RAMAC Virtual Array: Implementing Peer-to-Peer Remote Copy

Alison Pate Adrian Butler Ulrich Gohl
Alfonso Santiago Diane Williams

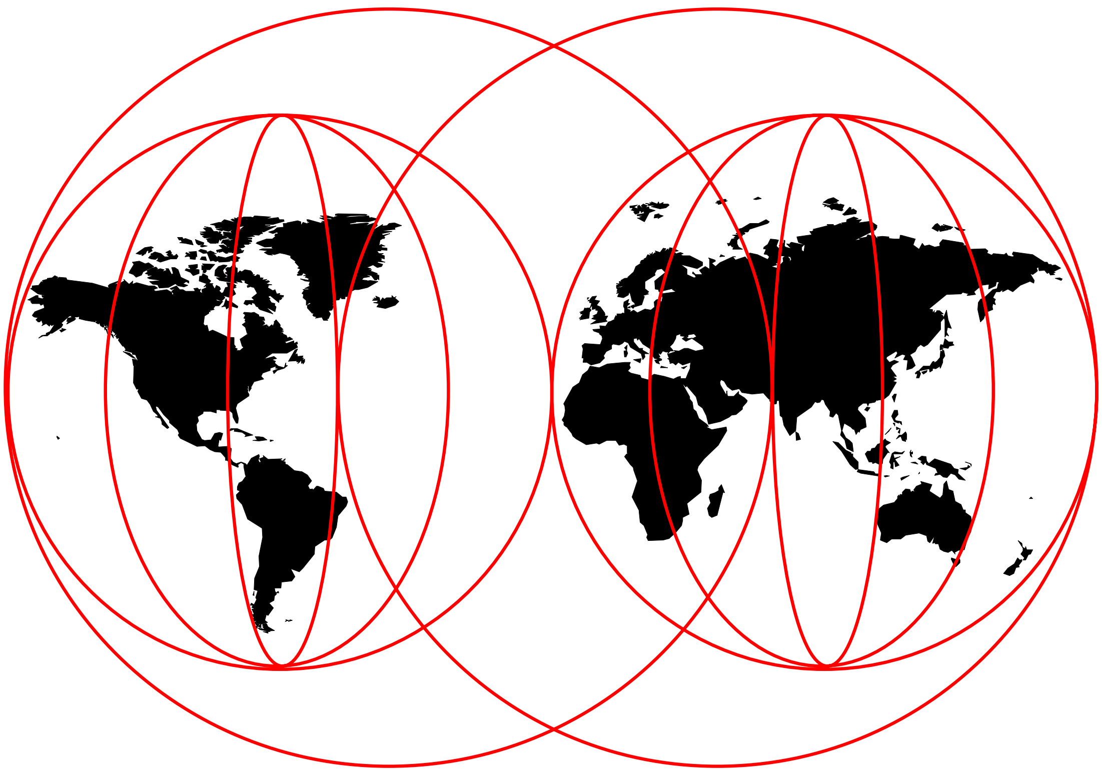

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5338-00

IBM

International Technical Support Organization

SG24-5338-00

RAMAC Virtual Array: Implementing Peer-to-Peer Remote Copy

December 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special Notices” on page 117.

First Edition (December 1998)

This edition applies to peer-to-peer remote copy (PPRC) on the RAMAC Virtual Array Model T82.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

Contents

Preface . . . v
The Team That Wrote This Redbook . . . v
Comments Welcome . . . vi

Chapter 1. Introduction . . . 1
1.1 The IBM 9393 RVA Model T82 . . . 1
1.2 Remote Copy . . . 3
1.3 Terminology Used in This Book . . . 4

Chapter 2. Planning for Remote Copy . . . 5
2.1 Characteristics of a Remote Copy Solution . . . 5
2.2 Avoiding the Rolling Disaster or Dying Scenario . . . 6
2.3 Planning for Remote Copy on the RVA . . . 7
2.4 Remote Copy Techniques on the RVA . . . 8
2.4.1 XRC on the RVA . . . 8
2.4.2 PPRC on the RVA . . . 9
2.5 Considerations on Using PPRC on the RVA . . . 10
2.5.1 Primary Device Considerations . . . 11
2.5.2 Secondary Device Considerations . . . 11
2.5.3 SnapShot Considerations . . . 11
2.6 Requirements for Implementing PPRC on the RVA . . . 11
2.6.1 Hardware Requirements . . . 11
2.6.2 Software Requirements . . . 13
2.7 Managing PPRC on the RVA . . . 13
2.7.1 Unidirectional and Bidirectional Copy . . . 14
2.7.2 Manage Snapshot and IXFP Activities . . . 14
2.7.3 PPRC Migration Manager . . . 14
2.8 Configuring PPRC on the RVA . . . 15
2.8.1 RVA ESCON Connectivity . . . 15
2.8.2 Considerations on Using the ESCON Director . . . 16
2.8.3 Point to Point or Direct Connection . . . 17
2.8.4 ESCON Director Connection . . . 18
2.9 PPRC of Open Systems Data . . . 18
2.10 Database Recovery Using PPRC . . . 19

Chapter 3. Configuration . . . 23
3.1 ESCON . . . 23
3.1.1 ESCON Connection Modes . . . 23
3.1.2 ESCON Specifications . . . 25
3.1.3 ESCON Distance Configurations with PPRC . . . 28
3.2 IBM RAMAC Virtual Array . . . 31
3.2.1 Serial Number . . . 31
3.2.2 Logical Control Unit . . . 31
3.2.3 Subsystem Identifier . . . 32
3.2.4 Channel Connection Address . . . 32
3.2.5 Director Device Connection . . . 33
3.2.6 ICE Cards . . . 35
3.2.7 System Adapter Identifier . . . 36
3.2.8 Logical Path . . . 38
3.2.9 Channel Register Sets . . . 38

3.2.10 CESTPATH Command . . . 39
3.3 PPRC . . . 44
3.3.1 Physical Configuration of the PPRC Links . . . 45
3.3.2 Logical Configuration of the PPRC Paths . . . 47

Chapter 4. Operational Considerations . . . 51
4.1 Hardware Characteristics . . . 51
4.2 Command Support . . . 52
4.2.1 CESTPATH: Establishing Paths . . . 52
4.2.2 CESTPAIR: Establishing Pairs . . . 53
4.2.3 CSUSPEND: Suspending Pairs . . . 54
4.2.4 CQUERY: Querying Status . . . 55
4.2.5 CGROUP: Controlling Volume Groups . . . 55
4.3 ICKDSF . . . 56
4.4 SnapShot . . . 57
4.5 Automation . . . 57
4.5.1 PPRC Migration Manager . . . 57
4.5.2 Error Recovery . . . 59
4.5.3 The Rolling Disaster . . . 61
4.5.4 Error Recovery in a GDPS . . . 61
4.5.5 Remote Copy Management Facility . . . 61
4.5.6 Monitoring Net Capacity Load . . . 61
4.5.7 Managing Secondary Devices . . . 61
4.5.8 Processing CQUERY Output . . . 62
4.5.9 Monitoring Pair Status . . . 63
4.5.10 Establishing PPRC Pairs . . . 65
4.5.11 The Application Programming Interface . . . 65
4.5.12 Error Situations . . . 65
4.6 Fallback Planning . . . 66
4.7 PPRC Dynamic Address Switching . . . 66
4.8 Licenced Internal Code Upgrades . . . 66
4.9 Systems Managed Storage Considerations . . . 66
4.10 Security . . . 67
4.11 PPRC and Planned System Outages . . . 67
4.12 Procedures and Documentation . . . 68
4.13 ESCON Manager . . . 68
4.14 Recommendations . . . 69

Chapter 5. Planning for PPRC Testing . . . 71
5.1 Producing a PPRC Link Failure . . . 71
5.1.1 Step-by-Step Action Plan . . . 71

Chapter 6. RVA PPRC Performance . . . 77
6.1 Performance Characteristics . . . 77
6.2 Performance Variables . . . 77
6.2.1 Configuring for Performance . . . 78
6.2.2 Distance between the RVA Primary and Secondary Subsystems . . . 78
6.3 Performance during Establish . . . 78
6.3.1 Link Recommendation to Optimize Establish Times . . . 78
6.3.2 Number of Concurrent Establish Pairs . . . 79
6.4 Workload Performance . . . 80
6.4.1 Database and Transaction Processing Workloads . . . 80
6.4.2 Sequential Workloads . . . 81
6.5 Predicting RVA PPRC Performance . . . 81

Chapter 7. Using PPRC on the RVA with VM/ESA and VSE/ESA 83
7.1 Considerations for VM/ESA 83
7.1.1 Configuring for the VM/ESA Environment 83
7.1.2 Operations for the VM/ESA Environment 84
7.2 Considerations for VSE/ESA 84
7.2.1 Configuring for a VSE/ESA Environment 85
7.2.2 Operations for the VSE/ESA Environment 85

Chapter 8. Using PPRC in a Database Environment 87
8.1 Using PPRC with IMS 87
8.2 IMS Recovery Process 88
8.3 Using PPRC with DB2 88
8.3.1 CRIT(Y) 88
8.3.2 CRIT(N) 90
8.3.3 PPRC Consistency Grouping Timer 91
8.3.4 Automation and DB2 92
8.4 Recommendations 92

Chapter 9. Using PPRC with SnapShot 95
9.1 Primary Devices of a PPRC Pair 95
9.2 Secondary Devices of a PPRC Pair 96
9.3 Using DFSMSdss 98
9.3.1 DFSMSdss SnapShot 98
9.3.2 Virtual Concurrent Copy 99
9.4 SnapShot before PPRC Resync 100
9.5 NCL on Secondary RVA Subsystems 100

Chapter 10. PPRC and SnapShot for System Testing 103
10.1 Cloning Systems with PPRC 103
10.2 SnapShot Recovery 104

Chapter 11. Geographically Dispersed Parallel Sysplex 105
11.1 Requirement for Time Consistency of Data 105
11.2 Remote Copy Management Facility 106
11.3 Planned Reconfiguration Support 107
11.4 Unplanned Reconfiguration Support 107
11.5 GDPS Service Offering 107

Appendix A. Sample PPRC Implementation Project Plan 109
A.1.1 Education on PPRC Functions and the RVA 109
A.1.2 Planning Hardware Requirements 109
A.1.3 Planning Application Recovery Requirements 110
A.1.4 Planning Database Recovery Requirements 111
A.1.5 Capacity Planning 111
A.1.6 Planned Outages 112
A.1.7 Functional Testing of PPRC 112
A.1.8 Testing Operational and Management Procedures 112
A.1.9 Testing Error Recovery Procedures 113
A.1.10 Performance Testing 113
A.1.11 Cutover to Production 113

Appendix B. Sample ESCON Configuration Worksheet 115

Appendix C. Special Notices 117

Appendix D. Related Publications 121
D.1 International Technical Support Organization Publications 121
D.2 Redbooks on CD-ROMs 121
D.3 Other Publications 121

How to Get ITSO Redbooks 123
IBM Redbook Fax Order Form 124

Index 125

ITSO Redbook Evaluation 127

Preface

This redbook provides a foundation for the implementation of peer-to-peer remote copy (PPRC) on the IBM 9393-T82 RAMAC Virtual Array (RVA). This book brings together the different papers and references documenting PPRC experiences with the IBM 3990-6 and the IBM 9390.

This redbook is intended for IBM client representatives, IBM technical specialists, and IBM customers who are planning to implement PPRC on the RVA.

In this book, we review the RVA and the principles of remote copy. We present the considerations to take into account in planning for and implementing PPRC on the RVA. We provide examples of RVA PPRC configurations and their restrictions and demonstrate how you can obtain the optimum performance in using these configurations.

We also provide pointers on the operation and management of PPRC and explain how to gain the maximum benefit from SnapShot while running PPRC on the RVA. We introduce you to Geographically Dispersed Parallel Sysplex (GDPS). Finally, we present a sample project plan that you can use as the basis for your PPRC planning and implementation.

The Team That Wrote This Redbook

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

Alison Pate is a project leader in the International Technical Support Organization, San Jose Center. She joined IBM in 1985 after completing an MSc in Information Technology. A large-systems storage specialist since 1990, Alison has eight years of practical experience in using and implementing IBM DASD solutions. She has acted as a consultant to some of the largest, leading-edge customers in the United Kingdom, providing technical support and guidance to their key business projects.

Adrian Butler is an Advisory IT Specialist in the United Kingdom. He has 18 years of experience in IT, having worked in operations, automated operations, and storage management.

Ulrich Gohl is a Senior HW Specialist in Germany. He joined IBM in 1978 as a hardware engineer. He has 16 years of practical experience in problem determination at many of the major large-systems MVS customers in Germany. Ulrich's areas of expertise include all types of ES/9000 and S/390 attachable IBM disk subsystems as well as the ES/9000 and S/390 systems themselves.

Alfonso Santiago Jr. is a Systems Services Representative in the Philippines. He has nine years of experience in Enterprise Systems platform. Alfonso holds a degree in Electronic and Communications Engineering.

Diane Williams is a Consulting IT Technical Specialist in the United States. Her background includes 2 years as an MVS and IMS systems programmer and 15 years as an IBM Systems Engineer. Diane has been in her current position supporting high-end disk products since 1993.

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

Janet Anglin, Storage Systems Division, San Jose
Robert Beiderbeck, Product Support Service, Frankfurt, Germany
Bernd Berresheim, European Product Support Group, Mainz, Germany
Don Chesarek, Storage Systems Division, San Jose
Pat Courtney, Storage Team, Marks and Spencer
Al Downs, Product Engineering, Tucson
Lloyd Johnson, Storage Systems Division, Tucson
Bob Kern, ISC Hursley, UK
Terri Leamon, Storage Systems Division, San Jose
Bruce McNutt, Storage Systems Division, San Jose
David Petersen, S/390 Parallel Sysplex, Washington
Joe Renier, Product Engineering, Tucson
Preston Rivers, Storage Systems Division, Tucson
Paul Suddath, Storage Systems Division, Tucson
Tom Taleck, Storage Systems Division, Tucson
Bill Worthington, Advanced Technical Support, San Jose

Thanks are also due to the many other collaborators and reviewers who contributed to the production of this document. In particular, we would like to thank the following people:

Eneo Baborsky, IBM Italy
Daniel Leplaideur, IBM France
Michael Milillo, Storage Technology Corporation
Carolyn Steward, IBM Australia
Peter Tallman, Storage Technology Corporation
Roland Wolf, IBM Germany

Comments Welcome

Your comments are important to us!

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

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 127 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Send your comments in an Internet note to `redbook@us.ibm.com`

Chapter 1. Introduction

This redbook provides a foundation for planning and implementing peer-to-peer remote copy (PPRC) on the IBM 9393-T82 RAMAC Virtual Array (RVA).

1.1 The IBM 9393 RVA Model T82

The IBM 9393 RAMAC Virtual Array Storage is a member of the IBM RAMAC Array Family of storage arrays. It is built on a virtual disk architecture and emulates both the IBM 3380 and 3390 direct access storage devices (DASD).

As shown in Figure 1 on page 2, all writes to the RVA are writes to both cache and NVS. If the track is in the cache (a cache hit), the RVA signals an I/O complete to the host, and the channel is released. If the track is not in cache, it is staged from the disk array. The track is then queued for transfer from cache to the disk array. The RVA uses an advanced storage subsystem architecture where logical volumes are indirectly mapped to physical volumes within the subsystem. Unlike traditional storage subsystems, the RVA maps user data across the subsystem.

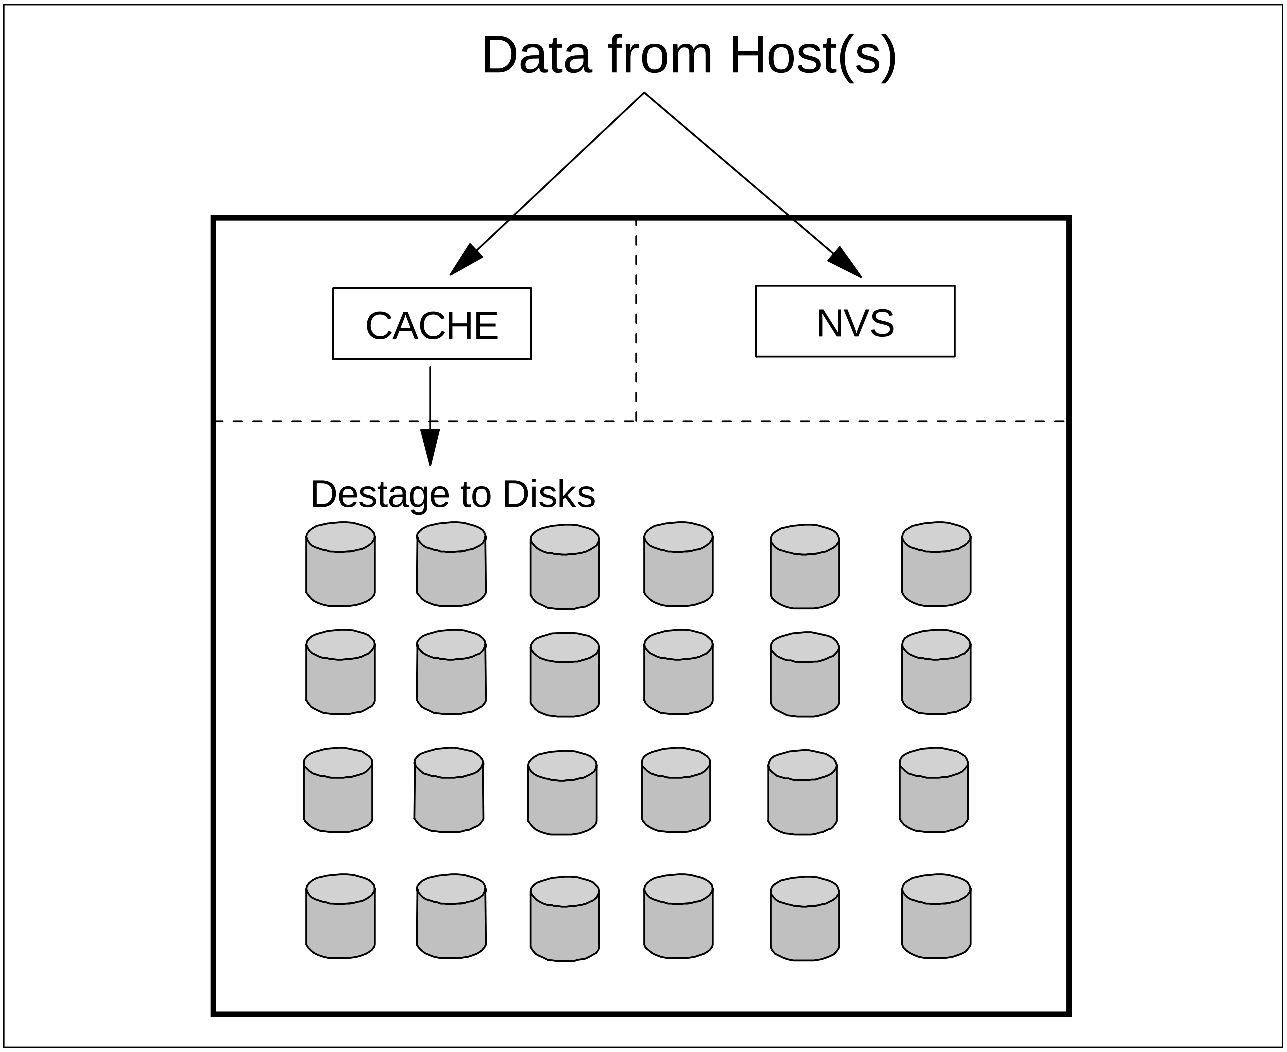

Figure 1. RVA Write Process

The RVA provides an advanced implementation of RAID 6 which delivers high availability with high performance in a variety of environments for both random and sequential workloads. The RVA offers simplified storage management and, with data compression, improves the price-to-performance ratio of your disk investment. With SnapShot, a new data duplication facility, you can realize dramatic savings in time and data storage.

The RVA 9393-T82 features the turbo shared memory for improved performance. It can have up to eight concurrent data transfers over ESCON channels and provides 128 logical paths for improved connectivity. PPRC is also supported on the RVA 9393-T81, which is a field upgrade from an RVA 9393-T43 to include the turbo shared memory and 8 paths.

The IBM Redbook, *IBM RAMAC Virtual Array*, SG24-4951 covers the characteristics and features of the RVA.

1.2 Remote Copy

Although remote copy provides a foundation on which to form your disaster recovery solution, it is also a useful tool for migration, data transfer, and data mining. In addition many people have found remote copy a valuable solution to copying production data for year 2000 testing.

The RVA PPRC data copy to the secondary subsystem is synchronous with the I/O operation of the primary subsystem. The steps, shown in Figure 2, are:

1. Write to the primary

   All writes to the RVA are writes to cache and nonvolatile storage (NVS). However, in the RVA, write data received from the channel is *compressed* before it is stored in the subsystem cache.

2. Channel end (CE)

   When the data has been written to cache and NVS, channel-end status is sent to the host.

3. Write compressed data to secondary

   As the update is completed, the primary sends full tracks of compressed data to the secondary. The RVA does not decompress the data to send it to the secondary. It sends compressed data tracks to be written directly to the cache and NVS of the secondary.

4. Acknowledgment - last track written

   The secondary subsystem informs the primary of a “write complete” of the last track.

5. Device end (DE)

   After the transfer is complete, the primary subsystem sends a device-end status to the host.

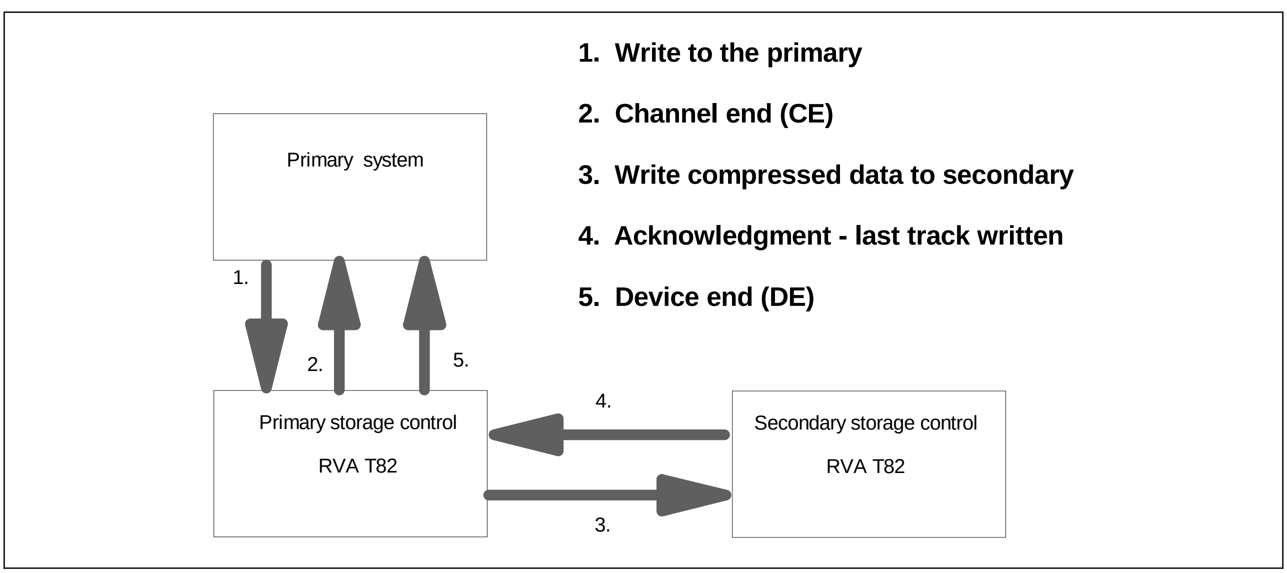

Figure 2. RVA Remote Copy Process

1.3 Terminology Used in This Book

For purposes of uniformity and clarity, we use the following terms in this book:

- Systems where the production or online data and applications run are referred to as the *primary site*.
- Subsystems where production data and applications reside and which are the primary of a PPRC pair are referred to as *primary subsystems*.
- Systems where the recovery or test data and applications run are referred to as the *secondary site*.
- Subsystems where copies of production data and applications reside, which are used to keep the data current, and are the secondary of a PPRC pair are referred to as *secondary subsystems*.
- The term *rolling disaster* is used to refer to the "dying scenario."
- *Simplex* state is the initial state of the volumes before any PPRC pairs have been defined. It is also the state of the volumes after the PPRC pairs have been terminated through using the CDELPAIR command.
- *Pending* state indicates that the volumes have been established as PPRC pairs and are currently copying data to the secondary devices. This is also the state of the PPRC pairs when they are reestablishing after they have been suspended.
- *Duplex* state indicates that PPRC is complete and the pairs are synchronized.
- *Suspended duplex* state occurs when the primary and secondary volumes cannot be synchronized or a host issued a CSUSPEND or a CGROUP FREEZE command.

Chapter 2. Planning for Remote Copy

To successfully plan for remote copy, keep in mind that remote copy is just a part of your overall disaster recovery plan. An awareness of this will greatly help you decide which remote copy technique you need to implement. Disaster recovery plan targets such as recovery period, priority applications, and financial impact are just some of the considerations that will help you evaluate which remote copy technique to implement to make full use of your current disk subsystem investment.

In this chapter, we do not go through the whole process of planning for both extended remote copy (XRC) and PPRC. Instead we briefly review planning for both remote copy techniques and then discuss planning PPRC for the RVA. The IBM Redbook, *Planning for IBM Remote Copy*, SG24-2595, covers the characteristics of XRC and PPRC and is a good reference for understanding the differences between the two techniques.

2.1 Characteristics of a Remote Copy Solution

Besides addressing your company's overall disaster recovery requirements, any real-time disaster recovery solution should:

- Be transparent to applications and systems
- Be able to freeze the secondary copy of the data at the moment of disaster
- Maintain the sequence of writes, across any number of volumes, control units, or hosts
- Allow an emergency restart after any disaster
- Allow simplified recovery scenarios
- Have options not to mirror the rolling disaster

As illustrated in Figure 3 on page 6, as part of your overall disaster recovery strategy, the goal of remote copy is to have a consistent concurrent copy of your primary subsystem.

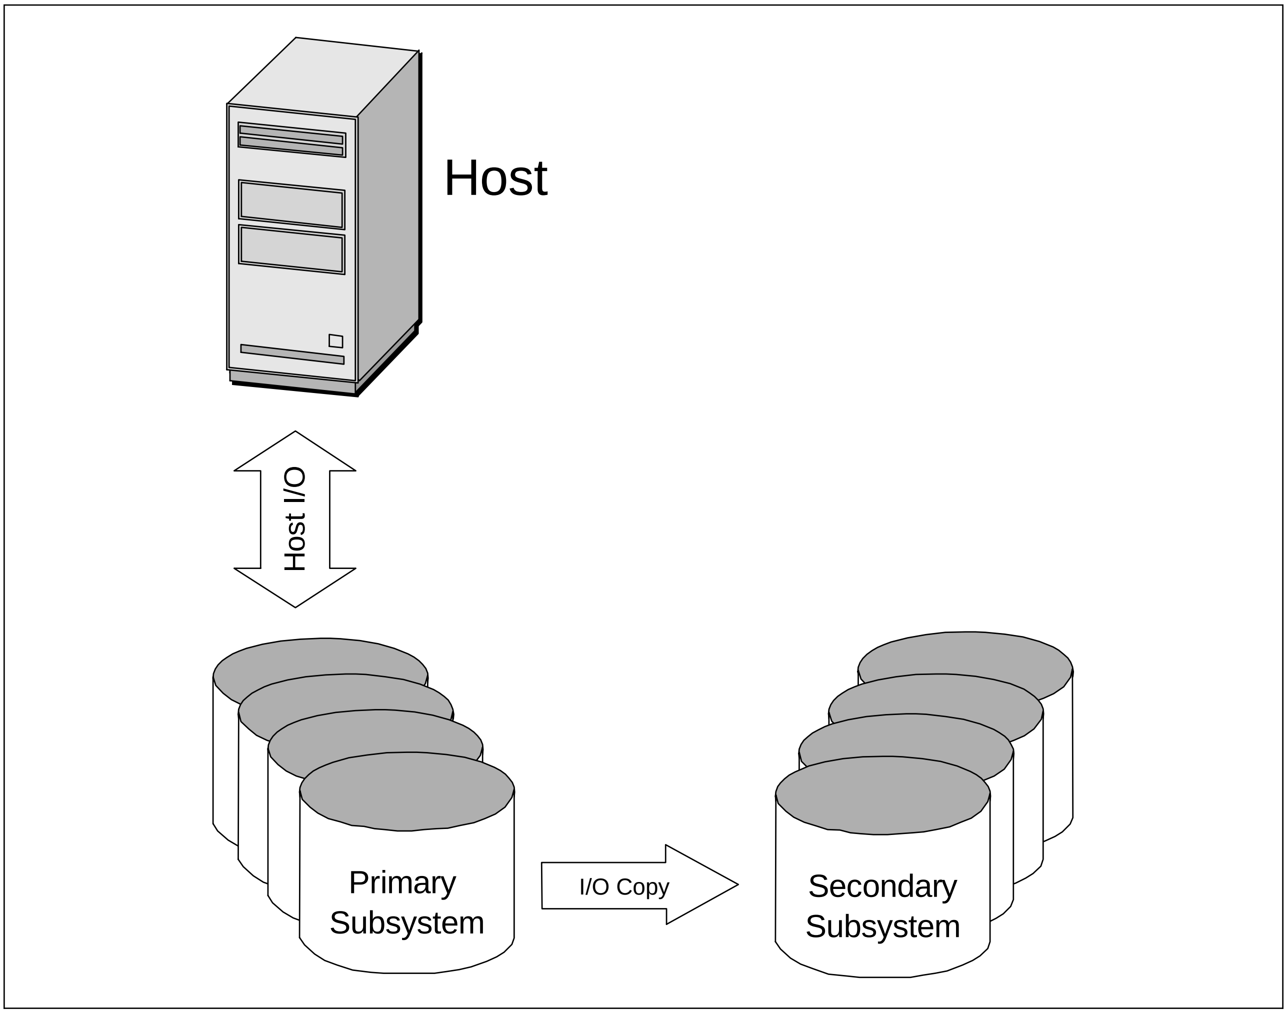

Figure 3. Goal of Remote Copy

2.2 Avoiding the Rolling Disaster or Dying Scenario

In real situations, your information technology complex is subject to many forms of problems that could lead to disaster. Situations such as fire, earthquake, and explosion threaten your entire complex. In disaster situations you never expect your entire complex to fail at the same moment. Failures tend to be intermittent and gradual, and disaster actually occurs over many seconds, even minutes. Because data may have been passed or lost in transition, data integrity is suspect. This situation is called the *rolling disaster* or the *dying scenario*.

Your remote copy solution must protect you from the many failures that occur in a short amount of time, and it must not reflect the rolling disaster.

2.3 Planning for Remote Copy on the RVA

Start by reviewing your company's disaster recovery strategy. From there you can base your decisions on the requirements of remote copy. Evaluate the following items:

1. Sites involved
   - Survey the distance between the sites.
   - Check whether the technology to connect them is available.
   - Examine the applications.
   - Assess how critical these applications are to your enterprise.
2. Applications that must be able to survive site outage
   - Categorize your installation's data and determine which data should be copied remotely.
   - Determine whether you will make a partial or full image copy of your subsystem, on the basis of the impact on your recoverability.
3. Management
   - Decide whether you want to manage your subsystems from the primary site, recovery site, or a third site.
   - Decide whether you want a unidirectional or bidirectional copy of your data between systems.
4. Time limit for the recovery site to be operational
   - Understand the financial impact of a prolonged recovery period.
5. Additional requirements
   - Determine the additional hardware or software you need to implement remote copy.
   - Evaluate changes required to your primary and secondary sites.
   - Determine whether you need additional staff to implement the changes.
   - Assess the skills needed to implement your plan.
   - Determine your need to set up additional recovery procedures and whether they will be automated.

Document your final disaster recovery requirements because this will be the basis of your implementation plan.

Understand the nature of a rolling disaster and its impact on your recovery. Remote copy must protect your data from literally hundreds of failures in a short amount of time. It must not exactly mirror the rolling disaster.

Finalize your choice between XRC and PPRC, as we explain in the next section.

2.4 Remote Copy Techniques on the RVA

Although the RVA can be used for both synchronous (PPRC) and asynchronous (XRC) copy, implementation is architecturally different from the remote copy using the 3990-6 and 9390 subsystems.

2.4.1 XRC on the RVA

The RVA can be used only as a secondary subsystem in an XRC configuration (see in Figure 4).

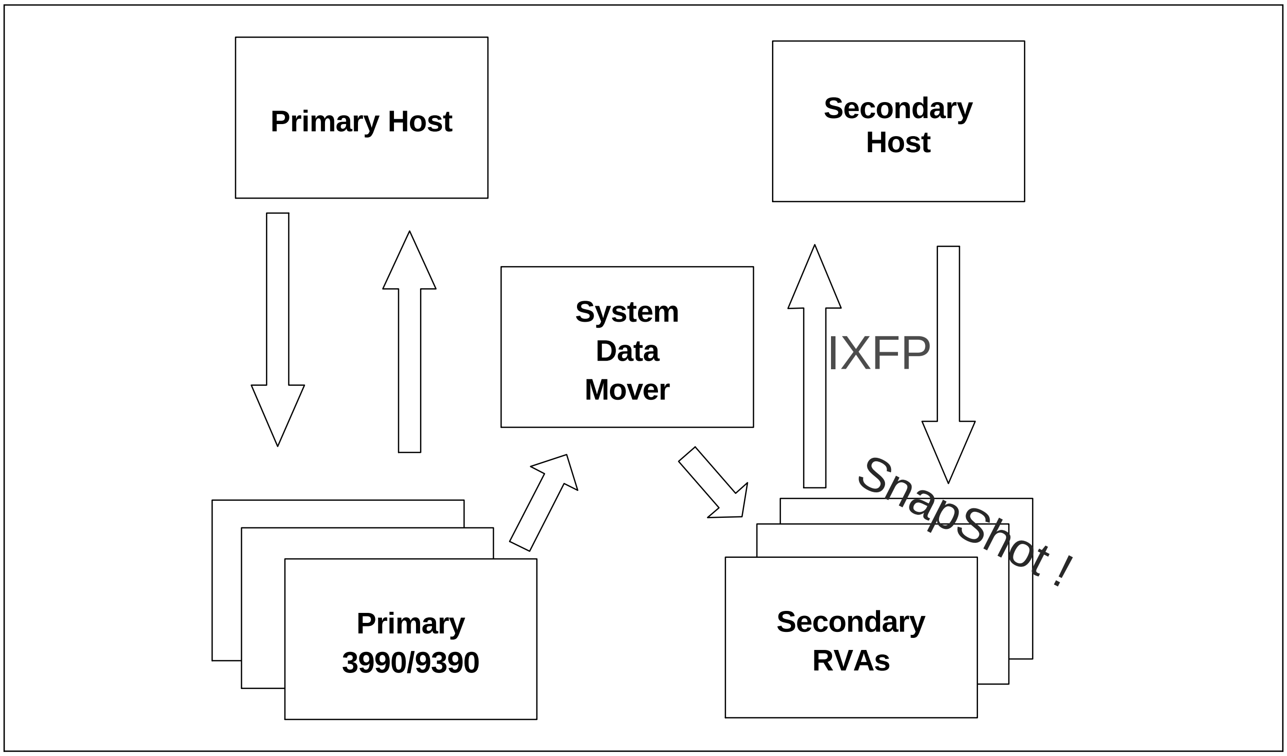

Figure 4. XRC on the RVA

One advantage of using XRC on the RVA is that you can run the IBM Extended Facilities Program (IXFP) and SnapShot on your secondary volumes. Through SnapShot you can have copies of your secondary volumes, which you can use for year 2000 testing, data mining, and other testing you may require. Although with XRC some data is lost in flight when the primary subsystems fail, customers have preferred XRC for the following reasons:

- It has minimal impact on I/O response times while maintaining data consistency between subsystems.
- It automatically maintains consistency across multiple storage controllers.
- It has a short recovery time.
- It can go beyond the limits of ESCON distance.
- It is a nonproprietary solution.

XRC is managed by a system data mover (SDM), requires OS/390, and uses CPU resources. You can manage XRC from the primary site, secondary site, or a third site depending on where you locate your SDM.

2.4.2 PPRC on the RVA

The RVA supports PPRC between T82 model subsystems, that is, both the primary and secondary subsystems must be RVA eight-path turbo subsystems as illustrated in Figure 5.

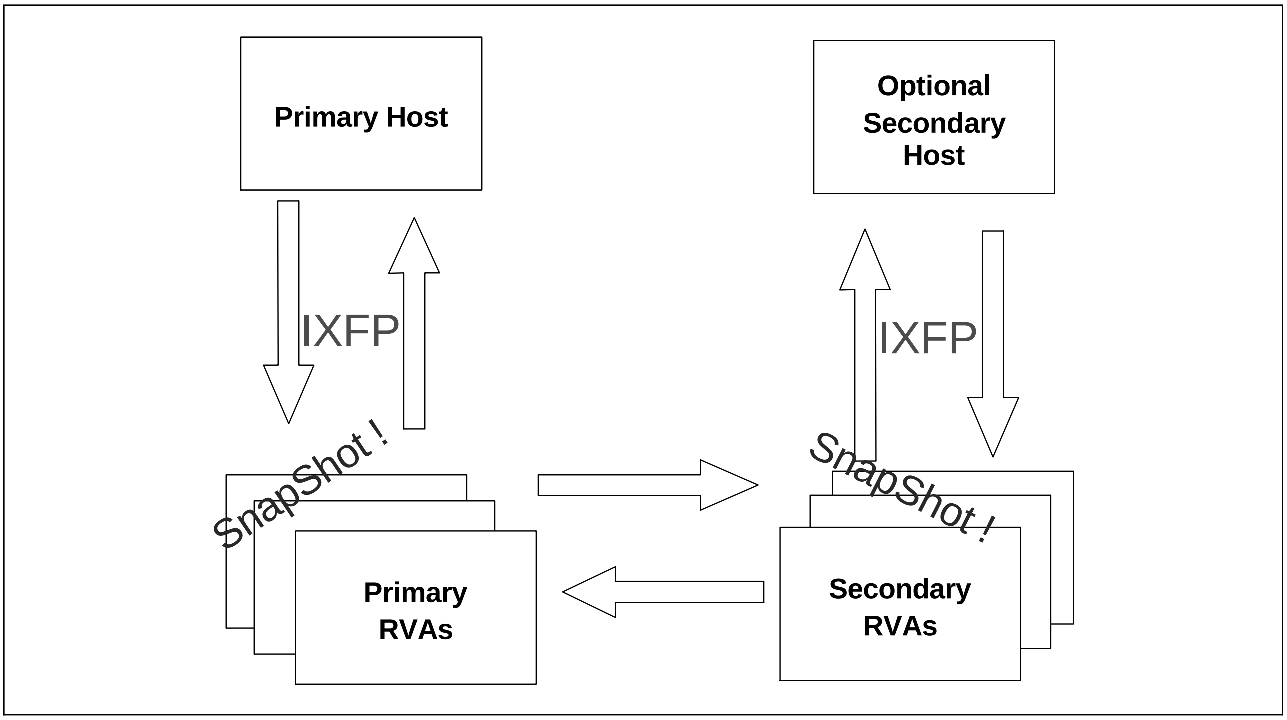

Figure 5. PPRC on the RVA

RVA PPRC is similar in configuration and operation to PPRC between 3990-6 and 9390 subsystems. However, you may encounter some architectural differences when running PPRC on the RVA. Be aware of these differences so that if necessary you can apply modifications to the applications that have automated PPRC procedures. PPRC generally requires automation to freeze the secondary volumes upon disaster and to provide cross-control-unit data consistency.

Nevertheless, PPRC has been the choice of customers who require data currency at a recovery site within ESCON distance. And with our Geographically Dispersed Parallel Sysplex (GDPS) service offering, you can have freeze capability, address data consistency, and simplified and fast recovery. For details on GDPS, see Chapter 11, “Geographically Dispersed Parallel Sysplex” on page 105.

You can also maintain your flexibility with the availability of IXFP and SnapShot on both your primary and secondary subsystems although you must follow certain rules when using SnapShot between your PPRC subsystems. Through SnapShot you can have copies of your secondary volumes, which you can use for year 2000 testing, data mining, and other testing.

2.5 Considerations on Using PPRC on the RVA

Before implementing PPRC, you must be aware of the differences between the RVA and 3990-6 storage subsystems. With the following differences in mind you can determine whether you have to modify your remote copy procedures:

- The RVA supports 3380 device types as well as 3390 device types.
- The RVA does not support parallel channels on subsystems that support PPRC.
- The fencing of the failed disk drive module (DDM), the reconstruction of the data to a spare DDM, and the replacement of the failed DDM by the CE using the regular maintenance package, can be performed while PPRC volumes are active on the RVA. PPRC does not affect normal maintenance procedures.
- PPRC on the RVA is not affected by subsystem cache settings, NVS, volume caching, or DASD fast write settings. There is no pinned data on an RVA.
- There are no PPRC switches on the RVA. On 3990 or 9390 the vital product data (VPD) is altered from the storage controller's installation and configuration panels. You therefore cannot:
  - Alter the setting of CRIT(YES). The setting is always CRIT(YES-paths). You cannot have CRIT(YES-all). We discuss how to mimic CRIT(YES-all) with a new parameter in 4.2.1, “CESTPATH: Establishing Paths” on page 52.
  - Change the consistency grouping timer. This switch specifies the time in seconds that the long busy state is issued for a device. On the RVA this duration is always 120 seconds.
  - Change consistency grouping on or off. This is controlled by the CGROUP parameter of the CESTPATH command.
- We generally recommend that a host be attached to the secondary subsystem. If “normal” operations are going to take place at the secondary subsystem, you need to have an Extended Control and Monitoring (ECAM) device at that secondary control unit.
- When the RVA generates a system information message (SIM), it is sent down *every* active path group, causing error recovery procedure (ERP) SIMs to be issued to *every host* attached to the subsystem.
- The RVA allows connectivity to 128 host channels in any combination of host processors. In an 8-ESCON channel adapter configuration, each channel adapter allows connectivity to 16 host channels. In a 16 ESCON channel adapter configuration, each channel adapter allows connectivity to 8 host channels. On the primary subsystem, channel adapters used by PPRC cannot service requests from hosts. Therefore, for each channel adapter used for PPRC from the primary subsystem, you will have 16 or 8 fewer logical channels (for 8 ESCON or 16 ESCON configurations, respectively) available for use by host processors to access RVA DASD.
- If you are running PPRC bidirectionally between two RVAs, the paths in each direction must be dedicated to the primary to secondary connection. On a 3990 or 9390 bidirectional links could share the same physical ESCON link as long as each primary to secondary connection occupies a separate logical path.

You must also observe certain conditions to ensure a smooth operation of PPRC on RVA for primary device, secondary device, and SnapShot.

2.5.1 Primary Device Considerations

You must ensure that devices are available on the subsystem for the SnapShot of devices that are not part of the PPRC pair. As a rule of thumb, you should not establish all 256 devices to be part of the PPRC pair if you want to use SnapShot.

You should not put primary devices in SnapShot target storage groups.

2.5.2 Secondary Device Considerations

You must ensure that devices are available on the secondary subsystem for use of SnapShot. You cannot use an ECAM device as a secondary device because you need an ECAM device to run IXFP on the secondary subsystem. This would limit the secondary subsystem to a maximum of 255 PPRC devices. If you are using GDPS, you cannot use a GDPS utility device as a secondary pair. You need one GDPS utility device per logical control unit. You can use the ECAM device as a GDPS utility device.

2.5.3 SnapShot Considerations

You must remember the following conditions when you plan to use SnapShot on RVA PPRC:

- You cannot read from or write to a secondary.
- You can SnapShot ***from*** a primary.
- You cannot SnapShot ***to*** a primary.
- You must have online non-PPRC volumes available as SnapShot targets.
- You cannot SnapShot to or from a secondary while it is in a SECONDARY status.
- SnapShot does not check PPRC status.

For an explanation of the above rules and how to manage SnapShot on RVA PPRC, see Chapter 9, “Using PPRC with SnapShot” on page 95.

2.6 Requirements for Implementing PPRC on the RVA

There are unique hardware and software requirements for implementing PPRC on the RVA.

2.6.1 Hardware Requirements

The following are the hardware requirements for both the primary and secondary RVA subsystems:

1. Model T82 with 8 or 16 channel link connections

Notes:

a. Model T42 is field upgradable to model T81. This is equivalent to a T82.

b. FC4000, which is mandatory for model T82, provides ICE cards on positions ICE00, ICE02, ICE10, and ICE12. This is the eight minimum channel links.

c. FC4002 provides four additional ICE cards for positions ICE01, ICE03, ICE11, and ICE13.

2. 16 MB shared memory (IXT2 cards)

3. PPRC support microcode

   Notes:

   a. Licenced Internal Code (LIC) upgrade to PPRC is nonconcurrent

   b. FC7001 provides the IXT2 cards and the PPRC support microcode

4. Both subsystems must have the Call Home feature active

In addition, we recommend that you have 4 GB of cache for RVA PPRC on both subsystems for optimal performance.

Other hardware requirements for implementing RVA PPRC are:

- You *must* dedicate the primary ESCON links used for PPRC.

  At least one ESCON link must exist between RVA subsystems in which you want to set up PPRC pairs. Any ESCON link is capable of addressing all volumes between a primary RVA subsystem and a secondary RVA subsystem. If you are using an ESCON Director, you do not have to have more links dedicated to PPRC from the primary because a single link can get to multiple subsystems (up to four).

  Additionally, whenever the last path is removed from PPRC use, that link is available for use by a host. This situation could cause operational inconsistencies when you try to reestablish paths.

  Note

  For links that are going to be used by PPRC, we recommend that:

  - You have at least one link from each cluster, to ensure connectivity. Two links per cluster provide better performance for systems with a lot of PPRC activity.
  - You not define links that are being used for PPRC in the I/O configuration definition (IOCD) or hardware configuration definition (HCD) for any host system.
  - In configurations where it is possible that the secondary RVA subsystem can take over as a primary RVA subsystem, you not include these links in the IOCP or HCD for any host system;
  - If a PPRC link is in the IOCP or HCD definition of a host system, you must configure the link offline to that host before any PPRC paths can be established using that link.

- Upgrades to the RVA LIC should not be disruptive to PPRC operations. You will be informed if an LIC upgrade is going to be disruptive.

- Changing an SSID is disruptive to PPRC operations. Once an SSID has been assigned to a logical subsystem, in order to change it you must delete all PPRC pairs and paths and remove the PPRC capability from the subsystem. We recommend that you not plan to change the SSID.
- Because the RVA supports both 3380 and 3390 device types, you must be careful that you PPRC to the same device geometry. You can PPRC from a smaller device to a larger one of the same geometry.

2.6.2 Software Requirements

There is no specific software that is mandatory for PPRC; however, the following software support is recommended:

- DFSMS 1.3 plus a PTF is required to use the API for the CGROUP command. This PTF also updates CQUERY with new status codes for the RVA.
- PTF for IXFP provides detailed text for RVA error messages.
- PTF for SnapShot provides specific error messages when using SnapShot with PPRC.
- PTF for ICKDSF adds support for the seven-digit serial numbers for PPRCOPY ESTPATH and QUERY. This PTF also includes the new RVA PPRC status information explanations to QUERY.
- To use PPRC commands available from ICKDSF you need PTFs for ICKDSF Release 16 to support this on the RVA. This is similar to the requirements for ICKDSF on the 3990-6. No changes are required for the general ICKDSF commands or for EREP.

For PTF numbers and the latest software support information, contact the Support Center and obtain the PSP bucket for 9393.

Make sure that you obtain informational APAR II08303 for PPRC.

2.7 Managing PPRC on the RVA

Decide from which site you should manage your PPRC subsystems and how you should access your data. Usually PPRC is managed from the primary site, but this implementation is not a hard and fast rule. It is the preferred practice because operations management is usually centralized on the primary site or because a secondary host is not installed.

Managing PPRC from the secondary site separates your PPRC-related activities from your usual operations activities at the primary site. You have to plan for PPRC operations activities be it in your primary or your secondary site. You have to weigh the advantages and disadvantages of separating the management of PPRC from your primary site, taking into account the skills involved, time saved, and resources needed. You also have to consider how you will set up any automation sequences you want to implement. We discuss RVA PPRC operations thoroughly in Chapter 4, “Operational Considerations” on page 51.

2.7.1 Unidirectional and Bidirectional Copy

Your overall disaster recovery plan dictates your need for either a unidirectional or bidirectional copy of your data. Unidirectional copy is usually done for systems that use the secondary site only as the recovery site in disaster situations or for testing. Bidirectional copy gives you flexibility to manage your PPRC pairs from both primary and secondary sites if your sites are to be each other's backup. Bidirectional copy also allows you to balance the I/O workload across all your RVAs. However, a bidirectional implementation may be more complex to manage in terms of automation and documentation. We recommend that all volumes you need to recover an application be at the same site. You must make sure that each primary to secondary connection has a separate physical link as well as a logical path. If you have only eight ESCON channels, you may need to upgrade to 16 channels to provide adequate paths for both PPRC and host connections.

2.7.2 Manage Snapshot and IXFP Activities

In planning your strategy of managing RVA PPRC, also consider whether you are planning to use both SnapShot and PPRC. If you use SnapShot, you might want to reserve some volumes for it, so do not plan to PPRC all 256 addresses.

Since you can use IXFP on your primary and secondary subsystems, you must also plan on managing deleted data space release (DDSR) on your subsystems. DDSR is automatic for secondary devices, because changes are propagated from the primary devices, but you may want to have a host attachment for IXFP monitoring of NCL.

2.7.3 PPRC Migration Manager

For IBM customers who have MVS systems, the PPRC Migration Manager for MVS is of great help and use in using PPRC to migrate production data to secondary subsystems. This software tool exploits PPRC, and many customers are using it not only for migration but also to move data for load balancing among systems and to create copies of data for use in year 2000 testing activities. It makes extensive use of the PPRC TSO commands and uses Interactive System Productivity Facility (ISPF) dialog functions. The latest service level adds new functional support required for use with the RVA.

The PPRC Migration Manager automatically recognizes whether it is developing jobs for use with an IBM 3990-6/9390 or an RVA. Keep in mind these differences when using the PPRC Migration Manager on RVA subsystems:

- The RVA requires the use of a seven-digit serial number. With this requirement, the 3990-6/9390 high order two digits will be 00.
- The LINK DETAIL panel for a 3990-6/9390 subsystem is different from the LINK DETAIL panel for an RVA. The correct panel is automatically presented to the user.
- The channel connection address (CCA) value of the RVA is a relative CCA value. Check the I/O configuration data set (IOCDS) definition.

The PPRC Migration Manager for MVS is available from the Internet. See 4.5.1, “PPRC Migration Manager” on page 57 for instructions to obtain the PPRC Migration Manager code and documentation.

2.8 Configuring PPRC on the RVA

There are two techniques for connecting your primary and secondary RVA subsystems:

- Point to point or direct connection
- ESCON Director connection

Each technique has distinct requirements which you must consider to help you decide which to implement. We recommend having a one-to-one logical control unit connection to keep your configuration simple. For a detailed discussion on RVA logical control unit and configuration, see Chapter 3, “Configuration” on page 23.

Important: Because paths used for PPRC are not defined in the IOCP or HCD, these paths are not known to the ESCON manager.

Physical configuration of PPRC on the RVA requires that you have sufficient knowledge of configuring and implementing ESCON links. You must include in your plan how you will manage your ESCON configuration. For a discussion of the characteristics of ESCON see the IBM Redbook, *ESCON Implementation Guide*, SG24-4662.

2.8.1 RVA ESCON Connectivity

The RVA allows connectivity to 128 logical paths by means of 32 channel register sets. The channel register sets are equally allocated to each ESCON port configured in the ICE cards installed.

An RVA with eight ESCON links has four channel register sets per link. Each link can be paired with four other physical secondary RVAs (see Figure 6). The channel register sets are numbered for illustration purposes only; they are not related to any logical or physical addresses.

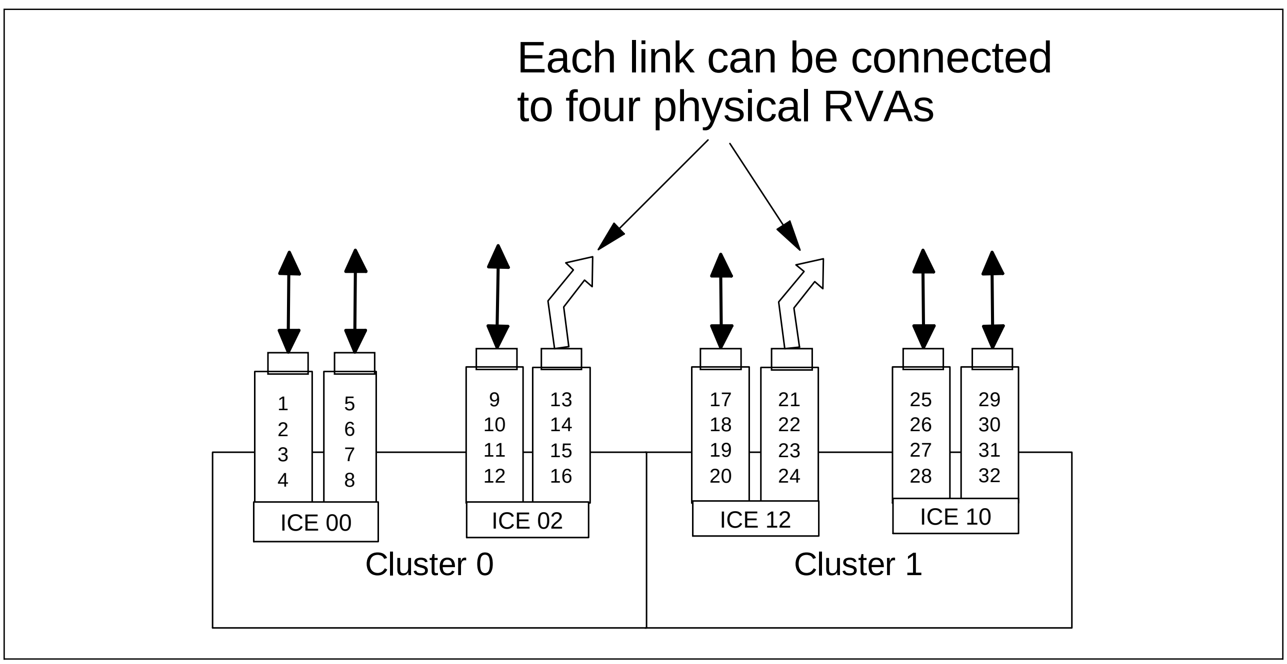

Figure 6. RVA with Eight ESCON Links

An RVA with 16 ESCON links has two channel register sets per link. Each link can be paired with two other physical secondary RVAs (see Figure 7 on page 16).

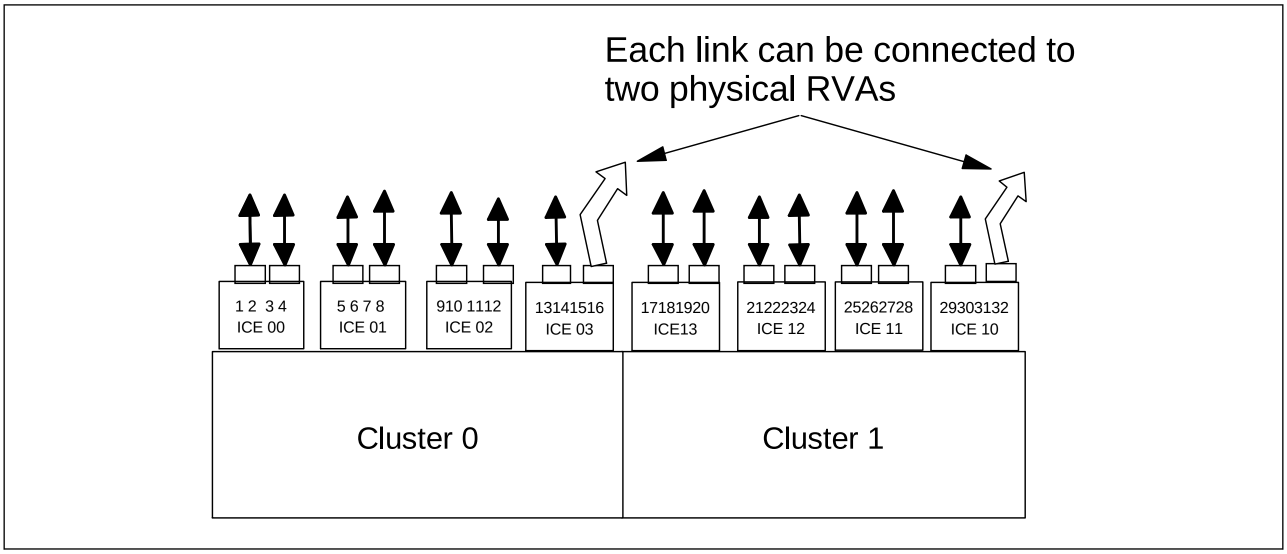

Figure 7. RVA with 16 ESCON Links

In Chapter 3, "Configuration" on page 23 we discuss channel register sets, modes, paths, and ICE cards.

2.8.2 Considerations on Using the ESCON Director

If you plan to use the ESCON Director, you can manage it using the ESCON Director console or the ESCON Manager program product.

You can have a static (dedicated) or dynamic connection through an ESCON Director. However, if the connections go through more than one ESCON Director, remember that only one connection can be configured as dynamic.

The ESCON Director temporarily connects ports for transmission. The connection can remain in effect while several frames are exchanged between the ports. To manage links and their operation, frames called *link-level frames* are transmitted. Frames that are associated with device operations are called *device-level frames*. As shown in Figure 8 on page 17, both the link-level and device-level frame link-headers contain the following:

- Destination address
- Source address
- Link control information

The destination address is the address that the switching function of the ESCON Director uses to determine the outbound port to which to "send" the data. Byte 0 of the destination address contains the destination link address when the channel path includes one dynamic connection. If a port has a static connection to another port, the address is not checked, and the data is passed through the destination port. Therefore, if a frame must travel through two ESCON Directors. only one can establish a dynamic connection.

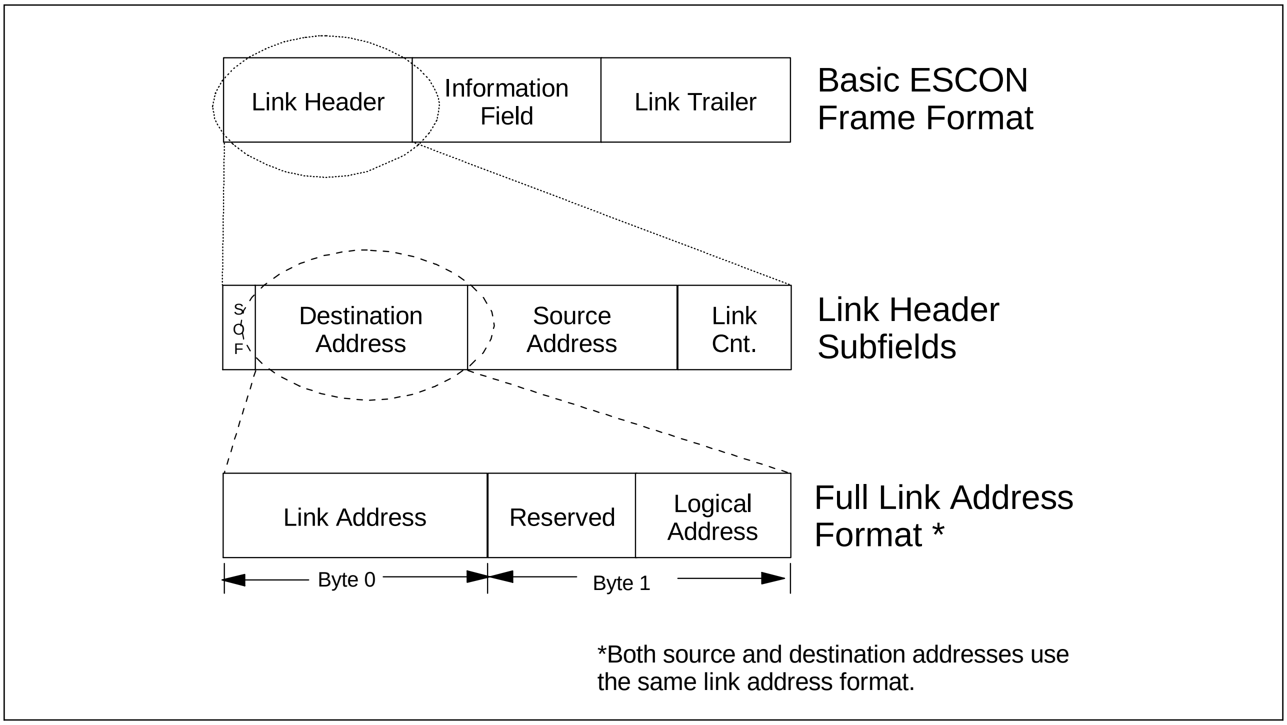

Figure 8. ESCON Address Frame Format

2.8.3 Point to Point or Direct Connection

You can use point to point connection if the systems involved are within ESCON distance and you do not want to use an ESCON Director (see Figure 9).

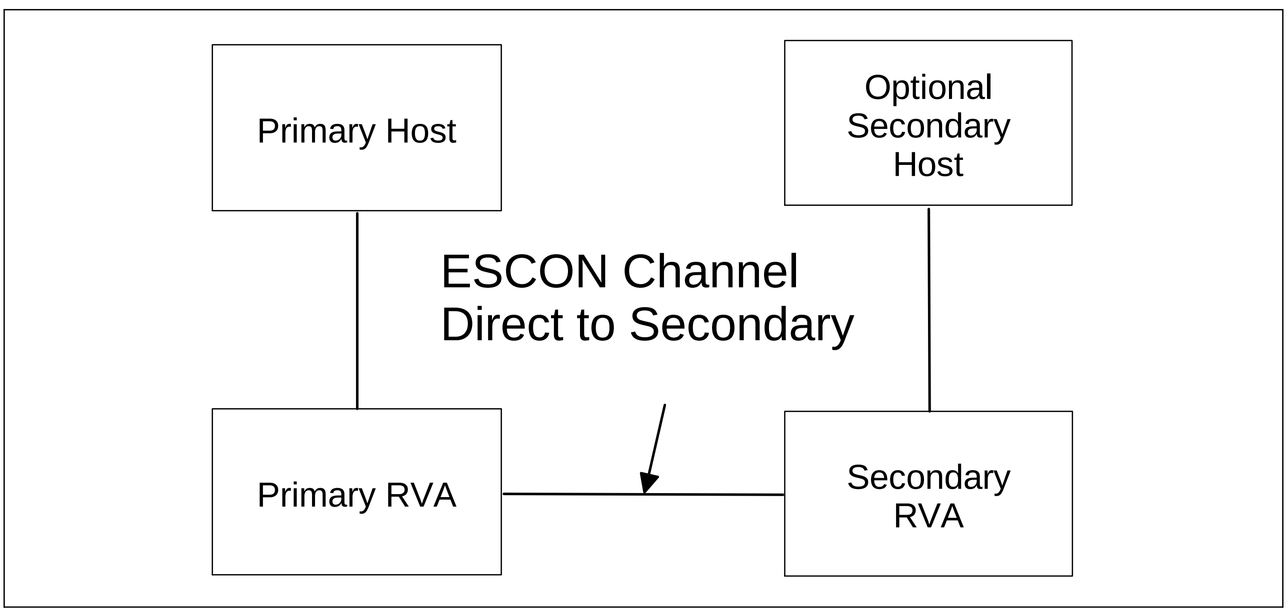

Figure 9. Point to Point Connection

A point to point connection is simple to configure and manage but in the long run may be a more expensive implementation because ESCON cables between the systems are dedicated to servicing PPRC alone. We discuss this further in 3.1.1.1, "Point to Point Connection" on page 23.

2.8.4 ESCON Director Connection

You can opt to use an ESCON Director between the primary and secondary subsystems of your PPRC pair (see Figure 10). You can have either a static or a dynamic connection using the ESCON Director.

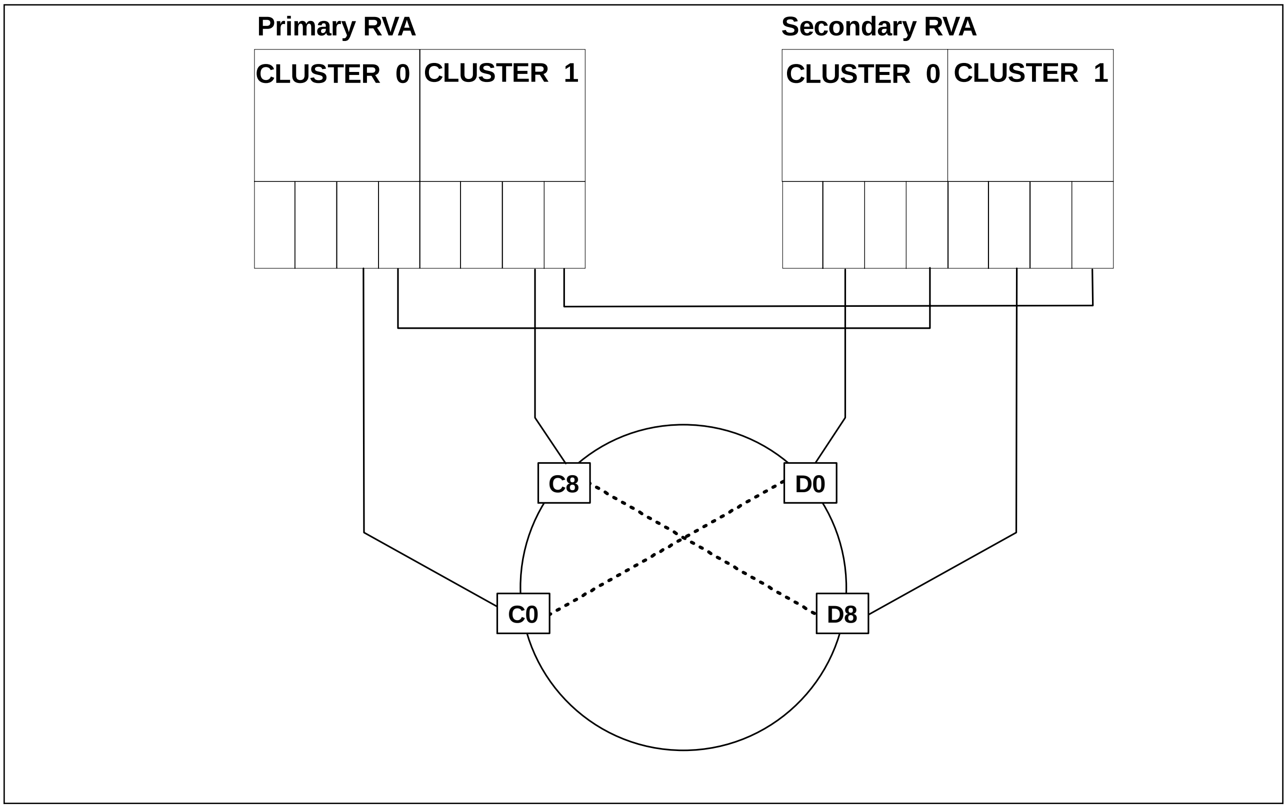

Figure 10. ESCON Director Connection

In 3.1.1.2, "Static Connection through an ESCON Director" on page 24 and 3.1.1.3, "Dynamic Connection through an ESCON Director" on page 24 we describe the differences between the two ESCON Director connection techniques.

By using the ESCON Director, you have added flexibility, provided you follow the requirements for ESCON link definitions in 2.6.1, "Hardware Requirements" on page 11.

Remember: When you use an ESCON Director, a link that the primary control unit uses for PPRC cannot be used to access devices from a host.

2.9 PPRC of Open Systems Data

You can store your Open Systems data in an RVA subsystem through the use of an IBM 9399 Cross Platform Extension (XPE). You can connect your Open Systems machines on the XPE. You can then connect the XPE directly to the RVA ESCON channel interface or through an ESCON Director (see Figure 11 on page 19). You can use PPRC to copy these XPE volumes to your secondary subsystem.

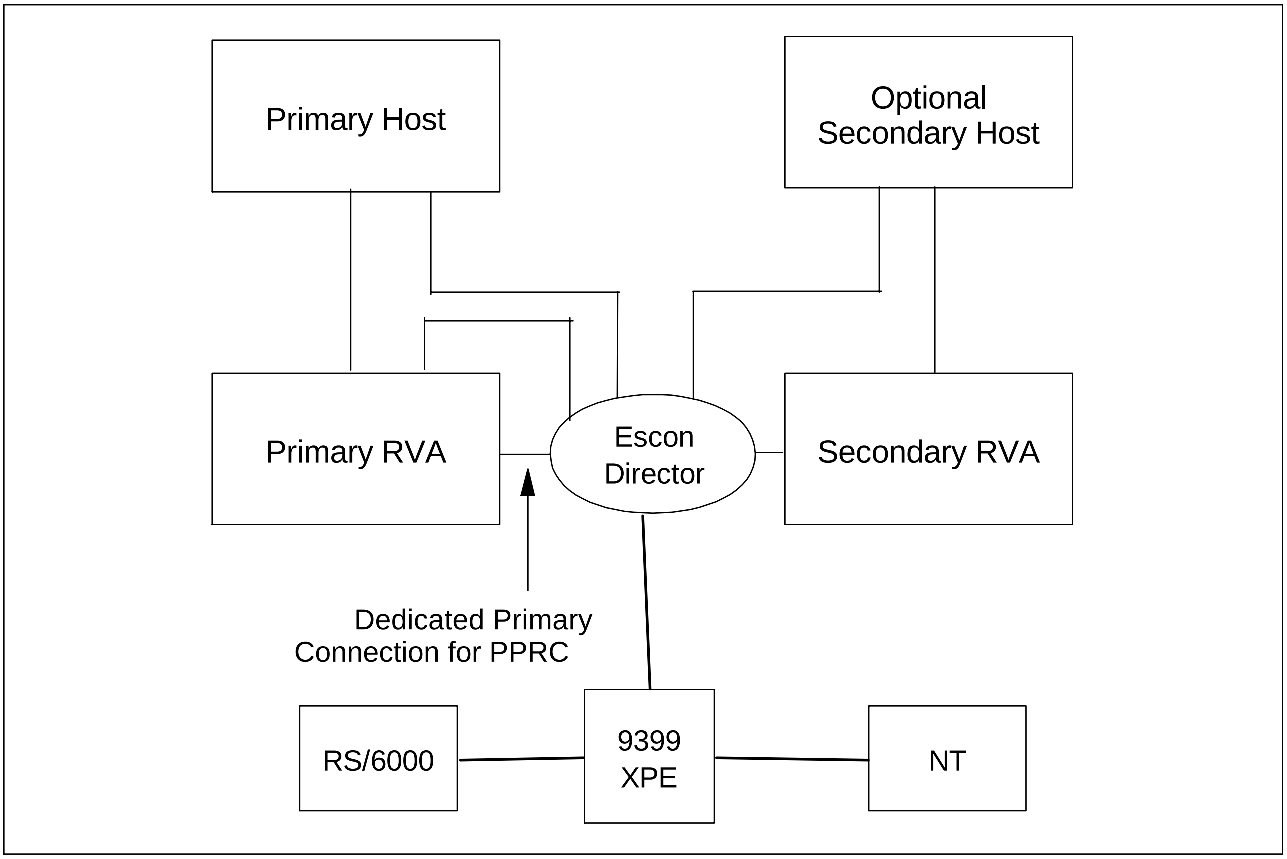

Figure 11. PPRC of Open Systems Data through XPE

The IBM Redbook, *Data Sharing: XPE Implementation Guide*, SG24-5256 covers XPE in detail.

2.10 Database Recovery Using PPRC

In this section we present some items that are important for you to consider when planning for RVA PPRC. We discuss database system recovery in detail in Chapter 8, "Using PPRC in a Database Environment" on page 87.

Database recovery involves restoring a database to its original condition after it is rendered invalid by some failure. The task of developing recovery procedures and performing recovery is an important one.

There are two ways to make your database current in case of failure at your primary site:

- Restart

  Restart consists of starting a database application after a failure without having to restore the database. This process is measured in minutes. In order to restart, your database must be in a consistent state. In a disaster recovery situation the point of consistency is at a point in time before the disaster occurred; any updates that occurred after the point in time will be lost. This data loss may be acceptable because a database restart is fast, so there is a business trade-off between the speed with which you get your system back, and the amount of data lost.

- Recover

  Recover consists of restoring the last set of image copy tapes and applying log changes to bring the database up to the point of failure. This process is measured in hours or even days. However, if you have all log data available, you can recover all updates that occurred up to the time of the disaster.

In any operating system, the sequence in which updates are made is what maintains the integrity of the data. If that sequence is changed, data corruption will occur. PPRC, together with the ERP, provides consistent copy across multiple database systems, such as IMS and DB2.

As Figure 12 shows, to maintain integrity any database system must follow this sequence:

1. Record on the log (or logs if the log data set is mirrored) that a database update will be done.
2. Write the updates to the database volumes and their corresponding pairs.
3. Mark that the update is complete on the log.

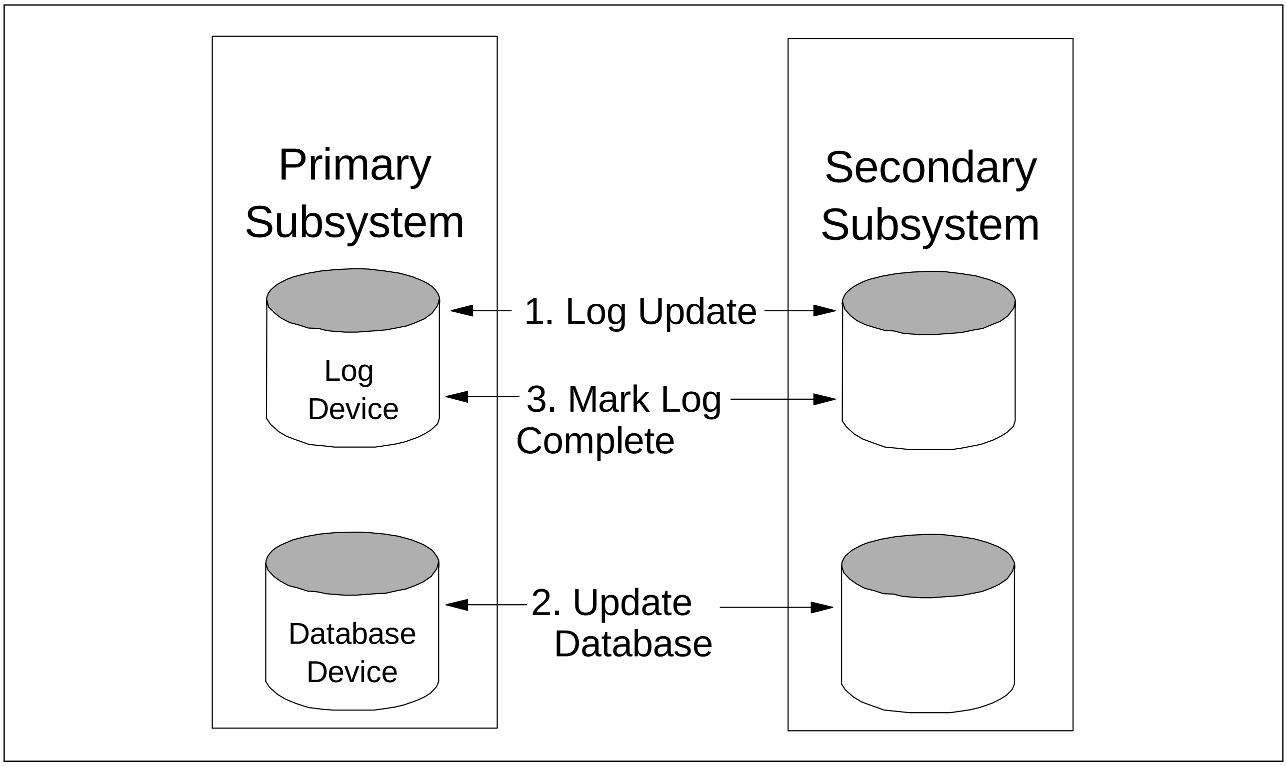

Figure 12. Maintaining Write Sequence Integrity

To preserve data integrity at your secondary site, it is essential that updates are written to the secondary site in the same sequence in which they are written at the primary site. This is called preserving the *write sequence integrity*. XRC maintains write-sequence consistency by timestamping and consistency group management. The issue in a synchronous remote copy implementation, such as PPRC, is to preserve write-sequence consistency across multiple volumes and subsystems.

Using CGROUP in combination with the DASD ERP and automation, you can preserve the write-sequence consistency of your database in the event of a disaster. In the event of an error, which may be the first sign of a rolling disaster, you can use CGROUP to freeze the secondary site and prevent any further updates from being written at the secondary site. Thus you can restart your applications, but any updates that were made after you froze your secondary site are not reflected at the secondary site. This could be considered to be a data loss.

GDPS allows you to freeze all updates at both the primary and secondary site in the event of an error. This option allows you to eliminate any potential data loss, but at the cost of an interruption to your production workload.

To ensure remote site recoverability, we recommend that you use the CGROUP command as described in Chapter 4, "Operational Considerations" on page 51, or the GDPS service offering. This command provides data integrity management of the secondary site in various error situations.

We also recommend that you take note of the recommendations and considerations in using PPRC with various database systems as described in Chapter 8, "Using PPRC in a Database Environment" on page 87.

Chapter 3. Configuration

In this chapter we present an overview of:

- ESCON
- RVA
- PPRC on the RVA

We assume you are already familiar with the concepts and operation of the RVA as well as the concepts and operation of PPRC on 3990-6/9390 subsystems. If you are not, we recommend that you read these redbooks: *IBM RAMAC Virtual Array*, SG24-4951, and *Planning for IBM Remote Copy*, SG24-2595. In this chapter we provide information about the RVA that is of interest only for PPRC implementations.

3.1 ESCON

ESCON is an IBM architecture and fiber optic technology. It allows a data center installation to replace electrical-copper cables with fiber-optic cables. Beyond this physical aspect, ESCON provides new capabilities to connect I/O equipment to a host processor at distances up to 43 km. Refer to 3.1.2, “ESCON Specifications” on page 25 for further details. With these increased distances you can establish new disaster recovery methods, such as remote copy.

In this section we explain the various connection modes in general. For detailed information about the different configuration options to deal with great distances, see 3.1.2, “ESCON Specifications” on page 25.

For detailed information about implementing ESCON, see the *ESCON Implementation Guide*, SG24-4662.

3.1.1 ESCON Connection Modes

You can connect an I/O control unit to a host channel through:

- Point to point connection
- An ESCON director as a static connection
- An ESCON director as a dynamic connection

3.1.1.1 Point to Point Connection

You can establish a point to point connection (Figure 13 on page 24) by connecting a fiber between a host channel and an I/O control unit. This is called a *physical link*. Through this physical link a host can communicate with only the devices attached to that I/O control unit. A host communicating with another I/O control unit needs reconnect of the fiber, creating a new physical link.

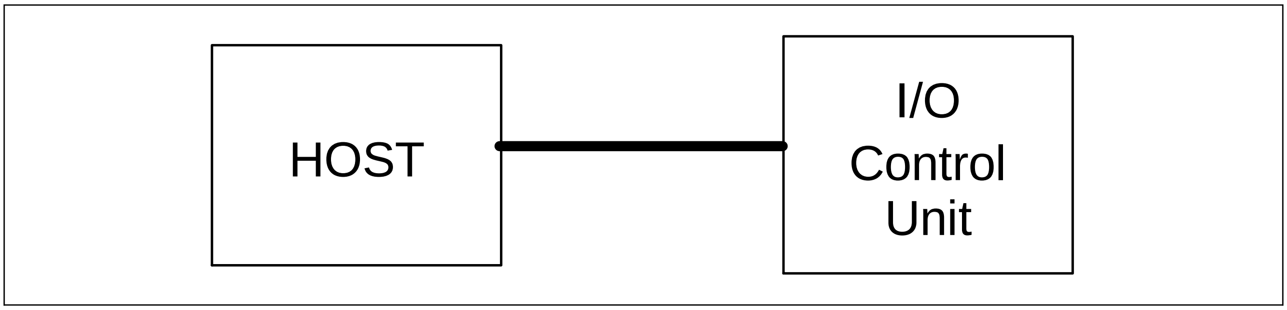

Figure 13. Point to Point Connection

3.1.1.2 Static Connection through an ESCON Director

A static connection through an ESCON director involves having one fiber connected from a host channel to an ESCON Director and a second fiber from an ESCON Director to an I/O control unit. In Figure 14, the host channel is connected through ESCON Director port D8 and ESCON Director port C3 to an I/O control unit. You must set up an ESCON Director for a static connection (dedicated connection) of these two ports, by using the ESCON Director console or the ESCON Manager.

A static connection through an ESCON Director is more flexible than a point to point connection. You can achieve a different connection by statically switching to another port, without reconnecting the fiber.

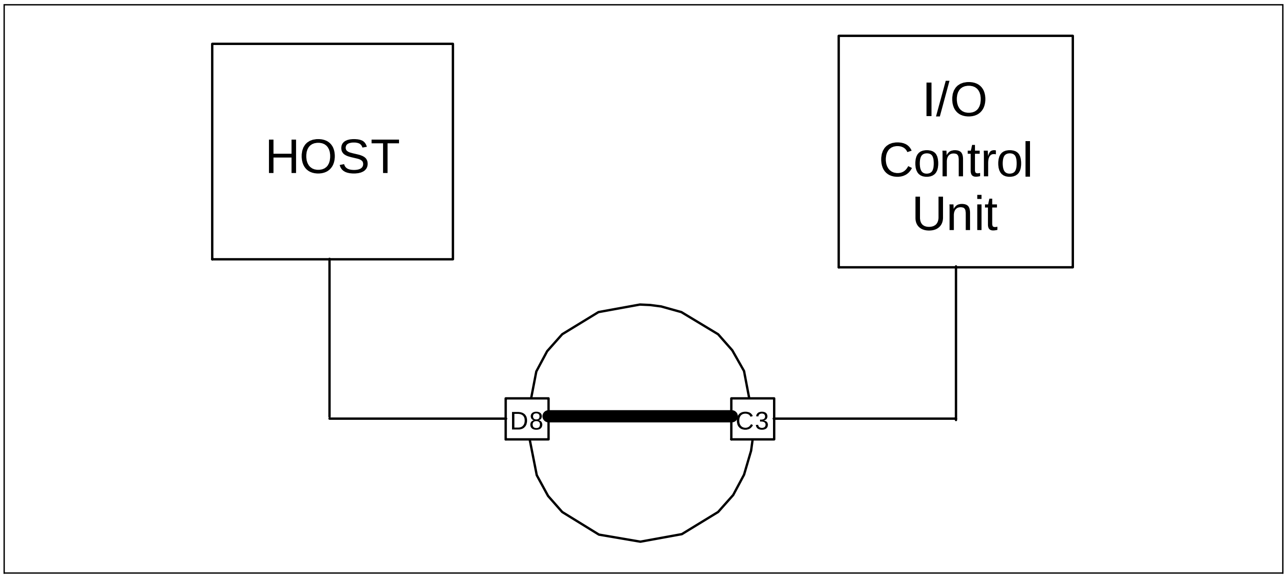

Figure 14. Static Connection through an ESCON Director

3.1.1.3 Dynamic Connection through an ESCON Director

A dynamic connection through an ESCON Director (Figure 15 on page 25) involves having the same fiber arrangement as for a static connection. However, instead of setting up an ESCON Director for a static connection, you have to code the connections to ESCON Director destination ports C3 and E7 in the link statement of the control unit macro in the IOCDS or equivalent HCD panel. In this mode a host can communicate with multiple devices behind different I/O control units by using the same physical link between the host channel and the ESCON

Director. The ESCON Director dynamically connects to the destination port, as required.

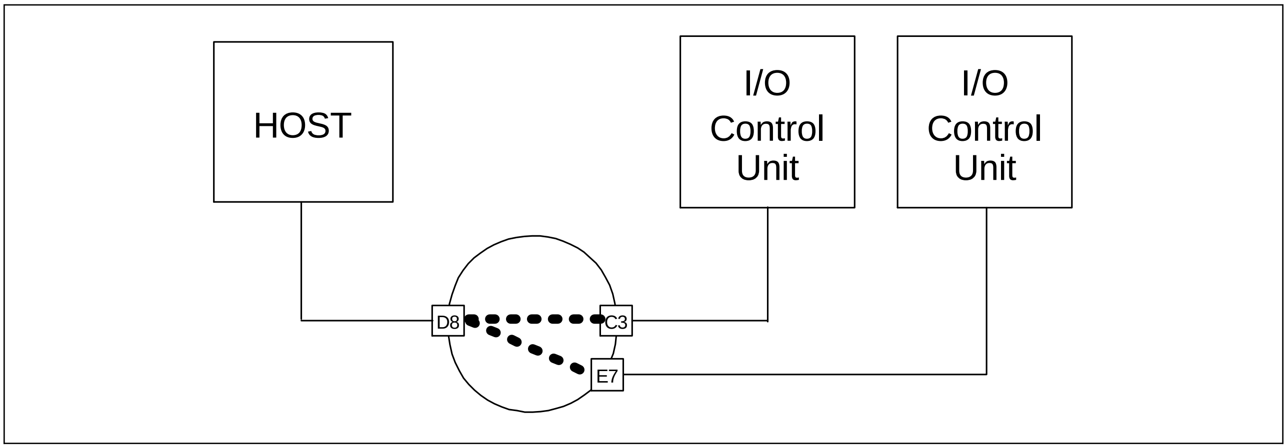

Figure 15. Dynamic Connection through an ESCON Director

Note

To connect a secondary RVA to a primary RVA you can use one connection mode or a mixture of modes, depending on your installation.

3.1.2 ESCON Specifications

Although ESCON provides capabilities to connect over long distances, the following factors affect the distance that can be achieved:

- Physical size of fiber
- Fiber transmission mode
- Connection mode

3.1.2.1 Physical Size of Fiber

An optical fiber functions as a kind of waveguide for light. It is usually made of silica glass. The fiber itself has a central core and a surrounding cladding of slightly different glass material. The physical size of an optical fiber is determined by the diameter of the core and cladding, expressed in microns (μ). One micron is a millionth (1/1,000,000) of a meter. A fiber optic cable having a core diameter of 62.5 μ and a cladding of 125 μ is designated as 62.5/125 μ optical fiber. Other fibers commonly used for ESCON are 50/125 μ and 9/125 μ.

The jumper cable has the same function as the pair of cables that make up a parallel channel. An ESCON jumper cable consists of two unidirectional fibers, one that carries data in one direction (transmit), and the other that carries data in the opposite direction (receive). The unidirectionality is due to the way the fibers are connected to transmitters and receivers, not to any property of the fiber itself.

Jumper cables are typically used to connect:

- A channel to a control unit
- A channel to an ESCON Director
- An ESCON Director to an ESCON Director

An ESCON Director to a control unit

An RVA to an RVA (PPRC)

ESCON cables, which include two fibers, are available from IBM in different standard lengths up to 122 m (400 ft). These are called single jumpers, because they connect one single link. ESCON multi jumpers, which include multiple fibers, are available from most cable suppliers. Figure 16 shows an ESCON single jumper cable.

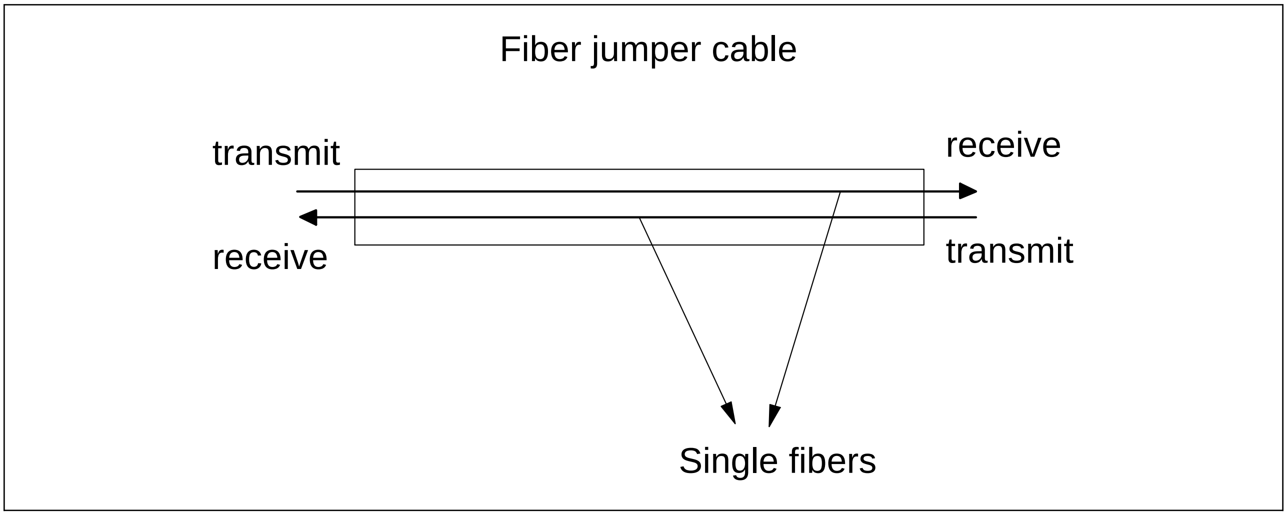

Figure 16. ESCON Single Jumper Cable

3.1.2.2 Fiber Transmission Mode

Two modes can be used to send light signals through an optical fiber: single mode (mono mode) or multimode. The optical fibers used are called accordingly single-mode or multimode fibers. These fibers have different physical dimensions and light transmission characteristics.

The intensity of the light decreases when the light passes through the fiber. The longer the fiber, the lower the light intensity. The light intensity attenuation results in length limitations of the physical link because the receiver needs a minimum level of light intensity to correctly detect the signals.

The multimode fibers supported by IBM are either 62.5/125 μ or 50/125 μ. The light source used for multimode fiber is usually a light emitting diode (LED). LED jumpers are orange, and their duplex connectors are black. The maximum distance for a multimode fiber link is 3 km if 62.5/125 μ fibers are used and 2 km if 50/125 μ fibers are used. The maximum distance for a multimode fiber link is 2 km if both 62.5/125 μ and 50/125 μ fibers are used.

Single-mode fiber usually has a core diameter of 8 to 10 μ and a cladding diameter of 125 μ. The light source used for single-mode fiber is a laser. IBM supports single-mode 9/125 μ fibers for use in an ESCON environment. The single-mode fiber must be used on links between machines with ESCON-Extended Distance Feature (XDF) adapters installed. These adapters are available only on ESCON Directors and on IBM 9036 ESCON remote channel extenders. They are no longer available on host channels. Fiber jumper cables of 9/125 μ for IBM ESCON Directors are supplied by IBM in standard lengths up to 122 m (400 ft), if XDF ports are ordered. The number of jumper cables supplied at no charge depends on the number of XDF ports installed. XDF jumpers are yellow, and their duplex connectors are grey. The maximum distance for a single-mode fiber link is 20 km.

3.1.2.3 Connection Mode

Because the light intensity is lower on increased distances, a kind of amplifier or repeater is needed to achieve distances longer than:

- 2 km for a 50/125 μ multimode fiber
- 3 km for a 62.5./125 μ multimode fiber
- 20 km for a 9/125 μ single-mode fiber

There are several options for extending this distance:

- ESCON Director

  The main purpose of an ESCON Director is to provide switching functions within an ESCON network. It also offers the possibility to extend the ESCON distance, up to 20 km on a physical link, by using an ESCON-XDF port. These ports are available as a feature on ESCON Directors.

  Depending on the transmission mode, configuration, adapters, and fibers used, the distance that can be reached may be less than 43 km. In addition to its switching function an ESCON Director acts as an amplifier or repeater, because the light signals are received on one port and transmitted by another, thus allowing distances of 2 km, 3 km, or 20 km to be reached.

- ESCON Remote Channel Extender

  The IBM 9036 ESCON remote channel extender is a signal repeater and fiber type converter that can be used to support ESCON-XDF configurations. The 9036 is available in various models, three of which are useful in a PPRC environment:

  – Model 1 - enables one connection between one ESCON multimode link and one ESCON single-mode (ESCON XDF) link
  – Model 2 - enables one connection between two ESCON single-mode (ESCON-XDF) links
  – Model 4 - combines the functions of Models 1 and 2. It can have up to 14 channel extender units within one box, thus giving a maximum of 14 connections. A configuration option is available to order the exact number of multimode to single-mode extender units and single-mode to single-mode extender units you want.

  The 9036 does not support multimode to multimode connection.

- Optical Wavelength Division Multiplexer

  The IBM 9729 is an optical multiplexer for high-speed serial data links. By dividing a single fiber into as many as 20 separate channels, it enables economical transmission of multiple bidirectional bit streams over a single fiber optic link between geographically separate locations. Each channel carries a separate bit stream, with each bit stream potentially using a different communication protocol, bit rate, and frame format. The channels can carry such protocols as ESCON, FDDI, Sysplex Timer, Coupling Links, OC-3, and Fast Ethernet. Two channels are needed for each ESCON connection. The two sites can be separated by as much as 50 km. A pair of 9729s is always needed, one at the primary site and one at the secondary site.

If your planning for PPRC results in the need for a large number of fiber connections between the primary and the secondary site, consider having a pair of 9729s. Installing a 9729 reduces the number of physical fibers needed between the primary and secondary site, which, depending on the cost of the fibers, may result in a great saving for you. The 9729 cannot be used in combination with ESCON-XDF on the same link, but the pair of 9729s can be separated by as much as 50 km, so there is no need for ESCON-XDF.

Because a fiber connection in an ESCON environment consists of two fibers, one 9729 can multiplex up to 10 ESCON connections. Figure 17 shows a 9729 configuration with eight ESCON link connections connecting two RVAs at the primary site with two RVAs at the secondary site. Each pair of RVAs is connected by four ESCON links, but only one fiber is needed to connect the two sites.

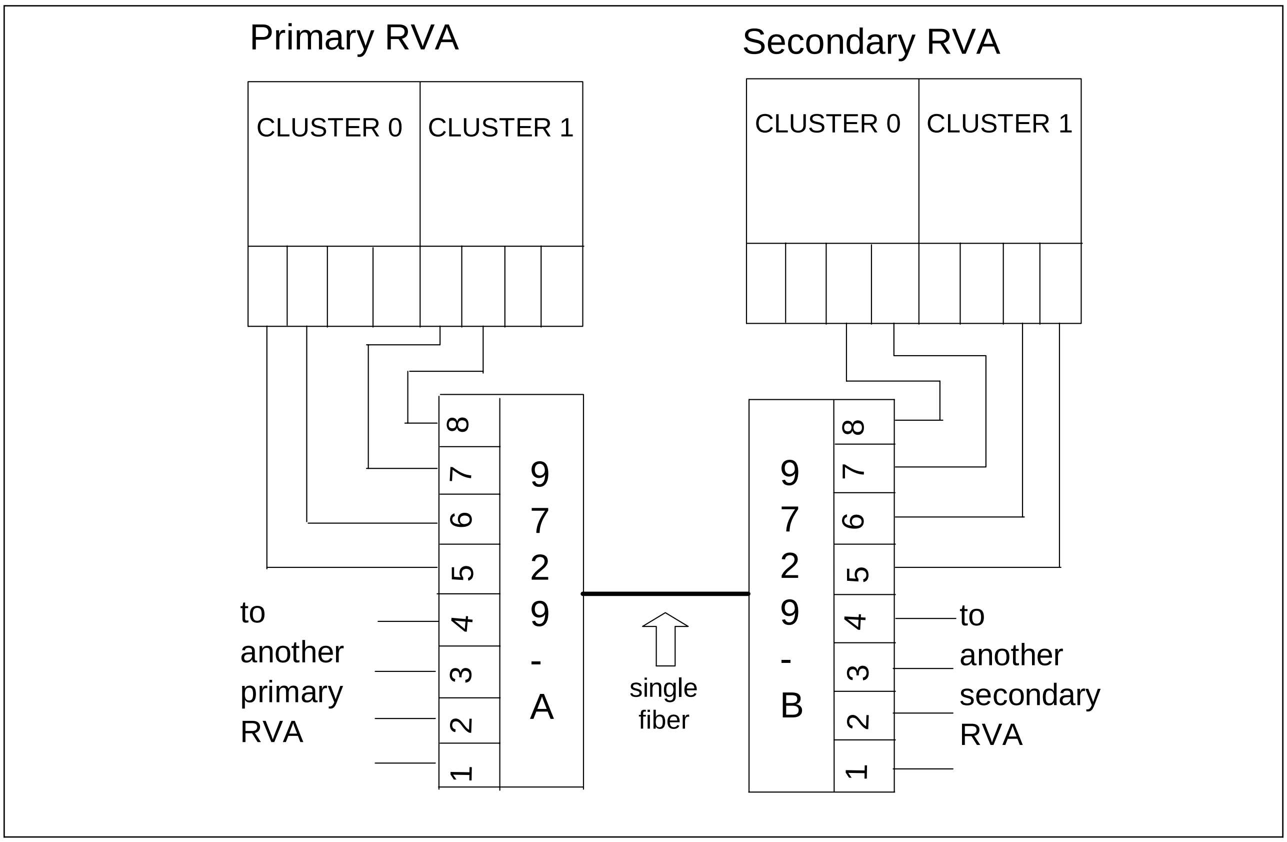

Figure 17. Sample 9729 Configuration with Eight ESCON Link Connections

3.1.3 ESCON Distance Configurations with PPRC

In this section we present the most common ESCON distance configurations.

3.1.3.1 ESCON Configuration up to 3 km

Figure 18 on page 29 shows a direct point to point connection between a primary and a secondary RVA. Because standard ESCON adapters are used in the RVA, this connection can only be a multimode connection. Depending on the fibers you use, the maximum distance can be 2 km (if 50/125 μ fibers are used) or 3 km (if 62.5/125 μ fibers are used).

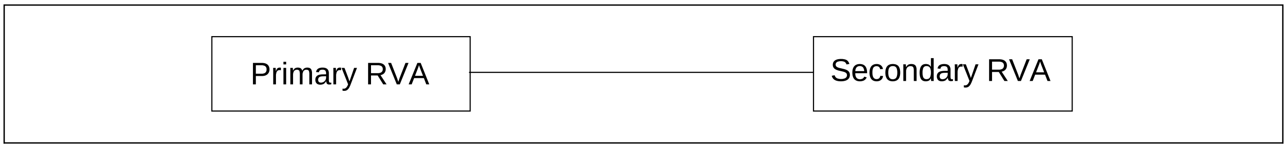

Figure 18. Point to Point ESCON Configuration

3.1.3.2 ESCON Configuration up to 6 km

Figure 19 shows a connection through an ESCON Director. In terms of distance, there is no difference between a static and a dynamic connection through an ESCON Director. The maximum distance between each RVA and the ESCON Director can be 2 km (50/125 μ fiber) or 3 km (62.5/125 μ fiber), resulting in a maximum distance between both RVAs of 4 km or 6 km, respectively.

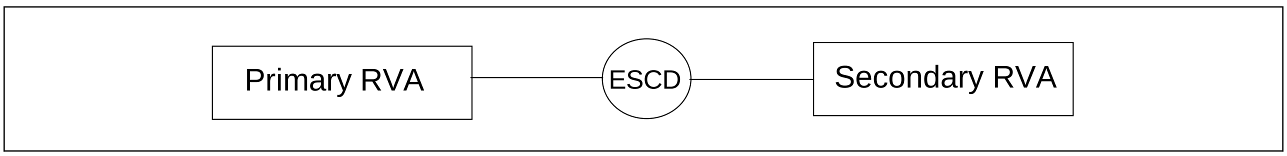

Figure 19. Configuration with One ESCON Director

3.1.3.3 ESCON Configuration up to 26 km

Figure 20 illustrates a configuration with two ESCON Directors. Remember, on such configurations, only one of the ESCON Directors can be configured for a dynamic connection. Distances of up to 26 km between the two RVAs can be realized with this configuration.

The distance between an RVA and an ESCON Director is 2 km (50/125 μ fiber) or 3 km (62.5/125 μ fiber).

The distance between the two ESCON Directors can be 2 km for 50/125 μ fiber, 3 km for 62.5./125 μ fiber, or 20 km for 9/125 μ fiber.

The 9/125 μ fiber requires an ESCON-XDF adapter installed on the ESCON Director.

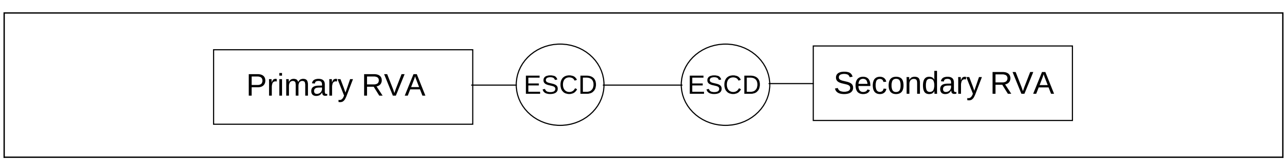

Figure 20. Configuration with Two ESCON Directors

3.1.3.4 ESCON Configuration up to 43 km

The configuration in Figure 21 on page 30 allows a maximum distance of 43 km. The 9036-001 connects one side to a standard ESCON adapter and the other side to an ESCON-XDF adapter, so the 9036-001 connects to the RVA at distances of up to 2 km (50/125 μ fiber) or 3 km (62.5/125 μ fiber). The ESCON Director, in this example, requires ESCON-XDF ports, so that the 9036-001 can connect to them. The XDF ports allow a distance of 20 km for 9/125 μ fiber on each of the physical links between the ESCON Director and the 9036-001. In theory this results in a maximum distance of 46 km (2 x 20 km + 2 x 3 km). Because the distance restriction for the RVA, as well as for the 3990-6/9390, is 43 km, the ESCON specification is superseded. Therefore the sum of the two multimode connections between the RVAs and the 9036-001s must not exceed 2 km for 50/125 μ fiber or 3 km for 62.5/125 μ fiber on such a configuration.

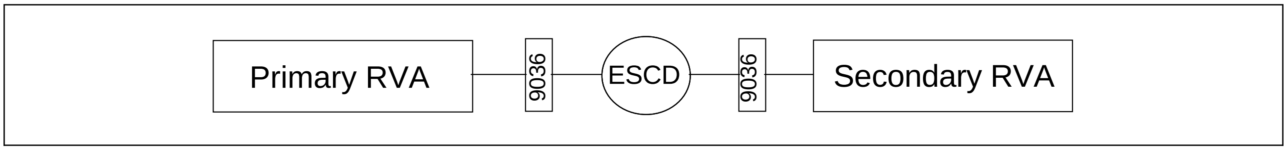

Figure 21. Configuration with One ESCON Director and Two 9036-001s

Figure 22 shows a configuration that allows a maximum distance of 43 km between the primary RVA and the secondary RVA. Assuming that both ESCON Directors have ESCON-XDF ports installed, a connection between two single-mode links is needed. Only the 9036-002 provides this connection capability. This configuration theoretically allows a distance of 46 km. Because the maximum distance for a PPRC connection between two RVAs is 43 km, the sum of the two multimode connections between the RVAs and each of the ESCON Directors must not exceed 2 km for 50/125 μ fiber or 3 km for 62.5/125 μ fiber.

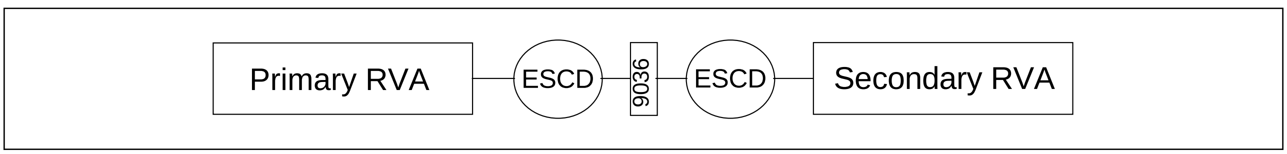

Figure 22. Configuration with Two ESCON Directors and One 9036-002

Figure 23 shows a configuration with one ESCON Director and two 9036-002s for a distance up to 43 km. Remember, the ESCON Director needs an ESCON-XDF adapter to connect to the 9036-002.

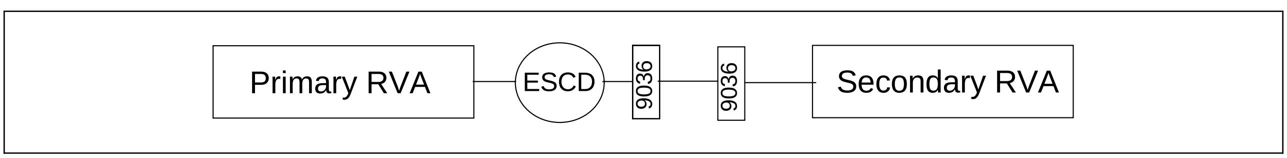

Figure 23. Configuration with One ESCON Director and Two 9036-002s

Figure 24 shows a configuration with a pair of 9729s installed. The distance between both 9729s can be up to 50 km. Nevertheless, 43 km is the maximum distance for a PPRC connection between two RVAs.

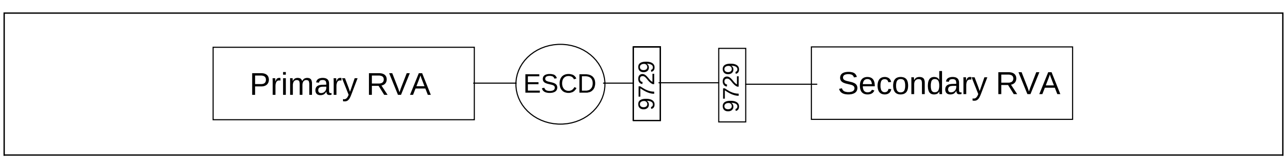

Figure 24. Configuration with One ESCON Director and Two 9729s

NOTE

Because PPRC on the RVA is less distance-sensitive than PPRC on the 3990-6 or 9390, it may be possible to attach RVAs at greater distances. You may be able to test this in your own installation with extended cable distances.

3.2 IBM RAMAC Virtual Array

PPRC is supported only on the RVA model T82. PPRC configuration and operation on the RVA are analogous to PPRC configuration and operation on 3990-6 and 9390 subsystems. The differences are documented in this book.

3.2.1 Serial Number

The serial number is a unique number given to each RVA at the plant. You can view it on screen CD21 (subsystem configuration) at the RVA display or by using IXFP. The format on the screen is S/N 0XXYYYYY, where XX is the plant code where the RVA was built, and YYYYY is the five-digit serial number. Figure 25 shows subsystem configuration screen CD21. The serial number and plant code must be known for the syntax of the CESTPATH command. See 3.2.10, “CESTPATH Command” on page 39 and 4.2.1, “CESTPATH: Establishing Paths” on page 52 for more information about the command and its syntax.

The serial number is also available from CQUERY or PPRCOPY QUERY commands.

```
Connected:  DACD           CSRC  Master: DACD              ISP-0
Status: CSE Entering Maintenance Menus          S/N 0XXYYYYY
-------------------------------------------------------------------------
                   Subsystem Configuration                           CD21

Site Name:     IBM       Subsystem Name:     TEST01   Site Loc No.: 22222

                         S/N: 0XXYYYYY         Cluster 0 Channels: 16
SSIDs: 00F1  00F2  00F3  00F4                  Cluster 1 Channels: 16
Release Lev: Trb 4.05.27                       ISP Ver:

Cache Size:      1024 MB   ---Capacities----Prod----Test----Overall---
NonVolatile:     8    MB   DA Capacity:     117.9     0.0   117.9 GB
Num Arrays: 2              % Net Load:        64.4    0.0    64.4 %
Global Spares:      2      % Coll Free Space: 32.5    0.0    32.5 %
Def Array Size:     13+2   % Uncoll Free Spc:  3.1    0.0     3.1 %
Date:  98/09/15
Time:  8 :46:38      SnapshotCopy  Maintenance  IXOF   ESCON
 =>
------------------------------------------------------------------------

  |HELP | |QUIK | |PREV |
  |     | |MENU | |MENU |
  |_____| |_____| |_____|
```

Figure 25. Subsystem Configuration Screen CD21

3.2.2 Logical Control Unit

The RVA can emulate up to four IBM 3990-3 subsystems. Each subsystem is referred to as a logical control unit (LCU). Each LCU can have a maximum of 64 devices and has to have a unique subsystem ID (SSID).

3.2.3 Subsystem Identifier

An SSID is a four-digit number given to each LCU in an RVA. The range is from 0000 to FFFF. The secondary LCU SSID cannot be specified as 0000. Each disk subsystem within a data processing center should have its own SSID. Subsystems attached to the same operating system must have different SSIDs. MVS checks during initial program load (IPL) that there are no duplicate SSIDs.

SSIDs are configured at installation time. Any SSID change, after the installation, is a disruptive task.

Once PPRC is enabled on the RVA, an SSID change requires that you delete the PPRC pairs, delete the paths, and disable the PPRC function. After you change the SSID, you can reenable PPRC and reestablish the pairs and the paths. With that in mind, make sure the values for the SSIDs are as you want to have them, before the PPRC function is enabled and in use.

3.2.4 Channel Connection Address

Because the RVA is a disk subsystem with a virtual architecture, the CCA is virtual too. Rather than having a CCA, you have a logical CCA and a relative CCA. The logical CCA is related to the functional device ID (FDID) within an RVA. It ranges from 00 to FF for 256 devices. The relative CCA is always relative to the LCU. As you are emulating a 3990-3, the CCA is always between X'00' - X'3F' for each logical control unit.

Note

The CESTPAIR command uses the relative CCA to identify the secondary volume with which you are going to establish a pair.

If the device numbers on your RVA have been generated as a contiguous group of 256 addresses (that is, AA00-AAFF), the relative CCA can be calculated from the last two digits of the device number (that is, 00-FF), as indicated in Table 1.

Table 1. Device Numbers and Relative Channel Connection Addresses

Last Two Digits of Device Number	Relative CCA
'00' - '3F'	x'00' - x'3F', LCU=0
'40' - '7F'	x'00' - x'3F' (subtract x'40'), LCU=1
'80' - 'BF'	x'00' - x'3F' (subtract x'80'), LCU=2
'C0' - 'FF'	x'00' - x'3F' (subtract x'C0'), LCU=3

The logical CCA in the MVS DEVSERV PATHS command is not to be used for the CESTPAIR command.

3.2.5 Director Device Connection

The director device connection (DDC) for the RVA does not have a direct relationship to any specific hardware component. It is a logical numbering of the devices within one LCU. Because the RVA emulates a 3990-3 with up to 64 devices per LCU, it can have four LCUs, each with a DDC range from 00 to 3F. Therefore, the DDC address alone does not allow you to identify a specific logical device within an RVA, until you know the LCU number.

Note

The DDC value is always the same as the relative CCA. Thus you can use the DDC value from DEVSERV PATHS to identify the secondary logical device in the CESTPAIR command.

The examples presented below show the result of the DEVSERV PATHS command for different IOCDS or HCD coding.

The IOCDS or HCD coding shown in Figure 26 provides a contiguous group of 128 addresses (100-17F).

```
*************************************************************************
IOCDS EXAMPLE 1
*************************************************************************
CHPID PATH=20,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=24,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=28,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=2C,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CNTLUNIT  CUNUMBR=520,PATH=(20,24),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
CNTLUNIT  CUNUMBR=521,PATH=(28,2C),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
IODEVICE  ADDRESS=(100,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(520,521),FEATURE=(SHARED),STATDET=Y
CNTLUNIT  CUNUMBR=522,PATH=(20,24),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
CNTLUNIT  CUNUMBR=523,PATH=(28,2C),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
IODEVICE  ADDRESS=(140,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(522,523),FEATURE=(SHARED),STATDET=Y
```

Figure 26. IOCDS or HCD Example: Contiguous Range of 128 Addresses

Figure 27 on page 34 shows an MVS DEVSERV PATHS command issued to devices attached to a control unit, configured as in the example in Figure 26.

```
************************************************************************
MVS DEVSERV PATHS COMMAND EXAMPLE 1
************************************************************************
DS P,100,1

RESULT:
IEE459I 16.20.55 DEVSERV PATHS
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
0100,33903 ,O,000,SYS100,20=+ 24=+ 28=+ 2C=+
     339038  0020  Y  YY.  YY.    N   SIMPLEX   00   00        3990-3
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE

DS P,140,1

RESULT:
IEE459I 16.20.56 DEVSERV PATHS
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
0140,33903 ,O,000,SYS140,20=+ 24=+ 28=+ 2C=+
     339038  0021  Y  YY.  YY.    N   SIMPLEX   40   00        3990-3
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE
```

Figure 27. Sample DEVSERV PATHS Command

The IOCDS or HCD coding shown in Figure 28 provides two address ranges (100-13F and 160-19F). Different device number ranges were assigned to the devices attached to LCU0 and LCU1.

```
************************************************************************
IOCDS EXAMPLE 2
************************************************************************
CHPID PATH=20,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=24,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=28,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=2C,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CNTLUNIT  CUNUMBR=520,PATH=(20,24),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
CNTLUNIT  CUNUMBR=521,PATH=(28,2C),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
IODEVICE  ADDRESS=(100,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(520,521),FEATURE=(SHARED),STATDET=Y
CNTLUNIT  CUNUMBR=522,PATH=(20,24),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
CNTLUNIT  CUNUMBR=523,PATH=(28,2C),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
IODEVICE  ADDRESS=(160,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(522,523),FEATURE=(SHARED),STATDET=Y
```

Figure 28. IOCDS or HCD Example: Two Address Ranges

The MVS DEVSERV PATHS command in Figure 29 on page 35 looks different from that in Figure 27. Because a noncontiguous device address range has been assigned to the devices attached to LCU0 and LCU1, device number 160 is the first device on LCU1. Its physical CCA is 40. The DDC starts with 00 for device number 160, because each LCU range is from 00 through 3F. The relative CCA for device 160 in this example is also 00.

```
************************************************************************
MVS DEVSERV PATHS COMMAND EXAMPLE 2
************************************************************************
DS P,100,1

RESULT:
IEE459I 16.20.55 DEVSERV PATHS
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
0100,33903 ,O,000,SYS100,20=+ 24=+ 28=+ 2C=+
     339038  0020  Y  YY.  YY.    N   SIMPLEX   00   00        3990-3
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE

DS P,160,1

RESULT:
IEE459I 16.20.56 DEVSERV PATHS
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
0160,33903 ,O,000,SYS140,20=+ 24=+ 28=+ 2C=+
     339038  0021  Y  YY.  YY.    N   SIMPLEX   40   00        3990-3
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE
```

Figure 29. Sample DEVSERV PATHS Command: Two Address Ranges

3.2.6 ICE Cards

The ESCON channel adapter cards on the RVA are called ICE cards. They are responsible for handling data traffic to and from the RVA. The RVA can be equipped with four or eight of them. Each ICE card on an RVA has been assigned a name according to its physical plug location. An RVA with four ICE cards has cards installed in locations ICE00, ICE02, ICE10, and ICE12. An RVA with eight ICE cards has additional cards installed in locations ICE01, ICE03, ICE11, and ICE13. Each card has two physical ports, an upper and a lower, to attach a fiber link. Each link is capable of addressing all volumes between a primary RVA and a secondary RVA subsystem. The ICE cards can be configured in either single port or dual port mode.

Note

The term *link* as used in this section, means the port on the ICE card and the attached fiber optic cable.

When configured in dual port mode, both ICE card links are considered active, and the logical paths are spread evenly across all of the links. Dual port mode is the default at installation. In single port mode, one of the two links is inactive and does not consume any logical paths. Switching to single port mode must be specified at installation; otherwise it requires an initial microcode load (IML) of the box.

The links on each ICE card can run in either of the following modes:

- Control unit mode
  - Control unit mode is used to process host requests and to receive data and commands from a PPRC primary RVA.
  - Control unit mode is the normal mode of the links.
- Channel mode
  - Channel mode is used to send data and commands to a secondary RVA.
  - Channel mode is set by using the CESTPATH command for specific links.

Note

Channel mode cannot be set if a logical path from a host is active on that link.

To remove a logical path from a host, use the MVS CF CHPID off command. On shared CHPID in LPAR mode, the command must be issued on each LPAR where the CHPID is defined and active.

Control unit mode cannot be set if PPRC paths are established on that link.

Path status is available from CQUERY PATHS. If you try to establish a PPRC path on a link that is in use by a host, the CQUERY PATHS command reports a status of 0A.

To remove a path, use the CDELPATH command. The CDELPATH command is not accepted if active PPRC pairs are using the path. An error condition does not necessarily terminate the connections.

3.2.7 System Adapter Identifier

To identify each link that is physically available on the ICE cards, each link is assigned a system adapter ID (SAID); see Table 2.

Table 2 (Page 1 of 2). System Adapter IDs

ICE Card Location	Link on ICE Card	Cluster	SAID Bytes
ICE00	upper	0	X'0000'
ICE00	lower	0	X'0001'
ICE01	upper	0	X'0020'
ICE01	lower	0	X'0021'
ICE02	upper	0	X'0040'
ICE02	lower	0	X'0041'
ICE03	upper	0	X'0060'
ICE03	lower	0	X'0061'
ICE10	upper	1	X'0010'

Table 2 (Page 2 of 2). System Adapter IDs			
ICE Card Location	**Link on ICE Card**	**Cluster**	**SAID Bytes**
ICE10	lower	1	X'0011'
ICE11	upper	1	X'0030'
ICE11	lower	1	X'0031'
ICE12	upper	1	X'0050'
ICE12	lower	1	X'0051'
ICE13	upper	1	X'0070'
ICE13	lower	1	X'0071'

The relationship that the SAID and ICE card location have on the channel adapter address depends on the number of ICE cards installed. Figure 30 shows the the channel adapter addresses for a four-card configuration assuming the ports are configured in dual port mode.

	Cluster 0		Cluster 1	
	ICE 00	ICE 02	ICE 12	ICE 10
Upper Ports	SAID=0000 CA=ABEF	SAID=0040 CA=IJMN	SAID=0050 CA=IJMN	SAID=0010 CA=ABEF
Lower Ports	SAID=0001 CA=CDGH	SAID=0041 CA=KLOP	SAID=0051 CA=KLOP	SAID=0011 CA=CDGH

Figure 30. Channel Adapter Addresses for a Four ICE Card Configuration

Figure 31 on page 38 shows the the channel adapter addresses for an eight-card configuration assuming the ports are configured in dual port mode.

	Cluster 0			
	ICE 00	ICE 01	ICE 02	ICE 03
Upper Ports	SAID=0000 CA=AB	SAID=0020 CA=EF	SAID=0040 CA=IJ	SAID=0060 CA=MN
Lower Ports	SAID=0001 CA=CD	SAID=0021 CA=GH	SAID=0041 CA=KL	SAID=0061 CA=OP

	Cluster 1			
	ICE 13	ICE 12	ICE 11	ICE 10
Upper Ports	SAID=0070 CA=MN	SAID=0050 CA=IJ	SAID=0030 CA=EF	SAID=0010 CA=AB
Lower Ports	SAID=0071 CA=OP	SAID=0051 CA=KL	SAID=0031 CA=GH	SAID=0011 CA=CD

Figure 31. Channel Adapter Addresses for an Eight ICE Card Configuration

3.2.8 Logical Path

To communicate with a device, a host channel has to have a logical path to it. The logical path is established on the host at power-on-reset time, depending on IOCDS or HCD coding, by using the physical link from the host down to the device. On the RVA, a primary LCU can communicate with a secondary LCU only if a logical path has been established. To establish a logical path, use the CESTPATH command. A logical path is needed for each LCU connecting to another LCU. The maximum number of logical paths on an RVA is 128. They are spread evenly across the links on the ICE cards. Multiple logical paths can share the same physical link. For each PPRC path that you establish, you lose from 4 to 32 logical paths depending on your configuration.

3.2.9 Channel Register Sets

Channel register sets are available on the RVA. They are spread evenly across the links. In control unit mode, the channel register sets manage the logical paths. In channel mode they manage the primary control unit's connection to a secondary RVA. Thirty-two channel register sets are available on an RVA.

An RVA with four ICE cards, configured in dual mode, provides eight links, with four channel register sets on each (see Table 3 on page 39). If an RVA with eight ICE cards is configured in dual mode, only two channel register sets are available on each link. To have a link on a primary RVA paired with other physical secondary

RVAs, one channel register set is needed for each physical secondary RVA. Thus you can have a maximum of four physical secondary RVAs paired on a link of an RVA with eight ESCON links. An RVA with 16 ESCON links can be paired with two other physical secondary RVAs per link.

When a link is in channel mode, the number of logical paths on the RVA that are available for host access is reduced. The number of available logical paths is reduced by 16, for each link in channel mode on an RVA with eight links. For each link in channel mode on an RVA with 16 ESCON links, the number of available logical paths is reduced by 8.

We recommend configuring the links on the ICE cards in dual mode, to minimize the number of logical paths lost.

Table 3. Number of Logical Paths and Channel Register Sets per ESCON Link

Number of ICE cards	ICE card Mode	Number of ESCON Links	Number of Logical Paths per Link	Number of Channel Register Sets per Link
4	single	4	32	8
4	dual	8	16	4
8	single	8	16	4
8	dual	16	8	2

3.2.10 CESTPATH Command

Use the CESTPATH command to establish logical paths between a primary and a secondary LCU. Each CESTPATH command can establish up to four paths between a primary and a secondary LCU.

Logical paths from a primary LCU to multiple secondary LCUs need separate CESTPATH commands issued, one for each primary to secondary LCU relationship.

Note

The CESTPATH command is a “replace” function. Specified paths that establish a link between a primary LCU and a secondary LCU replace any previously established PPRC paths between the LCU pair. Each time that you issue a CESTPATH command, the paths specified on it replace the paths established by the last CESTPATH command issued. You can use the CESTPATH command to selectively add and remove paths, but use caution to ensure that you do not remove paths that you want to keep.

3.2.10.1 Syntax

The syntax of the CESTPATH command is:

```
CESTPATH -- DEVN(device_number) -- PRIM(ssid serialno)
         -- SEC(ssid serialno) -- LINK(linkaddr)
```

The required parameters are:

- DEVN

  Specifies the four-digit device number of any simplex DISK volume attached to the primary LCU without directing any operation to the volume. You can also specify a PPRC primary volume once a path and pair have been established.

- PRIM

  Specifies the primary logical control unit SSID and serial number. The SSID is a four-digit value as described in 3.2.3, "Subsystem Identifier" on page 32. The serial number is the two-digit plant code followed by the actual serial number of the RVA to which this LCU belongs. Thus a seven-digit value has to be entered. See 3.2.1, "Serial Number" on page 31 to get the value for your RVA.

- SEC

  Specifies the secondary logical control unit SSID and serial number. For a description of the values, and how to get them, see under PRIM.

- LINK

  Specifies the addressing path from the primary logical control unit to the secondary logical control unit. Up to four path addresses can be specified with each LINK parameter. Each path address is an eight-digit hexadecimal address in the form X'aaaabbcc'.

  The LINK parameter address consists of:

  1. A primary RVA SAID (aaaa in this example) that uniquely identifies the physical location of the ICE card link, which physically connects to the secondary RVA or to an ESCON Director.

     Table 2 on page 36 provides the required information.

  2. A secondary logical control unit ESCON link destination address, bb in this example. The destination port on the ESCON Director, where the secondary RVA is physically connected, is assigned to bb. A destination port address is valid only if the primary RVA is dynamically connected to the secondary RVA through an ESCON Director. For static connections through an ESCON Director or point to point connections, the link destination address value is zero. Refer to 3.1.1, "ESCON Connection Modes" on page 23 for detailed information.

  3. A logical control unit number on the secondary RVA, (cc in this example). See 3.2.2, "Logical Control Unit" on page 31 for more details.

     The LCU number is the CUADD as specified in the IOCP or HCD configuration for the subsystem. Each CUADD contains up to 64 devices. Valid CUADDs for an RVA are 0, 1, 2, and 3. Thus cc is filled accordingly with 00, 01, 02, or 03.

If the device numbers on your RVA are generated as a contiguous range of 256 addresses (that is, AA00–AAFF), the LCU can be easily calculated from the last two digits of the device number (that is, 00–FF). Use Table 4 on page 41.

You can also calculate the logical control unit from the output of the MVS DEVSERV PATHS command, taking the CCA value instead of the last two digits of the device number in Table 4.

Table 4. Logical Channel Connection Addresses and Logical Control Unit Number

Logical CCA	Logical Control Unit Number
'00' - '3F'	LCU0
'40' - '7F'	LCU1
'80' - 'BF'	LCU2
'C0' - 'FF'	LCU3

3.2.10.2 Example

In this example, four paths are to be established between a primary LCU and secondary LCU on two RVAs. Figure 32 shows the RVA configuration.

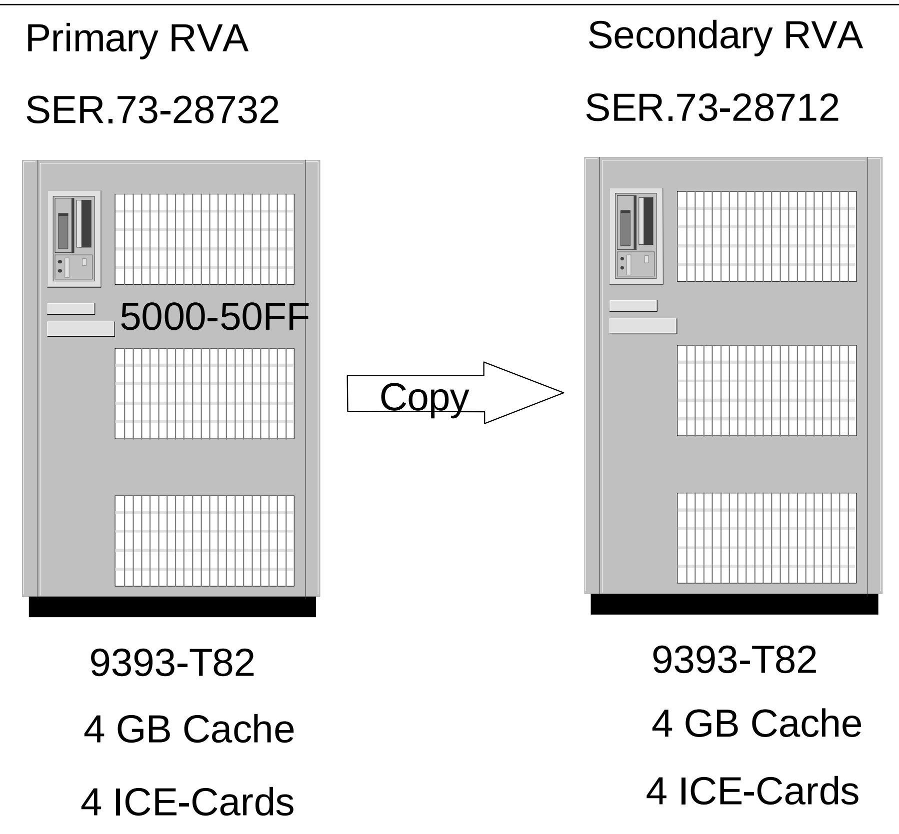

Figure 32. RVA Configuration

The two RVAs are connected by four physical links. Two links are direct connected and two are dynamically switched. Figure 33 on page 42 shows the connection details.

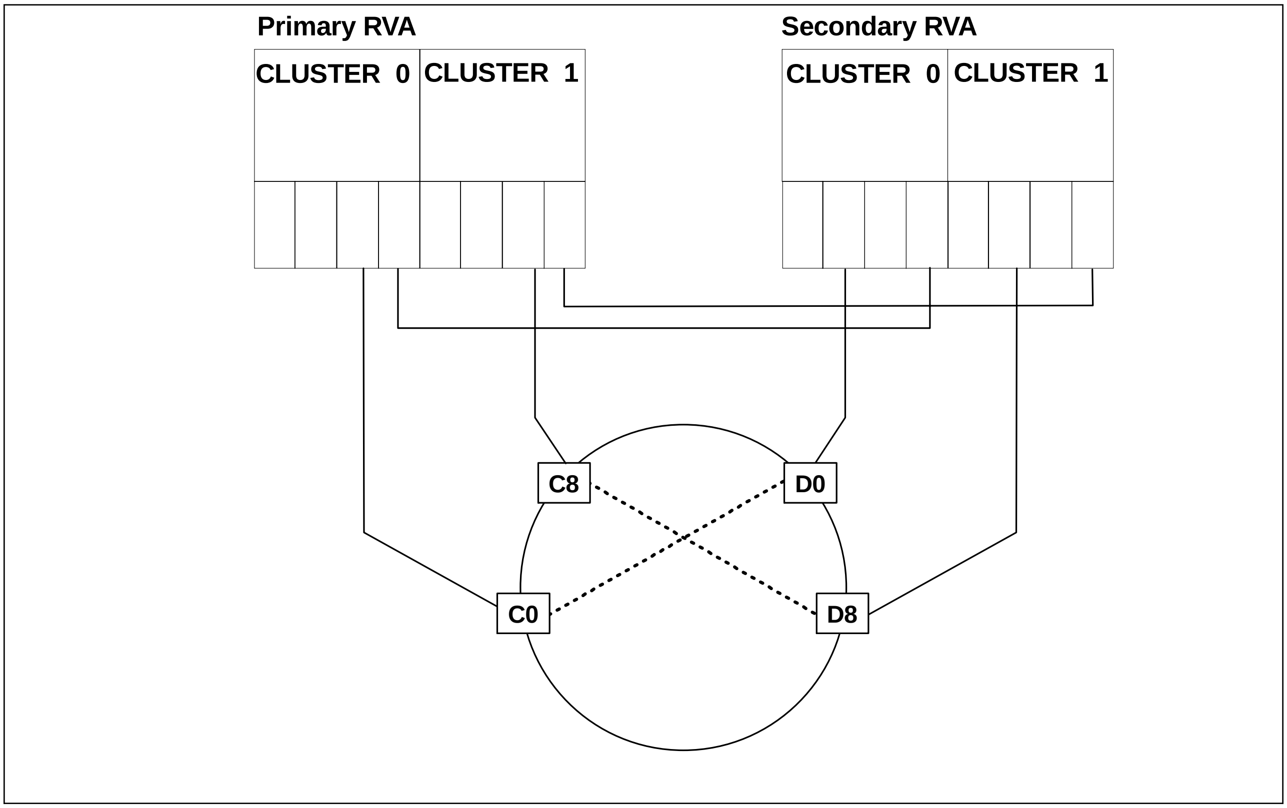

Figure 33. PPRC ESCON Connections

As you can see in Figure 33, on both clusters the second channel interfaces are used for the switched links, and the last interfaces on both clusters are used for the direct links. Assume CHPID 15 to be connected to Cluster 0, first interface, CHPID 29 to Cluster 0, third interface, CHPID 95 to Cluster 1, first interface, and CHPID AD to Cluster 1, third interface.

The IOCDS example in Figure 34 on page 43 shows the coding for the primary RVA with 256 logical devices spread across four LCUs. This coding results in device addresses 5000–50FF. These devices are available on LPAR PROD as well as on TEST. If you are not running LPAR mode on your system, the partition statement is not needed.

```
*************************************************************************
IOCDS EXAMPLE
*************************************************************************
CHPID PATH=15,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=29,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=95,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CHPID PATH=AD,TYPE=CNC,PARTITION=(PROD,TEST),SHARED
CNTLUNIT  CUNUMBR=520,PATH=(15,95),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
CNTLUNIT  CUNUMBR=521,PATH=(29,AD),
      UNIT=3990,UNITADD=((00,64)),CUADD=0
IODEVICE  ADDRESS=(5000,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(520,521),FEATURE=(SHARED),STATDET=Y
CNTLUNIT  CUNUMBR=522,PATH=(15,95),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
CNTLUNIT  CUNUMBR=523,PATH=(29,AD),
      UNIT=3990,UNITADD=((00,64)),CUADD=1
IODEVICE  ADDRESS=(5040,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(522,523),FEATURE=(SHARED),STATDET=Y
CNTLUNIT  CUNUMBR=524,PATH=(15,95),
      UNIT=3990,UNITADD=((00,64)),CUADD=2
CNTLUNIT  CUNUMBR=525,PATH=(29,AD),
      UNIT=3990,UNITADD=((00,64)),CUADD=2
IODEVICE  ADDRESS=(5080,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(524,525),FEATURE=(SHARED),STATDET=Y
CNTLUNIT  CUNUMBR=526,PATH=(15,95),
      UNIT=3990,UNITADD=((00,64)),CUADD=3
CNTLUNIT  CUNUMBR=527,PATH=(29,AD),
      UNIT=3990,UNITADD=((00,64)),CUADD=3
IODEVICE  ADDRESS=(50C0,64),UNIT=3390,UNITADD=00,
      CUNUMBR=(526,527),FEATURE=(SHARED),STATDET=Y
```

Figure 34. IOCDS Example

Figure 35 shows the output from a DEVSERV PATHS command issued to the primary RVA.

```
*************************************************************************
MVS DEVSERV PATHS COMMAND EXAMPLE
*************************************************************************
DS P,5000,1

RESULT:
IEE459I 16.20.55 DEVSERV PATHS
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
5000,33903 ,O,000,SYS100,15=+ 95=+ 29=+ AD=+
     339038  0040  Y  YY.  YY.    N   SIMPLEX   00   00        3990-3
************************ SYMBOL DEFINITIONS ************************
O = ONLINE                         + = PATH AVAILABLE
```

Figure 35. Sample DEVSERV PATHS Command Issued to Primary RVA

Figure 36 on page 44 shows the format of the CESTPATH command that is required to establish the PPRC connections between the primary and secondary LCUs.

```
************************************************************************
CESTPATH COMMAND EXAMPLE
************************************************************************

CESTPATH DEVN(X'5000') PRIM(X'0040' 7328732) SEC(X'0044' 7328712)
         LINK(X'0001D000' X'00410000' X'0011D800' X'00510000')
```

Figure 36. Sample CESTPATH Command to Establish PPRC Connections

Figure 37 shows the output from a CQUERY PATHS command issued to the primary LCU.

```
************************************************************************
CQUERY COMMAND EXAMPLE
************************************************************************
CQUERY DEVN(X'5000') PATHS

RESULT:
ANTP0050I CQUERY PATHS FORMATTED
*************** PPRC REMOTE COPY CQUERY - PATHS ********************
* PRIMARY UNIT: SERIAL#= 000007328732  SSID= 0040                 *
*              FIRST        SECOND        THIRD         FOURTH    *
*            SECONDARY     SECONDARY     SECONDARY     SECONDARY   *
*SERIAL NO: 000007328712  ............  ............  ............ *
*     SSID:     0044          0000          0000          0000     *
*    PATHS:      4             0             0             0       *
*          SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*          --------- --  --------- --  --------- --  --------- -- *
*       1: 0001 D000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       2: 0041 0000 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       3: 0011 D800 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       4: 0051 0000 01  ---- ---- 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 5000. COMPLETION CODE: 00
```

Figure 37. Sample CQUERY Command Issued to Primary Logical Control Unit

3.3 PPRC

To create PPRC pairs between RVA subsystems you must have ESCON links between the primary RVA and the secondary RVA (see 3.1, “ESCON” on page 23 for possible connection modes). At least one ESCON link must exist between RVA subsystems for which you want to set up PPRC pairs. To ensure connectivity we recommend that you have at least one link from each cluster.

> **NOTE**
>
> Any ESCON link is capable of addressing all volumes on an RVA, but you should configure your PPRC environment from a hardware point of view for redundancy. A single link failure should not disrupt your PPRC activities.
> .

3.3.1 Physical Configuration of the PPRC Links

Figure 38 shows a minimum PPRC configuration.

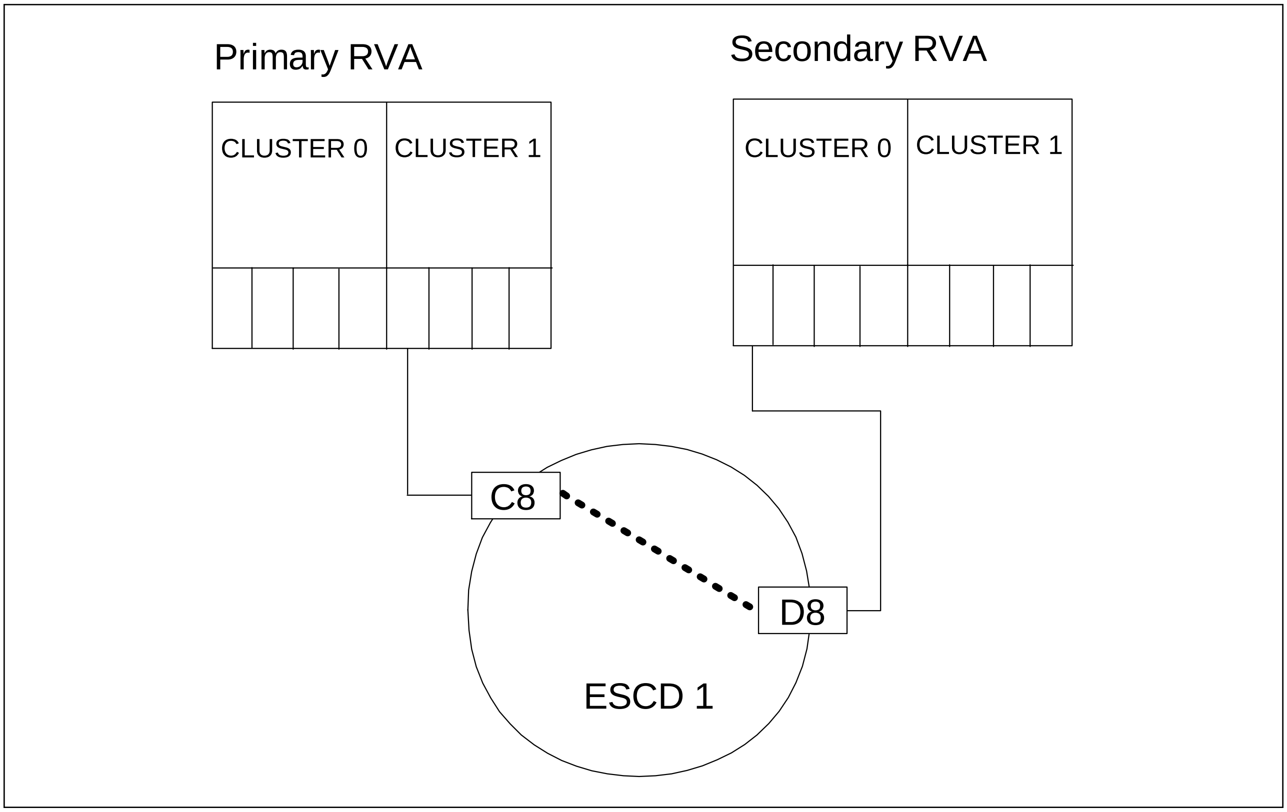

Figure 38. Sample Minimum PPRC Configuration

3.3.1.1 Configuration for Availability

The RVA provides an advanced implementation of RAID-6, which delivers high availability. To get your PPRC implementation to a high availability level as well, consider the following while planning your physical configuration:

- Spread the attached ESCON cables as much as possible over the ICE cards. This is valid for the host links as well as the PPRC links.
- If ESCON Directors are used, spread your links over multiple ESCON Directors.
- If you are using only one ESCON Director, do not use ports on the same card of the ESCON Director for redundant links. Spread the ports across different cards.
- If you are using a single ESCON Director, we recommend installing the availability feature.
- Route the ESCON cables through two different routes.

- Have available spare fibers on your trunk cable.

Figure 39 shows a configuration for availability.

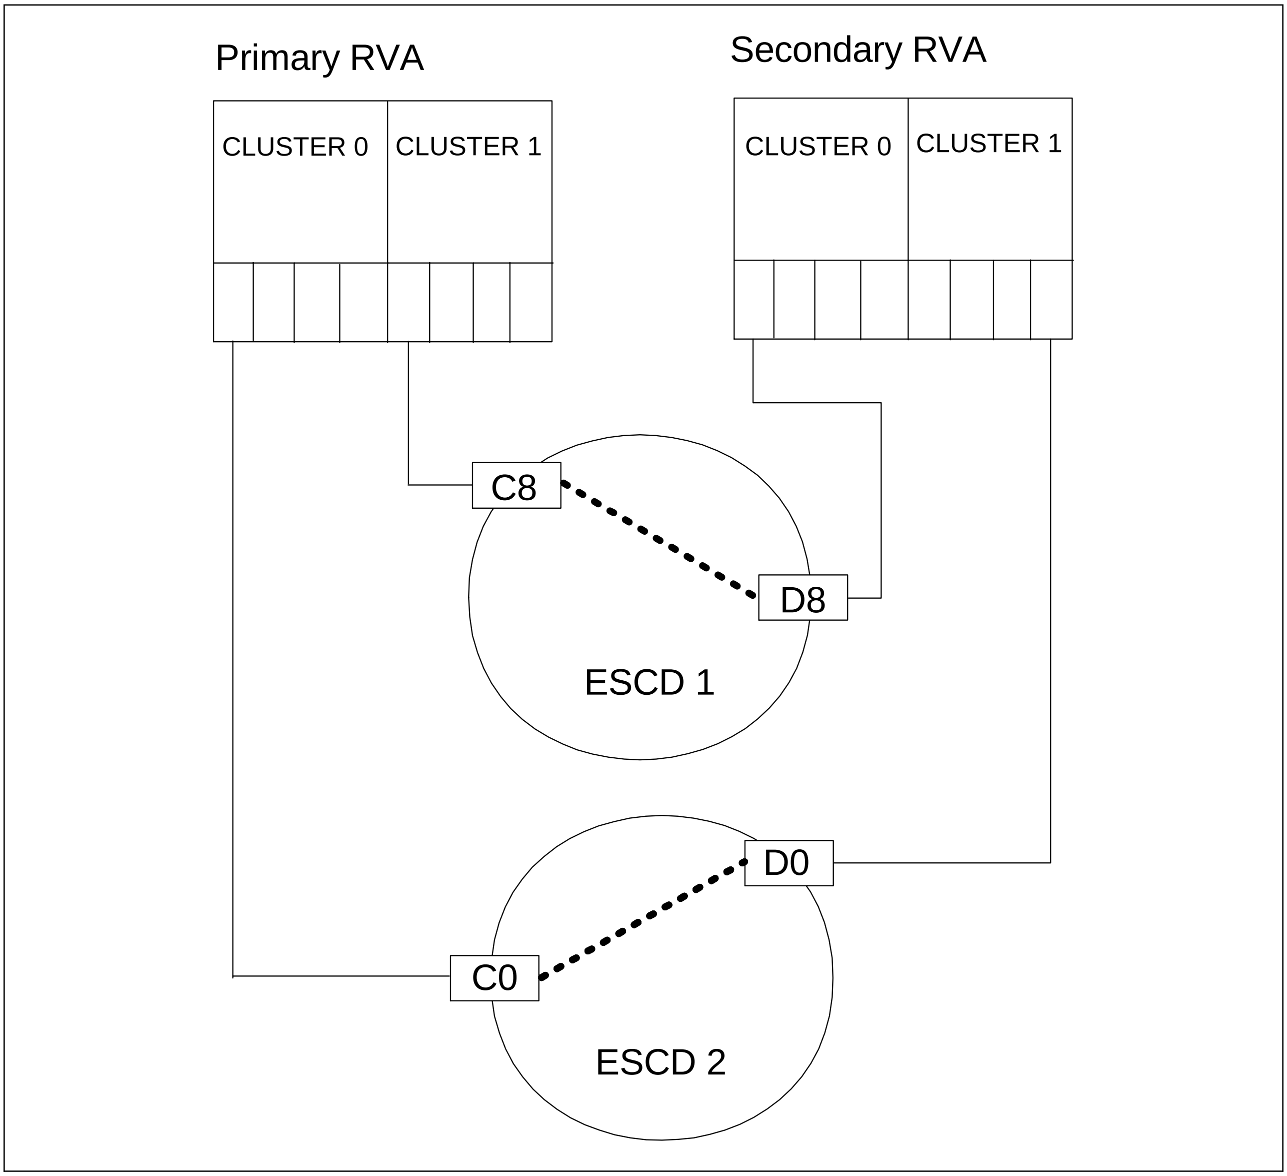

Figure 39. Sample of a PPRC Configuration for Availability

3.3.1.2 Configuration for Performance

Figure 40 on page 47 shows a configuration for performance with multiple links provided between the RVAs.

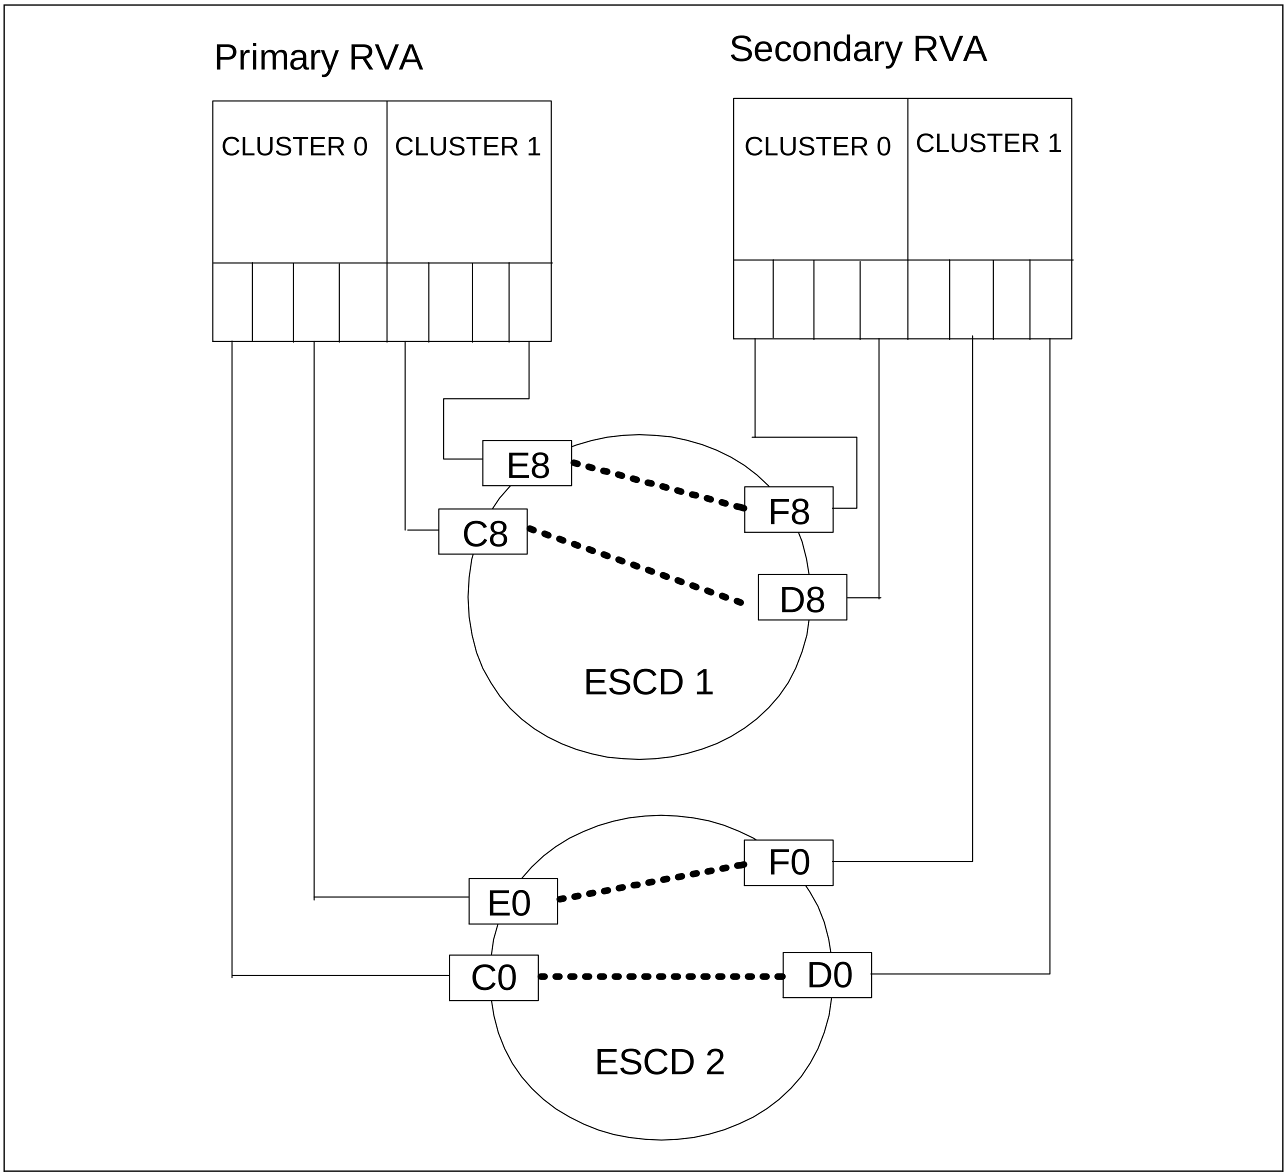

Figure 40. Sample PPRC Configuration for Performance

Of course this configuration could also be cross-configured across both ESCON directors to provide an optimal configuration for both performance and availability.

3.3.2 Logical Configuration of the PPRC Paths

To establish logical paths on the physical links, use the CESTPATH command. See 3.2.10, "CESTPATH Command" on page 39 for command details. The logical paths are established on the LCU level between the primary LCU and the secondary LCU. Figure 41 on page 48 shows the simplest arrangement. We recommend using a one-to-one relationship between the primary and secondary LCUs whenever a one-to-one relationship is possible.

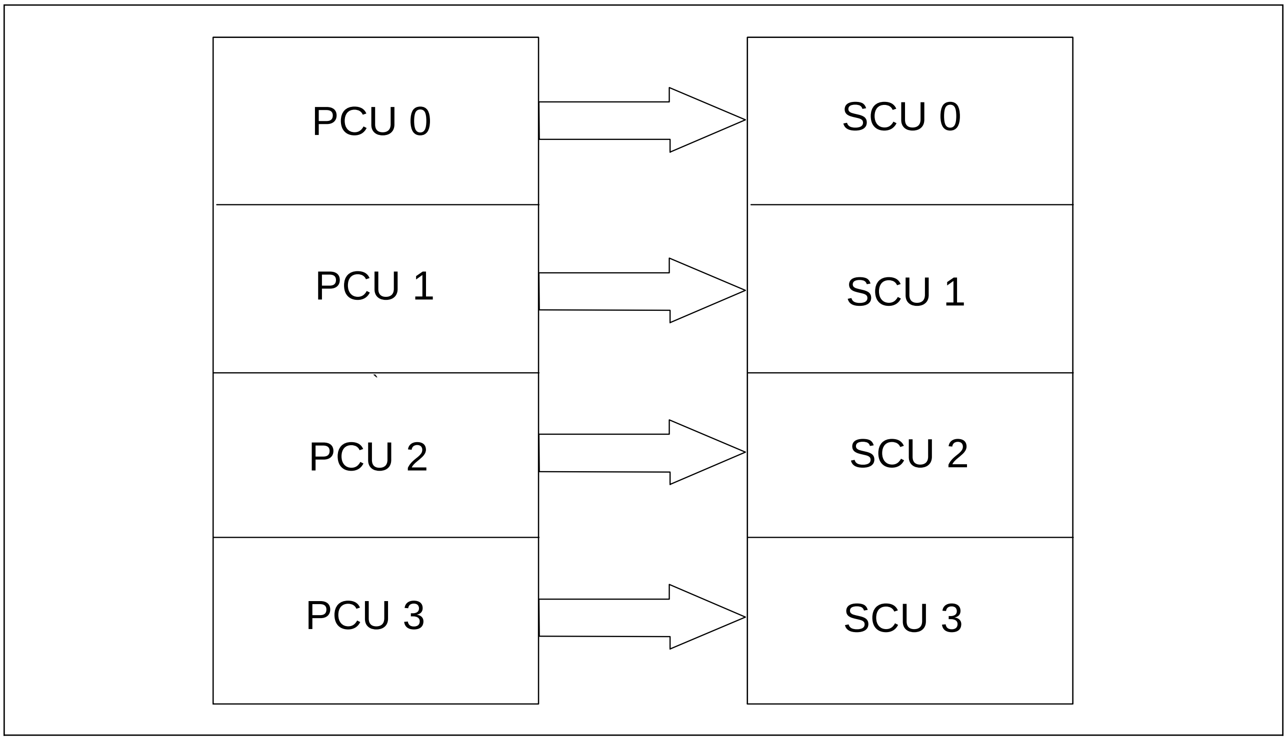

Figure 41. One-to-One Relationship between Logical Control Units

If you need addressability to all 256 devices on the secondary RVA from every LCU on the primary RVA, you need 16 CESTPATH commands, 4 for each primary LCU to secondary LCU relationship, as shown in Figure 42. We do not recommend such a configuration, however.

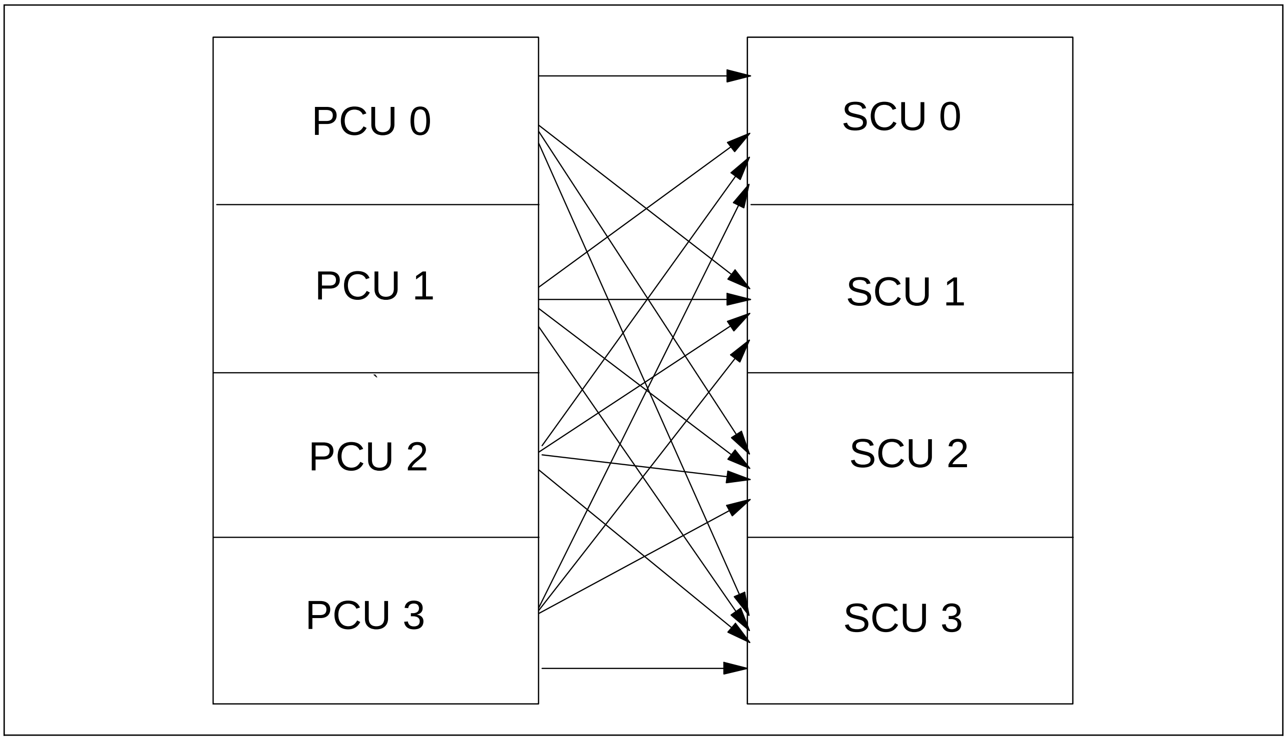

Figure 42. Any-to-Any Relationship between Logical Control Units

If you want to connect two RVAs, each with some primary and some secondary devices, you have to define a PPRC path in each direction. PPRC paths cannot share the same physical links or logical paths.

A bidirectional configuration requires more physical links between the RVAs and may require that you upgrade your RVA to provide more channel adapters, especially if you have an eight ESCON channel configuration. The complexity of managing a bidirectional environment increases, so it is even more important that you put in place good documentation and operational procedures to minimize any confusion.

Chapter 4. Operational Considerations

The operation of PPRC on the RVA is sometimes different from the operation of a 3990-6 or 9390 subsystem. In this chapter we explain the characteristics of the RVA and how they affect the operation of PPRC.

4.1 Hardware Characteristics

The following hardware characteristics may affect the operational support of PPRC on the RVA:

- The RVA supports 3380 as well as 3390 device types.
- PPRC on the RVA is not affected by subsystem cache settings, NVS, volume caching, or DASD fast write settings.
- We recommend attaching a host to the secondary subsystem. For operational support of a secondary subsystem, you need an ECAM device. Because a PPRC secondary volume can not be an ECAM device, the subsystem is limited to a maximum of 255 PPRC pairs.
- When the RVA generates a SIM, it is sent down every active path group. Thus ERP SIMs are issued to every connected host.
- As there are no host connections through the PPRC links, the failure of a single path is not reported. On the 3990 the failure of a single link is reported, if there is host access on the same link, by message IOS581E.
- Using PPRC on an RVA reduces the number of logical paths available for host I/O. See Chapter 3, “Configuration” on page 23. PPRC may also reduce the number of physical ESCON channels that are available for host I/O. If you do not have sufficient channel adapters available, you may not be able to drive the maximum concurrent I/Os through the RVA.
- You cannot obtain pathing information for the RVA through the ICKDSF ANALYZE command.
- There are no PPRC switches on the RVA. On 3990 or 9390, the VPD is altered from the storage controller's installation and configuration panels. You therefore cannot:
  - Alter the setting of CRIT(YES). The setting is always CRIT(YES-paths). You cannot have CRIT(YES-all).
  - Change the consistency grouping timer. This switch specifies the time in seconds that the long busy state is issued for a device. On the RVA this duration is always 120 sec.
  - Change consistency grouping on or off. This is controlled through the CGROUP parameter of the CESTPATH command.
- If you are running in a GDPS you need on utility device per logical control unit. This arrangement reduces the number of devices per logical control unit available for PPRC pairing. We also recommend that you do not share the GDPS utility device between MVS systems, to prevent the utility device from being RESERVED to one system when you are trying to issue FREEZE commands from another system.

4.2 Command Support

Initiation and support of PPRC on the RVA use the same TSO and ICKDSF commands as the 3990. These commands are detailed in the latest version of the *Remote Copy Administrator's Guide and Reference*, SC35-0169. However, be sure to consider the following points when building commands:

- For commands that require an SSID, you must specify the SSID of the new LCU.
- For commands that require a CCA, the CCA is relative to the LCU and will always be a value between x'00' and x'3F'.
- For commands that need the machine serial number, the serial number on an RVA is seven digits (not five).
- Some parameters are ignored.
- Standard ICKDSF restrictions for the RVA apply.

4.2.1 CESTPATH: Establishing Paths

The CESTPATH command is used to establish logical paths between the primary and secondary LCUs.

If you are going to establish pairs for all 256 devices on a primary RVA subsystem, you must issue at least four CESTPATH commands, one for each LCU. This is true if all the logical control units are in the same secondary RVA subsystem or in multiple secondary RVA subsystems. In theory you can have up to 16 logical paths from a single logical control unit. Therefore you could have up to 64 logical paths defined on a single RVA subsystem for PPRC. We do not recommend this configuration, however. Keep your configuration as simple as possible.

A new command parameter, CGROUP, has been created for the RVA to enable or disable the PPRC consistency grouping function. On a 3990-6 and 9390 this function is controlled through a VPD switch on the CE panel. Consistency grouping provides the ability to temporarily queue I/O to all PPRC volumes on a single control unit pairing when an error occurs. The I/O to the volumes is queued by issuing a long busy state on an IEA494I message with the state “SUSPENDING” and text “EXTENDED LONG BUSY STATE.” Using automation triggered from the message, you can suspend all pairings on one or more LCU pairings, using the CGROUP command while queuing the I/O temporarily with long busy states.

By specifying CGROUP(Y) you enable this function. The default is CGROUP(N). The CGROUP parameter is valid only on an RVA subsystem as it replaces the VPD switches. ICKDSF does not support this function.

If you have system-related volumes in PPRC pairings, we recommend that you place them on an LCU pairing with CGROUP(N) and CRIT(NO). With these settings, the devices will not go into a long busy state or issue a unit check. System-related volumes should also be excluded from any automation that issues the CGROUP FREEZE command.

Table 5 on page 53 shows the interactions of the CRIT and CGROUP parameter settings.

Table 5. Interactions of CRIT and CGROUP Parameters

	CRIT(Y)	CRIT(N)
CGROUP(Y)	• IEA494I with long busy state (SUSPENDING) • Pair in error suspended • Unit check on primary device • IEA494I when suspended	• IEA494I with long busy state (SUSPENDING) • Pair in error suspended • No unit check on primary device • IEA494I when suspended
CGROUP(N)	• No long busy state • Pair in error suspended • Unit check on primary device • IEA494I when suspended	• No long busy state • Pair in error suspended • No unit check on primary device • IEA494I when suspended

The IEA494I message is issued when a device has changed status. When the consistency grouping function is enabled through the CGROUP(Y) parameter, this message shows an additional state, “SUSPENDING,” and indicates that it is in the process of changing state from duplex to suspended. This message also shows “EXTENDED LONG BUSY STATE.” The IEA494I message with “SUSPENDING” state is issued only when an I/O is attempted to the device.

The IEA494I message with “SUSPENDED” state is issued for the in-error pair to all LPARS that have the primary volume online.

The IEA491E message is returned to the issuing system when the next I/O is attempted to the volume that is suspended.

When you issue the CGROUP FREEZE command to a device defined with CRIT(Y), the parameter is reset to CRIT(N).

4.2.2 CESTPAIR: Establishing Pairs

The CESTPAIR command is used to specify the primary and secondary volumes of a pair and initiate the copy process. Unlike PPRC on a 3990 controller, all establish pair commands issued to an RVA are attempted in parallel.

Two parameters are different from 3990/9390 when using PPRC on an RVA subsystem:

- The PACE parameter is ignored. See 6.3, “Performance during Establish” on page 78.
- The CRIT(YES) parameter is implemented only in the CRIT(YES-paths) mode.

If a secondary device cannot be accessed because paths are not available, and the pair was established with CRIT(Y), the pair is suspended. The IEA491E message is issued with a “SUSPEND” state to all LPARs that have the primary of the failing pair online. The primary devices are put in write-inhibit mode, and all subsequent writes to the primary are unit checked.

If there is no path between the primary and secondary subsystem, the same effect is seen on any other I/O to primary devices on that LCU. The PPRC pairs are suspended, and the ERP issues the IEA494I message with a state of "SUSPENDED" for each pair as it is suspended.

If you establish pairs with CRIT(N), a critical volume state unit check is not issued. The long busy state of the suspending pair allows you to halt your applications (and prevent dependent writes) long enough for your automation to detect that there is a problem and freeze all PPRC pairs to ensure consistency of the secondary site. Updates to the primary devices continue, so once links to the secondary site are restored you must CESTPAIR(RESYNC) all PPRC pairs.

The objective of CRIT(Y) is to stop the applications to ensure that no updates are written to the primary and therefore to the secondary sites. This is intended to preserve the consistency of your secondary site, if there is rolling disaster. All updates to the primary LCU are prevented, so there is an availability impact on your primary applications.

CRIT(Y) is not effective for all applications. For example, if your application spans multiple RVAs, some updates may continue on other RVA subsystems if the application does not completely stop. DB2, for example, does not always stop because one device is unavailable for writing. See 8.3, "Using PPRC with DB2" on page 88 for details.

For complete recoverability of your secondary site, we recommend that you exploit the capabilities of the CGROUP command to ensure consistency.

The CESTPAIR command cannot check if the secondary device:

- Is online or offline
- Is in use by another host
- Contains valid data

If you are unable to establish a pair, or if you get a suspend message and cannot determine the cause from CQUERY, we recommend that you analyze error log records at the secondary site to determine the cause of the error.

Detailed procedures, both manual and automated, must be used for all aspects of your PPRC operation.

4.2.3 CSUSPEND: Suspending Pairs

The CSUSPEND command is used to suspend PPRC operations between a volume pair.

The PRIMARY parameter is ignored for an RVA subsystem. The parameter allows ICKDSF media maintenance and is not applicable on an RVA.

The QUIESCE parameter is not supported for an RVA subsystem. Use the CGROUP command instead.

4.2.4 CQUERY: Querying Status

The CQUERY command is used to query the status of one volume of a PPRC pair or the status of all paths associated with the LCU for the named device number.

Table 6 lists the new path status conditions that the CQUERY command returns.

Table 6. New CQUERY Path Status Condition Codes

Code	Description
07	SSID mismatch
08	ESCON link is offline
09	Establish failed but will retry again when conditions change
0A	System adapter has a host path already established
0B	Path cannot be connected in the same cluster

You can establish PPRC paths between logical subsystems on the same RVA subsystem provided that the link is to the other cluster. When you issue a CQUERY command, a path status of x'0B' is returned if you attempt to establish a PPRC path on the same cluster.

The setting used on a particular CESTPATH command for CGROUP is reported on the CQUERY VOLUME command.

If a secondary device is suspended, the primary is not automatically notified. The next write attempt from the primary detects the error condition and suspends the primary. In your automation you should check the status of both primary and secondary volumes before any operations. In some cases a secondary may be suspended, while the primary is still in a duplex state because no I/O has been attempted to the secondary. If the cause of the secondary suspension is not clear, for example, CQUERY returns a status that cannot be understood, we recommend that you analyze error logs at the secondary site to determine the cause of failure.

4.2.5 CGROUP: Controlling Volume Groups

The CGROUP command controls operations for all PPRC volume pairs on a single primary and secondary LCU pairing. A single LCU can be paired with up to four other LCUs. Therefore, you may have to issue up to four CGROUP commands to suspend all pairs on a single primary LCU. For an RVA PPRC implementation, the CGROUP command relates to a LCU.

The CGROUP FREEZE command is not supported by ICKDSF.

The function enables you to maintain secondary volume update consistency and is particularly important in maintaining synchronization in a database subsystem and the ability to ensure restart at the secondary site in the event of a disaster.

You can use this command effectively by combining it with the CESTPATH command parameter, CGROUP. The CGROUP command with the FREEZE parameter queues all I/O to both primary and secondary volumes in the related LCU pairings, placing the volumes in SUSPENDED status. This command has to be issued for all LCU pairings containing related PPRC pairs. Once all of the commands have completed, you can use the CGROUP command with the RUN parameter to release the queued I/O and resume normal operations. However, if the command is not done within a control unit extended long busy state timeout window, all I/O is released independent of the command being issued. On the RVA this timeout window has been set to 2 min.

As with operations on a 3990, you must be aware of the potential impact of the CGROUP command on host systems if I/O is suspended to related system disks. As this command is issued from TSO, you could see related CGROUP commands not completing before the timeout. This failure would leave you exposed to corruption of your secondary site integrity. To avoid this problem you can split system disks onto a separate CESTPATH command definition and not issue a freeze for that particular logical control unit pairing.

The pairs remain suspended until the paths are reestablished with the CESTPATH command and the pairs are synchronized with the CESTPAIR command.

GDPS uses various Sysplex, automated operations manager (AOM), and DASD ERP messages to trigger a CGROUP FREEZE of all secondary site volumes. GDPS uses the CGROUP FREEZE/RUN capability to provide an integrated hardware and software solution to manage data integrity in a multisite environment. Using GDPS, you can maximize your ability to provide a database restart in the event of a planned or unplanned site switch.

4.3 ICKDSF

ICKDSF Release 16 plus PTFs supports PPRC operations in MVS, VM, VSE, and stand-alone environments.

The PPRC functions are supported through the PPRCOPY command. This command uses the DELPAIR, ESTPAIR, DELPATH, ESTPATH, SUSPEND, RECOVER, and QUERY parameters.

ICKDSF does not support consistency grouping. Therefore you cannot request any long busy states on failing devices or freeze control unit pairings, and thus you cannot use the CGROUP command or parameter.

Because of its internal architecture the RVA ignores media maintenance functions. Such functions as INSPECT have no meaning as the volumes the system sees are functional, not physical, volumes.

ICKDSF provides a subset of MVS PPRC commands for stand-alone use when no operating system is available. This subset could be used to issue RECOVER commands at a secondary site. When running standalone there is no ERP support, so any messages or processes relating to PPRC are not instigated.

The CONTROL CONFIGURE(DISPLAY) command is not supported on the RVA; use PPRCOPY QUERY instead.

When MVS runs as a VM guest, the related VM directory entry must have the STDEVOPT control statement that tells the control program (CP) that VM is authorized to issue PPRC commands. PPRC implementation in a VM environment is discussed in Chapter 7, “Using PPRC on the RVA with VM/ESA and VSE/ESA” on page 83.

For further information about ICKDSF and PPRC, refer to *ICKDSF R16 Refresh User's Guide*, GC35-0033.

4.4 SnapShot

The RVA enables you to use SnapShot functions with PPRC pairs. For considerations regarding the interaction of SnapShot and the volumes in a PPRC pair, see Chapter 9, “Using PPRC with SnapShot” on page 95.

4.5 Automation

Because there are some small differences between PPRC operations on the 3990 and on the RVA, you must review your operational procedures, either automated or manual.

4.5.1 PPRC Migration Manager

The PPRC Migration Manager uses an ISPF interface to give you an automation tool that enables you to interrogate DASD controllers and build the JCL and commands to establish and manage PPRC pairs.

The PPRC Migration Manager code is intended as a tool for migrating data but can be used as a building block for automating a disaster recovery solution using PPRC.

The PPRC Migration manager function is divided into two sections: PPRC global functions and PPRC cluster functions.

The PPRC global functions develop jobs to:

- Create a label and volume table of contents (VTOC) on each secondary volume
- Put the secondary 3990 subsystem cache and NVS and device cache and DASD fast write (DFW) into the correct state for PPRC operation
- Collect data needed to build the PPRC TSO commands
- Verify that PPRC pairs are active
- Return PPRC pairs to simplex
- Verify simplex status
- Clip simplex status
- Reformat VTOCs
- Overwrite every track on volumes with manufacturing data to “erase” any user data

The PPRC cluster functions manage logical clusters of devices by developing jobs to:

- Establish paths
- Establish PPRC device pairs
- Delete PPRC pairs
- Suspend PPRC pairs

- Synchronize suspended PPRC pairs
- Query PPRC pairing status

Some of these functions are not relevant for an RVA.

New functional support has been added to the PPRC Migration Manager for the RVA:

- It automatically recognizes whether it is developing jobs for use with an RVA.
- It builds all jobs with seven-digit serial numbers. For 3990s, the leading two digits are 00.
- The panels used to define PPRC links are matched to the terminology used for either the RVA or 3990.
- The prerequisites to using PPRC with the RVA are different from those required with a 3990.

The information required to complete the LINK DETAIL data entry panel is described in Chapter 3, “Configuration” on page 23, although much of the detail is obtained by the automation from the DSTAT00A program. This program is a part of the PPRC Migration Manager. The code that processes this panel checks to make sure the link information is consistent. It cannot check the ESCON link values. An error in selecting the RVA ICE interface or in specifying the correct ESCON Director port address will be detected only when you actually attempt to establish the PPRC links.

The PSTAT00B program is also a part of this migration process and is used to check on device status.

How to get the PPRC Migration Manager

The DSTAT00A and PSTAT00B programs and the PPRC Migration Manager documentation are available in softcopy from the Internet.

Point your Web browser to:

`ftp://www.redbooks.ibm.com/redbooks/sg245338`

Alternatively you can go to:

`www.redbooks.ibm.com`

select **SEARCH** and **Additional Redbook Materials**. The following program files must be transferred in binary format:

- VMIGCNTL XFRSEQ
- VMIGPNLS XFRSEQ
- VMIGREXX XFRSEQ

The documentation can be found in:

- VMIGROOT HTML

4.5.2 Error Recovery

Error messages are issued to alert you to any change in the status of your PPRC volumes. These error messages can be intercepted and acted upon by automation, either using your own automated operations applications or the GDPS service offering. You can use the CGROUP command to control error situations and to maintain consistency of your secondary site.

In a mixed environment of RVA and 3990, you have to consider creating procedures as standard as possible to cover both types. Keep procedures as simple as possible. The more complex the procedures or design, the more likely they are to fail or be misinterpreted.

By specifying CGROUP(YES) on the CESTPATH command, you are requesting the ERP to issue an "Extended Long Busy State" on message IEA494I to all attached systems. The IEA494I message indicates that the pairings are changing status.

If a system attempts I/O to the primary volume while the status is changing, it will receive the IEA494I message with the status of "SUSPENDING." When the pairs actually suspend, the IEA494I is broadcast to all attached systems where the primary volume is online with the status "SUSPENDED."

Messages and environmental triggers can be detected by automation and used to suspend secondary PPRC operations with the CGROUP command and the FREEZE parameter. The FREEZE command temporarily queues all I/O to both the primary and secondary devices, removes the paths, and suspends the pairs. Once automation detects channel end and/or device end from all the commands, it can issue the CGROUP command with the RUN parameter. This command releases all I/O to the primary device but leaves the secondary devices suspended at a point in time. The 2 min timer of the extended long busy state gives automation enough time to issue a CGROUP FREEZE command. I/O resumes against the devices after 2 min if a CGROUP RUN command is not received.

The IEA491E message indicates that a PPRC pair is suspended and gives a reason why the suspension occurred. IEA491E is returned to the issuing system when the first I/O is attempted to the primary volume of the suspended pair.

In your automation procedures you can dedicate one device per LCU as a utility device, which is used as the target of the PPRC commands. Having a dedicated utility device ensures that the device is immediately available to direct commands to and will not be reserved to another system. This is especially important where automation is controlled from an automation or network management focal point system rather than from the normal production systems.

Figure 43 on page 60 shows an example of such an automation cycle. In this example CRIT(Y) is coded on the CESTPAIR command, and CGROUP(Y) is coded on the CESTPATH command.

1. The controller detects an error. The controller goes into a long busy state.
2. The DASD ERP detects that the controller is in a long busy state because of a PPRC pair suspending. A unit check is issued.
3. The DASD ERP issues an IEA494I message with a long busy state.

4. Automation triggers a CGROUP command with the FREEZE parameter to all LCUs in the target freeze set. All I/O is temporarily queued for each targeted LCU pairing.
5. The last channel end is detected from all commands.
6. The CGROUP command with the RUN parameter is issued to release the I/O.

Normal operations resume on primary volumes now in suspended status.

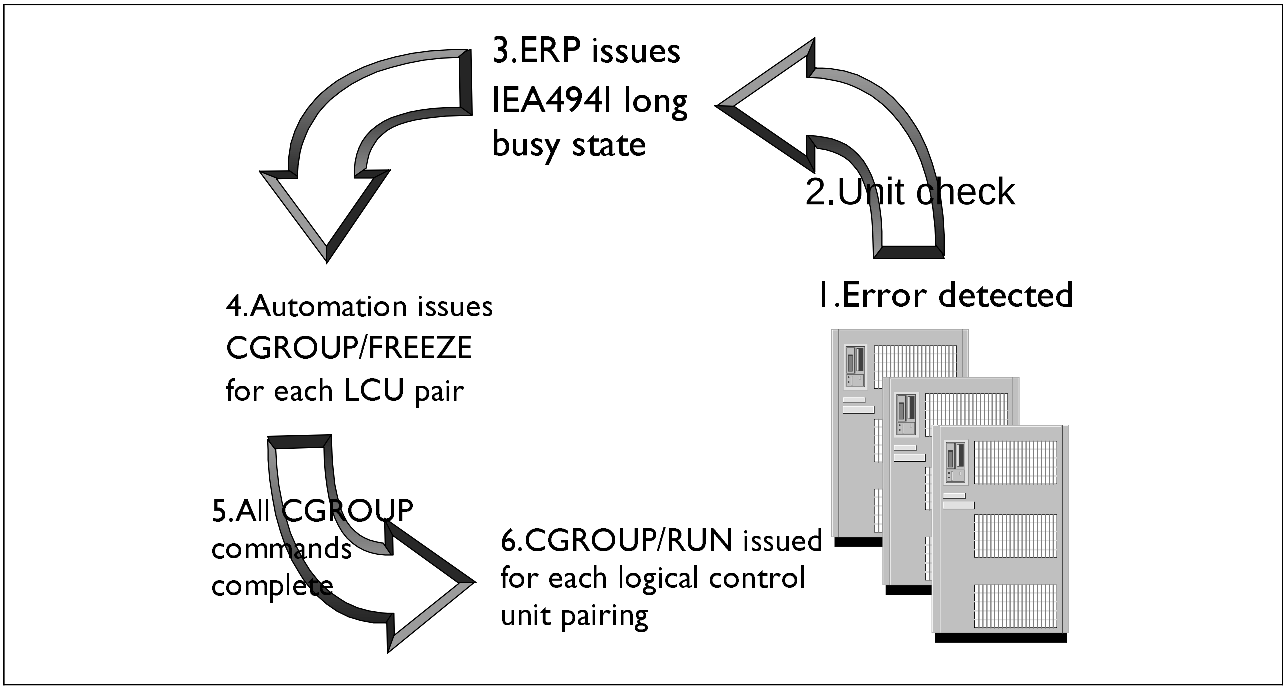

Figure 43. Sample CGROUP Automation Cycle

The way in which you implement the CGROUP command depends on:

- The impact the command has on primary systems if errors occur
- The reasons why you implemented PPRC
- The type of applications you are running

If in the event of a disaster you want to be able to perform a database restart rather than a database recovery, you must consider suspending all devices to a point in time and avoid the rolling disaster scenario. This decision will have an impact on the primary systems while the pairs are being suspended.

You have the option with GDPS to choose the best policy for your business requirement:

- FREEZE and go; resume processing on the primary system once pairs are frozen.
- FREEZE and stop; stop processing on the primary system until the cause of the FREEZE is identified.
- FREEZE and stop conditionally, choose whether or not to continue processing on the primary system depending on the reason for the FREEZE.

4.5.3 The Rolling Disaster

By suspending I/O to all PPRC pairs you maintain write sequence integrity. This integrity is particularly important in an environment that uses logging such as in IMS and DB2.

In a real disaster you will more than likely see your system fail over a period of time. The PPRC pairs do not want to reflect this. If your PPRC pairs fail over time, you will end up with data out of synchronization and a lengthy recovery using tape backups and database image copies. If your procedures allow the rolling disaster scenario to occur, you are wasting the investment in a DASD-based, real-time remote copy solution.

Keep in mind that, even if you have perfect recovery documentation, the secondary system will be useless if all of the required volumes are not included in the PPRC pairings or are not recoverable from other sources. We discuss this issue in Chapter 2, “Planning for Remote Copy” on page 5.

4.5.4 Error Recovery in a GDPS

GDPS, the multisite application availability service offering, uses storage subsystem triggers, host triggers including SYSPLEX messages and heartbeat, cross coupling facility (XCF) functions, and environmental triggers generated by the Automation Operations Manager. These triggers are detected by the GDPS automation code. This service offering provides a complete multisystem monitor and site resource switching capability from a single point of control.

4.5.5 Remote Copy Management Facility

The remote copy management facility (RCMF) is a subset of the GDPS service offering that provides an ISPF panel front end and automates various PPRC functions. RCMF provides a single point of control for the overall management of PPRC operations across multiple sites. RCMF is an IBM service offering.

4.5.6 Monitoring Net Capacity Load

Monitoring net capacity load (NCL) on the secondary RVA should be included in your installation's procedures, particularly where multiple LCUs or volumes are in PPRC pairs to a single RVA or logical control unit.

All data is physically copied from a primary to a secondary device, so data copied on the primary device with SnapShot will be physically copied to the secondary subsystem and not get the space-saving benefit of SnapShot. Therefore the NCL may be higher than expected on the secondary RVA subsystem. We discuss this issue further in 9.5, “NCL on Secondary RVA Subsystems” on page 100.

4.5.7 Managing Secondary Devices

We recommend that you put all of your PPRC secondary volumes in the IXFP device exclusion table. The device exclusion table removes volumes from direct IXFP processing; otherwise IXFP initialization or recycle causes an IXPF message to be issued for each secondary device.

All DDSR activity is mirrored on the secondary volume of a PPRC pairing as part of the PPRC process.

4.5.8 Processing CQUERY Output

If the CQUERY command is issued from a TSO account, the command output is returned to the user and the MVS operator console. If the command is issued from a batch job running IKJEFT01, the command output is returned to the SYSTSPRT DD statement and the MVS operator console. When automating, use the UNFORMAT parameter (Figure 44) to reduce the amount of output generated and therefore the number of lines parsed to the routines.

```
ANTP0025I CQUERY UNFORMATTED
 0101,PRIMARY,PENDING,ACTIVE,
 0010,01,000003234333,1010,01,000003254333,N,N,
 4,00050000,01,00150000,01,0004CE00,01,0014CF00,01,
 00000,00000,000,
```

Figure 44. Unformatted CQUERY Output

As the serial number field on this output is 12 digits, no fields will move position to accommodate the seven-character length of the RVA serial number. However, if your automation is expecting a five-digit number and uses a fixed substring length to extract the information, you have to alter your substring length to seven. All PPRC commands will work with seven digits defined, so leading zeros are acceptable.

The CQUERY output now has an extra field indicating the setting of the CGROUP subparameter of the CESTPATH command.

When processing the output from an unformatted CQUERY, be aware that the output is displayed in a fixed order but not a fixed position. If a field has no current value, it is not padded with blanks or zeros. Each field is delimited by a comma. Figure 45 shows an example of a CQUERY command with empty fields.

```
ANTP0025I CQUERY UNFORMATTED 365
 2200,PRIMARY,PENDING,ACTIVE,
 FF24,00,000001321029,2300,00,000001323562,N,N,
 1,00002B00,01,,,,,,,
 01747,03338,053,
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 2200. COMPLETION CODE: 00
```

Figure 45. Unformatted CQUERY Output with Empty Fields

Figure 46 on page 63 shows an example of how you can process the unformatted output of a CQUERY command. The PARSE command takes a line of the output contained in the RECORD variable and creates three new variables, FIELD_1, FIELD_2, and THE_REST. It takes the first and second output fields delimited by commas and places them in FIELD_1 and FIELD_2. The remaining information, commas included, is placed in THE_REST.

```
PARSE VALUE record WITH  field_1 ',' field_2 ','  the_rest;
```

Figure 46. REXX PARSE to Process Unformatted CQUERY Output

Figure 47 is an example of formatted output from the CQUERY command. It shows the new field reporting which setting was used for the CGROUP parameter on the related CESTPATH command.

```
ANTP0030I CQUERY VOLUME FORMATTED 079
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                         (PRIMARY)   (SECONDARY) *
*                                          SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 2200  PRIMARY..  PENDING...   ACTIVE..    FF24 00      2300 00   *
*      CRIT(NO).......       CGRPLB(NO). 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   1   0000 2B00   01    PATH ESTABLISHED...                      *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
* IF STATE = PENDING/SUSPEND:  FIRST CYL OUT OF SYNC =  01419      *
*                               LAST CYL OUT OF SYNC =  03338      *
*                              PERCENT OF COPY COMPLETE = 043%     *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 2200. COMPLETION CODE: 00
```

Figure 47. Formatted CQUERY Output

4.5.9 Monitoring Pair Status

As well as having event-driven automation, we recommend that you use timer-driven automation to check the status of PPRC pairs. Timer-driven automation picks up on problems that may have occurred when automation was unavailable or during the establish process. It also identifies devices that have been added to your environment but not added to the PPRC configuration.

As with all automation, this timer-driven process is easy to manage if you have good naming standards for volumes. This process could also run at the systems managed storage (SMS) storage group level.

The PPRC Migration Manager provides processes to do this monitoring.

In Figure 48 on page 64, IDCAMS DCOLLECT is used to provide a list of volumes that meet a specific naming standard. In this case, any volume with a prefix of *P* will be selected.

```
//DCOLLECT EXEC PGM=IDCAMS,TIME=1439
//SYSPRINT DD SYSOUT=A
//OUTDS    DD DSN=&&DCOLL,
//            DISP=(NEW,PASS),
//            SPACE=(TRK,(1,2),RLSE),
//            DCB=(DSORG=PS,RECFM=VB,LRECL=264)
//SYSIN    DD *
  DCOLLECT -
    OFILE(OUTDS) -
    VOLUME(P*) -
    NODATAINFO
/*
```

Figure 48. IDCAMS DCOLLECT Deck

The output can then be processed by a REXX routine.

Figure 49 shows an extract of a REXX routine that reads the output from a DCOLLECT step into variables and then outputs the generated JCL. Not all of the variables in the example would be used to generate PPRC jobs.

```
/* REXX EXAMPLE CODE TO PROCESS DCOLLECT OUTPUT */

"EXECIO * DISKR INPUT (STEM RECORD." /* Read from INPUT DD */

DO I = 1 TO RECORD.0

    VOLUME_NAME = SUBSTR(RECORD.I,25,6);
    DEVICE_TYPE = SUBSTR(RECORD.I,69,4);
    FREE_PERCT  = C2D(SUBSTR(RECORD.I,36,1));
    SGNAME_LGTH = C2D(SUBSTR(RECORD.I,81,2));
    STORE_GROUP = SUBSTR(RECORD.I,83,SGNAME_LGTH);
    DEV_FREESPC = C2D(SUBSTR(RECORD.I,37,4));
    DEV_ALLOC   = C2D(SUBSTR(RECORD.I,41,4)) / (1024*1024);
    DEV_CAPACTY = C2D(SUBSTR(RECORD.I,45,4)) / (1024*1024);

   ... GENERATE JCL .....

   QUEUE ... JCL STATEMENTS..........;

END;

"EXECIO * DISKW OUTPUT"    /* Write to OUTPUT DD */
```

Figure 49. Sample Extract of a REXX Program to Process DCOLLECT Output

For a full description of the DCOLLECT command and the format of the output, see *Access Method Services for ICF*, SC26-4906.

4.5.10 Establishing PPRC Pairs

The RVA attempts to establish pairings for all requests issued to it in parallel. Automation can be used to stream requests to best fit your requirements. We discuss the performance impact of establishing PPRC pairs in Chapter 6, "RVA PPRC Performance" on page 77.

4.5.11 The Application Programming Interface

You can use the application programming interface (API) macro, ANTRQST, to automate PPRC functions. For a description of this macro and its return codes, see *DFSMSdfp Advanced Services*, SC26-4921.

4.5.12 Error Situations

On a 3990 control unit, pairs may become suspended because of:

- Primary device write failure
- Secondary device write failure
- Secondary not ready, intervention required
- Primary subsystem failure
- Secondary subsystem failure
- Communications failure to the secondary

On an RVA, due to its internal architecture, single device failures are unlikely to occur and thus the RVA is much less likely to have interruptions to PPRC pairings.

If a hardware error occurs on a secondary RVA subsystem that has all of its devices offline, no SIMs will be seen.

We usually recommend that the ECAM device be offline and contain no data. In a PPRC configuration, as the ECAM device cannot be part of a PPRC pair, you may want to consider having the ECAM device online. With this approach SIMs can be issued to a secondary host.

You cannot guarantee that all errors will instigate a "Call home" response from the hardware.

When all paths fail between a primary and secondary subsystem, only the primary will issue error messages and go into suspended status. The secondary subsystem will not be aware that the paths have failed. The volumes on a secondary subsystem will remain in a duplex state. If you are monitoring the secondary device status, be aware that this situation can occur.

To perform a recovery of the secondary subsystem you will have to check the status of the secondary devices. For the whole system to be valid, all devices should be in a duplex state. If they are not, you have to investigate why and determine what recovery actions are required.

4.6 Fallback Planning

When preparing your disaster recovery plan, you must ensure you have plans for fallback. You should prepare plans for returning to your primary site in the event of an invocation of your disaster plan. An invocation of the disaster plan may be the result of a major disaster, an error situation that is less serious in its impact, or a full-scale test of the disaster plan.

4.7 PPRC Dynamic Address Switching

PPRC dynamic address switching (P/DAS) is a software function that enables you to redirect all primary I/O from one PPRC volume to its secondary volume. This function is fully supported on the RVA, provided you have set up the proper pathing. Therefore you can dynamically switch volumes between RVAs.

All considerations for the use of P/DAS on a 3990-6 apply to the RVA. For further information about P/DAS, see *P/DAS and Enhancements to the IBM 3990-6 and RAMAC Array Family*, SG24-4724.

4.8 Licenced Internal Code Upgrades

There are no special considerations for nondisruptive code load on an RVA with PPRC. Normal procedures should be followed, and PPRC pairs do not have to be suspended.

You will be informed if an upgrade is going to be disruptive. When an upgrade does require an IML of the hardware, all information in cache will be lost. You must therefore have your systems quiesced before you issue a suspend. (This does not mean a CSUSPEND QUIESCE.)

Once you have IMLed you have to run CESTPATH and CESTPAIR with the RESYNC parameter.

4.9 Systems Managed Storage Considerations

By using SMS you can group data into specific storage groups. You can then control what volumes are available as PPRC pairs. To avoid problems with attempts to snap to a PPRC primary disk, you can limit snapping to specific target storage groups.

You must have volumes that are PPRC-eligible separated by storage group from volumes that are not being used for PPRC. Any enabled volume in a storage group is eligible for allocation. If you include non-PPRC volumes in a storage group with PPRC volumes, you may end up with only part of an application's data sets or part of a multivolume data set or database at the secondary site.

If you add volumes to an SMS storage group, you must amend your procedures to include checks to see whether the storage group is being remote copied. If space problems occur, you must have procedures not only for adding volumes to the storage group but also for establishing the pairing. If this is not done, the secondary site will be corrupted and useless for recovery.

4.10 Security

PPRC commands are extremely powerful, so it is important that they are used correctly and directed to the correct devices. We recommend that you restrict access to these commands by either protecting the command library or defining resource profiles.

You can use two RACF profiles to control PPRC commands:

- STGADMIN.ANT.PPRC.COMMANDS

  This profile controls access to all PPRC commands, including the CQUERY.

- STGADMIN.ANT.PPRC.CQUERY

  This profile controls access only to the CQUERY command and therefore lets you give wider access to the query function.

4.11 PPRC and Planned System Outages

Because of commercial restraints, it is becoming less and less desirable to have downtime on systems. In a global economy, many companies require high service availability seven days a week.

PPRC can be used to allow continued operation at an alternative site during planned system outages such as major hardware upgrades and maintenance that requires powerdowns.

By utilizing automated procedures, such as PPRC Migration Manager, to generate jobs you can switch sites with only a short disruption to service. The process you could use would be similar to this:

1. Quiesce the system to be moved to guarantee no I/O on PPRC volumes.
2. Delete the PPRC pairs to return them to simplex states.
3. Delete the PPRC paths.
4. Establish PPRC paths from the secondary subsystem back to the primary subsystem.
5. Use the NOCOPY parameter of CESTPAIR to establish PPRC pairing from the secondary volumes back to the primary volumes. This means no copying is done as we know the volumes are identical and the establish processes for each will be very quick. This reverses the direction of the PPRC pairings.
6. Suspend all PPRC volumes. This breaks the pairings, but changes are recorded.
7. Shut down the original system and IPL the system on the secondary devices.

When you are ready to return your system to the original location:

1. Use the RESYNC parameter of CESTPAIR to establish pairs. This copies all changed data back to the original system.
2. Quiesce the system to guarantee no I/O on PPRC volumes.
3. Delete the PPRC pairs to return them to simplex states.
4. Delete the PPRC paths.

5. Establish PPRC paths from the original subsystem.
6. Use the NOCOPY parameter of CESTPAIR to establish PPRC pairing from the original volumes. This parameter means no copying is done as, again, we know the volumes are identical. This process has reversed the direction of the PPRC pairings back to the original configuration.

4.12 Procedures and Documentation

As with all operational procedures, your PPRC procedures should be documented and accessible to all people who may need to support your operation.

The documentation should include:

- A naming standard for volumes and SSIDs
- An up-to-date disk plan identifying PPRC volumes and candidates for PPRC pairing.
- Detailed configurations. See Chapter 3, “Configuration” on page 23
- Requirements for day-to-day support
- What to do in an emergency
- Preprepared jobs
- What the automated procedures do
- What to do if the automated processes fail

4.13 ESCON Manager

ESCON Manager is a program offering that uses a TSO ISPF interface to emulate the functionality of a hardware console attached directly to the ESCON Director. It checks the path status on all connected host systems when you attempt a disruptive operation. As PPRC links between RVAs cannot be seen from the host systems, neither ESCON Manager nor the hardware configuration manager (HCM) can check the connections.

The ESCON Director allows you to document links. We recommend that you use this function to define the PPRC links in the active matrix (see Figure 50 on page 69) and fully document them in your operational procedures.

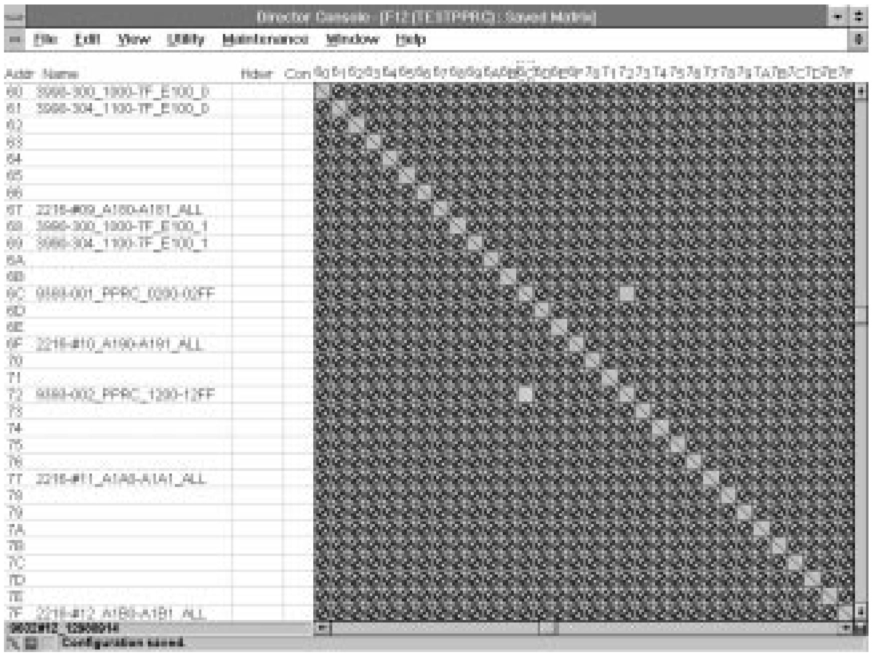

Figure 50. Sample ESCON Director Active Matrix

In the sample active matrix we have an ESCON link between two RVA subsystems using ports 72 and 6C on the ESCON Director. All other ports are prohibited. In the description text we have identified the RVA ID, that it is a PPRC link, and the address range used.

For further information about using ESCON Directors, see the *ESCON Implementation Guide*, SG24-4662.

4.14 Recommendations

In summary, we make the following operational recommendations:

- Keep the configuration as simple as possible.
- Avoid errors by automating as many processes as possible. Consider using an automation offering.
- Use both event-driven and time-driven monitoring.
- Attach a host to the secondary subsystems to report SIMs and allow monitoring.
- Control access to PPRC functions, using a security subsystem such as RACF.
- Have detailed documentation and procedures.
- Document PPRC paths in the active matrix of any ESCON Directors.

Chapter 5. Planning for PPRC Testing

In this chapter we provide an overview of possible test cases for a PPRC environment. These test cases will assist you in suspending a pair and the related MVS console messages, in verifying the operations of your automation, and in obtaining performance data for your environment under certain circumstances.

These procedures are not for removing logic cards or other hardware components from the RVA.

5.1 Producing a PPRC Link Failure

We recommend that you have at least two physical links between primary and secondary RVAs. A link error can be created by pulling one of the links at the ICE card. This causes the RVA to detect a link failure and to report it to a host. As long as other PPRC links are active between the primary LCU and the secondary LCU, the PPRC pairs will stay in duplex mode. In this action plan we have only one PPRC link, so the pair will be suspended.

The following factors can lead to a PPRC link failure:

- Failure of a physical cable on one link
- Failure of a channel adapter on either RVA that causes the logical paths on the link to drop
- ESCON Director failure

5.1.1 Step-by-Step Action Plan

Both device 22F0 and 23F0 are in simplex state and a PPRC link has not been established. You can verify this status with the CQUERY command (see Figure 51 and Figure 52 on page 72).

```
ANTP0030I CQUERY VOLUME FORMATTED 914
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                         (PRIMARY)   (SECONDARY) *
*                                          SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------- ---------  ----------  ----------- ---------    ---------  *
* 22F0  .........  SIMPLEX...   INACTIVE    FF27 30      .... ..   *
*       ...............       ........... 000001321029 ............*
* PATHS SAID/DEST STATUS: DESCRIPTION                             *
* ----- --------- ------  -------------------                     *
*   0   ---- ----   00    NO PATH...........                      *
*       ---- ----   00    NO PATH...........                      *
*       ---- ----   00    NO PATH...........                      *
*       ---- ----   00    NO PATH...........                      *
*******************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
```

Figure 51. CQUERY Volume Command: Device 22F0 in Simplex State

```
ANTP0030I CQUERY VOLUME FORMATTED 916
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                         (PRIMARY)   (SECONDARY) *
*                                          SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 23F0  .........  SIMPLEX...   INACTIVE    23C0 30      .... ..   *
*       ..............      ........... 000001323562 ............*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*******************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 23F0. COMPLETION CODE: 00
```

Figure 52. CQUERY Volume Command: Device 23F0 in Simplex State

We now establish one PPRC path.

Because the logical path was established from the LCU with SSID FF27 to the LCU with SSID 23C0, we did not get a path indication when we issued the CQUERY PATHS command to secondary device 23F0. Only when we issued the CQUERY PATHS command to primary device 22F0, did we get an indication of the established path.

Look at Figure 53 and Figure 54 on page 73 to become familiar with the system response to a CQUERY command.

```
ANTP0050I CQUERY PATHS FORMATTED 918
*************** PPRC REMOTE COPY CQUERY - PATHS ********************
* PRIMARY UNIT: SERIAL#= 000001321029  SSID= FF27                  *
*             FIRST        SECOND        THIRD         FOURTH      *
*           SECONDARY     SECONDARY     SECONDARY     SECONDARY    *
*SERIAL NO: 000001323562  ............  ............  ............ *
*     SSID:     23C0          0000          0000          0000     *
*    PATHS:      1             0             0             0       *
*          SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*          --------- --  --------- --  --------- --  --------- -- *
*       1: 0051 0003 01  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       2: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       3: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*       4: ---- ---- 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 22F0. COMPLETION CODE: 00
```

Figure 53. CQUERY PATHS Command Issued to Primary Device

```
ANTP0050I CQUERY PATHS FORMATTED 920
*************** PPRC REMOTE COPY CQUERY - PATHS ********************
* PRIMARY UNIT: SERIAL#= 000001323562  SSID= 23C0                  *
*              FIRST         SECOND        THIRD         FOURTH    *
*            SECONDARY     SECONDARY     SECONDARY     SECONDARY   *
*SERIAL NO: ............  ............  ............  ............ *
*    SSID:     0000          0000          0000          0000      *
*   PATHS:      0             0             0             0        *
*           SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*           --------- --  --------- --  --------- --  --------- -- *
*        1: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        2: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        3: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        4: ---- ---- 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 23F0. COMPLETION CODE: 00
```

Figure 54. CQUERY PATHS Command Issued to Secondary Device

Now establish a pair. When the copy process from the primary device to the secondary device is completed, the devices will be in full duplex mode. As indicated in Figure 55 and Figure 56 on page 74, the path status is active on the primary device as well as on the secondary device. On the primary you can also see the settings for:

- The CRIT parameter on the CESTPAIR command
- The CGRPLB parameter on the CESTPATH command

```
ANTP0030I CQUERY VOLUME FORMATTED 633
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                           SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------- ---------  ----------  ----------- ---------    ---------  *
* 22F0  PRIMARY..  DUPLEX....   ACTIVE..    FF27 30      23C0 30   *
*       CRIT(NO).......       CGRPLB(NO). 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   1   0051 0003   01    PATH ESTABLISHED...                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
```

Figure 55. CQUERY Command for Primary Device: Path Status Active

```
ANTP0030I CQUERY VOLUME FORMATTED 081
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                         (PRIMARY)   (SECONDARY) *
*                                          SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 23F0  SECONDARY  DUPLEX....   ACTIVE..    FF27 30      23C0 30   *
*       ...............       ........... 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 23F0. COMPLETION CODE: 00
```

Figure 56. CQUERY Command for Secondary Device: Path Status Active

Remove the ESCON link now. Figure 57 shows that the status of the PPRC path is inactive and that the primary device is suspended.

```
ANTP0030I CQUERY VOLUME FORMATTED 686
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                         (PRIMARY)   (SECONDARY) *
*                                          SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 22F0  PRIMARY..  SUSPEND(6)   INACTIVE    FF27 30      23C0 30   *
*       CRIT(NO).......       CGRPLB(NO). 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   1   0051 0003   09    (UNDETERMINED).....                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
*       ---- ----   00    NO PATH............                      *
* IF STATE = PENDING/SUSPEND:  FIRST CYL OUT OF SYNC =  00001      *
*                               LAST CYL OUT OF SYNC =  00435      *
*                              PERCENT OF COPY COMPLETE = 099%     *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
```

Figure 57. CQUERY Volume for Primary Device: Path Status Inactive

Figure 58 on page 75 shows that the state of the secondary device is still duplex. It is still duplex because the secondary LCU has no way of distinguishing between a lost link or a situation when there is no I/O from the primary LCU. In a disaster scenario when links are lost between primary and secondary control units, the status of all secondary devices should be duplex until a CRECOVER action is performed to ready the devices for use.

```
ANTP0030I CQUERY VOLUME FORMATTED 700
************** PPRC REMOTE COPY CQUERY - VOLUME *******************
*                                          (PRIMARY)   (SECONDARY) *
*                                           SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------- ---------  ----------  ----------- ---------    ---------  *
* 23F0  SECONDARY  DUPLEX....   ACTIVE..    FF27 30      23C0 30   *
*       ...............       ........... 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 23F0. COMPLETION CODE: 00
```

Figure 58. CQUERY Volume Command for Secondary Device: Duplex State

Figure 59 and Figure 60 on page 76 show the output of the CQUERY PATHS commands to both the primary and secondary LCUs.

```
ANTP0050I CQUERY PATHS FORMATTED 667
*************** PPRC REMOTE COPY CQUERY - PATHS *******************
* PRIMARY UNIT: SERIAL#= 000001321029  SSID= FF27                  *
*             FIRST       SECOND        THIRD         FOURTH       *
*           SECONDARY     SECONDARY     SECONDARY     SECONDARY    *
*SERIAL NO: 000001323562  ............  ............  ............ *
*     SSID:     23C0          0000          0000          0000     *
*    PATHS:      1             0             0             0       *
*           SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*           --------- --  --------- --  --------- --  --------- -- *
*        1: 0051 0003 09  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        2: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        3: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        4: ---- ---- 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 22F0. COMPLETION CODE: 00
```

Figure 59. CQUERY PATHS Command: Primary Device

The path status of 09 indicates that the establish failed because of the loss of the PPRC path, and that the establish will be retried again when conditions change.

```
ANTP0050I CQUERY PATHS FORMATTED 649
*************** PPRC REMOTE COPY CQUERY - PATHS ********************
* PRIMARY UNIT: SERIAL#= 000001323562  SSID= 23C0                  *
*             FIRST         SECOND        THIRD         FOURTH     *
*           SECONDARY     SECONDARY     SECONDARY     SECONDARY    *
*SERIAL NO: ............  ............  ............  ............ *
*     SSID:     0000          0000          0000          0000     *
*    PATHS:      1             0             0             0       *
*           SAID DEST S*  SAID DEST S*  SAID DEST S*  SAID DEST S* *
*           --------- --  --------- --  --------- --  --------- -- *
*        1: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        2: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        3: ---- ---- 00  ---- ---- 00  ---- ---- 00  ---- ---- 00 *
*        4: ---- ---- 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 23F0. COMPLETION CODE: 00
```

Figure 60. CQUERY PATHS Command: Secondary Device

Chapter 6. RVA PPRC Performance

In this chapter we discuss the performance characteristics of PPRC on an RVA, present some recommended configurations, and provide some advice on sizing a PPRC workload.

6.1 Performance Characteristics

PPRC for the RVA is designed to take advantage of the unique RVA architecture to provide the best possible performance. The following factors affect the performance of PPRC on the RVA when compared with performance of PPRC on 3990-6/9390 subsystems:

- The RVA always transfers full track and compressed data, whereas the 3990-6/9390 subsystems transfer data on a record basis. This reduces the amount of handshaking between the two subsystems.
- In the case of large sequential writes, the RVA can overlap the transfers to the secondary with transfers being received from the host. In contrast, the 3990-6/9390 wait for the transfer to the secondary to complete before receiving transfers from the host.
- Chained sequential transfers to the same track are consolidated into a single transfer to the RVA secondary, whereas the 3990-6/9390 sends each record separately.
- The ESCON frame size for the RVA is larger at 1024 bytes as compared to 256 bytes for the 3990-6/9390. The size of the buffer for the RVA is also larger than for the 3990-6/9390.
- Only the data known to the host is stored on the RVA (specifically, unallocated space, as well as allocated and unused space, uses no disk space). During establish of the PPRC pairs, the RVA takes advantage of this by sending only compressed and stored data to the secondary.
- The RVA has eight internal paths that can be active concurrently, whereas the 3990-6/9390 subsystems have four internal paths, with the exception of the 9390-2, which has eight internal paths.
- The RVA does not internally limit the number of establish pairs, whereas the 3990-6/9390 subsystems internally limit the number of concurrent establish pairs to two.

6.2 Performance Variables

Several variables may affect the performance you can achieve from the RVA in a PPRC environment. Some are in the area of configuration, such as ESCON distance, number of ESCON links, and cache size. Others deal with workload, such as data block size and transfer size. In addition, performance achieved is a function of the amount of write activity to the primary subsystem.

6.2.1 Configuring for Performance

Ideally the secondary subsystem in an RVA PPRC configuration should have capacity similar to that of the primary subsystem. The primary and secondary subsystems should have an equivalent number of links, disk capacity, and cache sizes. If SnapShot is also used, the secondary may require additional disk capacity.

The cache capacity of the secondary should be adequate capacity to handle the primary activity once the secondary is invoked for recovery purposes. Our recommendation for cache sizes of the primary and the secondary RVAs is 4 GB.

6.2.2 Distance between the RVA Primary and Secondary Subsystems

In planning for 3990-6/9390 PPRC implementations, you have to manage the distance between the links. The RVA, however, is less sensitive to distance because of the differences in ESCON frames size and buffer sizes. Although ESCON distance is a consideration for RVA PPRC performance, its impact is less than for 3990-6/9390 PPRC implementations. We recommend that you consider using four links, two per cluster, from the primary RVA to the secondary RVA. You may want to reduce the practical number of links you use after reviewing the discussion in 3.2.9, “Channel Register Sets” on page 38.

6.3 Performance during Establish

The RVA transfers only compressed, stored tracks of data from the primary to the secondary. In addition, the RVA sends data with a minimum of overhead. These factors make the establish time relatively fast.

During establish, the RVA transmits up to a cylinder of data across to the secondary without allowing interruption for host I/O. Also, the PACE parameter of the CESTPAIR is ignored. The volume is locked during the cylinder transfer. After the cylinder transfer is complete, an I/O is allowed to the volume. You should be aware of the potential queuing on the volume especially for volumes with high activity.

Distance has only a minor impact on the time to establish because of the frame sizes and buffer sizes.

6.3.1 Link Recommendation to Optimize Establish Times

Because links are not shared with host data transfers, you may be limited to the number of links available for PPRC. Where possible and to maximize performance during establish, we recommend that you use four links from the primary to the secondary, preferably two per cluster.

Remember that if you have only eight ESCON channel adapters and use some of them for PPRC, you will reduce the number of concurrent host I/Os that can be processed. This may limit the throughput of the RVA, and you may consider upgrading to a 16 channel configuration to provide the maximum of 8 host attached channels in addition to your PPRC links.

6.3.2 Number of Concurrent Establish Pairs

Because the RVA starts all establish pairs that it is given, you will want to limit the number of jobs that establish PPRC pairs through automation. Here we explain the use of the MSGREQ parameter with the CESTPAIR command and the corresponding syslog messages.

Figure 61 shows an example of the CESTPAIR command with the MSGREQ(NO) parameter.

```
//SYSTSIN  DD *
 CESTPAIR DEVN(X'22F0') PRIM(X'FF27' 1321029 X'30') -
          SEC(X'23C0' 1323562 X'30') -
          MSGREQ(NO)
```

Figure 61. CESTPAIR Command with MSGREQ(NO)

Figure 62 shows the sequence of messages received when you use MSGREQ(NO) on the CESTPAIR command. You can see that the ANTP0001I message is issued before the PPRC pair is in full duplex mode.

```
$HASP100 RBPPRC2A ON INTRDR          FROM TSU07743     RBPPRC2
IRR010I  USERID RBPPRC2  IS ASSIGNED TO THIS JOB.
$HASP373 RBPPRC2A STARTED - INIT   AL - CLASS A - SYS SYS8
ANTP0001I CESTPAIR COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
IEA494I 22F0,SL22F0,PPRC PAIR PENDING,SSID=FF27,CCA=F0
JOB=RBPPRC2A STEP=GO       PGM=IKJEFT01 CC=0000
$HASP395 RBPPRC2A ENDED
IEA494I 22F0,SL22F0,PPRC PAIR FULL DUPLEX,SSID=FF27,CCA=F0
```

Figure 62. Console Messages Issued for MSGREQ(NO)

Figure 63 shows a CESTPAIR command with the MSGREQ(YES) parameter.

```
//SYSTSIN  DD *
 CESTPAIR DEVN(X'22F0') PRIM(X'FF27' 1321029 X'30') -
          SEC(X'23C0' 1323562 X'30') -
          MSGREQ(YES)
```

Figure 63. CESTPAIR Command Using MSGREQ(YES)

In Figure 64 on page 80 you can see that the ANTP0001I message is issued after the PPRC pair is in full duplex mode.

```
$HASP100 RBPPRC2A ON INTRDR          FROM TSU07743     RBPPRC2
IRR010I  USERID RBPPRC2  IS ASSIGNED TO THIS JOB.
$HASP373 RBPPRC2A STARTED - INIT   AD - CLASS A - SYS SYS8
IEA494I 22F0,SL22F0,PPRC PAIR PENDING,SSID=FF27,CCA=F0
IEA494I 22F0,SL22F0,PPRC PAIR FULL DUPLEX,SSID=FF27,CCA=F0
ANTP0001I CESTPAIR COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
JOB=RBPPRC2A STEP=GO       PGM=IKJEFT01 CC=0000
$HASP395 RBPPRC2A ENDED
```

Figure 64. Console Messages Issued for MSGREQ(YES)

Because the PACE parameter of the CESTPAIR command is ignored for the RVA, you can use the MSGREQ(YES) parameter to control the number of concurrent establish pairs for each RVA. MSGREQ(YES) specifies that PPRC wait until the initial full volume copy operation is complete before issuing completion message ANTP0001I. If you are using automation, use this message as a trigger before submitting the subsequent establish pair jobs.

Select periods of low utilization when possible. To minimize the impact on host I/O, establish one PPRC pair at a time. The maximum number of establish pairs executing concurrently should be limited to four. Beyond four at a time, you experience the law of diminishing returns.

6.4 Workload Performance

In this section we review several types of workloads and explain how they may be affected when you add PPRC to the environment.

6.4.1 Database and Transaction Processing Workloads

There are many ways of characterizing workloads. Here is one way of categorizing them:

- Cache friendly - characterized by unusually high hit ratios. Adding PPRC to a cache friendly workload yields an almost negligible increase in response times and has little impact on throughput because of the low write activity and good caching characteristics.
- Cache hostile - characterized by unusually low hit ratios. The addition of PPRC to a cache hostile workload could affect response times by up to several milliseconds, but the impact on throughput will be moderate.
- Typical database - 4 KB block size, read/write ratio of 3:1, read hit ratio from 75% to 80%, and high levels of device skew. For a typical database workload, you should expect response time to be affected by the addition of PPRC by 1 to 2 ms from IBM's benchmarks at low loads, and a larger increase under stress conditions. The impact on throughput should be almost negligible.

The affect of PPRC on performance is the same both at zero distance and for distances up to 43 km. This is an improvement over 3990-6 or 9390 PPRC configurations, which are more distance sensitive.

For more details on RVA PPRC performance refer to the white paper entitled “RAMAC Virtual Array PPRC Performance,” which is available on MKTTOOLS.

6.4.2 Sequential Workloads

PPRC reduces the write sequential throughput for a large number of streams by up to 20% to 30%, depending on the distance, transfer size, and block size of the data. This is an improvement over the 3990-6/9390 PPRC implementations where the write sequential throughput was reduced more substantially. For single write streams with characteristics similar to the IMS write-ahead data set (WADS), CICS journals, and DB2 logs, some queuing is likely to occur.

To optimize performance in an environment with a lot of PPRC activity (for example, an environment where all 256 devices are PPRC primary devices), we recommend that you use four links, two per cluster, from the primary RVA to the secondary RVA. We also recommend that you use 4 GB cache sizes.

6.5 Predicting RVA PPRC Performance

When considering using PPRC on the RVA, make an assessment of the performance of the RVA and understand any current bottlenecks. Consider the following from a performance point of view:

- Volumes with low read to write ratios (high write content) may not have good performance. Consider spreading the data across additional volumes by moving data sets. Or, if the activity is to a single data set, consider making it a multivolume data set.
- A subsystem with high channel utilization may already be saturated or near the saturation point, and adding PPRC to it is very likely to impact performance. In a situation where you use eight ESCON ports, adding PPRC links without adding ports reduces subsystem throughput.
- An application with a critical, single write sequential stream to a volume (similar to a log data set) will be affected if this stream is made part of a PPRC pair. A careful trade-off of performance over recoverability must be made for this type of application.
- During batch periods that involve large amounts of write sequential work, batch throughput may be reduced by up to 20% to 30%. Consider spreading batch write sequential work across multiple subsystems.

DASD Magic for OS/2 is a modeling tool which can be used to predict the performance of PPRC on the RVA. Available to IBM marketing personnel on MKTTOOLS, it does high-level modeling of the effects of changes to disk subsystem configurations, such as new disk or different cache sizes. DASD Magic for OS/2 is intended to model random transaction workloads and does not model sequential (batch) workloads.

DASD Magic for OS/2 uses performance information from an existing configuration together with information about subsystem characteristics, to predict likely performance of the same workload using a different storage subsystem. Input is Report Management Facility (RMF) data and cache statistics (usually also obtained from RMF, but optionally from Cache Analysis Aid (CAA) or Cache RMF Reporter (CRR)). It provides a quick estimate of expected benefits and enables judgment of whether more detailed analysis is required.

Note

All analytic queuing models including DASD Magic for OS/2 cannot accurately model sequential workloads. Do not use DASD Magic for OS/2 to model a heavy (over 15%) sequential workload. The resulting projections will be overly optimistic and invalid.

Chapter 7. Using PPRC on the RVA with VM/ESA and VSE/ESA

You can implement RVA PPRC on your VM/ESA or VSE/ESA systems. PPRC control and management in these environments is done only through ICKDSF. However, ICKDSF contains only a subset of the commands and functions available in TSO on MVS. You must establish measures to control recoverability in the event of a disaster, for example, procedures to respond to failure scenarios that might otherwise be addressed by the TSO commands.

7.1 Considerations for VM/ESA

In this section we cover the considerations for using PPRC on the RVA in a VM environment.

PPRC is supported on VM/ESA 2.1.0 and above. When operating as a guest under VM/ESA 2.1.0 and above:

- All PPRC volumes must be either dedicated volumes or defined as fullpack minidisks (including DEVNO-defined minidisks).
- The VM guest directory must include an entry stating "STDEVOPT DASDSYS DATAMOVER":
  - The STDEVOPT statement specifies the optional storage device management functions available to a virtual machine.
  - The DASDSYS operand tells the control program (CP) whether the virtual machine is authorized to control and process concurrent copy and PPRC establish pair channel control words (CCWs).
  - The DATAMOVER option tells the CP whether the virtual machine is authorized to issue concurrent copy and PPRC establish pair CCWs.
- You have to order the following fixes for the PPRC error recovery procedures:

  VM APARs: VM59895, VM59926, VM59965, VM60121, VM60155

  OW18825 with MVS/ESA guest
- You must review APAR II08303 for the latest maintenance updates.
- You need ICKDSF Release 16 with APAR PN66541.

7.1.1 Configuring for the VM/ESA Environment

The configuration considerations for using PPRC on the RVA in a VM environment are the same as for MVS. Refer to the Chapter 3, “Configuration” on page 23.

In managing your ESCON configurations, remember that the ESCON Manager is not available for VM. You must make sure that your documentation shows the latest ESCON connections, and any changes must be reflected immediately. If you plan to use the ESCON Manager on a guest MVS system, read the special considerations in *Planning for ESCON Manager*, GC23-0423.

The real addresses are used as values for addresses in the ICKDSF commands. We recommend that you keep the VM virtual device number the same as the real device number for the devices for which you want to establish PPRC pairs, to facilitate entering the commands correctly. If you are unable to keep the VM virtual device number the same as the real device number, be sure to accurately document the mapping.

7.1.2 Operations for the VM/ESA Environment

ICKDSF is used to execute PPRC commands in VM/ESA. The PPRCOPY command provides the means to initiate and terminate the functions, suspend a copy operation, query the status of volumes, and enable a recovery host to ability to recover control of the secondary volume.

For the commands that require the RVA serial number, you have to provide the seven-digit serial number of the RVA. Before issuing the PPRCOPY ESTPATH command, it may be useful to issue the CONTROL CONFIGURE(DISPLAY) UNIT(CCUU) command to determine the subsystem ID and serial number, and the unit's CCA. Obtain the SAID values and DESTination address from the person who has physical cabling and connection information for the affected control unit and the ESCON Director, if any, at the primary and secondary locations.

Do not use the SERIAL LINK ADDRESS from the CONTROL CONFIGURE(DISPLAY) for the SAID or DEST parameters, as the values from this command refer to the connection between the processor and the RVA, not the connection between the RVAs.

For the PPRCOPY ESTPAIR command, the PACE parameter is ignored. In addition, the RVA will start to establish as many PPRC pairs as the PPRCOPY ESTPAIR commands you give it. Therefore, you want to limit the number of PPRCOPY ESTPAIR commands. We continue to suggest that you limit the number of concurrent establishes on a single RVA to four.

See 4.2.1, “CESTPATH: Establishing Paths” on page 52 for the use of the CRIT parameter. As a reminder, ICKDSF does not support consistency group timer functions. You should establish procedures or automation that addresses failures that will cause your pairs to be suspended. When you use PPRC for a guest, VM forwards relevant unit check messages to the guest. The guest must handle all operations to preserve consistency.

Use the PPRCOPY QUERY command with or without the PATHS parameter to get information about the status of the PPRC pairs. The output from the PPRCOPY QUERY command is returned to the ICKDSF user.

For full descriptions of the PPRCOPY commands under ICKDSF, refer to *ICKDSF R16 Refresh User's Guide*, GC35-0033.

7.2 Considerations for VSE/ESA

In this section we cover the considerations for using PPRC on the RVA in a VSE/ESA environment.

Consider the following requirements in planning PPRC for VSE/ESA:

- PPRC is supported on VSE/ESA 2.1.0.
- You have to order VSE/ESA APAR DY44407 for the PPRC error recovery procedures.

- You must review APAR II08303 for the latest maintenance updates.
- You need ICKDSF Release 16 with APAR PN66541.

> **Note**
>
> Remember that CGROUP is not supported in ICKDSF.

7.2.1 Configuring for a VSE/ESA Environment

The configuration considerations for using PPRC on the RVA in the VSE/ESA environment are the same as for MVS. Refer to the Chapter 3, “Configuration” on page 23.

In managing your ESCON configurations, remember that the ESCON Manager is not available for VSE/ESA. You must make sure that your documentation reflects the latest ESCON connections, and any changes must be updated immediately.

7.2.2 Operations for the VSE/ESA Environment

As in VM/ESA, ICKDSF is used to execute PPRC commands in VSE/ESA. Remember that CGROUP is not supported in ICKDSF. Establish measures or automation that address failures that will cause your pairs to be suspended.

Chapter 8. Using PPRC in a Database Environment

In this chapter we make various recommendations for maintaining the write sequence integrity of your database systems and building measures to protect your database systems against a rolling disaster. Most of our recommendations aim to give you a restartable copy of your data because a database restart is easier to do than a recover. Restart requires that the state of all the data at your secondary site is consistent. This is achieved in a PPRC environment by freezing all secondary volumes at the first occurrence of an error. However, you must establish whether it is more important to have a current or consistent copy of your data among your subsystems. Freezing your secondary volumes ensures that they are consistent, but any updates that were made to the primary volumes after the freeze are not reflected at the secondary site. Therefore in the event of a disaster some data may be lost.

If you do not freeze your secondary volumes on an error, you can continue to propagate updates from the primary to the secondary volumes. However you run the risk that your data at the secondary site will not be consistent and that you will have to restore previous backups and forward recover to your latest saved logs.

Using automation techniques, or the GDPS service offering, you can maintain data consistency, with minimum disruption to your production applications, and ensure a fast restart in the event of a disaster.

Refer to the following documentation to learn more about planning for database recovery:

- *Planning for IBM Remote Copy*, SG24-2595
- *Remote Copy Administrator's Guide and Reference*, SC35-0169
- *Disaster Recovery Library: S/390 Technology*, SG24-4210-01
- “Remote Copy Technical Update (RCTECH),” available on MKTOOLS from your IBM Representative.

8.1 Using PPRC with IMS

IMS provides DASD volume and data set duplexing facilities to protect against DASD failure. Because of the RVA's architecture, failure of a single device is very unlikely.

Some customers use the IMS dual-write capability for IMS log and RECON data sets to place a second copy at another site and then forward recover individual data sets. However, in the event of a disaster the recovery process may be very lengthy due to the large numbers of data sets that must be recovered.

There are many methods of maintaining a time-consistent copy of your data at a secondary site:

- XRC uses timestamping and consistency group management to maintain time consistency across multiple storage controllers using the SDM.
- IMS Remote Site Recovery (RSR) is a separately priced feature of IMS. It enables you to resume online operations from a secondary site with minimal delay and minimum loss of data. RSR is IMS specific and covers only IMS

data; however, it makes full use of IMS-specific functions and features. Therefore, if your application is IMS only, we recommend that you use IMS RSR to manage your database remote copy.

- PPRC with CGROUP allows you to automate the detection of errors and freeze all secondary volumes at a consistent point in time. GDPS provides this automation as part of a service offering.

Each of these techniques allows you to maintain data consistency at your secondary site and, in the event of a disaster, to restart your applications in the minimum time. However, each solution uses different techniques, so cannot be combined in a single application system as recovery cannot be coordinated between the different solutions.

You should also aim to copy all data to your secondary site, using your preferred technique. If you selectively copy some data sets or volumes, you may find during recovery that you are missing some crucial data. For this reason you should also plan to make available tape data at your recovery site.

8.2 IMS Recovery Process

If all pairs are full duplex, IMS restart can be done on the secondary site. IMS will manage the entire operation. Use the GDPS service offering, or automation of the CGROUP FREEZE command, to ensure that all volumes are consistent. This is a simple approach to ensure that in the event of a disaster you can restart your applications in the shortest time with minimum data loss.

If volumes are individually suspended, the DASD ERP issues message IEA491E. If CGROUP was used, this message contains the time stamp of the failure and the suspension. Data sets can then be forward recovered by using data from DBRC of the IMS Database Manager. However, if volumes are suspended or pending, this will have a considerable impact on the duration of forward recovery.
Time-stamp-based recovery is a complicated procedure in IMS, you can use the IMS Recovery Saver utility program to establish a time-stamp recovery position.

8.3 Using PPRC with DB2

It is essential that you understand the impact that PPRC will have on your primary site's operation and availability and the secondary site's recovery integrity. Significant changes to your automation and recovery processes may be required when you implement PPRC in a DB2 environment.

8.3.1 CRIT(Y)

When you use the CRIT(Y) option, the primary volume returns a permanent error to all I/O writes if there are no paths to the secondary PPRC volume. For most applications a permanent write error causes the application to terminate, but DB2's internal recovery allows processing to continue, which could lead to data integrity issues.

If a write to a DB2 log volume receives a permanent write error, DB2 can redirect the I/O to another log data set. If a write to a table space volume receives such an error, DB2 logs the write error page range and marks it for recovery.

To maintain DB2 integrity at the secondary system, all volumes that are or could be used by DB2 must be in PPRC pairings. In an SMS environment you can avoid this problem by ensuring that all volumes in a particular SMS storage group are in PPRC pairings. All data will therefore continue to be reflected at the secondary site.

If a volume has been logged as needing recovery, some tables and partitions may be affected. If the error is due to all paths failing between a controller pair, a growing number of pages and tablespaces could be flagged as subsequent writes encounter the PPRC permanent errors. As a result the primary site will require many DB2 recoveries, and the restart at your secondary site will require the same recoveries for all affected tablespaces. The DASD ERP issues an IEA491E message for each pair that suspends to the system initiating the I/O, and an IEA494I message is broadcast to all attached systems where the volume is online.

Figure 65 on page 90 shows the impact of using CRIT(Y):

1. DB2 in normal operation with PPRC pairs active
2. Write failure to secondary. Permanent I/O error issued by ERP.
3. DB2 marks pages for recovery and redirects I/O. PPRC pairing suspended.

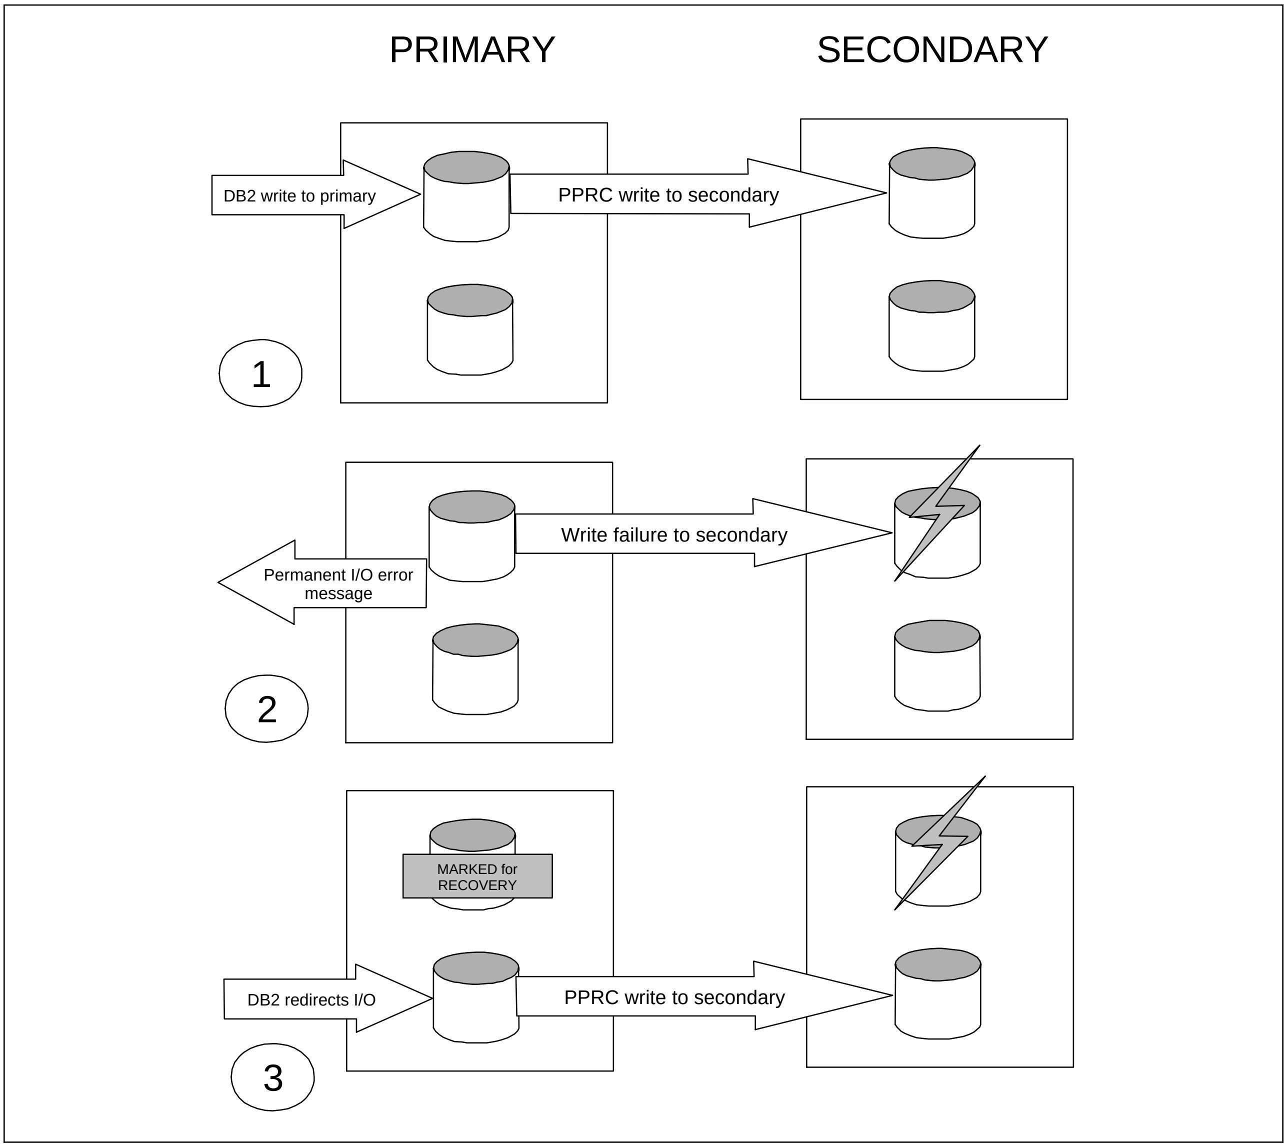

Figure 65. Sample Error Scenario Using CRIT(Y)

8.3.2 CRIT(N)

When you use the CRIT(N) option, the primary volume will return an error to a write if the data written cannot be copied to the secondary PPRC volume.

This error is not passed to DB2. The DASD ERP receives the write error from the storage controller. The sense information received identifies the error as a write failure to the secondary device, and a permanent error is not issued on the primary. The DASD ERP issues an IEA491E message for each pair that suspends to the system initiating the I/O, and an IEA494I message is broadcast to all attached systems where the volume is online.

The storage controller suspends the PPRC pairing. All subsequent writes to the primary should succeed as the copy to the secondary is not attempted. ERP retries the failed write operation. The application therefore continues with minimal impact.

The secondary volume no longer reflects any updates and is subsequently out of synchronization with the primary and other secondary volumes.

Figure 66 shows the impact of using CRIT(N):

1. DB2 in normal operation with PPRC pairs active
2. Write failure to secondary. Message issued by ERP.
3. DB2 I/O retried by ERP. Pair suspended.

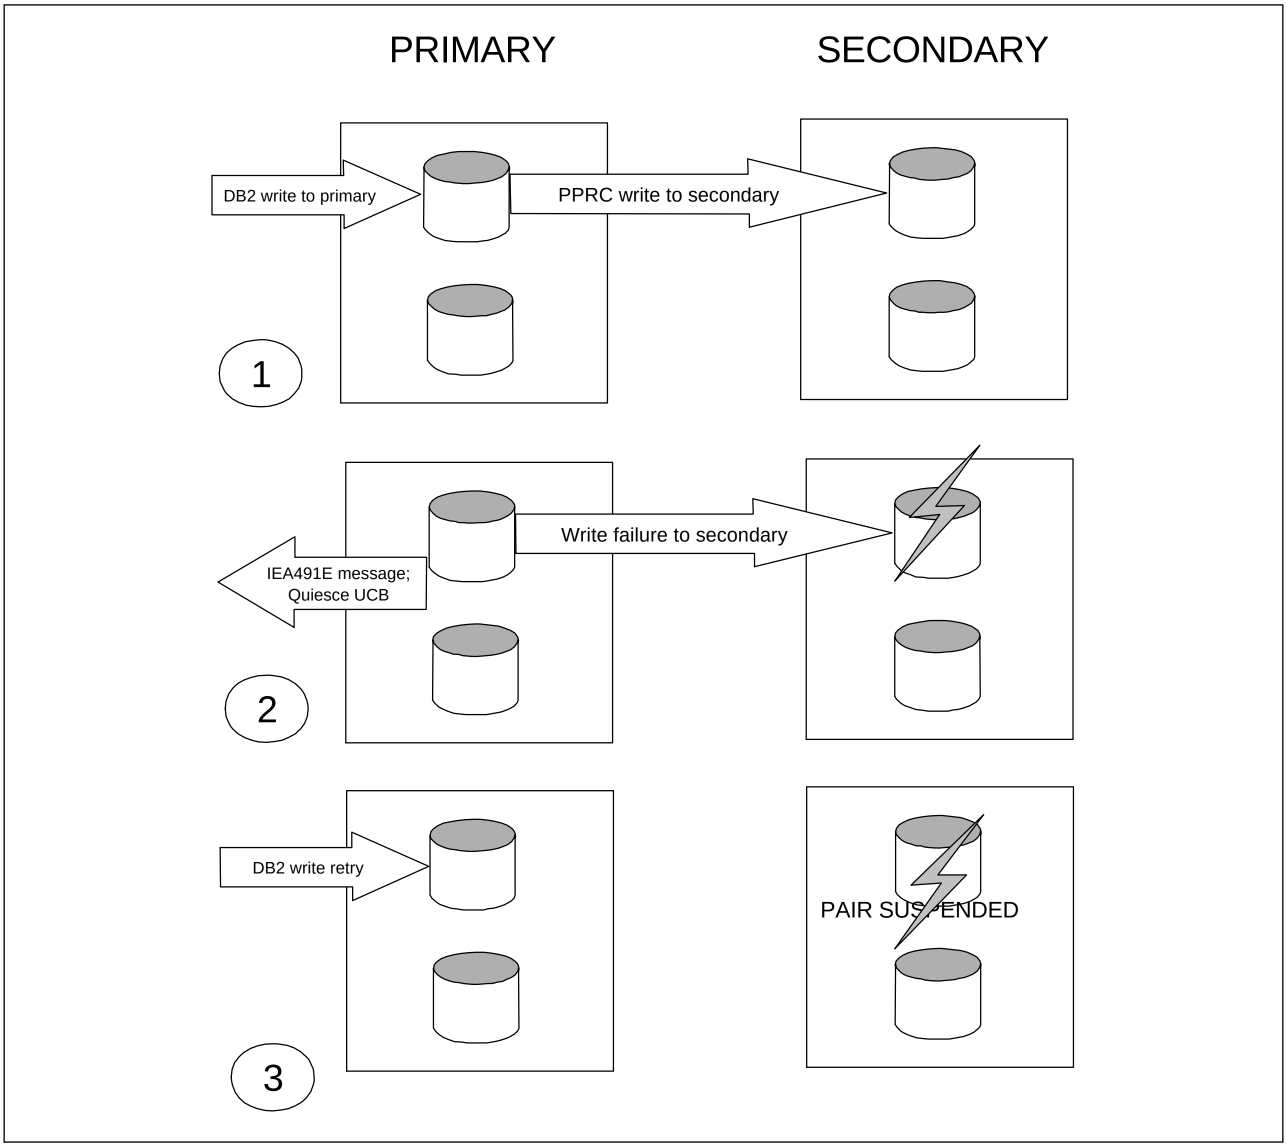

Figure 66. Sample Error Scenario Using CRIT(N)

8.3.3 PPRC Consistency Grouping Timer

The PPRC consistency grouping timer function gives you time to issue automated commands when a pairing has an error by putting the primary volume into a long busy state. This state is issued by the ERP and indicated on the IEA494I message, which indicates that the pairing is changing status and a long busy state has been instigated. This long busy state temporarily queues I/O to the PPRC pairs on the affected primary/secondary LCU pair.

On the 3990-6 and 9390 this is controlled through a VPD setting. On the RVA this function is provided through the CGROUP parameter of the CESTPATH command.

Table 5 on page 53 shows the interaction of the CRIT and CGROUP parameters.

8.3.4 Automation and DB2

Automated operations must identify when there is a PPRC failure and ensure that the secondary site knows that the secondary volume of a PPRC pairing has been suspended and will not be usable if the system is used for recovery. Automation then must perform damage limitation by deciding whether a major failure is in progress or a transient or isolated failure has occurred. You must avoid having the rolling disaster mirrored at your secondary subsystems to avoid many DB2 recoveries from image copies and logs.

If you have used the consistency grouping timer, your automation can use the CGROUP commands to temporarily suspend I/O to all other PPRC pairings. This creates a point of consistency at the secondary site.

8.4 Recommendations

When you use PPRC with DB2, the tradeoff is between impacting the application's continuous availability and maintaining a viable recovery system.

It is critical for DB2 restart to keep all volumes consistent. You can achieve integrity by having all volumes defined with CRIT(N) and the consistency grouping timer on. Use automation to issue CGROUP FREEZE commands to all relevant control units, logical or physical. Be aware of the impact this will have on any DB2 subsystems, however. The process will stop all I/O for up to 2 min while it suspends all PPRC pairings, but it does not cause any permanent write errors. With CRIT(N) you avoid placing pages and tablespaces on the primary site into recovery due to a unit check. The long busy state enables your automation time to freeze your DB2 subsystem to a point in time on the secondary subsystem and record changes until you want to resynchronize the pairings.

Figure 67 on page 93 shows the impact of using CRIT(N) and CGROUP(Y).

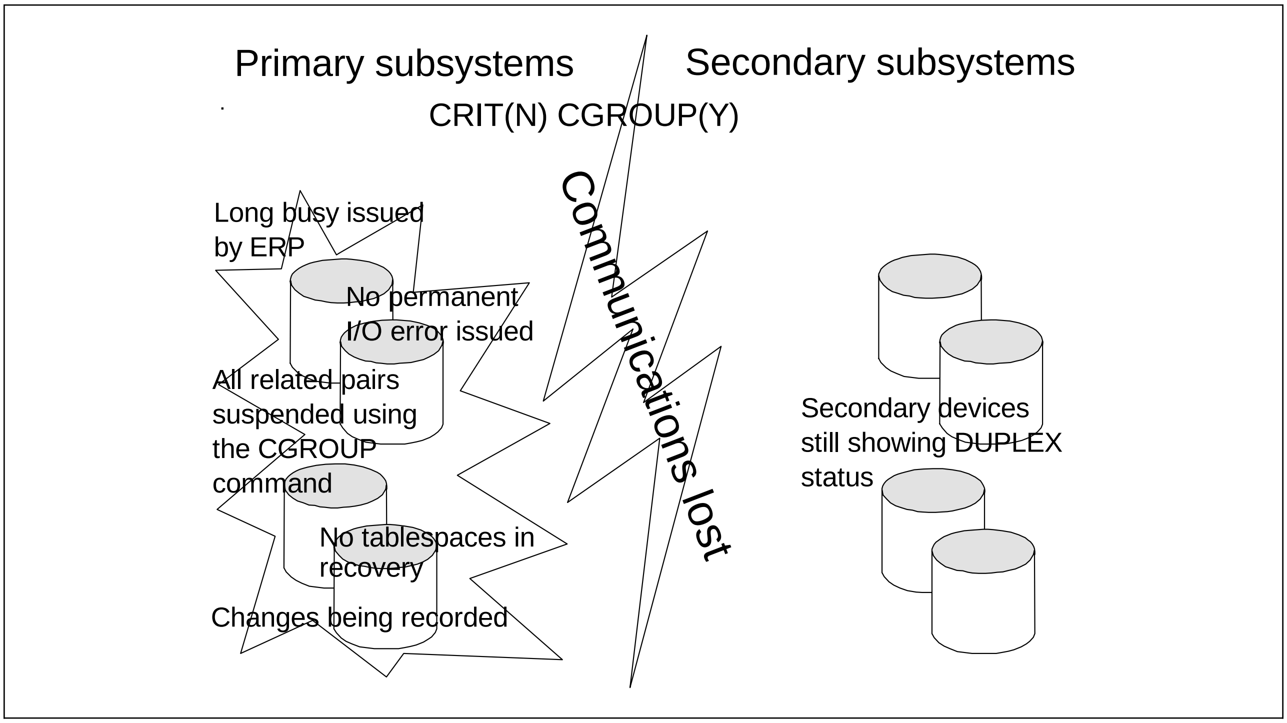

Figure 67. Using CRIT(N) and CGROUP(Y)

Your automation must be able to deal with all error situations and maintain data consistency. The automation solution in GDPS has been designed to meet this requirement.

Chapter 9. Using PPRC with SnapShot

On the RVA you can combine SnapShot with PPRC functions. There are considerations regarding the interaction of SnapShot with volumes that are part of a PPRC pair.

If you want to learn more about general SnapShot operations, see *Implementing SnapShot*, SG24-2241.

SnapShot requires that logical space be available on the same RVA subsystem where the source resides. If SnapShot is going to be used on an RVA subsystem that also has PPRC pairs, it is important to have some logical devices on that subsystem that are not part of a PPRC pair. You cannot, therefore, have all volumes of a subsystem in PPRC pairs and still use SnapShot.

You cannot SnapShot data onto any volume that is part of a PPRC pair. You can use SnapShot to copy data from a PPRC primary volume. See Table 7 for a summary of SnapShot possibilities with PPRC.

Table 7. SnapShot and PPRC Combinations

	RVA Simplex Volume	**PPRC Primary Volume**	**PPRC Secondary Volume**
SnapShot “Source” Volume	Yes	Yes	No (Note 1)
SnapShot “Target” Volume	Yes	No	No

Note: Use the procedure defined in 9.2, “Secondary Devices of a PPRC Pair” on page 96 to Snap from a PPRC secondary volume.

9.1 Primary Devices of a PPRC Pair

The primary device in a PPRC pair can be the source for a SnapShot. You can SnapShot any data from a primary volume onto a volume that is not part of a PPRC pair without modifying the PPRC state of the primary volume.

A PPRC primary device cannot be the target for SnapShot.

When dynamically allocating space for a SnapShot target, SnapShot uses all volumes available to it on a space available basis. It does not check PPRC status. Do not put PPRC volumes into a SnapShot target storage group or provide a volume list that includes a PPRC volume. SnapShot will fail if any part of a data set allocates a PPRC volume.

A PTF to SnapShot excludes PPRC volumes from selection when you use volume preferencing in an SMS environment. If you are not SMS managed, you must make sure that you exclude PPRC volumes from selection.

If you allocate a target data set onto a PPRC primary volume, the SNAP will fail, but, provided you specified the DATAMOVER parameter, SnapShot will invoke DFSMSdss to make the data set copy. If you do not specify a DATAMOVER the

Snap fails with message SIB4621E. We recommend that you always specify a DATAMOVER on your SnapShot commands.

A primary PPRC volume cannot be the target for a SnapShot command. If you need to make volume snaps of a primary device, the device must be returned to simplex state.

9.2 Secondary Devices of a PPRC Pair

You cannot do any read or write I/O to a secondary volume of a PPRC pair. You can, however, SnapShot a secondary volume by suspending the PPRC pair and returning the volume to simplex. Because the primary subsystem records changed tracks, the pairing can be quickly reestablished. A procedure similar to the following could be used:

1. CSUSPEND the pair.

   The primary and secondary devices are both suspended. Figure 68 shows the CQUERY output from device 22F0, which is a PPRC primary device.

```
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                      (PRIMARY)   (SECONDARY) *
*                                       SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 22F0  PRIMARY..  SUSPEND(4)   ACTIVE..    FF27 30      23C0 30   *
*       CRIT(NO).......       CGRPLB(NO). 000001321029 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   1   0051 0003   01    PATH ESTABLISHED...                      *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*                              PERCENT OF COPY COMPLETE = 100%     *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 22F0. COMPLETION CODE: 00
```

Figure 68. CQUERY Output from Suspended Primary Device

   Figure 69 on page 97 shows the output from a CQUERY command directed to 23F0, which is the secondary device.

```
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                           SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 23F0  SECONDARY  SUSPEND(5)   ACTIVE..    FF27 30      23C0 30   *
*       ...............       ........... ............ 000001323562*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
* SECONDARY WAS SUSPENDED (YMD/GMT): 1998-10-20 22.51.48.075595    *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 23F0. COMPLETION CODE: 00
```

Figure 69. CQUERY Output from Suspended Secondary Device

2. CRECOVER the secondary, returning it to simplex.

   The primary device stays in SUSPENDED state. The secondary device returns to SIMPLEX. Figure 70 shows QUERY output from device 23F0 which is the secondary device.

```
************** PPRC REMOTE COPY CQUERY - VOLUME ********************
*                                          (PRIMARY)   (SECONDARY) *
*                                           SSID CCA    SSID CCA   *
*DEVICE   LEVEL      STATE     PATH STATUS  SERIAL#     SERIAL#    *
*------ ---------  ----------  ----------- ---------    ---------  *
* 23F0  .........  SIMPLEX...   INACTIVE    23C0 30      .... ..   *
*       ...............       ........... 000001323562 ............*
* PATHS SAID/DEST STATUS: DESCRIPTION                              *
* ----- --------- ------  -------------------                      *
*   0   ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
*       ---- ----   00    NO PATH...........                       *
********************************************************************
ANTP0001I CQUERY COMMAND COMPLETED FOR DEVICE 23F0. COMPLETION CODE: 00
```

Figure 70. CQUERY Output from 23F0

3. Vary the secondary device online.

   If you want to bring the secondary device online to the same MVS system, you have to change the volser in the CRECOVER command, using the ID parameter. Figure 71 on page 98 shows a sample of the CRECOVER command to change the volser of the secondary from PPR000, which is the same as the primary volser, to SNP001 to allow the volume to be brought online.

```
//RBPPRC6A JOB ,MSGCLASS=H,NOTIFY=RBPPRC2
//GO EXEC PGM=IKJEFT01,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
 CRECOVER DEVN(X'23F0') PRIM(X'FF27' 1321029 X'30') -
          SEC(X'23C0' 1323562 X'30') -
          ID(PPR000 SNP000)
```

Figure 71. CRECOVER JCL to Change Volser of Secondary Device

4. SnapShot the secondary.
5. Vary the secondary device offline.
6. CESTPAIR RESYNC to copy any changed tracks to the secondary.

The CGROUP FREEZE command can be used in place of CSUSPEND to enable the SnapShot of all pairs on an LCU or across multiple LCUs. You then have to issue multiple CESTPAIR commands to synchronize the related pairs.

9.3 Using DFSMSdss

If you are using SnapShot to create inline batch backups, you cannot directly have these backups mirrored on the secondary system by using PPRC because you cannot SnapShot to a primary PPRC volume. However, you can use DFSMSdss to write to a PPRC primary volume, so using the virtual concurrent copy function you can still get the time-saving benefits of SnapShot in your batch window.

DFSMSdss can now invoke SnapShot through:

- DFSMSdss Snapshot
- Virtual concurrent copy

For a description of these functions, and considerations for thier implementation, see *Implementing DFSMSdss SnapShot and Virtual Concurrent Copy*, SG24-5268.

9.3.1 DFSMSdss SnapShot

When you use DFSMSdss COPY and the source and target are on the same RVA, providing no utility such as IDCAMS is invoked, the data set is snapped. No physical I/O occurs. In this case the data cannot be directed to a primary PPRC volume. If you are using DFSMSdss SnapShot to make backups, the output data sets should be on a volume that is not a PPRC primary.

However, if by mistake you direct the target data set to a PPRC primary volume, SnapShot is not invoked, and the data is physically copied by DFSMSdss.

9.3.2 Virtual Concurrent Copy

When you use DFSMSdss CONCURRENT COPY or DUMP with the source on an RVA, you can get the time-saving benefits of SnapShot by specifying the CONCURRENT keyword on the DUMP or COPY statement. SnapShot is internally invoked by a function called virtual concurrent copy.

Virtual concurrent copy snaps the source data set to an interim data set called the *working space data set* (WSDS). You need to make sure that you allocate the WSDS on a volume that is not a PPRC primary. DFSMSdss will move the data from the WSDS to the output data set using its traditional data mover techniques. Because the copy from the WSDS to the output data set invokes real I/O, it can be directed to a primary PPRC volume. The COPY or DUMP is logically complete when the snap completes. The data can then be made available for processing. The secondary site is not synchronized with the primary site until the physical copy is complete.

If the source and target volumes are on the same RVA, and DFSMSdss does not call a utility to move the data, SnapShot will be invoked for COPY even if you specify the CONCURRENT keyword. If you want to ensure that you get a data set snap rather than “real I/O,” then you should make sure that the source and target are on separate RVAs. Otherwise DFSMSdss SnapShot will result in a real DFSMS COPY without the time-saving benefits of SnapShot.

In Figure 72 on page 100 RVA A and B have primary volumes in PPRC pairings to secondary volumes on RVA C. If a DFSMSdss COPY CONCURRENT is specified to copy a data set from a primary PPRC volume on RVA A to a primary PPRC volume on RVA B, the following occurs:

1. The data is snapped to the WSDS on RVA A.
2. The data is physically copied to the volume on RVA B and is synchronously copied to the secondary volume on RVA C.

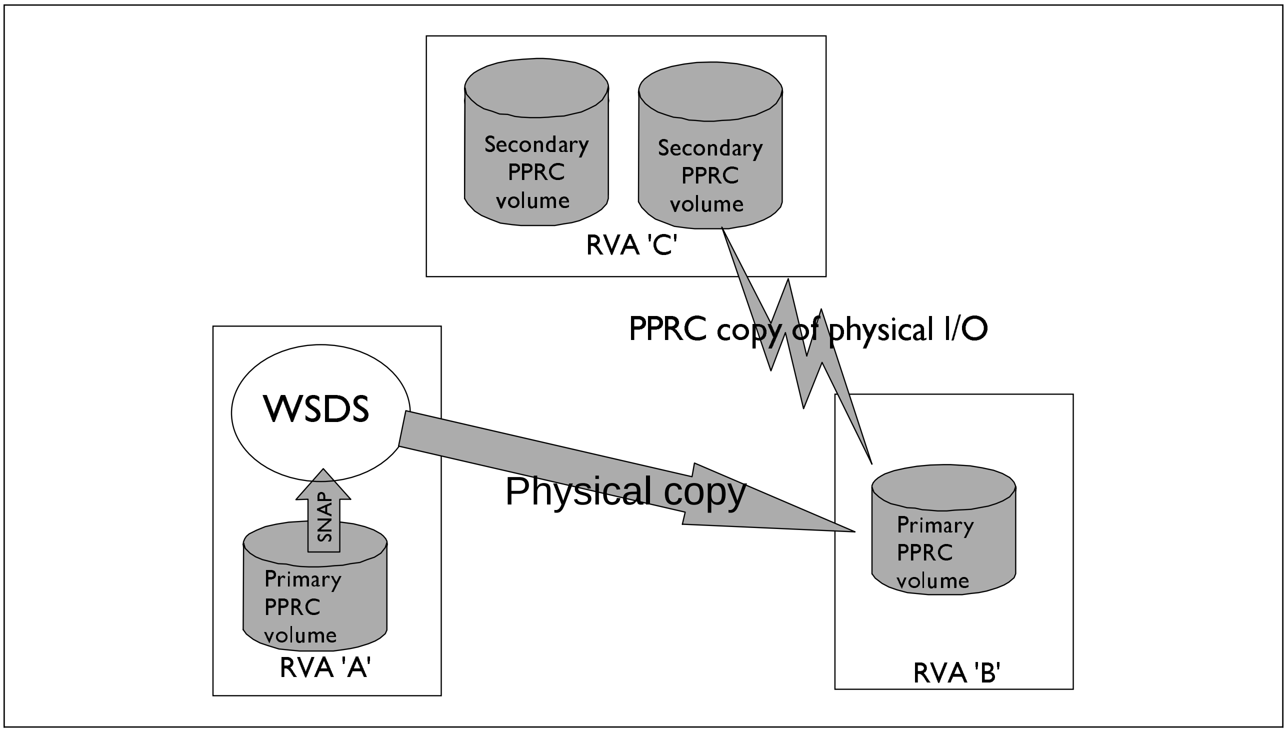

Figure 72. Virtual Concurrent Copy

You must preallocate the WSDS. As for a normal snap, the WSDS does not take up any physical space. If the source files are amended, they take up additional back-end storage. The space is not released until the physical move is complete.

For a more detailed explanation of virtual concurrent copy and the required hardware and software levels, see *Implementing DFSMSdss SnapShot and Virtual Concurrent Copy*, SG24-5268.

9.4 SnapShot before PPRC Resync

To ensure a recovery point during a PPRC RESYNC operation across multiple volumes, we recommend performing a SnapShot of all PPRC secondary volumes. If a disaster occurs during the RESYNC process, you still have a consistent copy of all the volumes to use for recovery.

A SnapShot of all secondary volumes can also be used to test your disaster recovery procedures without losing mirroring capability during the test.

9.5 NCL on Secondary RVA Subsystems

When you establish a PPRC pairing, all data is physically copied between the primary and secondary devices. Any data that has been duplicated using SnapShot on the primary device and may be sharing entries in the track number table (TNT) will have separate physical data at the secondary and therefore separate TNT entries.

Therefore, if you were to PPRC a whole RVA subsystem to a secondary RVA subsystem, the NCL on the secondary may be a lot higher. This could be critical if you are running at high levels of NCL on the primary subsystem. You can identify how much space you are saving from SnapShot, and how much you will need additionally on your secondary RVA from the IXFP LISTCFG reports. These reports give the total NCL and the Total Physical Space; the difference between them is your 'snapped' space.

Chapter 10. PPRC and SnapShot for System Testing

By combining PPRC and SnapShot, you have a unique tool for developing and testing new systems and applications.

10.1 Cloning Systems with PPRC

You can use PPRC to build a mirrored system for testing and development. There is no need for any direct channel connection between the mirror and “live” systems. Thus you can run the system in complete isolation. By using PPRC you avoid having to have a direct channel link from the production MVS to the test RVA to use a host application such as XRC or DFSMSdss, and you avoid tape dumps and restores.

If you want to maintain an exact duplicate of your production system, you have to have the same volume names. If all volumes can be seen by both production and test systems, at IPL the operator has to specify which volume should be brought online. Having duplicate volume names obviously presents a potential problem, which may be catastrophic if the wrong volume is brought online.

To make the initial copies, using either a host data mover or a tape dump, you need to have your system quiesced to guarantee a valid point-in-time copy. If you do not quiesce the system, catalogs and volumes could be out of synchronization. Recovery could be an unacceptably long process and have an impact on production applications if catalogs and data are not synchronized.

A tape restore process would also be lengthy. In some cases the tape subsystem is also kept isolated to avoid cross-system contamination and will more than likely be a manual operation.

You can use PPRC to create a mirror system. Once the pairs are synchronized, they can then be terminated to create the point-in-time system copy. Ideally you should quiesce the production workload, and suspend the PPRC pairs to create a consistent point in time copy. However, the quiesce is required only while you break the PPRC pairs. The quiesce could take place at a time of low or no system activity. Both the creation of the PPRC copies and the suspension of the pairs can be performed at times of low activity, with little or no impact on production applications. The PPRC links, once defined, can easily be reestablished to do further reloads from your production system if you need to maintain a test environment that closely reflects your production environment.

You can also use your disaster recovery site for application testing while maintaining disaster protection. Using SnapShot to copy your secondary volumes allows you to create a current copy of production data, without any disruption to production applications. Once you have snapped all the volumes, you can RESYNC your PPRC pairs and resume full mirroring.

10.2 SnapShot Recovery

In any testing, problems are frequently discovered, and the process must be rerun. In this situation you would restore the systems, make changes, and then rerun. This could be a lengthy and manual process.

You can use SnapShot to create quick and easy procedures for returning to the initial system image. By using snap back, you can test your application and system quickly and accurately.

You may also want to create multiple test environments for different development or test activities. In this case you can use SnapShot to create multiple copies of your data instantly and with minimal use of back-end disk storage. For example, you can use SnapShot to create images of your production data at key times, such as end-of-day, end-of-week, or end-of-month.

Remember, if you want to use SnapShot extensively for system testing, you must consider the impact on your available functional devices.

For additional information about using SnapShot for systems testing, see *Implementing SnapShot*, SG24-2241.

Chapter 11. Geographically Dispersed Parallel Sysplex

With the continued decrease in information technology (IT) costs and more emphasis on application availability and disaster recovery, enterprises are adopting a two-site strategy. IBM's S/390 multiple site application availability solution, the GDPS, integrates Parallel Sysplex technology and remote copy technology to enhance application availability and improve disaster recovery.

GDPS topology is a Parallel Sysplex cluster spread across two sites with all critical data mirrored between the sites. GDPS provides the capability to manage the remote copy configuration and storage subsystems, automates Parallel Sysplex operational tasks, and automates the recovery from a failure through a single point of control, thereby improving application availability. GDPS supports all transaction managers (for example, CICS and IMS) and database mangers (for example, DB2, IMS, and VSAM).

GDPS is enabled by means of the following key IBM technologies:

- Parallel Sysplex
- 9037-2 Sysplex Timer
- Coupling facility
- Systems Automation for OS/390
- Availability Operation Manager, or Site Manager 2.14 or higher
- RVA, 3990-6 or 9390
- PPRC
- 9729 Optical Wavelength Division Multiplexor (optional)

GDPS is a comprehensive offering that provides total IT business recovery through the management of processor, storage, and network resources across multiple sites. GDPS manages the application environment and the consistency of data along with the physical resources, providing full data integrity across volumes, subsystems, and sites. While optimizing the ability to perform a normal restart in the event of a site switch, GDPS minimizes the duration of the recovery window. Finally, GDPS uses PPRC to minimize or eliminate data loss and uses Parallel Sysplex cluster functions along with system automation to minimize the duration of the recovery window.

11.1 Requirement for Time Consistency of Data

Data consistency across all primary and secondary volumes spread across any number of storage subsystems is essential for data integrity and to enable a normal database restart in the event of a disaster. Data consistency is difficult to achieve for synchronous forms of remote copy because synchronous remote copy is entirely implemented within storage subsystem pairs. Data consistency is managed by the applications so that a consistent state of application data on disk is maintained if a failure occurs in the host processor, software, or storage subsystem. Applications where one write is dependent on the completion of another are said to have *dependent writes*.

Some common examples of dependent writes are:

- Database and the associated log file
- Catalogs and VTOCs
- VSAM index and data components

Because a disaster occurs over a period of time rather than at an instant in time, the data at the remote site must be managed so that cross-volume or subsystem data consistency is preserved during the intermittent or gradual failure.

GDPS is a multiple site management facility that is a combination of system code and automation that utilizes the capabilities of Parallel Sysplex technology, storage subsystem mirroring, and databases to manage processors, storage, and network resources. It is designed to minimize and potentially eliminate the impact of a disaster or planned site outage. It provides the ability to perform a controlled site switch for both planned and unplanned site outages, with no data loss, maintaining full data integrity across multiple volumes and storage subsystems and a normal DBMS restart (not DBMS recovery) at the recovery site. GDPS manages cross-site data consistency through a combination of storage subsystem, Sysplex and environmental triggers. GDPS provides the cross-site automation that works in combination with the disk error recovery procedures exploiting new storage subsystem functions such as CGROUP FREEZE to manage cross-site data consistency.

GDPS provides the following functions:

- Remote Copy Management Facility (RCMF) - automates the management of the remote copy infrastructure
- Planned reconfiguration support - automates a planned site switch
- Unplanned reconfiguration support - recovers from an OS/390, processor, storage subsystem, or site failure

11.2 Remote Copy Management Facility

The Remote Copy Management Facility (RCMF) is designed to simplify the storage administrator's remote copy management functions by managing the remote copy configuration rather than individual remote copy pairs. Storage management tasks include the initialization and monitoring of the PPRC volume pairs on the basis of policy and performing routine operations on installed storage subsystems.

RCMF can be useful in managing the RVA PPRC environment because it:

- Provides a central point of control through a full screen, tree-structured interface
- Provides the vehicle for initializing and maintaining the remote copy configuration
- Provides status monitoring and exception reporting
- Drives P/DAS for device migration

11.3 Planned Reconfiguration Support

GDPS planned reconfiguration support automates procedures performed by your operations center. The GPDS offering automates planned reconfigurations by controlling your policies (changes based on specified conditions or rules). Automation can be established for standard data center activities such as:

- Quiescing the system workload and removing the system from the Parallel Sysplex cluster by effectively stopping the system before the change window
- IPLing a system or starting it after the change window
- Quiescing the system workload, removing it from the Parallel Sysplex cluster, and re-IPLing the system to achieve a recycling of the system to pick up software maintenance
- Performing a complete planned site switch of the application, including network and DASD, from a single point of control.

For planned reconfigurations, GPDS offers a tested, proven site switch from a single point of control.

11.4 Unplanned Reconfiguration Support

GDPS is designed to minimize and potentially eliminate the amount of data loss and the duration of the recovery window in the event of a site failure. GDPS also minimizes the impact of and potentially masks an OS/390 system or processor failure on the basis of GDPS policy. GDPS uses PPRC to help minimize or eliminate data loss. Parallel Sysplex cluster functions along with automation are used to detect OS/390 system, processor, or site failures and to initiate recovery processing to help minimize the duration of the recovery window.

If an OS/390 system fails, the failed system is automatically removed from the Parallel Sysplex cluster, and re-IPLed in the same location if possible, and the workload is restarted. If it is not possible to restart the workload in the same location, GDPS will acquire resources based on policy and restart the workload.

In the event of a failure and based on policy, there is limited or no data loss because all critical data is synchronously mirrored through PPRC from the primary site to the recovery site. There is limited data loss if the production systems continue to update the primary copy of data after remote copy processing is suspended and if a subsequent disaster destroys some or all of the primary copy of data. No data is lost if the production systems do not make updates to the primary PPRC volumes after PPRC processing is suspended. GDPS provides various business policies to customize GDPS automation to meet your business requirements.

11.5 GDPS Service Offering

The GDPS solution is delivered as a service offering by IBM services.

With the RCMF service offering, the RCMF automation to manage the remote copy infrastructure is installed, the automation policy is customized, the automation is verified, operational education is provided for the enterprise.

With the GDPS service offering, an initial planning session is held to understand your application availability and disaster recovery objectives. Once these objectives are understood, GDPS automation to manage the multi-site infrastructure, perform routine tasks, and recover from failures is installed. The automation policy is customized, the automation is verified, and education is provided to the your operations staff.

For additional information about GDPS, contact your IBM representative.

Appendix A. Sample PPRC Implementation Project Plan

With this sample project plan, we intend to provide a building block for a customized RVA PPRC implementation plan. On the basis of the information in this sample plan, the project leader should be able to size the PPRC project in terms of skills, hardware and software requirements, management and operational procedures, and implementation time needed.

We present the objectives and considerations you have to pay attention to for the following topics of planning:

- A.1.1, " Education on PPRC Functions and the RVA"
- A.1.2, " Planning Hardware Requirements"
- A.1.3, " Planning Application Recovery Requirements" on page 110
- A.1.4, " Planning Database Recovery Requirements" on page 111
- A.1.5, " Capacity Planning" on page 111
- A.1.6, " Planned Outages" on page 112
- A.1.7, " Functional Testing of PPRC" on page 112
- A.1.8, " Testing Operational and Management Procedures" on page 112
- A.1.9, " Testing Error Recovery Procedures" on page 113
- A.1.10, " Performance Testing" on page 113
- A.1.11, " Cutover to Production" on page 113

This sample plan is based on *Remote Copy Installation Planning* by Andries J.M. de Jong and the *Remote Copy Technical Update* from IBM SSD in MKTTOOLS.

A.1.1 Education on PPRC Functions and the RVA

A.1.1.1 Objectives

Determine the level of understanding required to size the project, put in place the skills required, develop a tailored implementation plan, and successfully execute it.

A.1.1.2 Considerations

- Classes are scheduled throughout the year. You can obtain schedules from your local IBM Education Center.
- You can use as references the documentation that we refer to in this book. You can find the listing in Appendix D, "Related Publications" on page 121. Ask your IBM representative how you can obtain them.

A.1.2 Planning Hardware Requirements

A.1.2.1 Objectives

Determine the different requirements of the hardware you will use.

Evaluate the specifications regarding environment and power.

Identify the service providers you will coordinate with regarding power, lighting, cooling, communications, and other requirements.

A.1.2.2 Considerations

- The RVA is a dual power supply machine. Take advantage of this feature by connecting the power cords to separate power sources.
- The secondary RVA must be set to LOCAL power or must not be part of the remote power control of the secondary site host.
- The Call Home feature of the RVA must be active for both primary and secondary subsystems.
- Evaluate the benefit of using an ESCON Director, taking into consideration additional management procedures to implement it.
- If you are going to use ESCON Directors, you must ensure that they and their corresponding PC monitors are connected to power sources under an uninterruptible power source.
- Evaluate the distance between the primary and secondary sites and assess the available resources to connect them.
- Consider a redundancy connection for your physical links between sites so that connectivity can be maintained if one physical routing is damaged.

A.1.3 Planning Application Recovery Requirements

A.1.3.1 Objectives

Identify the applications that must be able to survive a site outage.

Separate your critical application data from noncritical data.

Decide whether you will make copy the entire system, or a subset of your applications.

A.1.3.2 Considerations

- Categorize your data into major categories such as system, catalogs, and databases, to determine which should be remotely copied.
- Determine whether you can run your applications in isolation.
- Determine whether you can restart your applications after a failure and whether the same procedures are still applicable.
- Evaluate your need to access application data on the primary site from the secondary site, considering its impact on response time and especially management complexity.
- Software support for PPRC is provided through PTFs to TSO for the commands and to EREP for error status and reporting. These PTFs are available on demand from the IBM Software Support Center. Request assistance from your IBM representative.
- Review any related APARs for applicable software maintenance and LIC requirements. Obtain information APAR II08303.

- Assess your need to implement P/DAS.

A.1.4 Planning Database Recovery Requirements

A.1.4.1 Objectives

Maintain write sequence integrity of database systems.

Establish measures to protect databases against a rolling disaster.

A.1.4.2 Considerations

- IMS considerations

  Multiple RECONs, WADSs, and OLDSs may be defined and spread across different storage controls.

  IMS databases should be defined and shadowed with the CGROUP(YES) parameter. Your procedures should then be able to handle the IEA494I message to issue a CGROUP/FREEZE to the logical control units.

  IMS system and library volumes can be maintained and shadowed with the CGROUP(YES) parameter. Your procedures should then be able to handle the IEA494I message to issue a CGROUP FREEZE to the logical control units.

  IMS archive data sets must have output at both the primary and secondary subsystems.

  PPRC is not usually used with DEDBs and MADSs.

- DB2 considerations

  Because it is critical to keep all DB2 volumes synchronized to ensure full recovery at the secondary site, all DB2 logs, BSDSs, and data volumes should be defined with the CGROUP(YES) parameter. Your procedures should then be able to handle the IEA494I message to issue a CGROUP FREEZE to the logical control units.

A.1.5 Capacity Planning

A.1.5.1 Objectives

Design a configuration that will have optimum throughput with stable performance to handle future applications and PPRC workload.

Come up with an expectation to evaluate test and production performance.

A.1.5.2 Considerations

- Determine whether you have adequately configured your RVA subsystems for primary and secondary, including cache sizes, pathing, and projected storage path utilization.
- Perform a performance evaluation of your application system operating with the planned PPRC configuration, taking into account first the critical workload.
- Consider using DASD Magic/2 for predicting the performance of PPRC on the RVA. Request assistance from your IBM Marketing personnel to obtain DASD Magic/2.

A.1.6 Planned Outages

A.1.6.1 Objective

Ensure that transfer to the secondary due to planned outages will be smoother and faster after implementing PPRC.

A.1.6.2 Considerations

- Assess your need for planned outages and set their schedules.
- Use your planned outages to test your disaster recovery process and PPRC recovery and management.
- Evaluate the preparedness of your personnel in a real problem situation.
- Evaluate the need to improve current recovery processes.

A.1.7 Functional Testing of PPRC

A.1.7.1 Objectives

Become familiar with PPRC commands and their system response.

Become familiar with information sources.

Test your procedures on moving back to primary.

Test fallback procedures.

A.1.7.2 Considerations

- Document resulting error conditions and their resolution.
- Record every step of your test activities and the results in each step.
- Consider repeating the test activities to become familiar with them.
- When you have repeated successful tests, inject errors to simulate different scenarios.
- Be aware that, for the RVA, sparing and concurrent repair will not terminate your PPRC pairs.

A.1.8 Testing Operational and Management Procedures

A.1.8.1 Objectives

Develop procedures for normal operations with the intent of automating as much as possible.

Protect your PPRC resources from unauthorized access.

Create clear documentation for procedures and configurations.

Educate all parties and personnel involved in the execution of these procedures.

A.1.8.2 Considerations

- Document manual recovery procedures to use if the automated process fails.
- Determine whether you need to have a controlling system or LPAR to monitor the status of your pairs.
- Consider a standard method of managing your ESCON links, especially if you are using an ESCON Director. Because PPRC paths are not known to the ESCON Manager, we recommend that you use a combination of the ESCON Manager and a documented ESCON link configuration or worksheet (see Appendix B, "Sample ESCON Configuration Worksheet" on page 115) in managing your ESCON links.
- Document your procedures and configurations and establish a mechanism that will trigger your change management process every time updates to these procedures and configurations occur.
- Set up a process, a procedure, or an alert for PPRC operations that would notify you of failures at the secondary and allow a primary volume to continue receiving updates when the copy cannot be sent to its secondary.
- Establish PPRC access authority levels, using your security system.

A.1.9 Testing Error Recovery Procedures

A.1.9.1 Objective

Evaluate the recovery procedures for both the primary and secondary sites.

A.1.9.2 Considerations

- Study your need to set up an automated process to check your PPRC status.
- Set up an automated process to alert your operations of failures.
- Establish the acceptable recovery time on the secondary and the return time to the primary.

A.1.10 Performance Testing

A.1.10.1 Objectives

Develop performance management procedures.

Test that applications will continue to perform satisfactorily after implementing PPRC.

A.1.10.2 Considerations

- Obtain DASD response time and cache statistics reports.
- Pay special attention to the difference in disconnect times between PPRC and non-PPRC runs.
- Remember to generate reports for both the primary and secondary devices.

A.1.11 Cutover to Production

A.1.11.1 Objective

Install and test the final configuration at the primary and secondary sites for PPRC operation.

A.1.11.2 Considerations

- Consider segmenting the process of taking the PPRC configuration in production and gradually expanding the number of active remote copy pairs.
- Monitor application response to avoid performance or availability problems.
- Establish the criteria for moving to the secondary and who will make the decision.

Appendix B. Sample ESCON Configuration Worksheet

Figure 73 on page 116 shows a sample worksheet to document your primary and secondary RVA SSIDs and ESCON link SAID addresses. The SSIDs and SAIDs are your input in the LINK parameter of the CESTPATH command.

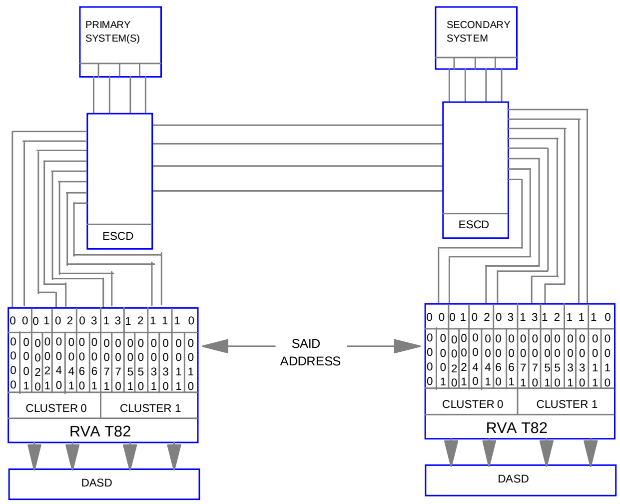

	PRI	SEC	PRI	SEC	PRI	SEC	PRI	SEC
SSID	_____	_____	______	______	______	_____	______	______
SERIAL	_____	_____	______	______	______	_____	______	______
	SAID FROM	DEST TO	SAID FROM	DEST TO	SAID FROM	DEST TO	SAID FROM	DEST TO
	X'00___	___00'	X'00______01'		X'00______02'		X'00______03'	
	X'00___	___00'	X'00______01'		X'00______02'		X'00______03'	
	X'00___	___00'	X'00______01'		X'00______02'		X'00______03'	
	X'00___	___00'	X'00______01'		X'00______02'		X'00______03'	

Figure 73. RVA ESCON Worksheet

Appendix C. Special Notices

This publication is intended to help IBM and customer technical and management personnel plan and execute an implementation of PPRC on the RAMAC Virtual Array Model T82. The information in this publication is not intended as the specification of any programming interfaces that are provided by PPRC. See the PUBLICATIONS section of the IBM Programming Announcement for PPRC for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

DATABASE 2	DB2
DFSMS	DFSMS/MVS
DFSMS/VM	DFSMSdfp
DFSMSdss	DFSMShsm
Enterprise Systems Connection Architecture	ESCON XDF
ESCON	Hardware Configuration Definition
IBM	IMS
IMS/ESA	MVS/ESA
NetView	RAMAC
Resource Measurement Facility	S/390
S/390 Parallel Enterprise Server	Virtual Machine/Enterprise Systems Architecture
VM/ESA	VSE/ESA

The following terms are trademarks of other companies:

SnapShot is a trademark and property of Storage Technology Corporation for a duplication product.

IXFP	Storage Technology Corporation
Extended Storage Architecture	Storage Technology Corporation
Iceberg	Storage Technology Corporation
XSA	Storage Technology Corporation
XSA/Reporter	Storage Technology Corporation
StorageTek	Storage Technology Corporation
SAS	SAS Institute, Incorporated
SAS/C	SAS Institute, Incorporated
SAS/GRAPH	SAS Institute, Incorporated
SAP and R/3	SAP AG

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 123.

- *Planning for IBM Remote Copy*, SG24-2595
- *IBM RAMAC Virtual Array*, SG24-4951
- *ESCON Implementation Guide*, SG24-4662
- *Implementing DFSMSdss SnapShot and Virtual Concurrent Copy*, SG24-5268
- *Implementing SnapShot*, SG24-2241
- *Using RVA and SnapShot for Business Intelligence Applications with OS/390 and DB2*, SG24-5333
- *Disaster Recovery Library: S/390 Technology*, SG24-4210-01
- *P/DAS and Enhancements to the 3990-6 and RAMAC Array Family*, SG24-4724
- *Fire in the Computer Room, What Now?*, SG24-4211-02

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

D.3 Other Publications

These publications are also relevant as further information sources:

- *Remote Copy Administrator's Guide and Reference*, SC35-0169
- *IOCP User's Guide and ESCON Channel-to-Channel Reference*, GC38-0401
- *IMS/ESA V5 General Information*, GC26-3467
- *IMS/ESA V5 Operations Guide*, SC26-8029
- *Planning for ESCON Manager*, GC23-0423

- *ICKDSF R16 Refresh User's Guide*, GC35-0033
- *Remote Copy Technical Update*, IBM/SSD
- *Remote Copy Installation Planning*, Andries J.M. de Jong
- *PPRC Migration Manager for MVS: User's Guide*, Don J. Chesarek

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** http://www.redbooks.ibm.com/

  Search for, view, download or order hardcopy/CD-ROMs redbooks from the redbooks web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders via e-mail including information from the redbook order form to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)	1-800-879-2755	
Canada (toll free)	1-800-IBM-4YOU	
Outside North America	(long distance charges apply)	
(+45) 4810-1320 - Danish	(+45) 4810-1220 - French	(+45) 4810-1270 - Norwegian
(+45) 4810-1420 - Dutch	(+45) 4810-1020 - German	(+45) 4810-1120 - Spanish
(+45) 4810-1540 - English	(+45) 4810-1620 - Italian	(+45) 4810-1170 - Swedish
(+45) 4810-1670 - Finnish		

This information was current at the time of publication, but is continually subject to change. The latest information for customers may be found at http://www.redbooks.ibm.com/ and for IBM employees at http://w3.itso.ibm.com/.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may also view redbook, residency and workshop announcements at http://inews.ibm.com/.

IBM Redbook Fax Order Form

Fax your redbook orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

Please send me the following:

Title	Order Number	Quantity

First name　Last name

Company

Address

City　Postal code　Country

Telephone number　Telefax number　VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date　Card issued to　Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Numerics

3990-6, PPRC characteristics
9036, description 27
9390, architecture
9729, description 27

A

application testing 103
automation
  DB2 92
  messages 52

C

cache
  configuration 78
  PPRC requirements 12
  settings 10
call home 12
CCA 32
CESTPAIR 53
CESTPATH 39, 44
CGROUP 10, 55
CGROUP FREEZE 52
CGROUP, automation 60
channel connection address
  See CCA
channel register set 38
Commands
  CESTPAIR 53
  CESTPATH 39, 44, 52
  CGROUP 10, 52, 55, 60
  CGROUP FREEZE 52
  CQUERY 36, 55, 62
  CRIT 10
  CSUSPEND 54
  DEVSERV 43
  ICKDSF 56
  PPRCOPY 56, 84
  protection 67
concurrent copy 99
configuration worksheet 115
consistency group timer 10, 51
consistency grouping 53
control unit address
  See CUADD
CQUERY 55, 62
CQUERY PATHS 36
CQUERY, output 63
CRIT 10
CRIT parameter 88
cross platform extension
  See XPE
CSUSPEND 54
CUADD 40

D

data consistency
  management 105
database
  DB2 88
  IMS 87
  integrity 20
  performance 80
  recommendations 87
  recovery 19, 87
  restart 19, 87
DB2
  automation 92
  CGROUP FREEZE 92
  consistency grouping 91
  CRIT parameter 88
  recommendations 88
  recovery 89
DDC 33
DDSR 61
deleted data space release
  See DDSR
DEVSERV 43
DFSMSdss snapshot 98
director device connection
  See DDC
disaster recovery
  consistency 54
  database 87
  fallback planning 66
  requirements 5
  scenarios 6
  SnapShot 100
dying scenario 6

E

ECAM 10
ERP 10, 54, 56
error messages 52
error recovery procedure
  See ERP
ESCON
  architecture 16
  cards 35

ESCON *(continued)*
channel extender 27
channel mode 36
channel register set 38
configuration 45
configuration worksheet 115
connection modes 23
connectivity 10
control unit mode 36
director 24
distance 27, 78
dual port mode 35
ESCON manager 68
fiber specification 25
jumper cable 25
logical path 38
multiplexer 27
performance 77, 78
single port mode 35
transmission mode 26
VM 84
XDF 26
extended control and monitoring
See ECAM
extended remote copy
See XRC

G
GDPS
description 105
error recovery 61
planned reconfiguration 107
service offering 105
unplanned reconfiguration 107
utility device 11, 51
volume groups 56

H
hardware configuration definition
See HCD
HCD 24

I
I/O configuration data set
See IOCDS
IBM extended facilities program
See IXFP
ICE cards 35
ICKDSF 56
IEA491E 59
IEA494I 59
IMS
recommendations 87

IMS *(continued)*
timestamp recovery 88
IMS remote site recovery
See RSR
IOCDS
ESCON Director 24
IXFP 11
DDSR 61
ECAM 65
exclusion table 61
IXFP. 61

L
LCU 31
LIC 12, 66
licenced internal code
See LIC
logical control unit
See LCU
long busy state 10

M
maintenance procedures 10
messages
automation 52
error processing 59
IEA491E 59
IEA494I 59
SIB4621E 96
migration manager, description 14
multimode fiber 26
MuxMaster 27

N
NCL 100
net capacity load
See NCL

O
operations
procedures 68
VM 83

P
P/DAS 66
PACE parameter 53, 78
Parallel Sysplex 105
performance 81
distance 78
establish 78
recommendations 81
sequential 81

performance *(continued)*
tools 81
pinned data 10
PPRC
architecture 9
automation 55, 57, 63, 87, 107
CCA 32
command support 52
data flow 3
data integrity 20
DB2 88
DDC 33
documentation 68
error recovery 59
errors 65
ESCON connection 42
ESCON requirements 10
hardware requirements 11
ICKDSF 83
implementation 109
IMS 87
LIC 66
logical path 38, 47
management 13, 55, 106
migration manager 14, 57
open systems 18
operations 51
paths 55
performance 77
procedures 68
secondary device 55, 61, 96
security 67
SMS 63, 66
SnapShot considerations 11
software 56
software requirements 13
systems cloning 103
transfers 77
VM support 56, 83
VSE/ESA 84
PPRC dynamic address switching
See PDAS
PPRC migration manager 57
PPRCOPY 84
project plan, sample 109

R

RACF 67
RCMF 61, 106
recovery
data consistency 21
distance 9
strategy 7
remote copy management facility
See RCMF
rolling disaster 6, 54, 61
RSR 88
RVA
addressing 34
architecture 1
cache 78
channel adapter 37
CUADD 40
ESCON cards 35
hardware characteristics 51
ICE cards 35
LCU 31
maintenance 56
PPRC characteristics 10
serial number 31

S

SAID 36
serial number 31, 52
service information message
See SIM
SIB4621E 96
SIMs 10, 51
single mode fiber 26
SMS 66, 95
SnapShot
allocation 95
DFSMSdss 98
disaster recovery 100
messages 96
NCL 100
PPRC considerations 11
PPRC interactions 95
primary device 95
recommendations 96
secondary device 96
SMS 95
systems cloning 103
virtual concurrent copy 98
volume preferencing 95
software support
APARS 13
VM/ESA 83
VSE/ESA 84
SSID 31, 52
subsystem ID
See SSID
system adapter identifier
See SAID
systems cloning 103
systems managed storage
See SMS

T

terminology 4

V

virtual concurrent copy 98
vital product data
  See VPD
VM
  commands 83
  data consistency 84
  ESCON 84
  PPRC support 56
VPD 10

W

wavelength division multiplexer 27
write sequence integrity 20, 87

X

XPE 18
XRC 8

ITSO Redbook Evaluation

RAMAC Virtual Array: Implementing Peer-to-Peer Remote Copy
SG24-5338-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and Fax it to: USA International Access Code + 1 914 432 8264 or:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
__Customer __Business Partner __Solution Developer __IBM employee
__None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____________

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

RAMAC Virtual Array: Implementing Peer-to-Peer Remote Copy

SG24-5338-00

SG24-5338-00
Printed in the U.S.A.