copy operation is feasible.

If other considerations are not a factor, use concurrent copy at times of lowest activity, especially lowest update activity. It may, however, be beneficial to use concurrent copy to back up data even during periods of higher I/O activity. In some cases, ensuring data availability may be more important than preserving levels of application performance.

Part 7. Appendixes

© Copyright IBM Corp. 1994, 2002

Appendix A. Advanced Copy Services Diagnostic Aids

The following sections describe diagnostic aids for PPRC and for the system data mover (SDM) functions XRC, snapshot, and concurrent copy. Each SDM diagnostic operation indicates which components (XRC, CC, snapshot, or the SDM itself) are applicable to that tool.

Restarting data mover address spaces (SDM, XRC, CC, Snapshot)

The XRC control address space, ANTAS000, and the concurrent copy address space, ANTMAIN, are automatically started during system IPL. The data mover automatically reinitializes either of these address spaces if the operator cancels it, or if either address space ends abnormally.

However, there are conditions that prevent the system from starting these address spaces at IPL time, or after an operator cancels the address spaces. For such cases, you can restart either (or both) of the ANTAS000 and ANTMAIN address spaces by submitting the following program:

```
//STARTANT JOB MSGLEVEL=(1,1),REGION=4096K
//STEP1    EXEC PGM=ANTSTRT
```

Diagnosing system data mover functions with the MVS MODIFY command

Use the MVS MODIFY command options that are described in this section to diagnose and repair XRC. The following command sections give descriptions for valid operations and operands:

Topic . . .	See page . . .
AUTO_READD operation (XRC)	296
DUMP operation (XRC, CC, Snapshot)	296
DVCDATA operation (XRC, CC)	297
LISTDVCS operation (XRC)	298
LISTSESS operation (XRC, CC)	298
RESTART operation (XRC)	299
RCVRSESS operation (SDM, XRC, CC)	300
SCDATA operation (XRC, CC)	300
SCTRAP operation (XRC)	301
TERMDVC operation (XRC)	302
TERMSESS operation (XRC, CC)	302
XENDDUMP operation (XRC)	304
XRCTRAP/SUPRDUMP operation (XRC)	304

Note: The publication *z/OS MVS System Messages, Vol 1 (ABA-AOM)* describes the system data mover messages (which begin with ANT) and error codes. Diagnostic procedures outlined in this section generate messages that begin with ANTX89.

© Copyright IBM Corp. 1994, 2002

If the messages are insufficient to diagnose your problem, issue the XQUERY TSO command to determine the status of the copy operation and volume pairs. If you need to report your problem to IBM for further diagnosis, please include the following information with your report:

- Merged console log from all attached systems. The SYSLOG contains a record of all XRC commands and messages, except for TSO commands.
- Merged SYS1.LOGREC data sets from all attached systems.
- For I/O related errors, a GTF trace of the failing command sequence on both the primary and secondary volumes.
- For functional errors, generate an SVC dump of the SDM address spaces ANTMAIN, ANTAS000, and ANTAS*nnn*. The value for the address space indicated by the ANTAS*nnn* variable can be ANTAS000 or any one of the address spaces numbered ANTAS001–ANTAS005.

  Note: To determine which ANTAS*nnn* address space is associated with the session, look at message ANTQ8200I, which appears in the header of XQUERY reports for the session.

Attention: The MODIFY command is abbreviated as F. The format of the MODIFY command is as follows:

```
F ANTASnnn,operation (optional_operands)
```

AUTO_READD operation (XRC)

This operation requests that the data mover automatically returns a volume pair back into the session if the error that the volume pair encountered is recoverable by reading a track from a primary volume. Such errors include No Record Found and Invalid Track Format errors, but do not include errors encountered *during* an AUTO_READD operation.

Example: Issuing the following activates the AUTO_READD function:

```
F ANTAS002,AUTO_READD ON
```

You can issue the command with the OFF option to disable the function.

DUMP operation (XRC, CC, Snapshot)

This operation obtains a dump of the XRC data mover address space and the associated trace data space. The SYS1.DUMP data set receives the output.

Examples: Examples of the dump operation follow:

- In order to dump the address space and trace data space for the XRC data mover function, use the command:

  ```
  F ANTAS003,DUMP
  ```

- In order to dump the address space and trace data space for concurrent copy, use the command:

  ```
  F ANTMAIN,DUMP
  ```

The DUMP operation does not affect the current state of the XRC data mover function. Dumping the data mover does, however, delay record update processing, and therefore increases the time until the dump operation is complete.

DVCDATA operation (XRC, CC)

This operation reports the storage control sessions that exist for a device. You may perform this operation regardless of whether XRC or concurrent copy is active on the system. You may only issue the MODIFY command with the DVCDATA operation from a system which has the ANTAS000 address space active.

Specify the device address for the device for which you want session information. The storage control Licensed Internal Code (LIC) must at least support concurrent copy.

Example: Issue the following command to list the existing sessions for the device F40:

```
F ANTAS000,DVCDATA F40
```

Result: Possible output in response to this command would be:

```
ANTX8965I F40 XA01
```

This means that storage control session 1 is an active XRC session.

The operation output contains: the system data mover message number, ANTX8965I; the specified device (in this case, it is F40); and an individual session report about each session for the specified device (in the example, it is XA01), which includes the following information:

- Type of storage control session
- Status of storage control session
- Storage control session number
- Host connection status

This information for each session has a format of *tsnnc*, where the following apply:

t C = concurrent copy session, X = XRC session.

s A = active session, S = suspended session, Q = quiesced session, T = timed-out session.

nn The storage control session number.

c The connection status of the storage control session, as known by the system that issued the command. If *c* is blank, this host processor owns the session. If *c* is “*”, another processor owns the session. If *c* is “?”, this host processor owns the session, but may or may not have access to the session.

Results: Typical responses for this command are:

- ANTX8965I F40 XA01—This means that storage control session 1 is an active XRC session.
- ANTX8965I F40 CA02—This means that storage control session 2 is an active concurrent copy session.
- ANTX8965I F40 - CA03 XQ04—This means that there are two storage control sessions: Session 3 is an active concurrent copy session, and Session 4 is a quiesced XRC session.
- ANTX8966I F40 NO SYSTEM DATA MOVER STORAGE CONTROL SESSIONS—This means that there are no sessions for the storage control that is associated with device F40.

Note: If errors occur, this command may also issue I/O error messages.

LISTDVCS operation (XRC)

This operation lists all of the devices that are part of a specific storage control session. A storage control session is comprised of all of the devices, on a particular storage control, that contain volumes that belong to a unique XRC session. The first device listed is the one that was used to initially set up the session.

The command syntax is as follows:

```
F ANTAS000,LISTDVCS dddd ss
```

The following definitions apply to the example above:

dddd The address of a device that is attached to the primary storage control you want to display. The storage control must have XRC-capable Licensed Internal Code (LIC) that supports this function.

ss The storage control session number (which you can obtain by issuing a F ANTAS000,LISTSESS *xxxx* command—see 298 for more information).

Example: You would issue the following command in order to list the devices owned by storage control session 04 of the storage control for device F40:

```
F ANTAS000,LISTDVCS 0F40 04
```

Result: A possible response to this command would be:

```
ANTX8927I 0F40 – XA04 – 25(0F40) 0F(....)
```

This means that session 04 is an active storage control session that has two devices. Device 0F40 relates to device channel connection address 25. The device that relates to device channel connection address 0F is offline.

The operation output contains: the system data mover message number, ANTX8927I; the address of the device (which is 0F40); information about the session (see “LISTSESS operation (XRC, CC)” for more information about the session status syntax); and multiple fields that provide information about each session for the specified device. The information for each session has a format of *cc(dddd)s*, where the following apply:

cc The channel connection address (CCA).

dddd The device number. If the device is offline, the value is “....”.

s An “s” indicates that the session is suspended. A blank indicates that the session is not suspended. This field is reported for ESS storage subsystems, only.

Results: Possible responses to the command include:

- `ANTX8927I 0F40 – XA04 – 25(0F40) 0F(....)`—This means that session 04 is an active storage control session that has two devices. Device 0F40 relates to device channel connection address 25. The device that relates to device channel connection address 0F is offline.
- `ANTX8928I 0F40 NO DEVICES FOR STORAGE CONTROL SESSION 04`—This means that storage control session 04 is not in the storage control.

LISTSESS operation (XRC, CC)

This operation reports the storage control sessions that exist for a storage control. You may perform this operation regardless of whether XRC or concurrent copy is

active on the system. You may only issue the MODIFY command with the LISTSESS operation from a system which has the ANTAS000 address space active.

Specify a device address for a device that is attached to the storage control for which you want session information. The storage control LIC must at least support concurrent copy.

As an example, issue the following to list the existing sessions for the storage control that is attached to device F40:

```
F ANTAS004,LISTSESS F40
```

Result: A possible response to this command is:

```
ANTX8913I F40 XA01
```

This means that storage control session 1 is an active XRC session.

The operation output contains multiple fields that provide information about each session on the specified device's storage control. The command returns the following information for each session:

- Type of storage control session
- Status of storage control session
- Storage control session number
- Host connection status

An individual session report has a format of *tsnnc*, where the following apply:

t C = concurrent copy session, X = XRC session.

s A = active session, S = suspended session, Q = quiesced session, T = timed-out session.

nn The storage control session number.

c The connection status of the storage control session, as known by the system that issued the command. If *c* is blank, this host processor owns the session. If *c* is "*", another processor owns the session. If *c* is "?", this host processor owns the session, but may or may not have access to the session.

Results: Other possible responses to this command are:

- `ANTX8913I F40 CA02`—This means that storage control session 2 is an active concurrent copy session.
- `ANTX8913I F40 - CA03 XQ04`—This means that there are two storage control sessions: Session 3 is an active concurrent copy session, and Session 4 is a quiesced XRC session.
- `ANTX8914I F40 NO SYSTEM DATA MOVER STORAGE CONTROL SESSIONS`—This means that there are no sessions for the storage control that is associated with device F40.

Note: If errors occur, this command may also issue I/O error messages.

RESTART operation (XRC)

The format of the RESTART command is as follows:

```
F ANTAS005,RESTART
```

This is a serviceability function that forces the XRC session to restart with minimal affect to the storage control session. Use this operation when it appears that the data mover has stalled or is not responding in the expected manner. The RESTART command, which you can only direct to the ANTAS*nnn* address space associated with the session you wish to restart, performs the following actions:

1. Generates an SVC dump of the ANTAS*nnn* address space.
2. Saves the software bitmaps into the state data set for all volumes in the XRC session that are attached to storage controls that support hardware bitmapping.
3. Suspends the XRC session and ends the ANTAS*nnn* address space.
4. Restarts the XRC session. When the ANTAS*nnn* address space has restarted, the data mover begins to drain updated records from the cache for storage controls that do not support hardware bitmaps.

After the XRC session has restarted, you can issue XADDPAIR commands to bring all of the suspended volumes back into the session. Storage controls that support hardware bitmapping resynchronize the volumes from the stored bitmap information.

RCVRSESS operation (SDM, XRC, CC)

This operation recovers a quiesced session from a storage control. (Alternatively, you can use the TERMSESS operation to perform this function.) The operation does not require XRC or concurrent copy to be active on the system. You may issue the MODIFY command with the RCVRSESS operation only from a system that has the XRC software installed and that has an active ANTAS000 address space. You must specify the following information:

- A device address. (The LIC in the device's storage control must at least support concurrent copy.)
- The quiesced storage control session number of the device that you want to recover.

Before you issue the MODIFY RCVRSESS command, you may wish to issue a MODIFY LISTSESS command to determine the storage control session number of the session that you want to recover. The RCVRSESS operation recovers a quiesced CC or XRC session, only.

Example: To recover quiesced storage control session 04 for the storage control attached to device F40 issue:

```
F ANTAS001,RCVRSESS F40 4
```

Result: A typical response from this command is:

```
ANTX8916I F40 STORAGE CONTROL SESSION 04 RECOVERED
```

SCDATA operation (XRC, CC)

This operation reports the status of a session on a primary storage control. You may perform this operation regardless of whether XRC or concurrent copy is active on the system. You can only issue the MODIFY SCDATA command from a system which has the ANTAS000 address space active.

The command syntax is as follows:

```
F ANTAS002,SCDATA ddd ss X
```

The following definitions apply to the example above:

ddd Specifies the address of a device that is attached to the primary storage

control you want to display. The storage control must have XRC-capable Licensed Internal Code (LIC) that supports this function.

ss Specifies the storage control session number (which you can obtain by issuing a F ANTAS000,LISTSESS *xxxx* command, see 298 for more information).

X Specifies that the channel extender status for a device be provided by way of the ANTX8970I message.

Result: If you specify the X parameter, the following example message displays the channel extender status for the device:

```
ANTX8970I F40 CHANNEL EXTENDER TYPE=CNT LEVEL=01 ALLOC=4000
    BUFFER=000B0000
```

Result: If there is no channel extender, the following message appears:

```
ANTX8970I 0F40 NO CHANNEL EXTENDER
```

Without the X parameter, the command returns a system data mover message of number ANTX8936I.

Example: If you issue the following command to list the current status of session 01 on the primary storage control that is associated with device F40:

```
F ANTAS002,SCDATA 0F40 01
```

Result: An example of the output is:

```
ANTX8936I 0F40 - SESSION=XA01 RESIDUAL=0005 NEXTDEV=0F40 MAXSIZE=2048
                 TIME(1997.234 12:24:32.273568) TIMEOUT(00.05.00)
```

The output of the `f antas000,scdata` command has been modified (with an asterisk) to show the utility device status. The following is a sample output message from the SCDATA command:

```
ANTX8936I 0F60 - SESSION=XA09 RESIDUAL=0000 NEXT DEV=0F60* MAXSIZE=E000
TIME(**********************) TIMEOUT(00.05.00)
```

Note: The asterisk after the NEXT DEV field indicates that this is a fixed utility device. Without the asterisk, the session is using a floating utility device.

SCTRAP operation (XRC)

This operation is a serviceability function that applies only to XRC support. The SCTRAP operation directs a storage control to dump its internal control queues when requested to do so by the XRC function. This information is useful for diagnosing storage control errors for storage controls that support this capability.

Example: An example of the SCTRAP operation follows:

```
F ANTAS003,SCTRAP ON
```

Specifying the ON operand allows XRC to request that a storage control dump its internal control queues when it encounters certain LIC-related errors. Specifying the OFF operand tells XRC, for certain error types, not to request that a storage control dump its internal control queues.

If you do not specify an operand, MVS displays the current status of the SCTRAP operation.

Attention: Turning SCTRAP on does not cause the storage control to dump its queues. Turning SCTRAP on only enables XRC to make the request if it encounters one of the predefined error conditions. There can be an application performance impact while the system dumps the storage control queues.

TERMDVC operation (XRC)

This operation ends the relationship between a device and an active XRC session for a storage control.

Guidelines: Follow these guidelines to use the TERMDVC operation:

- You can only issue the operation from a system that has the XRC function installed and has the ANTAS000 address space active.
- You cannot issue this operation if the device is the last device that is associated with the XRC session—the operation will not complete.
- You must issue the operation from the processor where the XRC storage control session is active, or was last active. If you cannot issue the command from that processor, use the TERMSESS operation to end the XRC session (see "TERMSESS operation (XRC, CC)").

You must specify the following information with the TERMDVC operation:

- A device address. The LIC in the device's storage control must at least support XRC and the TERMDVC function. You can use LISTDVCS to show the devices that are associated with the XRC session. The device address must be associated with the XRC session (see "LISTDVCS operation (XRC)" on page 298 for more information).
- The session number that is associated with the volume whose relationship you want to end.

Example: To end the relationship between device F40 and storage control session 4, issue:

```
F ANTAS004,TERMDVC F40 4
```

TERMSESS operation (XRC, CC)

This operation ends an active, quiesced, or timed-out XRC or concurrent copy session for a storage control. You may only issue the operation from a system that has the XRC function installed, and has the ANTAS000 address space active. TERMSESS provides four kinds of support, shown in Table 58 on page 303.

Table 58. Using the TERMESS operation

If you want . . .	Then . . .	Notes . . .
To end a single session	Specify: • A device address. For a session that is in either active status or timed-out status, you must issue the operation from the processor where the XRC session is active, or was last active. For a quiesced or suspended session, issue the operation from any processor. • The storage control session number of the session you want to end. **Example:** To end active storage control session number four, using device F40, issue: `F ANTAS005,TERMSESS F40 4`	The LIC in the device's storage control must at least support concurrent copy. The device address must belong to the session that you plan to end, and the device must be online to the processor that you are issuing the TERMSESS operation from. You can use LISTDVCS to show the devices that belong to the session.
To end all active or timed-out XRC sessions that were last active on the processor from which you issued the TERMSESS operation	Specify XRC. The TERMSESS XRC operation issues a verify prompt (REPLY Y | N) for the operation. You will also receive a verify prompt (REPLY Y | N) for each eligible storage control session unless you specify N with the operation. **Example:** For example, to end all active or timed-out XRC storage control sessions, and *not* be prompted for each individual session, issue: `F ANTAS005,TERMSESS XRC N`	At least one device that is associated with the XRC session must be online to the processor from which you issue the TERMSESS operation.
To end all quiesced or suspended XRC sessions, regardless of which processor the XRC session was last active on	Specify SUS. The TERMSESS SUS operation issues a verify prompt (REPLY Y | N) for the operation. You will also receive a verify prompt (REPLY Y | N) for each eligible storage control session, unless you specify N with the operation. **Example:** For example, to end all quiesced or suspended XRC storage control sessions with prompting for each individual session, issue: `F ANTAS005,TERMSESS SUS`	At least one device that is associated with the XRC session must be online to the processor from which you issue the TERMSESS operation.
To end all XRC or concurrent copy sessions that are in timed-out status	Specify TIM. The TERMESS TIM operation issues a verify prompt (REPLY Y | N) for the operation. You will also receive a verify prompt (REPLY Y | N) for each eligible storage control session, unless you specify N with the operation. **Example:** To end all timed-out XRC and concurrent copy storage control sessions, with prompting for each individual session, issue: `F ANTAS005,TERMESS TIM`	At least one device that is associated with the XRC session must be online to the processor from which the TERMSESS operation is issued.

XENDDUMP operation (XRC)

This is a serviceability function that forces XRC to generate a dump whenever a session-level XSUSPEND or XEND command is issued. XENDDUMP obtains diagnostic information before the ANTAS*nnn* address space ends. You might typically issue the command to capture diagnostic information during problem re-creation. Issue the MODIFY command with the XENDDUMP operation only from the ANTAS*nnn* address space. If you issue it from the ANTAS000 address space, XRC does not generate a dump.

Specifying the ON operand directs XRC to generate a dump whenever an XEND or XSUSPEND(session) command completes. Specifying the OFF operand directs XRC to not generate a dump. If you do not specify an operand, MVS displays the current status of the XENDDUMP operation.

Example: An example of this command follows:

```
F ANTAS001,XENDDUMP ON
```

XRCTRAP/SUPRDUMP operation (XRC)

The commands XRCTRAP and SUPRDUMP both perform the serviceability function that controls whether XRC generates a dump as a result of a software error.

Specifying the OFF operand on either command directs the data mover, if required, to dump to the SYS1.DUMP data set when it encounters a software error. Specifying the ON operand on either command suppresses dumps on software-related errors. The default is OFF.

If you do not specify an operand, MVS displays the current status of the XRCTRAP or SUPRDUMP operation.

Example: An example of the XRCTRAP operation follows:

```
F ANTAS002,XRCTRAP ON
```

Example: An example of the SUPRDUMP operation follows:

```
F ANTAS004,SUPRDUMP OFF
```

Peer-to-peer remote copy diagnostic aids

Peer-to-peer remote copy (PPRC) is a hardware feature of cached storage controls. TSO software operations enable you to control PPRC operations. ERP messages indicate when PPRC has suspended a volume pair, as well as the time and the reported hardware indicator for the suspension. The publication *z/OS MVS System Messages, Vol 1 (ABA-AOM)* describes the PPRC TSO and ERP messages.

If the messages are insufficient to diagnose your problem, issue the CQUERY TSO command to determine the status of the copy path and volume pair. If you need to report your problem to IBM for further diagnosis, please include the following information with your report:

- Merged console log from all attached systems. The SYSLOG contains a record of all PPRC operations and messages.
- Merged SYS1.LOGREC data sets from all attached systems.
- For I/O related errors, a GTF trace of the failing operation sequence for pairs involved with the error.

Appendix B. Configuring the Data Mover and PPRC Options on IBM 3990 and 9390 Storage Controls

Perform the following steps to configure the data mover and PPRC options on the IBM 3990 Model 6 and 9390 Storage Controls installation and configuration panels. Refer to *IBM Enterprise Storage Server Web Users Interface Guide* for information about configuring the copy services options of the Enterprise Storage Server (ESS).

1. From the START MENU, select F1: STORAGE CONTROL SERVICE MENU.

2. Next, select F3: INSTALL/RECONFIGURE MENU.

3. Next, select F4: RECONFIGURE THIS CONTROL UNIT.

4. Continue to press ENTER until screen 73S20 is displayed.

5. Select F2: NO, A VPD CHANGE IS NEEDED.

   Result: Screen 73S25 is displayed.

6. Use the keyboard arrows to position the cursor to the Data Mover line and enter a 1, as follows:

```
KEY A "1" TO SELECT EACH VPD FIELD GROUP TO BE CHANGED AND PRESS ENTER.    73S25
0:MODE/C.U.I.R./SSID   0:CHANNEL CONFIG.   0:CACHE/NVS/SCA SIZE   0:SIM OPTIONS
0:HARDWARE LEVELs  0:NVS BATTERY DATE  0:NUMBER OF DEVICES/control unit address
1:SERIAL NUMBER/DATA LINK/MODEM/ISOL. TRANS./DATA MOVER       EXIT:END
```

7. Press ENTER to view screen 73S26. Do not enter anything next to the options—press ENTER, you will exit the screen.

8. On screen 73S65, type a 1 next to Data Mover, and press ENTER.

   Result: Screen 73S67 is displayed.

9. Read the information on screen 73S67 and press F2 to continue.

   Result: Screen 73S75 is displayed.

10. On screen 73S75, select F1 for concurrent.

11. Press ENTER until screen 65S71 is displayed.

    Result: Screen 65S71 will display as follows:

```
KEY A "1" NEXT TO THE APPROPRIATE DATA MOVER OPTIONS.                 65S71
0: NEVER TIMEOUT ON XRC              0: PPRC CRIT HEAVY
0: NEVER TIMEOUT ON CONCURRENT COPY  0: PPRC CONSISTENCY GROUPING
0: FLOATING UTILITY DEVICE                          ENTER: CONTINUE
```

The data mover options that are displayed perform the following functions:

NEVER TIMEOUT ON XRC ON directs the data mover to ignore the

© Copyright IBM Corp. 1994, 2002

timeout value set for XRC. This is equivalent to setting an infinite timeout value.

NEVER TIMEOUT ON CONCURRENT COPY
ON directs the data mover to ignore the timeout value for concurrent copy. This is equivalent to setting an infinite timeout value.

FLOATING UTILITY DEVICE ON means that the floating utility device is active. OFF means that the fixed utility device is active.

PPRC CRIT HEAVY ON means that the PPRC session stops at an error if CRIT(YES) is specified.

PPRC CONSISTENCY GROUPING ON allows long-busy for the length of the PPRC Consistency Grouping Timer when an error occurs.

---

12. Type a `1` or `0` next to the desired options and press `ENTER`.
    Result: Screen 65S73 is displayed, as follows:

```
KEY THE TIMEOUT VALUE FOR THE DATA MOVER AND PRESS ENTER TO CONTINUE.      65S73
(NOTE:  THE VALUE 00000 WILL RESET THE TIMEOUT VALUE TO THE DEFAULT)

TIMEOUT VALUE(0 - 65535): xxxxx SECONDS
```

    The *data mover timer* is the timeout value for XRC, specified in seconds.

---

13. Type in a new value or zeros, then press `ENTER`.
    Result: Screen 65S74 is displayed, as follows:

```
KEY THE TIMEOUT VALUE FOR PPRC CONSISTENCY GROUPING AND PRESS ENTER.       65S74
(NOTE:  THE VALUE 00000 WILL RESET THE TIMEOUT VALUE TO THE DEFAULT)

TIMEOUT VALUE(0 - 65535): xxxxx SECONDS
```

    The *PPRC consistency grouping timer* is a value that specifies the length of time, in seconds, that the device remains in long-busy following an error.

---

14. Type in a new value or zeros, then press `ENTER`.
    Result: Screen 65S60 is displayed.

---

15. On screen 65S60, press `F2` to write the VPD changes to both storage and the hard disk.

---

16. Continue to press `ENTER` as directed by the maintenance panel until the soft copy process is complete.

---

17. Press `EXIT` until you exit the maps.

Appendix C. SMF Type 42 Records

The SMF type 42 records for extended remote copy (XRC) and concurrent copy (CC) sessions are summarized below. *MVS System Management Facilities* contains the complete layout of SMF type 42 records.

XRC information in SMF type 42 records

The system data mover (SDM) writes a System Management Facilities (SMF) record type 42 subtype 11 for extended remote copy session statistics whenever the SMF timer interval ends. This SMF XRC information can be used by the DFSMSdfp Optimizer to provide information for capacity planning and environmental tuning.

Note: XRC cannot process SMF records for physical sessions that have ended during the SMF interval. A session can be ended, for instance, due to a “lost” physical session.

The following three fields are included in the SMF type 42 record when a subtype 11 record is written:

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
36	24	SMF42XRO	4	binary	Offset to XRC statistics.
40	28	SMF42XRL	2	binary	Length of XRC statistics.
42	2A	SMF42XRN	2	binary	Number of XRC sessions.

Table 59 contains the SMF record type 42 information that applies to XRC. This information is grouped within a subtype 11 record, which is only included when there is an active XRC session. A volume must be in duplex state before SMF information is accumulated for that volume.

Table 59. Record layouts for subtype 11

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
The following relate to an XRC session:					
0	0	S42XRID	8	EBCDIC	Logical session ID.
8	8	S42XRTYP	8	EBCDIC	Session type: *XRC* = disaster recovery session, *MIGRATE* = migration session.
16	10	S42XRSSO	4	binary	Offset to first storage subsystem identifier (SSID) data section.
20	14	S42XRSSN	2	binary	Number of SSIDs for the session.
22	16	S42XRSSL	2	binary	Length of SSID data section.
24	18	S42XRVLO	4	binary	Offset to first volume data section.
28	1C	S42XRVLN	2	binary	Number of volume data sections.
30	1E	S42XRVLL	2	binary	Length of volume data sections.
The following relate to XRC SSID data:					
0	0	S42XRSNX	4	binary	Offset to next SSID data section (0 if last SSID).
4	4	S42XRSID	2	binary	SSID.

© Copyright IBM Corp. 1994, 2002

Table 59. Record layouts for subtype 11 (continued)

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
6	6	S42XRIDP	1	binary	Storage control session ID.
7	7		1		Reserved.
8	8	S42XRVSH	4	binary	Number of primary volumes associated with this SSID that are still included in the XRC session at the end of the SMF interval. This includes both volume pairs and utility volumes.
12	C	S42XRTPR	4	binary	Total number of primary data mover record updates read.
16	10	S42XRNWD	4	binary	Number of primary data mover reads with data.
20	14	S42XRNND	4	binary	Number of primary data mover reads with no data.
24	18	S42XRNLR	4	binary	Number of primary data mover record updates left to be read at the point in the SMF when information is collected.
28	1C	S42XRNFW	4	binary	Number of format writes.
32	20	S42XRNUW	4	binary	Number of update writes.
36	24	S42XRARS	4	binary	Average record update size of data processed, in bytes.
40	28		4		Reserved.
The following relate to XRC volumes:					
0	0	S42XRVLX	4	binary	Offset to next volume data section (0 if last volume).
4	4	S42XRVLV	6	EBCDIC	Volume serial number.
10	A	S42XRVAV	6	EBCDIC	Other volume serial number in pair.
16	10	S42XRVPS	1	EBCDIC	“P” if primary volume, “S” if secondary volume, “U” if utility volume.
17	11		1		Reserved.
18	12	S42XRVSS	2	binary	SSID associated with volume.
20	14	S42XRVPR	4	binary	Primary = total number of primary data mover record updates read. Secondary = total number of data mover records written.
24	18	S42XRVWD	4	binary	Primary = number of data mover reads with data. Secondary = 0.
28	1C	S42XRVND	4	binary	Primary = number of data mover reads with no data. Secondary = 0.
32	20	S42XRVFW	4	binary	Number of format writes.
36	24	S42XRVUW	4	binary	Number of update writes.
40	28	S42XRVRS	4	binary	Primary = average record size read, in bytes. Secondary = average record size written, in bytes.
44	2C	S42XRVRW	4	binary	Primary = number of read tracks. Secondary = number of write tracks.
48	30		4		Reserved.

Concurrent copy information in SMF type 42 records

The system data mover (SDM) writes a system management facility (SMF) type 42 subtype 4 record containing session statistics for each concurrent copy session when the session ends. Concurrent copy records contain the identifier *CC*.

The subtype 4 header section begins at offset 36 in the SMF record and contains the following fields. (The fields contain zeros if the SMF record was written for an extended sequential data set.)

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
36	24	SMF42CCO	4	binary	Offset to CC statistics.
40	28	SMF42CCL	2	binary	Length of CC statistics.
42	2A	SMF42CCN	2	binary	Number of CC session.

Concurrent copy SMF type 42 subtype 4 records

The concurrent copy header section starting at the offset contained in SMF42CCO provides information about the concurrent copy session:

Table 60. Record layouts for subtype 4

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
The following relate to a concurrent copy (CC) session:					
0	0	S42CCID	4	binary	Logical session ID.
4	4	S42CCRQS	2	EBCDIC	Session type: *CC* indicates concurrent copy.
6	6	S42CCTS	1	EBCDIC	Termination status of the session (*N* = normal termination, *A* = abnormal termination).
7	7		1		Reserved.
8	8	S42CCJNM	8	EBCDIC	Jobname that requested the CC session.
16	10	42CCJNO	8	EBCDIC	JES number of the job that requested the CC session.
24	18	S42CCSST	8	binary	The time (TOD timestamp) that the CC session started.
32	20	S42CCEIT	8	binary	The TOD that initialization completed for the CC session.
40	28	S42CCSET	8	binary	The TOD that the concurrent copy session ended.
48	30	S42CCSSO	4	binary	Offset to header information for the first subsystem in the CC session.
52	34	S42CCSSN	2	binary	Number of subsystems involved in the CC session.
54	36	S42CCSSL	2	binary	Length of the SSID header.
56	38		8		Reserved.
The following relate to the storage subsystem:					
0	0	S42CSNXT	4	binary	Offset to the next SSID header (0 if last SSID).
4	4	S42CSID	2	binary	SSID.
6	6	S42CSIDP	1	binary	Storage control session ID.

Table 60. Record layouts for subtype 4 (continued)

Offsets Dec.	Offsets Hex	Name	Length	Format	Description
7	7		1		Reserved.
8	8	S42CSMSF	4	binary	Maximum track threshold that was held in the storage control sidefile for the CC session.
12	C	S42CSVLO	4	binary	Offset to first volume in the CC session for this SSID.
16	10	S42CSVLN	2	binary	Number of volumes in the CC session associated with this SSID.
18	12	S42CSVLL	2	binary	Length of the volume section.
20	14		4		Reserved.
The following relate to volumes associated with the SSID in the previous section:					
0	0	S42CVLNX	4	binary	Offset to next volume for the subsystem (0 if last volume for the subsystem).
4	4	S42CVLSR	6	EBCDIC	Volume serial number for the volume.
10	A		2		Reserved.
12	C	S42CVLDV	1	binary	Device type (UCBTBYT4) for the volume.
13	D	S42CVLUA	3	EBCDIC	Unit address (UCBNAME) for the volume.
16	10	S42CVLTK	4	binary	Number of tracks DFSMSdss requested to be processed on the volume.
20	14	S42CVLRD	4	binary	Number of tracks read directly from the volume.
24	18	S42CVLRS	4	binary	Number of tracks changed since the logical complete message was issued but before DFSMSdss read the tracks (tracks buffered in a sidefile).
28	1C	S42CVLEP	4	binary	Number of I/Os used to read data from the volume for the session.
32	20		4		Reserved.

Appendix D. Advanced Copy Services Messages

You can find messages information for extended remote copy (XRC), coupled extended remote copy (CXRC), multiple extended remote copy (MXRC), peer-to-peer remote copy (PPRC), FlashCopy, SnapShot, concurrent copy (CC), and QuickCopy by:

- Reviewing *z/OS MVS System Messages, Vol 1 (ABA-AOM)*
- Using the LOOKAT tool (for more information, see “Using LookAt to look up message explanations” on page xii)

  The LOOKAT tool, which you can access by a Web site, is part of the z/OS online help facility. You may type a message number in the search field and the tool will return the message text, explanation, and the correct course of action to take.
- Accessing this Web site: http://www.storage.ibm.com/software/sms/sdm/index.htm

  This Web site reports the latest technical updates to Advanced Copy Services components XRC, CXRC, MXRC, PPRC, SnapShot, CC, QuickCopy, and FlashCopy. It references publications updates for these components that have not yet been publicized.

© Copyright IBM Corp. 1994, 2002

Appendix E. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen-readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen-readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Volume I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

© Copyright IBM Corp. 1994, 2002

Notices

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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.

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:

© Copyright IBM Corp. 1994, 2002

IBM Corporation
Information Enabling Requests
Dept. DZWA
5600 Cottle Road
San Jose, CA 95193 U.S.A.

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 information 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.

Programming Interface Information

This publication does not contain programming interface information information.

Trademarks

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

AIX
CICS
DFSMS/MVS
DFSMSdfp
DFSMSdss
DFSMShsm
DFSMSrmm
DFSORT
ESCON
IBM
IMS
MVS
MVS/DFP
MVS/ESA
MVS/XA
OS/390
Parallel sysplex
RACF
SecureWay
z/OS

NetView is a trademark of International Business Machines Corporation or Tivoli Systems Inc. in the United States, other countries, or both.

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

Glossary

This glossary defines technical terms and abbreviations used in this document. If you do not find the term you are looking for, refer to the index of this document, or view IBM Dictionary of Computing, located at: **http://www.ibm.com/networking/nsg/nsgmain.htm**.

A

ANTAS*nnn*. A generic address space identifier that refers to any one of the following five XRC session address spaces running concurrently in a single LPAR using CXRC or MXRC: ANTAS001, ANTAS002, ANTAS003, ANTAS004, or ANTAS005.

AOM. Asynchronous operations manager.

APAR. Authorized program analysis report.

asynchronous operation. A type of operation in which the remote copy XRC function copies updates to the secondary volume of an XRC pair at some time after the primary volume is updated. Contrast with *synchronous operation*.

ATTIME. A keyword for requesting deletion or suspension at a specific target time.

B

backup. The process of creating a copy of data to ensure against accidental loss.

C

cache. A random access electronic storage in selected storage controls used to retain frequently used data for faster access by the channel.

CCA. Channel connection address.

CCW. Channel command word.

CEC. Central electronics complex.

channel. (1) A path along which signals can be sent; for example, data channel and output channel. (2) A functional unit, controlled by the processor, that handles the transfer of data between processor storage and local peripheral equipment.

channel connection address (CCA). The input/output (I/O) address that uniquely identifies an I/O device to the channel during an I/O operation.

channel interface. The circuitry in a storage control that attaches storage paths to a host channel.

CLIST. TSO command list.

cluster. See *storage cluster*.

concurrent copy. A copy services function that produces a backup copy and allows concurrent access to data during the copy.

consistency group time. The time, expressed as a primary application system time-of-day (TOD) value, to which XRC secondary volumes have been updated. This term was previously referred to as “consistency time”.

control data set. A data set that contains consistent group information on the secondary volumes and the journal data set. It contains information necessary for recovery operations and acts as the table of contents for the session. The control data set keeps track of data written to secondary volumes, the location of unwritten data in the journal set, and which group to start recovery with.

control unit address (CUA). The high order bits of the storage control address, used to identify the storage control to the host system.

Note: The control unit address bits are set to zeros for ESCON attachments.

coupled extended remote copy (CXRC). An enhancement to XRC that provides synchronous copy operations in large environments, with an expanded number of primary storage controls and DASD volumes, in excess of those supported by a single data mover configuration. Installations may have configurations consisting of thousands of volumes in multiple XRC sessions, with coordination between the sessions to ensure that all volumes can be recovered to a consistent point in time. This greatly expands upon the ability of XRC to provide remote disaster recovery protection across an entire sysplex.

CUA. Control unit address.

CXRC. Coupled extended remote copy.

D

DASD. Direct access storage device.

DASD ERP. DASD error recovery procedures.

DASD fast write. A form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to the DASD. Both copies are retained in the storage control until the data is completely written to the DASD,

© Copyright IBM Corp. 1994, 2002

providing data integrity equivalent to writing directly to the DASD. DASD fast write is available with cached IBM storage controls.

DASD subsystem. A DASD storage control and its attached direct access storage devices.

data in transit. The update data on application system DASD volumes that is being sent to the recovery system for writing to DASD volumes on the recovery system.

data mover. See *system data mover*.

dependent write. An application I/O that depends upon a previous application I/O having completed.

destage. The asynchronous write of new or updated data from cache or nonvolatile storage to DASD. The fast write, dual copy, and remote copy functions destage data. See also *write hit*.

device address. Three or four hexadecimal digits that uniquely define a physical I/O device on a channel path in System/370 mode. The one or two leftmost digits are the address of the channel to which the device is attached. The two rightmost digits represent the unit address.

device ID. An 8-bit identifier that uniquely identifies a physical I/O device.

device number. Four hexadecimal digits that logically identify an I/O device in a System/370 Extended Architecture or Enterprise Systems Architecture/370 system.

Device Support Facilities program (ICKDSF). A program used to initialize DASD at installation and perform media maintenance.

DFDSS. Data Facility Data Set Services.

DFSMSdss. A functional component of DFSMS/MVS used to copy, dump, move, and restore data sets and volumes.

director. See *storage director* and *ESCON Director*.

disaster recovery. Recovery after a disaster, such as a fire, that destroys or otherwise disables a system. Disaster recovery techniques typically involve restoring data to a second (recovery) system, then using the recovery system in place of the destroyed or disabled application system. See also *recovery*, *backup*, and *recovery system*.

DRAIN. A keyword for requesting deletion or suspension when all existing record updates from the storage control cache have been cleared.

dump. A capture of valuable storage information at the time of an error.

dual copy. A high availability function made possible by the nonvolatile storage in cached IBM storage controls. Dual copy maintains two functionally identical copies of designated DASD volumes in the logical storage subsystem, and automatically updates both copies every time a write operation is issued to the dual copy logical volume.

duplex pair. A volume comprised of two physical devices within the same or different storage subsystems that are defined as a pair by a dual copy, PPRC, or XRC operation, and are in neither suspended nor pending state. The operation records the same data onto each volume.

E

ECSA. Extended common service area.

environmental data. Data that the storage control must report to the host; the data can be service information message (SIM) sense data, logging mode sense data, an error condition that prevents completion of an asynchronous operation, or a statistical counter overflow. The storage control reports the appropriate condition as unit check status to the host during a channel initiated selection. Sense byte 2, bit 3 (environmental data present) is set to 1.

Environmental Record Editing and Printing (EREP) program. The program that formats and prepares reports from the data contained in the error recording data set (ERDS).

EREP. Environmental Record Editing and Printing Program.

ERP. Error recovery procedure.

ESCD. ESCON Director.

ESCM. ESCON Manager.

ESCON. Enterprise System Connection. This is a set of IBM products and services that provides a dynamically connected environment within an enterprise.

ESCON Director (ESCD). A device that provides connectivity capability and control for attaching any two ESCON links to each other.

ESCON Manager (ESCM). A licensed program that provides host control and intersystem communication capability for ESCON Director connectivity operations.

ESS. Enterprise Storage Server.

extended remote copy (XRC). A hardware- and software-based remote copy service option that provides an asynchronous volume copy across storage subsystems for disaster recovery, device migration, and workload migration.

F

fiber optic cable. A fiber, or bundle of fibers, in a structure built to meet optic, mechanical, and environmental specifications.

fixed utility volume. A simplex volume assigned by the storage administrator to a logical storage subsystem to serve as working storage for XRC functions on that storage subsystem.

FlashCopy. A point-in-time copy services function that can quickly copy data from a source location to a target location.

floating utility volume. Any volume of a pool of simplex volumes assigned by the storage administrator to a logical storage subsystem to serve as dynamic storage for XRC functions on that storage subsystem.

G

GB. Gigabyte.

Geographically Dispersed Parallel Sysplex (GDPS). A multi-site application availability solution that provides the capability to manage remote copy configuration storage subsystems and automate Parallel Sysplex operational tasks. All GDPS functions can be performed from a single point of control, thereby simplifying system resource management. GDPS is designed to minimize and potentially eliminate the impact of any failure or planned site outage.

gigabyte. 1 073 741 824 bytes.

GTF. Generalized trace facility.

I

ICKDSF. See *Device Support Facilities program*.

identifier (ID). A sequence of bits or characters that identifies a program, device, storage control, or system.

IML. Initial microcode load.

initial microcode load (IML). The act of loading microcode.

I/O device. An addressable input/output unit, such as a direct access storage device, magnetic tape device, or printer.

IPL. Initial program load.

J

JCL. Job control language.

Job control language (JCL). A problem-oriented language used to identify the job or describe its requirements to an operating system.

journal data set. A checkpoint data set that contains work to be done. For XRC, the work to be done consists of all changed records from the primary volumes. Changed records are collected and formed into a “consistency group”, and then the group of updates is applied to the secondary volumes.

K

KB. Kilobyte.

keyword. A name that identifies a parameter in a command string. Keywords can be entered in their entirety or as abbreviations identified in the syntax diagram for the command.

kilobyte (KB). 1 024 bytes.

km. Kilometer.

L

LIC. Licensed Internal Code.

logical storage subsystem. A collection of addresses that are associated with the same logical subsystem.

LPAR. *L*ogical *PAR*tition. A logical segmentation of a mainframe’s memory and other resources that allows it to run its own copy of the operating system and associated applications. See also *virtual server.*

LSS. Logical storage subsystem.

M

master journal time. The minimum of the journal times reported by all logical sessions. This is the latest time to which all sessions can recover their data in the event of a disaster, and still provide coupled consistency. Because coupled consistency requires all volumes to be at the same time, no session is allowed to apply data that is after the master journal time.

master data set. The master data set ensures consistency among all XRC subsystems contained within the coupled XRC system

master session. A logical entity that is used to coordinate session commands and data consistency across multiple XRC sessions. A master session exists as long as there is an XRC session coupled to the master session.

MB. Megabyte.

megabyte (MB). 1 048 576 bytes.

MIP. Million instructions per second.

multiple data mover. A function of Coupled Extended Remote Copy (CXRC) that allows several logical data mover sessions that are running on independent MVS system images to be logically connected so that all volumes in all sessions are treated as if they are part of the same data mover session. This allows a data mover configuration to grow to support more volumes than can be supported by a single MVS image.

multiple extended remote copy (MXRC). An enhancement to XRC that allows you to run up to five XRC sessions within a single LPAR.

Multiple Virtual Storage (MVS). One of a family of IBM operating systems for the System/370 or System/390 processor, such as MVS/ESA.

MVS. Multiple Virtual Storage.

N

nonvolatile storage (NVS). Random access electronic storage with a backup battery power source, used to retain data during a power failure. Nonvolatile storage, accessible from all cached IBM storage clusters, stores data during DASD fast write, dual copy, and remote copy operations.

NVS. Nonvolatile storage.

O

operating system. Software that controls the execution of programs. An operating system may provide services such as resource allocation, scheduling, input/output control, and data management.

orphan data. Data that occurs between the last, safe backup for a recovery system and the time when the application system experiences a disaster. This data is lost when either the application system becomes available for use or when the recovery system is used in place of the application system.

P

partitioned data set extended (PDSE). A system-managed, page-formatted data set on direct access storage.

parallel access volume (PAV). Parallel access volume.

P/DAS. PPRC dynamic address switching.

PDSE. Partitioned data set extended.

peer-to-peer remote copy (PPRC). A hardware-based remote copy option that provides a synchronous volume copy across storage subsystems for disaster recovery, device migration, and workload migration.

pending. The initial state of a defined volume pair, before it becomes a duplex pair. During this state, the contents of the primary volume are copied to the secondary volume.

pinned data. Data that is held in a cached storage control, because of a permanent error condition, until it can be destaged to DASD or until it is explicitly discarded by a host command. Pinned data exists only when using fast write, dual copy, or remote copy functions.

port. (1) An access point for data entry or exit. (2) A receptacle on a device to which a cable for another device is attached.

PPRC. Peer-to-peer remote copy.

PPRC Dynamic Address Switching (P/DAS). A software function that provides the ability to dynamically redirect all application I/O from one PPRC volume to another PPRC volume.

primary device. One device of a dual copy or remote copy volume pair. All channel commands to the copy logical volume are directed to the primary device. The data on the primary device is duplicated on the secondary device. See also *secondary device*.

primary system. A system made up of one or more host systems that perform the main set of functions for an establishment. This is the system that updates the primary Disk volumes that are being copied by a copy services function. Also referred to as application system.

PTF. Program temporary fix.

Q

quiesce. To render a device temporarily inactive or disabled.

R

RACF. Resource access control facility.

RAS. Retrievability, availability, and serviceability.

read hit. When data requested by the read operation is in the cache.

read miss. When data requested by the read operation is not in the cache.

recover. To rebuild data after it has been damaged or destroyed.

recovery. The process of rebuilding data after it has been damaged or destroyed. In the case of remote copy, this involves applying data from secondary volume copies.

recovery system. A system that is used in place of a primary application system that is no longer available for use. Data from the application system must be available for use on the recovery system. This is usually accomplished through backup and recovery techniques, or through various DASD copying techniques, such as remote copy.

remote copy. A storage-based disaster recovery and workload migration function that can copy data in real time to a remote location. Two options of remote copy are available. See *peer-to-peer remote copy* and *extended remote copy*.

restore. Synonym for recover.

resynchronization. A track image copy from the primary volume to the secondary volume of only the tracks which have changed since the volume was last in duplex mode.

RVA. RAMAC Virtual Array Storage Subsystem.

S

SAID. System adapter identification.

SAM. Sequential access method.

secondary device. One of the devices in a dual copy or remote copy logical volume pair that contains a duplicate of the data on the primary device. Unlike the primary device, the secondary device may only accept a limited subset of channel commands.

service information message (SIM). A message, generated by a storage subsystem, that is the result of error event collection and analysis. A SIM indicates that some service action is required.

sidefile. A storage area used to maintain copies of tracks within a concurrent copy domain. A concurrent copy operation maintains a sidefile in storage control cache and another in processor storage.

SIM. Service information message.

simplex state. A volume is in the simplex state if it is not part of a dual copy or a remote copy volume pair. Ending a volume pair returns the two devices to the simplex state. In this case, there is no longer any capability for either automatic updates of the secondary device or for logging changes, as would be the case in a suspended state.

SMF. System Management Facilities.

SMS. Storage Management Subsystem.

SRM. System resources manager.

snapshot copy. A point-in-time copy services function that can quickly copy data from a source location to a target location.

SSID. Subsystem identifier.

stage. The process of writing data from a DASD to the cache.

state data set. A data set that contains status of the XRC session and of associated volumes that XRC is managing. The state data set is updated XADDPAIR, XDELPAIR, XSET, XSUSPEND, XRECOVER, or XEND command is issued, or whenever a volume state changes.

storage cluster. A power and service region that runs channel commands and controls the storage devices. Each storage cluster contains both channel and device interfaces. Storage clusters also perform the DASD control functions.

storage control. The component in a storage subsystem that handles interaction between processor channel and storage devices, runs channel commands, and controls storage devices.

STORAGE_CONTROL_DEFAULT. A specification used by several XRC commands and messages to refer to the timeout value specified in the maintenance panel of the associated storage control.

storage control session. A logical entity that is created for the purpose of processing updates to the XRC primary volumes. It is used to group sets of primary XRC volumes that are being processed by an XRC session within the storage control.

storage director. In an IBM storage control, a logical entity consisting of one or more physical storage paths in the same storage cluster. See also *storage path*.

Storage Management Subsystem (SMS). A component of MVS/DFP that is used to automate and centralize the management of storage by providing the storage administrator with control over data class, storage class, management class, storage group, aggregate group and automatic class selection routine definitions.

storage path. The hardware within the IBM storage control that transfers data between the DASD and a channel. See also *storage director*.

storage subsystem. A storage control and its attached storage devices.

string. A series of connected DASD units sharing the same A-unit (or head of string).

subsystem. See *DASD subsystem* or *storage subsystem*.

subsystem identifier (SSID). A user-assigned number that identifies a DASD subsystem. This number is set by the service representative at the time of installation and is included in the vital product data.

super consistency group. A set of consistency groups that were combined during secondary processing.

suspended state. When only one of the devices in a dual copy or remote copy volume pair is being updated because of either a permanent error condition or an authorized user command. All writes to the remaining functional device are logged. This allows for automatic resynchronization of both volumes when the volume pair is reset to the active duplex state.

synchronization. An initial volume copy. This is a track image copy of each primary track on the volume to the secondary volume.

synchronous operation. A type of operation in which the remote copy PPRC function copies updates to the secondary volume of a PPRC pair at the same time that the primary volume is updated. Contrast with *asynchronous operation*.

sysplex. A set of MVS or z/OS systems that are communicating and cooperating with each other through certain multisystem hardware components and software services, such as CXRC, to process workloads. This term is derived from “system complex”.

system data mover. A system that interacts with storage controls that have attached XRC primary volumes. The system data mover copies updates made to the XRC primary volumes to a set of XRC-managed secondary volumes.

system-managed data set. A data set that has been assigned a storage class.

T

TOD. Time of day.

Time Sharing Option (TSO). A System/370 operating system option that provides interactive time sharing from remote terminals.

timeout. The time in seconds that the storage control remains in a “long busy” condition before physical sessions are ended.

timestamp. The affixed value of the system time-of-day clock at a common point of reference for all write I/O operations directed to active XRC primary volumes. The UTC format is *yyyy.ddd hh:mm:ss.thmiju*.

TSO. Time Sharing Option.

U

Universal Time, Coordinated. Replaces Greenwich Mean Time (GMT) as a global time reference. The format is *yyyy.ddd hh:mm:ss.thmiju*.

utility volume. A volume that is available to be used by the extended remote copy function to perform data mover I/O for a primary site storage control’s XRC-related data.

UTC. Universal Time, Coordinated.

V

vital product data (VPD). Nonvolatile data that is stored in various locations in the DASD subsystem. It includes configuration data, machine serial number, and machine features.

virtual server. A logical partition.

volser. Volume serial number.

volume. The DASD space identified by a common serial number and accessed by any of a set of related addresses. See also *device*.

VPD. Vital product data.

VSAM. Virtual storage access method.

VTOC. Volume table of contents.

W

workload migration. The process of moving an application’s data from one set of DASD to another for the purpose of balancing performance needs, moving to new hardware, or temporarily relocating data.

write hit. A write operation where the data requested is in the cache.

write miss. A write operation where the data requested is not in the cache.

write update. A write operation that updates a direct access volume.

X

XDF. Extended distance feature (of ESCON).

XRC. Extended remote copy.

XRC planned-outage-capable. A storage subsystem with an LIC level that supports a software bitmap but not a hardware bitmap.

XRC session. A logical entity that is created for the purpose of processing XRC commands and coordinating the movement of data between primary and secondary volumes.

XRC unplanned-outage-capable. A storage subsystem with an LIC level that supports a hardware bitmap.

Index

Numerics

9036 ESCON Remote Channel Extender 34, 184, 187

A

accessibility 313
adding volumes
  as volume pairs 104
  in a volume group 105
  using the FULLCOPY or QUICKCOPY options 105
  using the XADDPAIR command 102
address space, XRC
  ANTAS000 22, 295
  ANTAS001 22
  ANTMAIN 295
  generating dumps of 156
  requirements 21
  starting 295
  system data mover 22
  virtual storage requirements 28
addressing the selection of utility devices 102
Advanced Copy Services functions, overview 3
alternate track restrictions 20, 182
ANTMAIN data sets 281
ANTRQST application programming interface 99
applying maintenance (XRC) 125
ATTIME parameter (XRC) 109
AUTO_READD option of MVS MODIFY command 296
avoiding potential deadlock conditions 114

B

balancing storage control configurations 196
benefits, remote copy 9
blocking option
  setting with ADDPAIR command 108

C

catalog
  copying 47, 188
  entry consistencies 236
CDELPAIR
  deleting volume pairs 231
  description 200
  example 201
  syntax 200
CDELPATH
  deleting paths 232
  description 201
  example 202
  syntax 201
central storage, CC requirements 288
CESTPAIR
  adding a volume pair 221
  description 202, 221
  example 205, 230
CESTPAIR *(continued)*
  syntax 202
CESTPATH
  description 205
  establishing paths 218
  example 207, 219, 230
  syntax 206
CGROUP
  description 208
  examples 208
  syntax 208
changing session recovery level (XRC) 129
changing the volume recovery level (XRC) 130
changing TIMEOUT for suspended sessions (XRC) 133
channel connection addresses, changing (XRC) 133
channel extenders
  functions 15
  with PPRC 185
  with XRC 32
channel usage (concurrent copy) 290
clipped to, description of 147
CLIST
  adding a path, volume pair, and querying PPRC 230
  adding an XRC volume pair 109
  deleting a PPRC volume pair 231
  deleting an XRC volume pair 110
commands
  FlashCopy
    FCESTABL 270
    FCQUERY 274
    FCWITHDR 275
  in automated procedures 41
  issued in a coupled environment
    XADVANCE 61
    XCOUPLE 63
    XEND 73
    XQUERY MASTER 80
    XSUSPEND 96
  PPRC
    CDELPAIR - deleting pairs 200
    CDELPATH - deleting paths 201
    CESTPAIR - establishing pairs 202
    CESTPATH - establishing paths 205
    CGROUP - controlling volume groups 208
    CQUERY - querying status 209
    CRECOVER - recovering data on recovery system 210
    CSUSPEND - suspending pairs 211
    overview 199
  using CLIST to add volume pairs 109
  using commands in various system modes 56
  XRC commands
    overview 55
    XADDPAIR - adding pairs 56, 104
    XDELPAIR - deleting pairs 66, 110
    XEND - ending a session 70, 125

© Copyright IBM Corp. 1994, 2002

commands *(continued)*
XRC commands *(continued)*
XQUERY - querying a session 75, 116
XRECOVER - recovering data on recovery system 82
XSET - modifying session parameters 84
XSTART - starting a session 89, 101
XSUSPEND - suspending volumes or sessions 91
common time reference (XRC) 20
comparing XRC and PPRC 12
concurrent copy
cache usage 288, 289
channel path usage 290
copying data 291
expanded storage 288
hardware requirements 288
logical completion 287
multisystem considerations 288
nonvolatile storge usage 288
overview 5
processor storage 288
processor support 288
resource requirements 288
restrictions 16
scheduling operations 291
session initialization 290
sharing subsystems 288
sidefiles 289
SMF information 309
software requirements 288
subsystem activity considerations 291
configuring
PPRC 184
storage controls 194
storage subsystem resources 194
XRC 30
connection
ESCON ports 34, 185, 187
establishing subsystem channels 34, 187
XRC Data Mover, establishing 35
considering the PPRC solution 184
control data sets
copying 47, 188
defining 42
specifying 46
controlling PPRC volume groups 208
copying
adding a volume with the default copy option 104
catalog, control data sets 47, 188
multivolume data sets 48, 193
operations, combining storage control 16
operations, XRC 101
copying data
dynamic vs. point-in-time 3
from 3380 DASD volumes to 3390 DASD volumes 12
coupled extended remote copy (CXRC)
applying software maintenance, example 159
benefits of 137
coupled HOLD status 150
coupled extended remote copy (CXRC) *(continued)*
coupling an XRC session that has been migrated 63
coupling XRC sessions, example 145
determining where to locate the system data mover 139
ending sessions 73
generating dumps of address space 156
handling a sympathetic suspension, example 152
recovering from a disaster, example 158
removing coupled sessions 156
type of environment used 13
using the XQUERY MASTER command 80
using the XSUSPEND command 96
CQUERY
description 209, 222, 227
example 210
formatted output 223, 227
options 222
path status codes 228
unformatted output 224, 227
querying paths 227
querying volumes 222
syntax 209
creating a backup copy of secondary volumes 147
CRECOVER
description 210
example 211
syntax 210
system IPL 215
volume failure states and recovery actions 234
CSUSPEND
description 211
example 214
suspending volume pair operations 230
syntax 213

D

DASD channel extenders
functions 15
with PPRC 185
with XRC 32
data migration
using remote copy 11
Data Mover
DVCDATA operation 297
establishing XRC connections 35
estimating XRC system resources 35
functions 22
hardware/software requirements for XRC 19
interacting with the storage control 22
LISTDVCS operation 298
LISTSESS operation 298
RCVRSESS operation 300
RESTART operation 299
session descriptions
active session 23, 24
inactive session 23
recovery session 23
setting VPD switch options 305

Data Mover *(continued)*
starting the address spaces 295
system configuration 31
TERMDVC operation 302
TERMSESS operation 302
data recovery
on the PPRC recovery system 233
on the XRC recovery system 161, 166
optimizing performance 195
recovery site, resuming operations 234
data sets
defining in a coupled environment 142
defining the master 144
guidelines when using journal, control, and state 144
DDR swap, compare to P/DAS 252
deadlock, avoiding 114
Define Extent access mode restrictions
for PPRC 182
for XRC 20
defining
data sets in a coupled environment 142
resource profiles in the RACF FACILITY class 41
determining which remote copy option to use 12
device
supported by PPRC 181
supported by XRC 20
DEVSERV
for PPRC 216
for XRC 136
DFSMSdss initializing concurrent copy 290
DFSMShsm and remote copy 27
diagnostics
for PPRC 304
for XRC 295
disability 313
disaster preparation (XRC) 126
disaster recovery
determining resources 32, 185
effect on primary storage subsystem 32, 186
ESCON connections 34, 187
examining from a system viewpoint 233
journal data sets
allocating 35, 45
defining 42
maximizing access 44, 52
specifying 43
recovery site capacity 33, 186
recovery system
performance 49
resource needs 34, 187
rolling disaster 233
storage subsystem channel connections 34, 187
storage subsystem, use of 33, 186
using remote copy, advantages of 10
XRC data mover connections 35
XRC system resources 35
DRAIN parameter (XRC) 109
dual copy function 16, 215, 251
DUMP option of MVS MODIFY command 296
DVCDATA option of MVS MODIFY command 297

E

Enterprise Storage Server (ESS)
using the copy services web interface 6
with FlashCopy 267
with PPRC 196
with XRC 50
error conditions, P/DAS 262
error recovery procedures (PPRC) 244
error recovery procedures (XRC) 177
ESCON
director restrictions 34, 188
establishing remote copy connections 34, 187
making changes for PPRC 215
multiple paths 221
subsystem path restrictions 182
establishing PPRC paths 218
examples
applying maintenance 125
CDELPAIR 201
CDELPATH 202
CESTPAIR 205
CESTPATH 207
CGROUP 208
changing session error recovery level 129
changing the volume recovery level 130
changing TIMEOUT for suspended sessions 133
changing volume channel connection addresses 133
CLIST
adding a path, volume pair, and querying PPRC 230
deleting a PPRC volume pair 231
deleting an XRC volume pair 110
CQUERY 210, 223, 227
CRECOVER 211
CSUSPEND 214
deleting a PPRC path 232
migrating data 127, 171
multiple paths, adding 221
P/DAS operations 263
preparing for disaster 126
RACF 41, 191
restarting an inactive XRC coupled session after an error 148
REXX 110
storage subsystem, establishing paths 218
suspending and resynchronizing a session 132
suspending and resynchronizing a volume pair 131
switching system modes 128
utility volumes 102
volume failure states and recovery actions 234
XADDPAIR 61, 108
XDELPAIR 70
XEND 73
XQUERY 80, 82
XRECOVER 84
XSET 89
XSTART 91
XSUSPEND 96
expanded storage, CC requirements 288

extended remote copy
accessing secondary volumes 124
active volume, description of 24
adding a utility volume 48, 115
adding a volume pair 56, 104, 114
applying maintenance 125
ATTIME parameter, description 109
balancing storage control configurations 52
changing session error recovery level 129
changing the volume recovery level 130
changing TIMEOUT for suspended sessions 133
changing volume channel connection addresses 133
choosing
reasons for 13
versus PPRC 12
command access, controlling 35, 40
components
data sets 25
primary system 21
storage control sessions 26
system data mover 22, 31
volumes 24
configuring 30, 32
connecting to the data mover 35
coupled environment
configuration report 122
master report 123
storage control 121
summary report 120
volume report 120
data sets
description of 25
specifying journal, control, and state data sets 42, 43
dedicating secondary volumes 26
definition of XRC 4, 9, 14
description of volumes 24, 25
diagnostic aids 295
DRAIN parameter, description 109
duplex volume pair, description of 25
ending a session 70, 125
error conditions, recovering from 175
error recovery procedures (XRC)
address space failure 178
comparing the parameters 175
environmental errors 177
processor failure 178
specifying the ERRORLEVEL parameter 58, 89, 175
XRC data set failure 178
evaluating, with DFSMShsm 27
how it differs from PPRC 12
how it works 14
implementing a configuration 32
in automated procedures 41
installing 37
introduction 13
maximum application impact, specifying 115
messages, pageable-storage 39
migrating data 127, 171
extended remote copy *(continued)*
modifying session parameters 84
overview 4, 14
pending volume pair, description of 25
performance 49
distributing workloads 52
limits on multiple sessions 20
maximizing journal data set access 44, 52
planning for XRC
bandwidth capacity 33
overall 19
preparing for disaster 126
querying sessions 75, 116
reading from active secondary volumes 124
recovering data on the recovery system 82
recovering from a disaster 10, 161
requirements
establishing secondary volumes 26
general failure recovery 177
hardware 19
host failure recovery 178
software 19, 28
storage subsystem failure recovery 177
resource access, controlling 39
restarting a session 101
restarting an inactive XRC coupled session after an error 148
restrictions 16, 20
seqcheck volume pair, description of 25
sessions, number of multiple 20
setting up the XRC environment 37
starting a session 89, 101
supported devices 20
suspended session, impact on storage control cache 113
suspended volume, description of 25
suspending and resynchronizing a session 132
suspending and resynchronizing a volume pair 131
suspending volumes or sessions 91, 112
switching system modes 128
system data mover
active session mode 23
address space 295
inactive session mode 23
recovery session mode 23
suspended session mode 24
system resources 29, 35
target volume, description of 24
TSO commands, issuing 55
uncoupled environment
configuration report 119, 123
set report 118, 121
storage control report 118, 119, 122
summary report 117
volume report 117
using commands in various system modes 56
utility volume
description of 24
specifying 48
volume group, description of 105

extended remote copy *(continued)*
volume pairs
adding 104, 114
copying multivolume data sets 48
deleting 66, 110
description of 24, 25
ICKDSF repair 39
identifying 48
volume synchronization process
consistency group time 161
timestamped writes 165
XRC commands
overview 55
XADDPAIR - adding pairs 56, 104
XDELPAIR - deleting pairs 66, 110
XEND - ending a session 70, 125
XQUERY - querying a session 75, 116
XRECOVER - recovering data on recovery system 82
XSET - modifying session parameters 84
XSTART - starting a session 89, 101
XSUSPEND - suspending volumes or sessions 91

F

FACILITY class profile names 41, 191
fixed utility devices
addressing the selection of 102
enabling support for 103
FlashCopy
commands 267
definition of FC....xxx% 274
FCESTABL 269
FCQUERY 269
FCWITHDR 269
restrictions of 271
controlling access to commands 269
determining hardware requirements 267
making cache available 267
overview 5
using the function 268
floating utility devices 103
FULLCOPY parameter 60

G

generating dumps
across multiple LPARs 157
in a single LPAR 156
Geographically Dispersed Parallel Sysplex (GDPS), description 319
glossary 317
group_name parameter (XRC) 59, 90
grouping volumes (XRC) 105

H

hardware
P/DAS requirements 252
hardware *(continued)*
PPRC
requirements 181
restrictions 182
supported devices 182
XRC
requirements 19
restrictions 20
supported devices 20
HOLD status
clearing a master session from 151
commands not allowed when master session is in 151
CXRC master session in 150
determining how master sessions enter 151
using the XCOUPLE command to release 155
using the XQUERY command to determine 150
host application requirements (XRC) 19
host error recovery procedures 243
implementing error recovery 244
preparing for error recovery 244
recovering from error conditions 175
how the XRC and PPRC options differ 12

I

IBM Storage Control
combining copy operations 16
configuring storage controls
balancing 52
dual copy function 12
establishing ESCON connections 34, 185, 187
interaction with the system data mover (XRC) 22
NVS and Cache status using DEVSERV and IDCAMS 216
recovery site 195
remote power sequence control cables 195
storage control session descriptions 26
VPD settings
for data mover and PPRC options 305
for SIM alerts 194
ICKDSF
limiting activity to PPRC volumes 191
PPRC support 181
repairing PPRC volumes 191
when to run 116
XRC procedures 39
IDCAMS
for PPRC 216
for XRC 136
IGF52xA messages (with P/DAS) 259
implementing remote copy 9
inactive sessions
restarting 147
using the ADVANCE command on 147
initializing concurrent copy 290
initializing volume pairs (XRC) 105
installing
PPRC 189
XRC 37

introduction to
concurrent copy 287
FlashCopy 267
snapshot 279
IXFP device support 279

J

JCL examples
to add a PPRC volume pair 222
to add an XRC volume pair 109
to add PPRC paths 222
to start PPRC 222
to start XRC 109
journal data sets
allocating 35, 45
bandwidth capacity 33
defining, XRC 42
guidelines 44
maximizing access 44, 52
migrating data with XRC 171
specifying 43

K

keyboard 313

L

LISTDVCS option of MVS MODIFY command 298
LISTSESS option of MVS MODIFY command 298

M

maintenance, applying (XRC) 125
manager, ESCON, making PPRC changes 215
managing
coupled XRC operations 145
PPRC operations 215
PPRC paths 221
XRC operations 136
master data set
association with the XCOUPLE command 144
changing the characteristics of 143
example of 144
placement of 144
master session
how it enters the coupled HOLD status 151
in a coupled HOLD status 150
restrictions of HOLD status 151
maximum sessions per storage control (XRC) 20
messages
for Advanced Copy Services components 311
URL for 311
migrating
advantages of using remote copy 12
data with PPRC 239
data with XRC 171
the workload with remote copy, overview 11
XRC session 63
modes, system data mover 23
MODIFY command
AUTO_READD operation 296
DUMP operation 296
DVCDATA operation 297
LISTDVCS operation 298
LISTSESS operation 298
RCVRSESS operation 300
RESTART operation 299
SCDATA operation 300
SCTRAP operation 301
TERMDVC operation 302
TERMSESS operation 302
XENDDUMP operation 304
XRCTRAP operation 304
moving data
and migrating data with PPRC 239
with the dual copy function 12
multiple extended remote copy
messages from 311
overview 3
MVS MODIFY command 296

O

operational scenarios (XRC)
applying maintenance 125
changing session error recovery level 129
changing the volume recovery level 130
changing TIMEOUT for suspended sessions 133
changing volume channel connection addresses 133
migrating data 127, 171
preparing for disaster 126
restarting an inactive XRC coupled session after an error 148
suspending and resynchronizing a session 132
suspending and resynchronizing a volume pair 131
switching system modes 128
options
CQUERY output for paths 227, 230
CQUERY output for volumes 222, 227
default copy 104
remote copy 12
overview
concurrent copy 287
FlashCopy 267
PPRC 15
snapshot 279
XRC 14

P

P/DAS (PPRC dynamic address switching)
actions for Switch Pair, and Swap 260
comparison to dual copy 251
definition of 251
device geometry considerations 260
environmental conditions 252
error conditions 262
example of operations 263

P/DAS (PPRC dynamic address switching) *(continued)*
hardware and software requirements 252
IGF520xA messages 259
operations
in nonsysplex, shared-DASD environments 254
in sysplex environments 256
overview 251
preparing for 252
PPRC path considerations 260
responding to messages 259
paths, establishing (PPRC) 205, 218
peer-to-peer remote copy
accessing resources 190
balancing storage control configurations 196
choosing
reasons for 14
versus XRC 12
command access, controlling 188
copying catalog and control data sets 188
copying multivolume data sets 193
data recovery errors, managing 236
deleting a path 232
deleting volume pair 231
diagnostic aids 304
distributing workloads 196
error recovery procedures (PPRC)
during PPRC data recovery 236
general failure 246
host failure 246
RAMAC device sparing 249
storage subsystem 247, 249
establishing paths
direct PPRC path 220
establishing the PPRC configuration 184
multiple paths 221
path status codes (CQUERY) 228, 230
storage control link 219
how it works 15
ICKDSF repair 191
identifying volume states
duplex 217, 233
pending 217
simplex 217
suspended 217
installing 189
introduction 14
managing errors during recovery 236
managing the system 215
migrating and moving
data 240
workload 239
overview 4, 15
performance 195
PPRC
CDELPAIR - deleting pairs 200
CDELPATH - deleting paths 201
CESTPAIR - establishing pairs 202
CESTPATH - establishing paths 205
CGROUP - controlling volume groups 208
CQUERY - querying status 209
peer-to-peer remote copy *(continued)*
PPRC *(continued)*
CRECOVER - recovering data on recovery system 210
CSUSPEND - suspending pairs 211
overview 199
RAMAC device sparing 249
recovering data on the recovery system
examining disaster recovery 233
system IPL volumes 215
the rolling disaster 233
recovering from error conditions
during PPRC data recovery 236
implementing recovery procedures 244
preparing for host error recovery 244
RAMAC device sparing procedure 249
recovery errors, managing 236
requirements
general failure recovery 246
hardware 181
host failure recovery 246
software 181
storage subsystem failure recovery 247
restrictions 16, 182
routing operational messages 200
setting up the PPRC environment 189
storage control VPD settings
for data mover and PPRC options 305
for SIM alerts 194
storage subsystem
channel connections 187
configuring resources 194
failure recovery procedures 247, 249
supported devices 182
suspending volume pair operations 230
using CLIST to add volume pairs 109
with DFSMShsm 27
performance
PPRC 195
XRC 49
balancing storage control 52
distributing workloads 52
maximizing access to journal data sets 44, 52
planning for remote copy
determining resources 32, 185
effect on primary storage subsystem 32, 186
ESCON connections 34, 187
journal data sets
allocating 35, 45
defining 42
maximizing access 44, 52
specifying 43
recovery site capacity 33, 186
recovery system
performance 49
resource needs 34, 187
storage subsystem channel connections 34, 187
storage subsystem, use of 33, 186
XRC data mover connections 35
XRC system resources 35
PPRC dynamic address switching (P/DAS) 251

preparing for disaster (XRC) 126
primary system, XRC 21
primary volume
  for PPRC
    identifying volume pairs 192
    operational considerations 182
  for XRC
    operational considerations 20
procedures
  automated, for PPRC TSO commands 193
  automated, for XRC TSO commands 41
  general failure recovery for PPRC 246
  host error recovery
    general failure recovery for XRC 177, 246
    host failure recovery 178, 246
    implementing recovery procedures 244
    preparing for recovery 244
    storage subsystem recovery 177, 247
  installing remote copy 37, 189
  migrating data 171, 240
  moving the workload 239
  resuming operations at recovery site 234
processor storage, CC requirements 288

Q

querying
  defaults from the state data set 118
  master session status 116
  status of an active XRC session 116
  status of each coupled session 146
  volume summary report 117
  XQUERY configuration report 119
  XQUERY summary report 117
querying an XRC session
  coupled environment
    configuration report 122
    master report 123
    storage control 121
    summary report 120
    volume report 120
  sessions, number of multiple 20
  uncoupled environment
    configuration report 119, 123
    set report 118, 121
    storage control report 118, 119, 122
    summary report 117
    volume report 117
QUICKCOPY parameter 60

R

RACF
  FACILITY class profile names 41, 191
  journal, control, and state data sets (XRC) 42
  putting TSO commands in a protected library 40, 190
RAMAC device sparing procedure (PPRC) 249
RCVRSESS option of MVS MODIFY command 300
record zero (R0) restrictions 20, 182
recovering data
  on the PPRC recovery system 233
  on the XRC recovery system 161, 166
  optimizing performance 195
  recovery site, resuming operations 234
recovery system resource needs 34, 187
remote copy
  accessing resources 190
  balancing storage control configurations 196
  benefits 9
  choosing
    reasons for 14
    versus XRC 12
  command access, controlling 188
  configurations resources 32, 185
  copying catalog and control data sets 188
  copying multivolume data sets 193
  deleting a path 232
  deleting volume pair 231
  determining resources 32, 185
  diagnostic aids 295, 304
  disaster recovery, overview 10
  distributing workloads 196
  effect on primary storage subsystem 32, 186
  error recovery procedures (PPRC)
    during PPRC data recovery 236
    general failure 246
    host failure 246
    RAMAC device sparing 249
    storage subsystem 247, 249
  ESCON connections 34, 187
  establishing paths
    direct PPRC path 220
    establishing the PPRC configuration 184
    multiple paths 221
    path status codes (CQUERY) 228, 230
    storage control link 219
  hardware requirements 181
  how it works 15
  ICKDSF repair 191
  identifying volume states
    duplex 217, 233
    pending 217
    simplex 217
    suspended 217
  implementing 9
  installing 189
  introduction 14
  journal data sets
    allocating 35, 45
    defining 42
    maximizing access 44, 52
    specifying 43
  managing the system 215
  migrating and moving
    data 240
    workload 239
  migrating workload, overview 11
  options 12
  overview 4, 15
  performance 195

remote copy *(continued)*
recovering data on the recovery system
examining disaster recovery 233
system IPL volumes 215
the rolling disaster 233
recovering from error conditions
during PPRC data recovery 236
implementing recovery procedures 244
preparing for host error recovery 244
RAMAC device sparing procedure 249
recovery site capacity 33, 186
recovery system
performance 49
resource needs 34, 187
requirements
general failure recovery 246
hardware 181
host failure recovery 246
software 181
storage 28, 30
storage subsystem failure recovery 247
resource needs 32, 185
restrictions 16, 182
hardware 16, 20, 182
software 20, 182
routing operational messages 200
setting up the PPRC environment 189
storage control VPD settings
for data mover and PPRC options 305
for SIM alerts 194
storage subsystem
channel connections 187
configuring resources 194
failure recovery procedures 247, 249
storage subsystem channel connections 34, 187
storage subsystem, use of 33, 186
supported devices 182
suspending volume pair operations 230
with DFSMShsm 27
XRC data mover connections 35
XRC system resources 35
removing
coupled sessions 156
using the XCOUPLE DELETE command 156
using the XCOUPLE PURGE command 156
HOLD status from a master session 155
XCOUPLE_FAILED status for a session 155
reports
coupled environment
configuration report 122
master report 123
storage control 121
summary report 120
volume report 120
sessions, number of multiple 20
uncoupled environment
configuration report 119, 123
set report 118, 121
storage control report 118, 119, 122
summary report 117
volume report 117
resource management facility data 216
resources needed by the recovery system 34, 187
RESTART option of MVS MODIFY command 299
restarting
an XRC session 101
inactive XRC coupled sessions 147
using the ADVANCE command 147
REXX example (XRC) 110
RVA (RAMAC Virtual Array)
with PPRC and snapshot 283
with snapshot 281

S

SCDATA option of MVS MODIFY command 300
scenarios
applying maintenance 125
changing session error recovery level 129
changing the volume recovery level 130
changing TIMEOUT for suspended sessions 133
changing volume channel connection addresses 133
migrating data 127, 171
preparing for disaster 126
restarting an inactive XRC coupled session after an error 148
suspending and resynchronizing a session 132
suspending and resynchronizing a volume pair 131
switching system modes 128
SCTRAP option of MVS MODIFY command 301
secondary volume
control data set 46
dedicating to PPRC 193
disaster recovery 233
failure states and recovery actions 234
for XRC
operational considerations 20
identifying states (PPRC) 216
parameter to copy allocated tracks of the primary volume to 60
parameter to copy each track of the primary volume to 60
protecting for PPRC usage 195
security
FACILITY class profile names 41, 191
journal, control, and state data sets (XRC) 42
PPRC 191
putting TSO commands in a protected library 40, 190
XRC 41
session modes, XRC system data mover
active 23
inactive 23
recovery 23
suspended 24
sessions
configured in a coupled XRC environment
active 138
coupled 138
inactive 138
interlocked 138

sessions *(continued)*
configured in a coupled XRC environment *(continued)*
master 138
suspended 138
XRC 138
coupling, max number 137
for XRC
operational considerations 20
storage control, states of
quiesced storage control 26
suspended storage control 26
setting up the environment
for PPRC 189
for XRC 37
sharing subsystems with concurrent copy 288
shortcut keys 313
sidefile (concurrent copy) 289
SMF information
for concurrent copy 309
for XRC 307
using the XADVANCE command for updating 147
snapshot
ANTMAIN data set 281
hardware requirements 280
I/O with PPRC secondary volumes 283
moving data 282
native support mode 282
overview 5, 9
restrictions with PPRC 283
software requirements 280
virtual-concurrent-copy mode 281
with PPRC and RVA 283
working-space data set 281
software
P/DAS requirements 252
PPRC requirements 181
XRC requirements 19
solution considerations, PPRC 184
state data set
guidelines 47
specifying 46
status of each coupled session, querying 146
STGADMIN.ANT.PPRC.COMMANDS 191
STGADMIN.ANT.XRC.COMMANDS 41
storage subsystem
capacity, recovery site 33, 186
channel connections, establishing 34, 187
concurrent copy considerations 288
configuring resources for PPRC 194
failure recovery procedures 177, 247, 249
matching the volume write rate 33, 186
primary and secondary example 218
sessions 26
use at recovery site 33, 186
use by primary volumes 32, 186
StorWatch ESS Copy Services, enabling 6
suspending and resynchronizing a session (XRC) 113, 132
suspending and resynchronizing a volume pair (XRC) 131
suspending PPRC volume pair operations 230
switching system modes (XRC) 128
synchronizing suspended state volumes 193
syntax for remote copy commands
CDELPAIR 200
CDELPATH 201
CESTPAIR 202
CESTPATH 206
CGROUP 208
CQUERY 209
CRECOVER 210
CSUSPEND 213
XADDPAIR 57
XDELPAIR 67
XEND 71, 73
XQUERY 76
XRECOVER 83
XSET 84
XSTART 89
XSUSPEND 92
SYS1.ANTMAIN.SNAPnnnn data sets 281
sysplex timer, requirement (XRC) 20
system data mover (SDM)
diagnostics 295
with snapshot 281
with XRC 22

T

telecommunication line channel extenders
functions 15
with PPRC 185
with XRC 32
TERMDVC option of MVS MODIFY command 302
TERMSESS option of MVS MODIFY command 302
tertiary volumes, using 111
TIMEOUT parameter, using for XRC 115
timestamping, correcting lack of (XRC) 21
TSO commands
in automated procedures 41
PPRC
CDELPAIR - deleting pairs 200
CDELPATH - deleting paths 201
CESTPAIR - establishing pairs 202
CESTPATH - establishing paths 205
CGROUP - controlling volume groups 208
CQUERY - querying status 209
CRECOVER - recovering data on recovery system 210
CSUSPEND - suspending pairs 211
overview 199
using CLIST to add volume pairs 109
using commands in various system modes 56
XRC commands
overview 55
XADDPAIR - adding pairs 56, 104
XDELPAIR - deleting pairs 66, 110
XEND - ending a session 70, 125
XQUERY - querying a session 75, 116
XRECOVER - recovering data on recovery system 82

TSO commands *(continued)*
XRC commands *(continued)*
XSET - modifying session parameters 84
XSTART - starting a session 89, 101
XSUSPEND - suspending volumes or sessions 91

U

uncoupling
an XRC session 96
using the XSUSPEND command 96
using
ICKDSF to change or repair volumes 39
tertiary XRC volumes 111
utility volume (XRC)
adding a new one 102
adding back a suspended one 115
deleting 110
description 24
examples 102
how to select 48
purpose 50, 102
syntax for ADDPAIR command 57

V

VM minivolume support 279
VM operations
for PPRC 183
for XRC 21
volume
alternate track restrictions 20, 182
grouping 105
managing PPRC volumes and paths
adding volume pair (CESTPAIR) 221
deleting a volume pair 231
path status codes 228, 230
querying (CQUERY) 222, 227
suspending pair operations 230
PPRC
examples of volume report output 223
ICKDSF repair 191
migrating data 240
multivolume data sets 193
pairs, identifying 192
record zero (R0) restrictions(XRC) 182
restrictions 182
synchronizing 193, 217
system IPL 215
write access mode restrictions 182
storage subsystem 32, 186
XRC
adding a utility volume 48
adding a volume pair 104
adding back a suspended utility volume 115
adding back a suspended volume pair 114
blocking options 108
dedicating secondary volumes 26
default copy option 104
deleting volume pairs 110
volume *(continued)*
XRC *(continued)*
ICKDSF repair 39
initializing pairs 105
migrating data 171
NOCOPY option 106
pairs, description of 24, 25
pairs, identifying 48
record zero (R0) restrictions 20
restrictions 20
suspended option 107
switching applications for migrating 172
write access mode restrictions 20
XRC tertiary 111
volume groups, controlling 208
VPD switch options 194, 305

W

Web interface for Copy Services 6
working-space data sets 281
workload migration, reasons for 11

X

XADDPAIR
copy options 105, 115
description 56, 104
example 61, 108
for use in data migration 171
syntax 57
XADVANCE
example 62
how it differs from the XRECOVER command 61
issuing to a master session 62
syntax 62
XCOUPLE
description 63
example 66
managing master session status 63
syntax for 64
adding XRC sessions to a master session 63
deleting XRC sessions from a master session 63
purging XRC sessions 63
removing HOLD status of the master session 63
XDELPAIR
description 66, 110
example 70
for use in data migration 172
syntax 67
XEND
description 70, 125
example 73
syntax 71, 73
using in a coupled environment 73
example 74
syntax for XEND 73
XENDDUMP option of MVS MODIFY command 304
XQUERY
coupled environment
configuration report 122

XQUERY *(continued)*
coupled environment *(continued)*
master report 123
storage control 121
summary report 120
volume report 120
description 75, 116
example 80, 82
for use in data migration 172
recover report 123
sessions, number of multiple 20
syntax 76
uncoupled environment
configuration report 119, 123
set report 118, 121
storage control report 118, 119, 122
summary report 117
volume report 117
XQUERY MASTER
example 82
syntax for 81
XRCTRAP option of MVS MODIFY command 304
XRCUTL parameter 57
XRECOVER
description 82
example 84
syntax 83
XSET
description 84
example 89
syntax 84
XSTART
description 89, 101
example 91
SESSION option 176
syntax 89
VOLUME option 175
XSUSPEND
description 91, 112
example 96
impact on cache resources 113
syntax 92
using in a coupled environment 96
example 98
syntax for XSUSPEND 97

Readers' Comments — We'd Like to Hear from You

z/OS
DFSMS Advanced Copy Services

Publication No. SC35-0428-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC35-0428-01

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

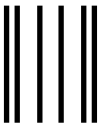

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

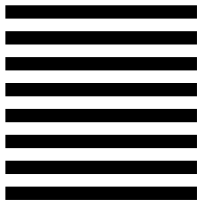

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9000 South Rita Road
TUCSON AZ 85775-4401

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Program Number: 5694-A01

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC35-0428-01