

# IBM System Storage Copy Services and IBM i:

## A Guide to Planning and Implementation

Discover IBM i 6.1 and DS8000 R3 Copy Services enhancements

Learn how to implement Copy Services through a GUI and DS CLI

Understand the setup for Metro Mirror and Global Mirror



Hernando Bedoya  
Nick Harris  
Ron Devroy  
Ingo Dimmer  
Adrian Froom  
Dave Lacey  
Veerendra Para  
Will Smith  
Herbert Velasquez  
Ario Wicaksono

**Redbooks**





International Technical Support Organization

**IBM System Storage Copy Services and IBM i:  
A Guide to Planning and Implementation**

September 2008

Archived

**Note:** Before using this information and the product it supports, read the information in “Notices” on page ix.

Archived

**Third Edition (September 2008)**

This edition applies to IBM i5/OS Version 6, Release 1.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	ix
Trademarks .....	x
<b>Preface</b> .....	xi
The team that wrote this book .....	xi
Become a published author .....	xiii
Comments welcome .....	xiv
<b>Part 1. Introduction</b> .....	1
<b>Chapter 1. Introduction to Copy Services and System i high availability</b> .....	3
1.1 Overview of the i5/OS architecture .....	4
1.1.1 Two-part primary operating system .....	4
1.1.2 Object-based system .....	4
1.1.3 Single-level storage .....	5
1.2 Software-based high availability solutions .....	5
1.2.1 Local journaling .....	6
1.2.2 Remote journaling .....	6
1.2.3 Object types not journaled .....	8
1.2.4 Replicating non-journaled object types .....	8
1.3 i5/OS clustering .....	9
1.3.1 Definition of a cluster .....	9
1.3.2 Cluster components .....	10
1.4 Auxiliary storage pools .....	13
1.4.1 Definition of an auxiliary storage pool .....	13
1.4.2 Definition of an independent user ASP .....	14
1.5 Cross-site mirroring concepts .....	14
1.5.1 Definition of cross-site mirroring .....	14
1.5.2 Definition of geographic mirroring .....	15
1.5.3 Failover and switchover .....	17
1.5.4 Supported and unsupported i5/OS object types .....	17
1.5.5 Benefits of cross-site mirroring .....	17
1.5.6 Limitations of cross-site mirroring .....	18
1.6 Copy Services based disaster recovery and high availability solutions .....	18
1.6.1 FlashCopy solutions .....	19
1.6.2 Metro Mirror and Global Mirror solutions .....	21
1.6.3 System i Copy Services usage considerations .....	23
1.6.4 System i Copy Services management solutions .....	24
<b>Chapter 2. System i external storage solution examples</b> .....	31
2.1 One-site System i external storage solution examples .....	32
2.1.1 System i5 model and all disk storage in external storage .....	32
2.1.2 System i model with internal load source and external storage .....	34
2.1.3 System i model with mixed internal and external storage .....	34
2.1.4 Migration of internal drives to external storage including load source .....	36
2.1.5 Migration of an external mirrored load source to a boot load source .....	37
2.1.6 Cloning i5/OS .....	38
2.1.7 Full system and IASP FlashCopy .....	38
2.2 Two-site System i external storage solution examples .....	40

2.2.1	The System i platform and external storage HA environments. . . . .	40
2.2.2	Metro Mirror examples . . . . .	41
2.2.3	Global Mirror examples. . . . .	43
2.2.4	Geographic mirroring with external storage . . . . .	47
<b>Part 2.</b>	<b>Planning and sizing . . . . .</b>	<b>49</b>
<b>Chapter 3.</b>	<b>i5/OS planning for external storage . . . . .</b>	<b>51</b>
3.1	Planning for external storage solutions . . . . .	52
3.2	Solution implementation considerations . . . . .	54
3.2.1	Planning considerations for boot from SAN . . . . .	54
3.2.2	Planning considerations for i5/OS multipath Fibre Channel attachment. . . . .	57
3.2.3	Planning considerations for Copy Services. . . . .	62
3.2.4	Planning storage consolidation from different servers . . . . .	66
3.2.5	Planning for SAN connectivity. . . . .	67
3.2.6	Planning for capacity. . . . .	67
3.2.7	Planning considerations for performance . . . . .	85
<b>Chapter 4.</b>	<b>Sizing external storage for i5/OS . . . . .</b>	<b>89</b>
4.1	General sizing discussion . . . . .	90
4.1.1	Flow of I/O operations. . . . .	90
4.1.2	Description of response times. . . . .	93
4.2	Rules of thumb . . . . .	94
4.2.1	Number of RAID ranks . . . . .	95
4.2.2	Number of Fibre Channel adapters. . . . .	97
4.2.3	Size and allocation of logical volumes . . . . .	100
4.2.4	Sharing ranks among multiple workloads . . . . .	102
4.2.5	Connecting using switches . . . . .	102
4.2.6	Sizing for multipath . . . . .	103
4.2.7	Sizing for applications in an IASP. . . . .	104
4.2.8	Sizing for space efficient FlashCopy. . . . .	104
4.3	Sizing tools . . . . .	106
4.3.1	Disk Magic . . . . .	106
4.3.2	IBM Systems Workload Estimator for i5/OS . . . . .	106
4.3.3	IBM System Storage Productivity Center for Disk. . . . .	107
4.4	Gathering information for sizing . . . . .	110
4.4.1	Typical workloads in i5/OS . . . . .	110
4.4.2	Identifying peak periods . . . . .	111
4.4.3	i5/OS Performance Tools . . . . .	111
4.5	Sizing examples with Disk Magic . . . . .	113
4.5.1	Sizing the System i5 with DS8000 for a customer with iSeries model 8xx and internal disks . . . . .	113
4.5.2	Sharing DS8100 ranks between two i5/OS systems (partitions). . . . .	137
4.5.3	Modeling System i5 and DS8100 for a batch job currently running Model 8xx and ESS 800. . . . .	150
4.5.4	Using IBM Systems Workload Estimator connection to Disk Magic: Modeling DS6000 and System i for an existing workload. . . . .	163
<b>Part 3.</b>	<b>Implementation and additional topics . . . . .</b>	<b>179</b>
<b>Chapter 5.</b>	<b>Implementing external storage with i5/OS . . . . .</b>	<b>181</b>
5.1	Supported environment. . . . .	182
5.1.1	Hardware . . . . .	182
5.1.2	Software . . . . .	182

5.2	Logical volume sizes . . . . .	183
5.3	Protected versus unprotected volumes . . . . .	184
5.3.1	Changing LUN protection . . . . .	184
5.4	Setting up an external load source unit . . . . .	185
5.4.1	Tagging the load source IOA . . . . .	185
5.4.2	Creating the external load source unit . . . . .	191
5.5	Adding volumes to the System i5 configuration . . . . .	195
5.5.1	Adding logical volumes using the 5250 interface . . . . .	196
5.5.2	Adding volumes to an independent auxiliary storage pool . . . . .	199
5.6	Adding multipath volumes to System i using a 5250 interface . . . . .	206
5.7	Adding volumes to System i using iSeries Navigator . . . . .	208
5.8	Managing multipath volumes using iSeries Navigator . . . . .	211
5.9	Changing from single path to multipath . . . . .	214
5.10	Protecting the external load source unit . . . . .	215
5.10.1	Setting up load source mirroring . . . . .	220
5.11	Migration from mirrored to multipath load source . . . . .	223
5.12	Migration considerations from IOP-based to IOP-less Fibre Channel . . . . .	241
5.12.1	IOP-less migration in a multipath configuration . . . . .	241
5.12.2	IOP-less migration in a mirroring configuration . . . . .	241
5.12.3	IOP-less migration in a configuration without path redundancy . . . . .	242
5.13	Resetting a lost multipath configuration . . . . .	242
5.13.1	Resetting the lost multipath configuration for V6R1 . . . . .	242
5.13.2	Resetting a lost multipath configuration for versions prior to V6R1 . . . . .	245
<b>Chapter 6. Implementing FlashCopy using the DS CLI . . . . .</b>		<b>253</b>
6.1	Overview of IBM System Storage DS CLI . . . . .	254
6.1.1	Installing and setting up the DS CLI . . . . .	254
6.2	Implementing traditional FlashCopy . . . . .	255
6.2.1	Turning off the source server . . . . .	256
6.2.2	Setting up the FlashCopy environment . . . . .	256
6.2.3	Performing an IPL of the source server . . . . .	258
6.2.4	Performing an IPL of the target server . . . . .	259
6.3	Implementing space efficient FlashCopy . . . . .	268
6.3.1	Configuring space efficient FlashCopy for a System i environment . . . . .	268
6.3.2	Using space efficient FlashCopy with i5/OS . . . . .	272
6.3.3	Reactions of System i partitions to a full repository . . . . .	275
<b>Chapter 7. Implementing Metro Mirror using the DS CLI . . . . .</b>		<b>279</b>
7.1	Overview of the test environment . . . . .	280
7.2	Creating a Metro Mirror environment setup . . . . .	281
7.2.1	Creating Peer-to-Peer Remote Copy paths . . . . .	282
7.2.2	Creating a Metro Mirror relationship . . . . .	287
7.3	Switching over the system from the local site to remote site . . . . .	290
7.3.1	Making the volumes available on the remote site . . . . .	291
7.3.2	Performing an IPL of the backup server on the remote site . . . . .	292
7.4	Switching back the system from the remote site to local site . . . . .	293
7.4.1	Starting Metro Mirror from the remote site to local site (reverse direction) . . . . .	293
7.4.2	Making the volumes available on the local site . . . . .	295
7.4.3	Performing an IPL of the production server on the local site . . . . .	297
7.4.4	Starting Metro Mirror from the local site to remote site (original direction) . . . . .	298
<b>Chapter 8. Implementing Global Mirror using the DS CLI . . . . .</b>		<b>301</b>
8.1	Overview of the test environment . . . . .	302
8.2	Creating a Global Mirror environment . . . . .	303

8.2.1	Creating Peer-to-Peer Remote Copy paths	304
8.2.2	Creating a Global Copy relationship	310
8.2.3	Creating a FlashCopy relationship	312
8.2.4	Creating a Global Mirror session	314
8.2.5	Starting a Global Mirror session	316
8.3	Switching over the system from the local site to remote site	321
8.3.1	Making the volumes available on the remote site	321
8.3.2	Checking and recovering the consistency group of the FlashCopy target volume	323
8.3.3	Reversing a FlashCopy relationship	326
8.3.4	Recreating a FlashCopy relationship	327
8.3.5	Performing an IPL of the backup server on the remote site	327
8.4	Switching back the system from the remote site to local site	328
8.4.1	Starting Global Copy from the remote site to local site (reverse direction)	329
8.4.2	Making the volumes available on the local site	331
8.4.3	Starting Global Copy from the local site to remote site (original direction)	333
8.4.4	Checking or restarting a Global Mirror session	335
8.4.5	Performing an IPL of the production server on the local site	335
	<b>Chapter 9. Copy Services scenarios</b>	<b>337</b>
9.1	Scenarios for System i using the DS GUI	338
9.1.1	Scenario background	338
9.1.2	Test scenarios	338
9.1.3	Accessing the DS GUI interface	338
9.1.4	System i5 models	341
9.1.5	IBM System Storage server	341
9.1.6	Test setup	342
9.2	FlashCopy scenario	343
9.3	Metro Mirror scenario	343
9.4	Global Mirror scenario	345
	<b>Chapter 10. Creating storage space for Copy Services using the DS GUI</b>	<b>347</b>
10.1	Creating an extent pool	348
10.2	Creating logical volumes	355
10.3	Creating a volume group	360
	<b>Chapter 11. Implementing FlashCopy using the DS GUI</b>	<b>363</b>
	<b>Chapter 12. Implementing Metro Mirror using the DS GUI</b>	<b>373</b>
12.1	Metro Mirror arrangement	374
12.2	Implementing Metro Mirror volume relationships	374
12.3	Displaying Metro Mirror volume properties	382
	<b>Chapter 13. Managing Copy Services in i5/OS environments using the DS GUI</b>	<b>387</b>
13.1	FlashCopy options	388
13.1.1	Make relationships persistent	388
13.1.2	Initiate background copy	388
13.1.3	Enable change recording	388
13.1.4	Permit FlashCopy to occur if target volume is online for host access	389
13.1.5	Establish target on existing Metro Mirror source	389
13.1.6	Inhibit writes to target volume	389
13.1.7	Fail relationship if space-efficient target volume becomes out of space	389
13.1.8	Write inhibit the source volume if space-efficient target volume becomes out of space	389



13.1.9	Sequence number for these relationships . . . . .	389
13.2	FlashCopy GUI . . . . .	390
13.2.1	Delete . . . . .	390
13.2.2	Initiate Background Copy . . . . .	392
13.2.3	Resync Target. . . . .	393
13.2.4	FlashCopy Reversible . . . . .	395
13.2.5	Reverse FlashCopy . . . . .	399
13.3	Metro Mirror GUI . . . . .	401
13.3.1	Recovery Failover . . . . .	401
13.3.2	Recovery Failback . . . . .	403
13.3.3	Suspend . . . . .	405
13.3.4	Resume. . . . .	406
13.4	Global Mirror GUI . . . . .	408
13.4.1	Create . . . . .	409
13.4.2	Delete . . . . .	412
13.4.3	Modify . . . . .	413
13.4.4	Pause . . . . .	415
13.4.5	Resume. . . . .	417
13.4.6	View session volumes. . . . .	418
13.4.7	Properties . . . . .	419
	<b>Chapter 14. Performance considerations . . . . .</b>	<b>423</b>
14.1	Configuration of the DS system . . . . .	424
14.2	Connectivity between the DS systems and System i environment . . . . .	424
14.2.1	Physical connections . . . . .	425
14.3	Connectivity between the DS systems . . . . .	425
14.3.1	Physical connections . . . . .	426
14.3.2	Logical connections . . . . .	427
14.3.3	Using independent ASPs with Metro Mirror . . . . .	427
	<b>Chapter 15. FlashCopy usage considerations . . . . .</b>	<b>431</b>
15.1	Using i5/OS quiesce for Copy Services . . . . .	432
15.2	Using BRMS and FlashCopy . . . . .	436
15.3	BRMS architecture . . . . .	438
15.4	Enabling BRMS to use FlashCopy . . . . .	438
15.4.1	Preliminary notification of FlashCopy mode . . . . .	439
15.4.2	Pre-backup step on backup system . . . . .	439
15.4.3	Setting the BRMS system state to backup system . . . . .	440
15.4.4	Setting the backup system to restricted state TCP/IP. . . . .	440
15.4.5	Changing hardware resource names on the backup system . . . . .	441
15.5	Performing the backup from the backup system. . . . .	441
15.6	Post FlashCopy steps . . . . .	442
15.6.1	Indicating that the BRMS backup activity is complete. . . . .	442
15.6.2	Sending QUSRBRM to the production system . . . . .	442
15.6.3	Indicating that the FlashCopy function is complete on the production system . . . . .	443
15.7	Daily maintenance in BRMS . . . . .	443
15.8	Printing recovery reports . . . . .	444
15.9	Recovering your entire system . . . . .	446
	<b>Appendix A. Troubleshooting . . . . .</b>	<b>449</b>
	Collection Services. . . . .	450
	Starting a performance collection . . . . .	450
	Checking the status of Performance Monitor . . . . .	456
	Management of Collection Services . . . . .	458

Performance Tools reports .....	459
Performance Explorer .....	471
DS8000 troubleshooting .....	471
<b>Appendix B. Installing the storage unit activation key using a DS GUI .....</b>	<b>473</b>
<b>Related publications .....</b>	<b>479</b>
IBM Redbooks publications .....	479
Other publications .....	479
Online resources .....	480
How to get IBM Redbooks publications .....	481
Help from IBM .....	481
<b>Index .....</b>	<b>483</b>

# Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

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

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

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

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

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

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


## COPYRIGHT LICENSE:

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

# Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

AIX®	IBM®	System i®
AS/400®	iSeries®	System p®
DataMirror®	Lotus®	System Storage™
DB2®	NetServer™	System Storage DS®
Domino®	OS/400®	System x™
DS6000™	PartnerWorld®	System z®
DS8000™	POWER5™	System/38™
Enterprise Storage Server®	POWER6™	z/OS®
eServer™	Redbooks®	zSeries®
FlashCopy®	Redbooks (logo)  ®	
i5/OS®	System i5®	

The following terms are trademarks of other companies:

Disk Magic, IntelliMagic, and the IntelliMagic logo are trademarks of IntelliMagic BV in the United States, other countries, or both.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Java, JRE, JVM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Pentium, Pentium 4, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

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

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

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

# Preface

This IBM® Redbooks® publication describes the implementation of IBM System Storage™ Copy Services with the IBM System i® platform using the IBM System Storage Disk Storage family and the Storage Management GUI and command-line interface. This book provides examples to create an IBM FlashCopy® environment that you can use for offline backup or testing. This book also provides examples to set up the following Copy Services products for disaster recovery:

- ▶ Metro Mirror
- ▶ Global Mirror

The newest release of this book accounts for the following new functions of IBM System i POWER6™, i5/OS® V6R1, and IBM System Storage DS8000™ Release 3:

- ▶ System i POWER6 IOP-less Fibre Channel
- ▶ i5/OS V6R1 multipath load source support
- ▶ i5/OS V6R1 quiesce for Copy Services
- ▶ i5/OS V6R1 High Availability Solutions Manager
- ▶ System i HMC V7
- ▶ DS8000 R3 space efficient FlashCopy
- ▶ DS8000 R3 storage pool striping
- ▶ DS8000 R3 System Storage Productivity Center
- ▶ DS8000 R3 Storage Manager GUI

## The team that wrote this book

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

**Nick Harris** is a Consulting IT Specialist for IBM System i. He spent the last nine years at the ITSO Rochester Center. He specializes in IBM System i and eServer™ iSeries® hardware, IBM i5/OS and IBM OS/400® software, logical partition (LPAR), high availability, external disk, Microsoft® Windows® integration, and Linux®. He writes IBM Redbooks publications and conducts classes at ITSO technical forums worldwide on all these subjects and how they are related to system design and server consolidation. Previously, Nick spent 13 years in the U.K. IBM AS/400® Business, where he worked with S/36, S/38, AS/400, and iSeries servers. You can contact him by sending e-mail to [niharris@us.ibm.com](mailto:niharris@us.ibm.com).

**Hernando Bedoya** is an IT Specialist at the ITSO Rochester Center. He writes extensively and teaches IBM classes worldwide in all areas of DB2® for i5/OS. Before joining the ITSO more than seven years ago, he worked for IBM Colombia as an IBM AS/400 IT Specialist doing pre-sales support for the Andean countries. He has 24 years of experience in the computing field and has taught database classes in Colombian universities. He holds a masters degree in computer science from EAFIT, Colombia. His areas of expertise are database technology, application development, and data warehousing. You can contact him by sending e-mail to [nbedoya@us.ibm.com](mailto:nbedoya@us.ibm.com).

**Ron Devroy** is a Software Support Specialist working in the IBM Rochester Support Center as a member of the Performance team. He also works as a virtual member of the external storage support team and is considered the Subject Matter Expert for external storage with his team. You can contact him by sending e-mail to [devroy@us.ibm.com](mailto:devroy@us.ibm.com).

**Ingo Dimmer** is an IBM Advisory IT Specialist for System i and a PMI Project Management Professional working in the IBM STG Europe storage support organization in Mainz, Germany. He has eight years of experience in enterprise storage support from working in IBM post-sales and pre-sales support. He holds a degree in Electrical Engineering from the Gerhard-Mercator University, Duisburg. His areas of expertise include System i external disk storage solutions, I/O performance, and tape encryption for which he has been an author of several whitepapers and IBM Redbooks publications. You can contact him by sending e-mail to [idimmer@de.ibm.com](mailto:idimmer@de.ibm.com).

**Adrian Froon** is a member of the IBM Custom Technology Center based in EMEA. He specializes in the design and implementation of external storage solutions, with an emphasis on the Copy Services Toolkit installation (FlashCopy and Metro Mirror). Adrian is also a key member of the Benchmark testing team that works on external storage for IBM European customers. You can contact him by sending e-mail to [Adrian.Froon@fr.ibm.com](mailto:Adrian.Froon@fr.ibm.com).

**Dave Lacey** is the lead technical specialist for System i5® and iSeries in Ireland and the Team Lead for the iSeries group in Ireland. Dave has also worked in IBM Service Delivery, Business Continuity and Recovery Services, and Service Delivery throughout Ireland. He has developed and taught courses on LPAR, IBM POWER5™, and Hardware Management Console. Dave has worked on the AS/400 since its launch in 1988. Prior to that, he was a software engineer for the IBM S/36 and S/38. You can contact him by sending e-mail to [dave\\_lacey@ie.ibm.com](mailto:dave_lacey@ie.ibm.com).

**Veerendra Para** is an advisory IT Specialist for System i in IBM Bangalore, India. His job responsibility includes planning, implementation, and support for all the iSeries platforms. He has nine years of experience in IT field. He has over six years of experience in AS/400 installations, networking, transition management, problem determination and resolution, and implementations at customer sites. He has worked for IBM Global Services and IBM SWG. He holds a diploma in Electronics and Communications. You can contact him by sending e-mail to [veerendra.para@in.ibm.com](mailto:veerendra.para@in.ibm.com) or [veepara2@in.ibm.com](mailto:veepara2@in.ibm.com).

**Will Smith** is the Team Leader of System i with IBM System Storage DS8000 performance group, based in Tucson for the past two years. Will has written two whitepapers on System i and DS8000 performance. The first whitepaper covers CPW and Save/Restore measurements. The second whitepaper covers PPRC Metro Mirror measurements in a System i environment. Will has experience in System i hardware and LPAR, software, and performance metrics. He is also an expert in DS8000 command structure, setup, and performance. You can contact him by sending e-mail to [wsmith@us.ibm.com](mailto:wsmith@us.ibm.com).

**Herbert Velasquez** works for GBM, an IBM Business Partner in Latin America. He has worked with the AS/400, iSeries, and now System i5 for 20 years. Herbert began as a Customer Engineer. He now works in a regional support role as a System Engineer who is responsible for designing and implementing solutions that involve LPAR and external storage for GBM customers. You can contact him by sending e-mail to [hvelasqu@gbm.net](mailto:hvelasqu@gbm.net).

**Ario Wicaksono** is an IT Specialist for System i at IBM Indonesia. He has two years of experience in Global Technology Services as System i support. His areas of expertise are System i hardware and software, external storage for System i, Hardware Management Console, and LPAR. He holds a degree in Electrical Engineering from the University of Indonesia. You can contact him by sending e-mail to [wario@id.ibm.com](mailto:wario@id.ibm.com).

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

Ginny McCright  
Jana Jamsek  
Mike Petrich  
Curt Schemmel  
Clark Anderson  
Joe Writz  
Scott Helt  
Jeff Palm  
Henry May  
Tom Crowley  
Andy Kulich  
Jim Lembke  
Lee La Frese  
Kevin Gibble  
Diane E Olson  
Jenny Dervin  
Adam Aslakson  
Steven Finnes  
Selwyn Dickey  
John Stroh  
Tim Klubertanz  
Dave Owen  
Scott Maxson  
Dawn May  
Sergrey Zhiganov  
Gerhard Pieper  
IBM Rochester Development Lab

Also thanks to the following people that shared written material from *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788:

Jana Jamsek  
Bertrand DufRASne  
International Support Center Organization, San Jose, California

## Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

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

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

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

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

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

[ibm.com/redbooks](http://ibm.com/redbooks)

- ▶ Send your comments in an e-mail to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

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



## Introduction

This book is divided into multiple sections. This part introduces Copy Services for System i and high availability concepts on System i. It also covers the different external storage solutions on System i.

This part includes the following chapters:

- ▶ Chapter 1, “Introduction to Copy Services and System i high availability” on page 3
- ▶ Chapter 2, “System i external storage solution examples” on page 31

Archived

# Introduction to Copy Services and System i high availability

As businesses become more conscious of the availability of IT systems, there is a move to bring IT procedures from tape backup recovery (cold site) to a disaster-recovery level or a high-availability (HA) level. The following components of the total availability time can be improved:

- ▶ Reduce the backup window using the IBM System Storage function of FlashCopy
- ▶ Reduce recovery time through IBM System Storage Metro Mirror or Global Mirror with boot from SAN
- ▶ Implement high availability through IBM System Storage Metro Mirror or Global Mirror with independent ASPs

With the introduction of IBM System Storage Copy Services in the IBM System i environment, an important hardware-based replication solution has been added to the possibilities to achieve a higher level of Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the System i platform. However, this solution does not remove the necessity for proper tape backups of IT systems and journaling within applications.

## 1.1 Overview of the i5/OS architecture

Several architectural features of the i5/OS distinguish it from other systems in the computing industry including:

- ▶ Two-part primary operating system
- ▶ Technology-independent machine interface (TIMI)
- ▶ Object-based system
- ▶ Single-level storage
- ▶ High degree of integration
- ▶ Multiple application program models
- ▶ High level of security
- ▶ Open standards

### 1.1.1 Two-part primary operating system

There are two components to the operating system software on a System i5 model— System Licensed Internal Code (SLIC) and i5/OS. This important distinction is unique in the industry in its completeness of implementation.

SLIC provides the TIMI, process control, resource management, integrated SQL database, security enforcement, network communications, file systems, storage management, JVM™, and other primitives. SLIC is a hardened, high-performance layer of software at the lowest level, similar to a UNIX® kernel, only far more functional.

i5/OS provides higher-level functions based on these services to users and applications. It also provides a vast range of high-level language (such as C/C++, COBOL, RPG, and FORTRAN) runtime functions. i5/OS interacts with the client-server graphical user interface (GUI), iSeries Navigator, or its new i5/OS V6R1 Web-based successor product called *IBM Systems Director Navigator for i5/OS*.

At a macro level, an entire logical partition (LPAR) running the traditional System i operating system can be referred to as *running i5/OS*. The name i5/OS can refer to either the combination of both parts of the operating system or more precisely just the “top” portion.

### 1.1.2 Object-based system

i5/OS keeps all information as *objects*. There are hundreds of object types, such as physical files, program objects, device descriptions, message queues, user profiles, and so forth. This object-based system is different from the simple byte-string, file-based manipulation used by many systems. Object-based design enables a powerful, yet manageable level of system integrity, reliability, and authorization constraints.

All programs and operating system information, such as user profiles, database files, programs, printer queues, and so on, have their associated object types stored with the information. In the i5/OS architecture, the *object type* determines how the contained information of the object can be used (which methods). For example, it is impossible to corrupt a program object by modifying its code sequence data as though it were a file. Because the system knows the object is a program, it only allows valid program operations (run and backup). Thus, with no write method, i5/OS program objects are, by design, highly virus resistant. Other kinds of objects include directories and simple stream data files residing in the Integrated File System (IFS), such as video and audio files. These stream-file objects provide familiar open, read, and write operations.

### 1.1.3 Single-level storage

i5/OS applications and the objects with which they interact all reside in large virtualized, single-level storage. That is, the entire system, including the objects that most other systems distinguish as “on disk” or “in memory” are all in single-level storage. Objects are designated as either permanent or temporary. *Permanent objects* exist across system IPLs (reboots). *Temporary objects* do not require such persistence. Essentially, the physical RAM on the server is a cache for this large, single-level storage space. Storage management, a component of SLIC, ensures that the objects that need to persist when the system is off are maintained in persistent storage. This is either magnetic hard disk or flash memory.

The benefit of providing a single, large address space, in which all objects on the system reside, is that applications do not need to tailor their memory usage to a specific machine configuration. In fact, due to single-level storage, i5/OS does not need to tailor such things as the sizes of disk cache versus paging space. This greatly facilitates the on-demand allocation of memory among LPARs.

This concept is an important concept when considering the advanced functions that are available with IBM System Storage Copy Services. When using storage-based replication due to the System i single-level storage architecture the granularity for replicating System i storage is either replicating all system ASP and user ASPs, sometimes referred to \*SYSBAS, together or replicating on an independent ASP (IASP) level. If you are planning to use IBM System Storage FlashCopy, all objects that exist in System i main memory must be purged to storage for creating a consistent image for either a normal IPL or normal IASP vary-on. This purging can only be achieved with either turning off the system or varying off the IASP before taking a FlashCopy. However, the new i5/OS V6R1 quiesce for Copy Services function eliminates the requirement to power down the system or vary off the IASP before taking a FlashCopy. This new quiesce function allows you to suspend the database I/O activity in an ASP. It still results in abnormal IPL processing in i5/OS, but because there are no more database inconsistencies, the long-running database recovery tasks to recover damaged objects are not required during the abnormal IPL processing.

## 1.2 Software-based high availability solutions

Software-based solutions offer an extremely different set of possibilities in regard to high availability and business continuity. In the following sections, the generalizations of these software-based high availability solutions are based on the functionality of IBM High Availability Business Partners (HABPs), such as DataMirror® and Vision Solutions.

This section discusses the features and functions that the specific individual business partner solutions have in common when compared to hardware-based replication solutions such as IBM System Storage solutions and cross-site mirroring (XSM) with IASPs.

Software-based high availability solutions are mostly based on i5/OS journaling. With journaling, you set up a *journal* and a *journal receiver*. Then, you define the physical files, data queues, data areas, or integrated file system objects that are to be journaled to this particular journal. Whenever a record is changed, a journal entry is written into the journal receiver, which contains information about the record that was changed, the file to which it belonged, which job changed it, the actual changes, and so forth. Journaling has been around since the IBM System/38™ platform. In fact, a lot of user applications are journaled for various purposes, such as keeping track of user activity against the file or for being able to roll back in case of a user error or a program error.

In our discussion, journaling is classified as *local journaling* or *remote journaling*. The difference is simply the manner in which the data is transferred between systems.

## 1.2.1 Local journaling

High availability solutions based on local journaling have a reader job on the source system that reads the journal entries for the files that are defined to be mirrored and that sends the changes across to the receiver job on the target system. Here, an apply process (job) applies the changes to the target database. The job that is used to transmit the changes from the source system to the target system is not a built-in journaling job, but part of the software-based high availability program or package.

Figure 1-1 shows a high availability solution with local journaling.

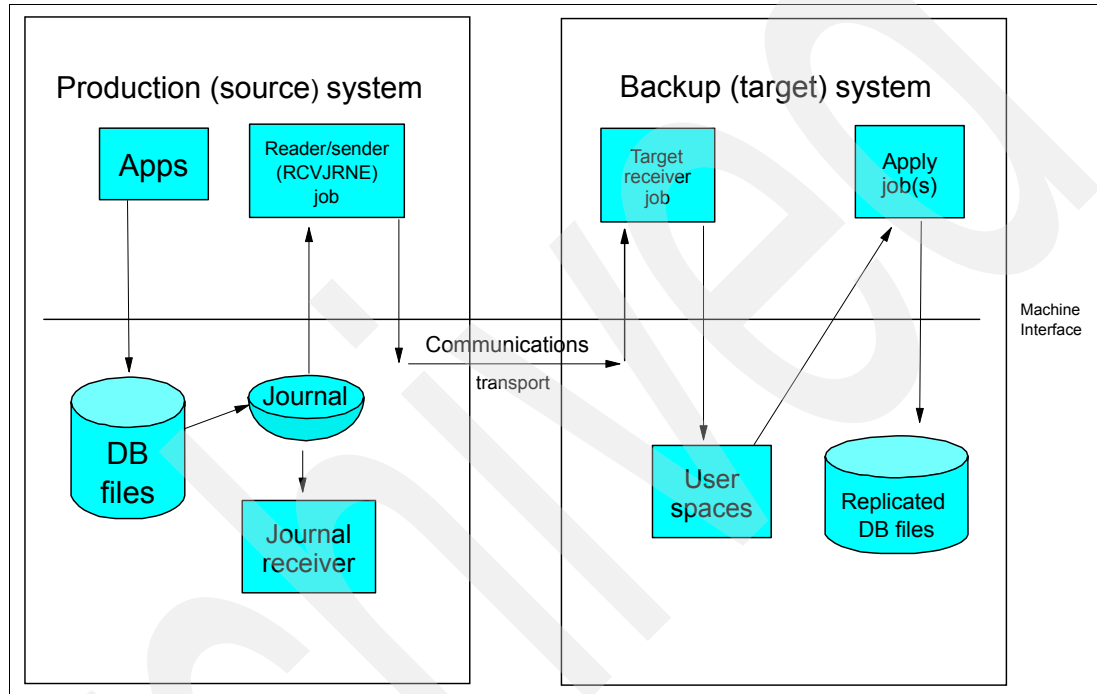


Figure 1-1 Example of a high availability solution with local journaling

## 1.2.2 Remote journaling

Since V4R2M0 of OS/400 (now i5/OS), the concept of remote journaling has enhanced communications between source and target systems. The changes can be sent more quickly from the source system to the target system than what is possible with local journaling and the use of a reader or sender job. Remote journaling is implemented at the Licensed Internal Code (LIC) layer, providing for faster processing between systems.

With remote journaling, you set up local journaling on your source system as you normally would. You then use the Add Remote Journal (ADDRMTJRN) command to associate your local journal to a remote journal, through the use of a relational database. When a transaction is put into the local journal receiver, it is sent immediately to the remote journal and its receiver through the communications path that is designated in the relational database directory entry.

Remote journaling allows you to establish journals and journal receivers on the target system that are associated with specific journals and journal receivers on the source system. After the remote journaling function is activated, the source system continuously replicates journal entries to the target system as described previously.

The remote journaling function is a part of the base OS/400 or i5/OS system and is not a separate product or feature.

Advantages of using remote journaling include:

- ▶ It lowers the processor consumption on the source machine by shifting the processing that is required to read the journal entries from the source system to the target system. Most of the workload is moved to the target system because the reader job that would normally read from the journal on the source system is moved to the target system.
- ▶ It eliminates the need to buffer journal entries to a temporary area before transmitting them from the source machine to the target machine. This translates into fewer disk writes and greater DASD efficiency on the source system.
- ▶ Because it is implemented at the LIC level, it improves significantly the replication performance of journal entries and allows database images to be sent to the target system in real time. This real-time operation is called *synchronous delivery mode*. If synchronous delivery mode is used, the journal entries are guaranteed to be in main storage on the target system prior to control being returned to the application on the source system.
- ▶ It allows the journal receiver save and restore operations to be moved to the target system. This way, the resource utilization on the source machine can be reduced.

For more information about remote journaling, refer to *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189.

Figure 1-2 shows an example of a high availability solution that uses remote journaling with a reader job on the target side.

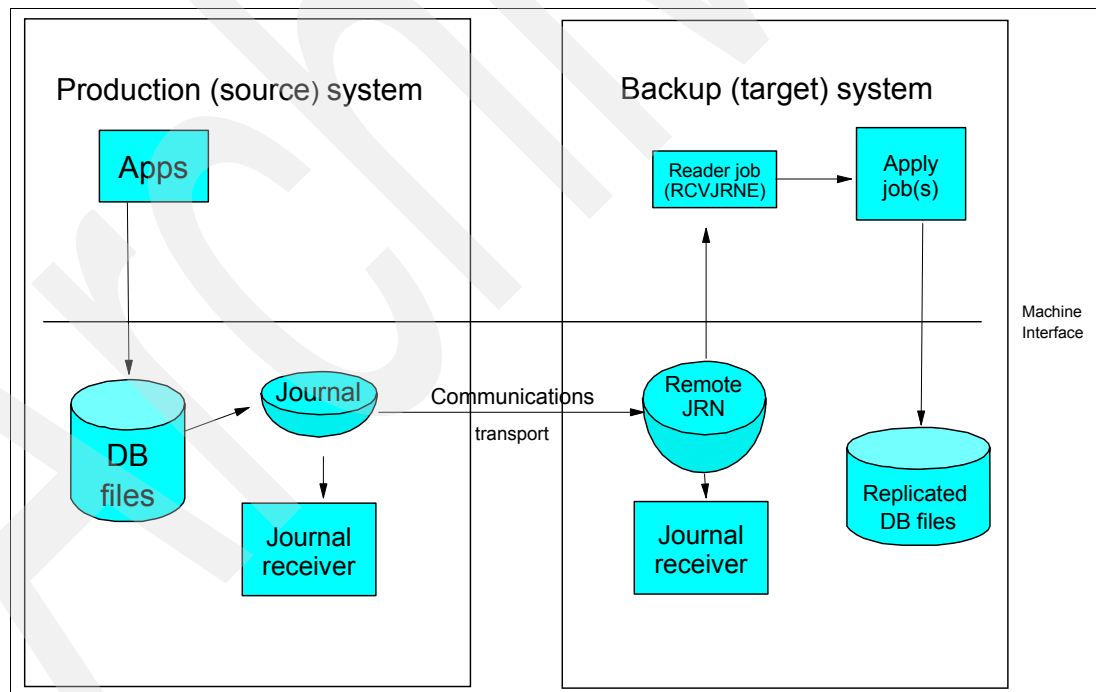


Figure 1-2 Example of a high availability solution with remote journaling

The remote journal function provides a much more efficient transport of journal entries than the traditional approach. In this scenario, when a user application makes changes to a database file, there is no need to buffer the resulting journal entries to a staging area on the production (source) system. Efficient system microcode is used instead to capture and transmit journal entries directly from the source system to the associated journals and journal

receivers on a target system. Much of the processing is done below the Machine Interface (MI). Therefore, more processor cycles are available on the production machine for other important tasks.

### 1.2.3 Object types not journaled

With journaling, whether local or remote, you can replicate changes from the source system to the target system for the following object types:

- ▶ Physical files
- ▶ Access paths
- ▶ Data areas
- ▶ Data queues
- ▶ Integrated file system objects

At V5R4 of i5/OS, only these object types are allowed to be journaled.

**Note:** With i5/OS V6R1, journaling is now supported on a library level to start journaling automatically if one these objects gets newly created in the journaled library.

However, a usable backup system usually requires more than just database and stream files. The backup system must have all of the applications and objects that are required to continue critical business tasks and operations.

Users also need access to the backup system. They need a user profile on the target system with the same attributes as that profile on the source system, and their devices must be able to connect to the target system.

The applications that a business requires for its daily operations dictate the other objects that are required on the backup system. Not all of the applications that are used during normal operations might be required on the backup system. In the event of an unplanned outage, the business can choose to run with a subset of those applications, which might allow the business to use a smaller system as the backup system or to reduce the impact of the additional users when the backup system is already used for other purposes.

The exact objects that comprise an application vary widely. Some of the object types that are commonly part of an application include:

- ▶ Authorization lists (\*AUTL)
- ▶ Job descriptions (\*JOBDD)
- ▶ Libraries (\*LIB)
- ▶ Programs (\*PGM)
- ▶ User spaces (\*USRSPC)

For many of the objects in the list, the content, attributes, and security of the object affect how the application operates. The objects must be continuously synchronized between the production and backup systems. For some objects, replicating the object content in near real time can be as important as replicating the database entries.

### 1.2.4 Replicating non-journaled object types

Most of the HABPs have solutions for mirroring non-journaled objects of the types listed in the previous section. This replication is typically done by using the system audit journal (QAUDJRN), which is configured to monitor for creations, deletions, and other object-related events. The HABP solution reads from the QAUDJRN journal. Based on its list of objects



defined to be mirrored, it sends the whole object to the target system using temporary save files or with ObjectConnect/400 (included in the operating system as option 22), if configured on the systems.

Figure 1-3 shows a general view of replication for non-journaled objects.

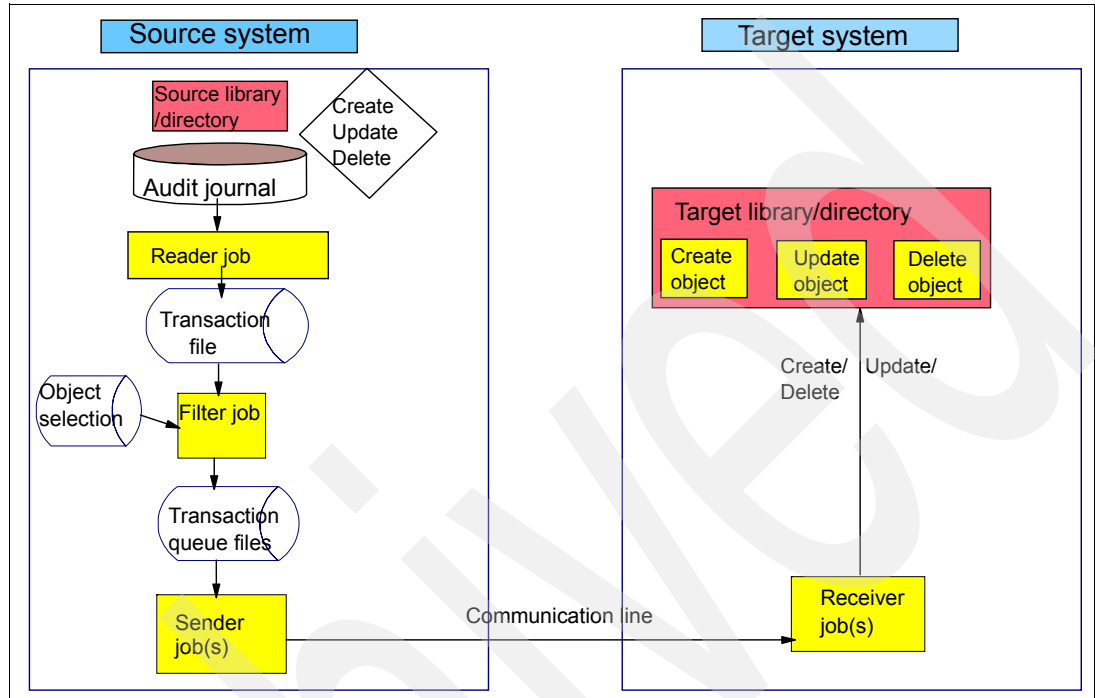


Figure 1-3 General view of mirroring non-database objects from source to target

## 1.3 i5/OS clustering

In this section, we describe briefly clustering, its basic components, and its concepts. We provide the basic elements that are required before you can configure IASPs, XSM, or geographic mirror.

For a complete discussion about clustering and how to set up a cluster, refer to *Clustering and IASPs for Higher Availability on the IBM eServer iSeries Server*, SG24-5194.

### 1.3.1 Definition of a cluster

A *cluster* is a collection of interconnected complete computers, or *nodes*, that appears on a network as a single machine. The cluster is managed as a single system or operating entity. It is designed specifically to tolerate component failures and to support the addition or subtraction of components in a way that is transparent to users.

The main purpose of clustering is to achieve *high availability*. High availability allows important production data and applications to be available during periods of planned system outages.

Clustering can also be used for *disaster recovery* implementations. Disaster recovery typically refers to ensuring that the same important production data and applications are available in the event of an unplanned system outage, caused many times by natural disasters.

Clustering becomes an important concept for both high availability and disaster recovery discussions.

*Cluster Resource Services*, a component of i5/OS, provides the following features:

- ▶ Tools to create and manage clusters, the ability to detect a failure within a cluster, and switchover and failover mechanisms to move work between cluster nodes for planned or unplanned outages.
- ▶ A common method for setting up object replication for nodes within a cluster.  
This includes the data objects and program objects necessary to run applications that are cluster-enabled.
- ▶ Mechanisms to switch automatically applications and users from a primary node to a backup node within a cluster for planned or unplanned outages.
- ▶ Heartbeat monitoring that uses a low-level message function to constantly ascertain that every node can communicate with other nodes in the cluster.

If a node fails or a break occurs in the network, heartbeat monitoring tries to re-establish communications. If communications cannot be re-established within a designated time, heartbeat monitoring reports the failure to the rest of the nodes within the cluster.

### 1.3.2 Cluster components

A cluster is made up of the following components:

- ▶ Cluster node
  - Primary node
  - Backup node
  - Replicate node
- ▶ Cluster resource group
  - Data resilient CRG (type-1)
  - Application resilient CRG (type-2)
  - Device resilient CRG (type-3)
  - Peer CRG
- ▶ Cluster resource services
- ▶ Cluster version
- ▶ Device domain
- ▶ Administrative domain
- ▶ Resilient resources
- ▶ Cluster management support and clients

Figure 1-4 shows the components of a cluster.

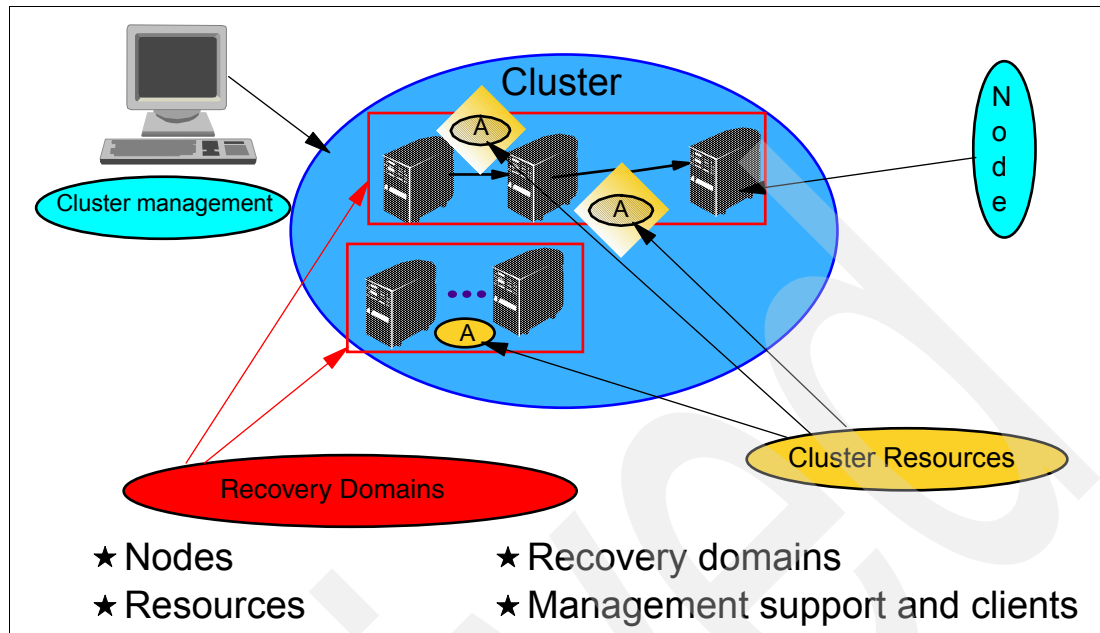


Figure 1-4 Cluster components

A *cluster node* is any System i model or partition that is a member of a cluster. Cluster communications that run over IP connections provide the communications path between cluster services on each node in the cluster. A cluster node can operate in one or more of the following roles:

- ▶ A *primary node*, which is the cluster node that is the primary point of access for cluster resources
- ▶ A *backup node*, which is a cluster node that can assume the primary role if the primary node fails or a manual switchover is initiated
- ▶ A *replicate node*, which is a cluster node that maintains copies of the cluster resources but is unable to assume the role of primary or backup

A cluster resource group (CRG) is an i5/OS external system object that is a set or group of cluster resources. The cluster resource group describes a *recovery domain*, a subset of cluster nodes that are grouped together in the CRG for a common purpose such as performing a recovery action or synchronizing events, and supplies the name of the cluster resource group exit program that manages cluster-related events for that group. One such event is moving users from one node to another node in case of a failure. Cluster resource group objects are defined either as data resilient, application resilient, or device resilient:

- ▶ A *data resilient CRG (type-1)* allows multiple copies of data to be maintained on more than one node in a cluster.
- ▶ An *application resilient CRG (type-2)* allows an application (program) to run on any of the nodes in a cluster.
- ▶ A *device resilient CRG (type-3)* allows a hardware resource to be switched between systems. The device CRG contains a list of device configuration objects used for clustering. Each object represents an IASP.

A *peer CRG*, which was newly introduced with i5/OS V5R4, defines nodes in the recovery domain with peer roles. It is used to represent the cluster administrative domain. It contains

monitored resource entries, for example user profiles, network attributes or system values that can be synchronized between the nodes in the CRG.

Cluster Resource Services is the set of OS/400 or i5/OS system service functions that support System i5 cluster implementations.

The cluster version identifies the communication level of the nodes in the cluster.

A *device domain* is a subset of cluster nodes across which a set of resilient devices, such as an IASP, can be shared. The sharing is not concurrent for each node, which means that only one node can use the resilient resource at one time. Through the configuration of the primary node, the secondary node is made aware of the individual hardware within the CRG and is “ready to receive the CRG” should the resilient resource be switched. A function of the device domain is to prevent conflicts that cause the failure of an attempt to switch a resilient device between systems.

Figure 1-5 shows a device domain with a primary node and a secondary node, as well as a switchable device (an IASP) that can be switched from Node 1 to Node 2.

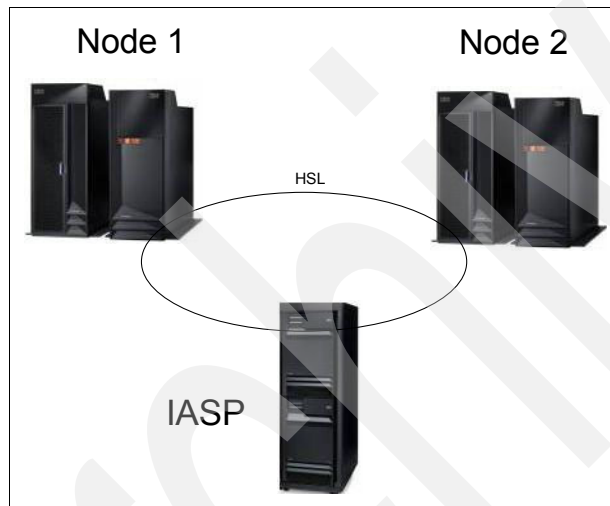


Figure 1-5 Example of a device domain

A *resilient resource* is a device, data, or an application that can be recovered if a node in the cluster fails.

*Resilient data* is data that is replicated, or copied, on more than one node in a cluster.

*Resilient applications* are applications that can be restarted on a different cluster node without requiring the clients to be reconfigured.

*Resilient devices* are physical resources, represented by a configuration object, such as a device description, that are accessible from more than one node in a cluster through the use of switched disk technology and independent disk pools.

### Cluster management support and clients

IBM provides a cluster management GUI that is accessible through iSeries Navigator or with i5/OS V6R1 through the IBM Systems Director Navigator for i5/OS and available through i5/OS option 41 (HA Switchable Resources). The utility allows you to create and manage a cluster that uses switchable IASPs and to ensure data availability. The cluster management GUI features a wizard that takes you through the creation of the cluster and all of its components (see Figure 1-6).

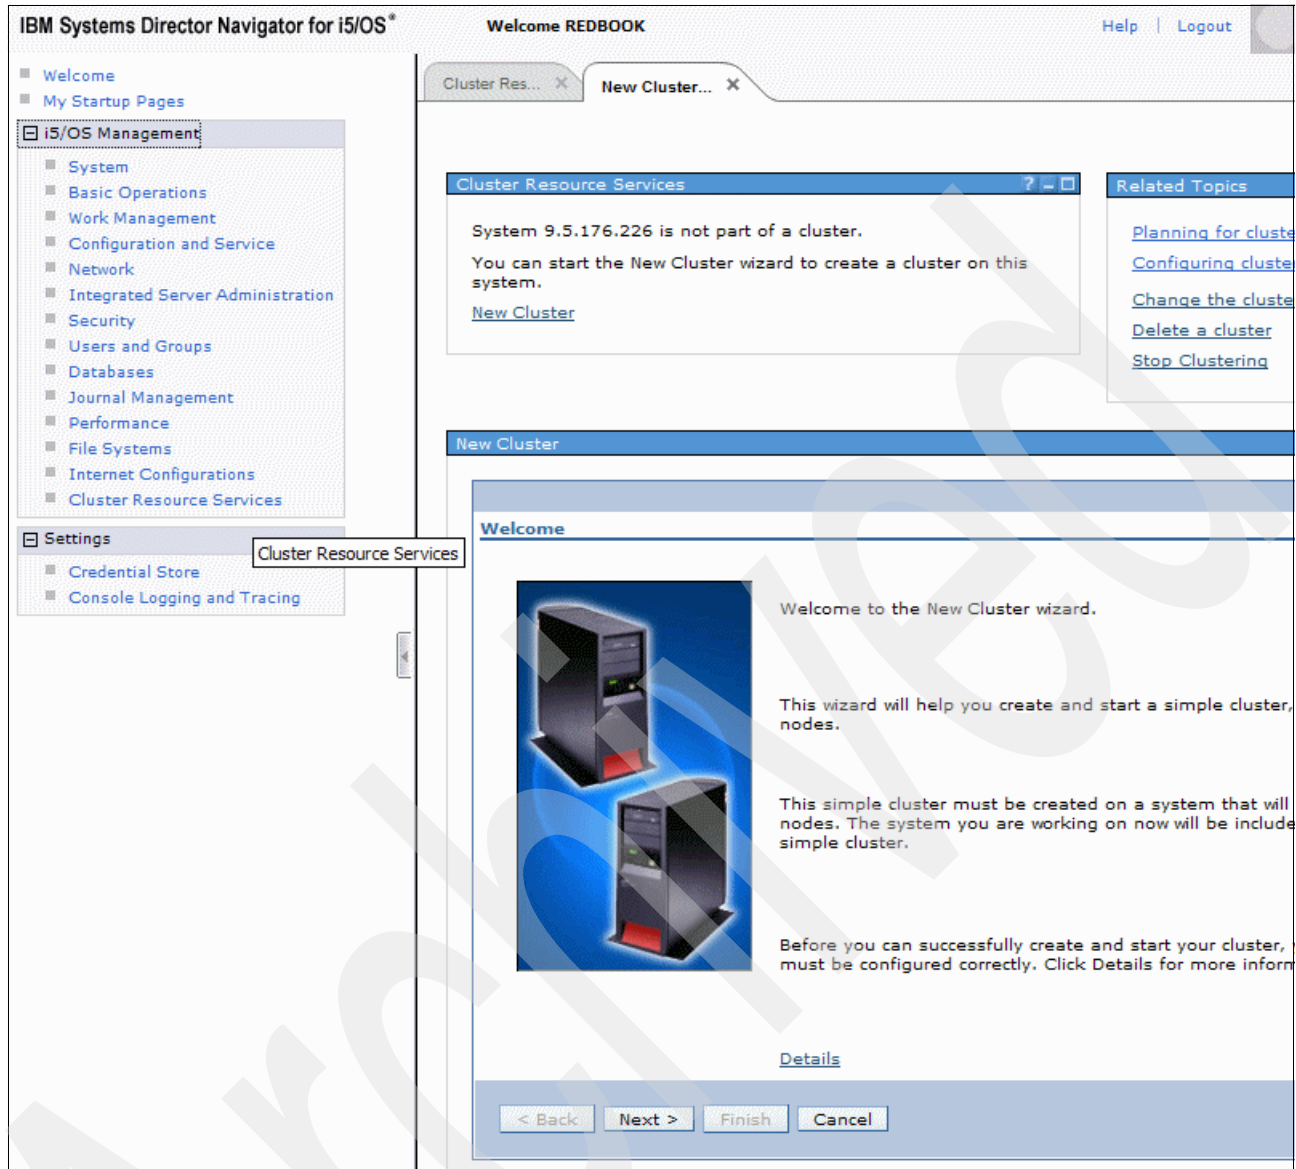


Figure 1-6 IBM Systems Director Navigator for i5/OS Cluster Resource Services

## 1.4 Auxiliary storage pools

In this section, we introduce the types of auxiliary storage pools and explain how they relate to clustering and i5/OS cross-site mirroring (XSM).

### 1.4.1 Definition of an auxiliary storage pool

*Auxiliary storage pools* (ASPs) have existed since the announcement of the AS/400 in 1988. These ASPs allow you to divide the total disk storage on the system into logical groups, or *disk pools*, in order to limit the impact of storage-device failures and to reduce recovery time. You can then isolate one or more applications or data in one or more ASPs, for various reasons related to backup and recovery, performance, or other purposes.

ASPs include the system ASP and user ASPs. The system ASP contains SLIC and i5/OS code. There is only one per system or partition, and it is always numbered 1.

User ASPs are any other ASPs defined on the system, other than the system ASP. Basic user ASPs are numbered 2 through 32. Independent user ASPs (IASPs) are numbered 33 through 255. Data in a basic user ASP is always accessible whenever the server is up and running.

## 1.4.2 Definition of an independent user ASP

Independent user ASPs (IASPs) are a type of user ASP, numbered 33 through 255. The system assigns the IASP number, where the user can choose the number for a basic ASP. IASPs are different from basic ASPs in several ways.

Independent ASPs are described in i5/OS with a device description (DEVD) and are identified by a device name. They can be used on a single system or switched between multiple systems or LPARs when the IASP is associated with a switchable hardware group, in clustering terminology known as a *device CRG*. When used on a single system, the IASP can be dynamically varied on or off without restarting the system which saves a lot of time and increases the flexibility offered by ASPs. In iSeries Navigator or its V6R1 Web-based version the IBM Systems Director Navigator for i5/OS, the IASP and its contents can be dynamically made available or unavailable to the system.

When used across multiple systems, clustering support with i5/OS option 41 (HA switchable resources) is required between the systems, and the cluster management GUI (see “Cluster management support and clients” on page 12) is used to switch the IASP across systems in the cluster. This is referred to as a *switchable IASP*. At any given time, the IASP can be used by only one of those systems. Multiple systems cannot simultaneously use the IASP.

The new i5/OS V6R1 disk encryption feature using the i5/OS option 45 (Encrypted ASP Enablement) allows to encrypt data on an ASP or IASP.

**Important:** When using disk encryption for switchable IASPs the master key needs to be set manually on each system in the device domain, and all systems need option 45 installed in order to vary on the IASP.

## 1.5 Cross-site mirroring concepts

In this section, we describe the general relationship between the XSM functions of clustering and auxiliary storage pools.

### 1.5.1 Definition of cross-site mirroring

*Cross-site mirroring* (XSM) is part of OS/400 or i5/OS Option 41 High Availability Switchable Resources. It provides the following features:

- ▶ Data resilience
  - Mirroring of an ASP group occurs from one location to a second location.
  - Switchover or automatic failover to the secondary copy happens in the event of an outage at the primary location.
- ▶ Extended capabilities for basic switchable IASPs
  - Addresses single point of failure.
  - Provides the possibility of multiple data copies.
  - Alleviates switchable tower connectivity restrictions.

- ▶ Site data resiliency protection in addition to high availability
  - The second copy of the IASP is kept at another “site.”
  - The other site can be geographically remote.
- ▶ Provides additional backup nodes for resilient data
  - Both copies of IASP can be stored in switchable devices.
  - Each copy can be switched between nodes locally.

XSM provides the ability to replicate changes made to the production copy of an IASP to a mirror copy of that IASP. As data is written to the production copy of an IASP, the operating system mirrors that data to a second copy of the IASP through another system. This process keeps multiple identical copies of the data.

Changes written to the production copy on the source system are guaranteed to be made in the same order to the mirror copy on the target system. If the production copy of the IASP fails or is shut down, you have a hot backup, in which case the mirror copy becomes the production copy.

The IASP used in XSM has the benefits of any other IASP, with its ability to be made available or unavailable (varied on or off), and you have greater flexibility for the following reasons:

- ▶ You can protect the production IASP and mirror IASP with the protection that you prefer, either disk unit mirroring or device parity protection (RAID-5 or RAID-6). Moreover, the production IASP and the mirror IASP are not required to have the same type of protection. While no protection is required for either IASP, we highly recommend some type of protection for most scenarios.
- ▶ You can set the threshold of the IASP to warn you when storage space is running low. The server sends a message, allowing you time to add more storage space or to delete unnecessary objects. Be aware that if the user ignores the warning and the production IASP becomes full, the application stops and objects cannot be created. With IASPs, there is no overflow of data into the system disk pool as opposed to using basic user ASPs.
- ▶ The mirror copy can be detached and then separately be made available to perform save operations, to create reports, or to perform data mining. However, when the mirror copy is reattached, prior to i5/OS V5R4 a full re-synchronization with the production copy is done and all modifications made to the detached copy are lost.

**Note:** The new XSM target site tracking function in i5/OS V6R1, available also for V5R4 using PTF MF40053, allows for partial synchronization from the source to the target site after a mirrored IASP copy is re-attached using the tracking option. In this case only pages that changed on the source or target site are sent to the mirrored copy at the target site. In contrast to V6R1 the V5R4 PTF still has the limitation that the detach with tracking must be done while the production IASP is offline.

The V5R4 source site tracking function allows for partial synchronization due to link communication problems only and does not cover the case for detaching a mirrored copy.

- ▶ If you configure the IASPs to be switchable, you increase your options to have more backup nodes that allow for failover and switchover methods.

## 1.5.2 Definition of geographic mirroring

Geographic mirroring has been made available in i5/OS V5R3. It is currently the only sub-function of XSM. The two terms are not interchangeable, however. *Geographic mirroring* specifically refers to System i server-based replication of IASP data on memory page level. XSM is a concept that describes replication of data at multiple sites.

Geographic mirroring is intended for use by clustered system environments and uses data port services. Data port services is Licensed Internal Code (LIC) that supports the transfer of large volumes of data between a source system and one of any specified target systems. This is a transport mechanism that communicates over TCP/IP. It provides both synchronous and asynchronous send modes. Be aware of the fact that, even in asynchronous mode, a local write waits for the data to reach main storage of the backup node before the write operation is considered complete.

While geographic mirroring is actively performed, users cannot access the mirror copy of the data.

Figure 1-7 and Figure 1-8 show a simple geographically mirrored IASP and an environment that also incorporates switchable IASPs at both sites.

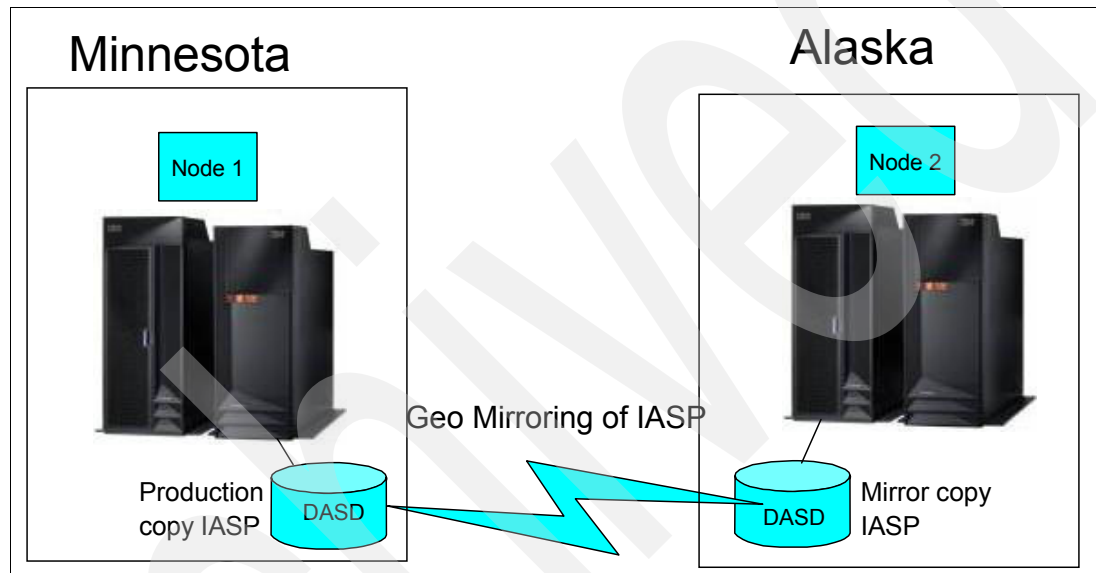


Figure 1-7 Example of geographic mirroring

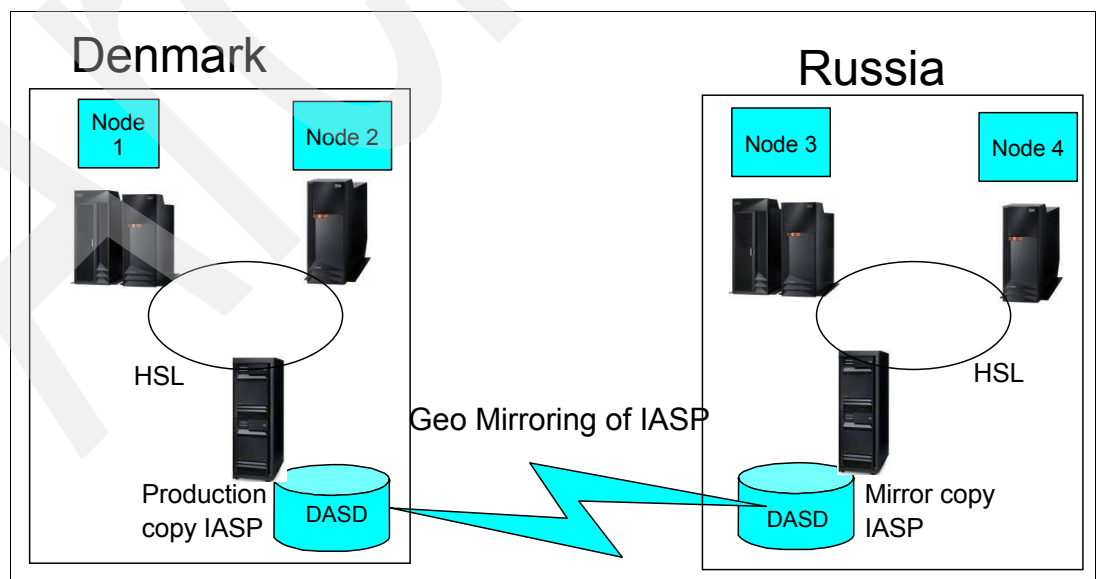


Figure 1-8 Example of geographic mirroring and switched IASPs



### 1.5.3 Failover and switchover

Two important concepts that are related to clustering and XSM are failover and switchover capabilities from the source system to the target system:

- ▶ A *failover* means that the source or primary system has failed and that the target or secondary system takes over. This term is used in reference to unplanned outages.
- ▶ A *switchover* is user-initiated and the user can perform a switchover if the primary system has to be shut down for maintenance, for example. In this case, production work is switched over to the target system (backup node), which takes over the role as the primary node.

### 1.5.4 Supported and unsupported i5/OS object types

Before you decide to base your high availability setup on XSM, you need to consider the object types that OS/400 or i5/OS allows you to put into an IASP. The list of supported object types changes with each new release. Therefore, you should review the i5/OS information center for your particular version of i5/OS to see the list of supported objects:

- ▶ V6R1

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaly/rzalysupportedunsupportedobjects.htm>

**Note:** New with i5/OS V6R1 is the support for JOBQ objects in IASPs which allow applications to be ported to IASPs with fewer changes. However the jobs in the JOBQs will not survive an IASP vary off/on so they will not be available when switching the IASP to a backup system.

- ▶ V5R4

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r4/topic/rzaly/rzalysupportedunsupportedobjects.htm>

- ▶ V5R3

<http://publib.boulder.ibm.com/infocenter/iseriess/v5r3/topic/rzaly/rzalysupportedunsupportedobjects.htm>

### 1.5.5 Benefits of cross-site mirroring

XSM offers the following benefits:

- ▶ XSM provides site disaster protection by keeping a copy of the IASP at another site, which can be geographically distant, by using the geographic mirror function. Having an additional copy at another remote site improves availability.
- ▶ XSM can provide several backup nodes. In addition to having a production copy and a mirrored copy, backup node possibilities are expanded when the IASP is configured as switchable in an expansion unit, on an I/O processor (IOP) on a shared bus, or on an IOP that is assigned to an I/O pool.

## 1.5.6 Limitations of cross-site mirroring

XSM has the following limitations:

- ▶ While XSM is active, you cannot access the mirror copy. This ensures that the data integrity of the mirror copy is maintained.
- ▶ If you detach the mirror copy to perform a save operation, to perform data mining, or to create reports, you must re-attach the mirror copy to resume XSM. With V5R3, this requires a full synchronization with the production copy after it is re-attached. This can be a lengthy process, possibly taking several hours, during which time your production system is unprotected.

Starting with V5R4 and a special PTF and natively with V6R1, you can also use Target Site Tracking, which allows for a partial re-synchronization and can significantly shorten the synchronization times.

- ▶ Not all object types can be mirrored using XSM. You have to maintain important objects, such as user profiles and authorization lists, on both systems by yourself. V5R4 introduced the cluster administrative domain to support this task.
- ▶ XSM can only be performed on objects in an IASP and not on objects in the system ASP or basic user ASPs.

## 1.6 Copy Services based disaster recovery and high availability solutions

With the support for SAN Load Source in System i introduced with i5/OS V5R3M5, it is now possible to have the entire disk space in an IBM System Storage environment. This provides new opportunities that were previously impossible for System i customers. iSeries or System i models that retain their load source drive as an internal physical disk drive unit in the central electronic complex (CEC) or a partition are unable to have the whole system copied and must have a mirrored pair of the load source located on the external storage. The recovery or attachment for disaster recovery or backup purposes with an internal load source disk unit is more complicated and time consuming than having the load source residing in your external disk subsystem with using boot from SAN.

Now you can create a complete copy of your entire system in moments using FlashCopy. You can then use this copy for a variety of purposes such as:

- ▶ Minimize your backup windows
- ▶ Protect yourself from a failure during an upgrade
- ▶ Use it as a fast way to provide yourself with a backup or test system.

You can accomplish all of these tasks by copying the entire direct access storage device (DASD) space with minimal impact to your production operations.

FlashCopy is generally not suitable for disaster recovery because due to its point-in-time copy nature it cannot provide continuous disaster recovery protection nor can it be used to copy data to a second external disk subsystem. To provide an off-site copy for disaster recovery purposes, use either Metro Mirror or Global Mirror depending on the distance between the two external disk subsystems (see 1.6, "Copy Services based disaster recovery and high availability solutions" on page 18).

## 1.6.1 FlashCopy solutions

FlashCopy is the process by which a point-in-time copy of a set of volumes (LUNs) is taken. A relationship is established between the source and target volumes. FlashCopy creates a copy of the source volume on the target volume. The target volume can be accessed as though all the data was copied physically. Unless you are using the new DS8000 Release 3 space efficient FlashCopy virtualization feature (refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120, it requires the same amount of disk storage within the Storage System product as the parent data.

A FlashCopy bitmap is created within DS8000 cache that keeps track of which tracks were already copied to the target and which were not copied. If the data on the original disk track is going to be changed and this track has not been copied to the target yet, to maintain the point-in-time copy state, the original source disk track is copied to the target first before the source track is changed.

Figure 1-9 shows the FlashCopy write I/O processing. The read I/O processing is rather straightforward as reads from the source are processed as though there were no FlashCopy relationship and reads from the target according to the FlashCopy bitmap are either derived from the target if the track has already been copied or are redirected to the source.

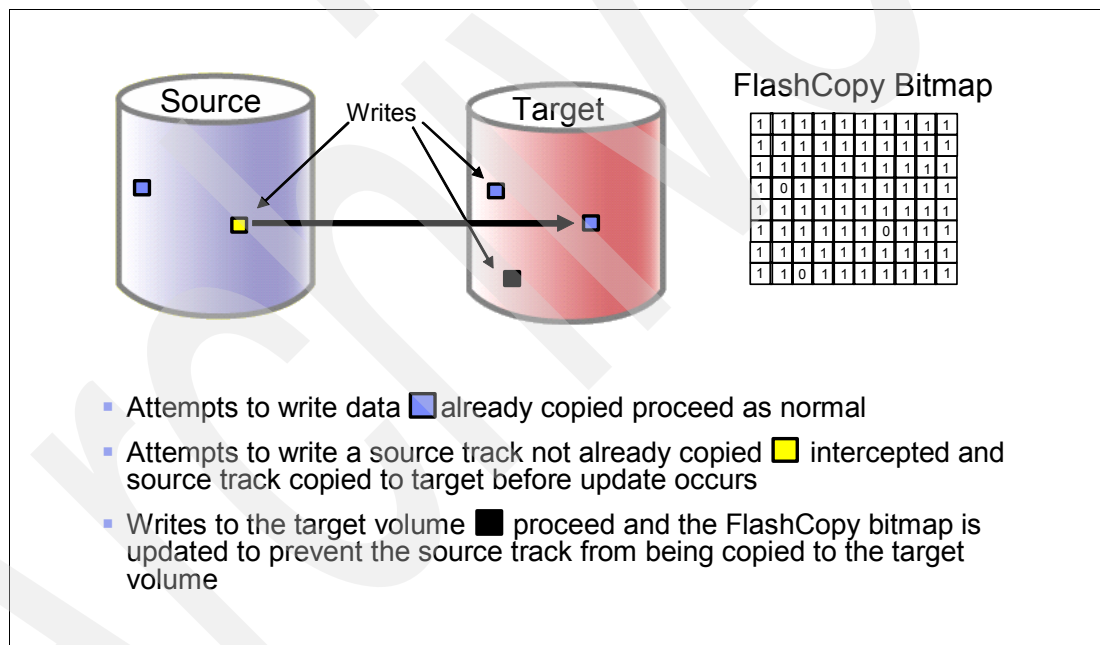


Figure 1-9 FlashCopy write I/O processing

When using the default FlashCopy full-copy option, the tracks are copied from the source to the target volume in the background, and the FlashCopy relationship ends when all tracks have been copied. For short-lived FlashCopy relationships where the source is not changed much over the time of the FlashCopy relationship, we recommend that you use the FlashCopy “no-copy” option, meaning that tracks are only copied to the target if they are going to be changed. The FlashCopy no-copy option is used typically for system backup purposes to limit the performance impact for the production source volumes.

To use FlashCopy, you must purchase the Point-In-Time Copy function authorization feature for your IBM System Storage disk subsystem. FlashCopy is suitable for the following operational environments:

- ▶ **Data backup system**

A FlashCopy of the production data allows the client to create backups with the shortest possible application outage. The main reason for data backup is to provide protection in case of source data loss due to disaster, hardware failure, software failure, or user errors.
- ▶ **Production backup system**

A FlashCopy of the production data allows data recovery from an older level of data. Recovery might be necessary due to a user error or a logical application error. The FlashCopy of the data can also be used by system operations to re-establish production in case of any server errors.
- ▶ **Test system**

Test environments created by FlashCopy can be used by the development team to test new application functions with real production data, thus speeding up the test setup process.
- ▶ **Data mining system**

A FlashCopy of the data can be used for data analysis, thus avoiding performance impacts for the production system due to long running data mining tasks.
- ▶ **Integration system**

New application releases (for example, SAP® releases) are likely to be tested prior to putting them onto a production server. By using FlashCopy, a copy of the production data can be established and used for integration tests. With the capability to reverse a FlashCopy, a previously created FlashCopy can be used within seconds to bring back production to the level of data it had at the time when the FlashCopy was taken.

### System backups using FlashCopy

Creating regular copies of the entire DASD space can be a part of the day-to-day tasks in order to minimize the downtime that is associated with taking backups. With FlashCopy, you can take a copy of the entire DASD space. After you shut down your system or use the new i5/OS V6R1 quiesce for Copy Services function to suspend your database write I/O, the actual FlashCopy relationship is created in milliseconds, after which you can immediately perform an initial program load (IPL) or resume your production system and return it to service while you perform your backup on a second system or partition. This significantly reduces the normal downtime for backup.

**Note:** The new DS8000 Release 3 space efficient FlashCopy virtualization function allowing you to significantly lower the amount of physical storage for the FlashCopy target volumes by thinly provisioning the target space proportional to the amount of write activity from the host fits very well for system backup scenarios with saving to tape.

You can also make a full backup of your system by using standard i5/OS commands with or without Save While Active (SWA). SWA requires the applications to be quiesced to some extent. Sometimes it is faster to go into a restricted state than to wait for the SWA checkpoint to be reached, which can take a considerable amount of time in a system with a complex library structure. When the backup is finished, the user subsystems must be started again.

A *warm FlashCopy* is another recently tested possibility. This method uses a combination of i5/OS independent auxiliary storage pools (IASPs) and FlashCopy. In this case, the system remains active, and only the IASP or the application on the IASP is varied off. This method

also uses the Copy Services Toolkit to automate FlashCopy and the attachment of the copied IASP to another System i server or partition that will perform the backup.

### **Using FlashCopy for system build**

Typically, the copy is used as a test system. Copying the whole system with FlashCopy avoids having to do a lengthy restore from tape.

With the ability to create a complete copy of the whole environment, you have a copy on disk that can be attached to a system or partition and you can perform an IPL normally. For example, if you have planned a release upgrade over a weekend, you can now create a clone of the entire environment on the same disk subsystem using FlashCopy immediately after doing the system shutdown and perform the upgrade on the original copy. If problems or delays occur, you can continue with the upgrade until just prior to the time that the service needs to be available for the users. If the maintenance is not completed, you can abort the maintenance and reattach the target copy representing the original state before the upgrade. Alternatively, you can do a FlashCopy *fast reverse restore* from the original production copy on the target volumes back to your production source LUNs and do a normal IPL, rather than having to do a full system restore.

Cloning a system can save a lot of time, not only for total system backups in connection with hardware or software upgrades, but also for other things such as creating a new test environment.

### **FlashCopy and single-level storage**

In the case of FlashCopy, you can avoid object damage completely only by turning off the System i server or partition or varying off the IASP. Running FlashCopy is a fast task (taking only a few milliseconds). The IPL processing of an i5/OS instance can also be fast (taking 15 minutes), but you must consider the ending and restarting of the application. Application end and restart can be relatively quick (5 to 10 minutes), but when added together, this time is often too long for a 24x7 operation.

With the new i5/OS V6R1 quiesce for Copy Services function, you can eliminate damage for database objects by suspending the database I/O activity for either \*SYSBAS or an IASP. Using this new function system, shutting down or varying off the IASP is not required before a taking a FlashCopy. The quiesce operation is not able to stop all System i host I/O, but it ensures the consistency of the database and avoids a lengthy database recovery when IPLing your system or varying on your IASP from the FlashCopy target volumes. The IPL or vary on of the FlashCopy target will still be abnormal, as though it would be taken with the application still running, which is called a *warm flash*. Using the quiesce function does not give a clean FlashCopy, which can still be achieved only with by shutting down the system or varying off the IASP. However, it ensures database consistency, can be an acceptable solution, and is much more favorable than performing a warm flash. For further information, refer to 15.1, “Using i5/OS quiesce for Copy Services” on page 432.

## **1.6.2 Metro Mirror and Global Mirror solutions**

*Metro Mirror* is the process by which a second copy is maintained on a second storage system. Metro Mirror uses synchronous data replication, which makes it impractical to use it over extended distances.

*Synchronous mirroring* means that each update to the source storage unit must also be updated in the target storage unit before the host gets the acknowledgement for the I/O to be complete. This update results in near perfect data consistency but can result in lag time between transactions. Metro Mirror copying supports a maximum distance of 300 km (186 miles). Delays in response times for Metro Mirror are proportional to the distance between the volumes. However, 100% of the source data is available at the recovery site when the copy operation is stopped.

Global Mirror processing provides a long-distance remote copy solution across two sites for open systems, z/OS®, or both open systems and z/OS data using asynchronous replication technology. Therefore, an additional copy of the data is required.

The Global Mirror function is designed to mirror data between volume pairs of a storage unit over greater distances without affecting overall performance. It is also designed to provide application consistent data at a recovery (or remote) site in case of a disaster at the local site. By creating a consistent set of remote volumes every few seconds, this function addresses the consistency problem that can be created when large databases and volumes span multiple storage units. With Global Mirror, the data at the remote site is maintained to be a point-in-time consistent copy of the data at the local site.

Global Mirror is based on existing Copy Services functions: Global Copy and FlashCopy. Global Mirror operations periodically invoke a point-in-time FlashCopy at the recovery site, at regular intervals, without disrupting the I/O to the source volume. Such operations result in a regular updating, nearly current data backup. Then, by grouping many volumes into a Global Mirror *session*, which is managed by the master storage unit, you can copy multiple volumes to the recovery site simultaneously while maintaining point-in-time consistency across those volumes.

Global Mirror processing is most often associated with disaster recovery or preparing for disaster recovery. However, you can also use it for everyday processing and data migration.

Consider using Global Mirror processing for the following reasons:

- ▶ Support for virtually unlimited distances between the local and remote sites, with the distance typically limited only by the capabilities of your network and the channel extension technology.  
This unlimited distance enables you to choose your remote site location based on business needs and enables site separation to add protection from localized disasters.
- ▶ A consistent and restartable copy of the data at the remote site, created with minimal impact to applications at your local site.
- ▶ Data currency, where your remote site might lag behind your local site by three to five seconds, minimizing the amount of data exposure in the event of an unplanned outage.  
The actual lag in data currency that you experience can depend upon a number of factors, including specific workload characteristics and bandwidth between the local and remote sites.
- ▶ Session support whereby data consistency at the remote site is internally managed across up to eight storage units that are located across the local and remote sites.
- ▶ Efficient synchronization of the local and remote sites with support for failover and failback modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage.

### 1.6.3 System i Copy Services usage considerations

In this section, we discuss some of the considerations to keep in mind when using Copy Services for System i.

#### FlashCopy usage considerations

Before i5/OS V6R1, using FlashCopy required that you shut down the system or vary off the IASP before creating a system or IASP image with FlashCopy to ensure that all of the modified data in main memory is flushed to disk.

**Note:** With the new i5/OS V6R1 quiesce for Copy Services function, shutting down the system or varying off the IASP are no longer required before taking a FlashCopy.

For further information about the new i5/OS V6R1 quiesce for Copy Services function, refer to 15.1, “Using i5/OS quiesce for Copy Services” on page 432.

#### Importance of using i5/OS journaling

Unlike taking controlled point-in-time copies using FlashCopy after a controlled quiesce or power down, with Metro Mirror and Global Mirror, which are *constantly* updating the target copy, you cannot be assured of having a clean starting point in a disaster scenario where the copy process was interrupted suddenly. There is no chance to preempt a disaster event with by shutting down the source system to flush objects from main storage. This issue applies to all environments, regardless of whether IASPs are used or not.

With both Metro Mirror and Global Mirror, you have a restartable copy, but the restart point is at the same point that the original system would be if an IPL was performed after the failure. The result is that all recovery on the target system includes abnormal IPL recovery. It is critical to employ application availability techniques such as journaling to accelerate and assist the recovery.

**Important:** As with all System i availability techniques, it is important to use i5/OS journaling to ensure that, even if objects remain in main memory, the journal receiver is written to disk. Consequently, it is copied to the disaster recovery site using Metro Mirror or Global Mirror and will be available on the disaster recovery server to apply changes to the database when the system is started.

With Metro Mirror, the recovery point is the same as the point at which the production system failed, that is a recovery point objective of zero (last transaction) is achieved. With Global Mirror, the recovery point is where the last consistency group was formed. By default Global Mirror consistency groups are formed continuously, as often as the environment allows, depending on the bandwidth and write I/O rate.

**Consistency group:** A *consistency group* is a function that can help create a consistent point-in-time copy across multiple logical unit numbers (LUNs) or volumes, and even across multiple IBM System Storage DS8000 systems.

## Points to remember about copying the entire DASD space

Remember that copying the entire DASD space creates a copy of the whole source system. Thus, you must take into consideration the following points:

- ▶ The copy is an exact copy of the original source system in every respect.
- ▶ The system name and network attributes are identical.
- ▶ The TCP/IP settings are identical.
- ▶ The BRMS network information is identical.
- ▶ User profiles and passwords are identical.
- ▶ The Job Schedule entries are identical.
- ▶ Relational database entries are identical.

You should be *extremely careful* when you activate a partition that has been built from a complete copy of the DASD space. In particular, you have to ensure that it does not connect automatically to the network, which can cause substantial problems within both the copy and its parent system.

You must ensure that your copy environment is customized correctly before attaching it to a network. Remember that booting from a SAN and copying the entire DASD space is not a high-availability solution, because it involves a large amount of subsequent work to make sure that the copy works in the environment where it is used.

### 1.6.4 System i Copy Services management solutions

There are different management tools available for IBM System Storage DS8000, DS6000™, or ESS model 800 Copy Services (see Figure 1-10 on page 25). Configuration tools for the native IBM System Storage DS® family exist, such as DS command-line interface (DS CLI), the DS Storage Manager GUI, and System Storage Productivity Center for Replication. From the System i perspective, these tools are stand-alone Copy Services management tools that provide no integration with System i clustering. Consequently, these tools are suitable for managing Copy Services with System i only for full-system disaster recovery solutions but not for System i disaster recovery and high-availability solutions using System i clustering with switchable independent ASPs.

The new i5/OS V6R1 High Availability Solutions Manager (HASM) or the System i Copy Services Toolkit are System i Copy Services management tools that provide a set of functions to combine PPRC, IASP, and i5/OS cluster services for coordinated switchover and failover processing through a cluster resource group (CRG), which is not provided by stand-alone Copy Services management tools such as System Storage Productivity Center for Replication or DS CLI. Both HASM and the toolkit require that you have installed i5/OS option 41 (HA switchable resources) and DS CLI on each system that participates in your high availability recovery domain. They provide the benefit of the Remote Copy function and coordinated switching of operations, which gives you good data resiliency capability if the replication is done synchronously.

**Note:** Using independent ASPs with Copy Services is only supported with using either HASM or the System i Copy Services Toolkit and a pre-sale and pre-install Solution Assurance Review is highly recommended or required.



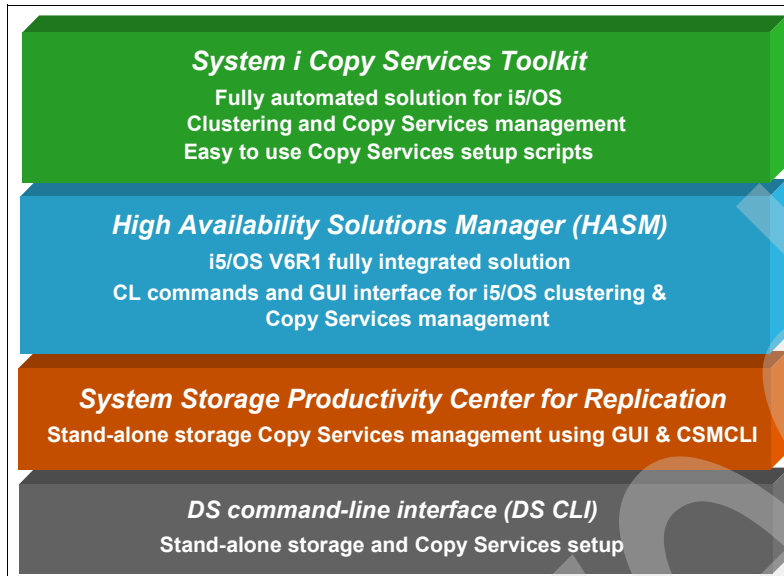


Figure 1-10 Enhanced functionality of integrated System i Copy Services Management Tools

## System Storage Productivity Center for Replication

Although System Storage Productivity Center for Replication provide no integration with System i, it reduces the complexity of Copy Services management by introducing the concept of copy sets and sessions:

- ▶ *Copy sets* represent volumes from the *same type* of Copy Services relationship that have a copy of the same data like either source and target of a FlashCopy volume pair or A, B, and C volumes of a Global Mirror volume relationship.
- ▶ *Sessions* are container entities for grouping copy sets together, for example on an application or host level, for easier management and achieving data consistency. Figure 1-11 shows an example of the sessions panel from the System Storage Productivity Center for Replication GUI.

Name	Status	Type	State	Active Host	Recoverable	Copy Sets
<a href="#">Webserver</a>	Severe	MM	Suspended	H1	Yes	16
<a href="#">DB2</a>	Normal	MM	Prepared	H1	Yes	16
<a href="#">GM1</a>	Normal	GM	Prepared	H1	Yes	32
<a href="#">DS8kMMFOFB</a>	Inactive	MM	Defined	H1	No	16
<a href="#">FlashDS8k</a>	Inactive	FC	Defined	H1	No	0
<a href="#">GM_FOFB</a>	Inactive	GM	Defined	H1	No	16
<a href="#">MGM_8k</a>	Inactive	MGM	Defined	H1	No	16
<a href="#">MM-Interfere1</a>	Inactive	MM	Defined	H1	No	2
<a href="#">MM-Interfere2</a>	Inactive	MM	Defined	H1	No	2

Figure 1-11 System Storage Productivity Center for Replication GUI: Sessions overview panel

System Storage Productivity Center for Replication implements Metro Mirror data consistency based on PPRC consistency groups and issuing a freeze operation triggered by a PPRC failure condition. By issuing a freeze operation against all logical subsystems (LSS) in the session, write I/O to all primary volumes of the session is temporarily halted, the primary volumes are suspended, and PPRC paths are removed to ensure crash-like data consistency at the secondary site. With the IBM System Storage disk subsystem sending a SNMP trap 200 for a PPRC consistency group volume pair error automation software can be used to unfreeze automatically the primary volumes or to initiate a failover to the secondary site. However, this kind of automation is not applicable for System i because of its single-level storage architecture, which makes no use of the IBM System Storage SNMP messages.

### **i5/OS High Availability Solution Manager**

IBM System i HASM is a new i5/OS V6R1 licensed program (5761-HAS) that provides a GUI, a command-line interface, and APIs for configuring and managing System i high-availability solutions using IASPs on either System i internal or external storage.

HASM combined with i5/OS cluster version 6 is the first end to end complete *native* and fully *integrated* i5/OS high availability solution. For managing external storage Copy Services HASM provides similar functionality to the System i Copy Services Toolkit. However compared to the Toolkit, HASM supports IASPs only, that is no full-system Copy Services, does not provide the Toolkit's Copy Services setup scripts to ease setting up the external storage Copy Services configuration and does not provide the level of switch-over automation with included PPRC state error checking.

You can implement high availability with the HASM GUI integrated in IBM Systems Director Navigator for i5/OS using either a solution-based approach or a task-based approach. The solution-based approach accessible from the High Availability Solutions Manager GUI navigation tree item guides you through verifying your environment as well as setting up and managing your chosen solution. Currently the solution-based approach supports the following configurations:

- ▶ Switched disk between logical partitions
- ▶ Switched disk between systems
- ▶ Switched disk with Geographic Mirroring (3-site solution)
- ▶ Cross-site mirroring with Geographic Mirroring

Solutions using external storage Copy Services for replicating an IASP are supported only using the task-based approach which allows you to design and build a customized high-availability solution for your business, using primarily the IBM Systems Director Navigator for i5/OS Cluster Resource Services and Disk Management interfaces.

Figure 1-12 shows the HASM GUI with the solution-based approach selected.

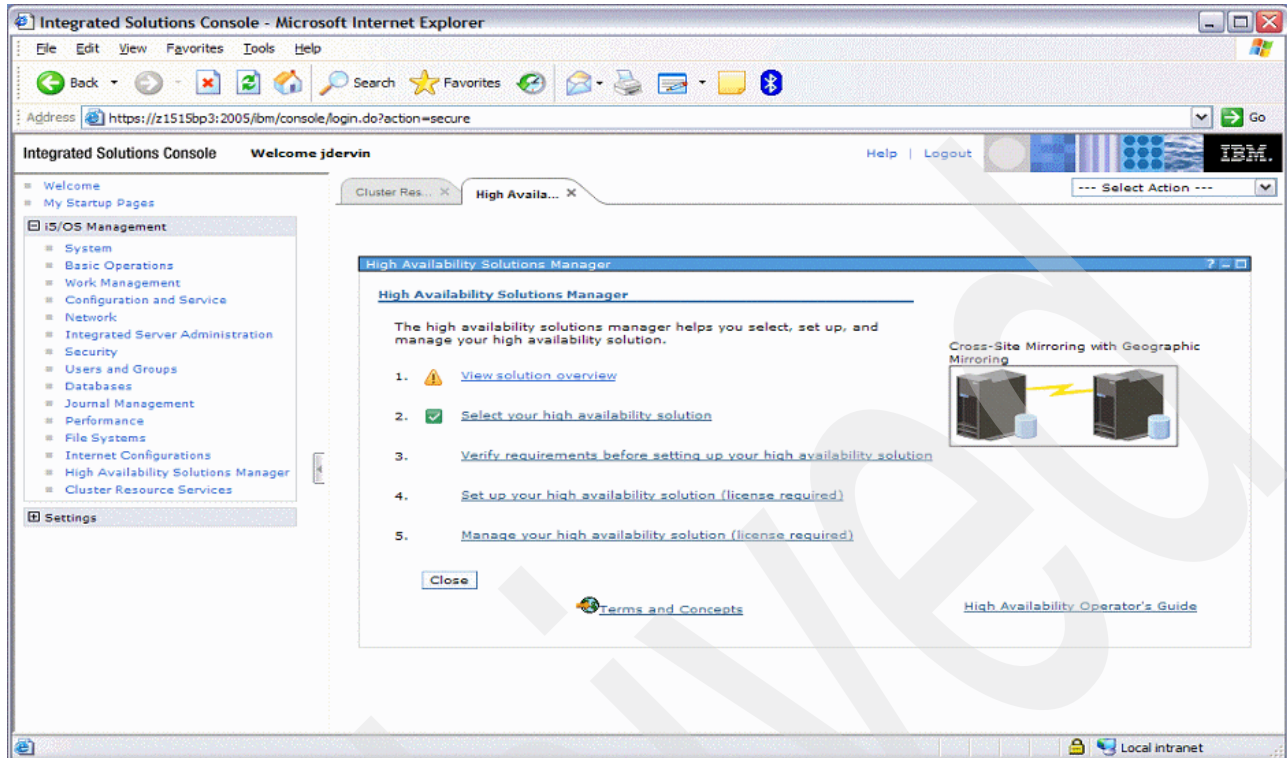


Figure 1-12 HASM GUI integrated in IBM Systems Director Navigator for i5/OS

## System i Copy Services Toolkit

The System i Copy Services Toolkit is a services offering from IBM STG lab services that was developed by IBM Rochester. It blends two technologies, using the System i availability architecture that is provided by IASPs along with the advanced functions that are provided by IBM System Storage Copy Services FlashCopy, Metro Mirror, and Global Mirror. The toolkit is a combination of management software to control the IASP environment and the services to implement it. It provides a fully-automated solution for System i and Copy Services management covering clustering with data CRGs, switch-/failover and also Copy Services

setup using provided scripts for DS CLI on i5/OS (see Figure 1-13). A customized extension is available for the toolkit to support a three-site synchronous/asynchronous external storage data replication solution similar to Metro Global Mirror.

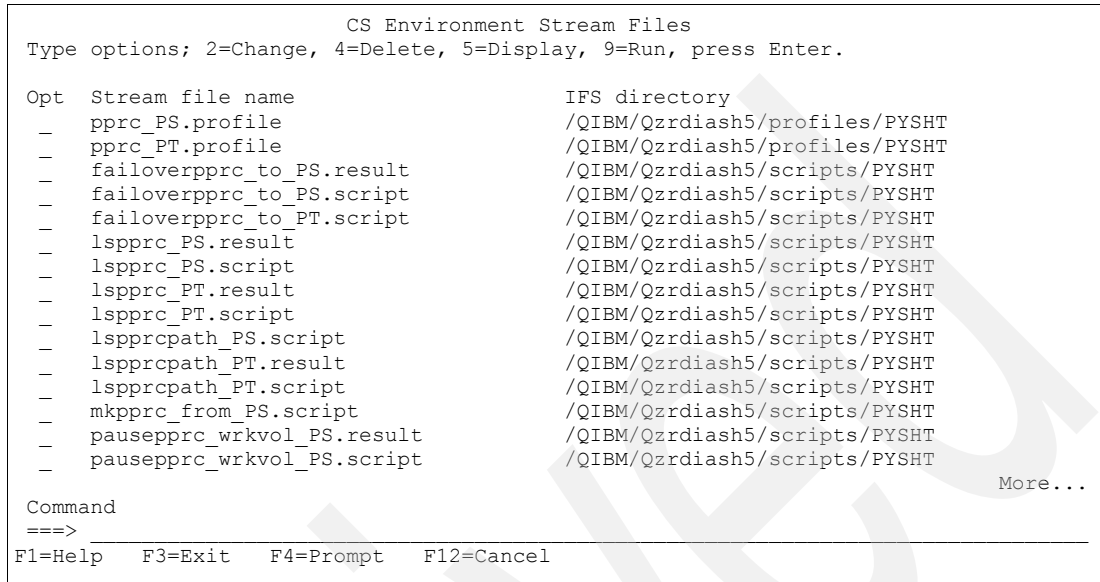


Figure 1-13 CS Toolkit: Copy Services scripts

Additionally a full-system FlashCopy Toolkit is available as a solution for FlashCopy backup automation from an i5/OS management LPAR through the HMC SSH interface. It completely automates the actual save process and is fully integrated with the BRMS steps for full system backups. That is, it will lock the production BRMS configuration before the flash and replicate the data back to production after the flash.

Both toolkit versions provide the Copy Services environment panels to manage FlashCopy and either full-system or IASP Metro Mirror or Global Mirror (see Figure 1-14), which can be used to manage a PPRC site switch-over for non-i5/OS systems.

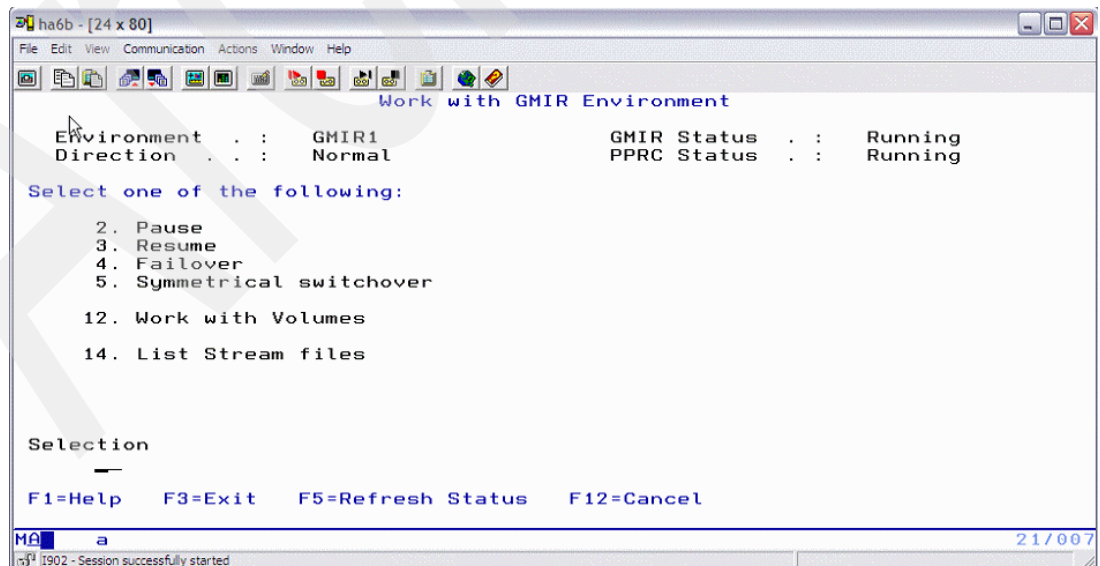


Figure 1-14 CS Toolkit: Work with Global Mirror Environment panel

For more information about the toolkit, contact the High Continuous Availability and Cluster group within the IBM System i Technology Center (iTC) by sending e-mail to [rhc1st@us.ibm.com](mailto:rhc1st@us.ibm.com) or IBM System Storage Advanced Technical Support.

Archived

Archived



## System i external storage solution examples

In this chapter, we discuss possible scenarios where System i environments and IBM System Storage DS solutions are connected. We start by providing basic examples for a 1-site local solution with using System i external storage optionally together with IASPs for high availability (HA) and FlashCopy for backup. Then, we discuss 2-site solutions that take advantage of remote data replication to a secondary site for disaster recovery (DR) and high availability. Using these example environments can guide you through the planning and implementation of external storage for i5/OS and the System i platform.

## 2.1 One-site System i external storage solution examples

Attaching an external disk to a System i platform is a relatively simple task if those who are performing the task understand the i5/OS operating system environment, the external storage environment, and the storage area network (SAN). The following sections show 1-site solution examples of implementing IBM System Storage DS8000, DS6000, or ESS model 800 external storage with System i.

### 2.1.1 System i5 model and all disk storage in external storage

In this scenario, the System i model has all its disk storage, including the load source, in the external storage server. Such a boot from SAN configuration is available only to HMC managed System i POWER5 or later servers and IBM System Storage model 800, DS6000 and DS8000 series. For further information about boot from SAN requirements, refer to 3.2.1, “Planning considerations for boot from SAN” on page 54)

Figure 2-1 shows the simplest example. The i5/OS load source is a logical unit number (LUN) in the DS model. To avoid single-points of failure for the storage attachment i5/OS multipathing should be implemented to the LUNs in the external storage server.

**Note:** With i5/OS V6R1 and later, multipathing is also supported for the external load source unit.

Prior to i5/OS V6R1 the external load source unit should be mirrored to another LUN on the external storage system to provide path protection for the load source. The System i model is connected to the DS model with Fibre Channel (FC) cables through a storage area network (SAN).

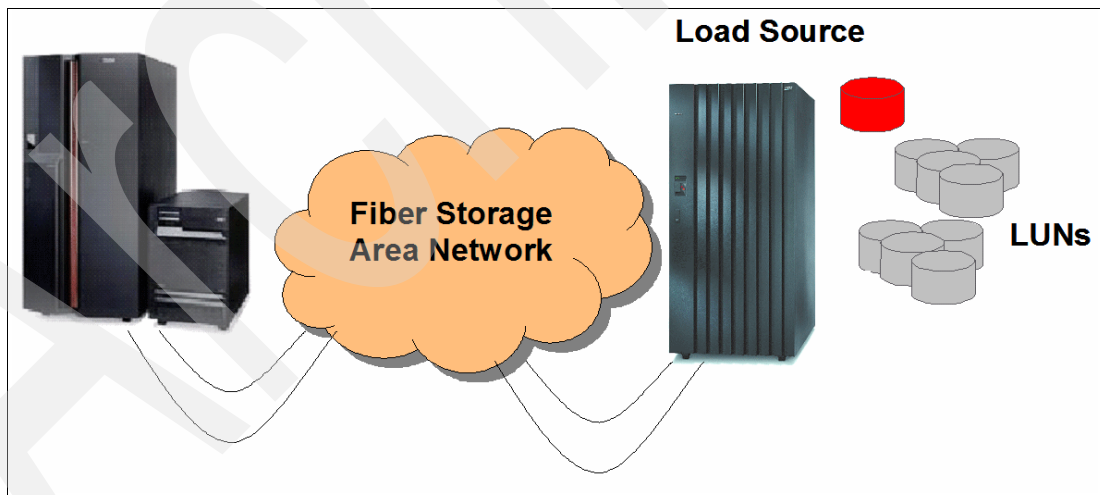


Figure 2-1 System i5 model and all disk storage in external storage



The FC connections through the SAN switched network are either through direct FC local connections or through a dark fiber up providing up to 10 km distance. Figure 2-2 is the same simple example, but the System i platform is divided into logical partitions (LPAR). Each LPAR has its own mirrored pair of LUNs in the DS model.

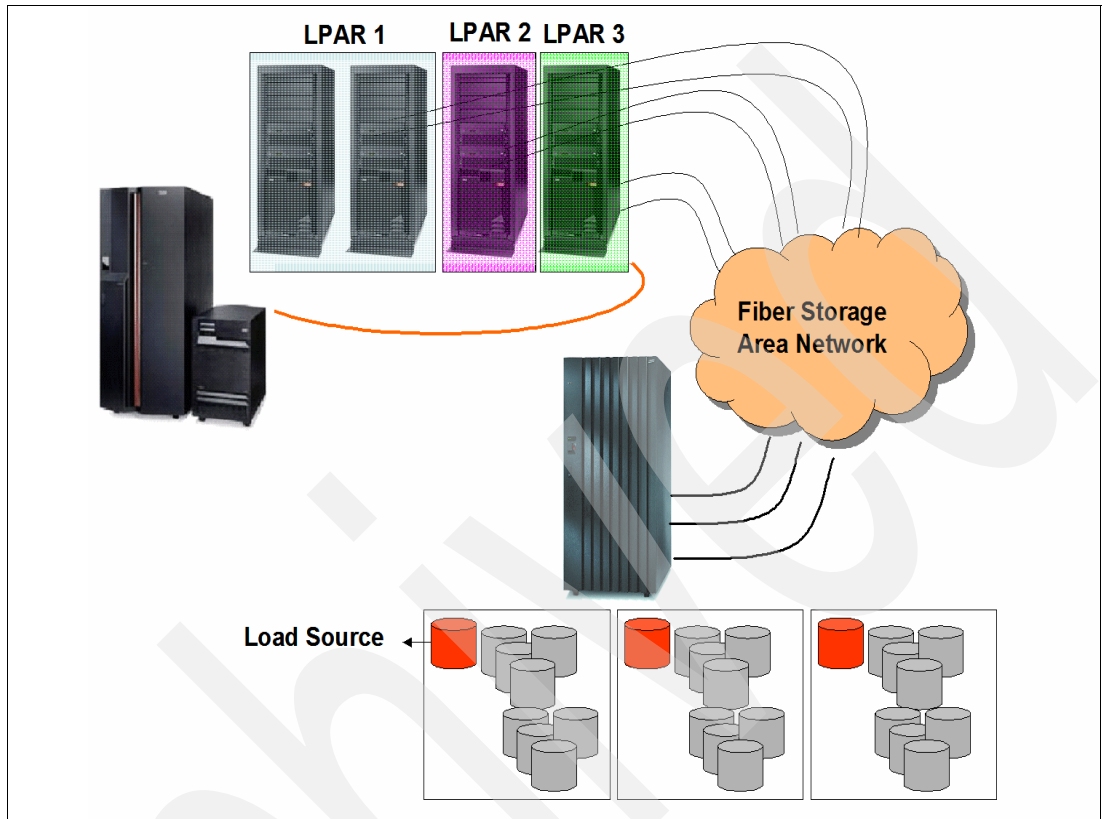


Figure 2-2 LPAR System i5 environment and all disk storage in external storage

## 2.1.2 System i model with internal load source and external storage

This example shows the previous scenario for implementing external disk with System i models but without boot from SAN. In this case the load source drive remains in either the System i central electronic complex (CEC) or the expansion tower where the system is logically partitioned. In this example, there are three logical partitions with the load source in an expansion tower. There is a remote load source in the external storage system for protection. Multipath can be implemented to all the LUNs in the external storage server.

Unless using switchable independent ASPs boot from SAN helps to significantly reduce the recovery time in case of a system failure by eliminating the requirement for a manual D-type IPL with remote load source recovery.

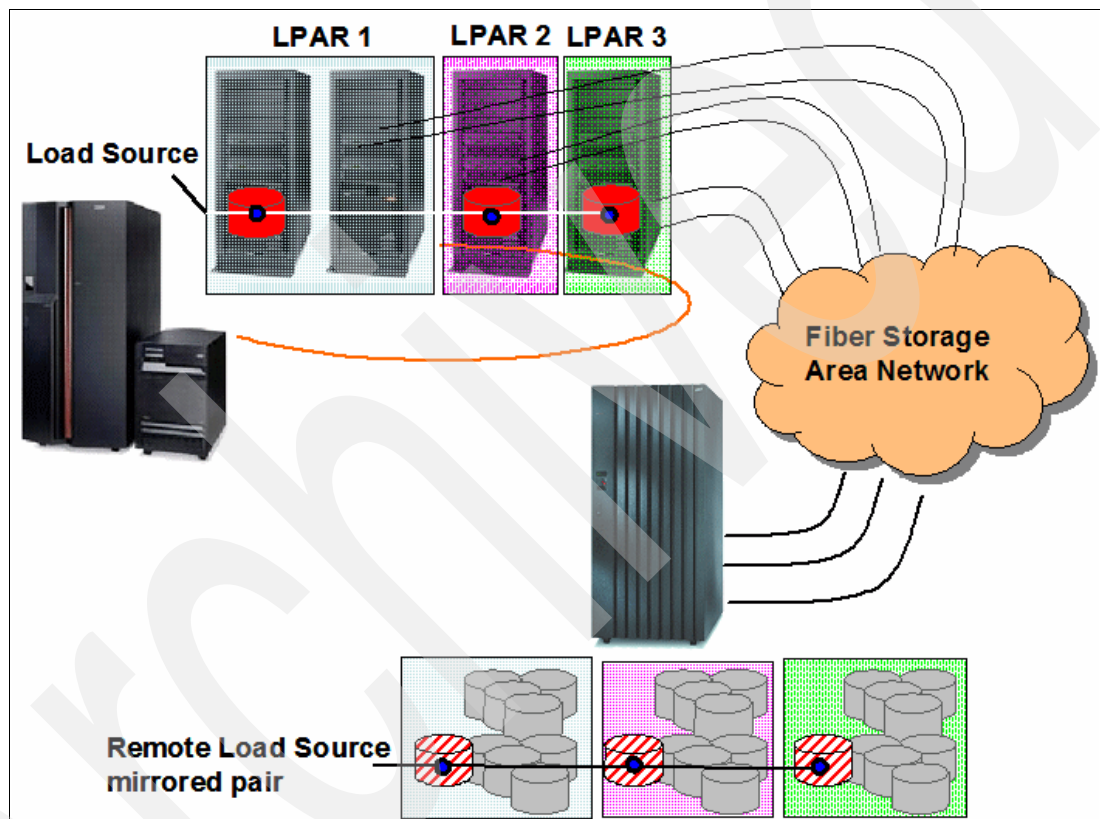


Figure 2-3 External disk with the System i5 internal load source drive

## 2.1.3 System i model with mixed internal and external storage

Examples of selected environments where the internal disk is retained in the System i model and additional disk is located in the external storage server include:

- ▶ A solution where the internal drives support \*SYSBAS storage and the external storage supports the IASP, which is similar to the example in “Metro Mirror with switchable IASP replication” on page 42.
- ▶ A solution where the internal drives are one half of the mirrored environment and the external storage LUNs are the other half, giving mirrored protection and distance capability.

- ▶ A solution that requires a considerable amount of space for archiving.

In Figure 2-4, the external disk is used typically for a user auxiliary storage pool (ASP) or an independent ASP (IASP). This ASP disk space can house the archive data, and his storage is fairly independent of the production environment.

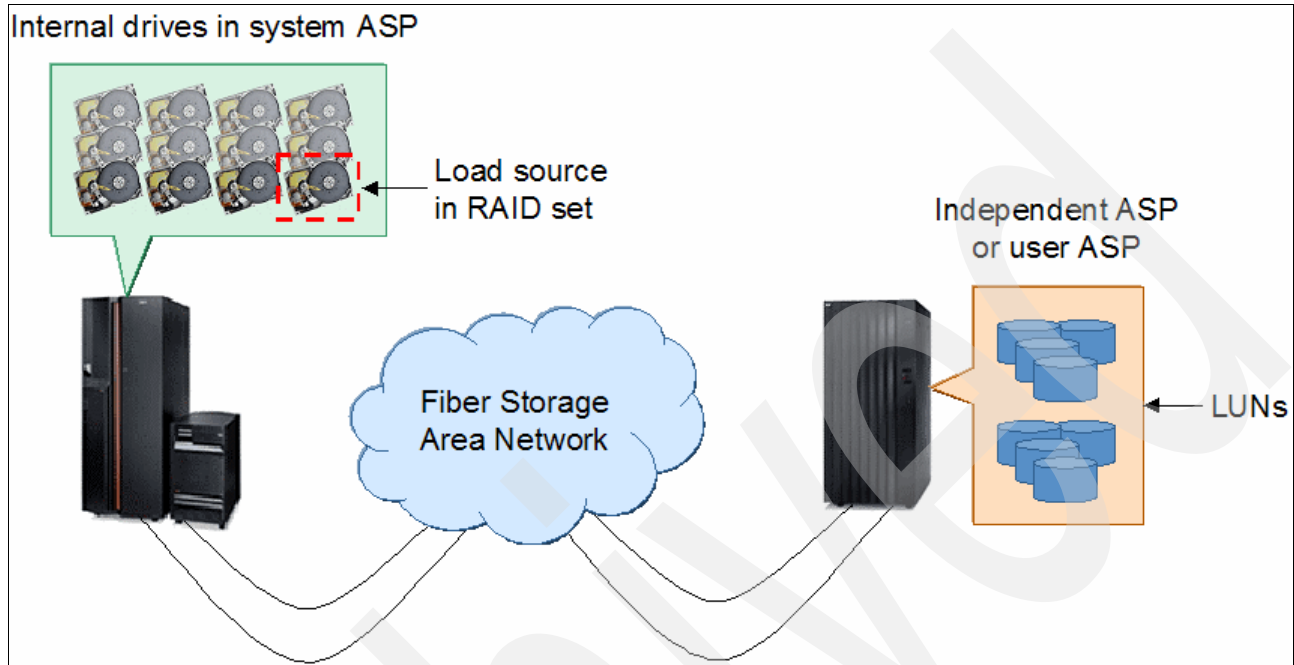


Figure 2-4 Mixed internal and external drives

It is possible to mix internal and external drives in the same ASP, but we do not recommend this mixing because performance management becomes difficult.

## 2.1.4 Migration of internal drives to external storage including load source

In this case, the customer has decided to adopt a consolidated storage strategy. The customer must have a path to migrate from their internal drives to the new external disk. For our example (shown in Figure 2-5), we assume the internal disk drives are all RAID protected.

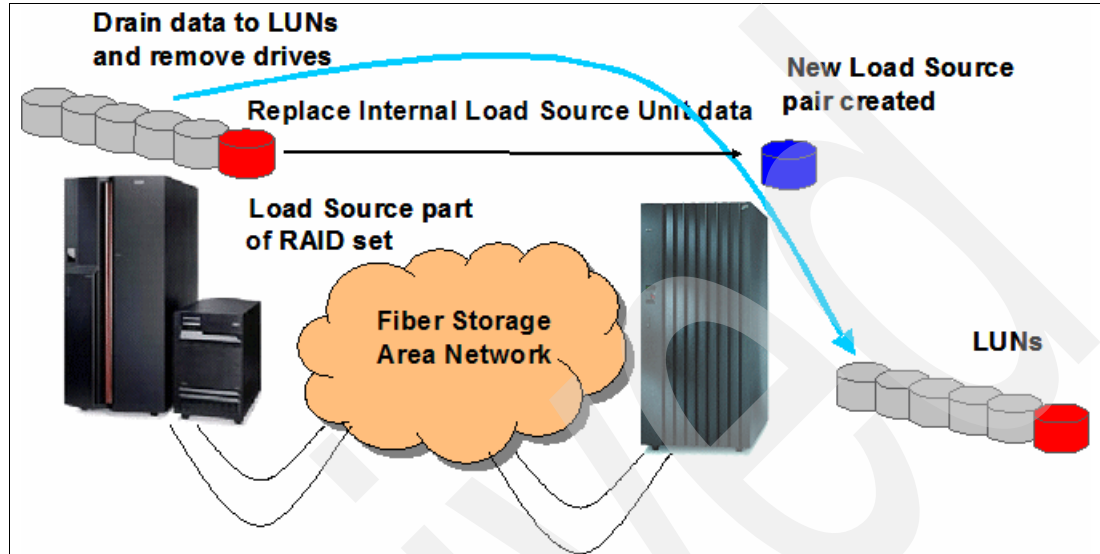


Figure 2-5 Migration from internal RAID protected disk drives to external storage

There are multiple techniques for implementing this migration.

One such technique is to add additional I/O hardware to the existing System i model to support the new external disk environment. This hardware can be an expansion tower, I/O loops (HSL or 12X), #2847 IOP-based or POWER6 IOP-less Fibre Channel IOAs for external load source support, other #2844 IOP-based or IOP-less FC adapters for the non-load source volumes.

The movement of data from internal to external storage is achieved by the Disk Migration While Active function (see Figure 2-5). Not all data is removed from the disk. Certain object types, such as temporary storage, journals and receivers, and integrated file system objects, are not moved. These objects are not removed until the disk is removed from the i5/OS configuration. The removal of disk drives is disruptive, because it has to be done from DST. The time to remove them depends on the amount of residual data left on the drive.

The removal method roughly follows this process:

1. Plan your data migration. When removing disks from the configuration, you must understand the RAID set arrangements to maintain protection.
2. Test the draining process outside the production environment to ensure that you are confident with the process.
3. Increase the load source drive to 17 GB or more, and load new operating system support.
4. Attach the new I/O and external storage.
5. Create LUNs in external storage.
6. Add LUNs to i5/OS.
7. Use the Disk Migrate While Active function (\*MOVDTA) of the Start ASP Balance (STRASPBAL) command on the drives that are to be drained.

8. For further details about this function, refer to *IBM eServer iSeries Migration: System Migration and Upgrades at V5R1 and V5R2*, SG24-6055.
9. Perform a manual IPL to DST, and remove the disks that have had the data drained from the i5/OS configuration.
10. Stop device parity protection for the load source RAID set.
11. Migrate the load source drive by copying the load source unit data.
12. Physically remove the old internal load source unit.
13. Change the I/O tagging to the new external load source.
14. Re-start device parity protection.

For detailed information about migrating an internal load source to boot from SAN, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

**Attention:** The Disk Migrate While Active function starts a job for every disk migration. These jobs can impact performance if many are started. If data is migrated from a disk and the disk is not removed from the configuration, a job is started. Do not start data moves on more drives than you can support without impacting your existing workload. Schedule the data movement outside normal business hours.

### 2.1.5 Migration of an external mirrored load source to a boot load source

In this example, the System i environment is already attached to an external storage server before the new boot from SAN support became available. Typically, in this environment, the solution includes the internal load source that is mirrored to an external pair, which is a similar-sized LUN in the external storage server.

This technique provides protection for the internal load source. The System i load source drive should always be protected either by RAID or mirroring.

To migrate from a remote mirrored load source to external mirrored load source (Figure 2-6):

1. Increase the size of your existing load source to 17 GB or greater.
2. Load the new i5/OS V5R3M5 or later operating system support for boot from SAN.
3. Create the new mirrored load source pair in the external storage server.
4. Turn off System i and change the load source I/O tagging to the remote external load source.
5. Remove the internal load source.
6. Perform a manual IPL to DST.
7. Use the replace configured unit function to replace the internal suspended load source with the new external load source.
8. Perform an IPL on the new external mirrored load source.

For detailed information about migrating an internal load source to boot from SAN, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

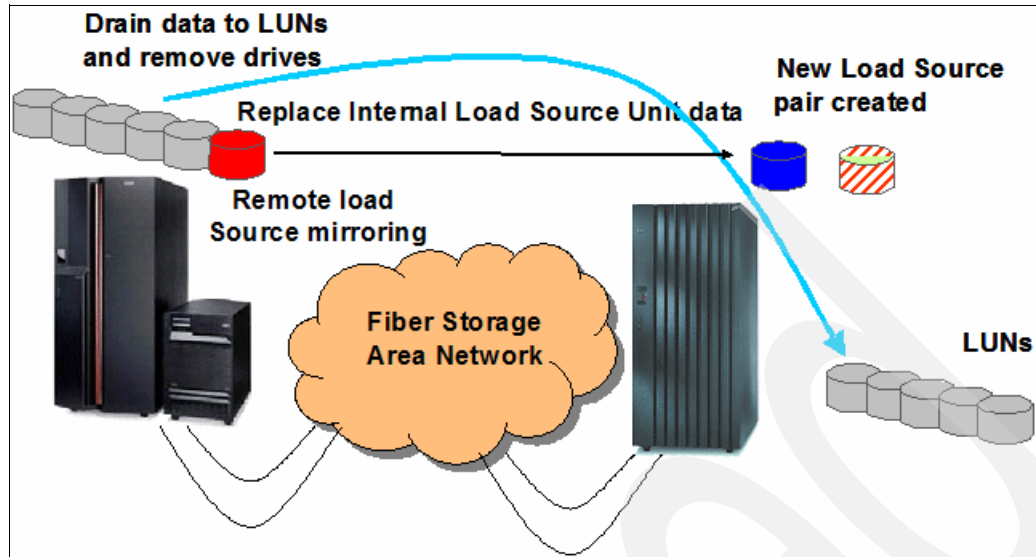


Figure 2-6 Migration of the load source to external storage

## 2.1.6 Cloning i5/OS

Cloning is a new concept for the System i platform since the introduction of boot from SAN with V5R3M5. Previously to create a new system image, you had to perform a full installation of the SLIC and i5/OS. When cloning i5/OS, you create an exact copy of the existing i5/OS system or partition. The copy can be attached to another System i model, a separate LPAR, or if the production system is powered off, the existing partition or system. After the copy is created, you can use it for offline backup, system testing, or migration.

Boot from SAN enables you to take advantage of some of the advanced features that are available with the DS8000 and DS6000 family, such as FlashCopy. It allows you to perform a point-in-time instantaneous copy of the data held on a LUN or group of LUNs. Therefore, when you have a system that has only SAN LUNs with no internal drives, you can create a clone of your system.

**Important:** When we refer to a *clone*, we are referring to a copy of a system that only uses SAN LUNs. Therefore, boot from SAN is a prerequisite.

## 2.1.7 Full system and IASP FlashCopy

FlashCopy allows you to take a system image for cloning and is also an ideal solution for increasing the availability of a System i production system by reducing the time for system backups.

To obtain a full system backup of i5/OS with FlashCopy, either a system shutdown or, since i5/OS V6R1, a quiesce is required to flush modified data from memory to disk. FlashCopy copies only the data on the disk. Therefore, a significant amount of data is left in memory, and extended database recovery is required if the FlashCopy is taken with the system running or not suspended.

**Note:** The new i5/OS V6R1 quiesce for Copy Services function (CHGASPACT) allows you to suspend all database I/O activity for \*SYSBAS and IASP devices before taking a FlashCopy system image eliminating the requirement to power down your system (see 15.1, “Using i5/OS quiesce for Copy Services” on page 432)

An alternative method to perform offline backups without a shutdown and IPL of your production system is using FlashCopy with IASPs, as shown in Figure 2-7. You might consider using an IASP FlashCopy backup solution for an environment that has no boot from SAN implementation or that is using IASPs in anyway for high availability. Because the production data is located in the IASP, the IASP can be varied off or because i5/OS V6R1 quiesced before taking a FlashCopy without shutting down the whole i5/OS system. It also has the advantage that no load source recovery is required.

**Note:** Temporary space includes QTEMP libraries, index build space, and so on. There is a statement of direction to allow spooled files in an IASP in the future.

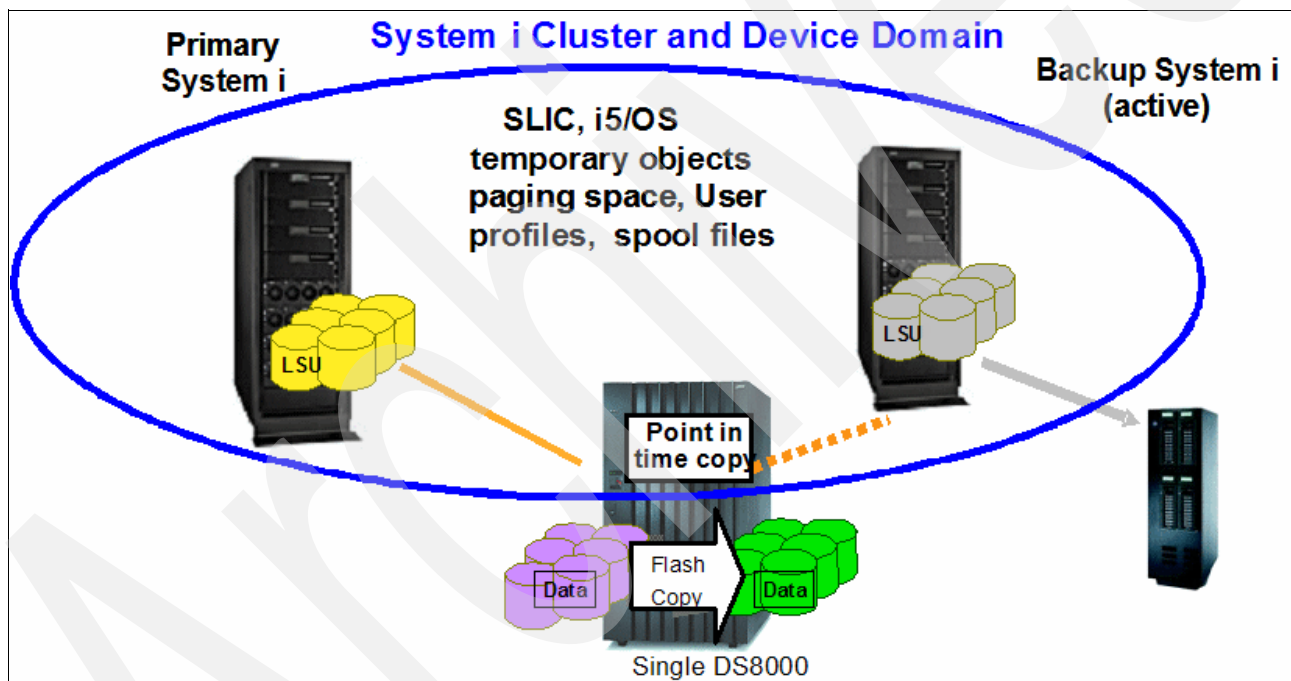


Figure 2-7 FlashCopy of IASP for offline backup

### Planning considerations

Keep in mind the following considerations:

- ▶ You must vary off or quiesce the IASP before the FlashCopy can be taken. Customer application data must be in an IASP environment in order to use FlashCopy. Using storage-based replication of IASPs requires using the System i Copy Services Toolkit or the new i5/OS V6R1 High Availability Solutions Manager (HASM) (see 1.6.4, “System i Copy Services management solutions” on page 24)
- ▶ Disk sizing for a system ASP is important because it requires the fastest disk on the system because this is where memory paging, index builds, and so on happen.

- ▶ An IPL is required of the backup system after a save to clean up the cluster management objects on the target system.
- ▶ A separate FC I/O processor (IOP) or I/O adapter (IOA) is required for each IASP on the target system. For more information, contact the High Continuous Availability and Cluster group within the IBM System i Technology Center (ITC) by sending e-mail to [rhc1st@us.ibm.com](mailto:rhc1st@us.ibm.com).

## 2.2 Two-site System i external storage solution examples

We now take a closer look at some examples of 2-site solutions that maintain a copy of your production data at a second remote site for disaster recovery (DR) and high availability (HA).

### 2.2.1 The System i platform and external storage HA environments

With huge demand for high available business systems, there are many instances where the System i platform collaborates with external storage servers to take advantage of both systems availability features and applications. The System i platform offers both RAID and mirrored protection for its disk subsystem. This function is provided by the operating system and I/O adapters (IOAs). Customers who have external storage can take advantage of the i5/OS mirroring function, which gives the possibility of separation by up to 10 km, giving this solution disaster recovery characteristics.

Figure 2-8 shows a System i model with internal drives that are one half of the mirror to an external storage server that is at a distance with a remote load source mirror and a set of LUNs that are mirrored to the internal drives.

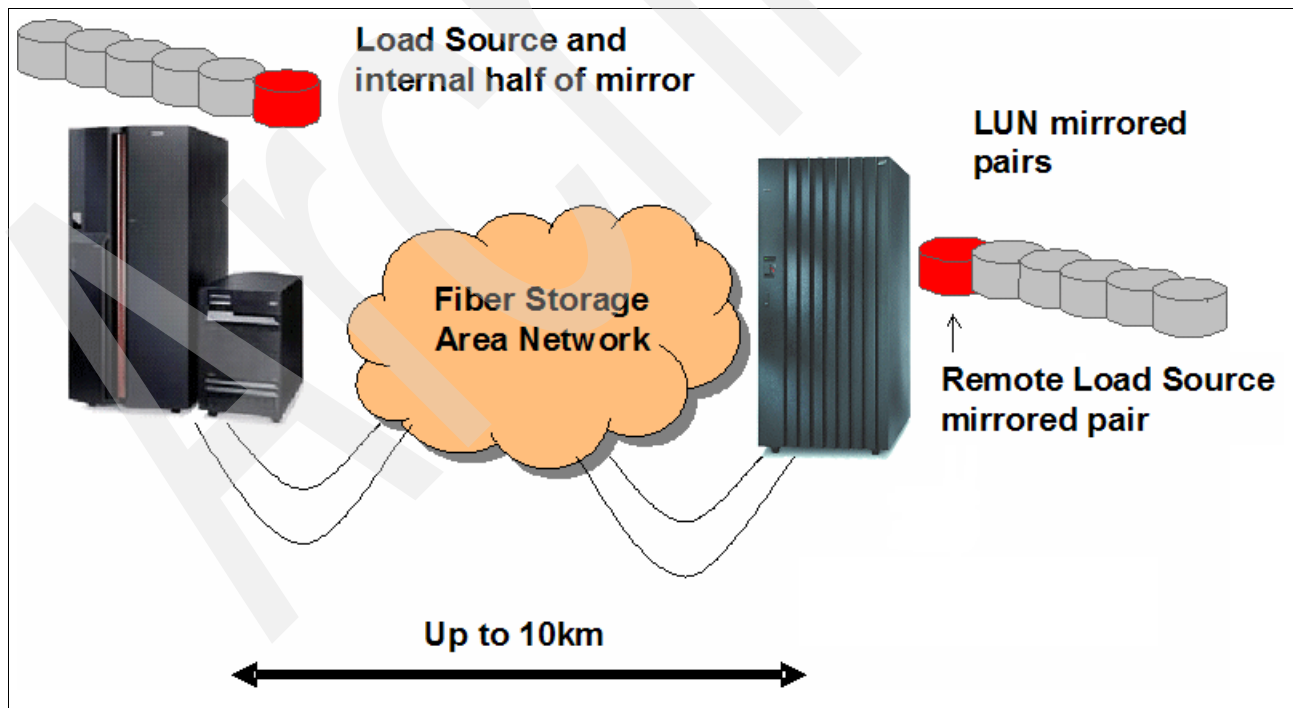


Figure 2-8 Internal to external mirroring for disaster recovery

If the production site has a disk hardware failure, the system can continue off the remote mirrored pairs. If a disaster occurs that causes the production site to be unavailable, it is



possible to IPL your recovery System i server from the attached remote LUNs. If your production system is running i5/OS V5R3M5 or later and your *recovery* system is configured for boot from SAN, it can directly IPL from the remote load source even without requiring a remote load source recovery.

**Restriction:** If using i5/OS mirroring for disaster recovery as we describe, your *production* system must *not* use boot from SAN because, at failback from your recovery to your production site, you cannot control which mirror side you want to be the active one.

## 2.2.2 Metro Mirror examples

In the following sections, we describe how to use Metro Mirror in conjunction with the System i platform.

### Metro Mirror and full system replication

Metro Mirror offers synchronous replication between two DS models or between a DS and ESS model 800. In the example shown in Figure 2-9 on page 42, two System i servers are separated by some distance to achieve a disaster recovery solution at the second site. This is a fairly simple arrangement to implement and manage. Synchronous replication is desirable because it ensures the integrity of the I/O traffic between the two storage complexes and provides a *recovery point of objective* (RPO) of zero (that is, no transaction gets lost). The data on the second DS system is not available to the second System i model while Metro Mirror replication is active, that is it must be turned off.

The main consideration with this solution is distance. The solution is limited by the distance between the two sites. Synchronous replication needs sufficient bandwidth to prevent latency in the I/O between the two sites. I/O latency can cause application performance problems. Testing is necessary to ensure that this solution is viable depending on a particular application's design and business throughput.

When you recover in the event of a failure, the IPL of your recovery system will always be an abnormal IPL of i5/OS on the remote site.

**Note:** Using i5/OS journaling for Metro Mirror or Global Mirror replication solutions is highly recommended to ensure transaction consistency and faster recovery.

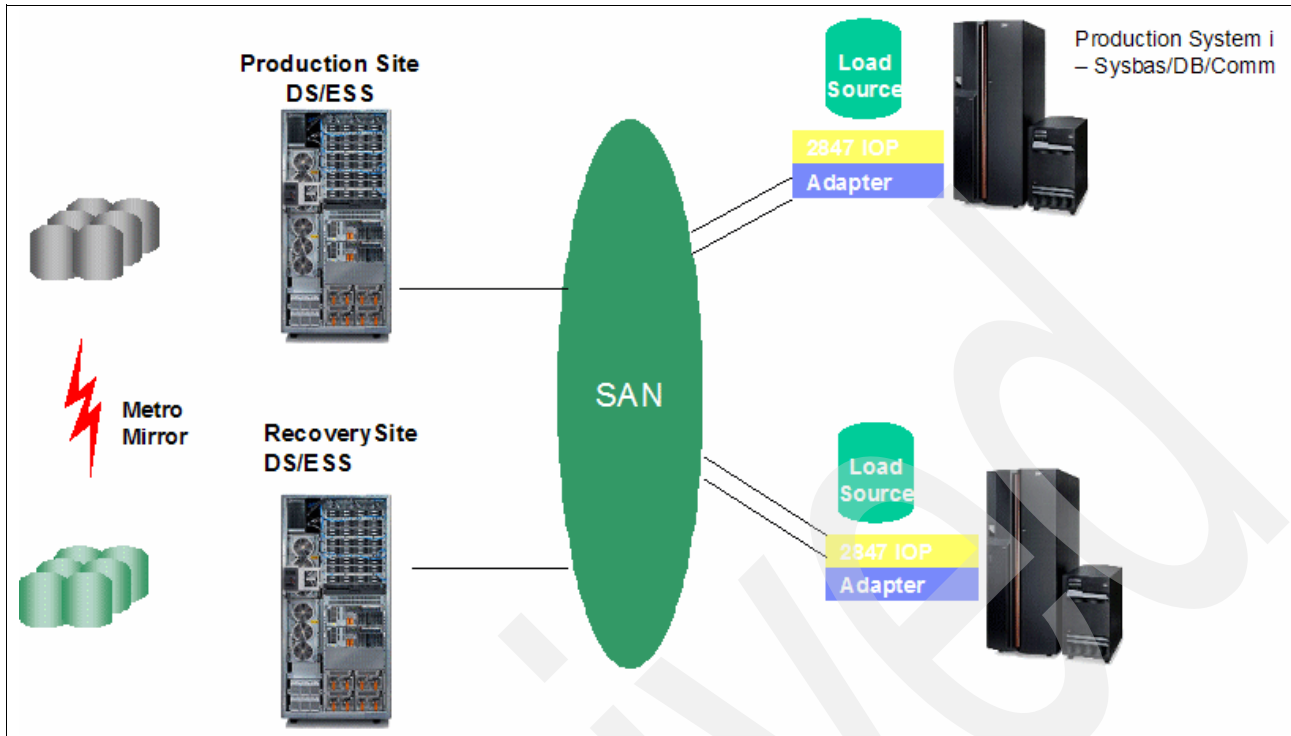


Figure 2-9 Metro Mirror full-system replication

### Metro Mirror with switchable IASP replication

In this example, we have the same DS configuration, but the data that we want to replicate is in an IASP located in the DS model (see Figure 2-10). The \*SYSBAS disks can be on an internal or external storage. Both the production system and the recovery system have to be a System i cluster device domain. The IASP is connected to the System i model through a switchable expansion tower in a device cluster resource group (CRG). While the backup system is powered on and can be running other applications, none of the data in the IASP environments that are shown is available to the backup system until a switchover occurs.

**Note:** Replicating switchable independent ASPs to a remote site provides both disaster recovery and high availability and is supported only with either using the System i Copy Services Toolkit or i5/OS V6R1 High Availability Solutions Manager (HASM).

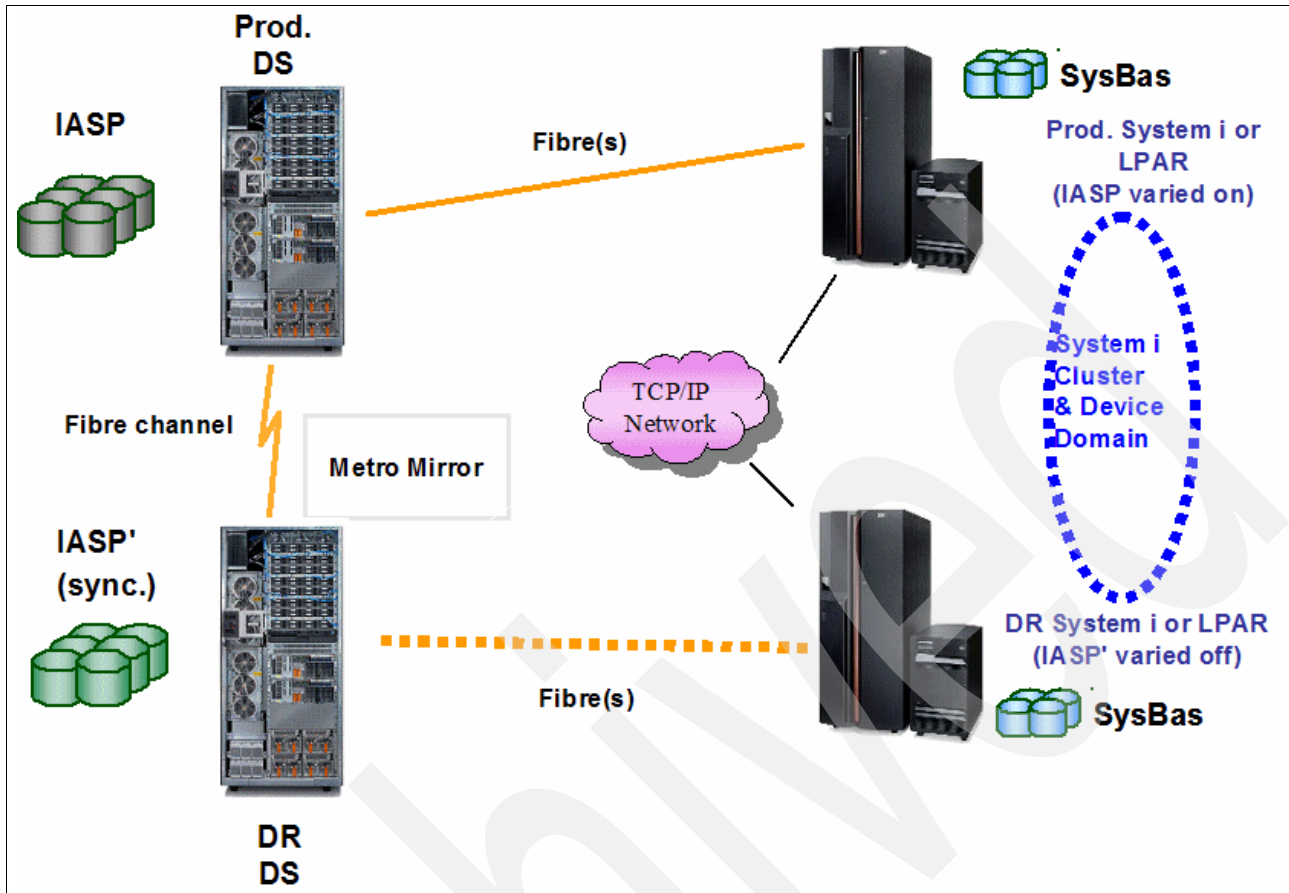


Figure 2-10 Metro Mirror IASP replication

Using switchable IASPs with Copy Services requires either the System i Copy Services Toolkit or the new i5/OS V6R1 High Availability Solutions Manager (HASM) for managing the failover or switchover. If there is a failure at the production site, i5/OS cluster management detects the failure and switches the IASP to the backup system. In this environment, we normally have only one copy of the IASP, but we are using Copy Services technology to create a second copy of the IASP at the remote site and provide distance.

The switchover and the recovery to the backup system are a relatively simple operation, which is a combination of i5/OS cluster services commands and DS command-line interface (CLI) commands. The IASP switch is cluster services passing the management over to the backup system. The backup IASP is then varied on the active backup system. During a disaster journal recovery attempts to recover or rollout any damaged objects. After the vary on action completes, the application is available. These functions are automated with the System i Copy Services Toolkit (see 1.6.4, “System i Copy Services management solutions” on page 24).

### 2.2.3 Global Mirror examples

In this section we present examples of using Global Mirror with the System i platform. Compared with Metro Mirror synchronous replication Global Mirror uses asynchronous replication of data consistency groups to allow for long-distance replication solutions with

guaranteeing data consistency. The design of Global Mirror prevents performance impacts to the production host provided that enough replication link bandwidth is available.

### Global Mirror and full system replication

In this example (Figure 2-11), no disk is located inside the production or backup system; all System i disk units are provided from the DS models. This is a disaster recovery environment. For this full-system replication scenario i5/OS clustering is not involved so there is no switchover.

All the data on the production system is asynchronously transmitted to the remote DS model. Asynchronous replication through Global Copy alone does not guarantee the order of the writes, and the remote production copy will lose consistency quickly. In order to guarantee data consistency Global Mirror creates consistency groups at regular intervals, by default as fast as the environment and the available bandwidth allows. FlashCopy is used at the remote site to save these consistency groups to ensure a consistent set of data is available at the remote site which is only a few seconds behind the production site, i.e. with using Global Mirror a recovery point objective (RPO) of only a few seconds can be achieved normally without any performance impact to the production site.

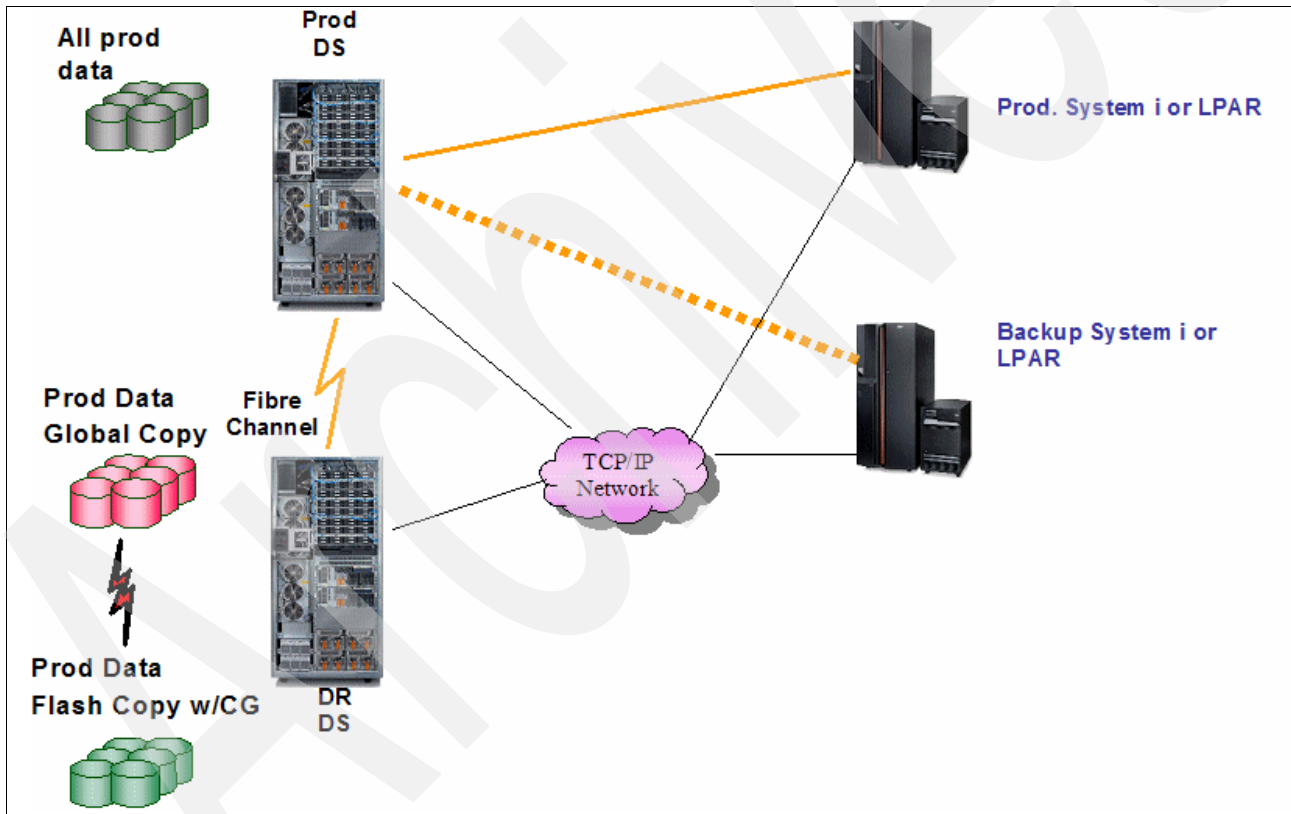


Figure 2-11 Global Mirror and the System i5 platform

This is an attractive solution because of the extreme distances that can be achieved with Global Mirror. However, it requires a proper sizing of the replication link bandwidth to ensure the RPO targets can be achieved, and testing should be performed to ensure the resulting image is usable.

### Global Mirror and switchable IASP replication

Global Mirror and switchable IASPs offer a new and exciting opportunity for a highly available environment. It enables customers to replicate their environment over an extremely long

distance without the use of traditional i5/OS replication software. This environment comes in two types, asymmetrical and symmetrical.

While Global Mirror can entail a fairly complex setup, the operation of this environment is simplified for i5/OS with the use of the System i Copy Services Toolkit, automating the switchover and failover or the IASP from production to backup.

### Asymmetrical replication

The configuration shown in Figure 2-12 provides both availability switching between the production system and the backup system. It also provides disaster recovery between either the production system or backup system, depending on which system has control when the disaster occurs, and the disaster recovery system. With the asymmetrical configuration, only one consistency group is setup, and it resides at the remote site. This means that you cannot do regular role swaps and reverse the I/O direction (disaster recovery to production).

In a normal operation, the IASP holds the application data and runs varied on to the production system. I/O is asynchronously replicated through Global Copy to the backup DS model maintaining a copy of the IASP. At regular intervals, FlashCopy is used to save the consistency groups created at repeated intervals by the Global Mirror algorithm. The consistency groups can be only a few seconds behind the production system, offering the opportunity for a fast recovery.

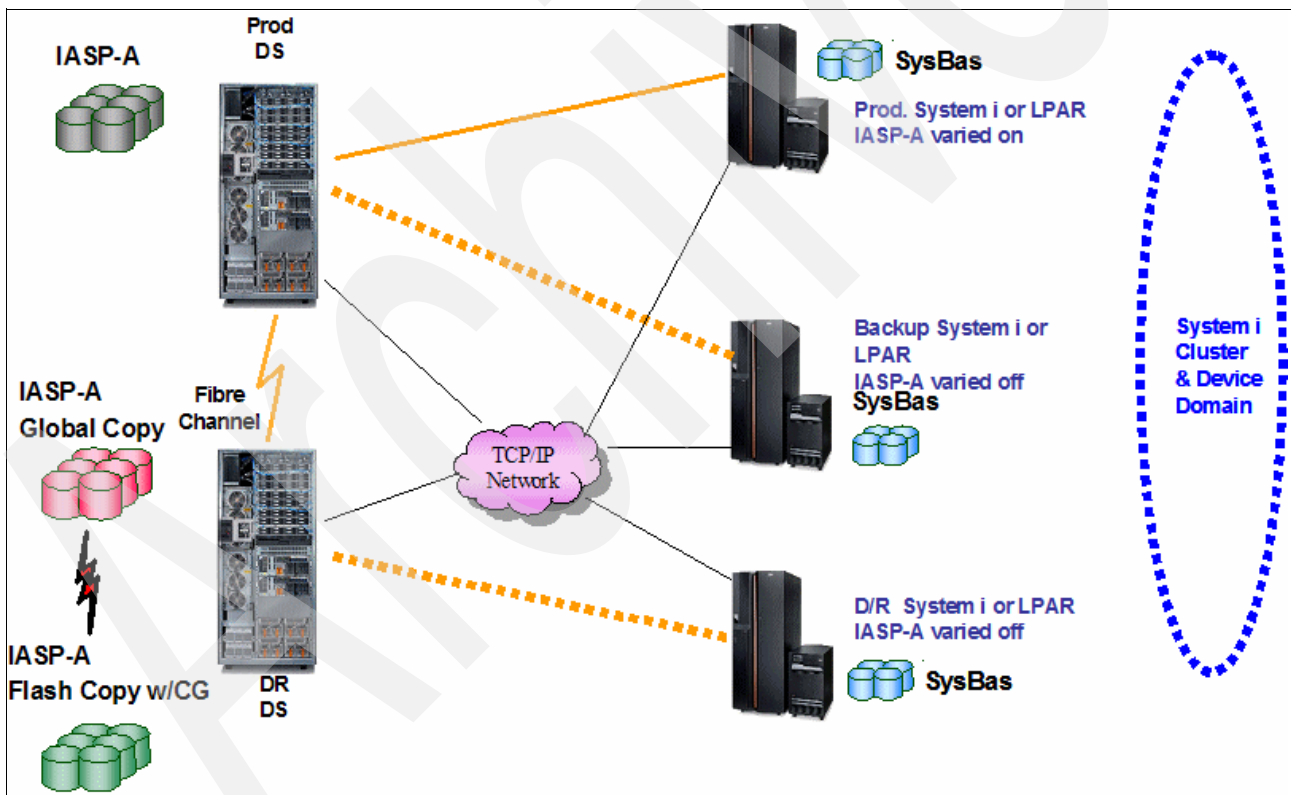


Figure 2-12 Global Mirror with asymmetrical IASP

Two primary operations can occur in this environment: switchover from production to backup and failover to backup. *Switchover from production to backup* does not involve the DS models in the previous example. It is simply a matter of running the System i Copy Services Toolkit switch PPRC (`swpprc`) command on the production system. The switch PPRC command varies off the IASP from the production system and varies it on to the backup system. Stopping the application on the production system and restarting it on the backup system

must also be considered in a switchover. This can be either a planned event where users simply log off, or it can be required for a programmatic activity to force users from the system. Both events can be automated with the use of i5/OS cluster resource group (CRG) exit programs.

The *failover to backup* configuration change is after a failure. In this case, you run the failover PPRC command (`failoverpprc`) on the backup system. Running this command allows the disaster recovery system to take over the production role, vary on the copy IASP as though it were the original, and restart the application. During vary on processing, journal recovery occurs. If the application does not use journaling, the vary on process is considerably long because the recovery process can fail due to damage and unrecoverable objects. You can restore these objects from backup tapes, but some data integrity analysis needs to occur, which can delay users who are allowed to access the application. This is similar to a disaster crash on a single system, where the same recover process needs to occur.

### Symmetrical replication

In this configuration, an additional FlashCopy consistency group is created on the source production DS model. It provides all the capabilities of asymmetrical replication, but adds the ability to do regular role swaps between the production and disaster recovery sites. When the role swaps occur with a configuration as shown in Figure 2-13, the backup system does not provide any planned switch capability for the disaster recovery site.

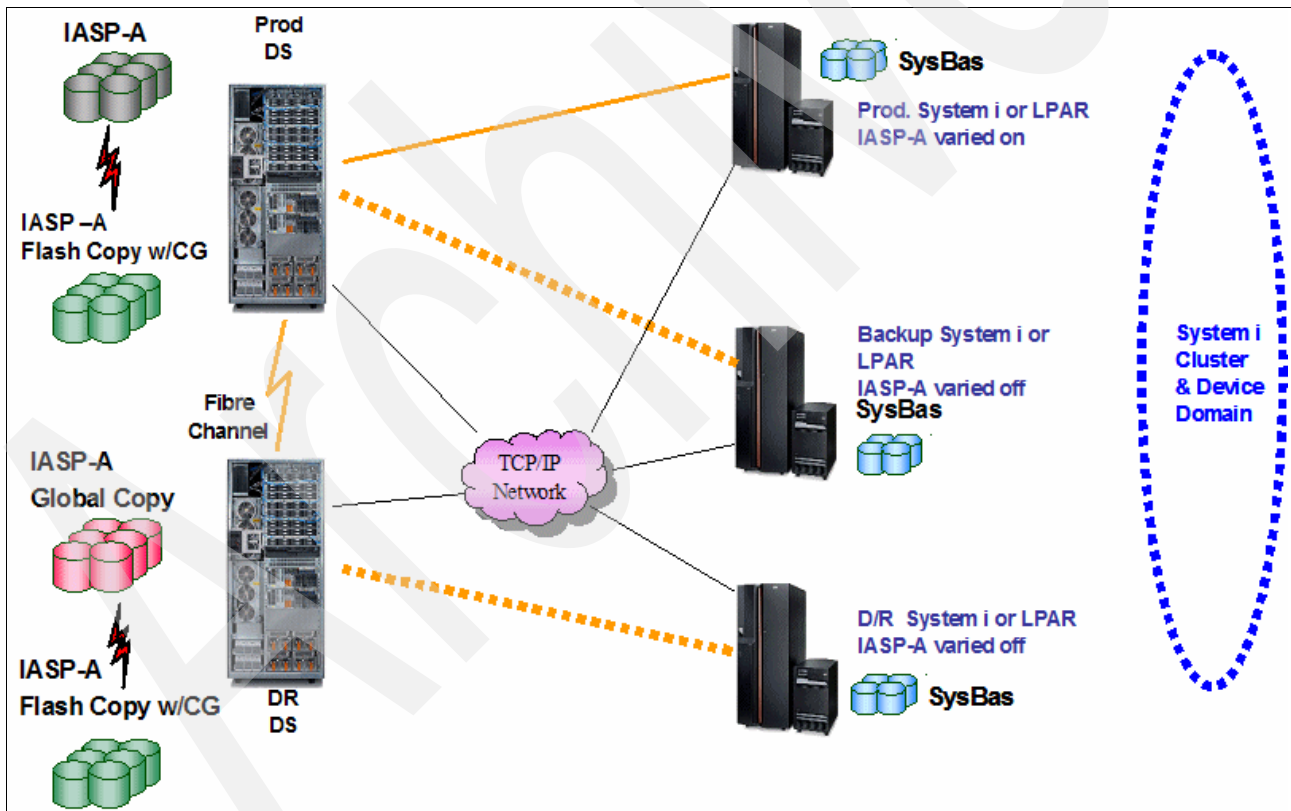


Figure 2-13 Metro Mirror symmetrical replication

In this configuration, there are multiple capabilities, local planned availability between the production and backup, and role swap or disaster recovery between the production and disaster recovery site. The planned availability switch between production and backup is the same as described in “Asymmetrical replication” on page 45, which does not involve the DS models.

If you are going to do a role swap between the production system and the disaster recovery site, you must also work with the DS models. Role swap involves the reversal of the flow of data between production DS and disaster recovery DS. While this is more complex, the tasks can be simply run from DS CLI and scripts. Either the System i Copy Services Toolkit or the i5/OS V6R1 High Availability Solutions Manager (HASM) is required for this solution. For more information about these System i Copy Services management tools, refer to 1.6.4, “System i Copy Services management solutions” on page 24.

## 2.2.4 Geographic mirroring with external storage

The geographic mirroring function of i5/OS cross-site mirroring (XSM) provides the ability to move a mirrored copy of data considerable distances from the production site. This removes the previous internal drive limitation of the copies only being separated by the maximum distance of the HSL copper or optical loops.

Figure 2-14 shows the internal drive solution for XSM. The replication between the source and target system is TCP/IP based, so considerable distance is achievable. Figure 2-14 also shows a local backup server, which enables an administrative (planned) switchover to occur if the primary system should need to be made unavailable for maintenance.

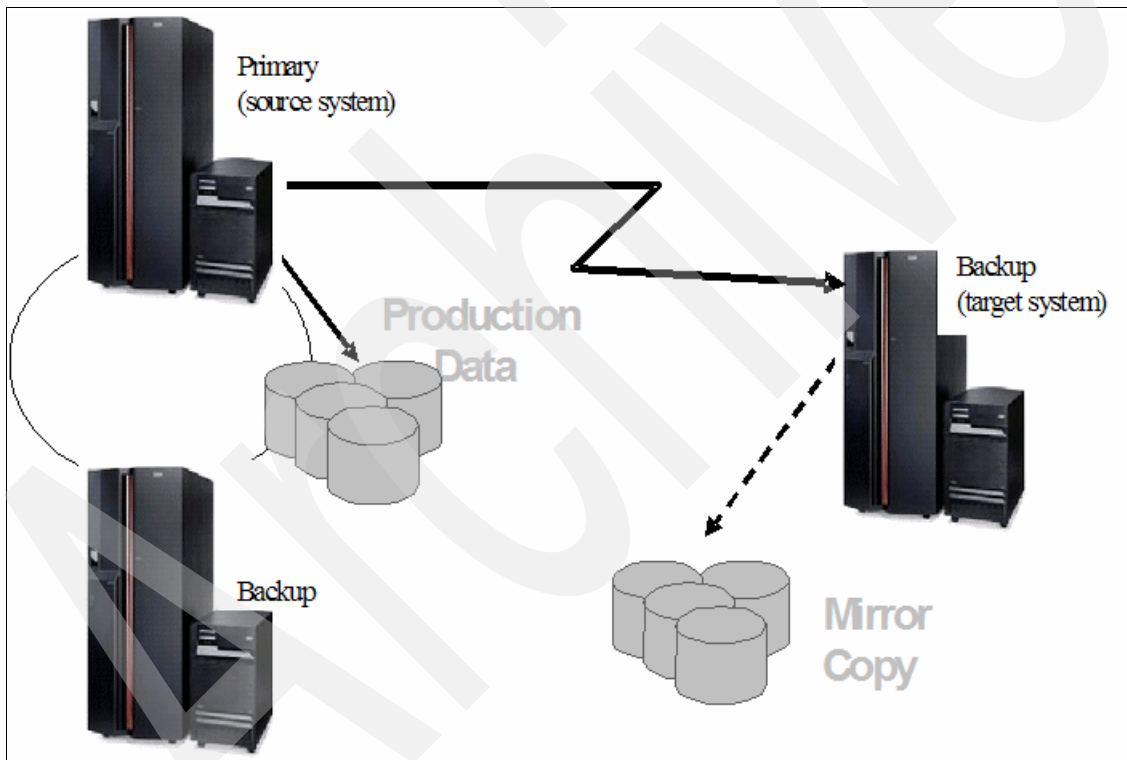


Figure 2-14 Geographic mirror with internal drives

Figure 2-15 shows a combination of the two disk technologies, with internal drives for the system ASP on each server with the IASPs located in the external storage server. In this instance, the expansion tower attached to the external disk storage becomes the switchable resource. Therefore, any I/O hardware that is needed by the source server should not be located in this tower. Functionally, this solution is the same as the internal solution.

If the load source and system base are located in the external storage system, it is possible to have all disks within the external storage system. Separation of the \*SYSBAS LUNs and the IASP LUNs and switchable tower are done at the expansion tower level.

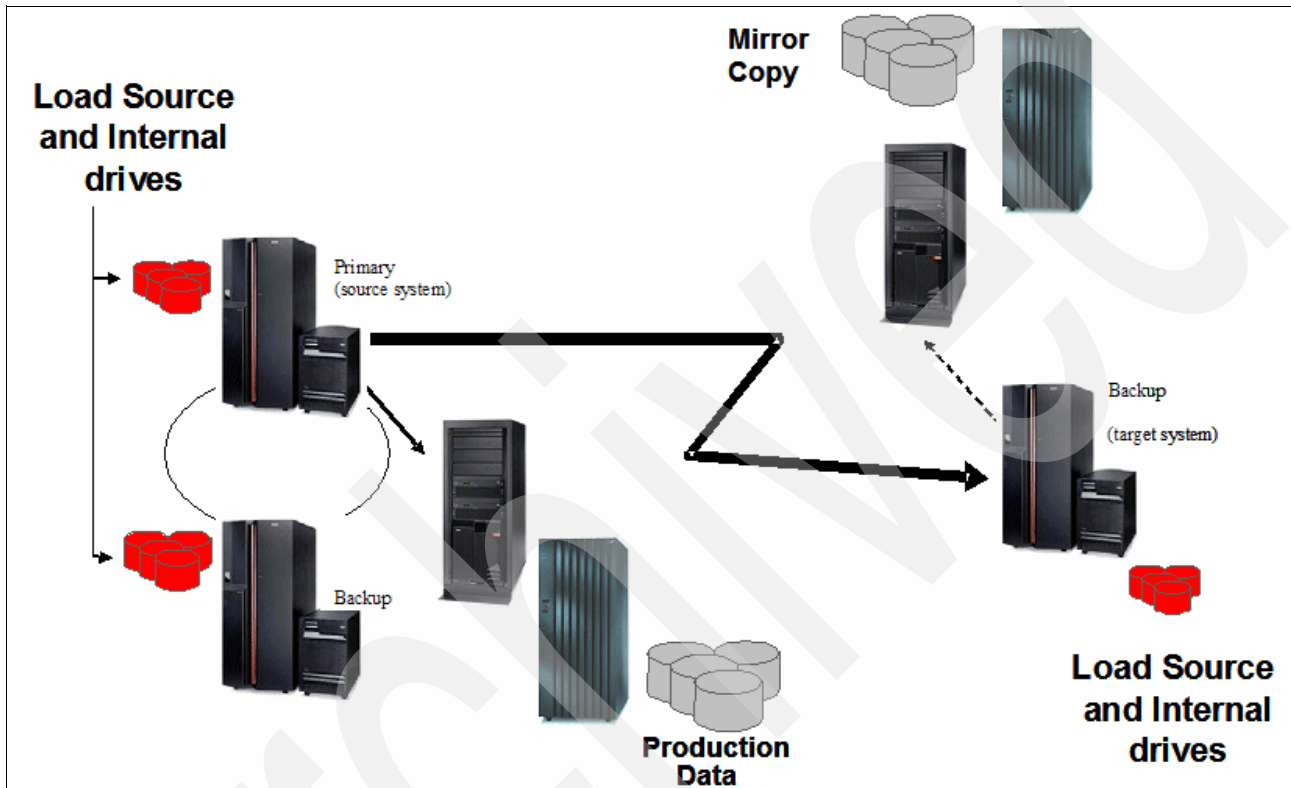


Figure 2-15 Geographic mirroring with a mix of internal and external drives



## Planning and sizing

In this part, we explain planning and sizing considerations for external storage on System i.

This part includes the following chapters:

- ▶ Chapter 3, “i5/OS planning for external storage” on page 51
- ▶ Chapter 4, “Sizing external storage for i5/OS” on page 89

Archived



## i5/OS planning for external storage

In this chapter, we discuss important planning considerations for setting up your i5/OS environment with Fibre Channel attached external IBM System Storage disk subsystems.

Good planning is essential for the successful setup and use of your server and storage subsystems. It ensures that you have met all of the prerequisites for your server and storage subsystems and everything you need to gain advantage from best practices for functionality, redundancy, performance, and availability.

Continue to use and customize the planning and implementation considerations based on your hardware setup and as recommended through the IBM Information Center documentation that is provided. Do not use the contents in this chapter as a substitute for completing your initial server setup (IBM System i or IBM System p® with i5/OS logical partitions), IBM System Storage subsystems, and configuration of the Hardware Management Console (HMC).

## 3.1 Planning for external storage solutions

To plan the implementation of ESS 800, DS6000, or DS8000 series correctly with the System i platform, you must plan for the solution that you want to implement. The solutions can vary, based on application and overall business continuity requirements. Some of these solutions are:

- ▶ Implement a SAN solution instead of integrated internal storage.
- ▶ Implement a disaster recovery solution by enabling IBM System Storage Copy Services functions for Disk Storage (DS).
- ▶ Minimize batch or backup job window using DS FlashCopy to provide a point-in-time replica of your source data.
- ▶ Implement a combination of the previous solutions to achieve disaster recovery and business continuity goals.

For example, you might want to use DS6000 or DS8000 storage for i5/OS, AIX®, and Linux, which reside on your System i servers. You might want to also implement a disaster recovery solution with Remote Mirror and Copy features, such as IBM System Storage Metro Mirror or Global Mirror, as well as plan to implement FlashCopy to minimize the backup window.

The flowchart in Figure 3-1 can assist you with the important planning steps that you need to consider based on your solution requirement. We strongly recommend that you evaluate the flow in this diagram and create the appropriate planning checklists for each of the solutions.

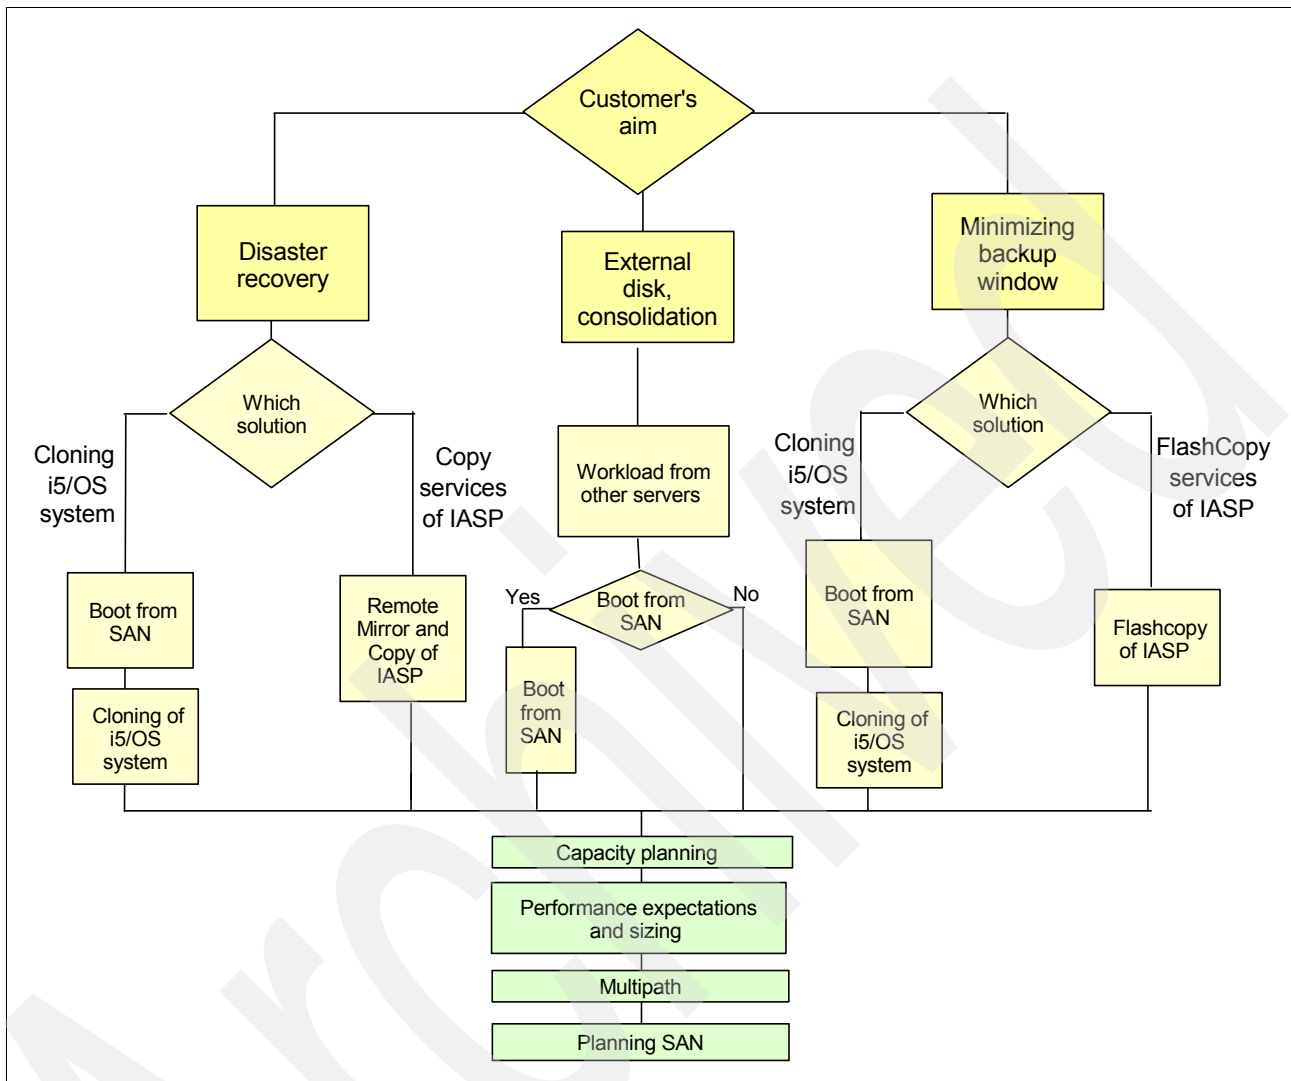


Figure 3-1 Pre-order planning

When planning for external storage solutions review the following planning considerations:

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

## 3.2 Solution implementation considerations

You need to consider multiple implementation considerations. The considerations vary based on the solution that you are trying to implement. In the following section, we highlight some of the important planning and implementation considerations in the following areas:

- ▶ Boot from SAN
- ▶ i5/OS multipath Fibre Channel attachment
- ▶ IBM System Storage Copy Services
- ▶ Storage consolidation from different servers
- ▶ SAN connectivity
- ▶ Capacity
- ▶ Performance

### 3.2.1 Planning considerations for boot from SAN

The deployment of boot support for the Fibre Channel (FC) i5/OS load source requires you to have minimum hardware and software configurations. In this section, we guide you through the required installation planning considerations when enabling boot from SAN support with the i5/OS load source being external within IBM System Storage disk subsystems.

Note that boot from SAN is required only if you are planning to externalize your i5/OS load source completely and to place all disk volumes that belong to that system or LPAR in the IBM System Storage subsystem. You might not need boot from SAN if you plan to use independent auxiliary storage pools (IASPs) with external storage, where the system objects (\*SYSBAS) could remain on System i integrated internal storage.

The new System i POWER6 IOP-less Fibre Channel cards 5749 or 5774 support boot from SAN for Fibre Channel attached IBM System Storage DS8000 models and tape drives. Refer to Table 3-1 and Table 3-2 for the minimum hardware and software requirements for IOP-less Fibre Channel and to 3.2.2, “Planning considerations for i5/OS multipath Fibre Channel attachment” on page 57 for further configuration planning information.

The 2847 I/O processor (IOP) introduced with the i5/OS V5R3M5 IOP-based Fibre Channel boot from SAN support is intended only to support boot capability for the disk unit of the FC i5/OS load source and up to 31 additional LUNs, in addition to the load source, attached using a 2766, 2787, or 5760 FC disk adapter. This IOP cannot be used as an alternate IPL device for booting from any other devices, such as a DVD-ROM, CD-ROM, or integrated internal load source. Also, the 2847 IOP cannot be used as a substitute for 2843 or 2844 IOP to drive non-FC storage, LAN, or any other System i adapters.

**Important:** The IBM Manufacturing Plant does not preload i5/OS and licensed programs on new orders or upgrades to existing System i models when the 2847 IOP is selected. You must install the system or partitions using the media that is supplied with the order, after you complete the set up the ESS 800, DS6000, or DS8000 series.

For information about more resources to assist with planning and implementation tasks, see “Related publications” on page 479.

### Minimum hardware requirements for IOP-less Fibre Channel

Table 3-3 highlights the minimum hardware configuration required for IOP-less Fibre Channel attachment of external disk storage subsystems. These attachment requirements implicitly include all the hardware requirements for boot from SAN with IOP-less Fibre Channel.

Table 3-1 Minimum IOP-less boot from SAN hardware requirements

Requirement	Complete
System i POWER6 server with PCIe or PCI-X I/O slots to support the I/O adapter (IOA) requirements	
5749 or 5774 Dual-port IOP-less Fibre Channel Disk Adapter (IOA) for attaching i5/OS storage to DS8000 series	
HMC attached to a System i POWER6 model	
IBM System Storage DS8000 series <sup>a</sup>	
Storage capacity in DS8000 series to define a LUN for the loadsource unit (with boot from SAN only)	

a. No DS6000 and no ESS models are supported for System i IOP-less Fibre Channel.

### Minimum software requirements for IOP-less Fibre Channel

When planning to use IOP-less Fibre Channel for SAN external disk storage, refer to Table 3-2 for the minimum required levels of software. These attachment requirements implicitly include all the software requirements for boot from SAN with IOP-less Fibre Channel.

Table 3-2 Minimum IOP-less boot from SAN software requirements

Requirement	Complete
i5/OS Version 6 Release 1 Modification 0 (V6R1M0)	
HMC firmware V7.3.1	
DS8000 microcode V2.4.3 However, we strongly recommend that you to install the latest level of FBM code available at the time of installation. Contact your IBM System Storage specialist for additional information.	

### Minimum hardware requirements for IOP-based boot from SAN

Table 3-3 highlights the minimum hardware configuration required to enable 2847 IOP-based boot support for an FC i5/OS load source.

Table 3-3 Minimum IOP-based boot from SAN hardware requirements

Requirement	Complete
2847 IOP for each server instance that requires a load source or for each LPAR that is enabled to boot i5/OS from Fibre Channel load source <sup>a</sup>	
When using i5/OS prior to V6R1 we recommend that the FC i5/OS load source is mirrored using i5/OS mirroring at an IOP level, with the remaining LUNs protected with i5/OS multipath I/O capabilities. For IOP-level redundancy, you need at least two 2847 IOPs and two FC adapters for each system image or LPAR.	

Requirement	Complete
System i POWER5 or POWER6 model, for POWER5 I/O slots in the system unit, expansion drawers or towers to support the IOP and I/O adapter (IOA) requirements, for POWER6 IOPs are only supported in HSL loop attached supported expansion drawers or towers	
System p models for i5/OS in an LPAR (9411-100) with I/O slots in expansion drawers or towers to support the IOP and IOA requirements	
2766, 2787, or 5760 Fibre Channel Disk Adapter (IOA) for attaching i5/OS storage to ESS 800, DS6000, or DS8000 series <sup>b</sup>	
HMC attached to a System i or p model	
IBM System Storage DS8000, DS6000 or Enterprise Storage Server® (ESS) 800 series	
A PC workstation to install DS6000 Storage Manager  The PC <i>must be</i> in the same subnet mask as the DS6000. The PC configuration must have a minimum of 700 MB disk, 512 MB of memory, and Intel® Pentium® 4 1.4 Ghz or more processor configuration.	
Additional storage capacity in ESS 800, DS6000 or DS8000 series, to define an additional LUN for mirroring the loadsource unit	

- a. The 2847 IOP is not supported on iSeries Models 8xx, any previous iSeries or AS/400 models, or any OEM hardware. Prior to i5/OS V6R1 the 2847 IOP does not support multipath for the i5/OS load source unit, but supports multipath for all other LUNs attached to this IOP.
- b. Each adapter requires a dedicated I/O processor. Use the 2847 IOP where one of the LUNs is an i5/OS load source. Use the 2844 PCI-X I/O processor for attaching additional LUNs through the 2766, 2787, or 5760 IOA. You cannot use the 2847 IOP as a substitute IOP to connect any other components.
- c. If multiple IP addresses are on the same DS6000 Storage Manager management console, the first network adapter must be on the same subnetwork as the DS6000.

### Minimum software requirements for IOP-based boot from SAN

When planning to install the 2847 IOP, you must consider several updates that need to be completed on the server, HMC, and system firmware. Table 3-4 lists the minimum levels of software that are required to enable support for the FC i5/OS load source using 2847 IOP.

Table 3-4 Minimum IOP-based boot from SAN software requirements

Requirement	Complete
i5/OS Licensed Internal Code (LIC): V5R3M5 (level RS 535-A or later)	
i5/OS Version 5 Release 3 Modification 0 (V5R3M0) Resave (level RS 530-10 or later)	
Program temporary fixes (PTFs) i5/OS and LIC: MF33328, MF33845, MF33437, MF33303, SI14550, SI14690, SI14755, or their supersedes	
System Firmware for System i or System p servers V2.3.5 or later	
HMC Firmware V5.1 or later	
DS6000 microcode: We strongly recommend that you install the latest level of Field Bill Material (FBM) code available at the time of installation. Go to the following Web page, and click <b>Downloadable files</b> to obtain more information about DS6000 microcode: <a href="http://www-03.ibm.com/servers/storage/support/disk/ds6800/downloading.html">http://www-03.ibm.com/servers/storage/support/disk/ds6800/downloading.html</a>	
DS6000 Storage Manager	



Requirement	Complete
DS8000 microcode. We strongly recommend that you to install the latest level of FBM code available at the time of installation. Contact your IBM System Storage specialist for additional information.	
ESS 800: 2.4.3.35 or later	

### 3.2.2 Planning considerations for i5/OS multipath Fibre Channel attachment

**Important:** With i5/OS V6R1, *multipath is now supported also for an external load source disk unit for both the older 2847 IOP-based and the new IOP-less Fibre Channel adapters.*

The new multipath function with i5/OS V6R1 eliminates the need with previous i5/OS V5R3M5 or V5R4 versions to mirror the external load source merely for the purpose of achieving path redundancy (see 5.10, “Protecting the external load source unit” on page 215).

Originally multipath support was added for System i external disks in V5R3 of i5/OS. Other platforms have a specific software component, such as the *Subsystem Device Driver (SDD)*. Multipath is part of the base operating system. With V5R3 and later, you can define up to eight connections from multiple I/O adapters on an iSeries or System i server to a single logical volume in the DS8000, DS6000 or ESS. Each connection for a multipath disk unit functions independently. Several connections provide redundancy by allowing disk storage to be used even if a single path fails.

Multipath is important for the System i platform because it provides greater resilience to storage area network (SAN) failures, which can be critical to i5/OS due to the single-level storage architecture. Multipath is not available for System i internal disk units, but the likelihood of path failure is much less with internal drives because there are fewer interference points. There is an increased likelihood of issues in a SAN-based I/O path because there are more potential points of failure, such as long fiber cables and SAN switches. There is also an increased possibility of human error occurring when performing such tasks as configuring switches, configuring external storage, or applying concurrent maintenance on DS6000 or ESS, which might make some I/O paths temporarily unavailable.

Many System i customers still have their entire environment on the system or user auxiliary storage pools (ASPs). Loss of access to any disk causes the system to enter a freeze state until the disk access problem gets resolved. Even a loss of a user ASP disk will eventually cause the system to stop. Independent ASPs (IASPs) provide isolation so that loss of disks in the IASP only affect users who access that IASP while the remainder of the system is unaffected. However, with multipath, even loss of a path to disk in an IASP will not cause an outage.

Prior to multipath, some customers used i5/OS mirroring to two sets of disks, either in the same or different external disk subsystems. This mirroring provided implicit dual path as long as the mirrored copy was connected to a different IOP or I/O adapter (IOA), bus, or I/O tower. However, this mirroring also required twice as much capacity for two copies of data. Because disk failure protection is already provided by RAID-5 or RAID-10 in the external disk subsystem, this was sometimes considered unnecessary.

With the combination of multipath and RAID-5 or RAID-10 protection in DS8000, DS6000, or ESS, you can provide full protection of the data paths and the data itself without the requirement for additional disks.

## Avoiding single points of failure

Figure 3-2 shows 15 single points of failure, excluding the System i model itself and the disk subsystem storage facility. Failure points 9-12 are not present if you do not use an *inter-switch link* (ISL) to extend your SAN. An outage to any one of these components (either planned or unplanned) causes the system to fail if IASPs are not used or causes the applications within an IASP to fail if IASPs are used.




Figure 3-2 Single points of failure

When implementing multipath, provide as much redundancy as possible. At a minimum, multipath requires two IOAs that connect the same logical volumes. Ideally, these should be on different buses, in different I/O racks in the System i environment, and if possible, on different high-speed link (HSL) or 12X loops. If a SAN is included, use separate switches in two different fabrics for each path. You should also use host adapters in different I/O drawer pairs in the DS6000 or DS8000 as shown in Figure 3-3.

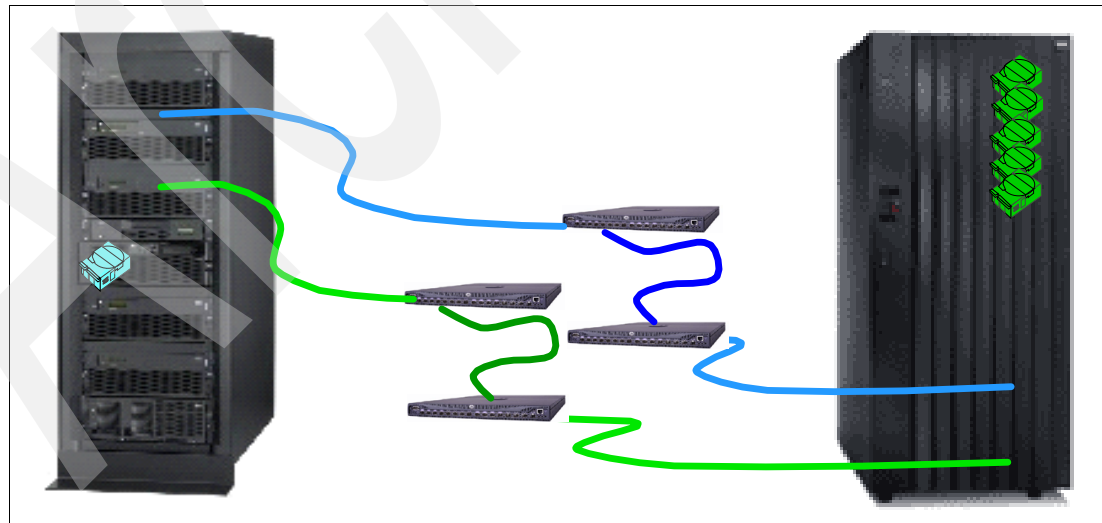


Figure 3-3 Multipath removes single points of failure

Unlike other systems that might support only two paths (dual-path), i5/OS V5R3 supports up to eight paths to the same logical volumes. At a minimum, you should use two paths, although some small performance benefits might be experienced with more paths. However, because i5/OS multipath spreads I/O across all available paths in a *round-robin* manner, there is no *load balancing*, only *load sharing*.

### Configuration planning

The System i platform has three IOP-based Fibre Channel I/O adapters that support DS8000, DS6000, and ESS model 800:

- ▶ FC 5760 / CCIN 280E 4 Gigabit Fibre Channel Disk Controller PCI-X
- ▶ FC 2787 / CCIN 2787 2 Gigabit Fibre Channel Disk Controller PCI-X (*withdrawn from marketing*)
- ▶ FC 2766 / CCIN 2766 2 Gigabit Fibre Channel Disk Controller PCI (*withdrawn from marketing*)

The following new System i POWER6 IOP-less Fibre Channel I/O adapters support DS8000 as external disk storage only:

- ▶ FC 5749 / CCIN 576B 4 Gigabit Dual-Port IOP-less Fibre Channel Controller PCI-X (see Figure 3-4)
- ▶ FC 5774 / CCIN 5774 4 Gigabit Dual-Port IOP-less Fibre Channel Controller PCIe (see Figure 3-5)

**Note:** The 5749/5774 IOP-less FC adapters are supported with System i POWER6 and i5/OS V6R1 or later only. They support both Fibre Channel attached disk and tape devices on the same adapter but not on the same port. As a new feature these adapters support D-mode IPL boot from a tape drive which should be either direct attached or, by proper SAN zoning, the only tape drive seen by the adapter. Otherwise, with multiple tape drives seen by the adapter, it picks only the first one that reported in and is loaded, and if it contains no valid IPL source, the IPL fails.




Figure 3-4 New 5749 IOP-less PCI-X Fibre Channel Disk Controller




Figure 3-5 New 5774 PCIe IOP-less Fibre Channel Disk Controller

**Important:** For *direct* attachment, that is point-to-point topology connections using no SAN switch, the IOP-less Fibre Channel adapters support only the Fibre Channel arbitrated loop (FC-AL) protocol. This support is different to the previous 2847 IOP-based FC adapters, which supported only the Fibre Channel switched-fabric (FC-SW) protocol, whether direct- or switch-connected, although other 2843 or 2844 IOP-based FC adapters support either FC-SW or FC-AL.

All these System i Fibre Channel I/O adapters can be used for multipath.

**Important:** Though there is no requirement for all paths of a multipath disk unit group to use the same type of adapter we strongly recommend to avoid mixing IOP-based and IOP-less FC I/O adapters within the same multipath group. In a multipath group with mixed IOP-based and IOP-less adapters the IOP-less adapter performance would be throttled by the lower performance IOP-based adapter due to the I/O being distributed by a round-robin algorithm across all paths of a multipath group.

The IOP-based single-port adapters can address up to 32 logical units (LUNs) while the dual-port IOP-less adapters support up to 64 LUNs per port.

Table 3-5 summarizes the key differences between IOP-based and IOP-less Fibre Channel.

Table 3-5 Key differences between IOP-based versus IOP-less Fibre Channel

Function	IOP-based	IOP-less
System i support	All models (#2847 requires POWER5 or later)	POWER6 models only
A / B mode IPL (boot from SAN)	Yes (with #2847 IOP only)	Yes
D mode IPL (from tape)	No	Yes

Function	IOP-based	IOP-less
<i>Direct-attach</i> protocol	FC-AL or FC-SW (FC-SW only for boot from SAN)	FC-AL
Disk LUNs per port	32	64
Max. concurrent I/Os	1	6
ESS (2105) support	Yes (#2847 supports ESS model 800 only)	No
DS6000 support	Yes	No
DS8000 support	Yes	Yes
Multipath Load Source	Yes (with V6R1)	Yes (with V6R1)

The System i i5/OS multipath implementation requires each path of a multipath group to be connected to a separate System i I/O adapter to be utilized as an active path. Attaching a System i I/O adapter to a switch and going from the switch to two different storage subsystem ports results in only one of the two paths between the switch and the storage subsystem being used with the second path only being used in case of a failure of the first one, sometimes referred to as *backup-link* and used to be a solution for higher redundancy with ESS external storage before i5/OS multipathing became available.

It is important to plan for multipath so that the two or more paths to the same set of LUNs use different hardware elements of connection, such as storage subsystem host adapters, SAN switches, System i I/O towers, and high-speed link (HSL) or 12X loops.

Good planning for multipath includes:

- ▶ Connections to the same set of LUNs through different DS host cards on DS8000
- ▶ Connections to the same set of LUNs through host adapters on different processors of DS6000
- ▶ Connections to the same set of LUNs through different SAN switches resp. fabrics
- ▶ Connections to the same set of LUNs on physically different IOA adapters, that is not using multipath to the same set of LUNs through the same dual-port IOP-less IOA adapter
- ▶ Placement of the IOA adapter pairs in the System i I/O tower that connects to the same set of LUNs, in different expansion towers and HSL/12X loops wherever possible

When deciding how many I/O adapters to use, your first priority should be to consider performance throughput of the IOA because this limit can be reached before the maximum number of logical units. See Chapter 4, “Sizing external storage for i5/OS” on page 89, for more information about sizing and performance guidelines.

For more information about implementing multipath, see Chapter 5, “Implementing external storage with i5/OS” on page 181.

### **Multipath rules for multiple System i models or partitions**

When you use multipath disk units, you must consider the implications of moving IOPs or IOAs with multipath connections between nodes. You must not split multipath connections between nodes, either by moving IOPs/IOAs between logical partitions (LPARs) or by switching expansion units between systems. If two different nodes both have connections to the same logical unit number (LUN) in the IBM System Storage disk subsystem, both nodes might potentially overwrite data from the other node.

System i single-level storage requires you to adhere to the following rules when you use multipath disk units in a multiple-system environment:

- ▶ If you move an IOP with a multipath connection to a different LPAR, you must also move all other IOPs with connections to the same disk unit to the same LPAR.
- ▶ When you make an expansion unit switchable, make sure that all multipath connections to a disk unit switch with the expansion unit.
- ▶ When you configure a switchable independent disk pool, make sure that all of the required IOPs for multipath disk units switch with the independent disk pool.

If a multipath configuration rule is violated, the system issues warnings or errors to alert you of the condition. It is important to pay attention when disk unit connections are reported missing. You want to prevent a situation where a node might overwrite data on a LUN that belongs to another node.

Disk unit connections might be missing for a variety of reasons, but especially if one of the preceding rules has been violated. If a connection for a multipath disk unit in any disk pool is found to be missing during an IPL or vary on, a message is sent to the QSYSOPR message queue.

If a connection is missing, and you confirm that the connection has been removed, you can update Hardware Service Manager to remove that resource. Hardware Service Manager is a tool to display and work with system hardware from both a logical and a packaging viewpoint, an aid for debugging IOPs and devices, and for fixing failing and missing hardware. You can access Hardware Service Manager in System Service Tools (SST) and Dedicated Service Tools (DST) by selecting the option to start a service tool.

### 3.2.3 Planning considerations for Copy Services

In this section, we discuss important planning considerations for implementing IBM System Storage Copy Services solutions for i5/OS.

#### FlashCopy storage configuration considerations

When planning for a FlashCopy implementation the following configuration guidelines from a DS storage system perspective help you to achieve good performance and smooth operations:

- ▶ Configure the FlashCopy source and target volumes within the same rankgroup, that is avoid cross-cluster FlashCopy relationships with the source being on an even LSS and the target on an odd LSS or vice versa.
- ▶ Use the same disk speed (preferably 15 KB RPM drives) for both source and target volumes.
- ▶ Use FlashCopy with the no-background copy (no-copy) option, instead of the default full-copy option, for short-lived FlashCopy relationships, such as for system backup to tape to limit the performance impact to your production system.
- ▶ If for the duration of the FlashCopy relationship your host write I/O workload causes more than about 20% of the data to be changed refrain from using space efficient FlashCopy and use regular FlashCopy instead.
- ▶ When planning to use the new DS8000 R3 space efficient FlashCopy function carefully size the storage space for your repository volumes to prevent running out of space causing the relationship to fail (see 4.2.8, “Sizing for space efficient FlashCopy” on page 104).

**Note:** The first release of space efficient FlashCopy with DS8000 R3 does not allow you to increase the repository capacity dynamically. That is, to increase the capacity, you will need to delete the repository storage space and re-create it with more physical capacity.

- ▶ For better space efficient FlashCopy write performance, you might consider using RAID10 for the target volumes as the writes to the shared repository volumes always have random I/O character (see 4.2.8, “Sizing for space efficient FlashCopy” on page 104).

## Planning for FlashCopy with i5/OS

The FlashCopy Copy Services function of IBM System Storage DS8000, DS6000 or ESS essentially enables you to create an i5/OS system image as an identical point-in-time replica of your entire storage space. This capability has become more realistic since i5/OS V5R3M5 with the advent of being able to place an i5/OS load source directly in a SAN storage subsystem.

By using FlashCopy for creating a duplicate i5/OS system image of your production system and IPLing another i5/OS LPAR from it running the backup to tape, you can increase the availability of your production system by reducing or eliminating down-times for system saves. FlashCopy can also assist you with having a backup image of your entire system configuration to which you can rollback easily in the event of a failure during a release migration or a major application upgrade.

**Important:** The new i5/OS V6R1 quiesce for Copy Services function (see 15.1, “Using i5/OS quiesce for Copy Services” on page 432) helps to ensure that you have modified data residing in main memory written to disks prior to creating an i5/OS image with FlashCopy. For i5/OS versions prior to V6R1, we recommend that you shut down the system completely (PWRDWN SYS) before you initiate a full-system FlashCopy. Ending subsystems or bringing the system to a restricted state does not guarantee that *all* contents of main storage will be written to the disk.

Keep in mind that creating an i5/OS image through FlashCopy is a point-in-time instance and thus should be used only for recovery of the production system only as a full backup for the production system image. Many of the objects, such as history logs, journal receivers, and journals, have different data history reflected in them and must not be restored to the production system.

You must not attach any copied LUNs to the original parent system unless they have been used on another partition first or initialized within the IBM System Storage subsystems. Failure to observe this restriction will have unpredictable results and can lead to loss of data. This is due to the fact that the copied LUNs are perfect copies of LUNs that are on the parent system. As such, the system would not be able to tell the difference between the original and the cloned LUN if they were attached to the same system.

As soon as you copy an i5/OS image, attach it to a separate partition that will own the LUNs that are associated with the copied image. By doing this, you make them safe to be reused again on the parent partition.

When planning to implement FlashCopy or Remote Mirror and Copy functions such as Metro Mirror and Global Mirror for copying of an i5/OS system consider the following points:

- ▶ Storage system licenses for use of Copy Services functions are required.
- ▶ Have a sizing exercise completed to ensure that your system and storage configuration is capable of handling the additional I/O requests. You also need to account for additional

memory, I/O, and disk storage requirements in the storage subsystem in addition to hardware resources at the system side.

- ▶ Ensure that the recovery system or backup partition is configured for boot from SAN to IPL from the copied i5/OS load source.
- ▶ Sufficient capacity (processor, memory, I/O towers, IOPs, IOAs, and storage) is reserved to bring up the target environment, either in an LPAR or on a separate system that is locally available in the same data center complex.
- ▶ When restarting the environment after attaching the copied LUNs, it is important to understand that, because these are identical copies of the LUNs in your production environment, all of the attributes that are unique to your production environment are also copied, such as network attributes, resource configuration, and system names. It is important that you perform a manual IPL when you first start the system or partition so that you can change the configuration attributes before the system fully starts. Examples of the changes that you need to perform are:
  - System Name, Local Location Name, and Default Location Name  
You need to change these attributes before you restart SNA or APPC communications, or prior to using BRMS.

**Tip:** You might want to create a “backup” startup program that you invoke during the restart of a cloned i5/OS image so that you can automate many of the configuration attribute changes that otherwise need manual intervention.

- TCPI/IP network attributes  
You need to reassign a new IP address for the new system and reconfigure any related attributes before the cloned image is added to the network, either for performing a full system save or for performing any read-only operations such as database queries or report printing.
- System name in the relational database directory entry  
You might need to update this entry using the WRKRDBDIRE command before you start any database activities that rely on these attributes.
- ▶ The hardware resource configuration will not match what is on the production system and needs to be updated prior to starting any network or tape connectivity.
- ▶ Remember that any jobs in the job queue will still be there, and any scheduled entries in the job scheduler will also be there. You might want to clear job queues or hold the job scheduler on the backup server to avoid any updates to the files, enabling you to have a true point-in-time instance of your production server.
- ▶ You must understand the usage of BRMS when saving from a FlashCopy image of your production system (see “Using Backup Recovery and Media Services with FlashCopy” on page 64.)

### **Using Backup Recovery and Media Services with FlashCopy**

In addition to the planning considerations discussed in the previous section, here we provide additional considerations for using BRMS with FlashCopy for which you need to plan:

- ▶ Enable BRMS on the production system to allow use of FlashCopy on the backup system.
- ▶ During the save operation on the backup machine (after FlashCopy has completed, and you have completed all of the IPL steps, including changes to the network attributes and system attributes), ensure that no backups are conducted on the production system. BRMS treats your backup as a point-in-time instance of your production system and



maintains the BRMS network and media information across other systems that share the media inventory.

- ▶ After BRMS has completed the save operation, complete the post backup options such as taking a full save of your QUSRBRM library and restoring it on the production system. You can do this by using either a tape drive or FTP to transfer the save file. This step is required to ensure that the BRMS management and media information is transferred back to the production system before you reuse the disk space associated with the FlashCopy instance. The restore of QUSRBRM back on the production system provides an accurate picture of the BRMS environment on the production system, which reflects the backups that were just performed on the clone system.
- ▶ After QUSRBRM is restored, indicate on the production system that the BRMS FlashCopy function is complete.

**Important:** If you have to restore your application data or libraries back on the production system, do not restore any journal receivers that are associated with that library. Use the OMTOBJ parameter during the restore library operation.

- ▶ BRMS for V5R3 has been enhanced to support FlashCopy by adding more options that can be initiated prior to starting the FlashCopy operation.

For more information about using BRMS with FlashCopy, see Chapter 15, “FlashCopy usage considerations” on page 431

### Planning for Remote Mirror and Copy with i5/OS

The Remote Mirror and Copy feature (formerly *PPRC*) copies data between volumes on two or more storage units. When your host system performs I/O update operations to the source volume, they are copied or mirrored to the target volume automatically. After you create a remote mirror and copy relationship between a source volume and target volume, the target volume continues to be updated with changes from the source volume until you remove the relationship between the volumes.

Note the following considerations when planning for Remote Mirror and Copy:

- ▶ Determine the recovery point objective (RPO) for your business and clearly understand the differences between synchronous storage-based data replication with Metro Mirror and asynchronous replication with Global Mirror, and Global Copy.
- ▶ When planning for a synchronous Metro Mirror solution, be aware of the maximum supported distance of 300 km and expect a delay of your write I/O of around 1 ms per 100 km distance.
- ▶ Have a sizing exercise completed to ensure that your system and storage configuration is capable of handling additional I/O requests, that your I/O performance expectations are met and that your network bandwidth supports your data replication traffic to meet your recovery point objective (RPO) targets.
- ▶ Acquire storage system licenses for the Copy Services functions to be implemented.
- ▶ Unless you are not replicating IASPs only, configure your System i production system and target system with boot from SAN for faster recovery times.
- ▶ Sufficient capacity (processor, memory, I/O towers, IOPs, IOAs, and storage) is reserved to bring up the target environment, either in an LPAR or on a separate system that is locally available in the same data center complex.

## Planning Remote Mirror and Copy with i5/OS IASPs

This solution involves the replication of data at the storage controller level to a second storage server using IBM System Storage DS8000, DS6000 or ESS. An independent auxiliary storage pool (IASP) is the basic unit of System i storage you can replicate using Copy Services. Using IASPs with Copy Services is supported by using a management tool such as the new i5/OS V6R1 High Availability Solutions Manager (HASM) licensed program product (5761-HAS) or the System i Copy Services Toolkit services offering.

One of the biggest advantages of using IASP is that you do not need to shut down the production server for switching over to your recovery system. A vary off of the IASP ensures that data is written to the disk prior to initiating a switchover. HASM or the toolkit enables you to attach the second copy to a backup server without an IPL. Replication of IASPs only instead of your whole System i storage space can also help you to reduce your network bandwidth requirements for data replication by excluding write I/O to temporary objects in \*SYSBAS. You also have the ability to combine this solution with other functions, such as FlashCopy, for additional benefits such as save window reduction.

Note the following considerations when planning for Remote Mirror and Copy of IASP:

- ▶ Complete the feasibility study for enabling your applications to take advantage of IASP. For the latest information about high availability and resources on IASP, refer to the System i high availability Web site:  
<http://www.ibm.com/eserver/iseriess/ha>
- ▶ Ensure that you have i5/OS 5722-SS1 option 41 - Switchable resources installed on your system and that you have set up an IASP environment.
- ▶ Keep in mind that journaling your database files is still required, even when your data is residing in an IASP.
- ▶ Objects that reside in \*SYSBAS, that is the disk space that is not IASP must be maintained at equal levels on both the production and backup systems. You can do this by using the software solutions offered by one of the High Availability Business Partners (HABPs).
- ▶ Set up IASPs and install your applications in IASP. After the application is prepared to run in an IASP and is tested, implement HASM or the System i Copy Services Toolkit, which is provided as a service offering from IBM STG lab services.

### 3.2.4 Planning storage consolidation from different servers

In this section, we refer only to the DS8000 or DS6000 because they are the new storage offerings. The assumptions are that the new consolidations will primarily happen by using the latest IBM offerings in the areas of IBM System Storage Disk Storage subsystems.

- ▶ Make sure that the DS8000 or DS6000 series is properly sized for open systems. For sizing of the i5/OS, refer to Chapter 4, "Sizing external storage for i5/OS" on page 89.
- ▶ For the DS8000 or DS6000 series with which you plan to share the i5/OS production workload and other servers, we recommend for performance reasons that you dedicate RAID ranks to i5/OS. Therefore, plan to allocate sufficient ranks for i5/OS workloads and separate them from being shared by other open systems.
- ▶ Ensure that the host ports that are to be shared between i5/OS and other open systems are sized adequately to account for combined I/O rates driven by all of the hosts.
- ▶ Understand which systems need disaster recovery solution and plan use of Remote Mirror and Copy functions accordingly.

- ▶ Understand which systems need FlashCopy and plan capacity for FlashCopy accordingly.
- ▶ Plan the multipath attachment optimized for high redundancy, refer to “Avoiding single points of failure” on page 58

### 3.2.5 Planning for SAN connectivity

When planning for System i SAN attached storage keep the following considerations in mind:

- ▶ Ensure that the FC switches are supported by your combination of IBM System Storage disk subsystem and System i model prior to ordering the SAN fabric. The best way to determine if a given SAN switch or director is supported is to check the System Storage Interoperation Center (SSIC) at:

<http://www-01.ibm.com/servers/storage/support/config/ess/index.jsp>

- ▶ We usually zone the switches so that multiple i5/OS FC adapters are in a zone with one storage subsystem host port.

**Note:** Avoid putting more than one storage subsystem host port into a switch zone with System i FC adapters. At any given time a System i FC adapter uses only one of the available storage ports in the switch zone whichever reports in first. A slack configuration of the SAN switch with multiple System i FC adapters having access to multiple storage ports can result in performance degradation by an excessive number of System i FC adapters accidentally sharing the same link to the storage port.

Refer to Chapter 4, “Sizing external storage for i5/OS” on page 89 for recommendations on the numbers of FC adapters per host ports.

- ▶ If the IBM System Storage disk subsystem is connected remotely to a System i host, or if local and remote storage subsystems are connected using SAN, plan for enough FC links to meet the I/O requirement of your workload.
- ▶ If extenders or dense wavelength division multiplexing (DWDMs) are used for remote connection, take into account their expected latency when planning for performance.
- ▶ If FC over IP is planned for remote connection, carefully plan for the IP bandwidth.

### 3.2.6 Planning for capacity

When planning the capacity of external disk storage for System i environments, ensure that you understand the difference between the following three capacity terms of the DS8000 and DS6000 series:

- ▶ Raw capacity
- ▶ Effective capacity
- ▶ Capacity usable for i5/OS

In this section, we explain these capacity terms and highlight the differences between them.

#### Raw capacity

Raw capacity of a DS, also referred to as *physical capacity*, is the capacity of all physical disk drives in a DS system including the spare drives. When calculating raw capacity, we do not take into account any capacity that is needed for parity information of RAID protection. We

simply multiply the number of disk drive modules (DDMs) by their capacity. Consider the example where a DS8000 has five disk drive enclosures (each enclosure has 16 disk drives) of 73 GB. Thus, the DDMs have 5.84 TB of raw capacity based on the following equation:

$$5 \times 16 \times 73 \text{ GB} = 5840 \text{ GB, which is 5.84 TB of raw capacity}$$

### Spare disk drives

In order to know effective capacity of a DS, you need to understand the rule for how spare disks are assigned. This rule is called the *sparing rule*.

### Device adapters in DS8000

A DS8100 (2-way processor) can contain up to four device adapter (DA) pairs. They are connected in the order 2, 0, 3, and 1, as we explain next. DA 2 is first used to connect the arrays. It is filled with arrays until it connects eight arrays (four array sites or two enclosures) in the base frame. Then DA pair 0 is used until it connects eight arrays in the base frame. If the expansion frame is present, DA 3 is used until it is filled with eight arrays in the expansion frame. Then DA 1 is used until it is filled with eight arrays in the expansion frame. If there are more arrays in the expansion frame, DA 2 is used again to connect them until it is filled with eight arrays in the expansion frame. Then DA 0 is used again.

Figure 3-6 illustrates this method.

**Note:** Figure 3-6 shows only front disk enclosures. However, there are actually as many back enclosures, that is up to eight enclosures per base frame and up to 16 enclosures per expansion frame.

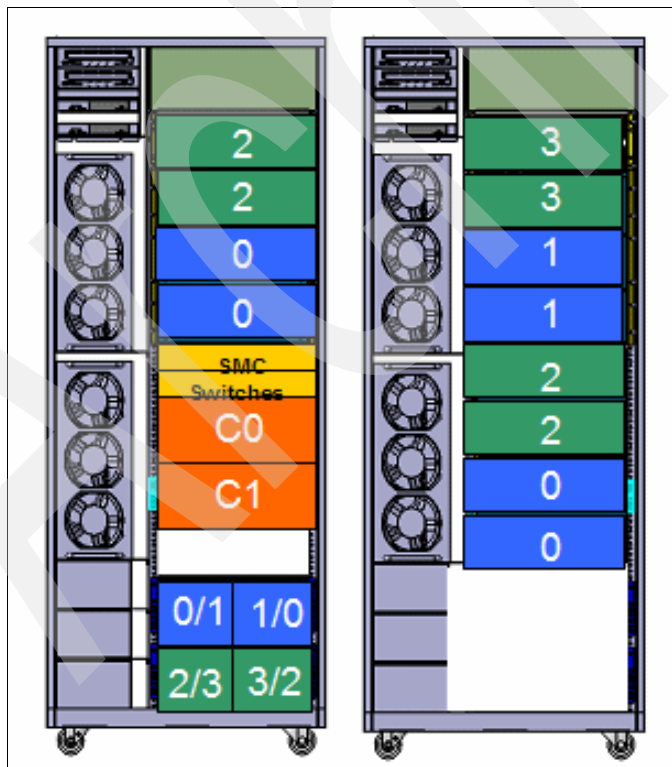


Figure 3-6 DS8100 device adapters

In DS8300 (4-way processors), there can be up to eight DA pairs. The pairs are connected in the following order: 2, 0, 4, 6, 7, 5, 3, and 1. They are connected to arrays in the same way as described for DS8100. All the DA pairs are filled with arrays until eight arrays per DA pair are reached. DA pair 0 and 2 are used for more than eight arrays if needed.

Figure 3-7 shows this method.

**Note:** Figure 3-7 shows only the front disk enclosures. However, there are actually as many back enclosures, that is up to eight enclosures per base frame and up to 16 enclosures per expansion frame.




Figure 3-7 Device adapters in DS8300

### Spares in DS8000

In DS8000, a minimum of one spare is required for each array site (or array) until the following conditions are met:

- ▶ Minimum of four spares per DA pair
- ▶ Minimum of four spares of the largest capacity array site on the DA pair
- ▶ Minimum of two spares of capacity and an RPM greater than or equal to the fastest array site of any given capacity on the DA pair

Knowing the rule of how DA pairs are used, we can determine the number of spares that are needed in a DS configuration and which RAID arrays will have a spare. If there are DDMs of a different size, more work is needed to calculate which arrays will have spares.

Consider the same example for which we calculate raw capacity in “Raw capacity” on page 67, and now calculate the spares. DS8100 with 10 array sites (10 arrays) of 73 GB DDMs, all of the arrays are RAID-5. Eight arrays are connected to DA pair 2. The first four arrays have a spare (6+P arrays) to fulfill the rule minimum of four spares for DA pair. The next four arrays on this DA pair are without a spare (7+P arrays). Two arrays are connected to DA pair 0. Both of them have a spare (6+P arrays) to fulfill the rule minimum of one spare per array site (array). Therefore in this DS configuration, there are six 6+P+S ranks and four 7+P ranks.

Figure 3-8 illustrates this example, which is a result of the DS CLI command `lsarray`.

Array	State	Data	RAIDtype	arsite	Rank	DA Pair	DDMcap (Decimal GB)
A0	Assigned	Normal	5 (6+P)	S1	R0	0	73.0
A1	Assigned	Normal	5 (6+P)	S2	R1	0	73.0
A2	Assigned	Normal	5 (6+P)	S3	R2	2	73.0
A3	Assigned	Normal	5 (6+P)	S4	R3	2	73.0
A4	Assigned	Normal	5 (6+P)	S5	R4	2	73.0
A5	Assigned	Normal	5 (6+P)	S6	R5	2	73.0
A6	Assigned	Normal	5 (7+P)	S7	R6	2	73.0
A7	Assigned	Normal	5 (7+P)	S8	R7	2	73.0
A8	Assigned	Normal	5 (7+P)	S9	R8	2	73.0
A9	Assigned	Normal	5 (7+P)	S10	R9	2	73.0

Figure 3-8 Sparing rule for DS8000

### Spares in DS6000

DS6000 has two device adapters or one device adapter pair that is used to connect disk drives in two FC loops, as shown in Figure 3-9 and Figure 3-10. In DS6000, a minimum of one spare is required for each array site until the following conditions are met:

- ▶ Minimum of two spares on each FC loop
- ▶ Minimum of two spares of the largest capacity array site on the FC loop
- ▶ Minimum of two spares of capacity and rpm greater than or equal to the fastest array site of any given capacity on the DA pair

Therefore, if only a single RAID-5 array is configured, then one spare is in the server enclosure. If two RAID-5 arrays are configured, two spares are present in the enclosure as shown in Figure 3-9. This figure shows the first expansion enclosure and its location on the second FC loop, which is separate from the server enclosure FC loop. Therefore the same sparing rules apply. That is, if the expansion enclosure has only one RAID-5 array, there is one spare. If two RAID arrays are configured in the expansion enclosure, then two spares are present.




Figure 3-9 DS6000 spares for RAID-5

Figure 3-10 shows an example of spares in DS6000 with RAID-10 arrays.

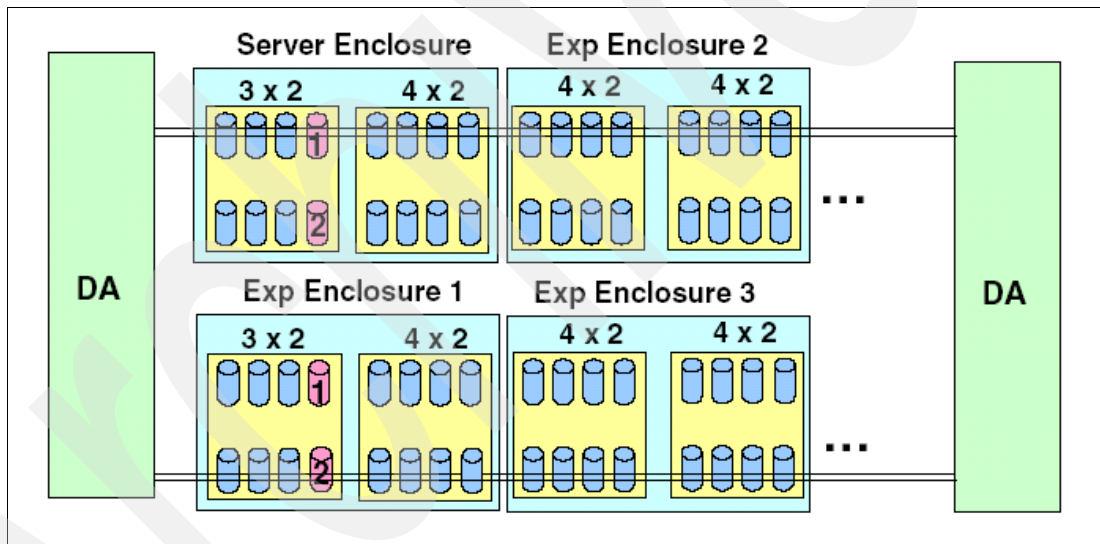


Figure 3-10 DS6000 spares for RAID-10

### Effective capacity

Effective capacity of a DS system is the amount of storage capacity that is available for the host system after the logical configuration of DS has been completed. However, the actual capacity that is visible by i5/OS is smaller than the effective capacity. Therefore, we discuss the actual usable capacity for i5/OS in “i5/OS LUNs and usable capacity for i5/OS” on page 73.

Effective capacity of a rank depends on the number of spare disks in the corresponding array and on the type of RAID protection of the array. When calculating effective capacity of a rank, we take into account the capacity of the spare disk, the capacity needed for RAID parity, the and capacity needed for metadata, which internally describes the logical to physical volume mapping. Also, effective capacity of a rank depends on the type of rank, either CKD or fixed block. Because i5/OS uses fixed block ranks, we limit our discussion to these ranks.

Table 3-6 shows the effective capacities of fixed block RAID ranks in DS8000 in decimal GB and binary GB. It also shows the number of extents that are created from a rank.

*Table 3-6 DS8000 RAID rank effective capacities*

RAID type	DDM cap.	Array form.	Extents	Binary GB	Decimal GB
RAID-5	73 GB	6+P+S	386	386	414.46
RAID-5	73 GB	7+P	450	450	483.18
RAID-5	146 GB	6+P+S	779	779	836.44
RAID-5	146 GB	7+P	909	909	976.03
RAID-5	300 GB	6+P+S	1582	1582	1698.66
RAID-5	300 GB	7+P	1844	1844	1979.98
RAID-10	73 GB	3+3+2S	192	192	206.16
RAID-10	73 GB	4+4	256	256	274.88
RAID-10	146 GB	3+3+2S	386	386	414.46
RAID-10	146 GB	4+4	519	519	557.27
RAID-10	300 GB	3+3+2S	785	785	842.89
RAID-10	300 GB	4+4	1048	1048	1125.28

Table 3-7 shows the effective capacity of fixed block 8-width RAID ranks in DS6000 in decimal GB and binary GB. It also shows the number of extents.

*Table 3-7 DS6000 8-width RAID rank effective capacity*

RAID type	DDM cap.	Array form.	Extents	Binary GB	Decimal GB
RAID-5	73 GB	6+P+S	382	382	410.17
RAID-5	73 GB	7+P	445	445	477.81
RAID-5	146 GB	6+P+S	773	773	830.00
RAID-5	146 GB	7+P	902	902	968.51
RAID-5	300 GB	6+P+S	1576	1576	1692.21
RAID-5	300 GB	7+P	1837	1837	1972.46
RAID-10	73 GB	3+3+2S	190	190	204.01
RAID-10	73 GB	4+4	254	254	272.73
RAID-10	146 GB	3+3+2S	386	386	414.46
RAID-10	146 GB	4+4	515	515	552.97
RAID-10	300 GB	3+3+2S	787	787	845.03
RAID-10	300 GB	4+4	1050	1050	1127.42



Table 3-8 shows the effective capacities of 4-width RAID ranks in DS6000.

Table 3-8 DS6000 4-width RAID rank effective capacity

RAID type	DDM cap.	Array form.	Extents	Binary GB	Decimal GB
RAID-5	73 GB	2+P+S	127	127	136.36
RAID-5	73 GB	3+P	190	190	204.01
RAID-5	146 GB	2+P+S	256	256	274.87
RAID-5	146 GB	3+P	386	386	414.46
RAID-5	300 GB	2+P+S	524	524	562.64
RAID-5	300 GB	3+P	787	787	845.03
RAID-10	73 GB	1+1+2S	62	62	66.57
RAID-10	73 GB	2+2	127	127	136.36
RAID-10	146 GB	1+1+2S	127	127	136.36
RAID-10	146 GB	2+2	256	256	274.87
RAID-10	300 GB	1+1+2S	261	261	280.24
RAID-10	300 GB	2+2	524	524	562.64

As an example, we calculate the effective capacity for the same DS configuration as we use in “Raw capacity” on page 67, and “Spare disk drives” on page 68. For a DS8100 with 10 RAID-5 ranks of 73 GB DDMs, six ranks are 6+P+S and four ranks are 7+P. The effective capacity is:

$$(6 \times 414.46 \text{ GB}) + (4 \times 483.18 \text{ GB}) = 4419.48 \text{ GB}$$

### i5/OS LUNs and usable capacity for i5/OS

i5/OS LUNs have fixed sizes and are composed of 520 byte blocks consisting of 512 usable bytes for data and 8 header bytes used by System i for storing metadata like the virtual address. The sizes of i5/OS LUNs that are expressed in decimal GB are 8.59 GB, 17.54 GB, 35.16 GB, and so on. These sizes expressed in binary GB are 8 GB, 16.34 GB, 32.75 GB, and so on.

A LUN on DS8000 and DS6000 is formed of so called extents of the size 1 binary GB. Because i5/OS LUN sizes expressed in binary GB are not whole multipliers of 1 GB, part of the space of an assigned extent will not be used but can also not be used for other LUNs.

Table 3-9 shows the models of i5/OS LUNs, their sizes in decimal GB, the number of extents they use, and the percentage of usable space (not wasted) in decimal GB for each LUN.

Table 3-9 i5/OS LUN sizes

Model, unprotected	Model, protected	i5/OS device size (decimal GB)	Number of extents	% of usable space (decimal GB)
A81	A01	* 8.59	8	100
A82	A02	17.54	17	96.14
A85	A05	35.16	33	99.24
A84	A04	70.56	66	99.57
A86	A06	141.1	132	99.57
A87	A07	* 282.2	263	99.95

**Important:** The supported logical volume sizes for a load source unit that is located on ESS Model 800 and DS6000 and DS8000 products are 17.54 GB, 35.16 GB, 70.56 GB, and 141.1 GB. Logical volumes of size 8.59 and 282.2 (noted by the asterisk (\*) in Table 3-9) are not supported as System i5 load source units, where the load source unit is to be located in the external storage server.

When defining a LUN for i5/OS, it is possible to specify whether the LUN is seen by i5/OS as RAID protected or as unprotected. You achieve this by specifying the correct model of i5/OS LUN. Models A0x are seen by i5/OS as *protected*, while models A8x are seen as *unprotected*. Here, *x* stands for 1, 2, 4, 5, 6, or 7.

The general recommendation is to define LUNs as protected models. However you must take into account that, whenever a LUN shall be mirrored by i5/OS mirroring, you must define it as unprotected. Whenever there will be mirrored and non-mirrored LUNs in the same ASP, define the LUNs that shall not be mirrored as *protected*. When mirroring on ASP is started, only the unprotected LUNs from this ASP are mirrored, but all the protected ones are left out of mirroring. This should be considered, e.g when using i5/OS prior to V6R1 when the load source used to be mirrored between an internal disk and a LUN or between two LUNs to provide path redundancy when multipathing was not supported yet for the load source unit.

LUNs are created in DS8000 or DS6000 storage from an extent pool which can contain one or more RAID ranks. For information about the number of available extents from a certain type of DS rank, see Table 3-6 on page 72, Table 3-7 on page 72, and Table 3-8 on page 73.

**Note:** We generally recommend to configure DS8000 or DS6000 storage with only one single rank per extent pool for System i host attachment. This ensures that storage space for a LUN is allocated from a single rank only which helps to better isolate potential performance problems. It also supports the recommendation to use dedicated ranks for System i server or LPARs not shared with other platform servers.

This implies that we also generally do *not* recommend to use the DS8000 Release 3 function of *storage pool striping* (also known as *extent rotation*) for System i host attachment. System i storage management already distributes its I/O as best as possible across the available LUNs in an auxiliary storage pool so that using extent rotation to distribute the storage space of a single LUN across multiple ranks is rather over-virtualization.

An i5/OS LUN uses a fixed number of extents. After a certain number of LUNs are created from an extent pool, usually some is space left. Usually, we define as much as possible LUNs of one size from an extent pool and optionally define LUNs of the next smaller size from the space remaining in the extent pool. We try to define the LUNs of as equal size as possible in order to have balanced I/O rate and consequently better performance.

Table 3-10 and Table 3-11 show possibilities for defining i5/OS LUNs in an extent pool.

*Table 3-10 LUNs from a 6+P+S rank of 73 GB DDMs (386 extents, 414.46 GB)*

<b>70 GB LUNs</b>	<b>35 GB LUNs</b>	<b>17 GB LUNs</b>	<b>8 GB LUNs</b>	<b>Used extents</b>	<b>Used decimal GB</b>
5	1	0	0	381	387.96
0	11	1	0	380	404.3
0	10	3	0	381	404.22
0	9	5	0	382	404.14
0	8	7	0	383	404.06
0	0	22	1	382	394.47
0	0	21	3	381	394.11
0	0	20	5	380	393.75
0	0	19	7	379	393.39
0	0	18	10	386	401.62
0	0	17	12	385	401.26
0	0	16	14	384	400.9

*Table 3-11 LUNs from a 7+P rank of 73 GB DDMs (450 extents, 483.18 GB)*

<b>70 GB LUNs</b>	<b>35 GB LUNs</b>	<b>17 GB LUNs</b>	<b>8 GB LUNs</b>	<b>Used extents</b>	<b>Used decimal GB</b>
6	1	0	0	429	458.52
0	13	1	0	446	474.62
0	12	3	0	447	474.54
0	11	5	0	448	474.46
0	10	7	0	449	474.38
0	0	26	1	450	464.63
0	0	25	3	449	464.27
0	0	24	5	448	463.91
0	0	23	7	447	463.55
0	0	22	9	446	463.19
0	0	21	11	446	462.83
0	0	20	13	444	462.47

Use the following equation to determine the number of LUNs of a given size that one extent pool can contain:

$$\text{number of extents in extent pool} - (\text{number of LUNs} \times \text{number of extents in a LUN}) = \text{residual}$$

Optionally, repeat the same operation to define smaller LUNs from the residual.

The capacity that is available to i5/OS is the number of defined LUNs multiplied by the capacity of each LUN. If for example the DS8000 is configured with six 6+P+S ranks and four 7+P ranks of 73 GB DDMs then from each 6+P+S rank, we define 11 35 GB LUNs and one 17 GB LUN, and from each 7+P rank we define 13 35 GB LUNs and one 17 GB LUN. The capacity available to i5/OS is:

$$(6 \times 404.03 \text{ GB}) + (4 \times 474.62 \text{ GB}) = 4322.66 \text{ GB}$$

### **Capacity Magic**

Capacity Magic, from IntelliMagic™ (Netherlands), is a Windows-based tool that calculates the raw and effective capacity of DS8000, DS6000 or ESS model 800 based on the input of the number of ranks, type of DDMs, and RAID type. The input parameters can be entered through a graphical user interface. The output of Capacity Magic is a detailed report and a graphical representation of capacity.

For more information, refer to:

<http://www.intellimagic.net/en/product.phtml?p=Capacity+Magic>

## Example of using Capacity Magic

In this example, we plan a DS8100 with 9 TB of effective capacity in RAID-5. We use Capacity Magic to calculate the needed raw capacity and to present the structure of spares and parity disks. The process that we use is as follows:

1. Launch Capacity Magic.
2. In the Welcome to Capacity Magic for Windows window (Figure 3-11), specify the type of planned storage system and the desired way to create a Capacity Magic project. In our example, we select **DS6000 and DS8000 Configuration Wizard** and select **OK** to guide us through the Capacity Magic configuration.

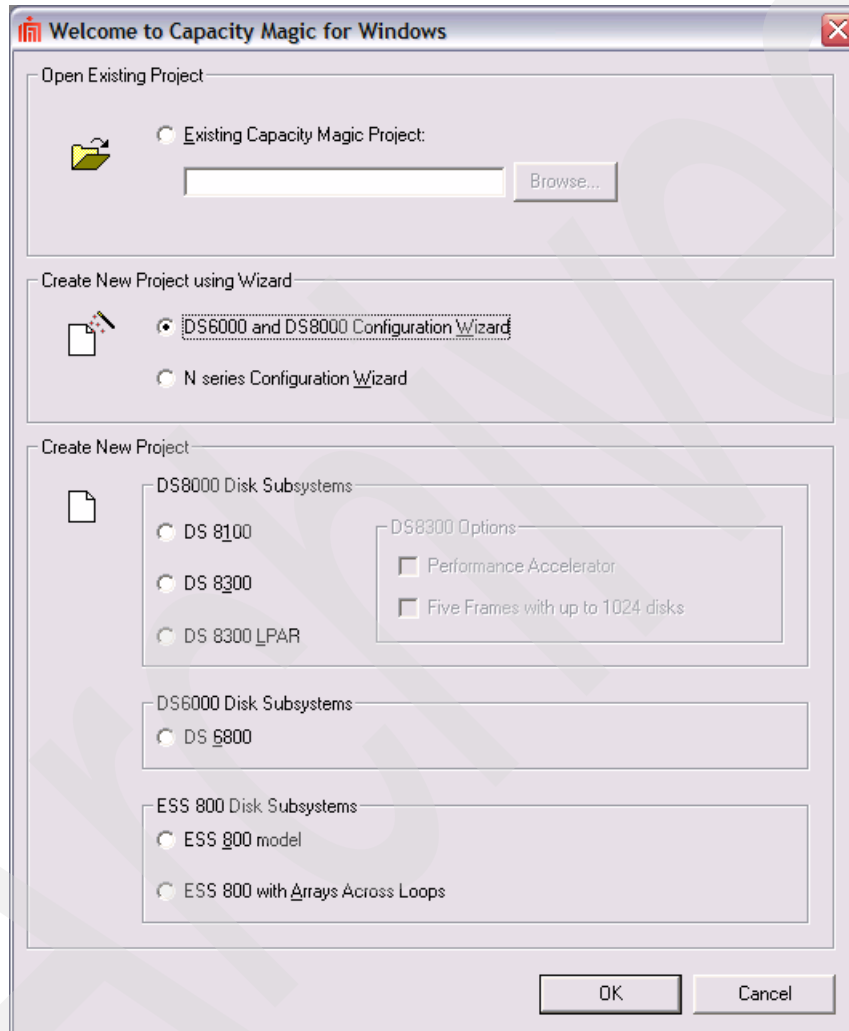


Figure 3-11 Selecting the type of storage system

3. After clicking **Next** in the Wizard's informational window, select which model of DS to use (Figure 3-12). In our example, we select **DS 8100 - 2-way** and click **Next**.




Figure 3-12 Specifying the DS model for Capacity Magic

4. Select the way in which you plan to define the extent pools. For System i attachment, we define **1 Extent Pool for each RAID rank** (see Figure 3-13). Click **Next**.

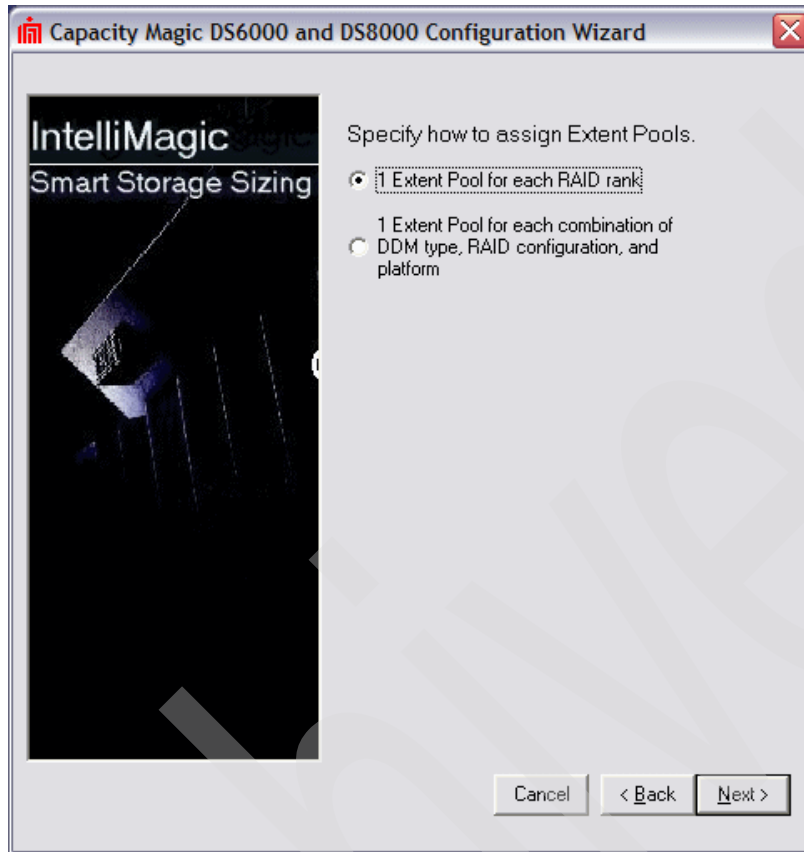


Figure 3-13 Method for defining the extent pool

5. Select the type of host system. In this example, we plan DS8000 for i5/OS, so we select **iSeries** (see Figure 3-14). Then, click **Next**.

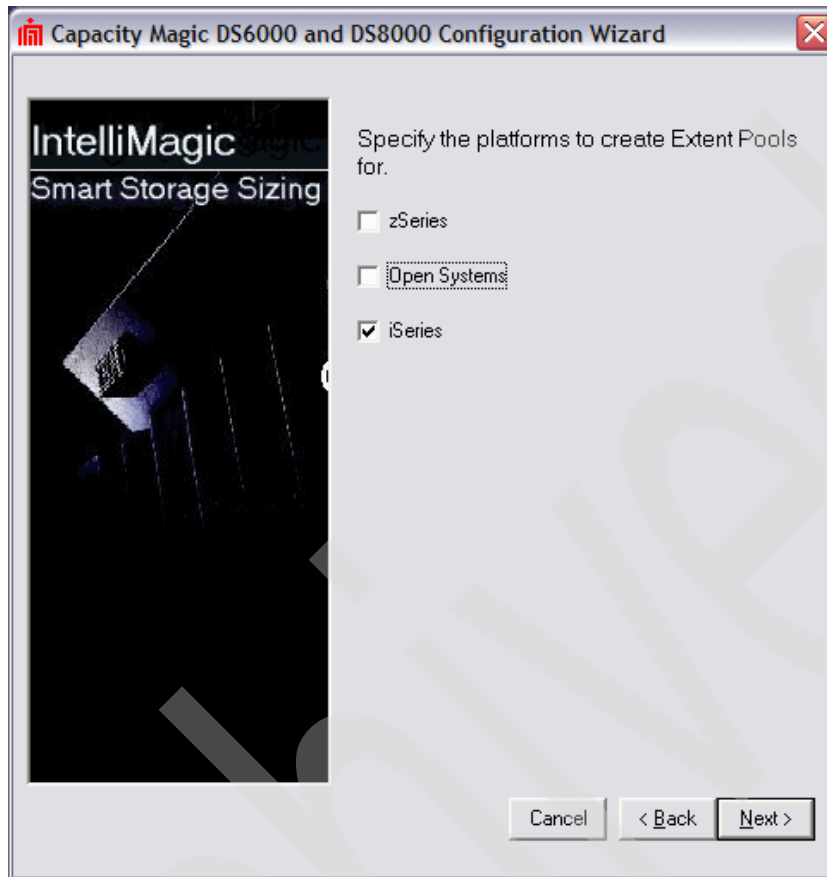


Figure 3-14 Specifying the host system



- Specify the type of DDMs and the type of RAID protection. As shown in Figure 3-15, observe that 73 GB DDMs and RAID-5 are already inserted as the default. In our example, we leave the default values. Click **Next**.

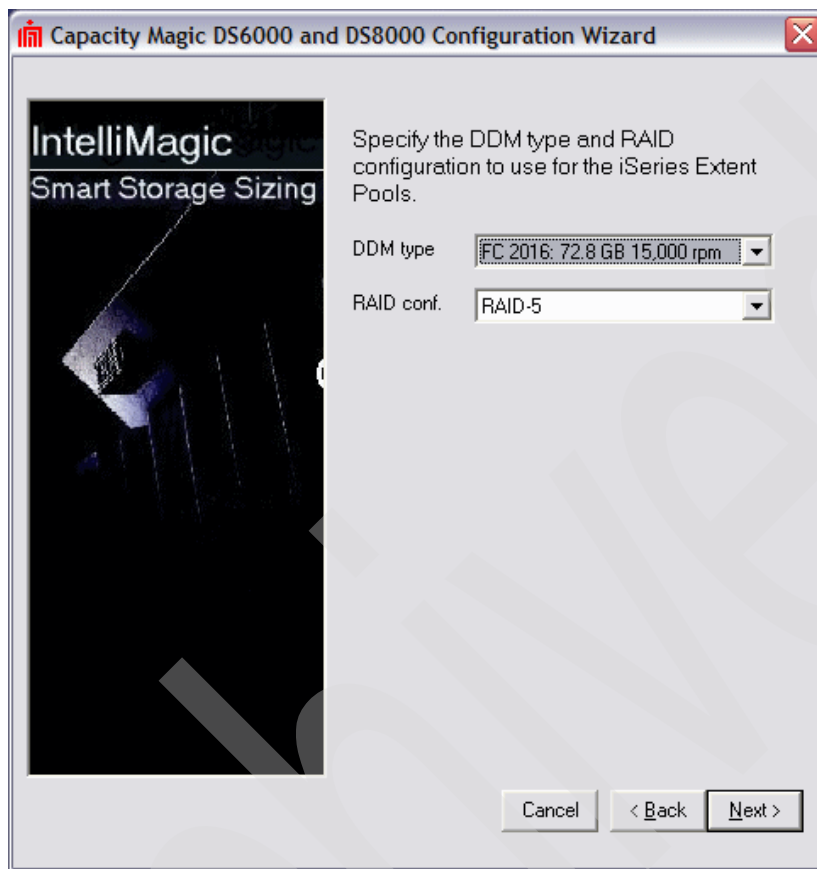


Figure 3-15 Specifying the DDM type and RAID type

7. Specify the desired effective capacity. In our example, we need effective capacity 9 TB, so we enter this value (as shown in Figure 3-16). Note that usable capacity for i5/OS is smaller than the effective capacity, as explained in “i5/OS LUNs and usable capacity for i5/OS” on page 73. Click **Next**.

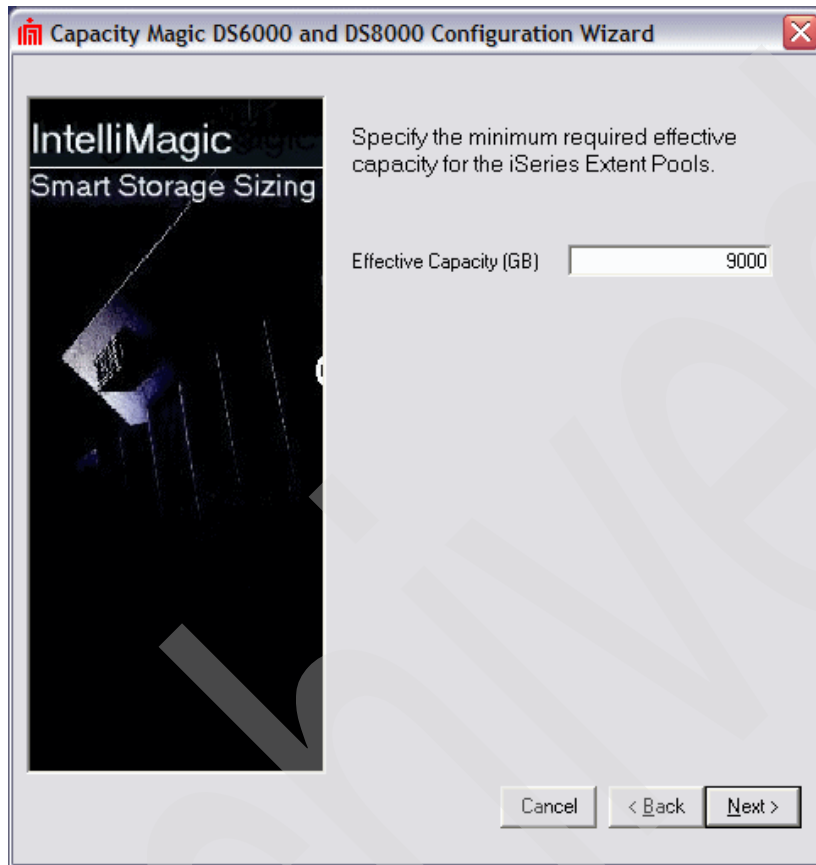


Figure 3-16 Inserting the desired effective capacity

8. Next, review the selected configuration and click **Finish** to continue, as shown in Figure 3-17.

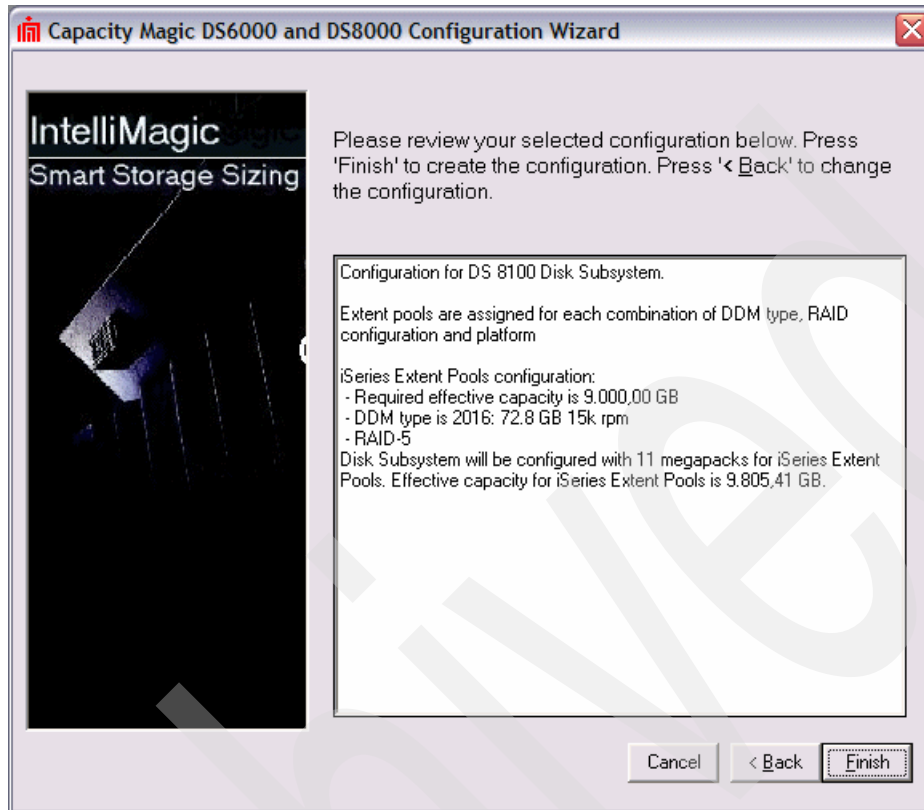


Figure 3-17 Reviewing the selected configuration

- The graphical output displays, which contains the needed number of arrays, spares, and parity disks (see Figure 3-18). Click **Report** table on the right.

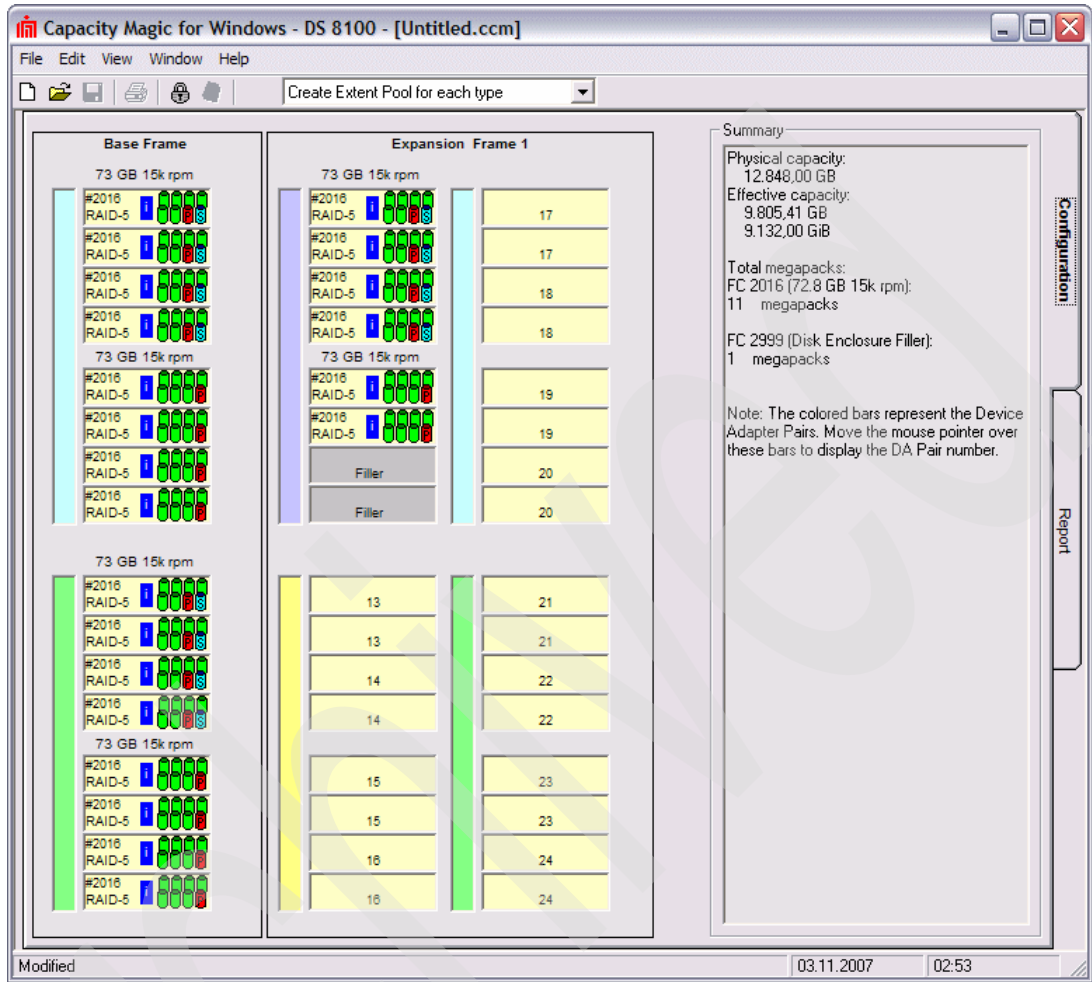


Figure 3-18 Graphical output of Capacity Magic

A detailed report displays of the needed drive sets (*megapacks*), including disk enclosure fillers, number of extents, raw capacity, effective capacity, and so on. Figure 3-19 shows a part of this report.

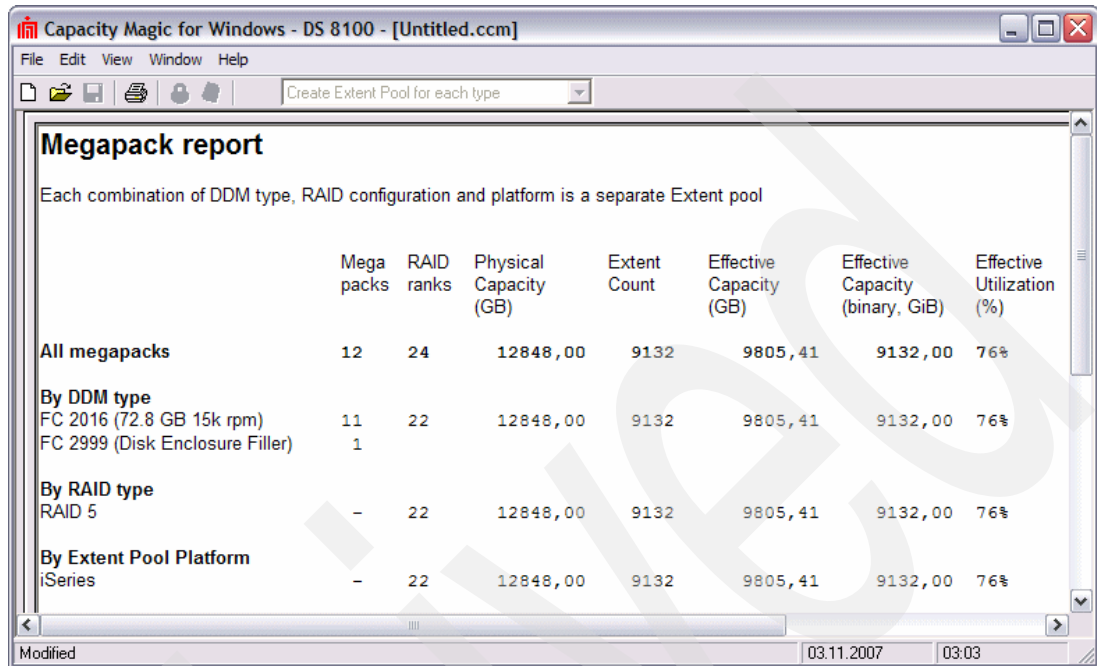


Figure 3-19 Capacity Magic report

### 3.2.7 Planning considerations for performance

It is extremely important that you plan and size both the System i platform and the DS8000, DS6000, or ESS series, properly for an i5/OS workload. You should start by understanding the critical I/O performance periods and the performance expectations. For some customers, response time during transaction workload is critical. For other customers, reduction in the overall batch runs might be important or reduction in the overall save time can be important.

It is equally important to ensure that the sizing requirements for your SAN configuration also take into account the additional resources required when enabling advanced Copy Services functions such as FlashCopy or PPRC. This is particularly important if you are planning to enable synchronous Metro Mirror storage-based replication or space efficient FlashCopy.

**Attention:** You must correctly size the Copy Services functions that are enabled at the system level to account for additional I/O resources, bandwidth, memory, and storage capacity. The use of these functions, either synchronously or asynchronously, can impact the overall performance of your system. To reduce the overhead of not duplicating the temporary objects that are created in the system libraries, such as QTEMP, consider using IASP with Copy Services functions.

We recommend that you obtain i5/OS performance reports from data that is collected during critical workload periods and size the DS8000 or DS6000 accordingly, for every System i environment or i5/OS LPAR that you want to attach to a SAN configuration. For information about how to size IBM System Storage external disk subsystems for i5/OS workloads see Chapter 4, “Sizing external storage for i5/OS” on page 89.

Where performance data is not available, we recommend that you use one of the IBM Benchmark Centers, either in Rochester or France. For more information, see:

[http://www-03.ibm.com/systems/services/benchmarkcenter/servers/benchmark\\_i.html](http://www-03.ibm.com/systems/services/benchmarkcenter/servers/benchmark_i.html)

### PCI I/O card placement rules

Implementation of the PCI architecture provides flexibility in the placement of IOPs and IOAs in IBM System i Models 515, 520, 525, 550, 570, and 595.

With PCI-X, the maximum bus speed is increased to 133 MHz from a PCI maximum of 66 MHz. PCI-X is backward compatible and can run at slower speeds, which means that you can plug a PCI-X adapter into a PCI slot and it runs at the PCI speed, not the PCI-X speed. This can result in a more efficient use of card slots but potentially for the tradeoff of less performance.

Increased configuration flexibility reinforces a requirement to understand the detailed configuration rules. For more information, see *PCI, PCI-X, PCI-X DDR, and PCIe Placement Rules for IBM System i Models*, REDP-4011, at:

<http://www.redbooks.ibm.com/redpieces/abstracts/redp4011.html?Open>

**Attention:** If the configuration rules and restrictions are not fully understood and followed, it is possible to create a hardware configuration that does not work, marginally works, or quits working when a system is upgraded to future software releases.

Follow these plugging rules for the #5760, #2787, and #2766 Fibre Channel Disk Controllers:

- ▶ Each of these adapters requires a dedicated IOP. No other IOAs are allowed on that IOP.
- ▶ For best performance, place these 64-bit adapters in 64-bit PCI-X slots. They can be plugged into 32-bit or 64-bit PCI slots but the performance might not be optimized.
- ▶ If these adapters are heavily used, we recommend that you have only one per Multi-Adapter Bridge (MAB) boundary.

In general spread any Fibre Channel disk controller IOAs as evenly as possible among the attached I/O towers and spread I/O towers as evenly as possible among the I/O loops.

Refer to the recommendations in Table 3-12 for limiting the number of FC adapters per System i I/O half-loop to prevent performance degradation due to congestion on the loop.

Table 3-12 System i I/O loop Fibre Channel adapter recommendations

I/O Half-Loop	Maximum number of #2766/#2787 for transaction/ sequential workload	Maximum number of #5760 for transaction/ sequential workload	Maximum number of IOP-less <sup>a</sup> for transaction/ sequential workload
HSL-1 (1 Gbps) (~400 MBps effective bandwidth unidirectional)	8 / 4	8 / 3	not supported
HSL-2 (2 Gbps) (~750 MBps effective bandwidth unidirectional)	14 / 8	14 / 6	2 / 2
12X (3 Gbps) (~1.2 GBps effective bandwidth unidirectional)	not supported	not supported	3 / 3

a. IOP-less FC dual-port adapters are #5774 and #5749

Our sizing is based on an I/O half-loop concept because, as shown in Figure 3-20, a physically closed I/O loop with one or more I/O towers is actually used by the system as two I/O half-loops. There is an exception to this though only for older iSeries hardware prior to POWER5 where a single I/O tower per loop configuration resulted in only one half-loop being actively used. As can be seen with three I/O towers in a loop, one half-loop will get two, the other half-loop will get I/O tower. The PHYP bringup code determines which half-loop gets the extra I/O tower.

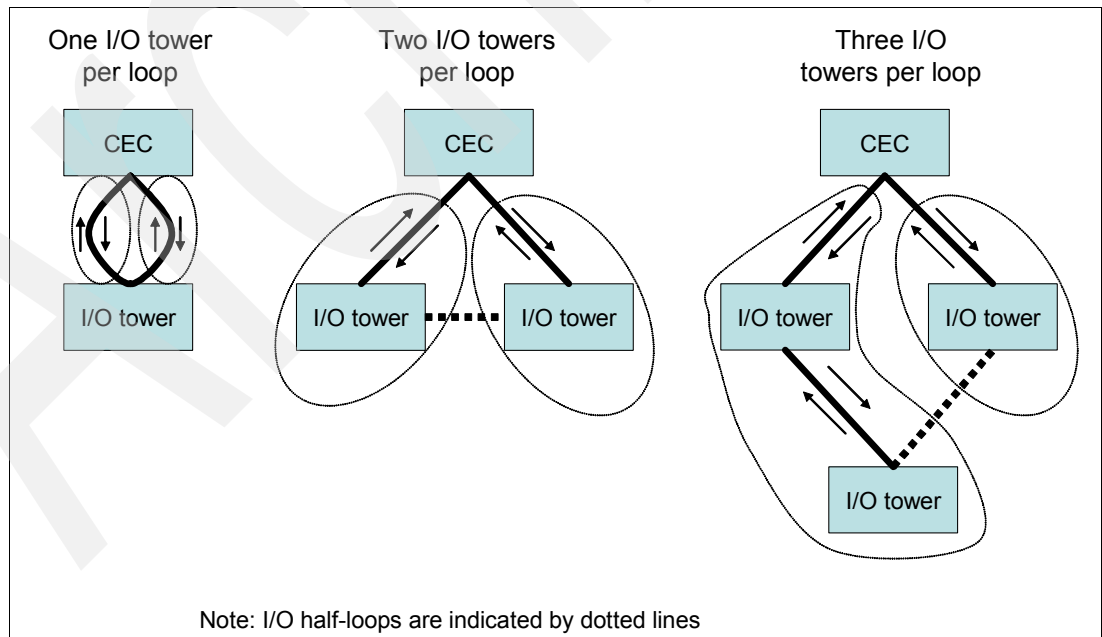


Figure 3-20 I/O half-loop concept

With the System i POWER6 12X loop technology, the parallel bus data-width is increased from previous 8 bits used by HSL-1 and HSL-2 to 12 bits, which is where the name 12X

comes from referring to the number of wires used for data transfer. In addition, with 12X the clock rate is increased to 2.5 GHz compared to 2.0 GHz of the previous HSL-2 technology.

For using System i POWER6 with 12X loops for external storage attachment plan for using it with GX slot P1-C9 (right one from behind) in the CEC which in contrast to its neighbor GX slot P1-C8 does not need to share bandwidth with the CEC's internal slots.

Archived



## Sizing external storage for i5/OS

IBM System Storage Disk Storage series provides maximum flexibility for multiserver storage consolidation and can be designed to meet a variety of customer requirements, including both large storage capacity and good performance. Most customers' System i models typically have mixed application workloads with varying I/O characteristics, such as high I/O rates from transaction workload, as well as large capacity requirements for slower workloads, like data archival.

Fully understanding your customer's i5/OS workload I/O characteristics and then using specific recommended analysis and sizing techniques to configure a DS and System i solution is key to meeting the customer's storage performance and capacity expectations. A properly sized and configured DS system on a System i model provides the customer with an optimized solution for their storage requirements. However, configurations that are drawn without care of proper planning or understanding of workload requirements can result in poor performance and even customer impact events.

In this chapter, we describe how to size a DS system for the System i platform. We present the rules of thumb and describe several tools to help with the sizing tasks.

For good performance of a DS system with i5/OS workload, it is important to provide enough resources, such as disk arms and FC adapters. Therefore, we recommend that you follow the general sizing guidelines or rules of thumb even before you use the Disk Magic™ tool for modeling performance of a DS system with the System i5 platform.

Figure 4-1 illustrates the recommended steps to follow when sizing a DS system for an i5/OS workload.

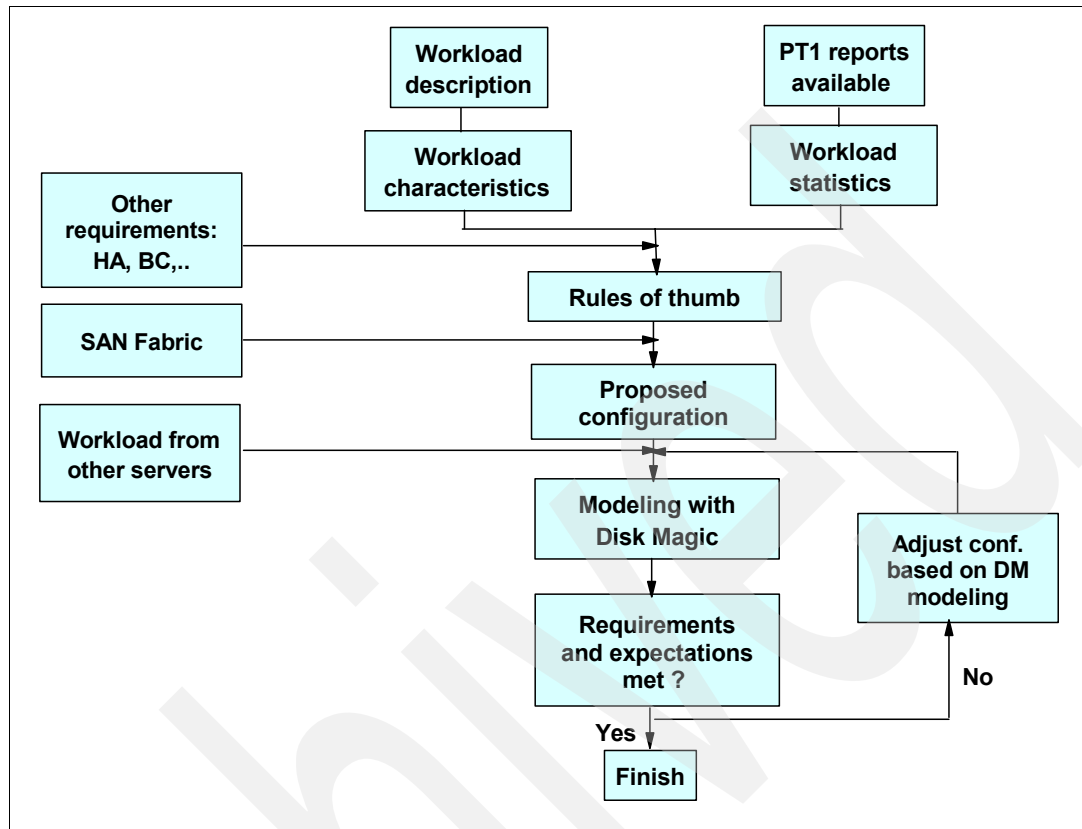


Figure 4-1 Sizing a DS system for i5/OS

## 4.1 General sizing discussion

To better understand the sizing guidelines, in this section, first we briefly describe the I/O flow operation in the System i5 and DS systems. Then we explain how disk response time relates to application response time.

### 4.1.1 Flow of I/O operations

The System i platform with i5/OS uses the same architectural component that is used by the iSeries and AS/400 platform, *single-level storage*. It sees all disk space and the main memory as one storage area. It uses the same set of 64-bit virtual addresses to cover both main memory and disk space. Paging in this virtual address space is performed in 4 KB memory pages.

Figure 4-2 illustrates the concept of single-level storage.

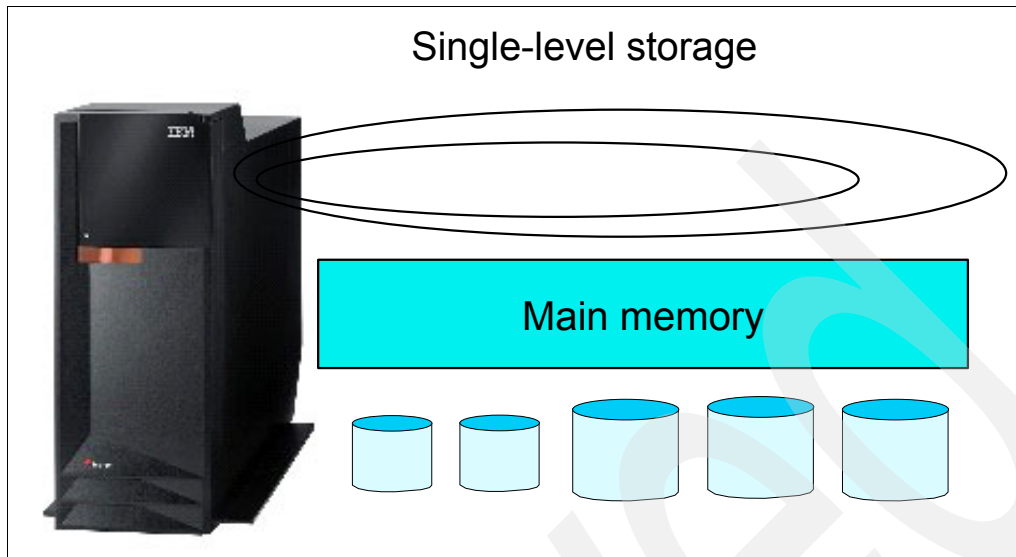


Figure 4-2 Single-level storage

When the application performs an I/O operation, the portion of the program that contains read or write instructions is first brought into main memory where the instructions are then executed.

With the read request, the virtual addresses of the needed record are resolved, and for each needed page, storage management first looks to see if it is in the main memory. If the page is there, it is used to resolve the read request. However, if the corresponding page is not in main memory, a *page fault* is encountered and it must be retrieved from disk. When a page is retrieved, it replaces another page in memory that recently was not used; the replaced page is paged out (destaged) to disk.

Similarly writing a new record or updating an existing record is done in main memory, and the affected pages are marked as changed. A changed page normally remains in main memory until it is written to disk as a result of a page fault. Pages are also written to disk when a file is closed or when write-to-disk is forced by a user through commands and parameters. Also, database journals are written to the disk.

When a page must be retrieved from disk or a page is written to disk, *System Licensed Internal Code* (SLIC) storage management translates the virtual address to a real address of a disk location and builds an I/O request to disk. The amount of data that is transferred to disk at one I/O request is called a *blocksize* or *transfer size*. From the way reads and writes are performed in single-level storage, you would expect that the amount of transferred data is always one page or 4 KB. In fact, data is usually blocked by the i5/OS database to minimize disk I/O requests and transferred in blocks that are larger than 4 KB. The blocking of transferred data is done based on the attributes of database files, the amount that a file extends, user commands, the usage of expert cache, and so on.

Figure 4-3 shows how i5/OS storage management handles read and write operations.

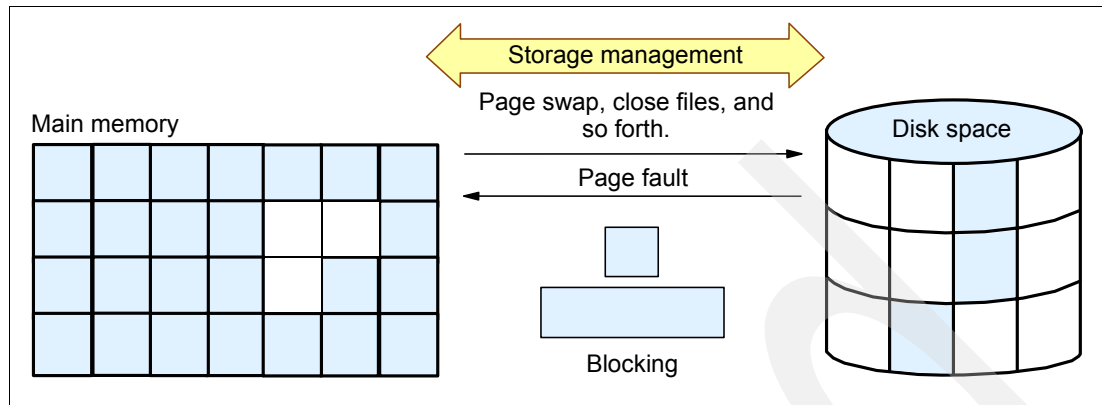


Figure 4-3 Handling I/O operations

An I/O request to disk is created by the I/O adapter (IOA) device driver (DD), which for System i POWER6 now resides in SLIC instead of inside the I/O processor (IOP). It proceeds through the RIO bus to the Fibre Channel IOA, which is used to connect to the external storage subsystem. Each IOA accesses a set of logical volumes, logical unit numbers (LUNs), in a DS system; each LUN is seen by i5/OS as a disk unit. Therefore, the I/O request for a certain System i disk (LUN) goes to an IOA to which a particular LUN is assigned; I/O requests for a LUN are queued in IOA. From IOA, the request proceeds through an FC connection to a host adapter in the DS system. The FC connection topology between IOAs and storage system host adapters can be point-to-point or can be done using switches.

In a DS system, an I/O request is received by the host adapter. From the host adapter, a message is sent to the DS processor that is requesting access to a disk track that is specified for that I/O operation. The following actions are then performed for a read or write operation:

- ▶ **Read operation:** A directory lookup is performed if the requested track is in cache. If the requested track is not found in the cache, the corresponding disk track is staged to cache. The setup of the address translation is performed to map the cache image to the host adapter PCI space. The data is then transferred from cache to host adapter and further to the host connection, and a message is sent indicating that the transfer is completed.
- ▶ **Write operation:** A directory lookup is performed if the requested track is in cache. If the requested track is not in cache, segments in the write cache are allocated for the track image. Setup of the address translation is performed to map the write cache image pages to the host adapter PCI space. The data is then transferred through DMA from the host adapter memory to the two redundant write cache instances, and a message is sent indicating that the transfer is completed.

Figure 4-4 shows the described I/O flow between System i POWER6 and a DS8000 storage system with the previous IOP.

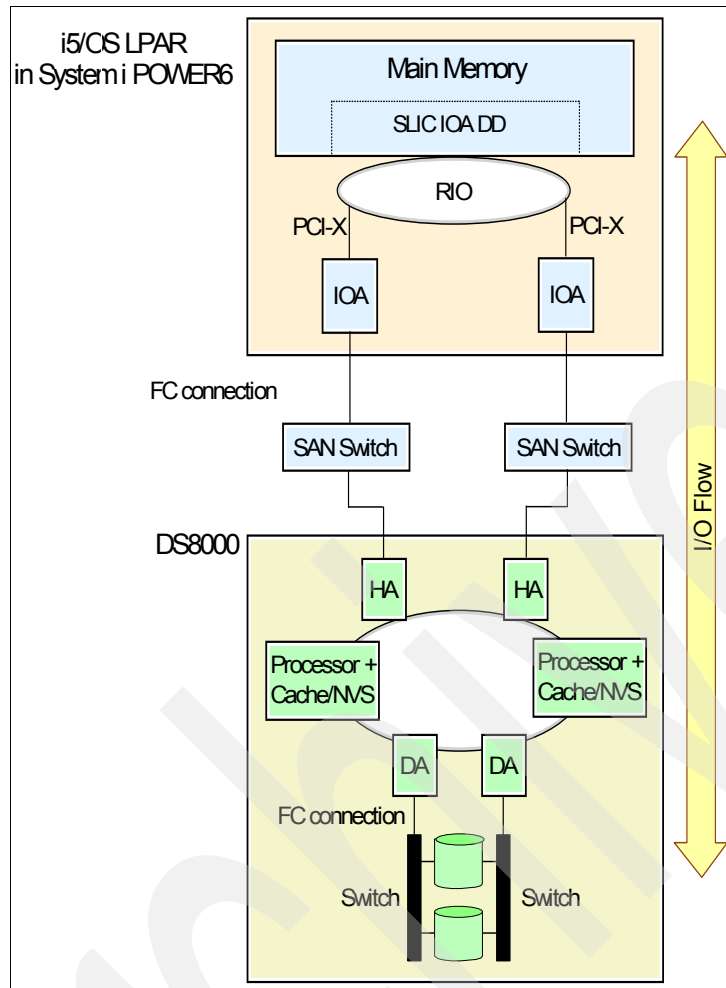


Figure 4-4 System i POWER6 external storage I/O flow

#### 4.1.2 Description of response times

When sizing, it is important to understand how performance of the disk subsystem influences application performance. To explain this, we first describe the critical performance times:

- ▶ **Application response time:** The response time of an application transaction. This time is usually critical for the customer.
- ▶ **Duration of batch job:** Batch jobs usually run during the night; the duration of a batch job is critical for the customer, because it must be finished before regular daily transactions start.
- ▶ **Disk response time:** The time that is needed for a disk I/O operation to complete which includes the service time for actual I/O processing and the wait time for potential I/O queuing on the System i host. For IOP-based IOAs disk response time is derived from sampling at the IOP level so that this data was only representative for at least around five I/O per second. With System i POWER6 IOP-less IOAs the disk response time is really measured in SLIC.

Single-level storage makes main memory work as a big cache. Reads are done from pages in main memory, and requests to disk are done only when the needed page is not there. Writes are done to main memory, and write operations to disk are performed only as a result of swap or file close, and so on. Therefore, application response time depends not only on disk response time but on many other factors, such as how large the i5/OS storage pool is for the application, how frequently the applications closes files, whether it uses journaling, and so on. These factors differ from application to application. Thus, it is difficult to give a general rule about how disk response time influences application response time or duration of a batch job.

Performance measurements were done in IBM Rochester that show how disk response time relates to throughput. These measurements show the number of transactions per second for a database workload. This workload is used as an approximation for an i5/OS transaction workload. The measurements were performed for different configurations of DS6000 connected to the System i platform and different workloads. The graphs in Figure 4-5 show disk response time at workloads for 25, 50, 75, 100, and 125 database users.

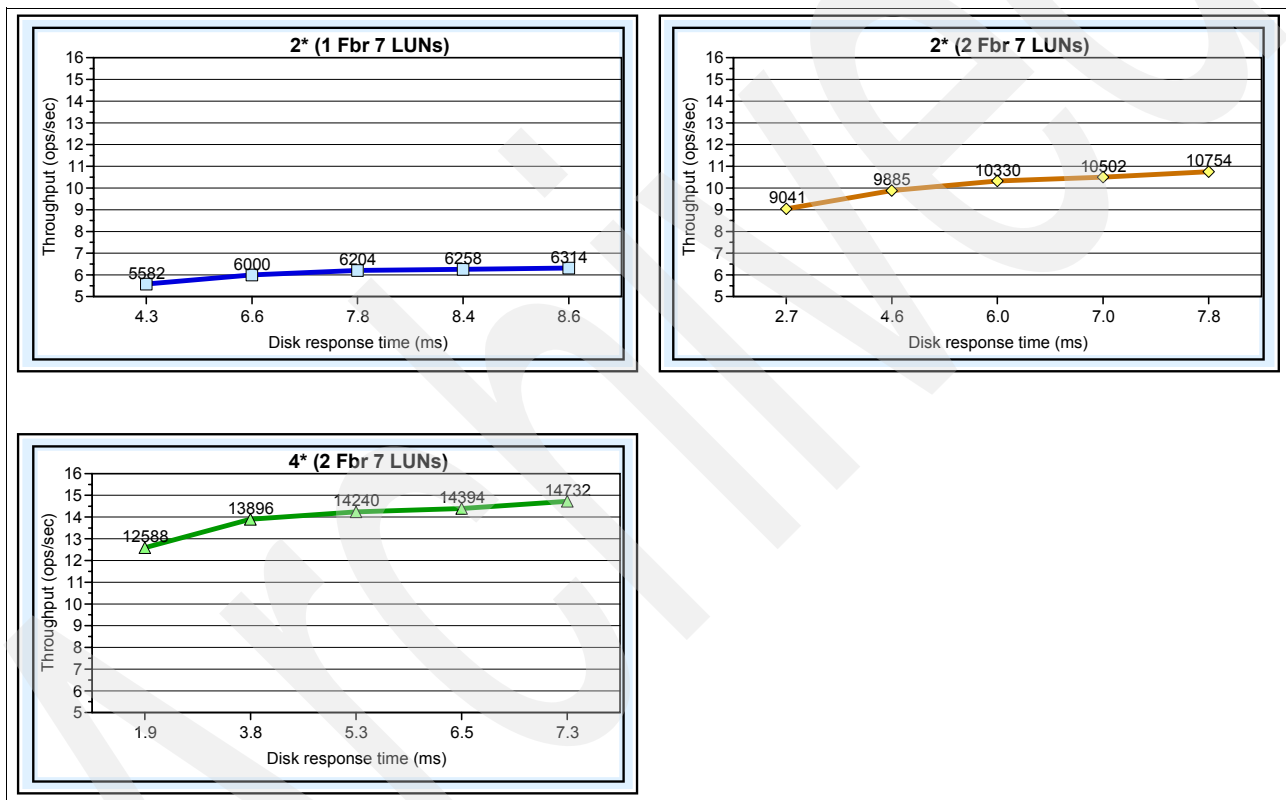


Figure 4-5 Disk response time at different database workloads

From the three graphs, notice that as we increase the number of FC adapters and LUNs, we gain more throughput. If we merely increased the throughput for a given configuration, we can see the disk response time grow sharply.

## 4.2 Rules of thumb

When sizing on i5/OS workload for external storage, we recommend that you use some sizing rules of thumb even before you start your external storage performance modeling with the Disk Magic sizing tool (see 4.3.1, “Disk Magic” on page 106). This way, we ensure that the

basic performance requirements are met and eliminate future performance bottlenecks as much as possible.

Through these rules of thumb, we determine the following characteristics of a DS6000 or DS8000 system, a System i model, and a storage area network (SAN) configuration:

- ▶ The number of RAID ranks in DS6000 or DS8000
- ▶ The number of System i5 Fibre Channel adapters
- ▶ The size of System i LUNs to create and how to spread them over extent pools
- ▶ The manner in which to share a DS6000 or DS8000 among multiple System i models or between a System i environment and another workload
- ▶ When connecting through SAN switches, the number of System i FC adapters to connect to one FC port in DS6000 or DS8000

### 4.2.1 Number of RAID ranks

For a typical System i transaction workload which due to its largely random I/O character is rather cache unfriendly it is extremely important to provide i5/OS with enough disk arms to achieve good application performance.

When a page or a block of data is written to disk space, storage management spreads it over multiple disks. By spreading data over multiple disks, it is achieved that multiple disk arms work in parallel for any request to this piece of data, so writes and reads are done faster.

When using external storage with i5/OS what SLIC storage management sees as a “physical” disk unit is actually a logical unit (LUN) composed of multiple stripes of a RAID rank in the IBM DS storage subsystem (see Figure 4-6). A LUN uses multiple disk arms in parallel depending on the width of the used RAID rank. For example, the LUNs configured on a single DS8000 RAID5 rank use six or seven disk arms in parallel, while with evenly distributing these LUNs over two ranks twice as much disk arm are used.

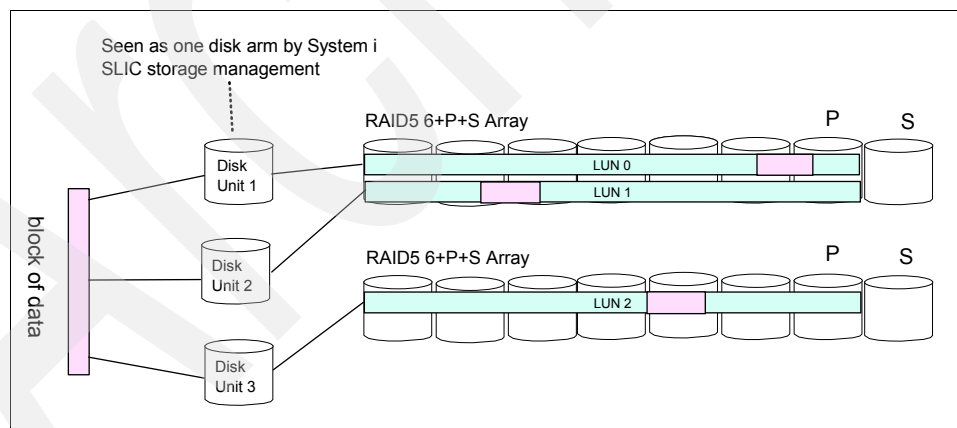


Figure 4-6 Usage of disk arms

Typically the number of physical disk arms that should be made available for a performance critical i5/OS transaction workload is prevailing over the capacity requirements.

**Important:** Determining the number of RAID ranks for a System i external storage solution by looking at how many ranks of a given physical DDM size and RAID protection level would be required for the desired storage capacity typically does *not* satisfy the performance requirements of System i workload.

The rule of thumb how many ranks are needed for on i5/OS workload is based on performance measurements for one DS6000 and DS8000 RAID rank. Our calculations are based on 15 KB RPM disk drive modules (DDMs).

**Note:** We generally do not recommend using lower speed 10 KB RPM drives for i5/OS workload.

The calculation for the recommended number of RAID ranks is as follows, providing that reads per second and writes per second of an i5/OS workload are known:

- ▶ A RAID-5 rank of 8 \* 15 KB RPM DDMs without a spare disk (7+P rank) is capable of maximum 1700 disk operations per second at 100% utilization without cache hits. This is valid for both DS8000 and DS6000.
- ▶ We take into account a recommended 40% utilization of a rank so the rank can handle 40% of 1700 = 680 disk operations per second. From the same measurement we can calculate maximum number of disk operations per second for other RAID ranks by calculating disk operations per second for one disk drive and then multiplying them by the number of active drives in a rank. For example, a RAID-5 rank with a spare disk (6+P+S rank) can handle maximum  $1700 / 7 * 6 = 1458$  disk operations per second. At recommended 40% utilization it can handle 583 disk operations per second.
- ▶ We calculate disk operations of i5/OS workload so that we take into account percentage of read cache hits, percentage of write cache hits, and the fact that each write operation in RAID-5 results in 4 disk operations (RAID-5 write penalty). If cache hits are not known, we make a save assumption of 20% read cache hits and 30% write cache hits. We use the following formula:

$$\text{disk operations} = (\text{reads/sec} - \text{read cache hits}) + 4 * (\text{writes/sec} - \text{write cache hits})$$

As an example, a workload of 1000 reads per second and 700 writes per second results in:

$$(1000 - 20\% \text{ of } 1000) + 4 * (700 - 30\% \text{ of } 700) = 2760 \text{ disk operations/sec}$$

- ▶ To obtain the needed number of ranks, we divide disk operations per second of i5/OS workload by the maximum I/O rate one rank can handle at 40% utilization.

As an example, for workload with previously calculated 2760 disk operations per second, we need the following number of 7+P raid-5 ranks:

$$2760 / 680 = 4$$

So, we recommend to use 4 ranks in DS for this workload.

A handy reference for determining the recommended number of RAID ranks for a known System i workload is provided by the table in Table 4-1 on page 97, which shows the I/O capabilities of different RAID-5 and RAID-10 rank configurations. The I/O capability numbers in the two columns for the host I/O workload examples of 70/30 and 50/50 read/write ratios imply no cache hits and 40% rank utilization. If the System i workload is similar to one of the



two listed read/write ratios a rough estimate for the number of recommended RAID ranks can simply be determined by dividing the total System i I/O workload by the listed I/O capability for the corresponding RAID rank configuration.

Table 4-1 DS8000/DS6000 RAID rank capabilities

RAID rank type	Disk I/O per second	Host I/O per second (70% read)	Host I/O per second (50% read)
RAID-5 15 KB (7 + P)	1700	358	272
RAID-5 10 KB (7 + P)	1100	252	176
RAID-5 15 KB (6 + P + S)	1458	313	238
RAID-5 10 KB (6 + P + S)	943	199	151
RAID-10 15 KB (3 + 3 + 2S)	1275	392	340
RAID-10 10 KB (3 + 3 + 2S)	825	254	220
RAID-10 15 KB (4 + 4)	1700	523	453
RAID-10 15 KB (4 + 4)	1100	338	293

## 4.2.2 Number of Fibre Channel adapters

For connecting the System i platform to IBM System Storage disk subsystems, 4 Gb IOP-based single-port Fibre Channel (FC) adapters with System i feature code 5760 or for System i POWER6 models only the new IOP-less dual-port FC adapters 5749 or 5774 are used. Also older 2 Gb IOP-based single-port FC adapters with System i feature number 2766 or 2787 can be used for this connection together with an IOP 2843, 2844, or 2847. Refer to 3.2, “Solution implementation considerations” on page 54 for further information about planning the attachment of System i to external disk storage.

Similar to the number of ranks, to avoid potential I/O performance bottleneck due to undersized configurations it is also important to properly size the number of Fibre Channel adapters used for System i external storage attachment. To better understand this sizing, we present a short description of the data flow through IOPs and the FC adapter (IOA).

A block of data in main memory consists of an 8 byte header and actual data that is 512 bytes long. When the block of data is written from main memory to external storage or read to main memory from external storage, requests are first sent to the IOA device driver which converts the requests to generate a corresponding SCSI command understood by the disk unit resp. storage system. The IOA device driver either resides within the IOP for IOP-based IOAs or within SLIC for IOP-less IOAs. In addition, data descriptor lists (DDLs) tell the IOA where in system memory the data and headers reside. See Figure 4-7.

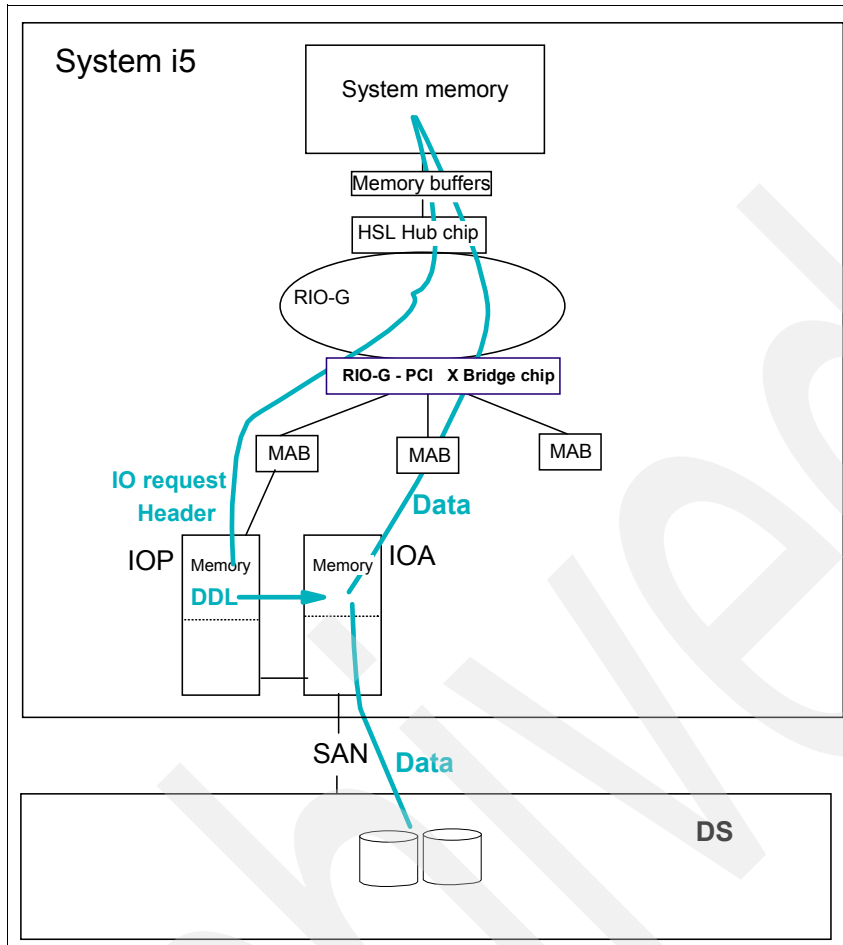


Figure 4-7 Data flow through IOP and IOA

With IOP-less Fibre Channel architectural changes in the process of getting the eight headers for a 4 KB page out of or to main memory by packing them into just one DMA request reduce the latency for disk I/O operations and put less burden on the PCI-X.

You need to size the number of FC adapters carefully for the throughput capability of an adapter. Here, you must also take into account the capability of the IOP and the PCI connection between the adapter and IOP.

We performed several measurements in the testing for this book, by which we can size the capability of an adapter in terms of maximal I/O per second at different block sizes or maximal MBps. Table 4-2 shows the results of measuring maximal I/O per second for different System i Fibre Channel adapters and the I/O capability at 70% utilization which is relevant for sizing the number of required System i FC adapters for a known transactional I/O workload.

Table 4-2 Maximal I/O per second per Fibre Channel IOA

IOP/IOA	Maximal I/O per second per port	I/O per second per port at 70% utilization
IOP-less 5749 or 5774	15000	10500
2844 IOP / 5760 IOA	3900	3200
2844 IOP / 2787 IOA	3650	2555

Table 4-3 shows the maximum throughput for System i Fibre Channel adapters based on measurement of large 256 KB block sequential transfers and typical transaction workload with rather small 14 KB block transfers.

Table 4-3 Maximum adapter throughput

FC adapter	Maximum sequential throughput per port	Maximum transaction workload throughput per port
IOP-less 5749 or 5774	310 MBps	250 MBps
2844 IOP / 5760 IOA	140 MBps	45 - 54 MBps
2844 IOP / 2787 IOA	92 MBps	34 - 54 MBps

When using *IOP-based* FC adapters there is another reason why the number of FC adapters is important for performance. With IOP-based FC adapters only one I/O operation per path to a LUN can be done at a time, so I/O requests could queue up in each LUN queue in the IOP resulting in undesired I/O wait time. SLIC storage management allows a maximum of six I/O requests in an IOP queue per LUN and path. By using more FC adapters for adding paths to a LUN the number of active I/O and the number of available IOP LUN I/O queues can be increased.

**Note:** For *IOP-based* Fibre Channel using more FC adapters for multipath with adding more paths to a LUN can help to significantly reduce the disk I/O wait time.

With *IOP-less* Fibre Channel support the limit of one active I/O per LUN per path has been removed and up to six active I/O per path and LUN are now supported. This inherently provides six times better I/O concurrency compared to previous IOP-based Fibre Channel technology and makes multipath for IOP-less a function primarily for redundancy which less potential performance benefits compared to IOP-based Fibre Channel technology.

When a System i customer plans for external storage, the customer usually decides first how much disk capacity is needed and then asks how many FC adapters will be necessary to handle the planned capacity. It is useful to have a rule of thumb to determine how much disk capacity to plan per FC adapter. We calculate this by using the access density of an i5/OS workload. The *access density* of a workload is the number of I/O per second per GB and denotes how “dense” I/O operations are on available disk space.

To calculate the capacity per FC adapter, we take the maximum I/O per second that an adapter can handle at 70% utilization (see Table 4-2). We divide the maximal number of I/O per second by access density to get the capacity per FC adapter. We recommend that LUN utilization does not exceed 40%. Therefore, we apply 40% to the calculated capacity.

Consider this example. An i5/OS workload has an access density of 1.4 I/O per second per GB. Adapter 5760 at IOP 2844 is capable of a maximum of 3200 I/O per second at 70% utilization. Therefore, it handles the capacity 2285 GB, that is:

$$3200 / 1.4 = 2285 \text{ GB}$$

After applying 40% for LUN utilization, the sized capacity per adapter is 40% of 2285 GB which is:

$$2285 * 40\% = 914 \text{ GB}$$

In addition to a proper sizing of the number of FC adapters to use for external storage attachment also following the guidelines for placing IOPs and FC adapters (IOAs) in the System i platform (see 3.2.7, “Planning considerations for performance” on page 85).

### 4.2.3 Size and allocation of logical volumes

A logical volume is seen by i5/OS as a disk drive, but in fact, it is composed of multiple data stripes taken from one RAID rank in the IBM System Storage disk subsystem, as shown in Figure 4-6 on page 95. From the DS perspective, the size of the LUN does not affect its performance.

With IOP-based Fibre Channel LUN size considerations are very important from System i perspective because of the limit of one active I/O per path per LUN. (We discuss this limitation in 4.2.2, “Number of Fibre Channel adapters” on page 97 and mention multipath for IOP-based Fibre Channel as a solution that can reduce the wait time as each additional path to a LUN enables one more active I/O to this LUN.) For the same reason of increasing the amount of System i active I/O with IOP-based Fibre Channel, we rather recommend using more smaller LUNs than fewer larger LUNs.

**Note:** As a rule of thumb for IOP-based Fibre Channel, we recommend choosing the LUN size so that the ratio between the capacity of the DDM and LUN is at least two LUNs per capacity of the DDM.

With 73 GB DDM capacity in the DS system, a customer can define 35.1 GB LUNs. For even better performance 17.54 GB LUNs can be considered.

IOP-less Fibre Channel supports up to six active I/O for each path to a LUN so compared to IOP-based Fibre Channel there is no stringent requirement anymore to use small LUN sizes for better performance.

**Note:** With IOP-based Fibre Channel we generally recommend using a LUN size of 70.56 GB, that is protected and unprotected volume model A04/A84, when configuring LUNs on external storage.

Currently the only exception, for when we recommend using larger LUN sizes than 70.56 GB, is when the customer anticipates a low capacity usage within the System i auxiliary storage pool (ASP). For a low ASP capacity usage, using larger LUNs can provide better performance by reducing the data fragmentation on the disk subsystem's RAID array resulting in less disk arm movements as illustrated in Figure 4-8.

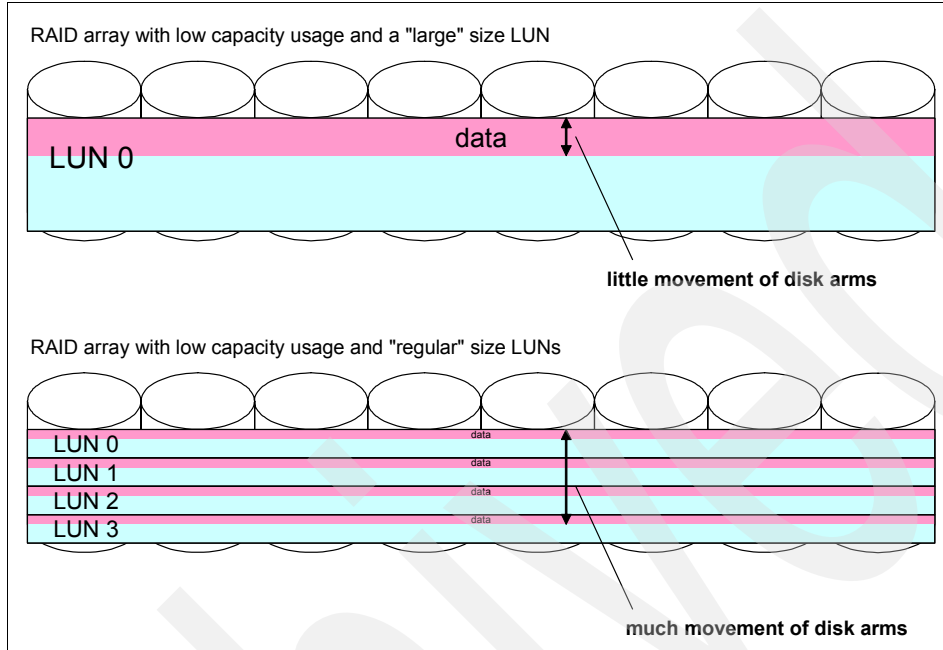


Figure 4-8 RAID array data distribution

When allocating the LUNs for i5/OS, consider the following guidelines for better performance:

- ▶ Balance the activity between the two DS processors, referred to as cluster0 and cluster1, as much as possible. Because each cluster has separate memory buses and cache, this maximizes the use of those resources.

In the DS system, an extent pool has an affinity to either cluster0 or cluster1. We define it by specifying a *rank group* for a particular extent pool with rank group 0 served by cluster0 and rank group 1 served by cluster1. Therefore, define the same amount of extent pools in rank group 0 as in rank group 1 for the i5/OS workload and allocate the LUNs evenly among them.

**Recommendation:** We recommend that you to define one extent pool from one rank to keep better evidence of LUNs and to ensure that LUNs are spread evenly between the two processors.

- ▶ Balance the activity of a critical application among device adapters in the DS system. When choosing extent pools (ranks) for a critical application, make sure that they are evenly served by as much as possible by device adapters.

In the DS system, we define a *volume group* that is a group of LUNs that are assigned to one System i FC adapter or to multiple FC adapters in a multipath configuration. Create a volume group so that it contains LUNs from the *same* rank group, that is do not mix logical subsystem (LSS) LUNs server by cluster0 and odd LSS LUNs served by cluster1 on the same System i host adapter. This multipath configuration helps to optimize sequential read performance with making most efficient usage of the available DS8000 RIO loop bandwidth.

## 4.2.4 Sharing ranks among multiple workloads

Looking from one angle, sharing a DS rank among multiple workloads can improve performance because workloads that share a rank do not use the disk arms at the same time. When one workload is idle, the other can use disk arms of a rank. It appears to each workload that it uses all disk arms of a rank. By sharing multiple ranks among workloads, we provide each workload with more disk arms compared to dedicating fewer ranks to a workload.

A heavy workload might get hold of disk arms in a rank and cache for almost all the time, so the other workload will rarely have a chance to use them. Alternatively, if two heavy critical workloads share a rank, they can prevent each other from using disk arms and cache at the times when both are busy.

Therefore, we recommend that you dedicate ranks for a heavy critical i5/OS workload such as SAP or banking applications. When the other workload does not exceed 10% of the workload from your critical i5/OS application, consider sharing the ranks.

Consider sharing ranks among multiple i5/OS systems or among i5/OS and open systems when the workloads are less important and not I/O intensive. For example, testing and developing, mail, and so on can share ranks with other systems.

## 4.2.5 Connecting using switches

When connecting System i FC adapters using a SAN switch to a storage subsystem, refer to 3.2.5, “Planning for SAN connectivity” on page 67 for information about how to zone SAN switches. Implementing a proper SAN switch zoning is crucial to help prevent performance degradation caused by potential link congestion problems. Still the question usually arises regarding the number of FC adapters to plan for attaching to one DS port to ensure good performance.

With DS8000, we recommend that you size up to four 2 Gb FC adapters per one 2 Gb DS port. In DS6000, consider sizing two System i FC adapters for one DS port. Figure 4-9 shows an example of SAN switch zoning for four System i FC adapters accessing one DS8000 host port.

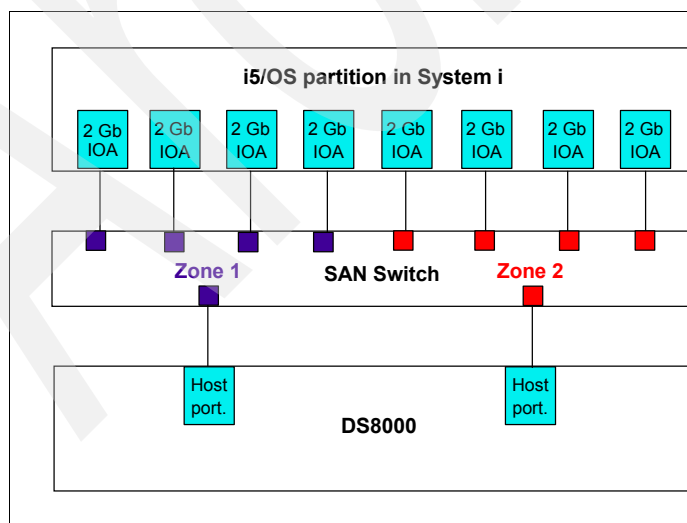


Figure 4-9 Connecting a System i environment to a DS system using SAN switches

Consider the following guidelines for connecting System i 4 Gb FC IOAs to 4 Gb adapters in the DS8000:

- ▶ Connect one 4 Gb IOA port to one port on DS8000, provided that all four ports of the DS8000 adapter card are used.
- ▶ Connect two 4 Gb IOA ports to one port in DS8000, provided that only two ports of the DS8000 adapter card are used.

## 4.2.6 Sizing for multipath

Multipath enables up to eight different paths to a set of LUNs. To ensure redundancy for each path, separate System i FC adapters and usually separate physical connections to DS are used. For *IOP-based* Fibre Channel IOAs multipath does not provide only high availability in case one path fails but also provides better performance compared to a single path connection due to more active I/O and higher I/O throughput. Regarding this, how much better does a set of LUNs in IOP-based multipath perform compared to a single path?

Figure 4-10 shows the disk response time measurements of the same database workload running in a single path and dual path at different I/O rates. The blue line represents a single path, and the yellow line represents dual path.

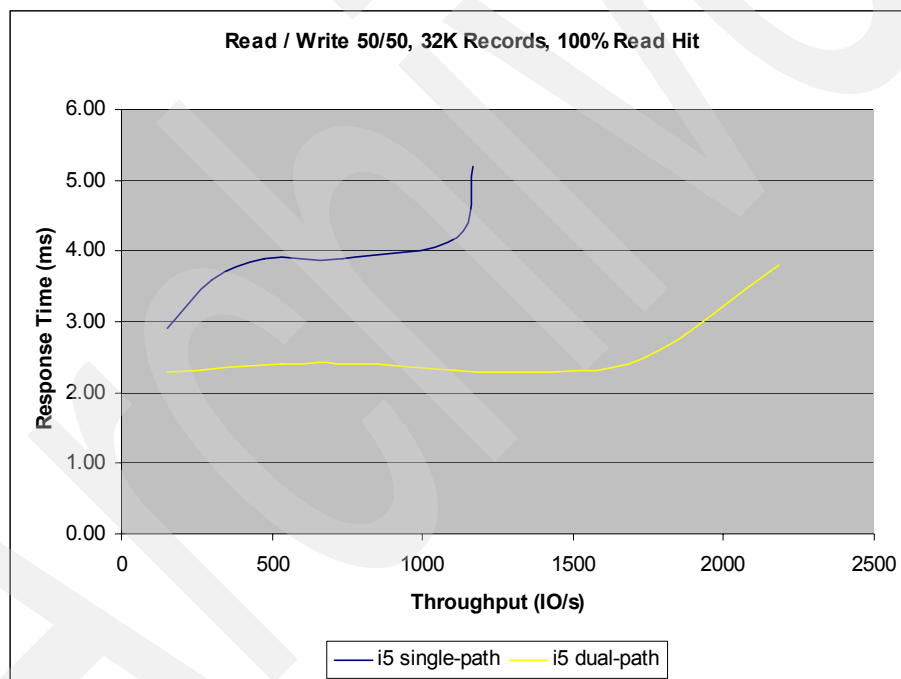


Figure 4-10 Single path versus dual path performance

The response time in IOP with a single path starts to increase drastically at about 1200 I/O per second. With two paths, it starts to increase at about 1800 I/Os per second. From this, we can make a rough rule of thumb that for IOP-based Fibre Channel multipath with two paths is capable of 50% more I/Os than a single path and provides significantly shorter wait time than a single path. Disk response time consists of service time and wait time. With multipath, only the wait time is improved, while it does not influence service time. With IOP-less Fibre Channel allowing six times as much active I/O as IOP-based Fibre Channel the performance improvement by using multipath is of minor importance and multipath is primarily used for redundancy.

The sizing tool Disk Magic takes the performance improvement due to multipath into account and is planned to be updated for modelling System i IOP-less Fibre Channel performance.

For more information about how to plan for multipath, refer to 3.2.2, “Planning considerations for i5/OS multipath Fibre Channel attachment” on page 57.

## 4.2.7 Sizing for applications in an IASP

To implement a high availability or disaster recovery solution for an application using independent auxiliary storage pools (IASPs) or using IASPs for other purposes such as server consolidation, we recommend that you size external storage for IASP and \*SYSBAS separately.

i5/OS performance reports—*Resource report - Disk utilization* and *System report - Disk utilization*—show the average number of I/O per second for both IASP and \*SYSBAS. To see how many I/Os per second actually go to an IASP, we recommend that you look at the *System report - Resource utilization*. This report shows the database reads per second and writes per second for each application job, as shown in Figure 4-12.

Average	U, 34	U, UU	1, 8	U, 5	1, U	834						
Job Type	CPU Util	Tns /Hour Rate	Active Jobs Per Interval	Total I/O	DBR	Synchronous DBW	NDBR	NDBW	DBR	Asynchronous DBW	NDBR	NDBW
DDM Server	0,0	0	3	0,1	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
PassThru	0,0	676	1	0,3	0,0	0,0	0,1	0,0	0,0	0,0	0,0	0,0
Batch	9,8	0	10	530,2	12,2	2,9	314,7	6,2	112,3	11,0	68,0	2,5
AutoStart	0,0	0	0	3,9	0,0	0,2	0,6	0,5	0,0	0,2	1,5	0,6
COLLECTION	0,0	0	0	0,7	0,1	0,1	0,3	0,0	0,0	0,1	0,0	0,0
DIRSRV	0,0	0	0	0,4	0,0	0,0	0,2	0,2	0,0	0,0	0,0	0,0
FTP	0,0	0	0	0,1	0,0	0,0	0,1	0,0	0,0	0,0	0,0	0,0
TELNET	0,0	0	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
MGMTCENTRAL	0,0	0	0	0,2	0,0	0,0	0,2	0,0	0,0	0,0	0,0	0,0
Average	10,1	677	14	536,1	12,4	3,3	316,5	7,1	112,4	11,4	69,5	3,1

Figure 4-11 Database reads and writes

Add the database reads per second (synchronous DBR and asynchronous DBR) and the database writes per second (synchronous DBW and asynchronous DBW) of all application jobs in IASP. Then, you can obtain reads per second and writes per second of IASP. Calculate the number of reads per second and writes per second for \*SYSBAS so that you subtract the reads per second of the IASP from the overall reads per second and subtract the writes per second of the IASP from the overall writes per second.

To allocate LUNs for IASP and \*SYSBAS, we recommend that you first create the LUNs for the IASP and spread them across available ranks in the DS system. From the left free space on each rank, define (smaller) LUNs for using as \*SYSBAS disk units. The reasoning for this approach is that the first LUNs created on a RAID rank are created on the outer cylinders of the disk drives which provide a higher data rate than the inner cylinders.

## 4.2.8 Sizing for space efficient FlashCopy

Properly sizing the storage space for the repository volume that provides the physical storage capacity for the space efficient volumes within the same extent pool is very important. Proper sizing prevents you from running out of physical storage space and causes the space efficient FlashCopy relationship to fail, which causes performance issues due to an undersized number of disk arms for the repository volume.



The following sizing approach can help you prevent this undesired situation:

1. Use i5/OS Performance Tools to collect a *resource report* for disk utilization from the production system, which accesses the FlashCopy source volumes, and the backup system, which accesses the FlashCopy target volumes (see 4.4.3, “i5/OS Performance Tools” on page 111).
2. Determine the amount of write I/O activity from the production and backup system for the expected duration of the FlashCopy relationship, that is the duration of the system save to tape.
3. Assuming that one track (64 KB) is moved to the repository for each write I/O and 33% of all writes are re-writes to the same track, calculate 50% contingency for the recommended space for the repository capacity as follows:

$$\text{Recommended repository capacity [GB]} = \text{write I/O/s} \times 67\% \times \text{FlashCopy active time [s]} \times 64 \text{ KB/I/O} / (1048576 \text{ KB/GB}) \times 150\%$$

For example, let us assume an i5/OS partition with a total disk space of 1.125 TB, a system save duration of 3 hours, and a given System i workload of 300 write I/O per second.

The recommended repository size is then is as follows:

$$300 \text{ I/O/s} \times 67\% \times 10800 \text{ s} \times 64 \text{ KB/I/O} / 1048576 \text{ GB/KB} \times 150\% = 199 \text{ GB}$$

So, the repository capacity needs to be 18% of its virtual capacity of 1.125 TB for the copy of the production system space.

### Sizing the repository number of disk arms

Because the workload to the shared repository volume has random I/O character, it is also important to provide enough physical disk arms for the repository volume space to ensure adequate performance.

To calculate the recommended number of physical disk arms for the repository volume space depending on your write I/O workload in tracks per second (at 50% disk utilization), refer to Table 4-4.

Table 4-4 Recommended number of physical arms

Disk Configuration	Tracks per second per disk arm
RAID5 15K RPM	25
RAID5 10K RPM	18
RAID10 15K RPM	50
RAID10 10K RPM	36

For example, if you are using RAID5 with 15 KB RPM drives and 600 I/O per second, your production host peak write I/O throughput during the active time of the space efficient FlashCopy relationship is 600 I/O per second x 67% (accounting for 33% re-writes), corresponding to 402 tracks per second and resulting in a recommended number of disk arms as follows:

$$402 \text{ tracks per second} / (25 \text{ tracks per second}) / \text{disk arm} = 16 \text{ disk arms of 15 KB RPM disks with RAID5}$$

## 4.3 Sizing tools

Several tools are available for sizing and performance measurements of the System i5 platform with external storage. In this section, we present the most important tools. Some of these tools are System i sizing tools such as the Workload Estimator. Other tools, such as Disk Magic, are IBM System Storage DS or Enterprise Storage Server (ESS) performance tools.

### 4.3.1 Disk Magic

Disk Magic is a tool for sizing and modeling disk systems for various servers. You can use it to model IBM and other disk systems attached to IBM System i, System p, System z®, System x™, and other servers. Disk Magic is developed by IntelliMagic and is available for download as follows:

- ▶ For IBM employees:

[http://w3-1.ibm.com/sales/systems/portal/\\_s.155/254?navID=f220s380&geoID=A11&prodID=Disk&docID=SSD5D00689DF4](http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s380&geoID=A11&prodID=Disk&docID=SSD5D00689DF4)

- ▶ For business partners: Sign on to the IBM PartnerWorld® Web site and search for Disk Magic:

[http://www-1.ibm.com/partnerworld/pwhome.nsf/weblook/index\\_emea\\_en.html](http://www-1.ibm.com/partnerworld/pwhome.nsf/weblook/index_emea_en.html)

To use Disk Magic for sizing the System i platform with the DS system, you need the following i5/OS Performance Tools reports:

- ▶ Resource report: Disk utilization section
- ▶ System report: Disk utilization section
- ▶ Optional: System report: Storage utilization section
- ▶ Component report: Disk Activity section

For instructions on how to use Disk Magic to size System i with a DS system, refer to 4.5, “Sizing examples with Disk Magic” on page 113, which presents several examples of using Disk Magic for the System i platform. The Disk Magic Web site also provides a Disk Magic Learning Guide that you can download, which contains a few step-by-step examples for using Disk Magic for modelling external storage performance.

### 4.3.2 IBM Systems Workload Estimator for i5/OS

IBM Systems Workload Estimator (WLE) is a tool that provides sizing recommendations for System i or iSeries models that are running one or more workloads. WLE recommends model, memory, and disk requirements that are necessary to meet reasonable performance expectations, based on inserted existing workloads or planned workloads.

To use WLE, you select one or more workloads from an existing selection list and answer a series of questions about each workload. Based on the answers, WLE generates a recommendation and shows the predicted processor utilization.

WLE also provides the capability to model external storage for recommended System i hardware. When the recommended System i models are shown in WLE, you can choose to directly invoke Disk Magic and model external storage for this workload. Therefore, you can obtain both recommendations for System i hardware and recommendations for external storage in the same run of WLE combined with Disk Magic.

For an example of how to use WLE with Disk Magic, see 4.5.4, “Using IBM Systems Workload Estimator connection to Disk Magic: Modeling DS6000 and System i for an existing workload” on page 163.

### 4.3.3 IBM System Storage Productivity Center for Disk

IBM System Storage Productivity Center is an integrated set of software components that provides end-to-end storage management, from the host and application to the target storage device in a heterogeneous platform environment. This software offering provides disk and tape library configuration and management, performance management, SAN fabric management and configuration, and host-centered usage reporting and monitoring from the perspective of the database application or file system.

IBM System Storage Productivity Center is comprised of the following elements:

- ▶ A *data* component: IBM System Storage Productivity Center for Data
- ▶ A *fabric* component: IBM System Storage Productivity Center for Fabric
- ▶ A *disk* component: IBM System Storage Productivity Center for Disk
- ▶ A *replication* component: IBM System Storage Productivity Center for Replication

IBM System Storage Productivity Center for Disk enables the device configuration and management of SAN-attached devices from a single console. In addition, it includes performance capabilities to monitor and manage the performance of the disks.

The functions of System Storage Productivity Center for Disk performance include:

- ▶ Collect and store performance data and provide alerts
- ▶ Provide graphical performance reports
- ▶ Help optimize storage allocation
- ▶ Provide volume contention analysis

When using System Storage Productivity Center for Disk to monitor a System i workload on DS8000 or DS6000, we recommend that you inspect the following information:

- ▶ Read I/O Rate (sequential)
- ▶ Read I/O Rate (overall)
- ▶ Write I/O Rate (normal)
- ▶ Read Cache Hit Percentage (overall)
- ▶ Write Response Time
- ▶ Overall Response Time
- ▶ Read Transfer Size
- ▶ Write Transfer Size
- ▶ Cache to Disk Transfer Rate
- ▶ Write-cache Delay Percentage
- ▶ Write-cache Delay I/O (I/O delayed due to NVS overflow)
- ▶ Backend Read Response Time
- ▶ Port Send Data Rate
- ▶ Port Receive Data Rate
- ▶ Total Port Data Rate (should be balanced among ports)
- ▶ Port Receive Response Time
- ▶ I/O per rank
- ▶ Response time per rank
- ▶ Response time per volumes

Figure 4-12 shows the read and write rate of the System Storage Productivity Center graph.

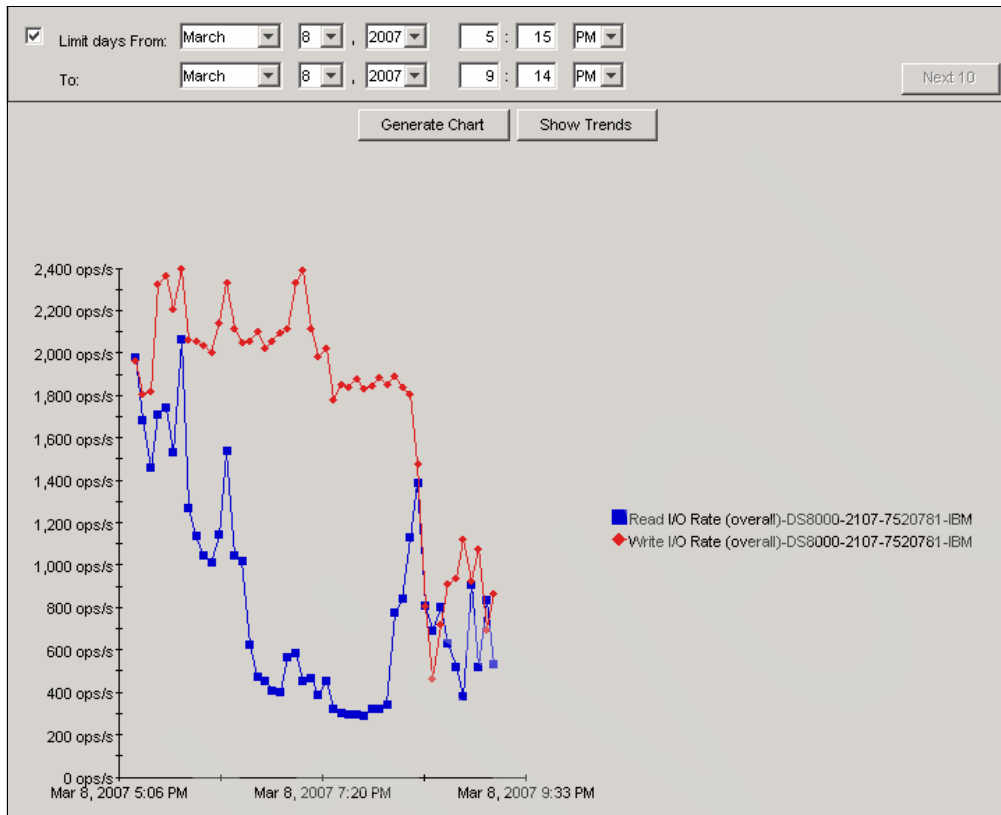


Figure 4-12 Read and write rate

Figure 4-13 shows the cache hit percentage of the System Storage Productivity Center graph.

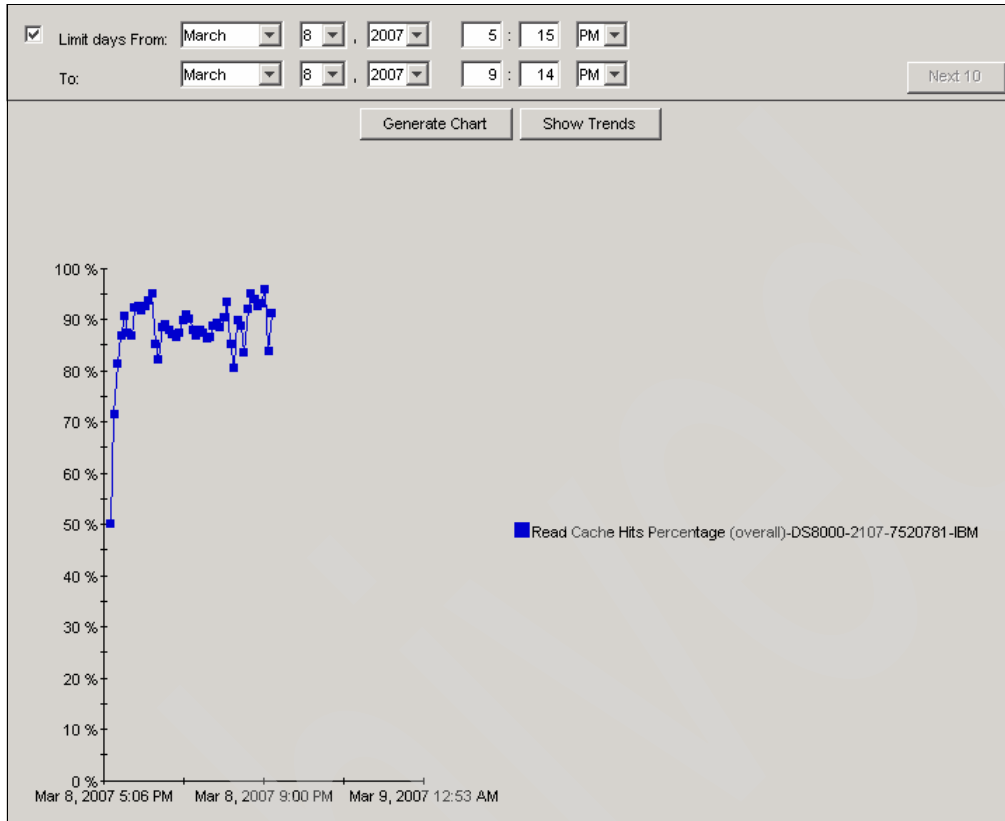


Figure 4-13 Read cache hit percentage

Figure 4-14 shows the write cache delay percentage of the System Storage Productivity Center graph.

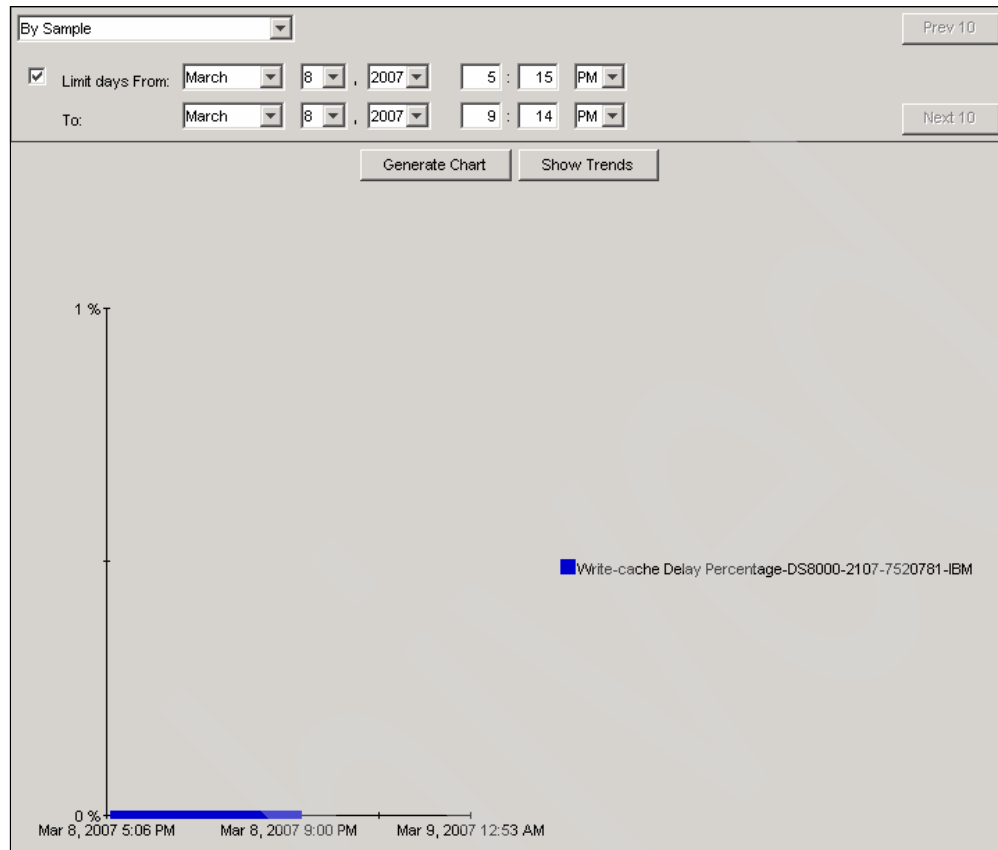


Figure 4-14 Write cache delay percentage

## 4.4 Gathering information for sizing

In this section, we discuss the methods and techniques for acquiring data for sizing the storage solution.

### 4.4.1 Typical workloads in i5/OS

To correctly size the DS system for the System i platform, it is important to know the characteristics of the workload that use the DS disk space. Many System i customer applications tend to follow the same patterns as the System i benchmark commercial processing workload (CPW). These applications typically have many jobs that run brief transactions with database operations.

Other applications tend to follow the same patterns as the System i benchmark compute intensive workload (CIW). These applications typically have fewer jobs running transactions that spend a substantial amount of time in the application itself. An example of such a workload is Lotus® Domino® Mail and Calendar.

In general, System i batch workloads can be I/O or compute intensive. For I/O intensive batch applications, the overall batch performance is dependent on the speed of the disk subsystem.

For compute-intensive batch jobs, the run time likely depends on the processor power of the System i platform. For many customers, batch workloads run with large block sizes.

Typically batch jobs run during the night. For some environments, it is important that these jobs finish on time to enable timely starting of the daily transaction application. The amount of time that a batch job takes is called a *batch window*.

## 4.4.2 Identifying peak periods

To size a DS system for an i5/OS system, we recommend that you identify one or two peak periods, each of them lasting one hour, and collect performance data during these periods. For instructions how to collect performance data and produce reports, refer to 4.4.3, “i5/OS Performance Tools” on page 111.

In many cases, you know when the peak periods or the most critical periods occur. If you know when these times are, collect performance data during these periods. In some cases, you might not know when the peak periods occur. In such a case, we recommend that you collect performance data during a 24-hour period and in different time periods, for example, during end-of-week and end-of-month jobs.

After the data is collected, produce a Resource report with a disk utilization section and use the following guidelines to identify peak periods:

- ▶ Look for one hour with the most I/O per seconds. You can insert the report into a spreadsheet, calculate the hourly average of I/O per second, and look for the maximum of the hourly average. Figure 4-15 shows part of such a spreadsheet.
- ▶ For many customers, performance data shows patterns in block sizes, with significantly different block sizes in different periods of time. If this is so, calculate the hourly average of the block sizes and use the hour with the maximal block sizes as the second peak.
- ▶ If you identified two peak periods, size the DS system so that both are accommodated.

	IO/sec	Reads/sec	Writes/sec	Blksize	High util	
31. jan 23-0	237,58	102,18	135,35	14,29	11,08	
1. feb 0-1	506,93	312,96	193,93	28,71	12,18	
1. feb 1-2	321,59	212,95	108,58	31,16	7,47	
1. feb 2-3	112,06	57,04	54,97	21,11	3,00	
1. feb 3-4	230,35	163,13	67,18	30,66	4,53	
1. feb 4-5	244,73	128,20	116,49	25,20	5,48	
1. feb 5-6	115,22	27,37	87,81	16,03	2,70	
1. feb 6-7	708,03	570,53	137,44	11,28	13,86	
1. feb 7-8	481,61	340,03	141,54	14,38	8,31	
1. feb 8-9	877,73	582,63	295,04	17,21	14,47	
1. feb 9-10	831,39	518,83	312,54	22,95	14,80	*****
1. feb 10-11	631,93	488,35	143,52	17,51	9,95	

Figure 4-15 Identifying the peak period for the System Storage Productivity Center

## 4.4.3 i5/OS Performance Tools

To use the sizing rules of thumb and Disk Magic, you need the following performance reports from i5/OS:

- ▶ *System Report*: Disk Utilization and Storage Pool utilization sections
- ▶ *Component Report*: Disk Activity section
- ▶ *Resource Report*: Disk Utilization section

To produce the System i5 performance reports that are needed for sizing the DS system:

1. Install the licensed program Performance Tools 5722-PT1 on i5/OS.
2. On the i5/OS command line, enter the GO PERFORM command.
3. In the IBM Performance Tools for i5/OS panel that opens, select **2. Collect Performance Data** as shown in Figure 4-16.

```
PERFORM                IBM Performance Tools for i5/OS                System:  RCHLTTN1
Select one of the following:
    1. Select type of status
    2. Collect performance data
    3. Print performance report
    5. Performance utilities
    6. Configure and manage tools
    7. Display performance data
    8. System activity
    9. Performance graphics
   10. Advisor
    70. Related commands
Selection or command
====> 2
F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel  F13=Information Assistant
F16=System main menu
```

Figure 4-16 *PERFORM* menu panel

4. On the Collect Performance Data panel, select **1. Start Collecting Data**.
5. On the Start Collecting Data panel, specify the collection interval as 15 minutes or 5 minutes, and press Enter. i5/OS starts collecting the performance data.
6. After a period of time, on the Collect Performance Data panel, select **2. Stop collecting data**.
7. On the IBM Performance Tools for iSeries panel, select **3. Print performance Report**.
8. On the Print Performance Report - Sample Data panel, make sure that the listed library at the field Library is the one to which you collected data. You might need to change the name of the library. For the member, select **1. System report**, and press Enter.
9. On the Select Section for Report panel, select **Disk Utilization and Storage Pool Utilization**, and press Enter.
10. On the Select Categories for Report panel, select **Time Interval**.
11. On the next panel, you can select the intervals for the report. If you collected data for 24 hours and then identified a peak hour, select only the intervals of this particular hour. Select the intervals and press Enter. This job starts and produces report in a spooled file.
12. On Print Performance Report - Sample Data panel, for member, select **2. Component report**.



13. On the Select Section for Report panel, select **Disk Activity**, and then select **Time Interval**. Then select all intervals or just the intervals of the peak period. Press Enter to start the job for report.
14. On the Print Performance report - Sample Data panel, for member, select **5. Resource report**.
15. On the Select Section for Report panel, select **Disk Utilization** and then select **Time Interval**. Then select all intervals or just the intervals of the peak period. Press Enter to start the job for the report.
16. To insert the reports into Disk Magic, transfer the reports from the spooled file to a PC using iSeries Navigator.
17. In iSeries Navigator, expand the **i5/OS system** on which the reports are located. Expand **Basic Operations** and double-click **Printer output**.
18. Performance reports in the spooled file are shown on the right side of the panel. Copy and paste the necessary reports to your PC.

## 4.5 Sizing examples with Disk Magic

In this section, we describe three examples of using Disk Magic to size the DS system for the System i platform.

### 4.5.1 Sizing the System i5 with DS8000 for a customer with iSeries model 8xx and internal disks

In this example, DS8000 is sized for a customer's production workload. The customer is currently running a host workload on internal disks; performance reports from a peak period are available. For instruction on how to produce performance reports, refer to 4.4.3, "i5/OS Performance Tools" on page 111.

To size DS8000 using Disk Magic:

1. On the Welcome to Disk Magic panel (Figure 4-17), select **Open and iSeries Automated Input (\*.IOSTAT, \*.TXT, \*.CSV)** and click **OK**.

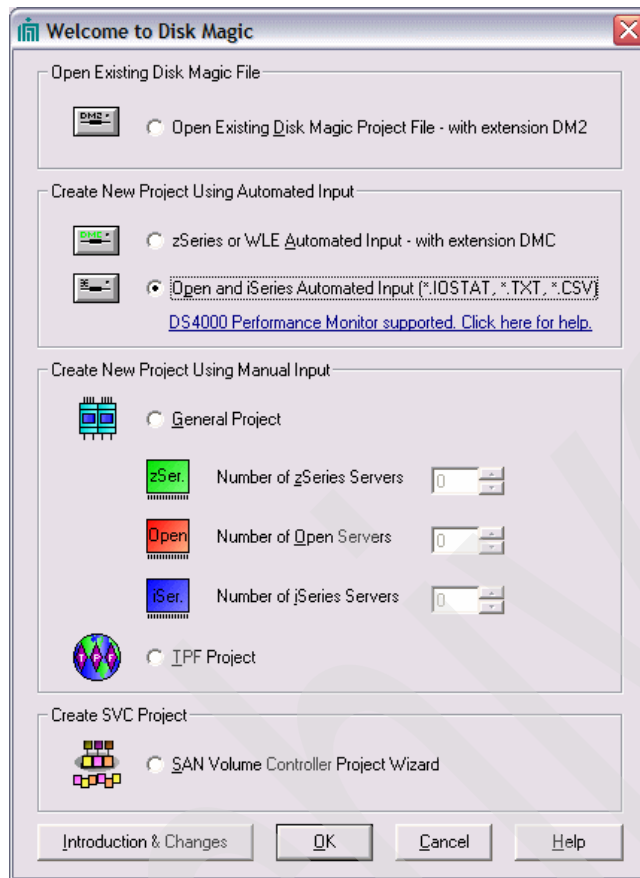


Figure 4-17 Disk Magic Welcome dialog box

2. In the Open window (Figure 4-18), choose the directory that contains the performance reports, select altogether the corresponding system, resource and component report files and click **Open**.

You can also concatenate all necessary iSeries performance reports into one file and insert it into Disk Magic. In this example, both System report - Storage pool utilization and System report - Disk utilization are concatenated into one System report file.

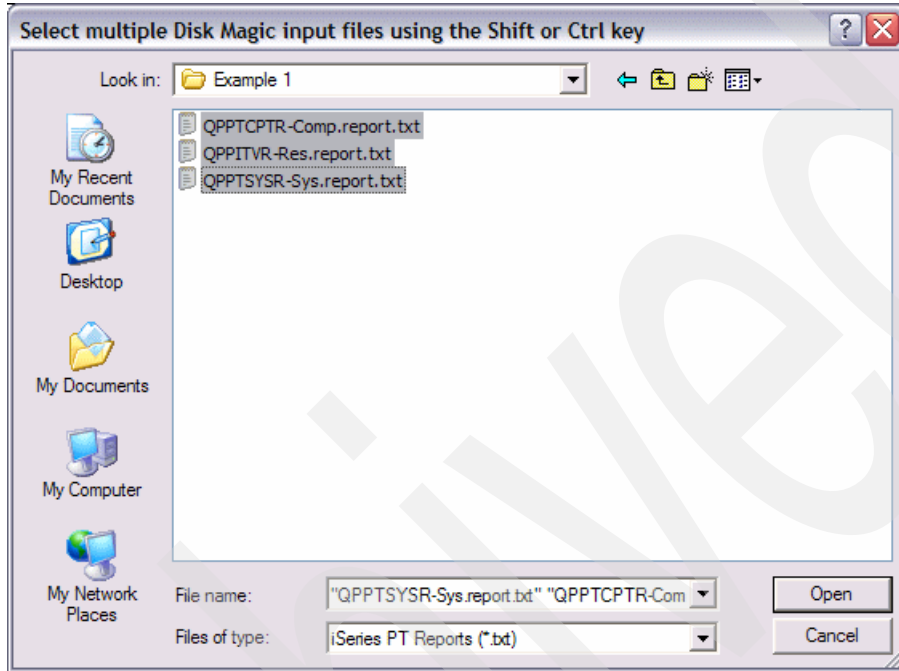


Figure 4-18 Inserting PT reports to Disk Magic

3. Disk Magic shows you an overview of the read performance files in the Multiple File Open - File Overview window as shown in Figure 4-19. By default Disk Magic accounts for different ASPs by treating them as separate I/O workloads allowing to model the external storage performance on an ASP level.

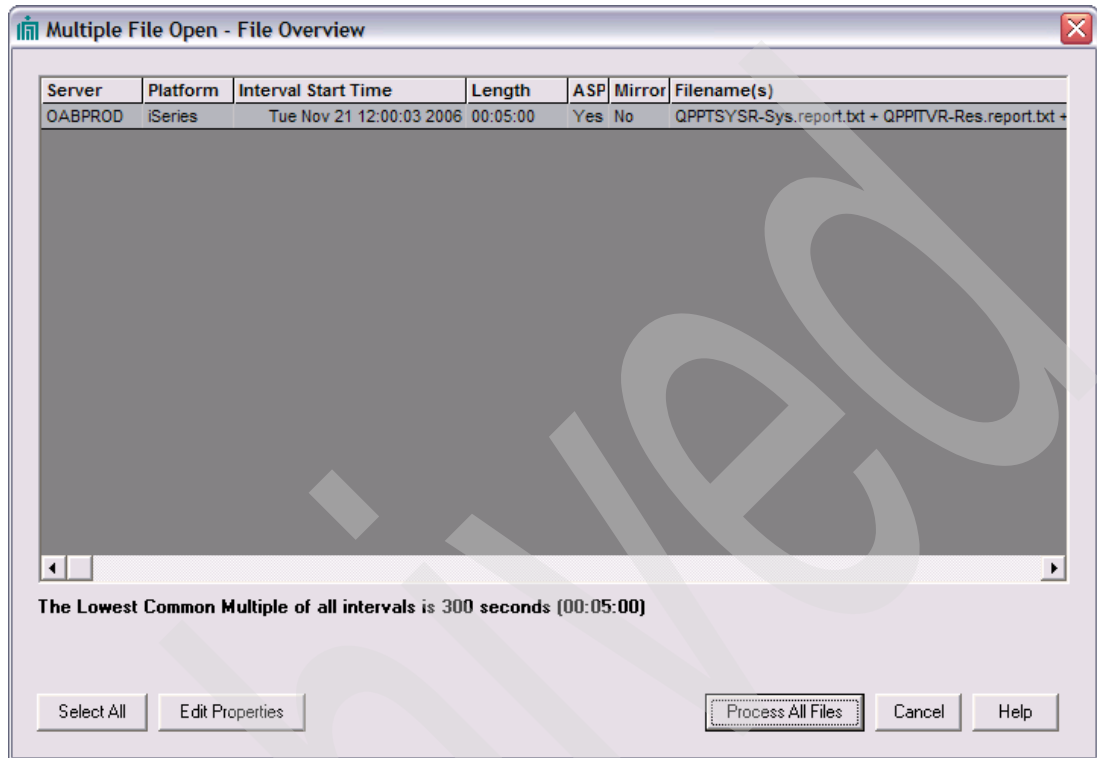


Figure 4-19 Disk Magic - File Overview dialog box

If you want to model your external storage solution with a system I/O workload aggregated from all fASPs or if you want to continue using potentially configured i5/OS mirroring with external storage:

- a. Click **Edit Properties**.
- b. Click **Discern ASP level**.
- c. Select **Keep mirroring, if applicable**,
- d. Click **OK** as shown in Figure 4-20.

Otherwise, click **Process All Files** (in Figure 4-19) to continue.

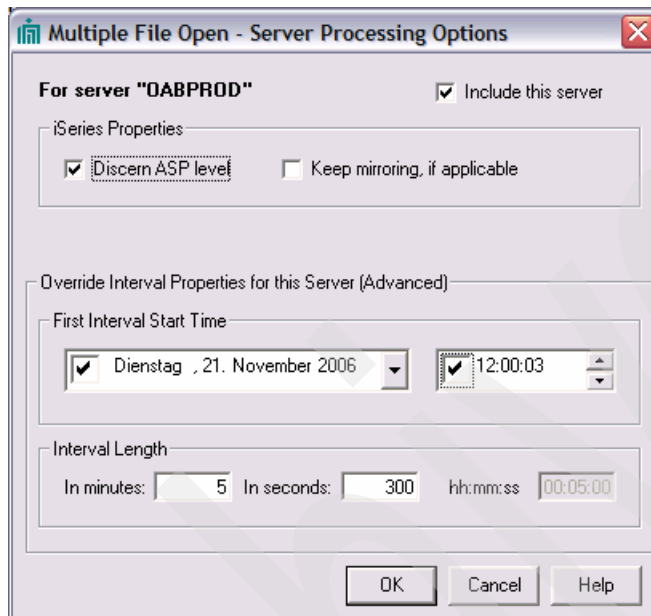


Figure 4-20 Disk Magic - Server processing options

While inserting reports, Disk Magic might show a warning message about inconsistent interval start and stop times (see Figure 4-21).

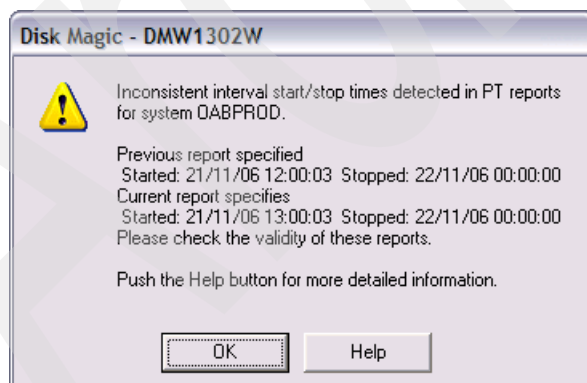


Figure 4-21 Inconsistent start/stop times message

One cause for inconsistent start and stop times might be that the customer gives you performance reports for 24 hours, and you select a one-hour peak period from them. Then the customer produces reports again and selects only the interval of the peak period from the collected data. In such reports, the start and stop time of the collection does not match the start and stop time of produced reports. The reports are correct, and you can ignore this warning. However, there can be other instances where inconsistent reports are

inserted by mistake, so we recommend that you resolve this issue by getting a set of consistent reports.

- After successful processing of the performance report files Disk Magic shows the I/O Load Summary as shown in Figure 4-22. Click **Create Model** to proceed with the external storage performance modeling.

Interval Start Time	Servers	I/O Rate	Read%	Write%	R/W Ratio	MB/s	W MB/s	R MB/s	Serv Time	kB/W	kB/R
Tue Nov 21 17:20:03 2006	1	3,434.3	90.0	10.0	9.0	14.7	1.5	13.2	0.2	4.4	4.4
Tue Nov 21 17:25:03 2006	1	2,417.8	84.7	15.3	5.5	23.0	3.5	19.5	0.3	9.8	9.8
Tue Nov 21 17:30:03 2006	1	768.3	68.4	31.6	2.2	12.2	3.9	8.4	0.8	16.3	16.3
Tue Nov 21 17:35:03 2006	1	414.4	49.1	50.9	1.0	6.6	3.4	3.3	1.1	16.4	16.4
Tue Nov 21 17:40:03 2006	1	1,502.6	53.8	46.2	1.2	14.1	6.5	7.6	0.4	9.6	9.6
Tue Nov 21 17:45:03 2006	1	605.2	18.6	81.4	0.2	13.7	11.1	2.5	0.8	23.1	23.1
Tue Nov 21 17:50:03 2006	1	493.4	29.0	71.0	0.4	32.0	22.7	9.3	1.2	66.4	66.4
Tue Nov 21 17:55:03 2006	1	583.8	27.2	72.8	0.4	30.7	22.4	8.4	1.6	53.9	53.9
Tue Nov 21 18:00:03 2006	1	612.4	25.5	74.5	0.3	32.7	24.4	8.3	1.1	54.7	54.7
Tue Nov 21 18:05:03 2006	1	702.3	22.2	77.8	0.3	30.4	23.7	6.8	1.0	44.4	44.4
Tue Nov 21 18:10:03 2006	1	497.5	13.2	86.8	0.2	10.6	9.2	1.4	0.5	21.9	21.9
Tue Nov 21 18:15:03 2006	1	933.0	20.4	79.6	0.3	7.3	5.9	1.5	0.5	8.1	8.1
Tue Nov 21 18:20:03 2006	1	1,639.6	32.9	67.1	0.5	11.3	7.6	3.7	0.5	7.1	7.1
Tue Nov 21 18:25:03 2006	1	608.2	11.2	88.8	0.1	4.6	4.1	0.5	0.3	7.7	7.7
Tue Nov 21 18:30:03 2006	1	503.3	7.3	92.7	0.1	3.7	3.4	0.3	0.2	7.4	7.4
Tue Nov 21 18:35:03 2006	1	610.0	6.9	93.1	0.1	4.6	4.3	0.3	0.2	7.7	7.7
Tue Nov 21 18:40:03 2006	1	431.1	7.1	92.9	0.1	3.2	3.0	0.2	0.2	7.7	7.7
Tue Nov 21 18:45:03 2006	1	479.6	7.6	92.4	0.1	3.7	3.4	0.3	0.2	8.0	8.0
Tue Nov 21 18:50:03 2006	1	603.5	6.8	93.2	0.1	4.4	4.1	0.3	0.2	7.5	7.5

Figure 4-22 Disk Magic - Successfully imported performance reports

- In the TreeView panel in Disk Magic, observe the following two icons (Figure 4-23):
  - Example1 denotes a workload.
  - iSeries1 denotes a disk subsystem for this workload.

Double-click **iSeries1**.




Figure 4-23 Selecting the disk subsystem

6. The Disk Subsystem - iSeries1 panel displays, which contains data about the current workload on the internal disks. The General tab shows the current type of disks (Figure 4-24).

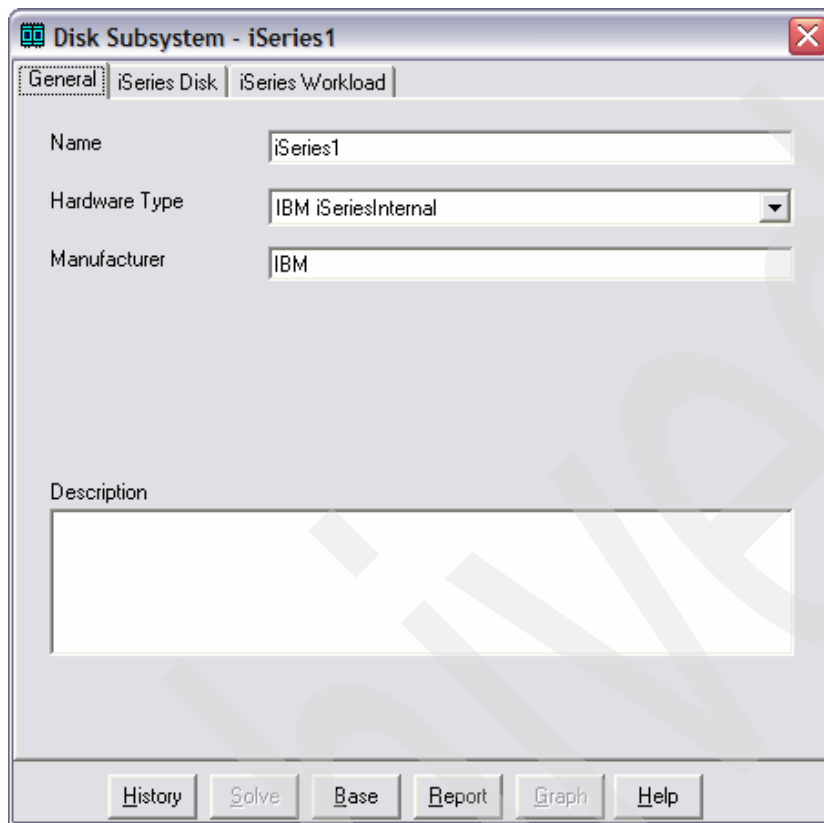


Figure 4-24 Disk Subsystem window: General tab

The iSeries Disk tab on the Disk Subsystem - iSeries1 window shows the current capacity and number of disk devices (Figure 4-25).

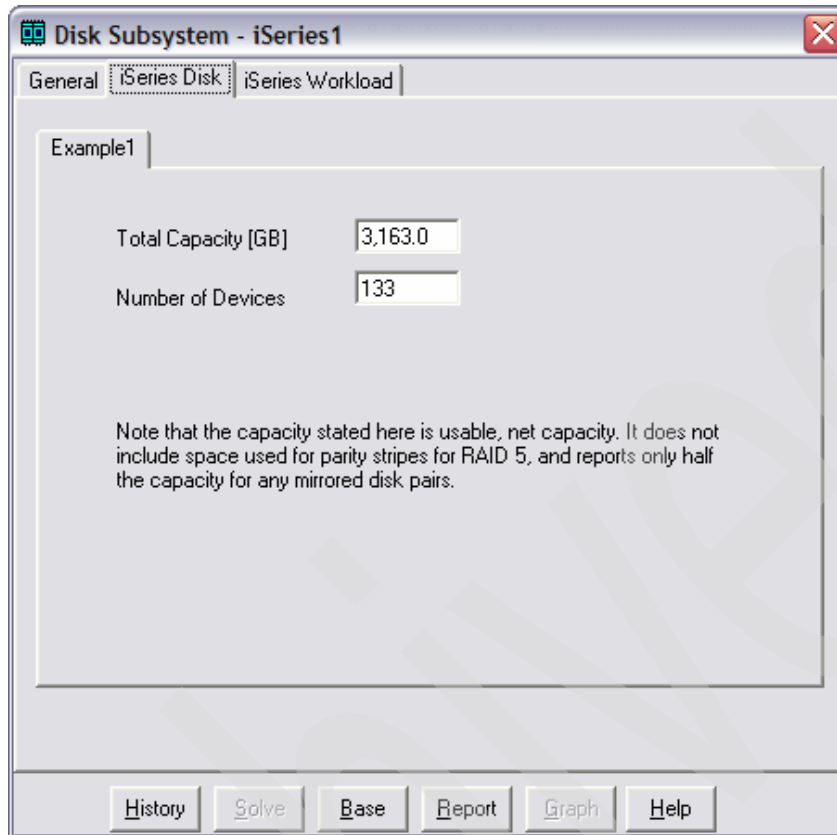


Figure 4-25 Disk subsystem window: iSeries Disk tab



The iSeries Workload tab on the same panel (Figure 4-26) shows the characteristics of the iSeries workload. These include reads per sec, writes per sec, block size, and reported current disk service time and wait time.

- a. Click the **Cache Statistics** button.

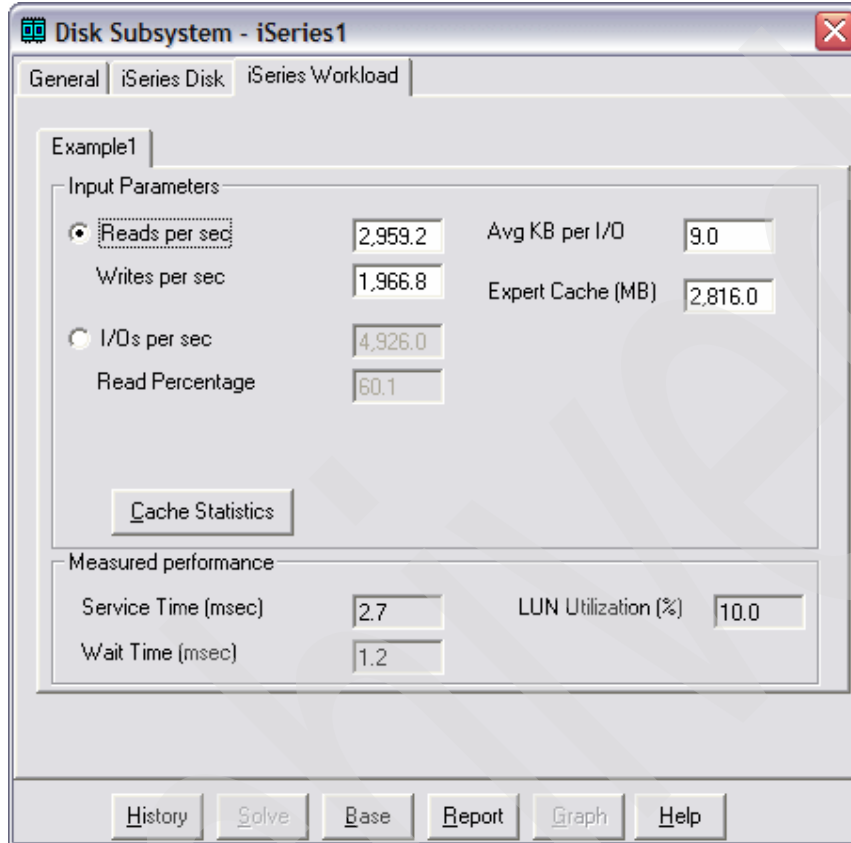


Figure 4-26 Disk Subsystem: iSeries Workload tab

- b. You can observe the current percentage of cache read hits and write efficiency as shown in Figure 4-27. Click **OK** to return to the iSeries Workload tab.




Figure 4-27 Cache statistics of workload on internal disk

- c. Click **Base** to save the current disk subsystem as a base for Disk Magic modeling.

- d. Disk Magic informs you that the base is created successfully, as shown in Figure 4-28. Click **OK** to save the base.



Figure 4-28 Saving the base

7. Insert the planned DS configuration in the disk subsystem model by inserting the relevant values on each tab, as shown in the next steps. In this example, we insert the following planned configuration:

- DS8100 with 32 GB cache
- 12 FC adapters in System i5 in multipath, two paths for each set of LUNs
- Six FC ports in DS8100
- Eight ranks of 73 GB DDMs used for the System i5 workload
- 182 LUNs of size 17.54 GB

To insert the planned DS configuration information:

- a. On the General tab in the Disk Subsystem - iSeries 1 window, choose the type of planned DS for Hardware Type (Figure 4-29).

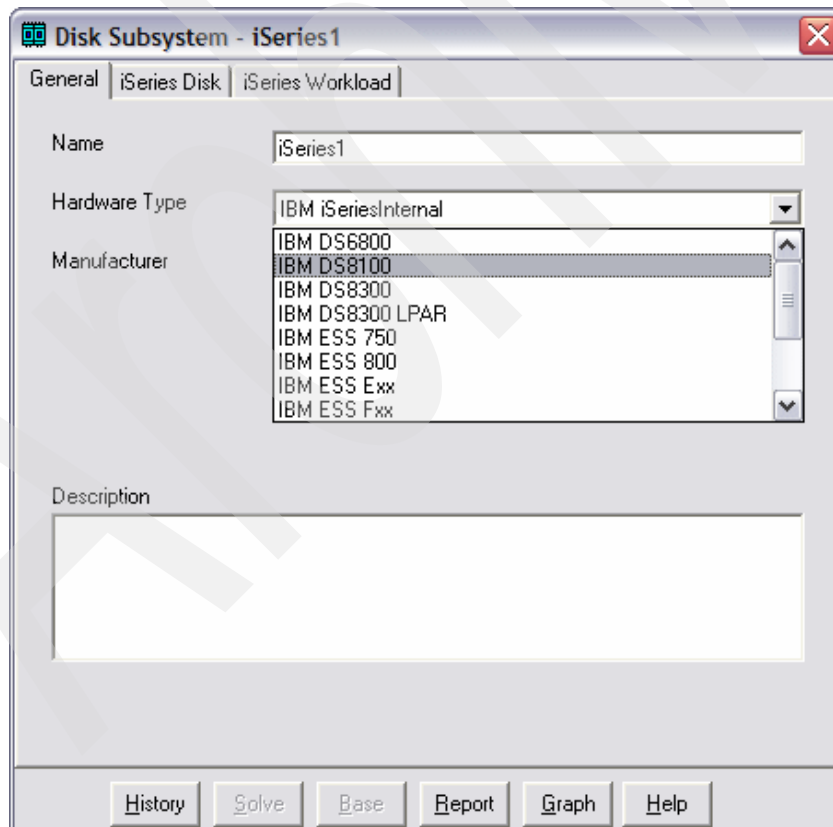


Figure 4-29 Inserting a planned type of DS

Notice that the General tab interface changes as shown in Figure 4-30. If you use multipath, select **Multipath with iSeries**. In our example, we use multipath, so we select this box. Notice that the Interfaces tab is added as soon as you select DS8000 as a disk subsystem.

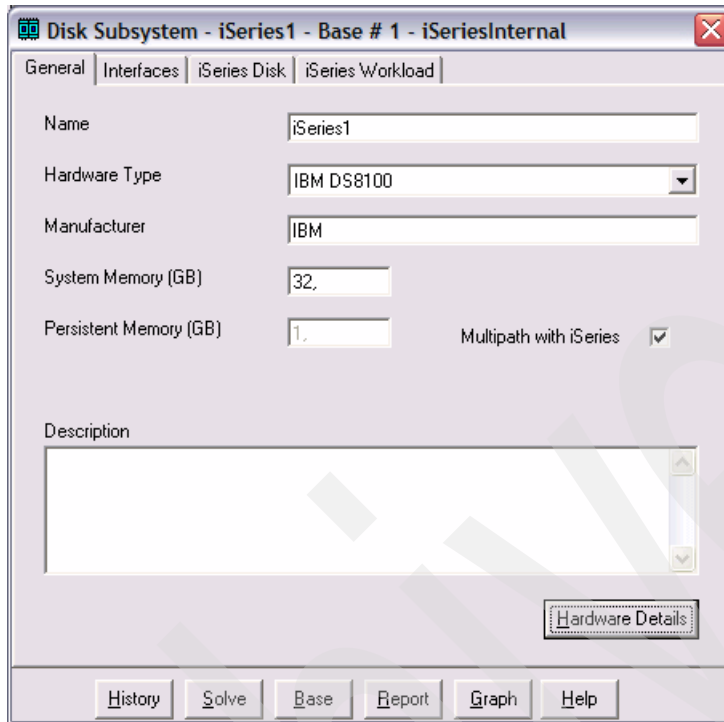


Figure 4-30 Disk Magic: Selecting the hardware and specifying multipath

- b. Click the **Hardware Details** button. In the Hardware Details window (Figure 4-31), for System Memory, choose the planned amount of cache, and for Fibre Host Adapters, enter the planned number of host adapters, and click **OK**.

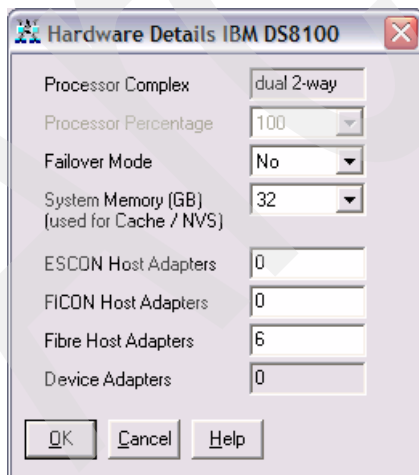


Figure 4-31 Disk Magic: Specifying hardware details of DS

- c. Next, in the Disk Subsystem - iSeries1 window, select the Interfaces tab, as shown in Figure 4-32. On the Interfaces tab, under the From Disk Subsystem tab, click the **Edit** button.

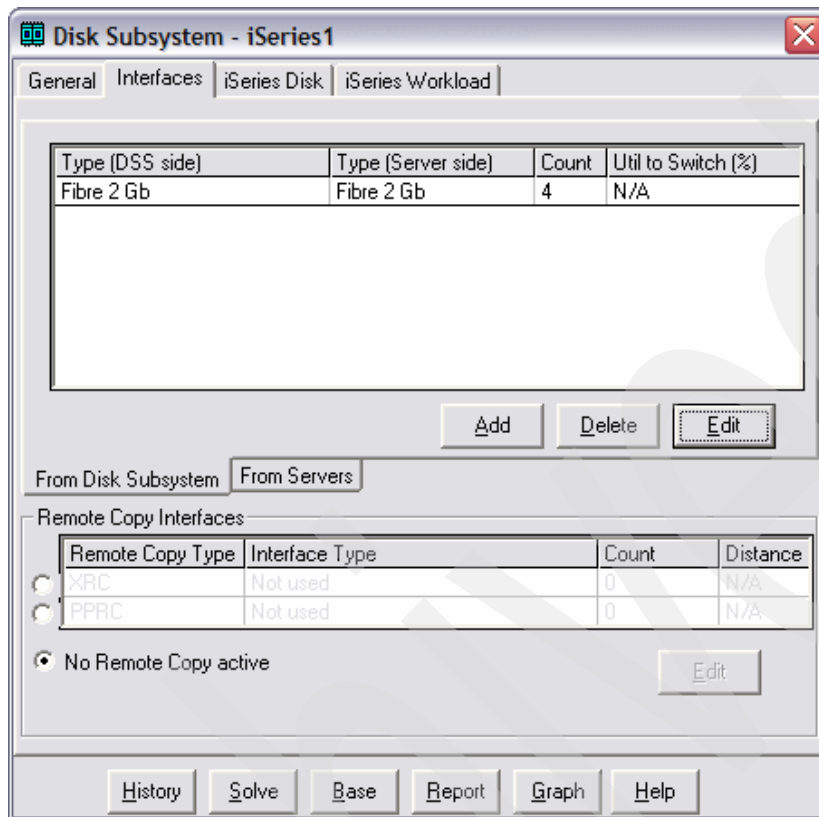


Figure 4-32 Specifying the DS host ports: Interfaces tab

- d. In the Edit Interfaces for Disk Subsystem window (Figure 4-33), for Count, enter the planned number of DS ports, and click **OK**.




Figure 4-33 Inserting the DS host ports: Edit Interfaces for Disk Subsystem

- e. Back on the Interfaces tab (Figure 4-32), select the From Servers tab, and click **Edit**. In the Edit Interfaces window (Figure 4-34), enter the number of planned System i5 FC adapters. Click **OK**.

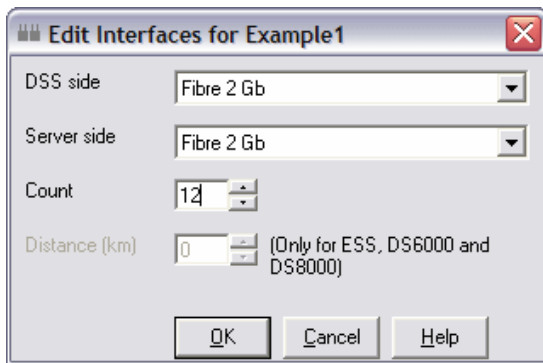


Figure 4-34 Inserting the System i5 FC adapters

- f. Next, in the Disk Subsystem - iSeries1 window, select the iSeries Disk tab, as shown in Figure 4-35. Notice that Disk Magic uses the reported capacity on internal disks as the default capacity on DS. Click **Edit**.

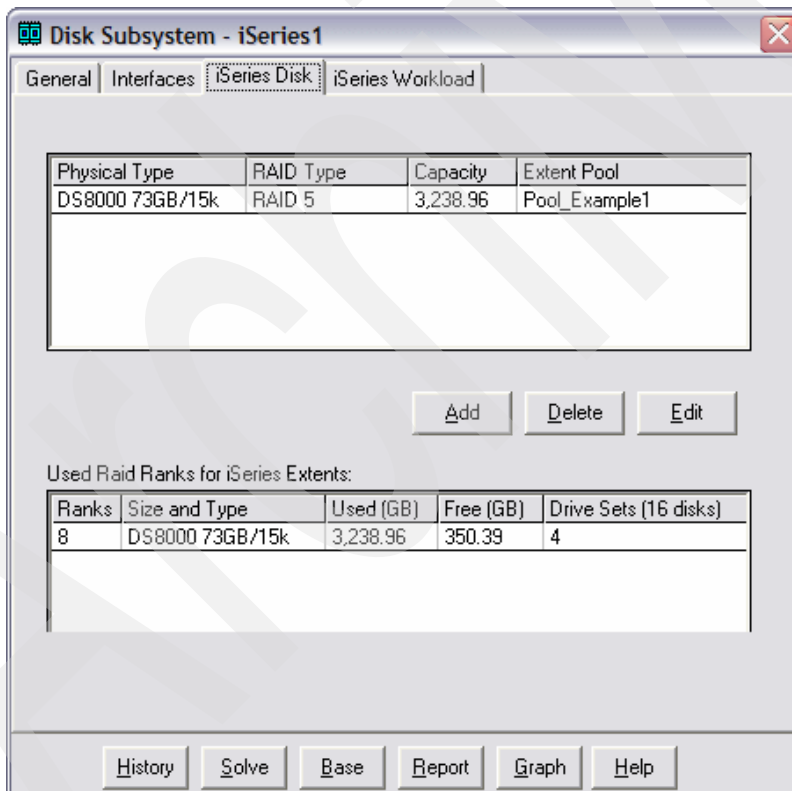


Figure 4-35 iSeries disk

- g. In the Edit a Disk Type window (Figure 4-36), enter the desired capacity to achieve modeling of the planned number of ranks.

In our example, we enter the capacity of planned eight ranks. Each RAID-5 rank with a spare disk (6+P+S rank) has 415 GB of effective capacity, and a RAID-5 rank without a spare disk (7+P ranks) has 483 GB of effective capacity. For Disk Magic modeling, we assume that only 6+P ranks are used, so we plan for 8 x 415 GB = 3320 GB capacity.

Refer to 3.2.6, “Planning for capacity” on page 67 for more information about available capacity.

The actual capacity used by i5/OS is specified in the Workload window. The capacity might be lower than the capacity that was specified in this panel. This is so because you cannot allocate all available capacity i5/OS. If you do, the capacity used by i5/OS will be lower because of fixed LUN sizes for i5/OS. Refer to Chapter 3, “i5/OS planning for external storage” on page 51 for more information about LUN sizes.

Observe that 73 GB DDMs and RAID-5 protection are the default values in this panel. Notice also that a default extent pool for iSeries workload is created in Disk Magic.




Figure 4-36 Inserting the capacity for the planned number of ranks

After you insert the capacity for the planned number of ranks, the iSeries Disk tab shows the correct number of planned ranks (see Figure 4-37).

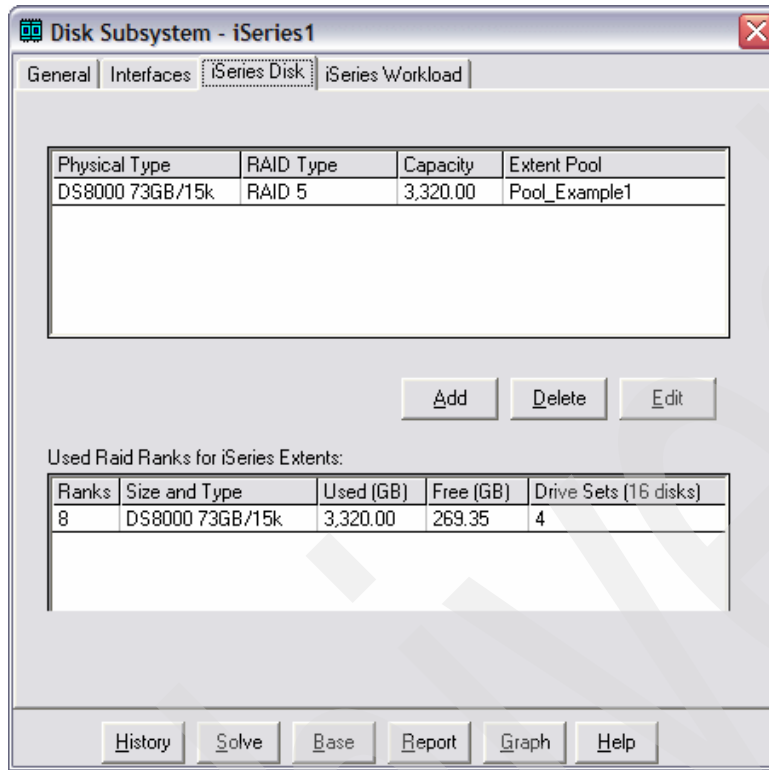


Figure 4-37 Planned number of ranks

- h. Finally, select the iSeries Workload tab. Specify the planned number of LUNs and the usable capacity for i5/OS.

In our example, we use 182 x 17.54 GB LUNs, so the usable capacity for i5/OS is 3192 GB (see Figure 4-38). We recommend that you create one extent pool from one DS rank. Nevertheless in Disk Magic, you can model one extent pool that contains all planned ranks, because modeled values do not depend on the way in which extent pools are specified in Disk Magic. In our example, we use only the extent pool created by Disk Magic as the default.

Click **Cache Statistics**.

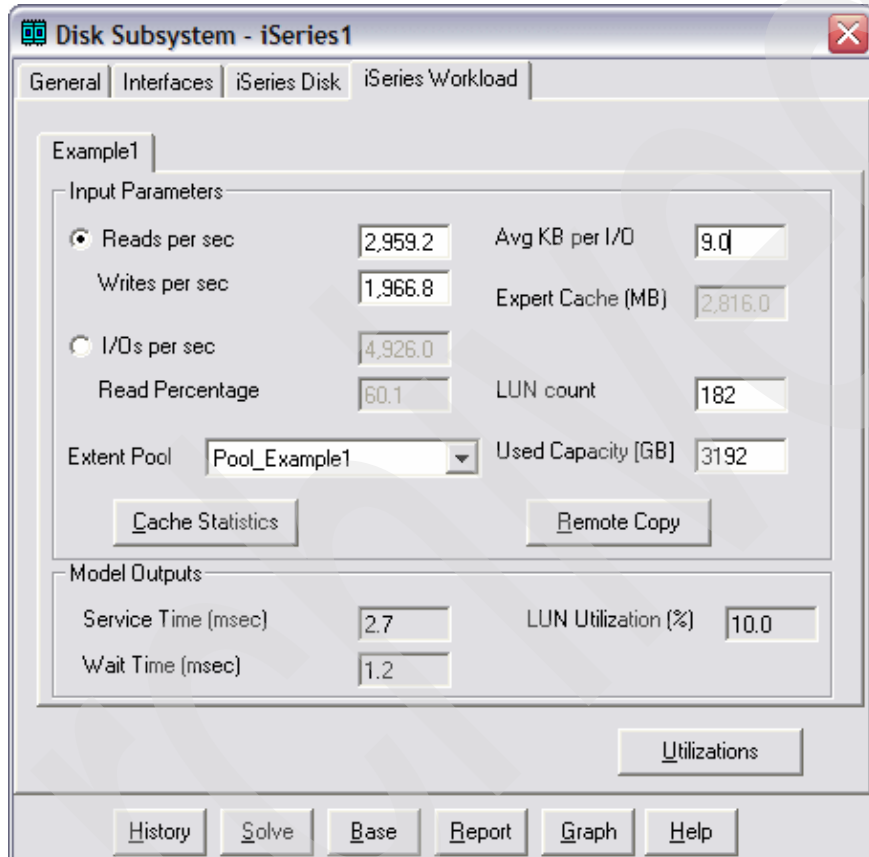


Figure 4-38 Planned number of LUNs



- i. In the Cache Statistics for Host window (see Figure 4-39), notice that Disk Magic models cache usage on DS8000 automatically based on the reported current cache usage on internal disks. Click **OK**.

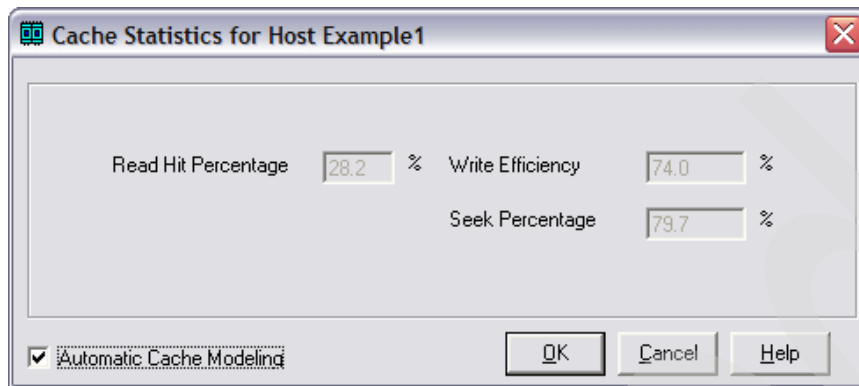


Figure 4-39 Automatic cache modeling

8. After you enter the planned values of the DS configuration, in the Disk Subsystem - iSeries1 panel (Figure 4-38), click **Solve**.
9. A Disk Magic message displays indicating that the model of planned scenario is successfully solved (Figure 4-40). Click **OK** to solve the model of iSeries or i5/OS workload on DS.




Figure 4-40 Solving the model of planned scenario

10. After you solve the model of the planned scenario, on the iSeries Workload tab (Figure 4-41), notice the modeled disk service time and wait time. Click **Utilizations**.

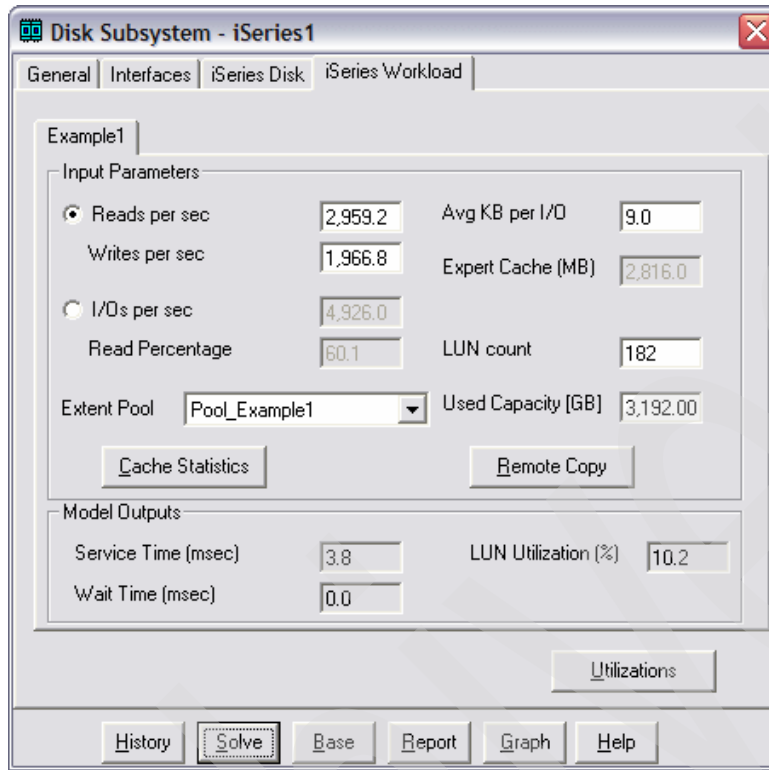


Figure 4-41 Modeled disk service time and wait time

11. In the Utilizations IBM DS8100 window (Figure 4-42), observe the modeled utilization of physical disk drives or hard disk drives (HDDs), DS device adapters, LUNs, FC ports in DS, and so on.

In our example, none of the utilization values exceeds the recommended maximal value. However, the HDD utilization of 32% approaches the recommended threshold of 40%. Thus, you need to consider additional ranks if you intend to grow the workload. Click **OK**.

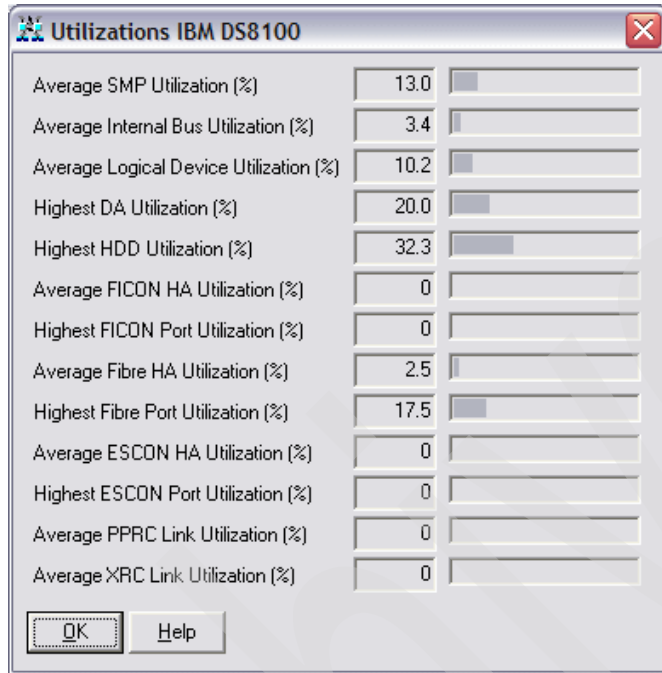


Figure 4-42 Modeled utilizations

12. On iSeries Workload tab (Figure 4-41), click **Cache Statistics**. In the Cache Statistics for Host window (Figure 4-43), notice the modeled cache values on DS. In our example, the modeled read cache percentage is higher than the current read cache percentage with internal disks, but modeled write cache efficiency on DS is about the same as current rather high write cache percentage. Notice also that the modeled disk seek percentage dropped to almost half of the reported seek percentage on internal disks.

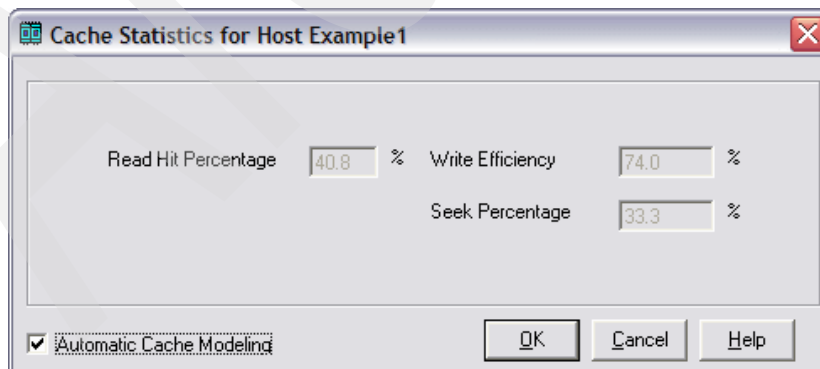


Figure 4-43 Modeled cache hits

You can also see modeled utilizations, disk service, wait times, and cache percentages in the Disk Magic log, as shown in Figure 4-44.

Cache Size / Backstore Sensitivity		6.0								
Advanced DS6000/DS8000 Outputs:										
Processor Utilization:	13.0%									
Highest HDD Utilization:	32.3%									
Back End Interface Utilization:	20.0%									
Internal Bus Utilization:	3.4%									
Avg. Host Adapter Utilization:	2.5%									
Avg. Host Interface Utilization:	17.5%									
Extent Pool	Type	HDD	RAID	Devices	GBytes	Log.Type				
Pool_Example1	FBiSeries	73GB/15k	RAID 5	182	3320.0	LUN				
Extent Pool	I/O Rate	IOSQ	Pend	Conn	Disc	Resp Time	Highest HDD Util			
Pool_Example1	4926.0	0.0	---	---	---	3.8	32.3%			
iSeries Server	I/O Rate	Transfer Size (KB)	Serv Time	Wait Time	Read Perc	Read Hit%	Write Hit%	Write Eff %	LUN Cnt	LUN Util%
Average	4926	9.0	3.8	0.0	60	41	100	74	182	10
Example1	4926	9.0	3.8	0.0	60	41	100	74	182	10

Figure 4-44 Modeled values in the Disk Magic log

13. You can use Disk Magic to model the critical values for planned growth of a customer's workload, which can be predicted to a point at which the current DS configuration no longer meet performance requirements and the customer must consider additional ranks, FC adapters, and so on. To model DS for growth of the workload:

a. In the Disk Subsystem - iSeries1 window, click **Graph**. In the Graph Options window (Figure 4-45), select the following options:

- For Graph Data, choose **Response Time in ms**.
- For Graph Type, select **Line**.
- For Range Type, select **I/O Rate**.

Observe that the values for range of I/O rate are already filled with default values, starting from current I/O rate. In our example, we predict a growth rate of three times larger than the current I/O rate, increasing by 1000 I/O per second at a time. Therefore, we insert 14800 in the To field and 1000 in the By field.

b. Click **Plot**.

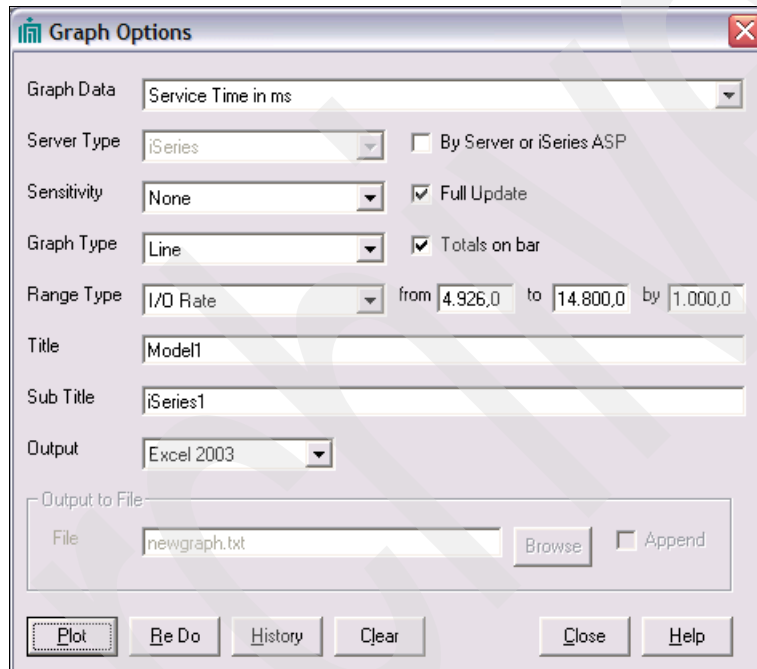


Figure 4-45 Graph options for disk response time

A spreadsheet is created that contains a graph with the predicted disk response time (service time + wait time) for I/O rate growth. Figure 4-46 shows the graph for our example. Notice that at about 9000 I/Os per second, the predicted response time will exceed 5 ms, which we consider as a high limit for good response time. At about 12000 I/Os per second, disk response time will go over 7 ms and start to drastically increase.

The customer can increase the I/O rate to about 9000 I/Os per second with a disk response time that is still acceptable. If the customer increases the I/O rate even more, the disk response time increases accordingly, but at about 12000 I/Os per second, the current DS configuration is saturated.




Figure 4-46 Disk response time at I/O growth

14. Next, produce the graph of HDD utilizations at workload growth.

- a. In the Disk Subsystem - iSeries1 window, on the iSeries Workload tab, click **Graph**. In the Graph Options window (Figure 4-47):
  - For Graph Data, select **Highest HDD Utilization (%)**.
  - For Graph Type, select **Line**.
  - For Range Type, select **I/O Rate** and select the appropriate range values. In our example, we use the same I/O rate values as for disk response time.
- b. Click **Plot**.

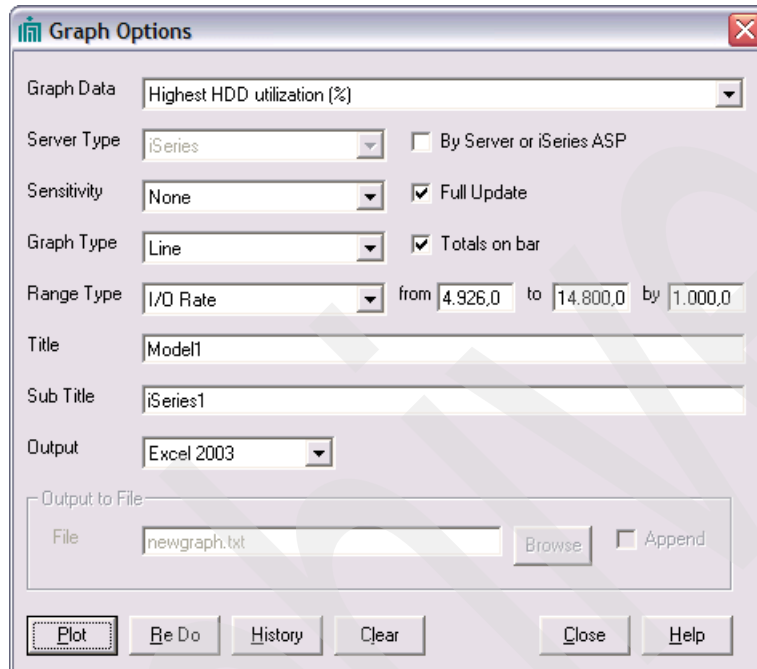


Figure 4-47 Graph options for HDD utilization

A spreadsheet is generated with the desired graph. Figure 4-48 shows the graph for our example. Notice that the recommended 40% HDD utilization is exceeded at about 6000 I/O per second, and 70% is exceeded at about 11000 I/O per second, which confirms that the current configuration is saturated at 11000 to 12000 I/O per second.




Figure 4-48 HDD utilization at I/O rate growth

After the installing the System i5 platform and DS8100, the customer used initially six ranks and 10 FC adapters in multipath for the production workload. Because an iSeries model 825 replaced a System i5 model, the I/O characteristics of the production workload changed, because of higher processor power and larger memory pool in the System i5 model. The production workload produces 230 reads per second and 1523 writes per second. Also, the actual service times and wait times do not exceed one millisecond.



## 4.5.2 Sharing DS8100 ranks between two i5/OS systems (partitions)

In this example, we use Disk Magic to model two i5/OS workloads that share the same extent pool in DS8000. To model this scenario with Disk Magic:

1. Insert into Disk Magic reports of the first workload as described in 4.5.1, “Sizing the System i5 with DS8000 for a customer with iSeries model 8xx and internal disks” on page 113.
2. After reports of the first i5/OS system are inserted, add the reports for the other system. In the Disk Magic TreeView panel, right-click the disk subsystem icon, and select **Add Reports** as shown in Figure 4-49.

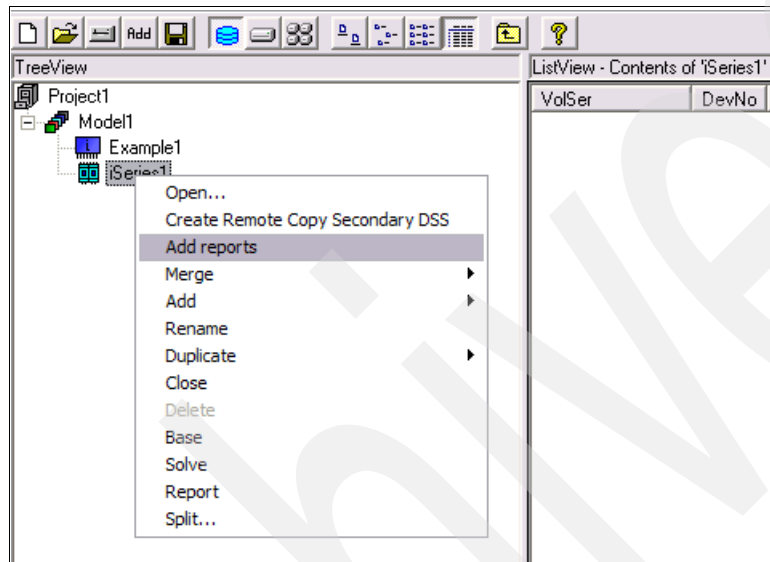


Figure 4-49 Adding reports from the other system

3. In the Open window (Figure 4-50), select the reports of another workload to insert, and click **Open**.

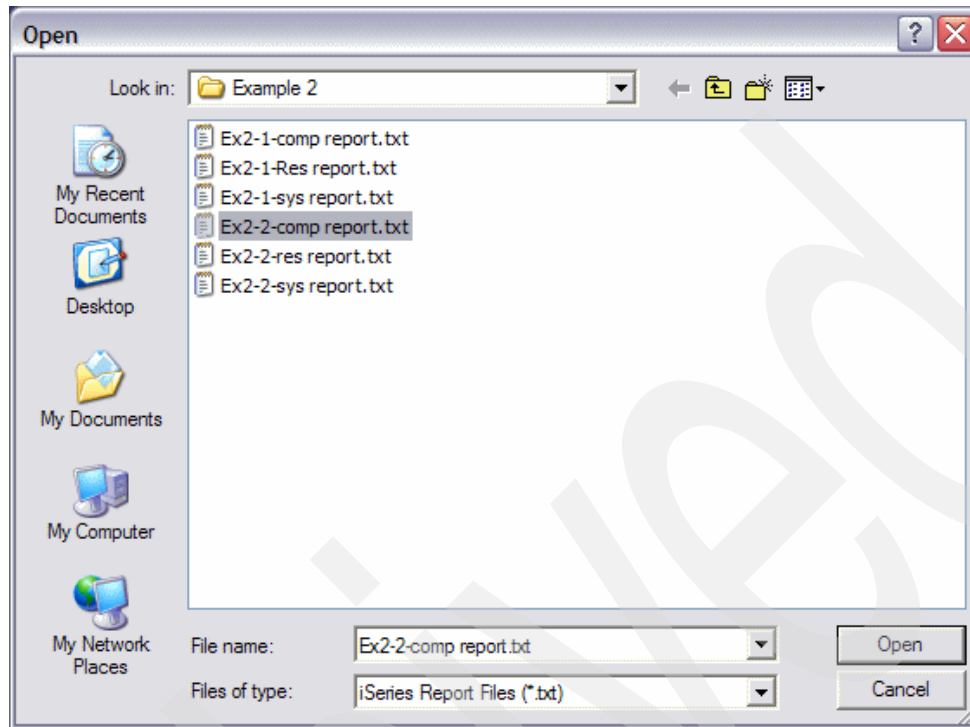


Figure 4-50 Inserting reports from another i5/OS system

4. After the reports of the second system are inserted, observe that the models for both workloads are present in TreeView panel as shown in Figure 4-51. Double-click the **iSeries** disk subsystem.

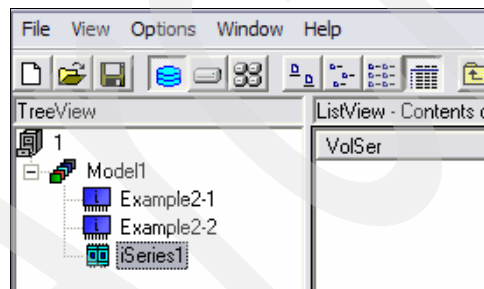


Figure 4-51 Models for both systems

5. In the Disk Subsystem - iSeries1 window (Figure 4-52), select the iSeries Disk tab. Notice that the two subtabs on the iSeries Disk tab and that each shows the current capacity for the internal disks of one workload.
  - a. Click the Example2-1 tab, and observe the current capacity for the first i5/OS workload.

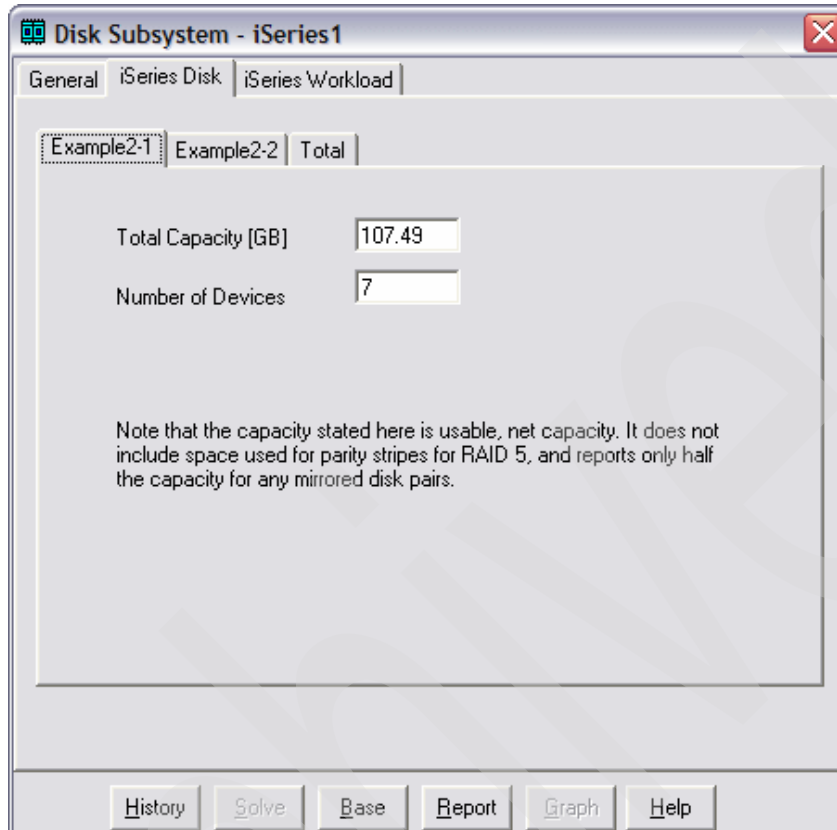


Figure 4-52 Current capacity of the first i5/OS system

- b. Select the Example2-2 tab to see the workload of the second i5/OS system (Figure 4-53).

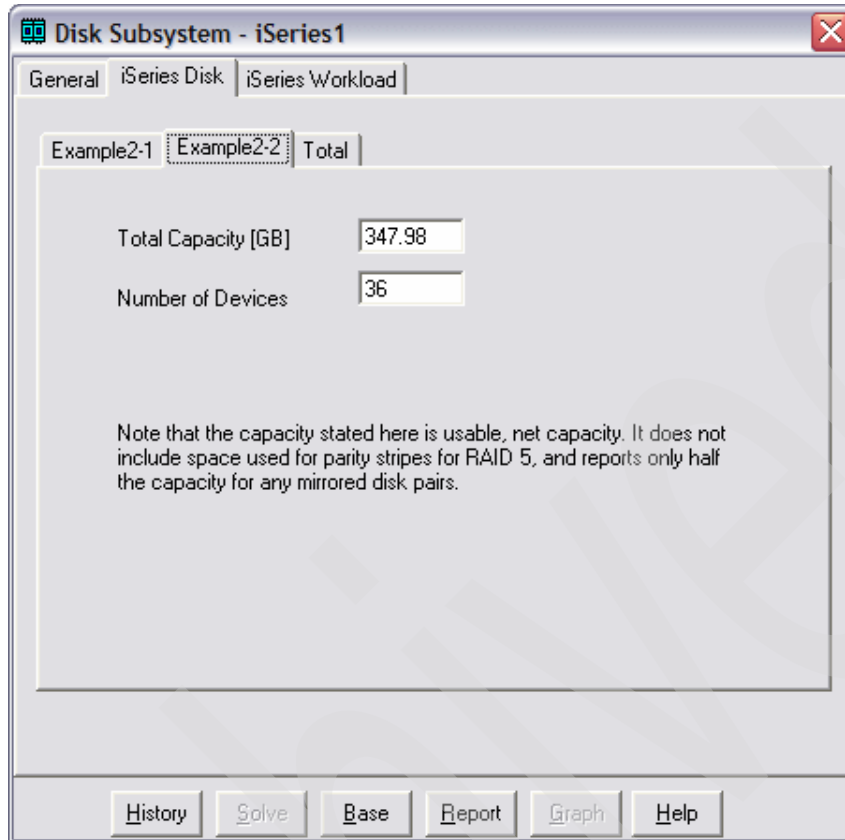


Figure 4-53 Workload characteristics of each system

6. Select the iSeries Workload tab, and click **Cache Statistics**. The Cache Statistics for Host window opens and shows the current cache usage. Figure 4-54 shows the cache usage of the second i5/OS system. Click **OK**.

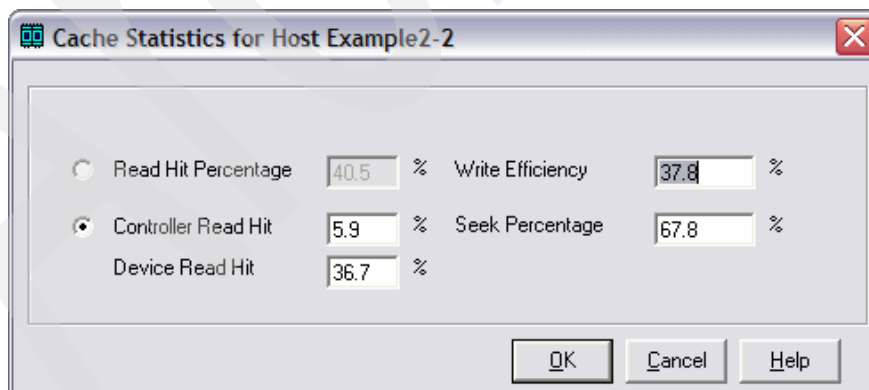


Figure 4-54 Current cache usage of workloads

7. In the Disk Subsystem - iSeries1 window, click **Base** to save the current configuration of both i5/OS systems as a base for further modeling.

8. After the base is saved, model the external disk subsystem for both workloads:
- In the Disk subsystem - iSeries1 window, select the General tab. For Hardware type, select the desired disk system. In our example, we select **DS8100** and **Multipath with iSeries**, as shown in Figure 4-55.

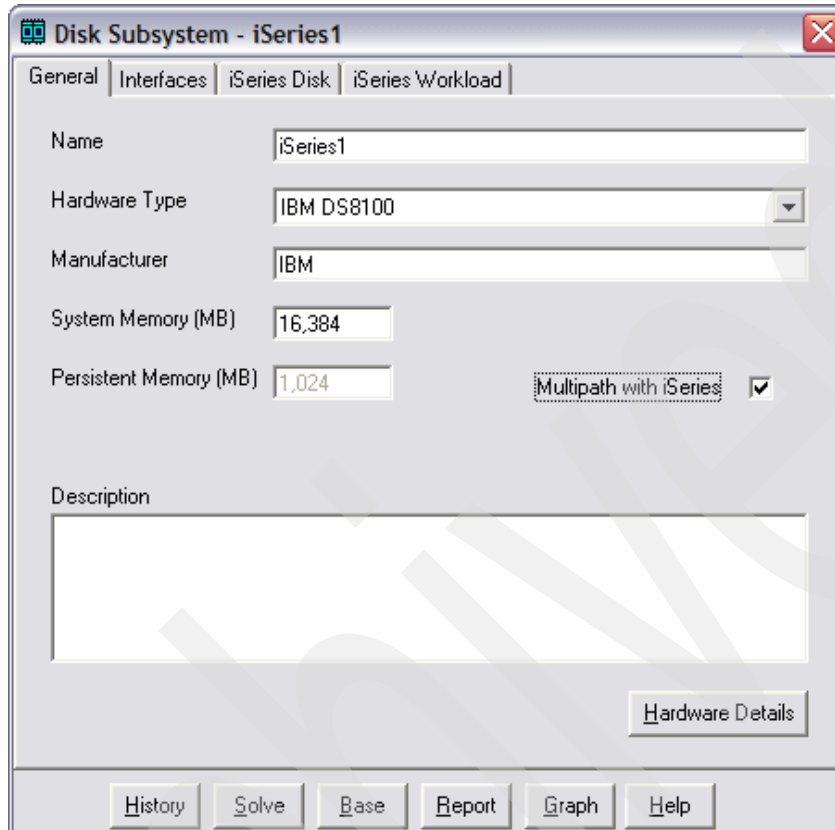


Figure 4-55 Selecting the external disk subsystem

In our example, we plan the following configurations for each i5/OS workload:

- Workload Example2-1: 12 LUNs of size 17 GB and 2 System i5 FC adapters in multipath
- Workload Example2-2: 22 LUNs of size 17 GB and 2 System i5 FC adapters in multipath

The four System i5 FC adapters is connected to two DS host ports using switches.

- b. To model the number of System i5 adapters, select the Interfaces tab, and then select the From Servers tab. You see the current workloads with the four default interfaces (see Figure 4-56). For each workload, highlight the workload, and click **Edit**.

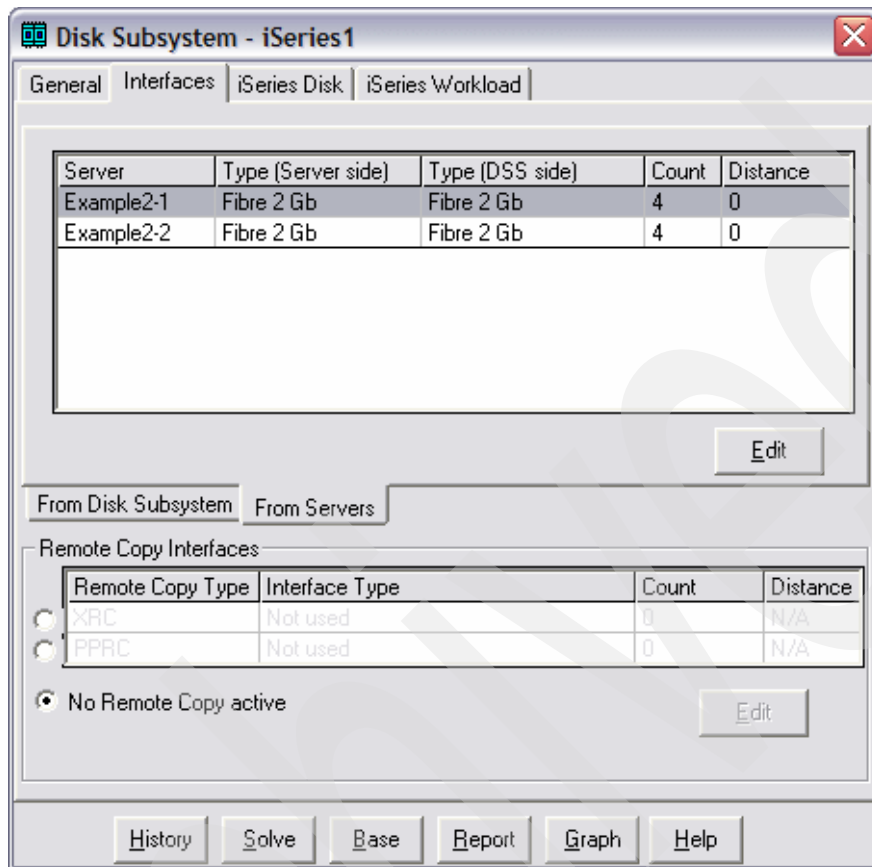


Figure 4-56 Current interfaces

- c. In the Edit Interfaces window (Figure 4-57), change the number of interfaces as planned, and click **OK**.




Figure 4-57 Insert planned no of System i5 adapters

- d. To model the number of DS host ports, select the Interfaces tab, and then select the From Disk Subsystem tab. You see the interfaces from DS8100. Click **Edit**, and insert the planned number of DS host ports. Click **OK**.

- e. In the Disk Subsystem - iSeries1 window, select the iSeries Disk tab. Notice that Disk Magic creates an extent pool for each i5/OS system automatically. Each extent pool contains the same capacity that is reported for internal disks. See Figure 4-58.

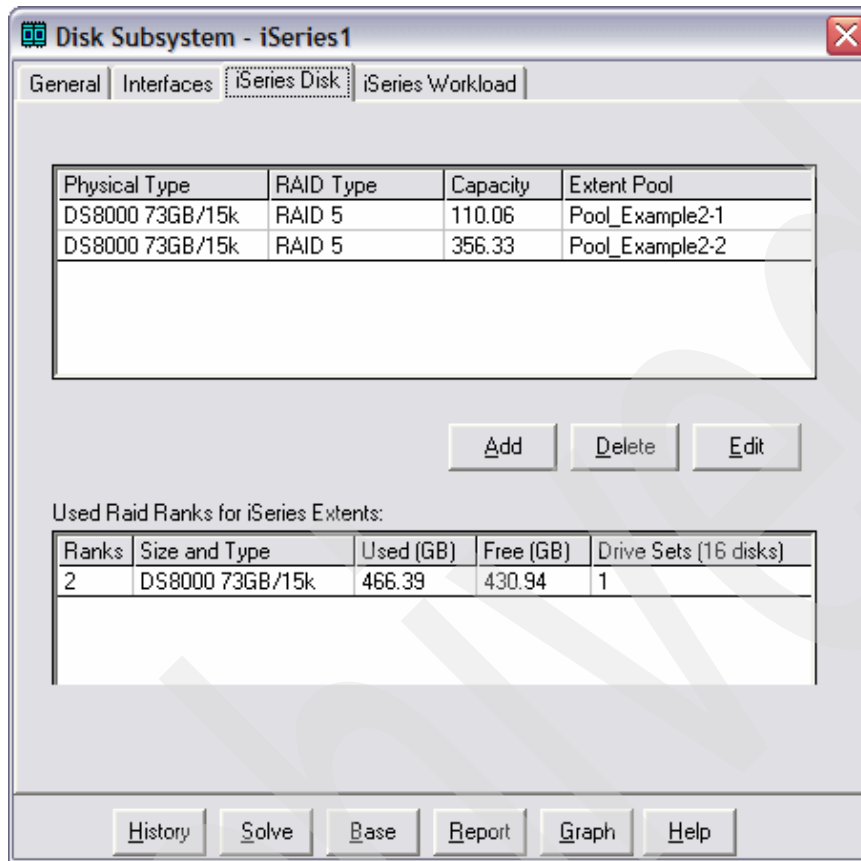


Figure 4-58 Current capacity in the extent pools

In our example, we plan to share two ranks between the two i5/OS systems, so we do not want a separate extent pool for each i5/OS system. Instead, we want one extent pool for both systems.

- f. On the iSeries Disk tab, click the **Add** button. In the Add a Disk Type window (Figure 4-59), in the Capacity (GB) field, enter the needed capacity of the new extent pool. For Extent Pool, select **Add New**.

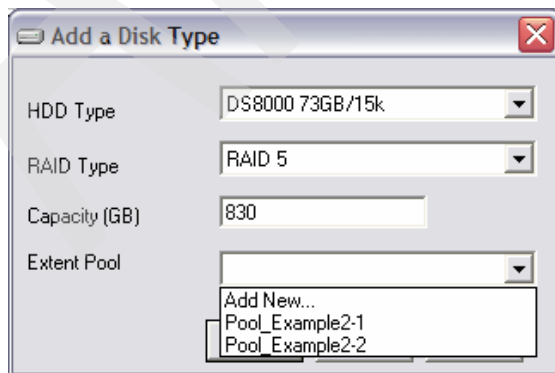


Figure 4-59 Creating an extent pool to share between the two workloads

- g. In the Specify Extent Pool name window (Figure 4-60), enter the name of the new extent pool, and click **OK**.




Figure 4-60 Name of the new extent pool

- h. The iSeries Disk tab shows the new extent pool along with the two previous extent pools (Figure 4-61). Select each extent pool, and click **Delete**.

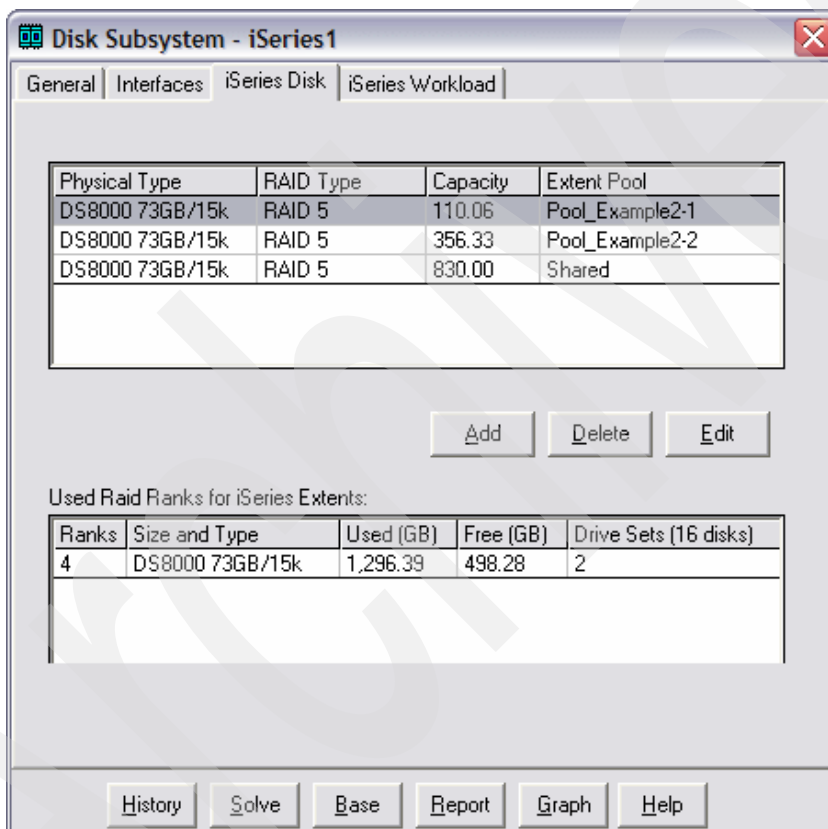


Figure 4-61 Deleting the previous extent pools



After you delete both of the previous extent pools, only the new extent pool named *Shared* is shown on the iSeries Disk tab, as shown in Figure 4-62.

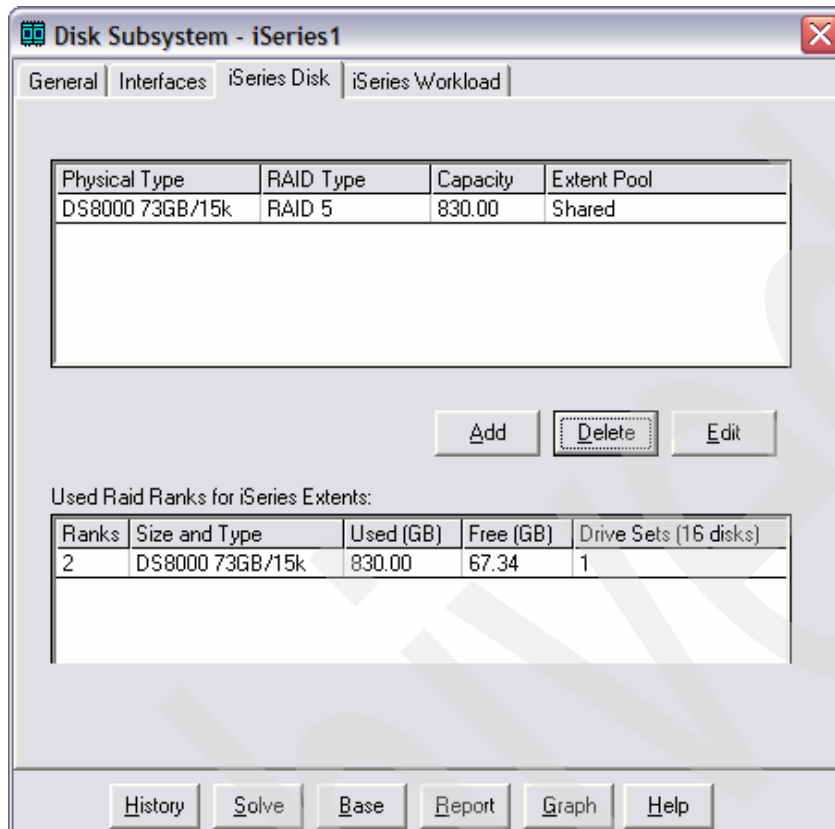


Figure 4-62 Only the Shared extent pool is available

- i. In the Disk Subsystem - iSeries1 window, select the iSeries Workload tab (Figure 4-63). Then, select the tab with the name of the first workload, which in this case is Example2-1. Complete the following information:
  - For Extent Pool, select the pool name **Shared**.
  - In the LUN count field, enter the planned number of LUNs for the first i5/OS system.
  - In the Used Capacity (GB) field, enter the usable capacity for i5/OS.

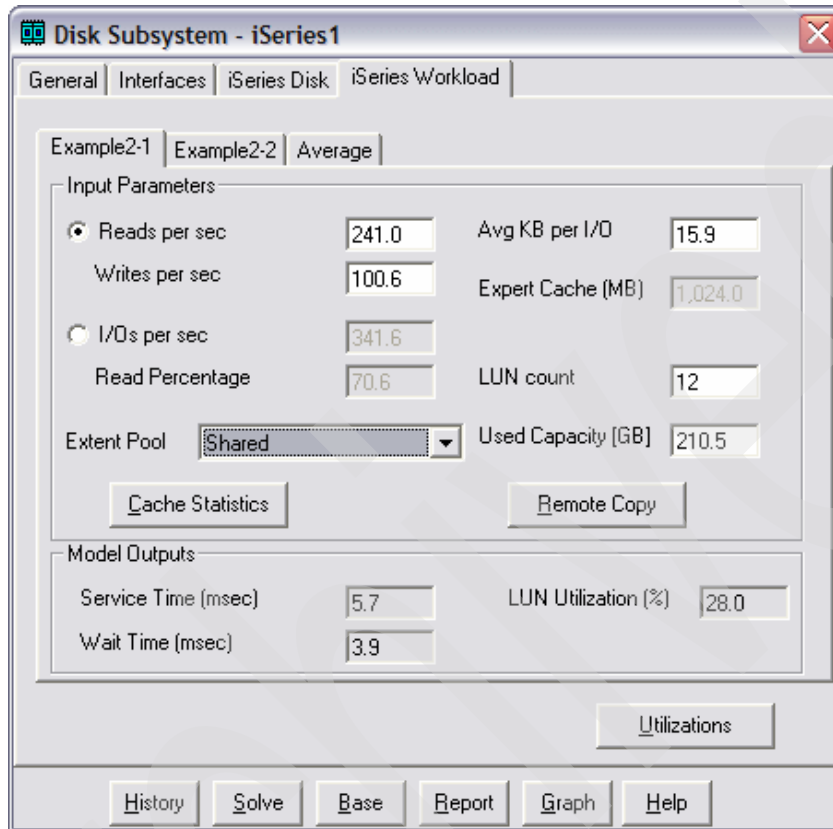


Figure 4-63 Inserting the values for first workload

- j. Select the tab with the name of the second i5/OS workload, which in this case is Example2-2 (Figure 4-64). Then, complete the following information:
- For Extent Pool, select the extent pool named **Shared**.
  - For LUN count, enter the planned number of LUNs.
  - For Used Capacity, enter the amount of usable capacity.

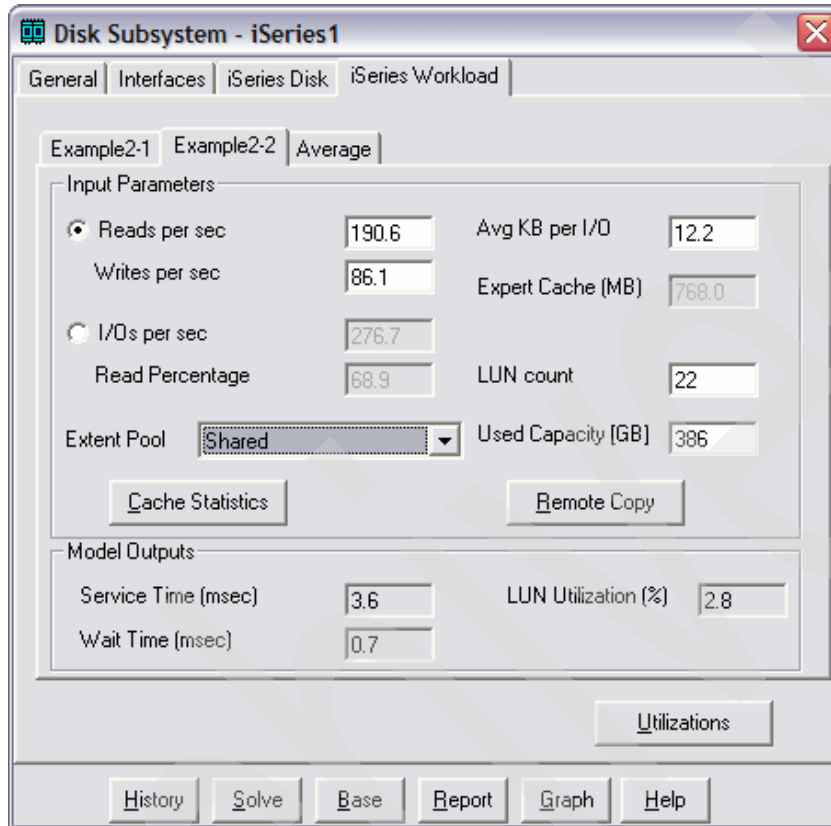


Figure 4-64 Inserting values for the second workload

- k. In the Disk Subsystem - iSeries1 window, click **Solve** to solve the modeled DS configuration.

- I. Then, select the iSeries Workload tab. Click the tab with the name of the first workload. Notice the modeled disk service time and wait time, as shown in Figure 4-65.

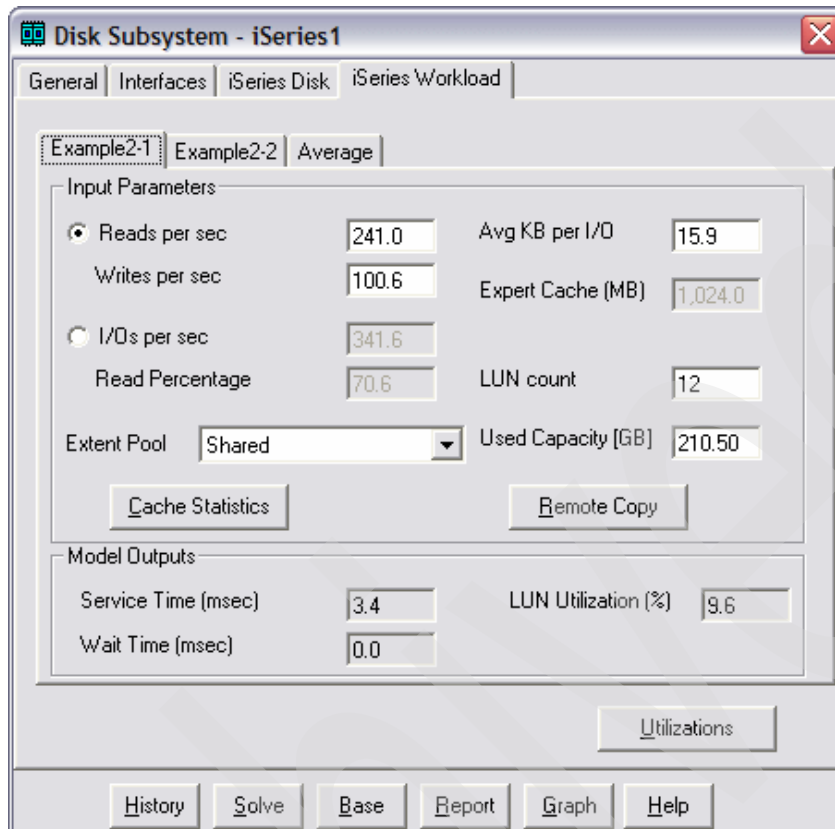


Figure 4-65 Modeled service time and wait time for the first workload

- m. Click the tab with the name of second workload, which in this case is Example2-2. Notice the modeled disk service time and wait time, as shown in Figure 4-66.
- n. Select the Average tab, and then click **Utilizations**.




Figure 4-66 Modeled service time and wait time for the second workload

- o. In the Utilizations IBM 8100 window, observe the modeled utilizations of DDMs (HDDs), FC adapters, and average utilization of LUNs for both workloads. See Figure 4-67.

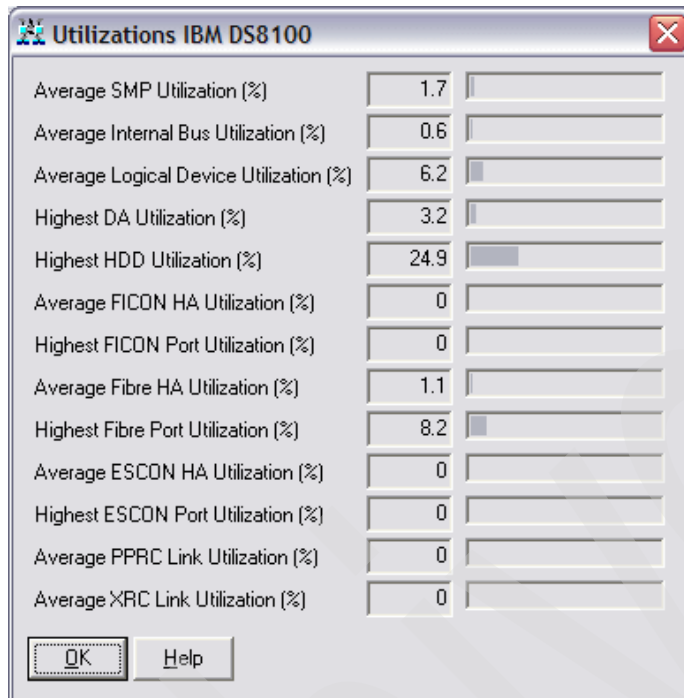


Figure 4-67 Modeled utilizations

### 4.5.3 Modeling System i5 and DS8100 for a batch job currently running Model 8xx and ESS 800

In this example, we describe the sizing of DS8100 for a batch job that currently runs on iSeries Model 825 with ESS 800. The needed performance reports are available, except for System report - Storage pool utilization, which is optional for modeling with Disk Magic.

To size a DS system for a workload that currently runs on ESS 800:

1. Insert an iSeries performance reports from current the workload to Disk Magic. For instructions about how to insert performance reports into Disk Magic, see 4.5.1, "Sizing the System i5 with DS8000 for a customer with iSeries model 8xx and internal disks" on page 113.

2. After you insert the performance reports, Disk Magic creates one disk subsystem for the I/O rate and capacity part of the iSeries workload that runs on ESS 800, and one disk subsystem for the part of the workload that runs on internal disks, as shown in Figure 4-68.

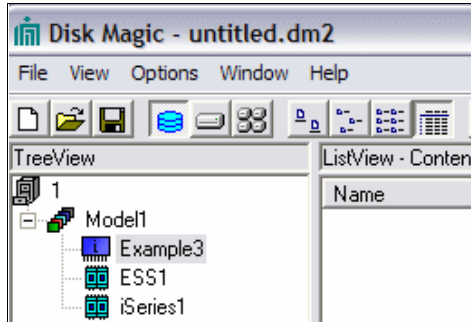


Figure 4-68 Disk Magic model for iSeries with external disk

3. Double-click **iSeries1**.
4. In the Disk Subsystem - iSeries1 window (Figure 4-69), select the iSeries Disk tab.



Figure 4-69 Subsystem for internal disks

As shown in Figure 4-70, notice the capacity that is used by the part of the workload on internal disks. In our example, the customer has only four 6 GB internal drives. Disk Magic does not take one of the drives into account because it considers it to be a mirrored load source. Therefore, three of them are in this model.

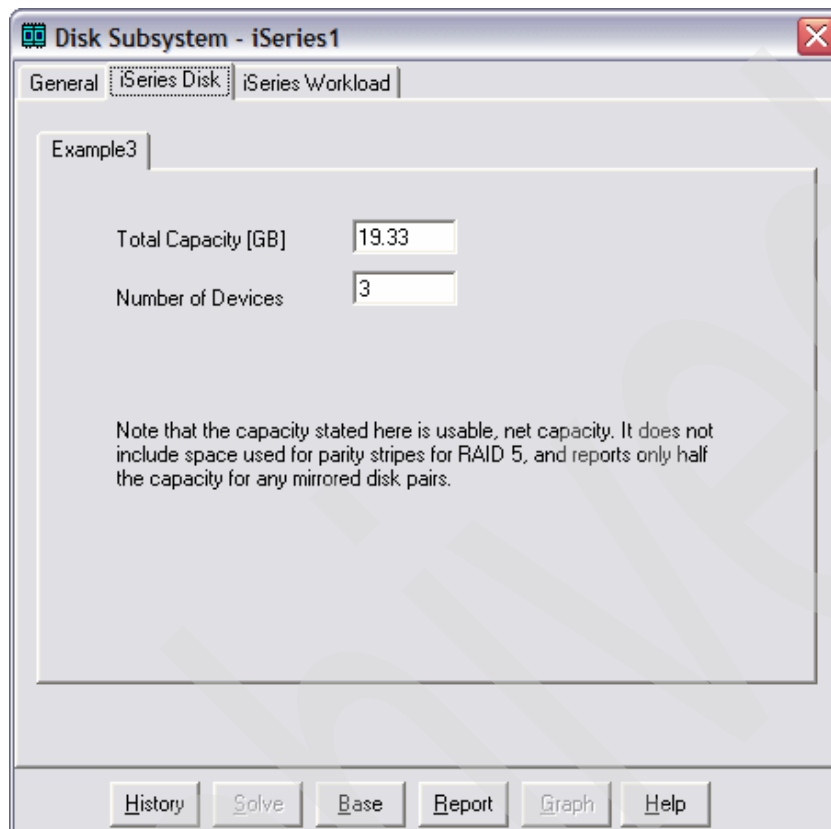


Figure 4-70 Internal capacity



5. Select the iSeries Workload tab. Notice the I/O rate on the internal disks as shown in Figure 4-71. In our example, a low I/O rate is used for the internal disks.

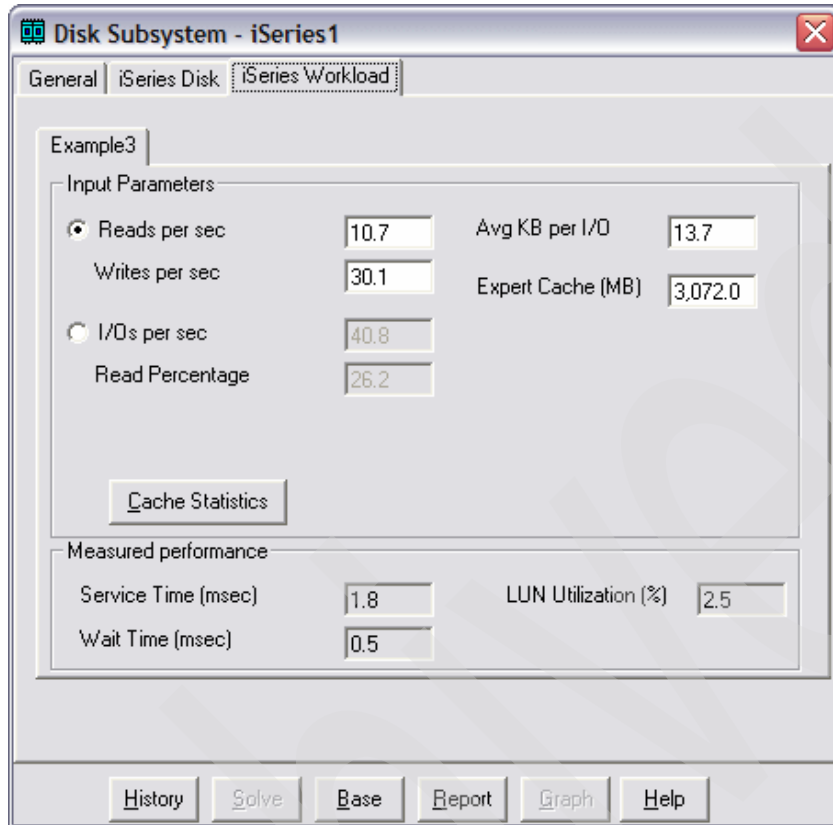


Figure 4-71 Workload on internal disks

6. In the Disk Subsystems - iSeries1 window, click **Base** to save the base for internal disks.
7. In the TreeView panel, double-click the **ESS1** icon.

The Disk Subsystem - ESS1 window opens (Figure 4-72) and shows the model for the part of capacity and workload on ESS 800.

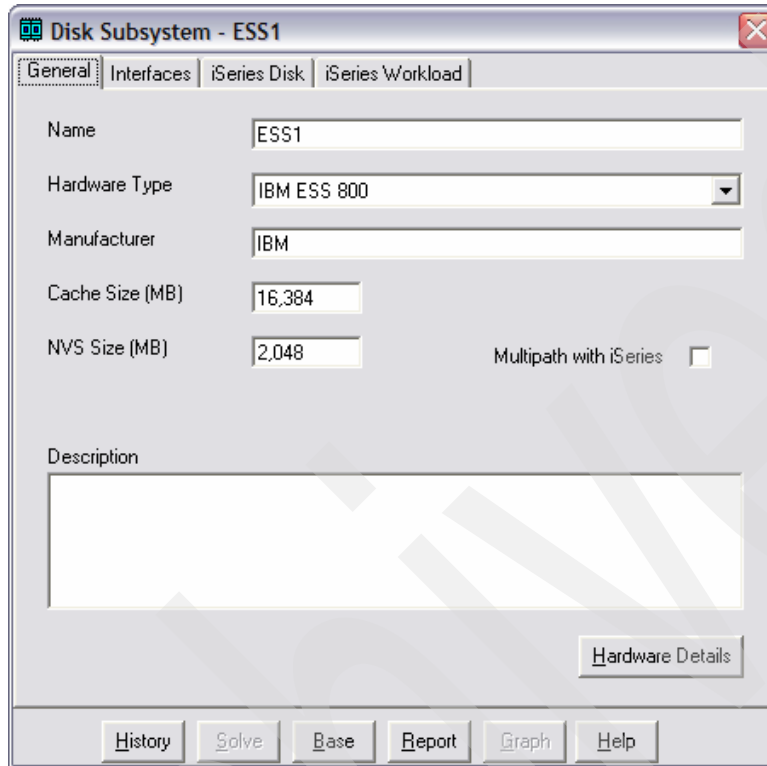


Figure 4-72 Workload on the ESS

8. Adjust the model for the currently used ESS 800 so that it reflects the correct number of ranks, size of DDMs, and FC adapters as described in the steps that follow. In our example, the existing ESS 800 contains 8 GB cache, 12 ranks of 73 GB 15 KB rpm DDMs and four FC adapters with feature number 2766, so we enter these values for disk subsystem ESS1. To adjust the model:
  - a. Select the General tab, and click **Hardware Details**.
  - b. The ESS Configuration Details window (Figure 4-73 on page 155) opens. Replace the default values with the correct values for the existing ESS. In our example, we use four FC adapters and 8 GB of cache, so we do not change the default values. Click **OK**.

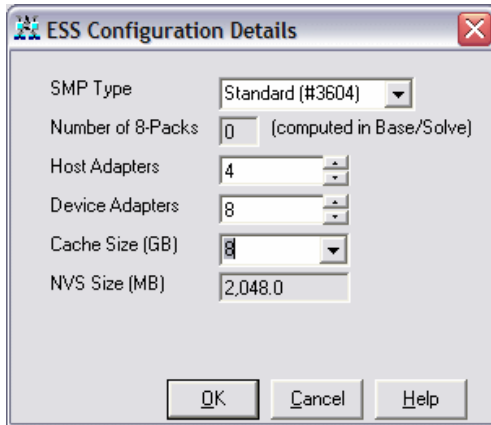


Figure 4-73 Hardware details for the existing ESS 800

- c. Select the Interfaces tab, and click the **From Disk Subsystem** subtab. Click **Edit**.
- d. The Edit Interfaces for Disk Subsystem window (Figure 4-74) opens. Enter the correct values for the current ESS 800. In our example, the customer uses four host ports from ESS, so we do not change the default value of 4. However, we change the type of adapters for Server side to **Fibre 1 Gb** to reflect the existing iSeries adapter 2766. Click **OK**.




Figure 4-74 Insert Interfaces for Disk Subsystem

- e. On the Interfaces tab, click the **From Servers** tab and click **Edit**.
- f. In the Edit Interfaces window (Figure 4-75), enter the current number and type of iSeries FC adapters. In our example, we use four iSeries 2766 adapters, so we leave the default value of 4. However, for Server side, we change the type of adapters to **Fibre 1 Gb** to reflect the current adapters 2766.

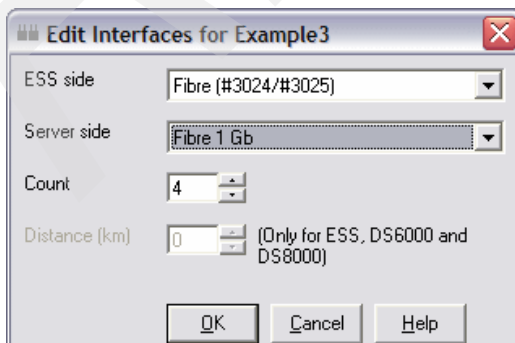


Figure 4-75 Current iSeries FC adapters

g. In the Disk Subsystem - ESS1 window, select the iSeries Disk tab.

On the iSeries Disk tab, observe that current capacity and the number of LUNs are inserted by Disk Magic and that 36 GB 15 KB rpm ranks are used as default for Physical Device Type. If necessary, select another value for Physical Device Type to reflect the current ESS configuration.

In our example, we select **ESS 73 GB/15000** because the customer currently uses 73 GB 15 KB rpm DDMs on the ESS. Observe that the number of used ranks change when we change the type of DDMs.

In some cases, it might be necessary to configure more ranks for performance than are required for capacity. Disk Magic can validate the proposed configuration. We recommend that you use Capacity Magic for capacity planning because Disk Magic does not take sparing into account.

In our example, Disk Magic models only two ranks for the customer's workload, as shown in Figure 4-76. With the DS systems, we can model less capacity used for a System i5 model than is the total capacity of used ranks.

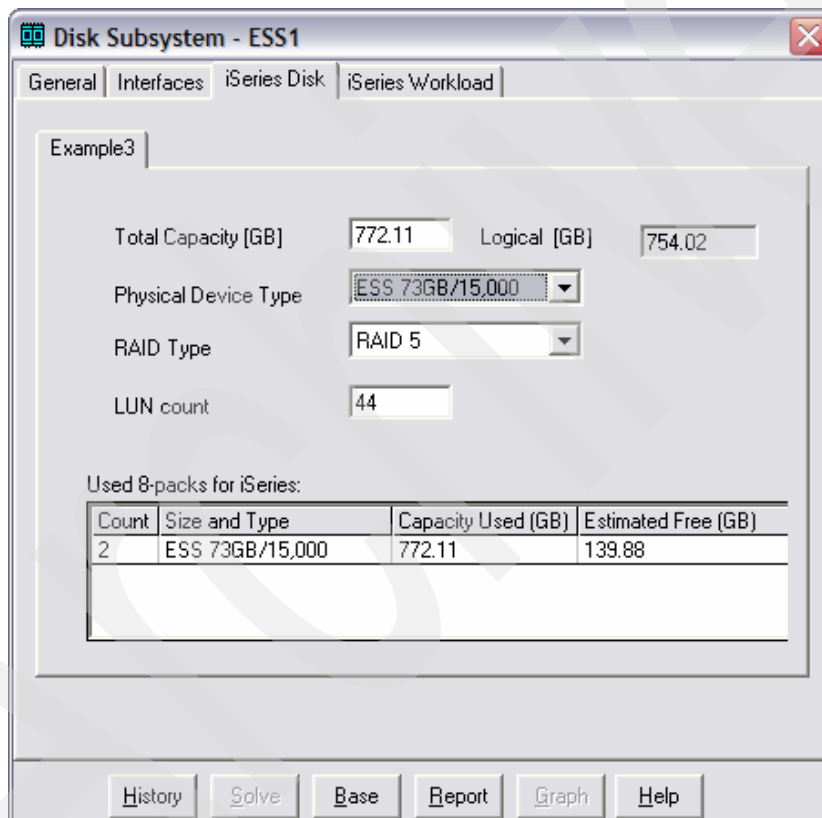


Figure 4-76 Current capacity on the ESS

- h. Select the iSeries Workload tab. Notice that the current I/O rate and block sizes are inserted by Disk Magic as shown in Figure 4-77.

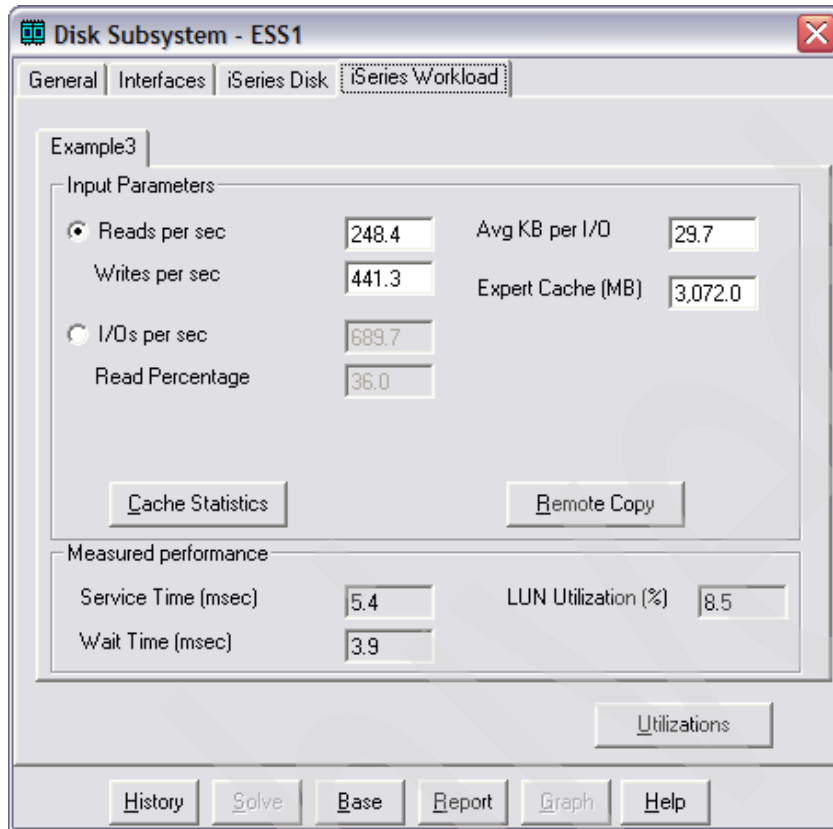


Figure 4-77 Current workload

- i. On the iSeries Workload tab, click **Cache Statistics**. In the Cache Statistics for Host window (Figure 4-78), notice the currently used cache percentages. Click **OK**.

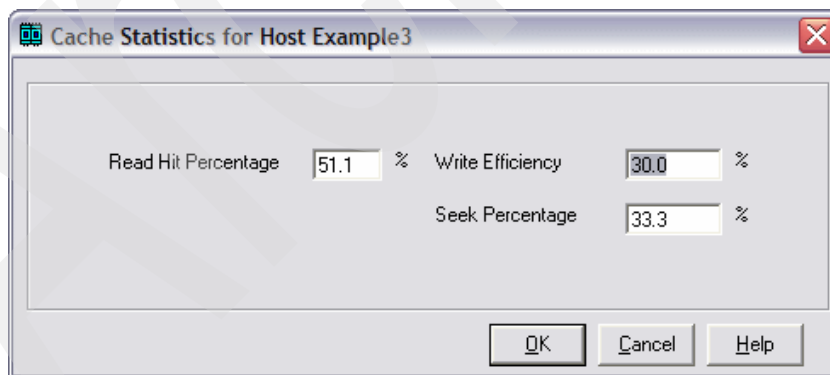


Figure 4-78 Current cache usage

- j. In the Disk Subsystem - ESS1 window, click **Base** to save the current model of ESS.

9. Next, insert the planned values for the DS system in the Disk Subsystem - ESS1 window.
  - a. Select the General tab. For Hardware Type, select the planned model of the DS system. In our example, we select **DS8100**, which is planned for this customer. See Figure 4-79.

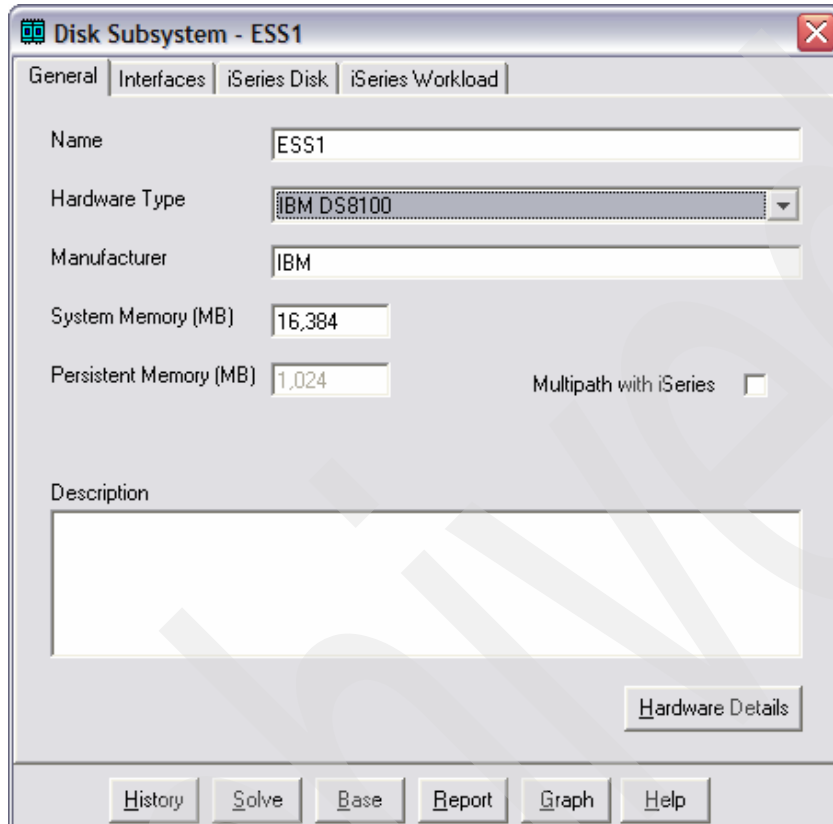


Figure 4-79 Planned hardware type

- b. Click **Hardware Details**. In the Hardware Details IBM DS8100 window (Figure 4-80), enter the values for the planned DS system. In our example, the customer uses four DS FC host ports, so we enter 4 for Fibre Host Adapters. Click **OK**.

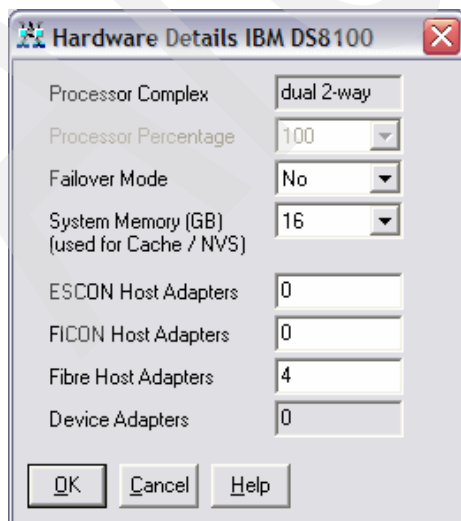


Figure 4-80 Hardware details of planned DS

- c. Select the Interfaces tab. Select the From Disk Subsystem tab and click **Edit**.
- d. The Edit Interfaces for Disk Subsystem window (Figure 4-81) opens. Enter the planned number and type of DS host ports. In our example, the customer plans on four DS ports and four adapters with feature number 2787 in the System i5 model. Therefore, we leave the default value for Count. However, for Server side, we change the type to **Fibre 2 Gb**. Click **OK**.

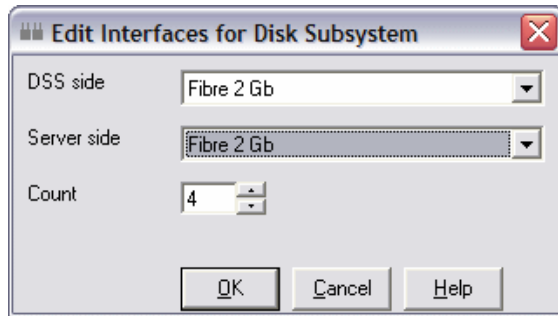


Figure 4-81 Planned DS ports

- e. On the Interfaces tab, select the From Servers tab and click **Edit**. The Edit Interfaces window (Figure 4-82) opens. Enter the planned number and type of System i5 FC adapters. In our example, the customer plans for four FC adapters 2787, so we leave the default value of 4 for Count. However, for Server side, we select **Fibre 2 Gb**. Click **OK**.

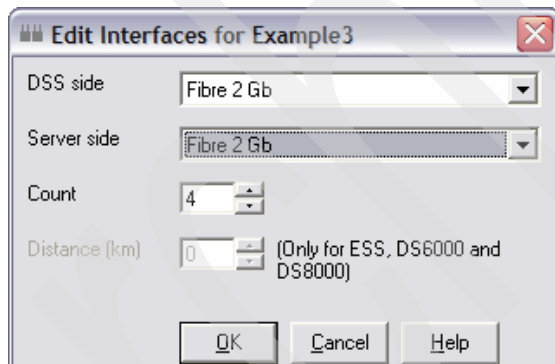


Figure 4-82 Planned System i5 FC adapters

- f. Select the iSeries Disk tab. Notice that an extent pool is already created with the same capacity as is used on ESS. See Figure 4-83. Click **Edit**.

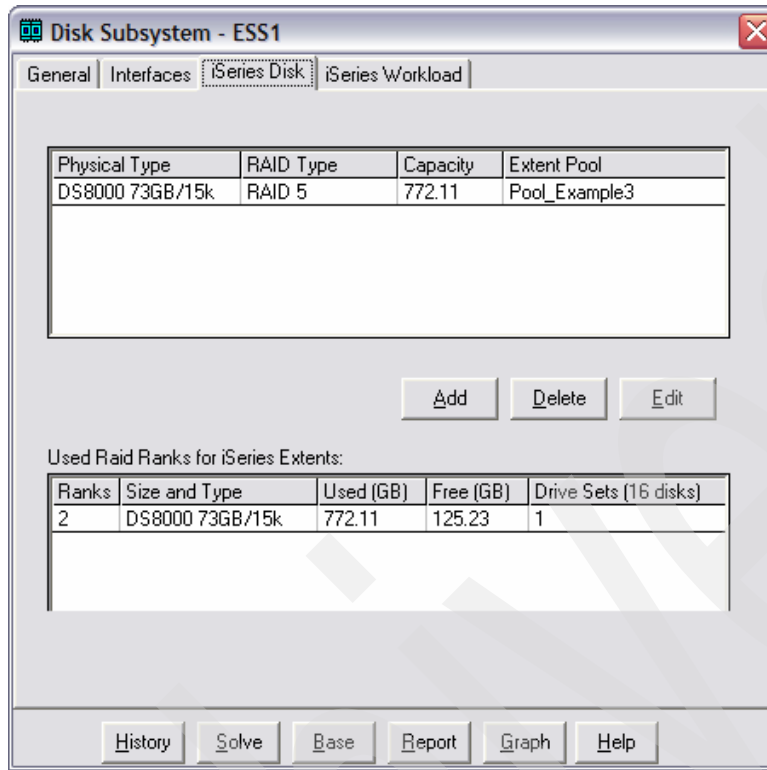


Figure 4-83 Planned capacity -1

- g. In the Edit a Disk Type panel (Figure 4-84), enter the capacity that corresponds to the desired number of ranks for Capacity. Observe that 73 GB 15 KB rpm ranks are already inserted as the default for HDD Type.

In our example, the customer plans nine ranks. The available capacity of one RAID-5 73 GB rank with spare (6+P+S rank) is 414.46 GB. We enter a capacity of 3730 (9 x 414.46 GB = 3730 GB), and click **OK**.




Figure 4-84 Planned number of ranks



- h. Select the iSeries Workload tab. Enter the planned number of LUNs and the amount of capacity that is used by the System i5 model. Notice that the extent pool for the i5/OS workload is already specified for Extent Pool.

In our example, the customer plans for 113 of 17 GB LUNs, so we enter 113 for LUN count. We also enter 1982 (using the equation  $113 \times 17.54 = 1982$  GB) for Used Capacity. See Figure 4-85.




Figure 4-85 Planned capacity-2

- i. On the iSeries Workload tab, click **Cache Statistics**. In the Cache Statistics for Host window (Figure 4-86), notice that the box Automatic cache Modeling is selected. This indicates that Disk Magic will model cache percentages automatically for DS8100 based on the reported values from performance reports for the currently used ESS 800. Note that write cache efficiency reported in performance reports is not correct for ESS 800, so Disk Magic uses a default value 30%.

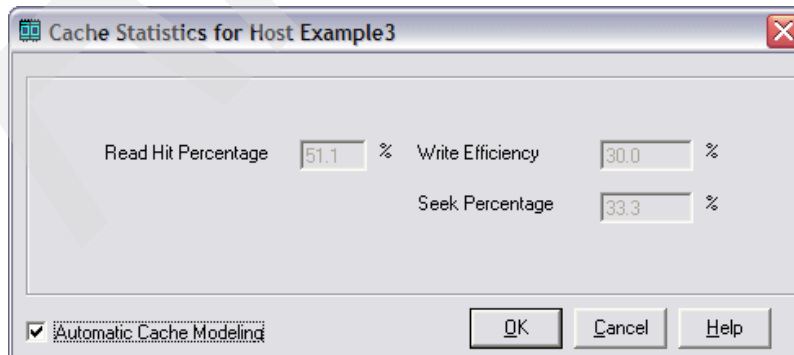


Figure 4-86 Automatic cache modeling

- j. In the Disk Subsystem - ESS1 window, click **Solve** to ensure that the planned DS configuration is modeled for the current workload. On the iSeries Workload tab (Figure 4-87), notice the modeled disk service time and wait time. In our example, the modeled service time is 3.8 ms, and the modeled wait time is 0.4 milliseconds.

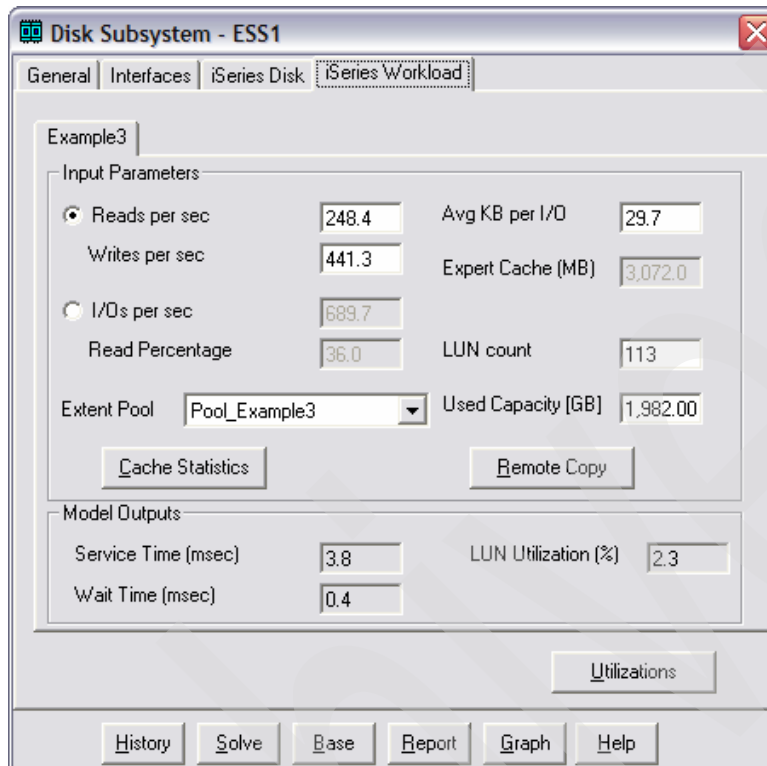


Figure 4-87 Modeled service time and wait time

- k. On the iSeries Workload tab, click **Utilizations**. Notice the modeled utilization of HDDs, DS FC ports, LUNs, and so on, as shown in Figure 4-88. In our example, the modeled utilizations are rather low so the customer can grow the workload to a certain extent without needing additional hardware in the DS system.

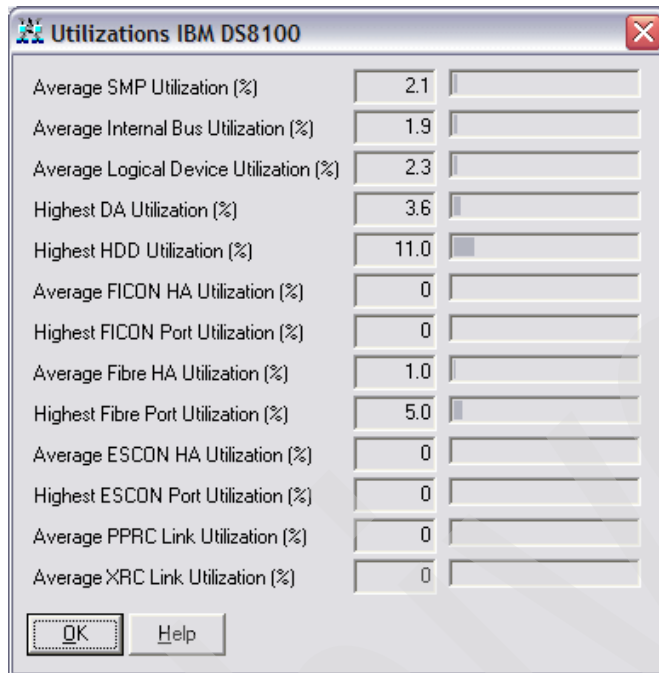


Figure 4-88 Modeled utilizations

In our example, the customer migration from iSeries model 825 to a System i5 model was performed at the same time as the installation of DS8100. Therefore, the number of I/Os per second and the cache values differ from the ones that were used by Disk Magic. The actual disk response times were lower than the modeled ones. The actual reported disk service time is 2.2 ms, and disk wait time is 1.4 ms.

#### 4.5.4 Using IBM Systems Workload Estimator connection to Disk Magic: Modeling DS6000 and System i for an existing workload

In this example, we present usage of IBM Systems Workload Estimator (WLE) together with Disk Magic, to size a System i server and DS6000 for an existing workload that runs on iSeries model 870 and internal disks. To perform this, you must have the following i5/OS Performance Tools reports:

- ▶ System report
  - Workload
  - Resource utilization
  - Storage Pool utilization
  - Disk utilization
- ▶ Resource report
  - Disk utilization
  - IOP utilizations
- ▶ Information about currently used disk adapters

To size System i5 and DS6000 with WLE and Disk Magic:

1. Start IBM Systems WLE by accessing the following Web page:  
<http://www-912.ibm.com/wle/EstimatorServlet>
2. On the License Agreement page, read the license agreement and then click **I Accept** if you accept the terms of the agreement.
3. On the User Demographic Information page, provide your demographic information and click **Continue**.
4. A panel displays as shown in Figure 4-89. To size an existing workload, click **Workload: Add** in the blue tab at the top of the panel.

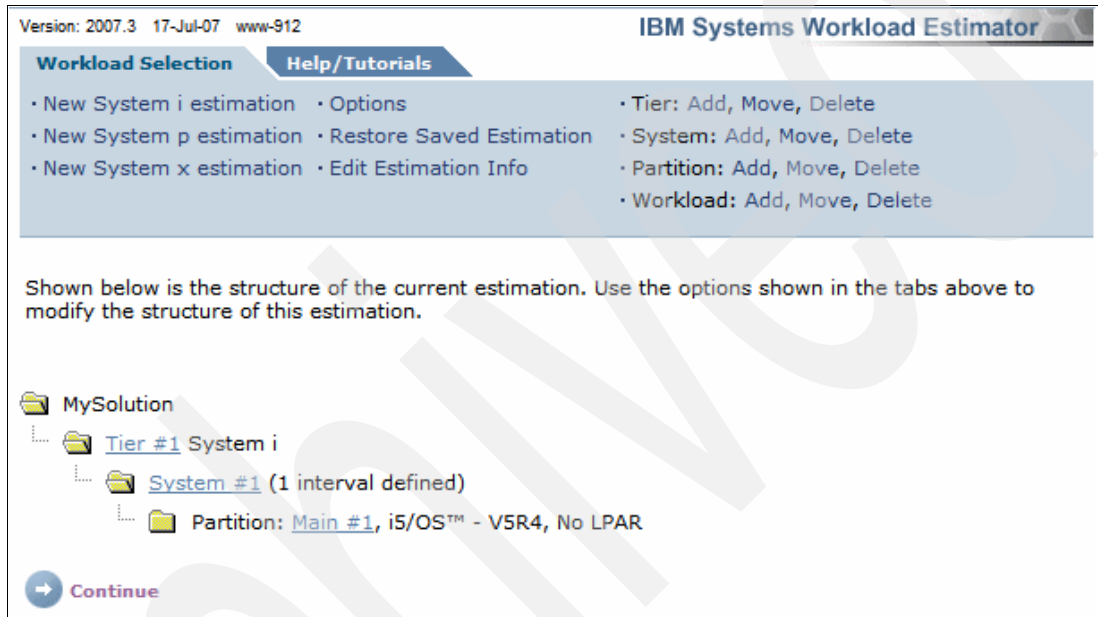


Figure 4-89 WLE initial panel

5. In the Workload Selection panel (Figure 4-90), for Add Workload, select **Existing** and click **Go**.

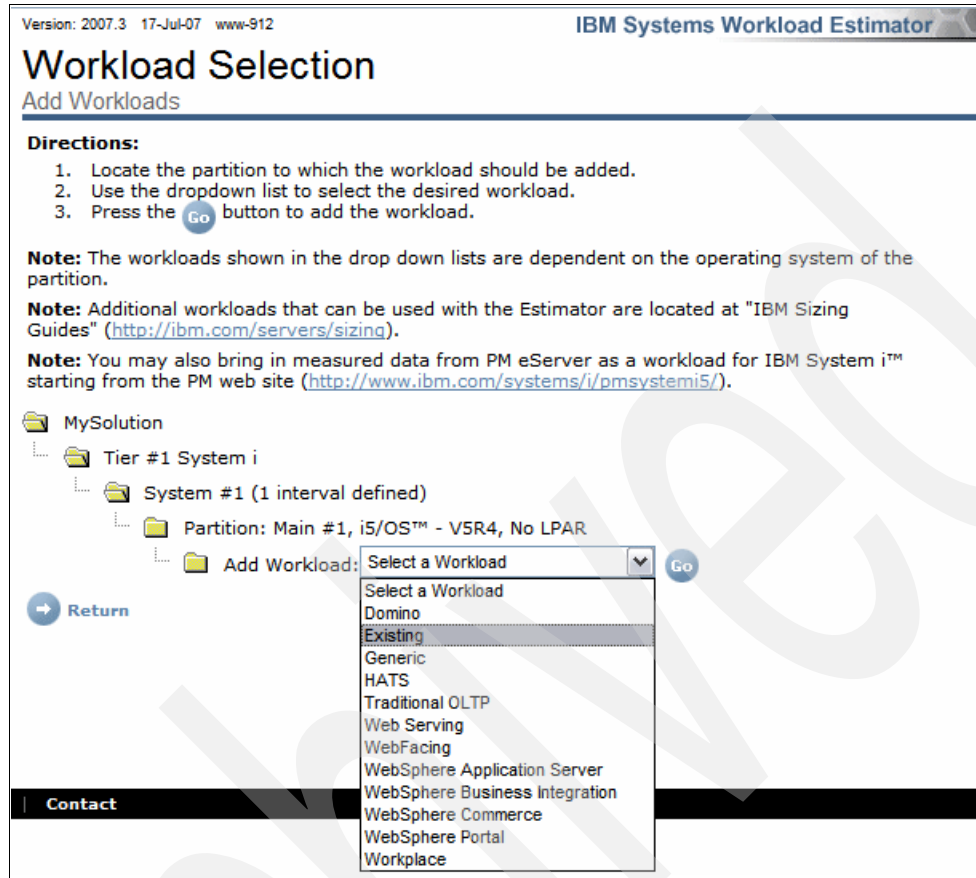


Figure 4-90 Workload selection

- In the next panel (Figure 4-91), you can add another workload. Notice that Existing workload #1 that you selected in the previous panel is shown. Do not select another workload. Click **Return**.

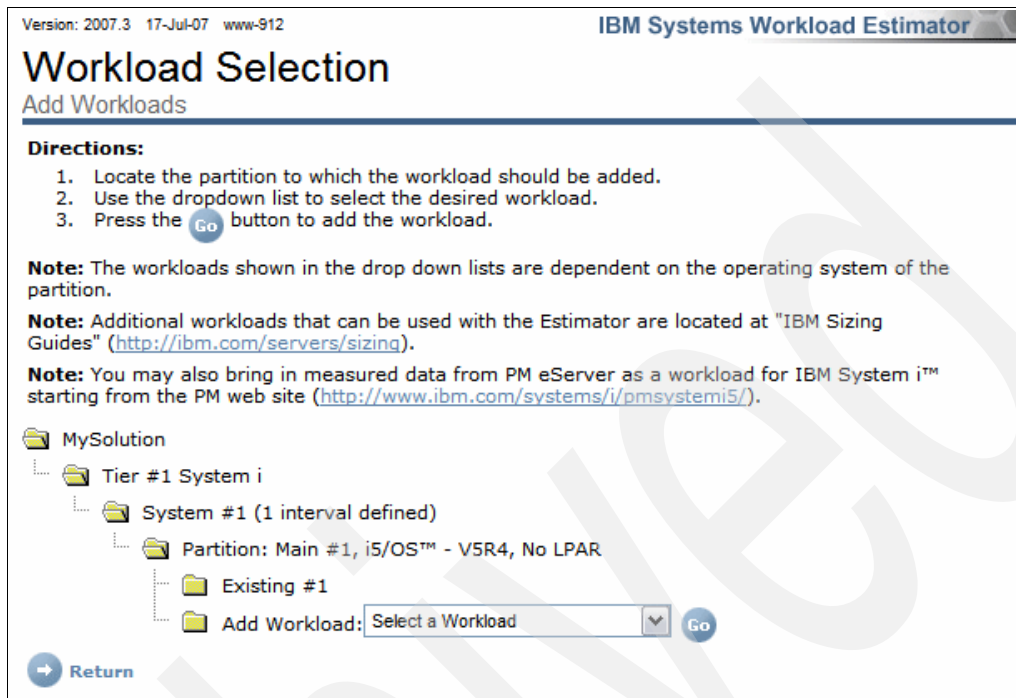


Figure 4-91 Selecting another workload

- You return to the initial panel, which contains the Existing #1 workload (see Figure 4-92). Click **Continue**.

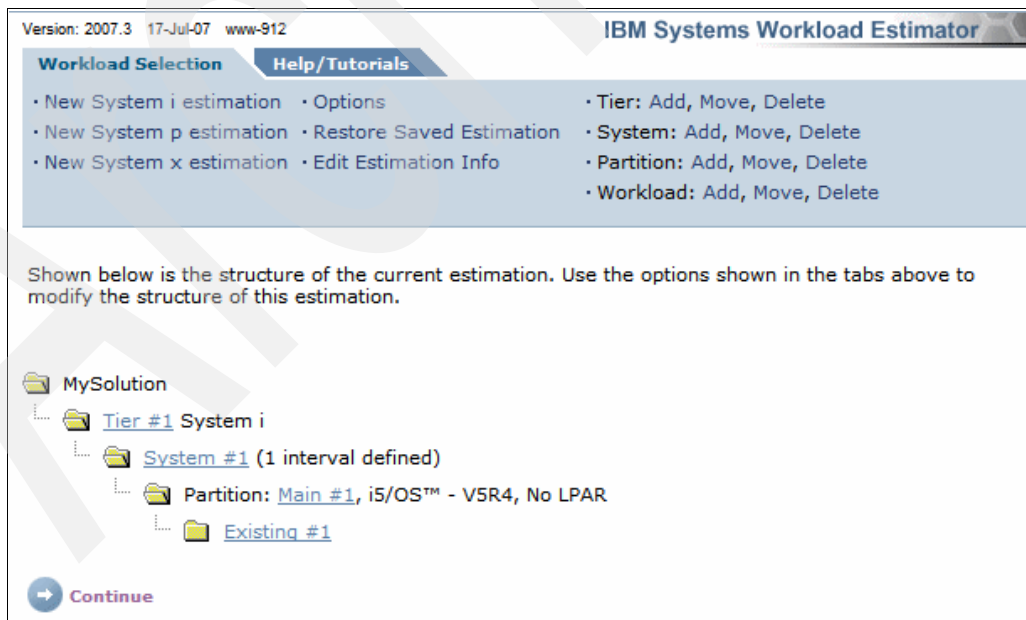


Figure 4-92 Initial panel with the Existing #1 workload

- In the Existing #1 - Existing System Workload Definition panel (Figure 4-93), enter the hardware and characteristics of the existing workload as described in the next steps.

## Existing #1

Existing System Workload Definition

Partition Name: Main #1  
i5/OS™ - V5R4  
No LPAR

**Please note:** The information requested is for an **Existing** workload

	Processor Model	Interactive Feature	Feature Code	Processor CPW	Interactive CPW
1. <a href="#">Processor Model</a>	150-0591	N/A	0591	27	13

2. [Total CPU Utilization](#)
3. [Interactive Utilization](#)
4. [Processing Cores Activated](#) (For Processor on Demand models)
5. [Represent a Logical Partition](#)
6. [Accelerator Feature used](#)
7. [Memory \(MB\)](#)
8. [Disk Configuration:](#)

Disk Group Name:  1: Group A \*

Disk Attachment Type:

Disk Drive Speed:

Storage Protection:

[Arms](#)

[Disk Busy Percentage](#)

[Storage \(GB\)](#)

[Read Ops Per Second](#)

[Bytes Per Read Op](#)

[Write Ops Per Second](#)

[Bytes Per Write Op](#)

Either ops per second and disk op size or arms and disk busy must be specified.
9. [Target Storage Protection](#)
10. Do you want to add [journal with commitment control](#)?
 

Yes  
 No
11. [Additional Characteristics:](#)
12. Will this workload be changed to use [WebFacing or HATS Support](#)?
 

WebFacing  
 HATS  
 No

Figure 4-93 Inserting the characteristics of the existing workload

- a. For Processor Model, select the model and processor features of the iSeries system on which the existing workload runs. First, obtain this information from the System report (see Figure 4-94).

System Report				19-07-05 12:01:07
Workload				Page 0001
Panter 14 7 2005 14:00 t/m 14:45				
Member . . . : Q195000004	Model/Serial . . : 870/xxxxxx	Main storage . . : 4096,0 MB	Started . . . . : 14-07-05 00:00:06	
Library . . . : QMPGPANT	System name . . : Example 4	Version/Release : 5/ 2,0	Stopped . . . . : 15-07-05 00:00:00	
Partition ID : 002	Feature Code . . : 7433-2489-7433			
QPFRAJ . . . : 2	QDYNPTYSCD . . : 1	QDYNPTYADJ . . . : 1		
Interactive Workload				

Figure 4-94 Model characteristics from the System report

- b. Next to Processor model, select the corresponding model and features (see Figure 4-95).

Partition Name: **Main #1**  
i5/OS™ - V5R4  
No LPAR

**Existing #1**  
Existing System Workload Definition

Please note: The information requested is for an Existing workload

	Model	Interactive Feature	Processor Feature Code	Processor CPW	Interactive CPW	
1. <a href="#">Processor Model</a>	150-0591	N/A	0591	27	13	▼
2. <a href="#">Total CPU Utilization</a>	840-2461	1546	26D6	20200	10000	▲
	840-2461	1547	26D7	20200	16500	
3. <a href="#">Interactive Utilization</a>	840-2461	1548	26D8	20200	20200	
	870-2486	7419	7419	11500	0	
4. <a href="#">Processing Cores Available</a>	870-2486	7421	7421	11500	11500	
	870-2486	7436	7436	11500	11500	
5. <a href="#">Represent a Logical Processor</a>	870-2489	7431	7431	7700	0	No ▼
6. <a href="#">Accelerator Feature</a>	870-2489	7433	7433	7700	11180	No ▼
	870-2489	7435	7435	7700	7700	
7. <a href="#">Memory (MB)</a>	890-0197	N/A	0197	29300	0	
	890-0198	N/A	0198	37400	0	
8. <a href="#">Disk Configuration:</a>	890-2487	1576	2AF0	20000	120	
	890-2487	1577	2AF1	20000	240	
	890-2487	1578	2AF2	20000	560	
	890-2487	1579	2AF3	20000	1050	
	890-2487	1581	2AF5	20000	2000	
	890-2487	1583	2AF7	20000	4550	
	890-2487	1585	2AF9	20000	10000	
	890-2487	1587	2AFB	20000	16500	
	890-2487	1588	2AFC	20000	20200	
	890-2488	1576	2AD0	29300	120	
	890-2488	1577	2AD1	29300	240	
	890-2488	1578	2AD2	29300	560	
	890-2488	1579	2AD3	29300	1050	
	890-2488	1581	2AD5	29300	2000	
	890-2488	1583	2AD7	29300	4550	
	890-2488	1585	2AD9	29300	10000	
	890-2488	1587	2ADB	29300	16500	
	890-2488	1588	2ADC	29300	20200	
	890-2488	1591	2ADF	29300	29300	

1: Group A \*

Add new Group    Remove this Group

Either ops per second and disk op size or arms and disk busy must be specified.

Figure 4-95 Selecting the model and features



- c. Obtain the total CPU utilization and Interactive CPU utilization data from the System report - Workload (see Figure 4-96).

NETSERVER	1	0	0	0	0	0,0000	0,0
Total	489	3.091.609	647.250	13.859	7.794		
Average						0,0003	858,7
Total CPU Utilization . . . . .							53,9
Total CPU Utilization (Interactive Feature) . . .							11,2
Total CPU Utilization (Database Capability) . . .							15,7

Figure 4-96 CPU utilization

- d. Obtain memory data from the System report in the Main Storage field (see Figure 4-94 on page 168).
- e. Insert these values into the Total CPU Utilization, Interactive Utilization, and Memory (MB) fields. If the workload runs in a partition, specify the number of processors for this partition and select **Yes** for Represent a Logical partition. See Figure 4-97.

	Model	Interactive Feature	Feature Code	Processor CPW	Interactive CPW
1. <u>Processor Model</u>	870-2489	7433	7433	7700	11180
2. <u>Total CPU Utilization</u>					53.9
3. <u>Interactive Utilization</u>					11.2
4. <u>Processing Cores Activated</u> (For Processor on Demand models)					2.0
5. <u>Represent a Logical Partition</u>					Yes
6. <u>Accelerator Feature used</u>					No
7. <u>Memory (MB)</u>					4096

Figure 4-97 Existing hardware

- f. In the Disk Configuration fields (see Figure 4-98), specify as many groups as there are different internal disk types on the system. In our example, we have only one disk type, so we use only one group. If necessary, you can add other groups by clicking **Add New Group**.

8. 8. Disk Configuration:

Disk Group Name: Group 1

Disk Attachment Type: IOA Feature #2780

Disk Drive Speed: 15,000 RPM

Storage Protection: RAID-5

Arms: 0

Disk Busy Percentage: 24

Storage (GB): 844.0

Read Ops Per Second: 112

Bytes Per Read Op: 10

Write Ops Per Second: 218

Bytes Per Write Op: 10

Add new Group Remove this Group

9. 9. Target Storage Protection

10. Do you want to add journal with commitment control?  Yes  No

11. 11. Additional Characteristics:

12. Will this workload be changed to use WebFacing or HATS Support?  WebFacing  HATS  No

Back Continue

Figure 4-98 Disk configuration

- g. Obtain the current IOA feature and RAID protection used from the iSeries configuration. Obtain the Drive Type and number of disk units from the System report - Disk Utilization (Figure 4-99).

Unit Unit	Name	Size Type	IOP (M)	IOP Util	Dsk CPU Name	ASP Rsc Util	ASP Name	--Percent-- ID	Op Full	Per Util	K Per Second	- Average I/O	Time Per Service	I/O Wait	-- Response
0001	DD004	4326	30.769	0,7	CMB01	0,6		1	59,0	1,8	14,98	9,7	.0012	.0002	.0014
0002	DD003	4326	26.373	0,7	CMB01	0,6		1	59,0	1,6	13,72	10,0	.0011	.0002	.0013
0003	DD011	4326	30.769	0,7	CMB01	0,6		1	59,0	1,6	11,83	11,7	.0013	.0003	.0016
0004	DD005	4326	30.769	0,7	CMB01	0,6		1	59,0	1,7	16,49	8,2	.0010	.0000	.0010
0005	DD009	4326	30.769	0,7	CMB01	0,6		1	59,0	1,5	15,17	9,5	.0009	.0002	.0011
0006	DD010	4326	26.373	0,7	CMB01	0,6		1	59,0	1,3	15,90	9,3	.0008	.0001	.0009
0007	DD007	4326	26.373	0,7	CMB01	0,6		1	59,0	1,2	11,42	10,2	.0010	.0001	.0011
0008	DD012	4326	30.769	0,7	CMB01	0,6		1	59,0	1,5	10,22	10,8	.0014	.0003	.0017
0009	DD008	4326	30.769	0,7	CMB01	0,6		1	59,0	1,5	15,67	9,0	.0009	.0001	.0010
0010	DD001	4326	26.373	0,7	CMB01	0,6		1	59,0	1,5	15,20	8,7	.0009	.0002	.0011
0011	DD006	4326	30.769	0,7	CMB01	0,6		1	59,0	1,7	21,17	8,3	.0008	.0000	.0008

Figure 4-99 Disk units

- h. In the Storage (GB) field, insert the number of disk units multiplied by the size of a unit. In our example, we have 24 of disk feature 4326, which is a 15 KB RPM 35.16 GB internal disk drive. They are connected through IOA 2780.

- i. In the Storage field, insert the total current disk capacity, by multiplying the capacity of one disk unit with the number of disks. In our example, there are 24 x 35.16 GB disk units, so we insert in the Storage field, 24 x 35.16 GB = 844 GB (see Figure 4-98).

You can also click the WLE **Help/Tutorials** tab for instructions on how to obtain the necessary values to enter in the WLE.

- j. Obtain the Read Ops Per Second value from the Resource report - Disk utilization (see Figure 4-100).

Average Itv End	Average Average I/O /Sec	Average Reads /Sec	Writes /Sec	High K Per I/O	High Avg Util	High High Util	Disk Util Unit	Srv Time	Srv ce Unit d
14:00	357,4	111,5	245,8	11,7	1,6	2,3	0019	,0016	0019 5
14:15	327,4	104,5	222,8	10,4	1,5	2,2	0002	,0016	0002 5
14:30	501,6	188,0	313,6	8,7	2,1	2,9	0005	,0014	0016 1
14:45	132,0	44,2	87,8	6,5	0,6	1,0	0020	,0000	415.487
Average:	329,6	112,0	217,5	9,7	1,5				

Figure 4-100 I/O per second and block size

- k. If the workload is small or if WebFacing or HATS is used, specify the values for in the Additional Characteristics and WebFacing or HATS Support fields. Refer to WLE Help for more information about these fields.
  - l. The System reports are shown in one block size (size of operation) for both reads and writes, so insert this size for both operations. Click **Continue** (see Figure 4-98).
9. The Selected System - Choose Base System panel displays as shown in Figure 4-101. Here you can limit your selection to an existing system, or you can use WLE to size any system for the inserted workload. In our example, we use WLE to size any system. We click the two **Select** buttons.

Version: 2007.3 17-Jul-07 www-912 IBM Systems Workload Estimator

## Selected System

### Choose Base System

---

An existing system has been chosen as one of the workloads. Please make a selection below.

If you chose one of the existing systems, then after all the processing requirements of the various workloads have been calculated, the system selection algorithms will attempt to give preference to the selected system or a system on its upgrade path.

	Workload Name	System Model	Selected Status
<input type="button" value="Select"/>	Existing #1	870-2489/7433	
<input type="button" value="Select"/>	Do not limit selection based on <a href="#">upgrade path information</a> from any existing system. Size to any possible system.		

Figure 4-101 Selecting the sizing to size

10. The Selected System panel displays, as shown in Figure 4-102, on which two recommended models are shown. One model is intended as an *immediate* solution, and the other is meant to *accommodate* the workload growth. You can choose other models and features from Model/Feature and observe the predicted utilization with the existing workload. To size external storage with Disk Magic, click the **External Storage** link in the blue tab area at the top of the Selected System panel.

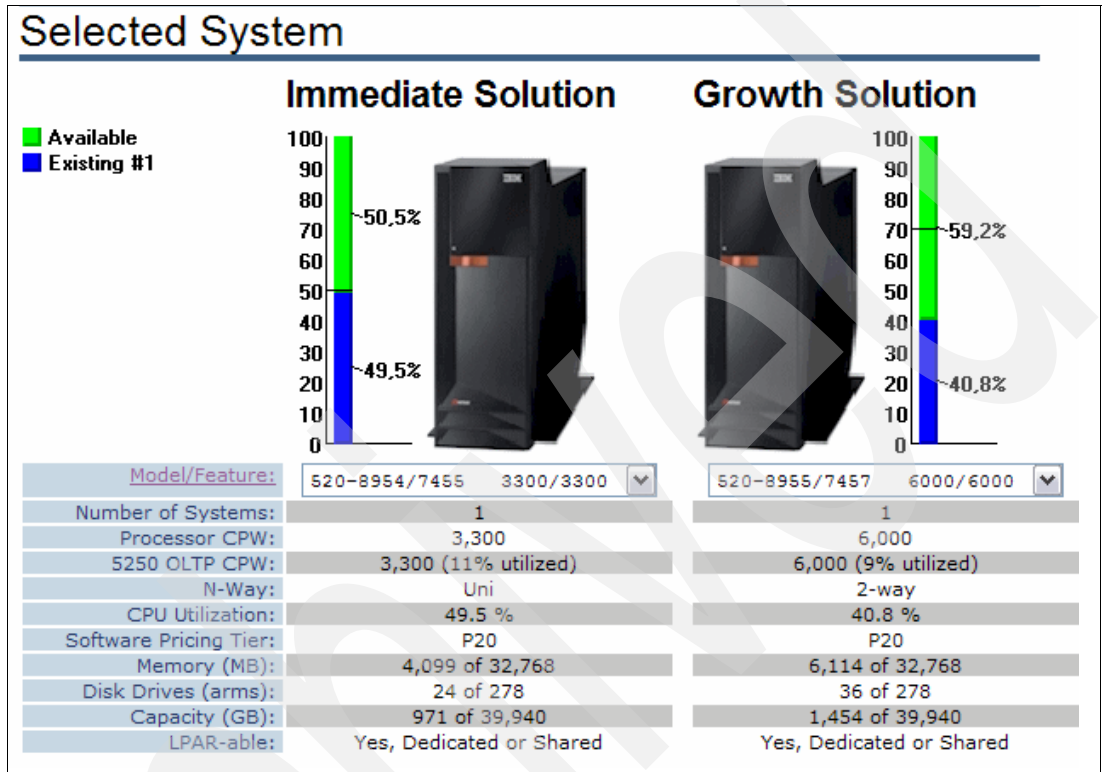


Figure 4-102 Selected system

11. The Selected System - External Storage Sizing Information panel displays as shown in Figure 4-103. For Which system, select either **Immediate** or **Growth** for the system for which you want to size external storage. In our example, we select **Immediate** to size our external storage. Then click **Download Now**.

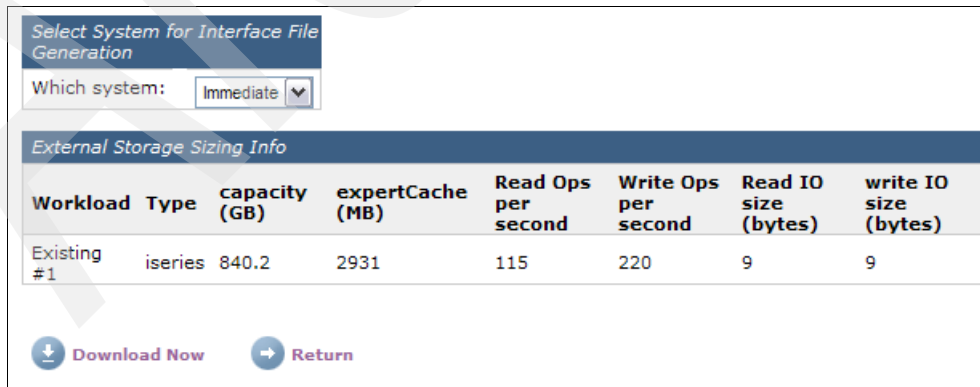


Figure 4-103 Selecting a system to size external storage

12. The File Download window opens. You can choose to start Disk Magic immediately for the sized workload (by clicking **Open**), or you can choose to save the Disk Magic command file and use it later (by clicking **Save**). In our example, we want to start Disk Magic immediately, so we click **Open**.

**Important:** At this point, to start Disk Magic, you must have Disk Magic installed.

13. Disk Magic starts with the workload modeled with WLE (see Figure 4-104). Observe that the workload Existing #1 is already shown under TreeView. Double-click **dss1**.

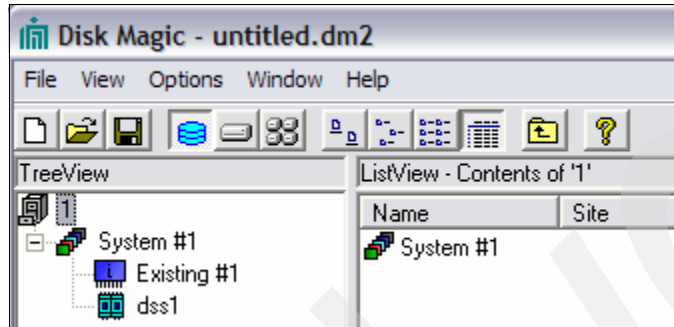


Figure 4-104 Disk Magic model

14. The Disk Subsystem - dss1 window (Figure 4-105) opens, displaying the General tab. Follow these steps:

- a. To size DS6800 for the Existing #1 workload, from Hardware Type, select **DS6800**. We highly recommend that you use multipath with DS6800. To model multipath, select **Multipath with iSeries**.

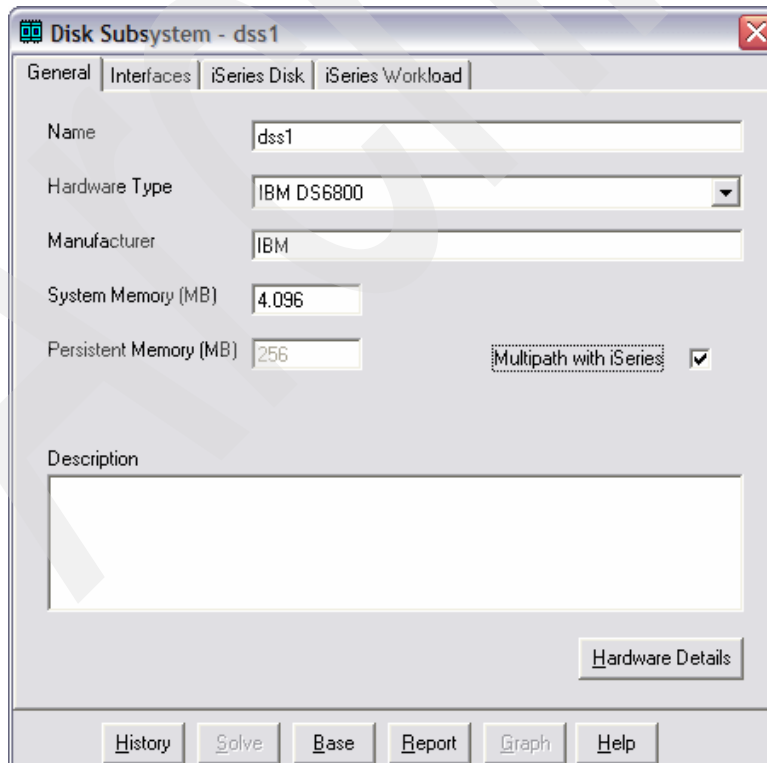


Figure 4-105 Selecting DS6800 and Multipath with iSeries

- b. Select the Interfaces tab and then select the From Servers tab (see Figure 4-106). Observe that four interfaces from servers with workload Existing #1 are already configured as the default. In our example, we plan four System i5 FC adapters in multipath so we leave this default value. If necessary, you can change it by clicking **Edit** and specifying the number of interfaces.

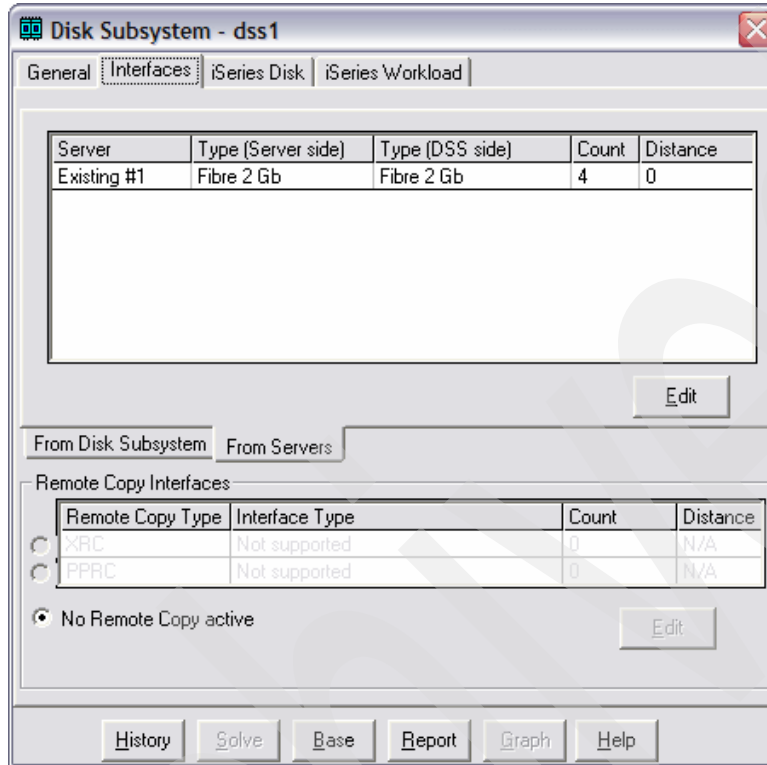


Figure 4-106 Interfaces from the server

- c. Click the From Disk Subsystem tab. Notice that four interfaces from DS6000 are configured as the default. In our example, we use two DS6000 host ports for connecting to the System i5 platform, so we change the number of interfaces. Click **Edit** to open the Edit Interfaces for Disk Subsystem window. In the Count field, enter the number of planned DS6000 ports. Click **OK**. In our example, we insert two ports as shown in Figure 4-107.




Figure 4-107 Interfaces from the DS

- d. In the Disk Subsystem - dss1 window, click the iSeries Disk tab (Figure 4-108). Observe that an extent pool is already configured for the Existing #1 workload. Its capacity is equal to the capacity that you specified in the Storage field of the WLE.

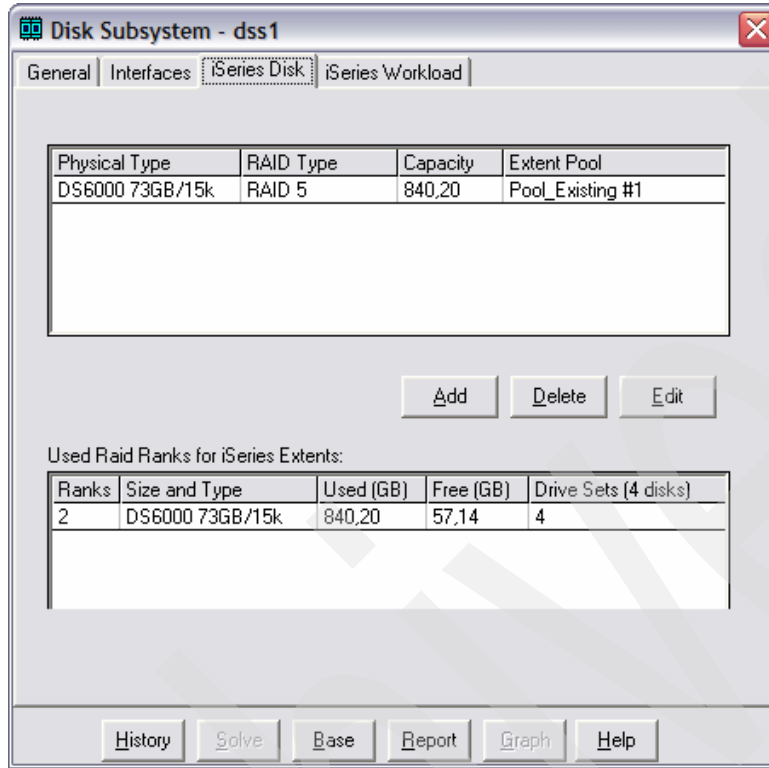


Figure 4-108 Capacity in Disk Magic

- e. In the Disk Subsystem - dss1 window, select the iSeries Workload tab. Notice that the number of reads per second and writes per second, the number of LUNs, and the capacity are specified based on values that you inserted in WLE. You might want to check the modeled expert cache size, by comparing it to the sum of all expert cache storage pools in the System report (Figure 4-109).

Pool ID	Expert Cache	Size (KB)	Act Lvl	CPU Util	Number Tns	Average Response	----- DB ----- Fault Pages	---- Non-DB ---- Fault Pages	Act- Wait
01	0	808.300	0	28,5	0	0,00	0,0 0,0	0,3 1,0	257 0
*02	3	1.812.504	147	15,7	825	0,31	3,8 17,9	32,7 138,8	624 0
*03	3	1.299.488	48	9,6	4.674	0,56	2,4 13,0	28,1 107,0	198 0
04	3	121.244	5	0,0	0	0,00	0,0 0,0	0,0 0,0	0 0
Total		4.041.536		53,9	5.499		6,3 31,0	61,2 246,9	1.080 0

Figure 4-109 Expert cache

- f. Enter the block size that was used for WLE, if needed (Figure 4-110). Click **Cache Statistics**.

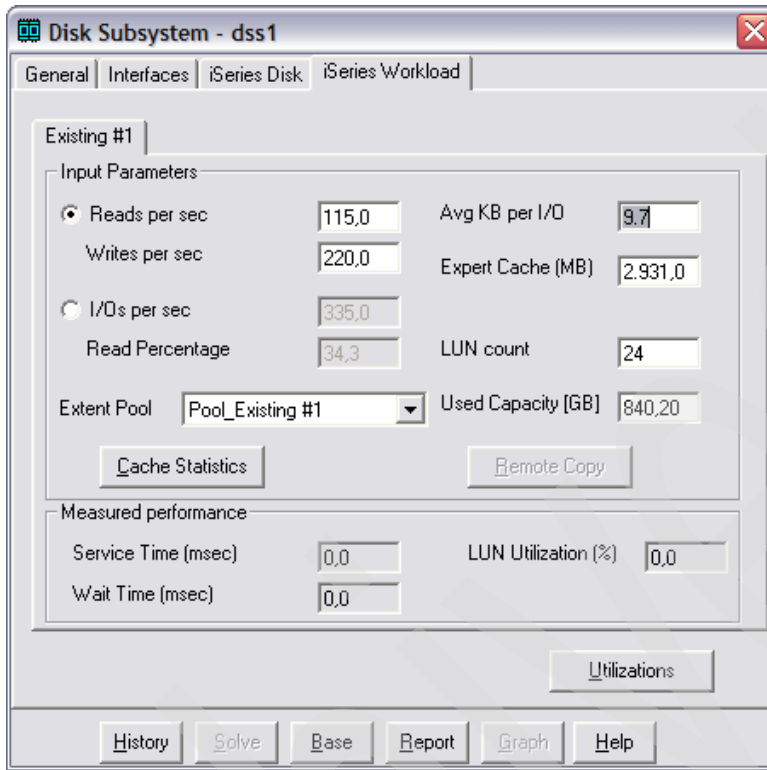


Figure 4-110 Workload in Disk Magic

- g. The Cache Statistics for Host Existing #1 window (Figure 4-111) opens. Notice that the cache statistics are already specified in the Disk Magic model. For more conservative sizing, you might want to change them to lower values, such as 20% read cache and 30% write cache. Then, click **OK**.

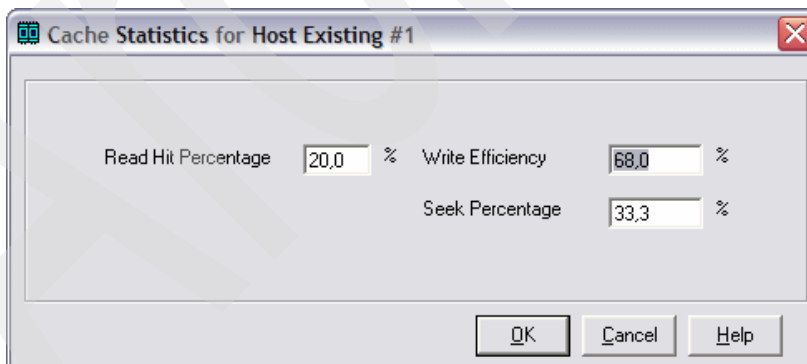


Figure 4-111 Cache values in Disk Magic



- h. On the Disk Subsystem - dss1 window, click **Base** to save the current model as base. After the base is saved successfully, notice the modeled disk service time and wait time, as shown in Figure 4-112.

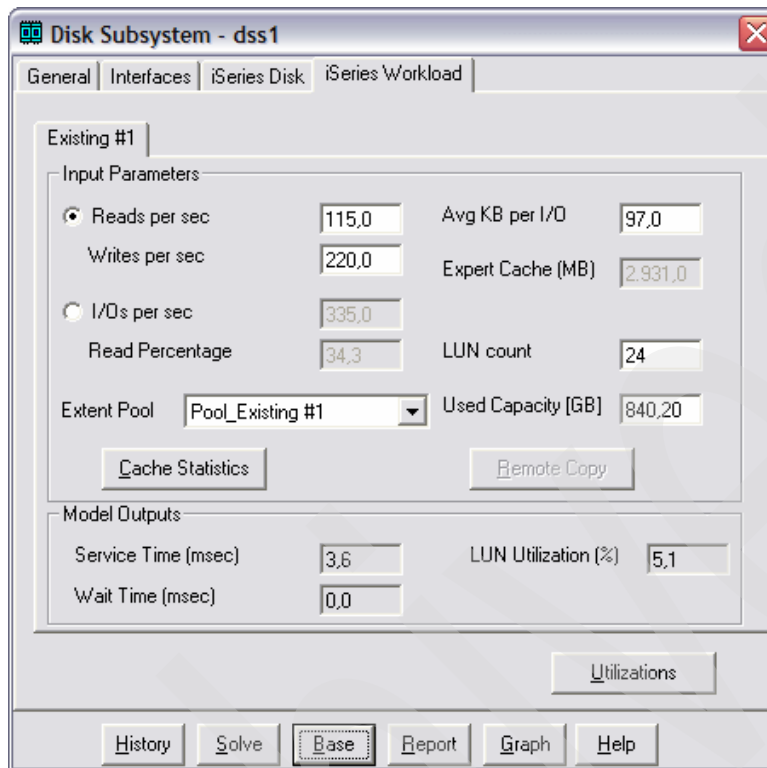


Figure 4-112 Modeled service and wait times

- i. On the iSeries Workload tab, click **Utilizations**. The Utilizations IBM DS6800 window (Figure 4-113) opens. Observe the modeled utilizations for the existing workload. In our example, the modeled hard disk drive (HDD) utilization and LUN utilization are far below the limits that are recommended for good performance. There is room for growth in the modeled DS configuration.

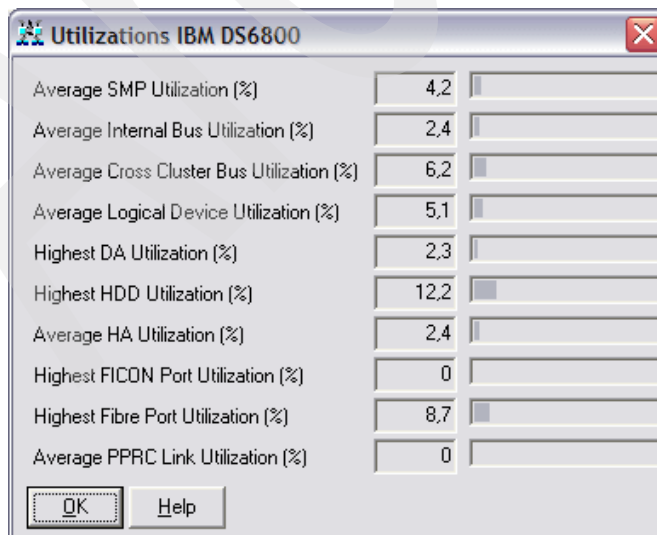


Figure 4-113 Modeled utilizations

Archived

## Implementation and additional topics

This part covers different implementation methods and additional topics concerning external storage on System i. It has the following chapters:

- ▶ Chapter 5, “Implementing external storage with i5/OS” on page 181
- ▶ Chapter 6, “Implementing FlashCopy using the DS CLI” on page 253
- ▶ Chapter 7, “Implementing Metro Mirror using the DS CLI” on page 279
- ▶ Chapter 8, “Implementing Global Mirror using the DS CLI” on page 301
- ▶ Chapter 9, “Copy Services scenarios” on page 337
- ▶ Chapter 10, “Creating storage space for Copy Services using the DS GUI” on page 347
- ▶ Chapter 11, “Implementing FlashCopy using the DS GUI” on page 363
- ▶ Chapter 12, “Implementing Metro Mirror using the DS GUI” on page 373
- ▶ Chapter 13, “Managing Copy Services in i5/OS environments using the DS GUI” on page 387
- ▶ Chapter 14, “Performance considerations” on page 423
- ▶ Chapter 15, “FlashCopy usage considerations” on page 431

Archived



## Implementing external storage with i5/OS

In this chapter, we discuss the supported environment for external storage including the logical volumes that are supported and the protection methods that are available. We also show how to add the logical volumes to the System i environment.

## 5.1 Supported environment

Not all hardware and software combinations for i5/OS support DS8000 and DS6000. This section describes the hardware and software prerequisites for attaching DS8000 and DS6000.

### 5.1.1 Hardware

DS8000, DS6000, and ESS model 800 are supported on all System i models that support Fibre Channel (FC) attachment for external storage. Fibre channel was supported on all iSeries 8xx models and later. AS/400 models 7xx and earlier only supported SCSI attachment for external storage so they cannot support DS8000 or DS6000.

The following IOP-based FC adapters for System i support DS8000 and DS6000:

- ▶ 2766 2 Gb Fibre Channel Disk Controller PCI
- ▶ 2787 2 Gb Fibre Channel Disk Controller PCI-X
- ▶ 5760 4 Gb Fibre Channel Disk Controller PCI-X

Each of these adapters requires its own dedicated I/O processor.

With System i POWER6 new IOP-less FC adapters are available which only support IBM System Storage DS8000 on LIC level 2.4.3 or later for external disk storage attachment:

- ▶ 5749 IOP-less 4 Gb dual-port Fibre Channel Disk Controller PCI-X
- ▶ 5774 IOP-less 4 Gb dual-port Fibre Channel Disk Controller PCIe

For further planning information with these System i FC adapters, refer to 3.2, “Solution implementation considerations” on page 54.

For information about current hardware requirements, including support for switches, refer to:

[http://www-1.ibm.com/servers/eserver/series/storage/storage\\_hw.html](http://www-1.ibm.com/servers/eserver/series/storage/storage_hw.html)

To support boot from SAN with the load source unit on external storage, either the #2847 I/O processor (IOP) or an IOP-less FC adapter is required.

**Restriction:** Prior to i5/OS V6R1 the #2847 IOP for SAN load source does not support multipath for the load source unit but does support multipath for all other logical unit numbers (LUNs) attached to this I/O processor (IOP). See 5.10, “Protecting the external load source unit” on page 215 for more information.

### 5.1.2 Software

The iSeries or System i environment must be running V5R3, V5R4 or V6R1 of i5/OS. In addition the following PTFs are required:

- ▶ V5R3
  - MF33328
  - MF33845
  - MF33437
  - MF33303
  - SI14690
  - SI14755
  - SI14550

- ▶ V5R3M5 and later
  - Load source must be at least 17.54 GB

**Important:**

- ▶ The #2847 PCI-X IOP for SAN load source requires i5/OS V5R3M5 or later.
- ▶ The #5760 FC I/O adapter (IOA) requires V5R3M0 resave RSI or V5R3M5 RSB with C6045530 or later (ref. #5761 APAR II14169) and for System i5 firmware level SF235\_160 or later
- ▶ The #5749/#5774 IOP-less FC IOA is supported on System i POWER6 models only

Prior to attaching a DS8000, DS6000, or ESS model 800 system to a System i model, check for the latest PTFs, which probably have superseded the minimum requirements listed previously.

**Note:** We generally recommend installing one of the latest i5/OS cumulative PTFs (cumPTFs) before attaching IBM System Storage external disk storage subsystems to System i.

## 5.2 Logical volume sizes

i5/OS is supported on the DS8000 and DS6000 system using fixed block storage. Unlike other Open Systems that use the fixed block architecture, i5/OS supports only specific volume sizes, which might not be an exact number of extents. In general the LUN sizes relate to the volume sizes available with System i internal disk devices. i5/OS volumes are defined in decimal GB (10<sup>9</sup>bytes).

Table 5-1 indicates the number of extents that are required for different System i volume sizes. The value xxxx represents 1750 for DS6000 and 2107 for DS8000.

Table 5-1 i5/OS logical volume sizes

Model type		i5/OS Device size (GB)	Number of logical block addresses	Extents	Unusable space (GiB) <sup>a</sup>	Usable space%
Protected	Unprotected					
xxxx-A01	xxxx-A81	8.59	16,777,216	8	0.00	100.00
xxxx-A02 <sup>b</sup>	xxxx-A82	17.54	34,275,328	17	0.66	96.14
xxxx-A05 <sup>b</sup>	xxxx-A85	35.16	68,681,728	33	0.25	99.24
xxxx-A04 <sup>b</sup>	xxxx-A84	70.56	137,822,208	66	0.28	99.57
xxxx-A06 <sup>b</sup>	xxxx-A86	141.12	275,644,416	132	0.56	99.57
xxxx-A07	xxxx-A87	282.25	551,288,832	263	0.13	99.95

a. GiB represents “Binary GB” (2<sup>30</sup> bytes) and GB represents “Decimal GB” (10<sup>9</sup> bytes).

b. Only Ax2, Ax4, Ax5, and Ax6 models are supported as external load source unit LUNs.

When creating the logical volumes for use with i5/OS, in almost every case, the i5/OS device size does not match a whole number of extents, so some space remains unused. Use the values in Table 5-1 in conjunction with extent pools to see how much space will be wasted for your specific configuration. Also, note that the #2766, #2787, and #5760 Fibre Channel Disk

Adapters used by the System i platform can only address up to 32 LUNs while the IOP-less FC adapter #5749 and #5774 support up to 64 LUNs per port.

For more information about sizing guidelines for i5/OS, refer to Chapter 4, “Sizing external storage for i5/OS” on page 89.

## 5.3 Protected versus unprotected volumes

When defining i5/OS logical volumes, you must decide whether these should be *protected* or *unprotected* volume models. This protection mode is simply a SCSI Inquiry data notification to i5/OS and does not mean that the data is protected or unprotected. In reality, all DS8000 or DS6000 LUNs are protected, either by RAID-5 or RAID-10. An *unprotected* volume is available for i5/OS to mirror that volume to another volume of equal capacity, either internal or external. Unless you intend to use i5/OS (host-based) mirroring, you should define your logical volumes as protected.

Under some circumstances, you might want to mirror the i5/OS internal load source unit to a LUN in the DS8000 or DS6000 storage system. In this case, define only one LUN as unprotected. Otherwise, when mirroring is started to mirror the load source unit to the DS6000 or DS8000 LUN, i5/OS attempts to mirror all unprotected volumes.

**Important:** Prior to i5/OS V6R1, we strongly recommend that if you use an external load source unit that you use i5/OS mirroring to another LUN in external storage system to provide path protection for the external load source unit (see 5.10, “Protecting the external load source unit” on page 215).

### 5.3.1 Changing LUN protection

Although it is possible to change a volume from protected to unprotected (or vice versa) using the DS command-line interface (CLI) `chfbvo1` command, you need to be *extremely careful* when changing LUN protection.

**Attention:** Changing the LUN protection of a System i volume is only supported for non-configured volumes, that is volumes not a part of the System i auxiliary storage pool configuration.

If the volume is configured, that is within an auxiliary storage pool (ASP) configuration, do not change the protection. In this case if you want to change the protection, you must remove the volume from the ASP configuration first and add it back later after having changed its protection mode. This process is unlike ESS models E20, F20, and 800 where from storage side no dynamic change of the LUN protection mode is supported so that the logical volume would have to be deleted, requiring the entire array that contains the logical volume to be reformatted, and created new with the desired other volume protection mode.

**Important:** Removing a logical volume from the System i configuration is an i5/OS disruptive task if the LUN is in the system auxiliary storage pool (ASP) or user ASPs 2 through 32 because it requires an initial program load (IPL) of i5/OS to completely remove the volume from the i5/OS configuration. However volumes can be removed from an independent ASP (IASP) with the IASP varied off without performing an IPL on the system. This is no difference from removing an internal disk from an i5/OS configuration.



## 5.4 Setting up an external load source unit

The new #5749 and #5774 IOP-less Fibre Channel IOAs for System i POWER6 allow to perform an IPL from a LUN in the IBM System Storage DS8000 series.

The #2847 PCI-X IOP for SAN load source allows a System i to perform an IPL from a LUN in a DS6000, DS8000, or ESS model 800. This IOP supports only a single FC IOA. No other IOAs are supported.

### Restrictions:

- ▶ The new IOP-less Fibre Channel IOAs #5749 or and #5774 support for *direct* attachment the FC-AL protocol only.
- ▶ For the #2847 IOP driven IOAs Point-to-Point (also known as FC-SW and SCSI-FCP) is the only support protocol. You must *not* define the host connection (DS CLI) or the Host Attachment (Storage Manager GUI) as FC-AL because this prevents you from using the system.

Creating a new load source unit on external storage is similar to creating one on an internal drive. However, instead of tagging a RAID disk controller for the internal load source unit, you must tag your load source IOA for the SAN load source.

**Note:** With System i SLIC V5R4M5 and later all buses and IOPs are booted in the D-mode IPL environment and if no existing loadsource disk unit is found, a list of eligible disk units (of the correct capacity) displays for the user to select the disk to use as the loadsource disk.

For previous SLIC versions, we recommend that you assign only your designated load source LUN to your load source IOA first to make sure that this is the LUN chosen by the system for your load source at SLIC install. Then, assign the other LUNs to your load source IOA afterwards.

### 5.4.1 Tagging the load source IOA

Even if you are only creating a system with one partition, you must use a Hardware Management Console (HMC) to tag the load source IOA. This tells the system which IOA to use when building the load source unit during the D-mode IPL SLIC installation. The external load source unit does not work on a system without an HMC.

On the HMC, set the tagged load source unit to the FC Disk Controller that is controlling your new external load source unit. On the HMC, follow these steps:

1. Select the partition name with which you are working. Then, select **Tasks** → **Configuration** → **Manage Profiles** as shown in Figure 5-1.

**Note:** For below HMC V7, right-click the partition profile name and select **Properties**.




Figure 5-1 Selecting the HMC partition profile properties

2. Select **Actions** → **Edit** as shown in Figure 5-2

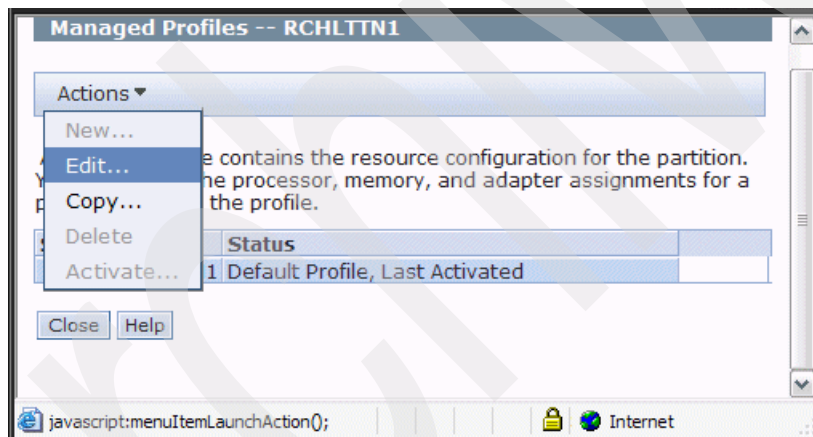


Figure 5-2 Managed Profiles

3. In the Logical Partition Profile Properties window (Figure 5-3), select the Tagged I/O tab.

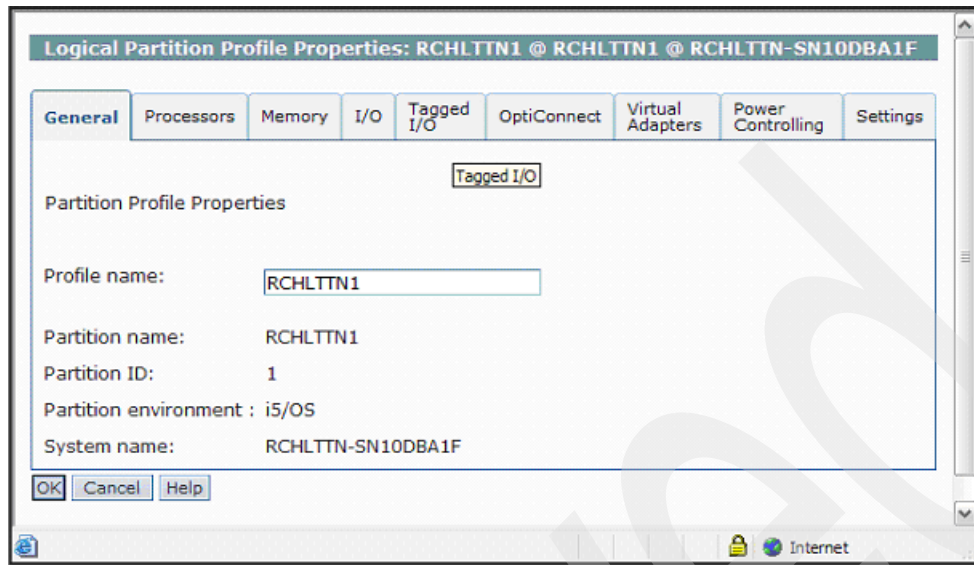


Figure 5-3 Logical Partition Profile Properties window

4. On the Tagged I/O tab, click the **Select** button that corresponds to the load source as shown in Figure 5-4.

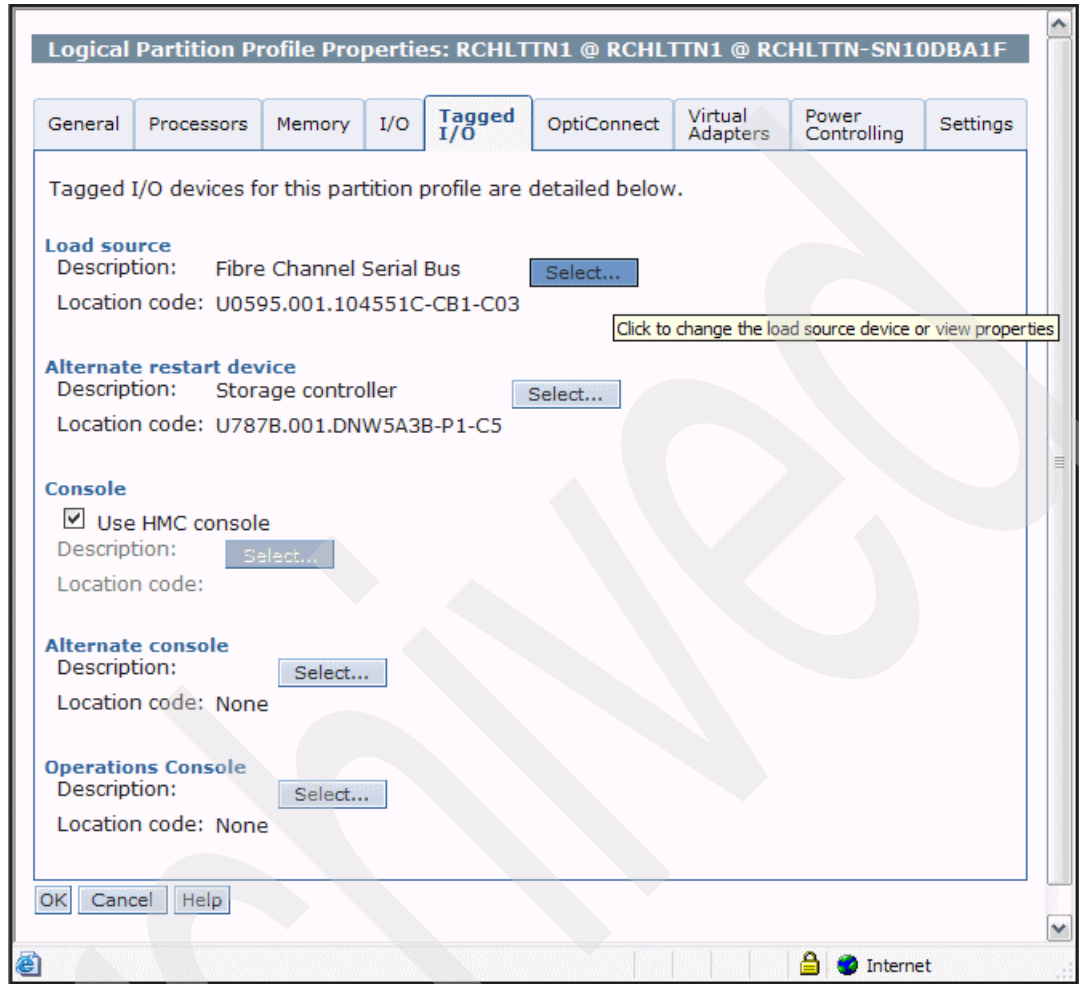


Figure 5-4 Tagged I/O properties

5. In the Load Source Device window (Figure 5-5), select the IOA to which your new load source unit is assigned. Click **OK**.

**Profile I/O devices**

Select	Unit	Bus	Slot	Added	Description	Location Code
<input type="radio"/>	U0595.001.104551C	13	C01	Required	I/O Processor	U0595.001.104551C-CB1-C01
<input type="radio"/>	U0595.001.104551C	13	C02	Required	Empty slot	U0595.001.104551C-CB1-C02
<input checked="" type="radio"/>	U0595.001.104551C	13	C03	Required	Fibre Channel Serial Bus	U0595.001.104551C-CB1-C03
<input type="radio"/>	U0595.001.104551C	13	C04	Required	Empty slot	U0595.001.104551C-CB1-C04
<input type="radio"/>	U5294.002.109672C	17	C11	Required	PCI I/O Processor	U5294.002.109672C-CB1-C11
<input type="radio"/>	U5294.002.109672C	17	C12	Required	PCI 1Gbps Ethernet UTP	U5294.002.109672C-CB1-C12
<input type="radio"/>	U5294.002.109672C	17	C13	Required	PCI I/O Processor	U5294.002.109672C-CB1-C13
<input type="radio"/>	U5294.002.109672C	17	C14	Required	Storage controller	U5294.002.109672C-CB1-C14
<input type="radio"/>	U5294.002.109672C	17	C15	Required	Empty slot	U5294.002.109672C-CB1-C15
<input type="radio"/>	U5294.002.109672C	18	C01	Required	PCI I/O Processor	U5294.002.109672C-CB1-C01
<input type="radio"/>	U5294.002.109672C	18	C02	Required	PCI-X Ultra RAID Disk Controller	U5294.002.109672C-CB1-C02
<input type="radio"/>	U5294.002.109672C	18	C03	Required	PCI RAID Disk Unit Controller	U5294.002.109672C-CB1-C03
<input type="radio"/>	U5294.002.109672C	18	C04	Required	PCI 1Gbps Ethernet UTP	U5294.002.109672C-CB1-C04
<input type="radio"/>	U5294.002.109672C	19	C05	Required	I/O Processor	U5294.002.109672C-CB1-C05
<input type="radio"/>	U5294.002.109672C	19	C06	Required	Fibre Channel Serial Bus	U5294.002.109672C-CB1-C06
<input type="radio"/>	U5294.002.109672C	19	C07	Required	Empty slot	U5294.002.109672C-CB1-C07
<input type="radio"/>	U5294.002.109672C	19	C08	Required	Empty slot	U5294.002.109672C-CB1-C08
<input type="radio"/>	U5294.002.109672C	19	C09	Required	Empty slot	U5294.002.109672C-CB1-C09

Load source: Fibre Channel Serial Bus

Location code: U0595.001.104551C-CB1-C03

Figure 5-5 Tagging the load source unit

6. Change the partition to do a manual IPL as follows:
  - a. Select **Tasks** → **Properties** from the drop-down menu as shown in Figure 5-6.

**Contents of: RCHLTTN-SN10DBA1F**

Select	Name	ID	Status	Pr	Ur	Environment	Reference Code
<input checked="" type="checkbox"/>	RCHLTTN1	1	Running			i5/OS	00000000
<input type="checkbox"/>	RCHLTTN2	2	Not Activated			i5/OS	00000000

Total: 2 Filtered: 2 Selected: 0

Tasks ▾  
 Properties  
 Change Default Profile  
 Operations  
 Configuration  
 Dynamic Logical Partitioning  
 Serviceability

Figure 5-6 Selecting the Properties option

**Note:** For below HMC V7, right-click the partition name and select **Properties**.

b. In the Partition Properties window (Figure 5-7), select the Settings tab.

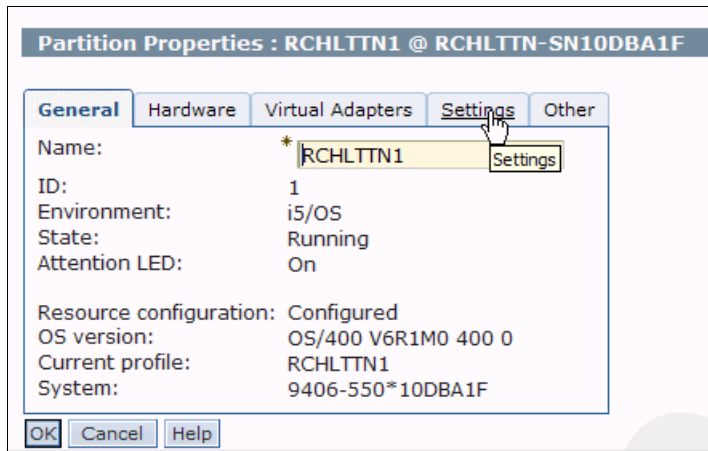


Figure 5-7 Partition Properties window

c. On the Settings tab, for Keylock position, select **Manual** as shown in Figure 5-8.

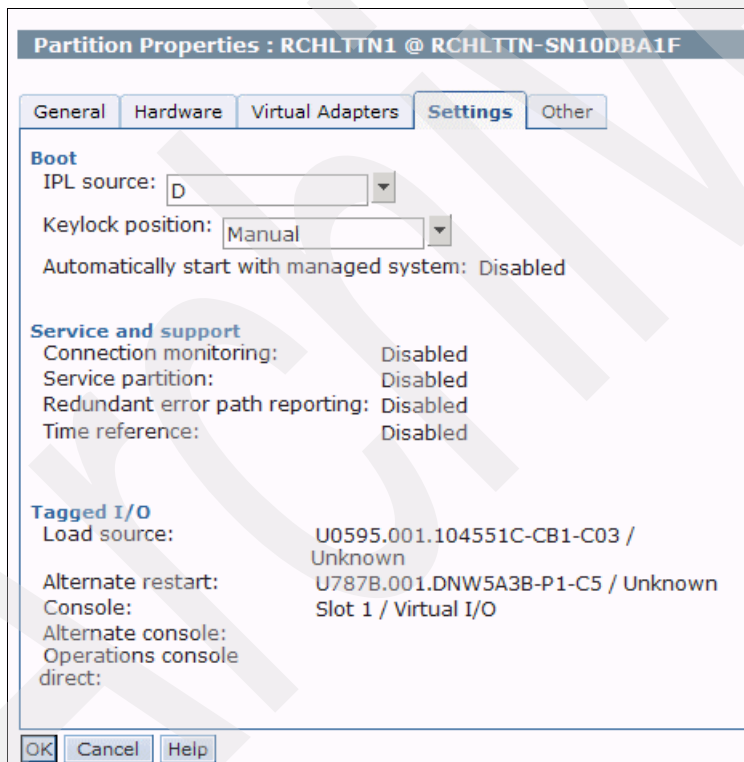


Figure 5-8 Setting the IPL type

## 5.4.2 Creating the external load source unit

After you tag the load source IOA, the installation process is the same as installing on an internal load source unit. Follow these steps:

1. Insert the I\_BASE SLIC CD into the alternate IPL DVD-ROM device and perform a D-mode IPL by selecting the partition and choosing **Tasks** → **Operations** → **Activate** as shown in Figure 5-9.

**Note:** For below HMC V7, right-click the partition, select **Properties**, and click **Activate**.




Figure 5-9 Activating a partition

2. In the Activate Logical Partition window (Figure 5-10), select the partition profile to be used and click **OK**.

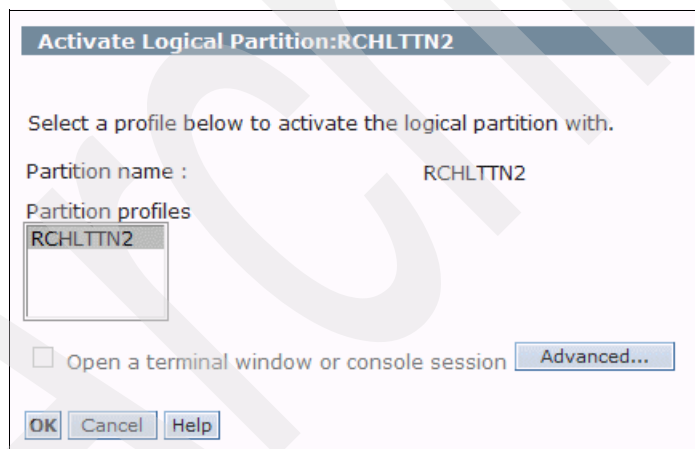


Figure 5-10 Selecting the profile for activation

In the HMC, a status window displays, which closes when the task is complete and the partition is activated. Wait for the Dedicated Service Tools (DST) panel to open.

3. After the system has done an IPL to DST, select **3. Use Dedicated Service Tools (DST)**.

4. On the OS/400 logo panel (Figure 5-11), enter the language feature code.




Figure 5-11 OS/400 logo panel

5. On the Confirm Language Group panel (Figure 5-12), press Enter to confirm the language code.

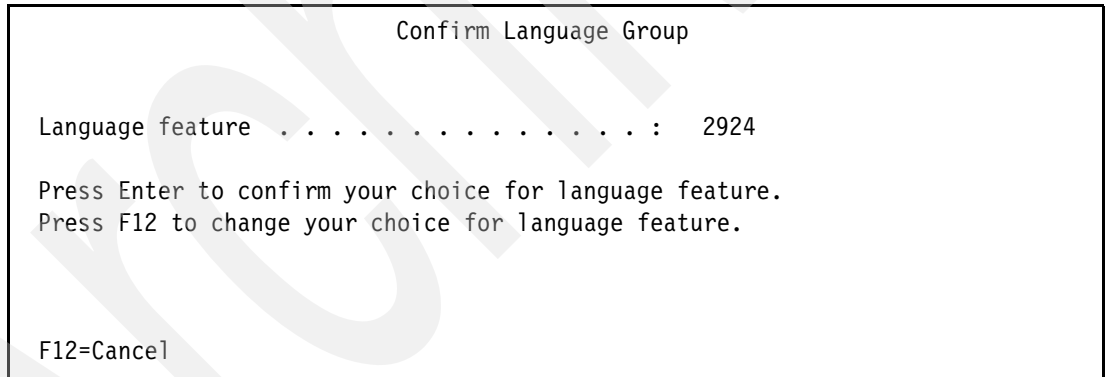


Figure 5-12 Confirming the language feature



6. On the Install Licensed Internal Code panel (Figure 5-13), select **1. Install Licensed Internal Code**.

Install Licensed Internal Code

System: G1016730

Select one of the following:

- 1. Install Licensed Internal Code**
- 2. Work with Dedicated Service Tools (DST)
- 3. Define alternate installation device

Selection  
1

Figure 5-13 Install Licensed Internal Code panel

7. The next panel shows the volume that is selected as the external load source unit and a list of options for installing the Licensed Internal Code (see Figure 5-14). Select **2. Install Licensed Internal Code and Initialize System**.

Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:

Serial Number	Type	Model	I/O Bus	Controller	Device
30-02000	1750	A85	0	1	1

Select one of the following:

- 1. Restore Licensed Internal Code
- 2. Install Licensed Internal Code and Initialize system**
- 3. Install Licensed Internal Code and Recover Configuration
- 4. Install Licensed Internal Code and Restore Disk Unit Data
- 5. Install Licensed Internal Code and Upgrade Load Source

Selection  
2

F3=Exit                  F12=Cancel

Figure 5-14 Install Licensed Internal Code options

8. On the Confirmation panel, read the warning message that displays (as shown in Figure 5-15) and press F10=Continue when you are sure that you want to proceed.

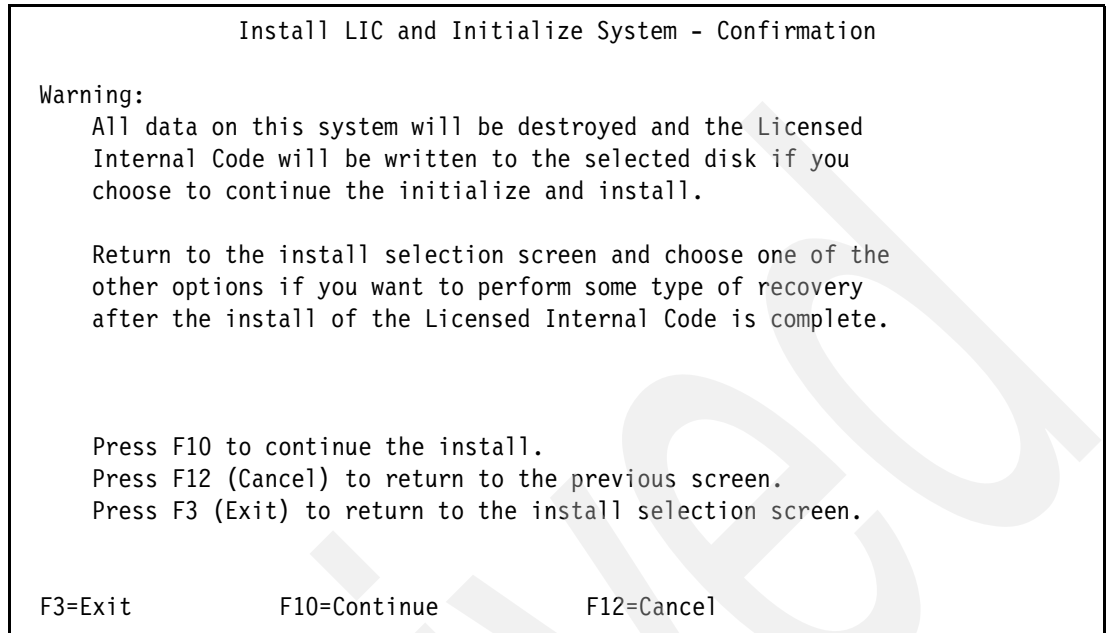


Figure 5-15 Confirmation warning

9. The Initialize the Disk - Status panel displays for a short time (see Figure 5-16). Unlike internal drives, formatting external LUNs on DS8000 and DS6000 is a task that is run by the storage system in the background, that is the task might complete faster than you expect.




Figure 5-16 Initialize the Disk - Status panel

When the logical formatting has finished, you see the Install Licensed Internal Code - Status panel as shown in Figure 5-17.




Figure 5-17 Install Licensed Internal Code status

When the Install Licensed Internal Code process is complete, the system does another IPL to DST automatically. You have now built an external load source unit.

## 5.5 Adding volumes to the System i5 configuration

After the logical volumes are created and assigned to the host, they appear as *non-configured units* to i5/OS. It can take some time for i5/OS to recognize the logical volumes after they are created. At this stage, they are used in exactly the same way as non-configured internal units. There is nothing particular to external logical volumes as far as i5/OS is concerned. You should use the same functions for adding logical units to an ASP as you would for internal disks.

Adding disk units to the configuration can be done either by using the 5250 interface with Dedicated Service Tools (DST) or System Service Tools (SST) or with iSeries Navigator.

## 5.5.1 Adding logical volumes using the 5250 interface

To add a logical volume in the DS8000 or DS6000 to the system ASP, using System Service Tools (SST), follow these steps:

1. Enter the command STRSST and sign on System Service Tools.
2. In the System Service Tools (SST) panel (Figure 5-18), select **3. Work with disk units**.

```
System Service Tools (SST)

Select one of the following:

1. Start a service tool
2. Work with active service tools
3. Work with disk units
4. Work with diskette data recovery
5. Work with system partitions
6. Work with system capacity
7. Work with system security
8. Work with service tools user IDs

Selection
 3

F3=Exit      F10=Command entry  F12=Cancel
```

Figure 5-18 System Service Tools menu

3. In the Work with Disk Units panel (Figure 5-19), select **2. Work with disk configuration**.

```
Work with Disk Units

Select one of the following:

1. Display disk configuration
2. Work with disk configuration
3. Work with disk unit recovery

Selection
 2

F3=Exit      F12=Cancel
```

Figure 5-19 Work with Disk Units panel

- When adding disk units to a configuration, you can add them as empty units by selecting Option 2, or you can allow i5/OS to balance the data across all the disk units. Normally, we recommend that you balance the data. In the Work with Disk Configuration panel (Figure 5-20), select **8. Add units to ASPs and balance data**.

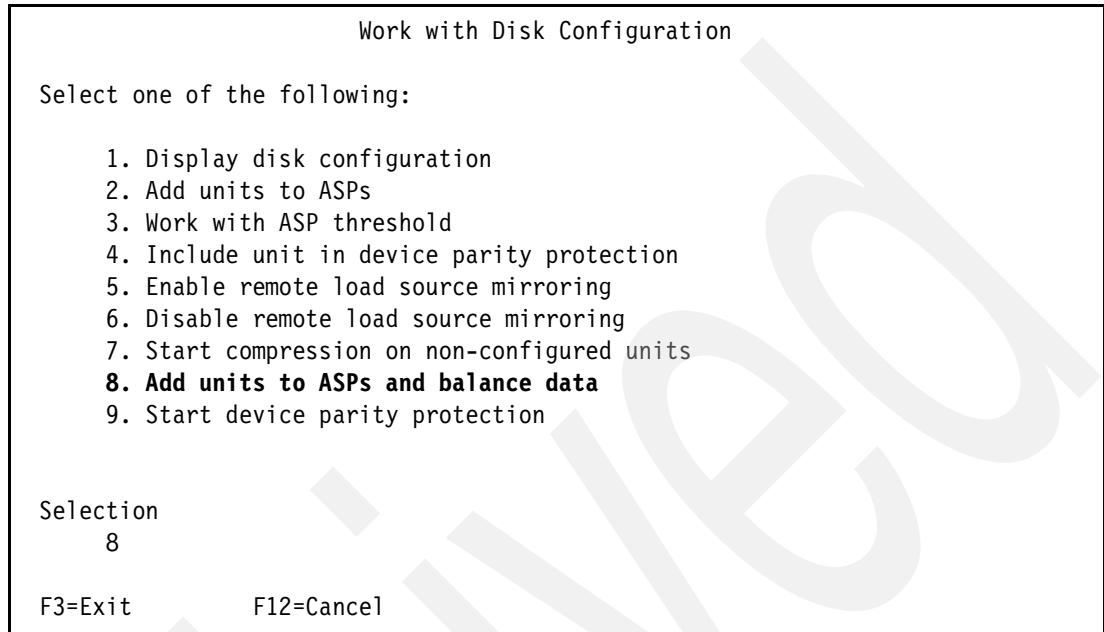


Figure 5-20 Work with Disk Configuration panel

- In the Specify ASPs to Add Units to panel (Figure 5-21), specify the ASP number next to the desired units. Here, we specify 1 for ASP, which is the System ASP. Press Enter.




Figure 5-21 Specify ASPs to Add Units to panel

- In the Confirm Add Units panel (Figure 5-22), review the information and verify that everything is correct. If the information is correct, press Enter to continue. Depending on the number of units that you are adding, this step can take some time to complete.

Confirm Add Units

Add will take several minutes for each unit. The system will have the displayed protection after the unit(s) are added.

Press Enter to confirm your choice for Add units.  
 Press F9=Capacity Information to display the resulting capacity.  
 Press F12=Cancel to return and change your choice.

ASP	Unit	Serial Number	Type	Model	Resource Name	Protection
1						Unprotected
	1	02-89058	6717	074	DD004	Device Parity
	2	68-0CA4E32	6717	074	DD003	Device Parity
	3	68-0C9F8CA	6717	074	DD002	Device Parity
	4	68-0CA5D96	6717	074	DD001	Device Parity
	5	75-1118707	2107	A85	DD006	Unprotected

F9=Resulting Capacity                      F12=Cancel

Figure 5-22 Confirm Add Units panel

- After the units are added, view your disk configuration to verify the capacity and data protection.

## 5.5.2 Adding volumes to an independent auxiliary storage pool

IASPs can be defined as *switchable* or *private*. Disks must be added to an IASP using the iSeries Navigator. That is, you cannot manage your IASP disk configuration from the 5250 interface. In this example, we add a logical volume to a private (non-switchable) IASP. Follow these steps:

1. Start iSeries Navigator. Figure 5-23 shows the initial window.

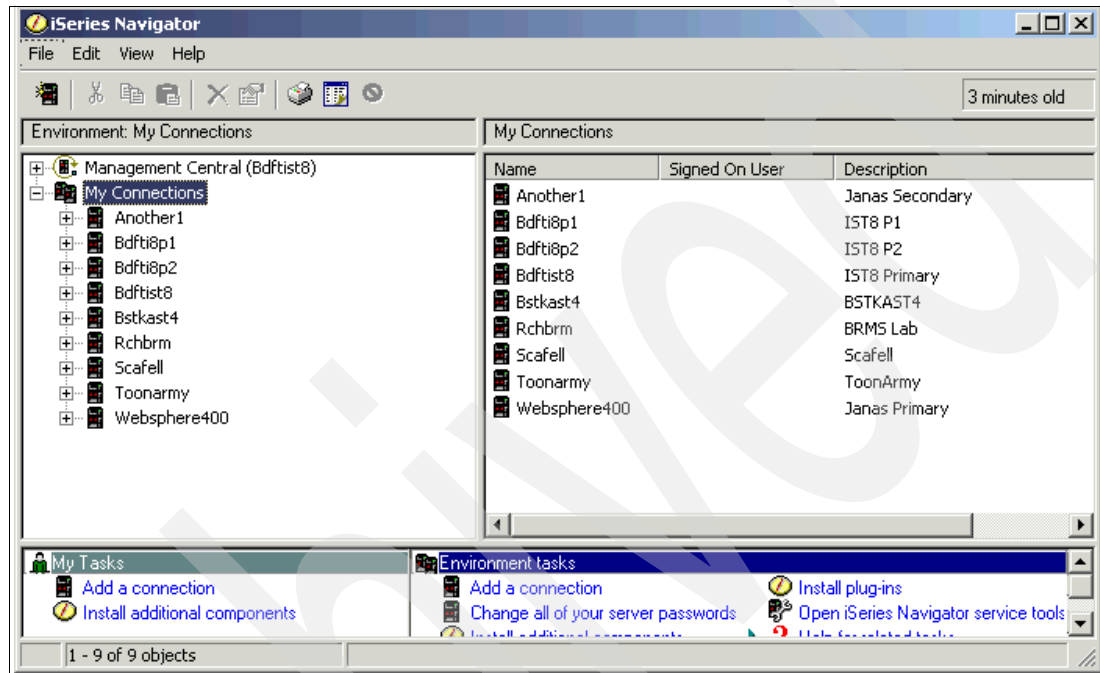


Figure 5-23 iSeries Navigator initial window

- Expand the iSeries to which you want to add the logical volume and sign on to that server. Then expand **Configuration and Service** → **Hardware** → **Disk Units** (see Figure 5-24).




Figure 5-24 Series Navigator Disk Units

- Sign on to SST. Enter your Service tools ID and password and then click **OK**.



4. Under Disk Units, right-click **Disk Pools**, and select **New Disk Pool** as shown in Figure 5-25.

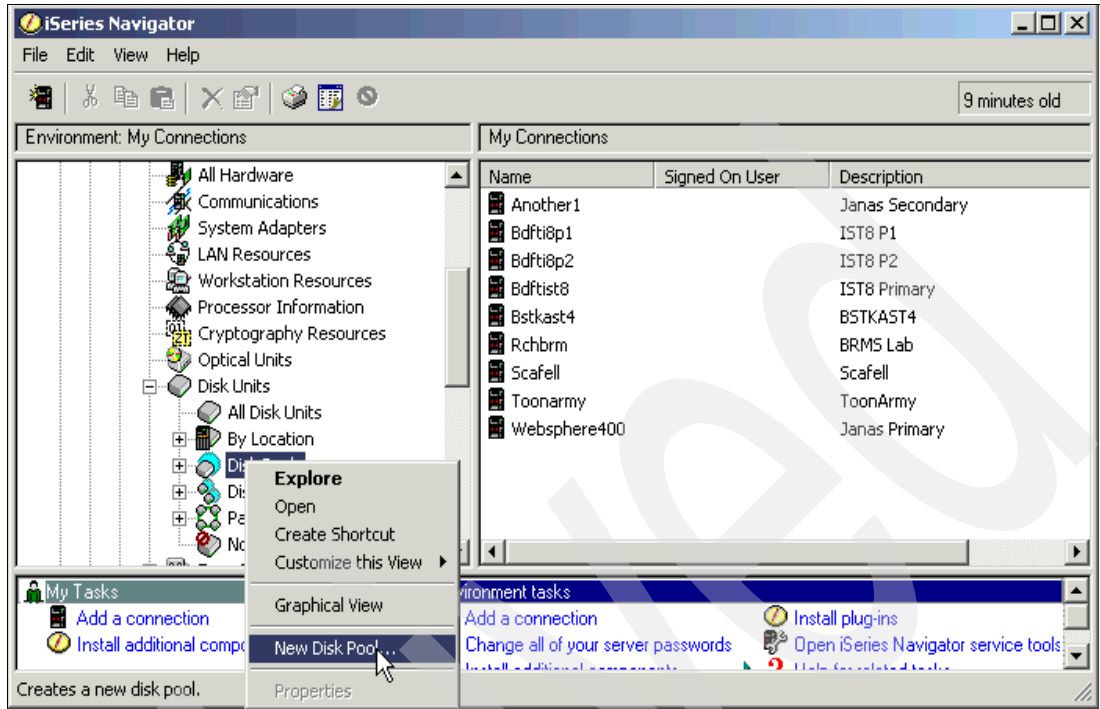


Figure 5-25 Creating a new disk pool

5. The New Disk Pool wizard opens. Figure 5-26 shows the Welcome window. Click **Next**.

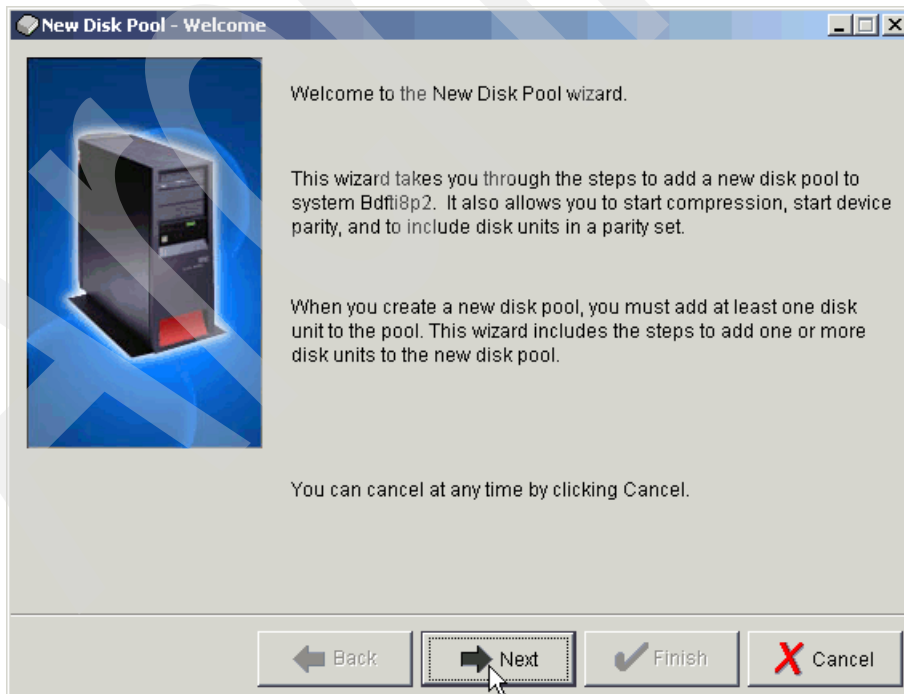


Figure 5-26 New Disk Pool - Welcome window

6. In the New Disk Pool window (Figure 5-27):
  - a. For Type of disk pool, select **Primary**.
  - b. For Disk pool, type the new disk pool name.
  - c. Leave Database set to the default of **Generated by the system**.
  - d. Ensure that the disk protection method matches the type of logical volume that you are adding. If you leave it deselected, you will see all available disks.
  - e. Select **OK** to continue.

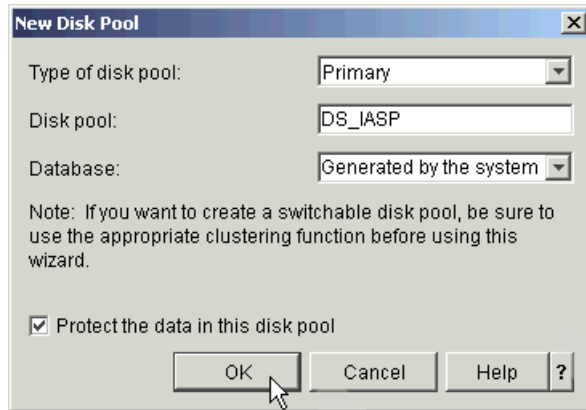


Figure 5-27 Defining a new disk pool

7. The New Disk Pool - Select Disk Pool window (Figure 5-28) summarizes the disk pool configuration. Review the configuration and click **Next**.

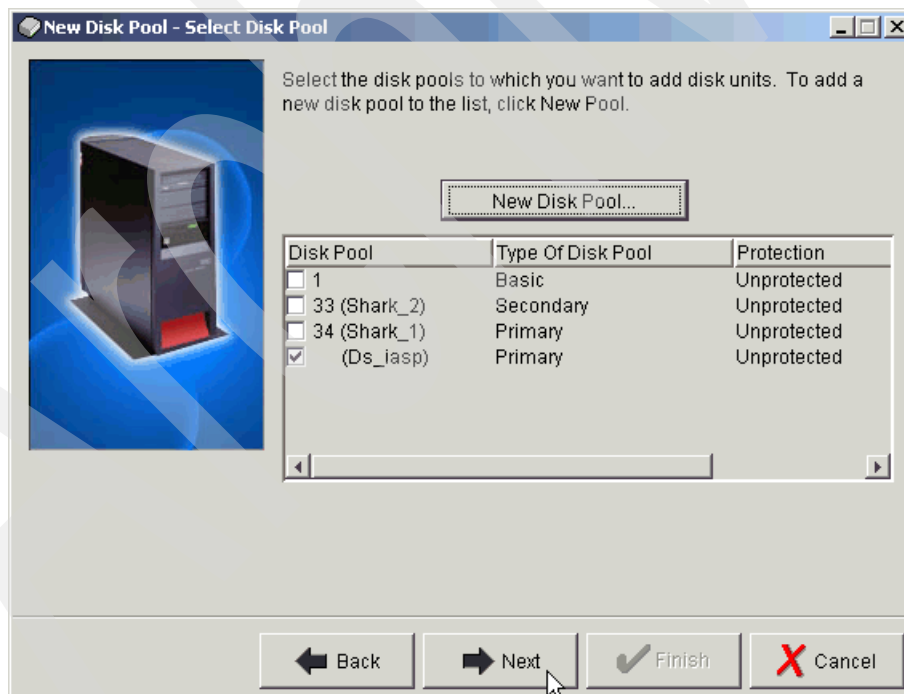


Figure 5-28 Confirming the disk pool configuration

- In the New Disk Pool - Add to Disk Pool window (Figure 5-29), click **Add Disks** to add disks to the new disk pool.

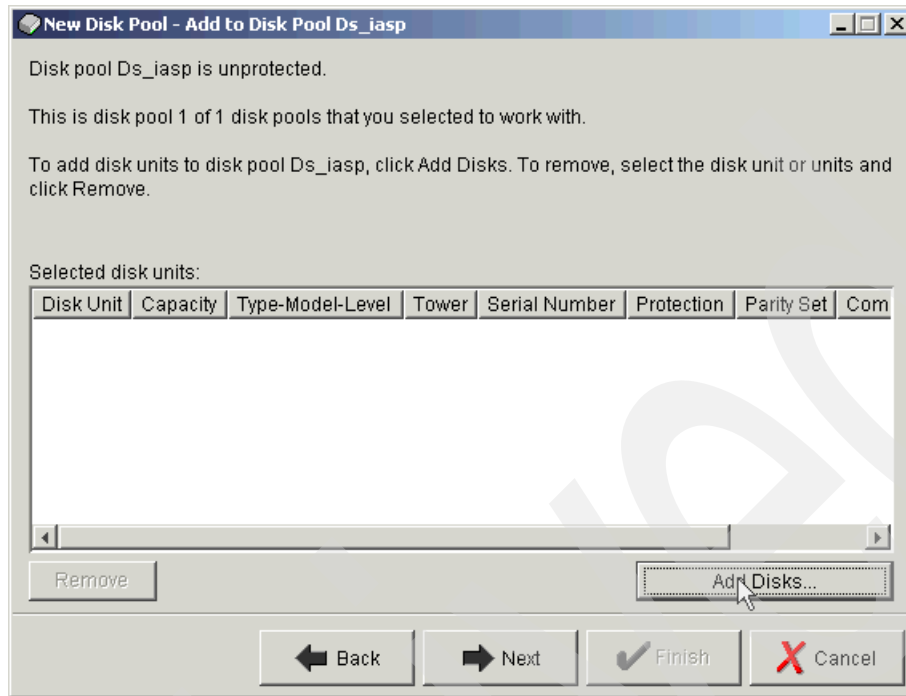


Figure 5-29 Adding disks to the disk pool

- The Disk Pool - Add Disks window lists the non-configured units. Highlight the disk or disks that you want to add to the disk pool, and click **Add**, as shown in Figure 5-30.

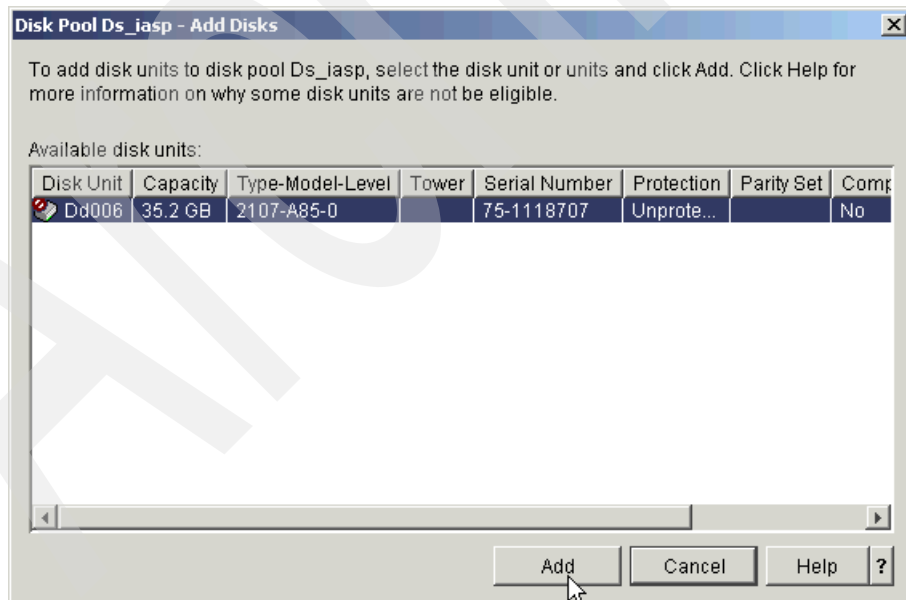


Figure 5-30 Choosing the disks to add to the disk pool

10. The next window confirms the selection (see Figure 5-31). Click **Next** to continue.




Figure 5-31 Confirming the disks to be added to the disk pool

11. In the New Disk Pool - Summary window, review the summary of the configuration. Click **Finish** to add the disks to the disk pool, as shown in Figure 5-32.

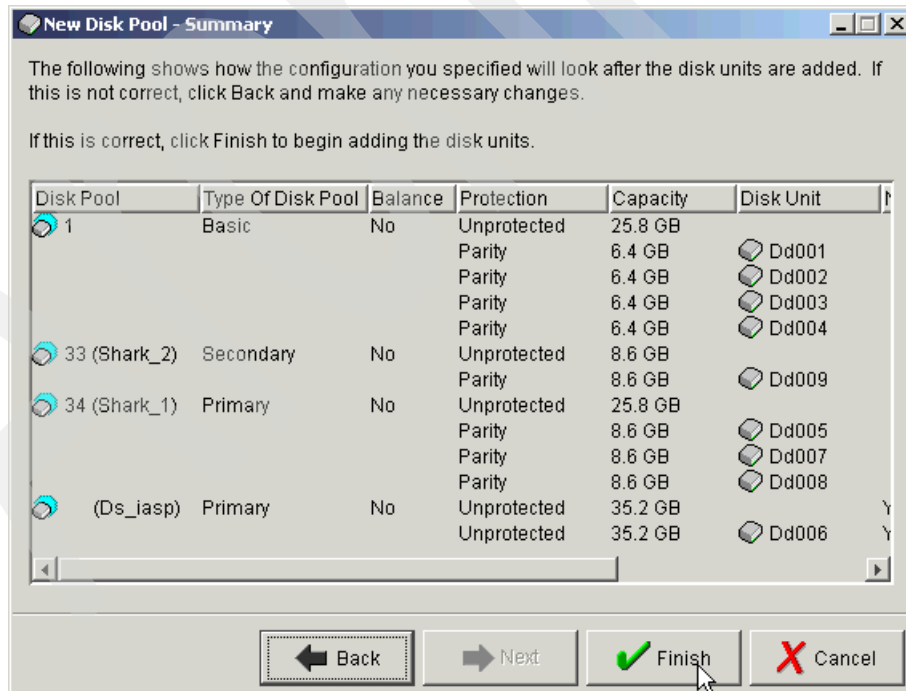


Figure 5-32 New Disk Pool - Summary window

12. Take note of and respond to any messages that display. After you take any necessary action regarding any messages, you see the New Disk Pool Status window (Figure 5-33), which shows the progress. This step might take some time, depending on the number and size of the logical units that are being added.




Figure 5-33 New Disk Pool Status

13. When the process is complete, a message window displays. Click **OK** as shown in Figure 5-34.




Figure 5-34 Disks added successfully to the disk pool

14. In iSeries Navigator, you can see the new disk pool under Disk Pools (see Figure 5-35).

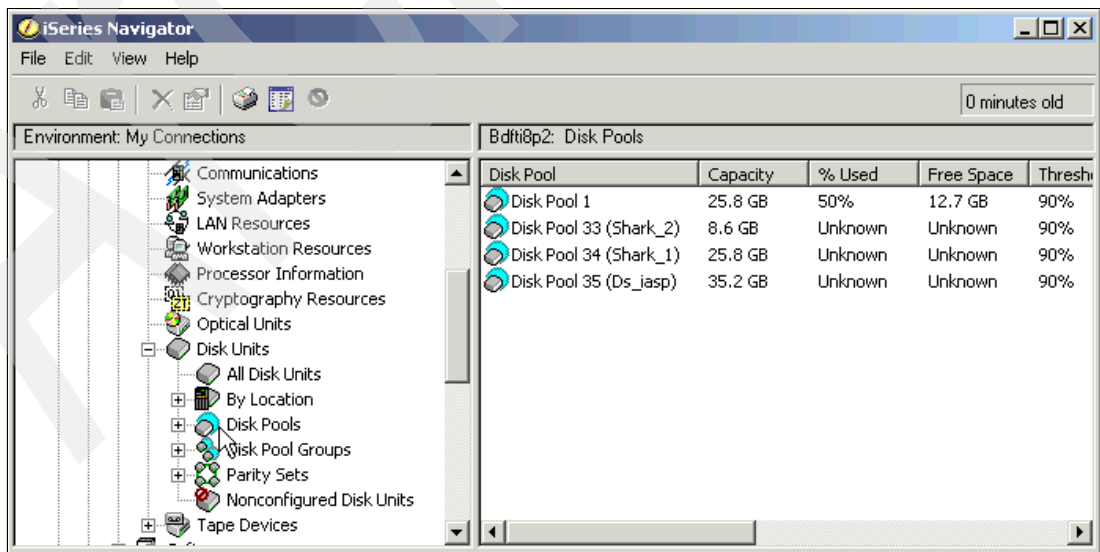


Figure 5-35 New disk pool shown in iSeries Navigator

15. To see the logical volume, expand **Configuration and Service** → **Hardware** → **Disk Pools** and select the disk pool that you just created. See Figure 5-36.




Figure 5-36 New logical volume in iSeries Navigator

## 5.6 Adding multipath volumes to System i using a 5250 interface

If you are using the 5250 interface, sign on to SST and perform the following steps:

1. On the first panel, select **3. Work with disk units**.
2. On the next panel, select **2. Work with disk configuration**.
3. On the next panel, select **8. Add units to ASPs and balance data**.
4. In the Specify ASPs to Add Units to panel (Figure 5-37), the values in the Resource Name column show DDxxx for single path volumes and DMPxxx for those which have more than one path. In this example, the 2107-A85 logical volume with serial number 75-1118707 is available through more than one path and reports in as DMP135.

Specify the ASP to which you want to add the multipath volumes.

**Note:** For multipath volumes, only one path is shown. For the additional paths, see 5.8, “Managing multipath volumes using iSeries Navigator” on page 211.




Figure 5-37 Adding multipath volumes to an ASP

- On the Confirm Add Units panel (Figure 5-38), check the configuration details. If the details are correct, press Enter.




Figure 5-38 Confirm Add Units panel

## 5.7 Adding volumes to System i using iSeries Navigator

You can use iSeries Navigator to add volumes to the system ASP, user ASPs, or IASPs. In this example, we add a multipath logical volume to a private (non-switchable) IASP. The same principles apply when adding multipath volumes to the system ASP or user ASPs.

1. Follow the steps in 5.5.2, “Adding volumes to an independent auxiliary storage pool” on page 199. When you reach the point where you select the volumes to add, a panel similar to the panel that is shown in Figure 5-39 displays. Multipath volumes appear as *DMPxxx*. Highlight the disk or disks that you want to add to the disk pool and click **Add**.

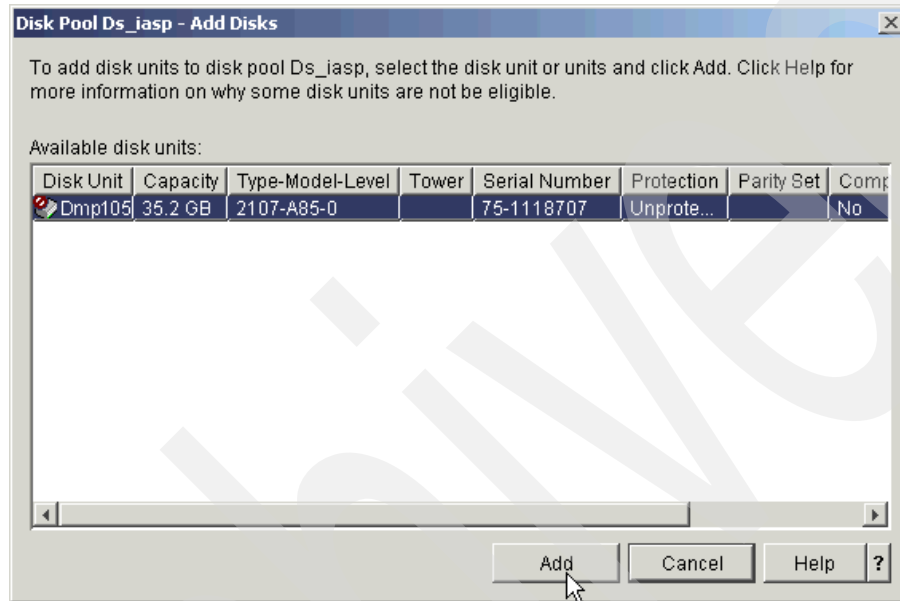


Figure 5-39 Adding a multipath volume

**Note:** For multipath volumes, only one path is shown. To see the additional paths, see 5.8, “Managing multipath volumes using iSeries Navigator” on page 211.



- The remaining steps are identical to those in 5.5.2, “Adding volumes to an independent auxiliary storage pool” on page 199.  
When you have completed the steps, you can see the new disk pool in iSeries Navigator under Disk Pools (see Figure 5-40).

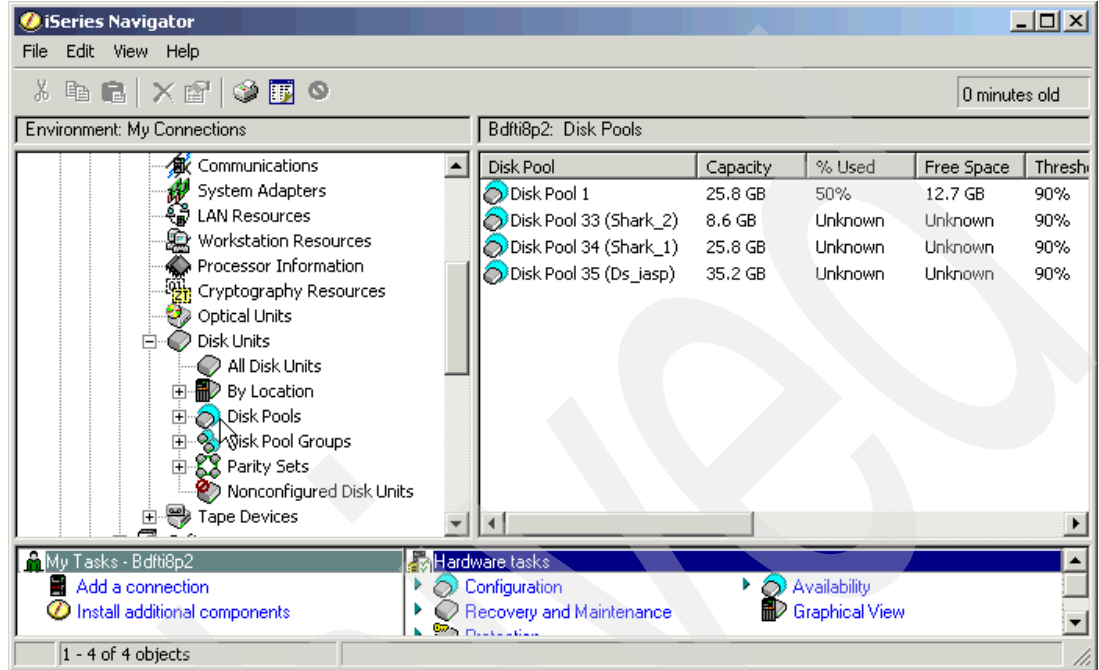


Figure 5-40 New disk pool in iSeries Navigator

3. To see the logical volume, expand **Configuration and Service** → **Hardware** → **Disk Units** → **Disk Pools**, and click the disk pool that you just created as shown in Figure 5-41.

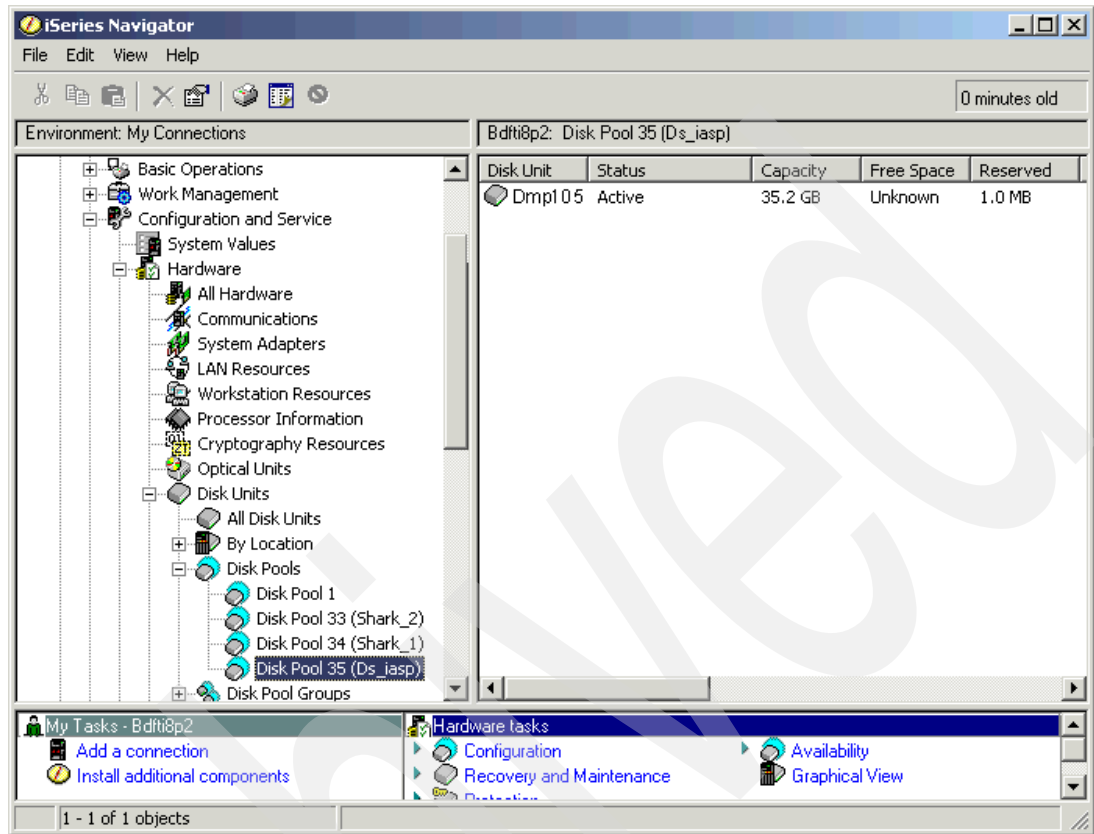


Figure 5-41 New logical volume shown in iSeries Navigator

## 5.8 Managing multipath volumes using iSeries Navigator

All units are initially created with a prefix of DD. As soon as the system detects that there is more than one path to a specific logical unit, it automatically assigns a unique resource name with a prefix of DMP for both the initial path and any additional paths.

When using the standard disk panels in iSeries Navigator, only a single path, the initial path, is shown. To see the additional paths follow these steps:

1. To see the number of paths available for a logical unit, open iSeries Navigator and expand **Configuration and Service** → **Hardware** → **Disk Units**. As shown in Figure 5-42, the number of paths for each unit is in the *Number of Connections* column (far right side of the panel). In this example, there are eight connections for each of the multipath units.




Figure 5-42 Example of multipath logical units

- To see the other connections to a logical unit, right-click a unit, and select **Properties**, as shown in Figure 5-43.

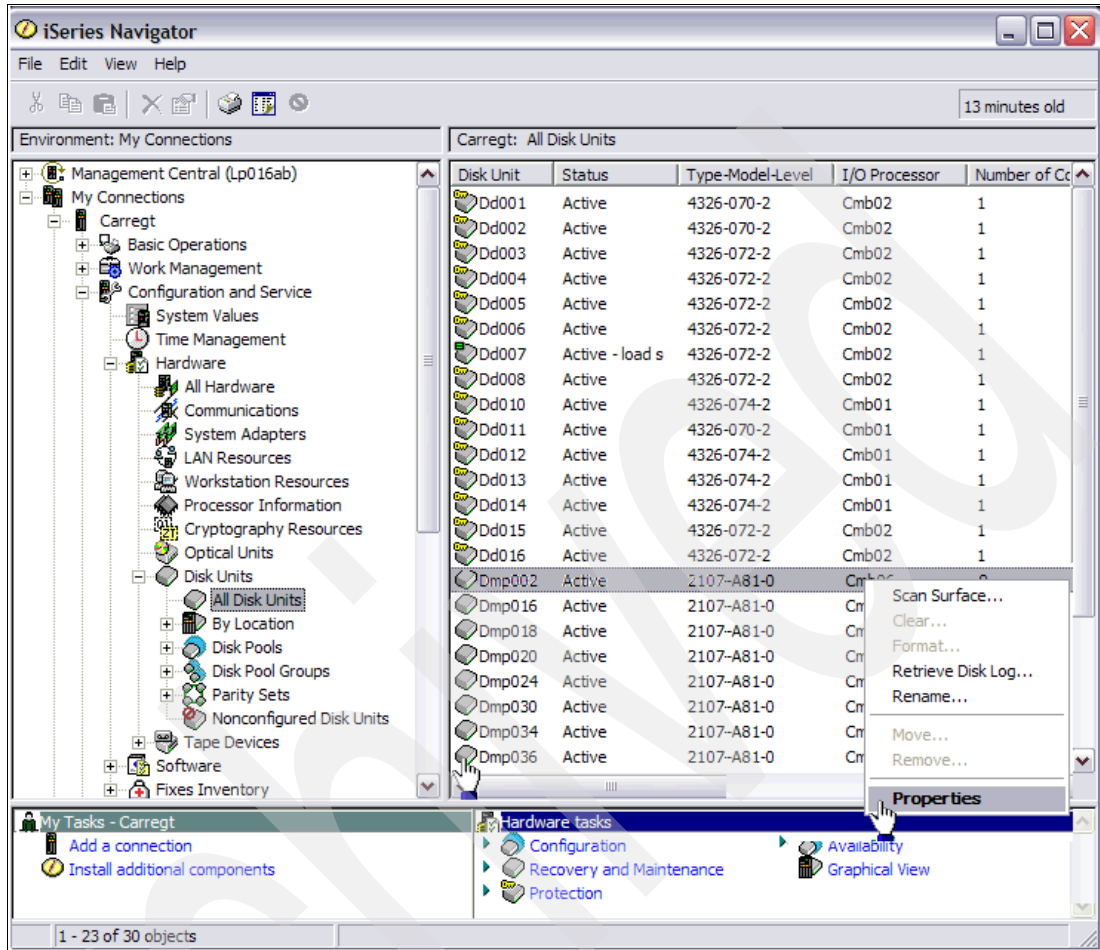


Figure 5-43 Selecting properties for a multipath logical unit

3. In the Properties window (Figure 5-44), you see the General tab for the selected unit. The first path is shown as Device 1 in the Storage section of the dialog box.



Figure 5-44 Multipath logical unit properties

To see the other paths to this unit, click the **Connections** tab, where the other seven connections for this logical unit are displayed, as shown in Figure 5-45.

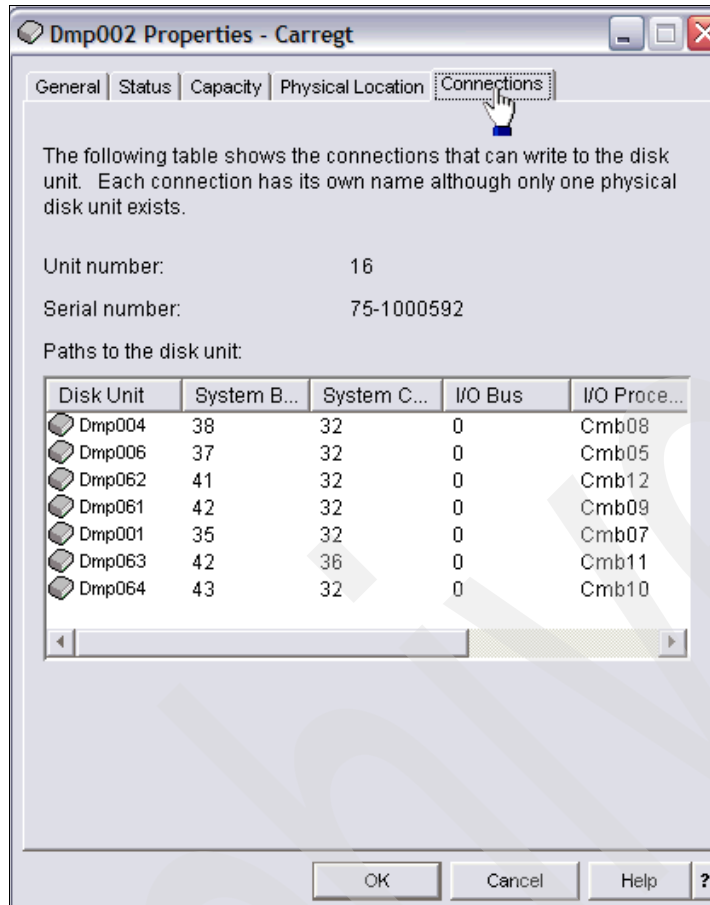


Figure 5-45 Multipath connections

## 5.9 Changing from single path to multipath

If you have an existing configuration where the logical units were assigned only to one Fibre Channel I/O adapter, you can change to multipath easily. Simply assign the logical units in the DS8000 or DS6000 system to another System i I/O adapter. Then the existing DDxxx resource names change automatically to DMPxxx, and new DMPyyy resources are created for the new path.

Figure 5-46 shows an example where 48 logical volumes are configured in the DS8000. The first 24 of these being in one DS volume group are assigned using a host adapter in the top left I/O drawer in the DS8000 to a Fibre Channel (FC) I/O adapter in the first iSeries I/O tower or rack. The next 24 logical volumes within another DS volume group are assigned using a host adapter in the lower left I/O drawer in the DS8000 to an FC I/O adapter on a different bus in the first iSeries I/O tower or rack. This is a valid single path configuration.

To implement multipath, the first group of 24 logical volumes is also assigned to an iSeries FC I/O adapter in the second iSeries I/O tower or rack through a host adapter in the lower right I/O drawer in the DS8000. The second group of 24 logical volumes is also assigned to an FC I/O adapter on a different bus in the second iSeries I/O tower or rack through a host adapter in the upper right I/O drawer.

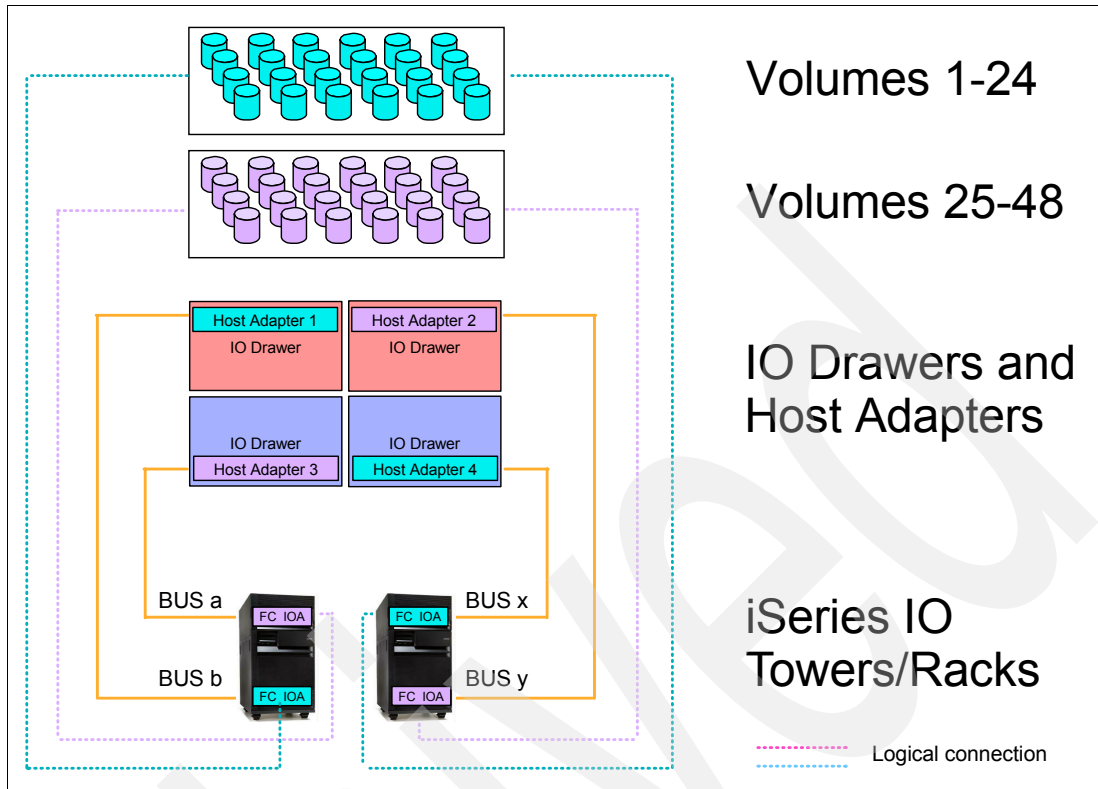


Figure 5-46 Example of multipath with the iSeries server

## 5.10 Protecting the external load source unit

With i5/OS V6R1 multipath is now also supported for the SAN load source unit for both #2847 IOP-based and #5749 or #5774 IOP-less Fibre Channel adapters. Therewith the load source unit data is not only data protected within the external storage unit, either by RAID-5 or RAID-10, but also protected against I/O path failures. Implementation for i5/OS LUN multipathing is achieved simply by configuring logical volumes from storage side to at least two System i Fibre Channel I/O adapters as discussed in 5.9, “Changing from single path to multipath” on page 214.

**Note:** For the remainder of this section, we focus on implementing load source mirroring for an #2847 IOP-based SAN load source prior to i5/OS V6R1.

Prior to i5/OS V6R1, the #2847 PCI-X IOP for SAN load source did not support multipath for the external load source unit. To provide path protection for the external load source unit prior to V6R1 it has to be mirrored using i5/OS mirroring. Therefore, the two LUNs used for mirroring the external load source across two #2847 IOP-based Fibre Channel adapters (ideally in different I/O towers to provide highly redundant path protection) are created as unprotected LUN models.

To mirror the load source unit, unless you are using SLIC V5R4M5 or later (see 5.4, “Setting up an external load source unit” on page 185) initially assign only one LUN to the IOA that is tagged as the load source unit IOA. Other LUNs, including the “mirror mate” for the load source unit, should be assigned to another #2847 IOP-based IOA as shown in Figure 5-47. The simplest way to do this is to create two volume groups on the DS8000 or DS6000. The first volume group (shown on the left) contains only the load source unit and is assigned to the #2847 tagged as the load source IOA. The second volume group (shown on the right) contains the load source unit mirror mate plus the remaining LUNs, which eventually will have multipaths. This volume group is assigned to the second #2847 IOP-based IOA.

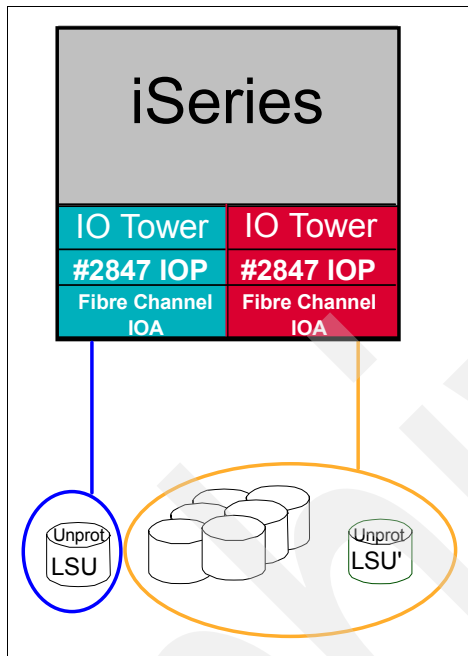


Figure 5-47 Initial LUN allocation



After you have loaded SLIC onto the load source unit, you can assign the remaining LUNs to the second #2847 IOP-based IOA to provide multipath as shown in Figure 5-48 by assigning those LUNs that will have multipaths to the volume group on the left.




Figure 5-48 Final LUN allocation

If you have more LUNs that require more IOPs and IOAs, you can assign these to volume groups with already using a multipath configuration as shown in Figure 5-49. It is important to ensure that your load source unit initially is the only volume assigned to the #2487 IOP-based

IOA that is tagged in Hardware Management Console (HMC) as the load source IOA. Our example including SAN switches shows a configuration with two redundant SAN switches to avoid a single-point of failure.

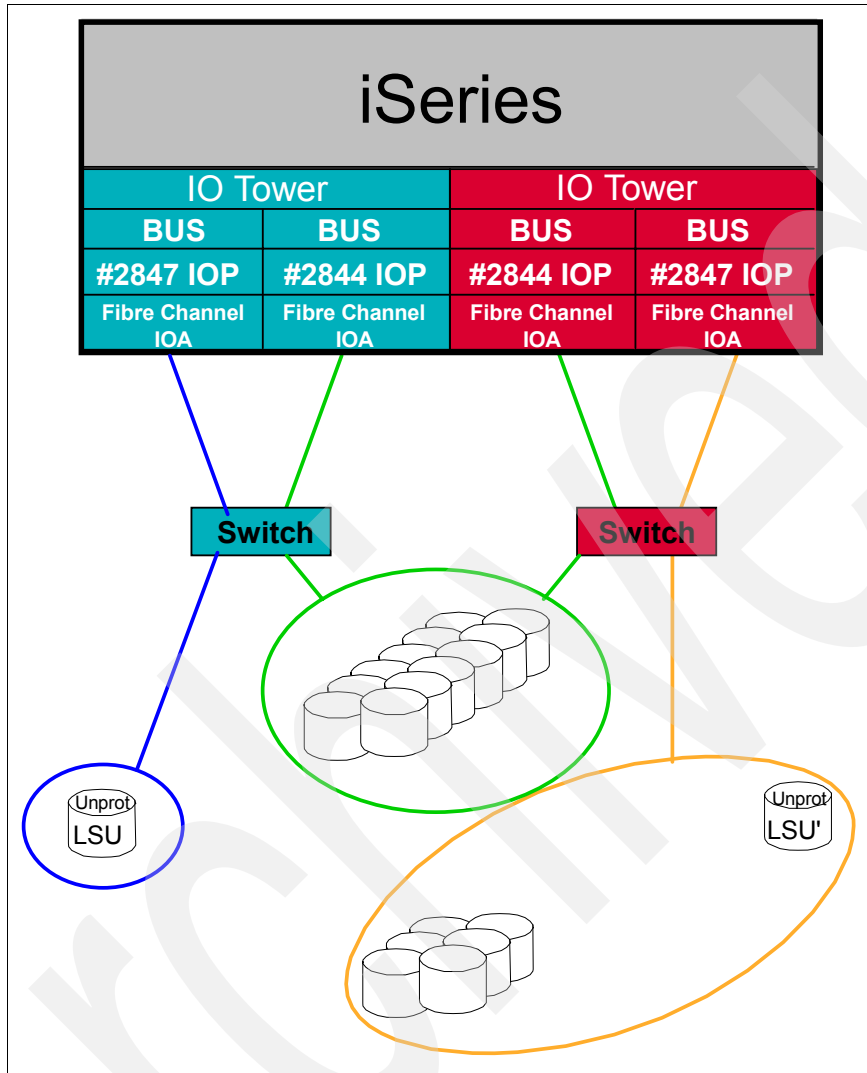


Figure 5-49 Initial LUN allocation with additional multipath LUNs

After SLIC is loaded on the load source unit, you can assign the multipath LUNs to the #2847 tagged as the load source unit by adding them to the volume group (on the left in Figure 5-50), which initially only contained the load source unit.

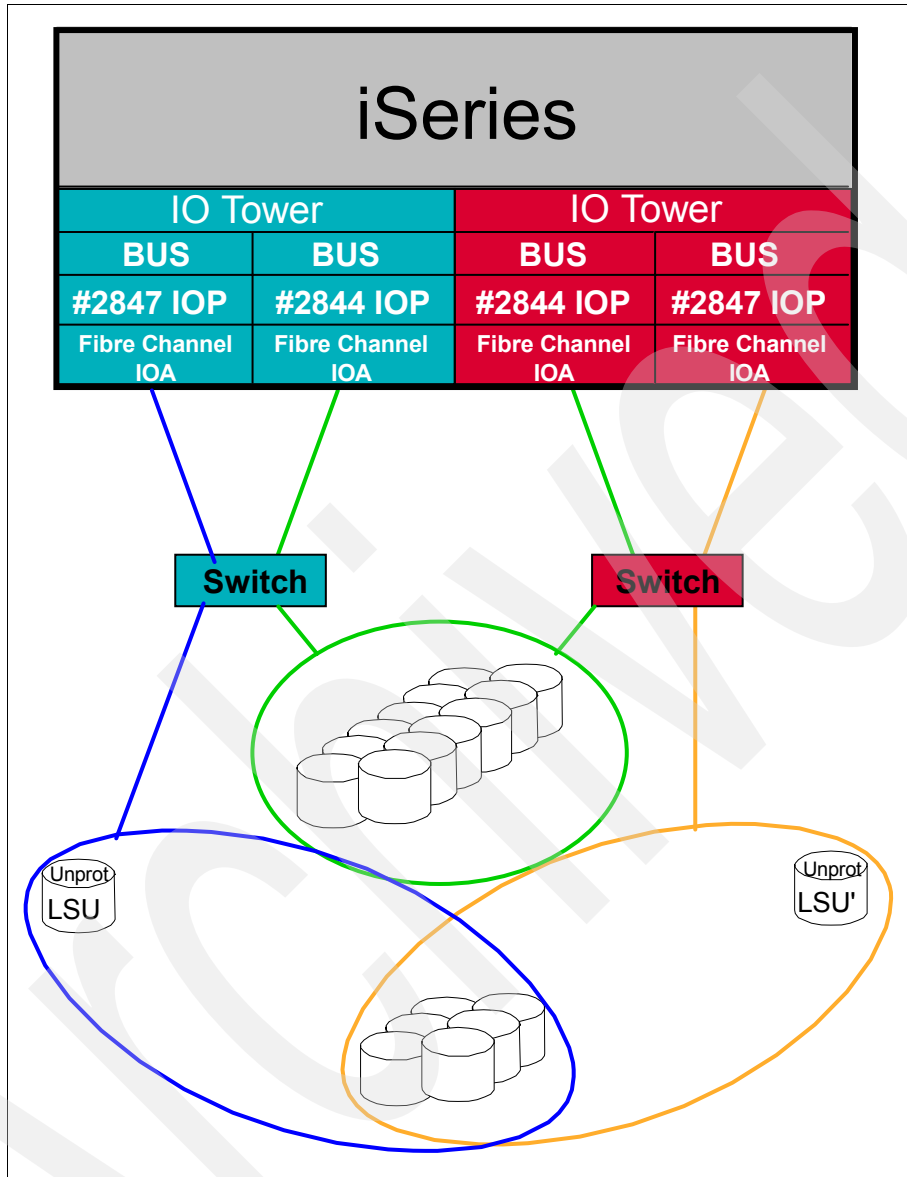


Figure 5-50 Final LUN allocation with additional multipath LUNs

## 5.10.1 Setting up load source mirroring

After you create the LUN to be set up as the remote load source unit pair, this LUN and any other LUNs are identified by SLIC and displayed under non-configured units in DST and SST. To set up load source mirroring on the System i5 platform, you must use DST:

1. From the DST menu (Figure 5-51), select **4. Work with disk units**.

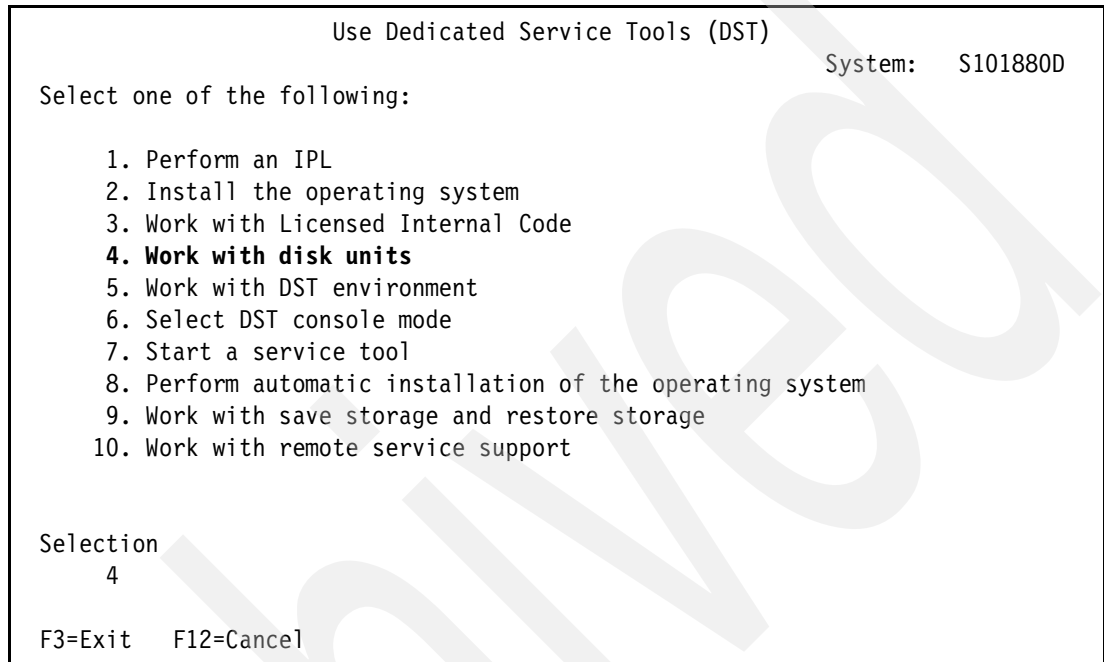


Figure 5-51 Using Dedicated Service Tools panel

2. From the Work with Disk Units menu (Figure 5-52), select **1. Work with disk configuration**.

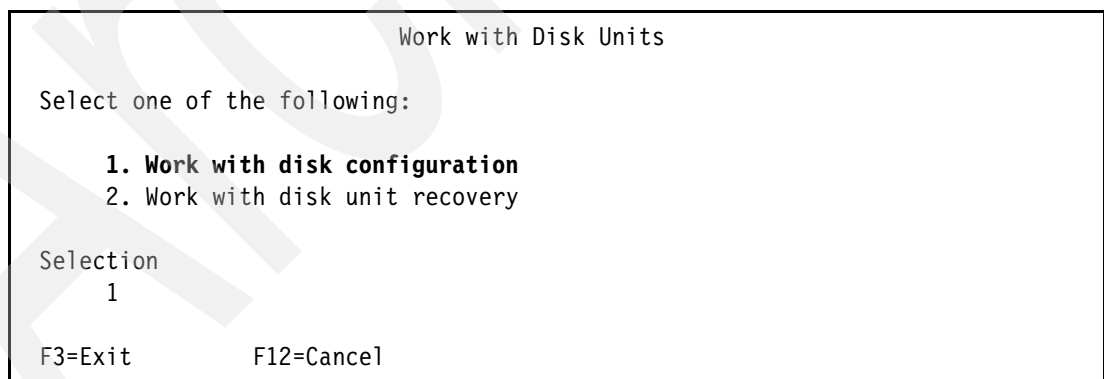


Figure 5-52 Working with Disk Units panel

3. From the Work with Disk Configuration menu (Figure 5-53), select **4. Work with mirrored protection**.

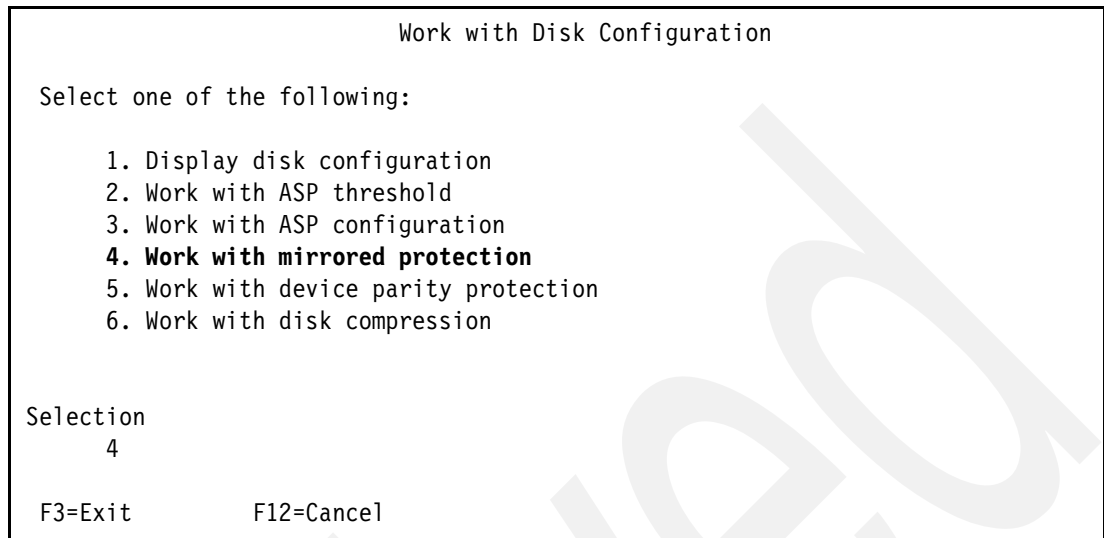


Figure 5-53 Work with Disk Configuration panel

4. From the Work with mirrored protection menu (Figure 5-54), select **4. Enable remote load source mirroring**. This option does not perform the remote load source mirroring but tells the system that you want to mirror the load source when mirroring is started.

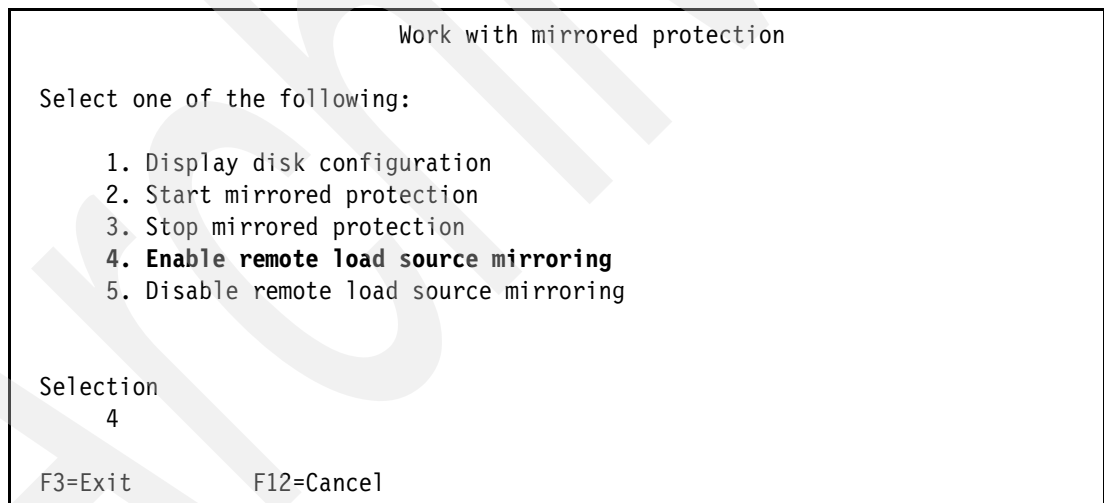


Figure 5-54 Setting up remote load source mirroring

5. In the Enable Remote Load Source Mirroring confirmation panel (Figure 5-55), press Enter to confirm that you want to enable remote load source mirroring.

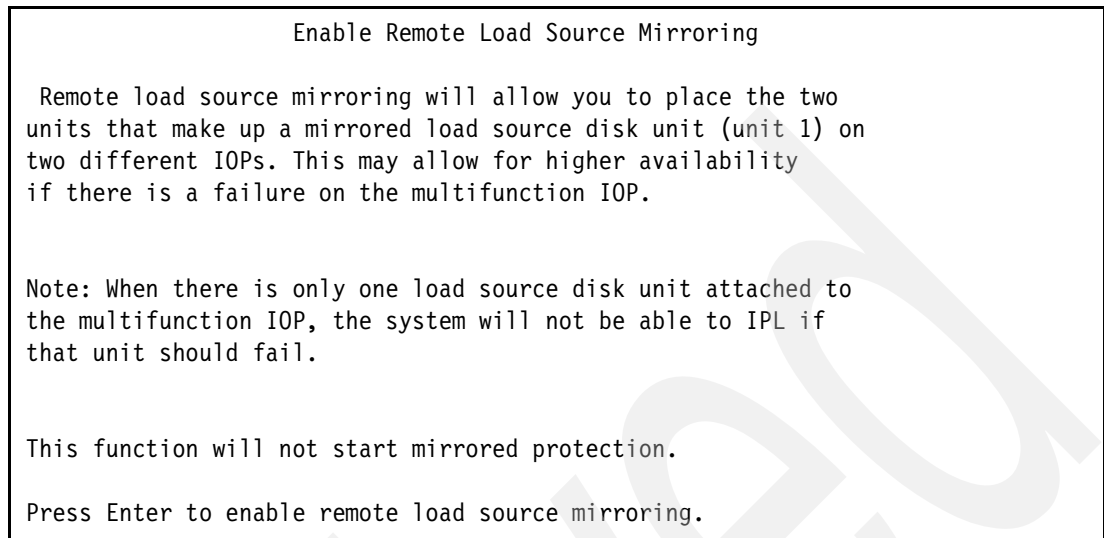


Figure 5-55 Enable Remote Load Source Mirroring panel

6. In the Work with mirrored protection panel, you see a message at the bottom of the panel, indicating that remote load source mirroring is enabled (Figure 5-56). Select **2. Start mirrored protection**, for the load source unit.

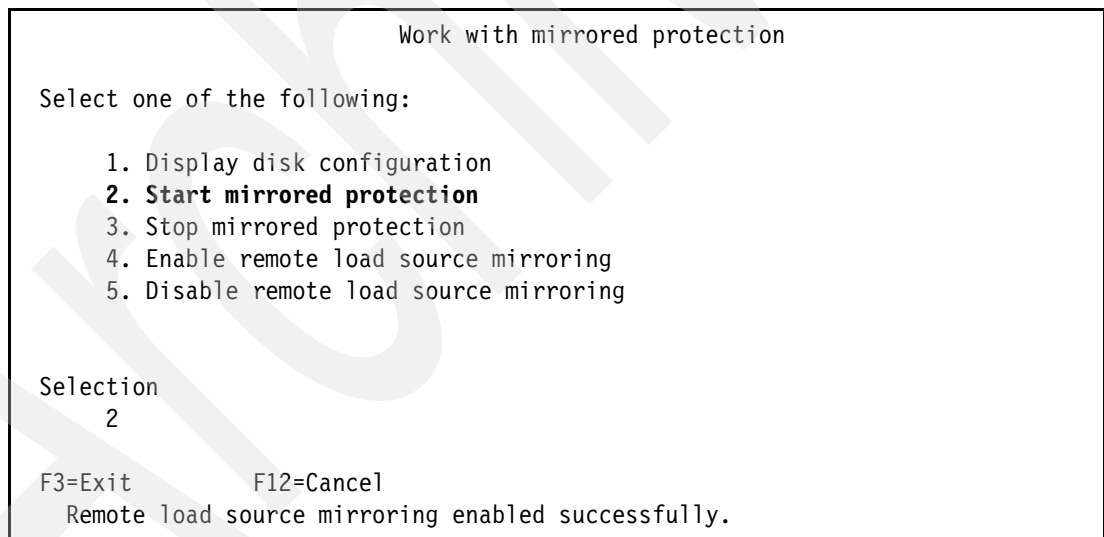


Figure 5-56 Confirmation that remote load source mirroring is enabled

7. In the Work with mirrored protection menu, select **1. Display disk configuration**, and then select **1. Display disk configuration status**.

Figure 5-57 shows the two unprotected LUNs (model A85) for the load source unit and its mirror mate as disk serial number 30-1000000 and 30-1100000. You can also see that there are four more protected LUNs (model A05) that are protected by multipath because their resource names begin with *DMP*.

Display Disk Configuration Status						
ASP	Unit	Serial Number	Type	Model	Resource Name	Status
1						Mirrored
	1	30-1000000	1750	A85	DD001	Active
	1	30-1100000	1750	A85	DD004	Active
	2	30-1001000	1750	A05	DMP002	DPY/Active
	3	30-1002000	1750	A05	DMP004	DPY/Active
	5	30-1101000	1750	A05	DMP006	DPY/Active
	6	30-1102000	1750	A05	DMP008	DPY/Active

Press Enter to continue.

F3=Exit                    F5=Refresh                    F9=Display disk unit details  
 F11=Disk configuration capacity    F12=Cancel

Figure 5-57 Unprotected load source unit ready to start remote load source mirroring

8. When the remote load source mirroring task is finished, perform an IPL on the system to start mirroring the data from the source unit to the target. This process is done during the database recovery phase of the IPL.

## 5.11 Migration from mirrored to multipath load source

With the new i5/OS V6R1 release System i supports multipath to the load source LUN for both 2847 IOP-based or IOP-less IOAs.

**Note:** This migration procedure is a *disruptive* procedure because it involves stopping mirrored protection and optionally changing the LUN protection mode for the load source unit.

To migrate from a mirrored external load source unit to a multipath load source unit, follow these steps:

1. Enter STRSST to start System Service Tools from the i5/OS command line.
2. Select **3. Work with disk units**.
3. Select **2. Work with disk configuration**.
4. Select **1. Display disk configuration**.

5. Select **1. Display disk configuration status** to look at your currently mirrored external load source LUNs. Take note of the two serial numbers for your mirrored load source unit 1 (105E951 and 1060951 in our example) because you will need these numbers later for changing the DS storage system configuration to a multipath setup.
6. Press F12 to exit from the Display Disk Configuration Status screen, as shown in Figure 5-58.

Display Disk Configuration Status					
ASP Unit	Serial Number	Type	Model	Resource Name	Status
1	50-105E951	2107	A85	DD001	Active
1	50-1060951	2107	A85	DD002	Active
2	50-1061951	2107	A05	DMP003	RAID 5/Active
3	50-105F951	2107	A05	DMP001	RAID 5/Active

Figure 5-58 Displaying mirrored disks

7. Select **6. Disable remote load source mirroring** to turn off the remote load source mirroring function as shown in Figure 5-59.

**Note:** Turning off the remote load mirroring function does not stop the mirrored protection. However, disabling this function is required to actually allow stop mirroring in a later step.

Work with Disk Configuration	
Select one of the following:	
1.	Display disk configuration
2.	Add units to ASPs
3.	Work with ASP threshold
4.	Add units to ASPs and balance data
5.	Enable remote load source mirroring
6.	<b>Disable remote load source mirroring</b>
7.	Start compression on non-configured units
8.	Work with device parity protection
9.	Start hot spare
10.	Stop hot spare
Selection	6
F3=Exit	F12=Cancel

Figure 5-59 Disable remote load source mirroring



8. Press Enter to confirm your action in the Disable Remote Load Source Mirroring panel, as shown in Figure 5-60.

```
Disable Remote Load Source Mirroring

Remote load source mirroring is currently enabled.  You
selected to turn this function off.  This may require that both
units that make up your mirrored load source disk unit (unit 1)
be attached to the same IOP.

This function will not stop mirrored protection.

Press Enter to disable remote load source mirroring.

F3=Exit          F12=Cancel
```

Figure 5-60 Disable Remote Load Source Mirroring confirmation screen

9. A completion message displays, as shown in Figure 5-61.

```
Work with Disk Configuration

Select one of the following:

1. Display disk configuration
2. Add units to ASPs
3. Work with ASP threshold
4. Add units to ASPs and balance data
5. Enable remote load source mirroring
6. Disable remote load source mirroring
7. Start compression on non-configured units
8. Work with device parity protection
9. Start hot spare
10. Stop hot spare

Selection

F3=Exit          F12=Cancel
Remote load source mirroring disabled successfully.
```

Figure 5-61 Message after disabling the remote load source mirroring

10. To stop mirror protection, set your system to B-type manual mode IPL, and re-IPL the system. When you get to the Dedicated Service Tools (DST) panel, continue with these steps.

11. Select **4. Work with disk units** as shown in Figure 5-62.

```
Use Dedicated Service Tools (DST)
System: RCHLTTN1

Select one of the following:

1. Perform an IPL
2. Install the operating system
3. Work with Licensed Internal Code
4. Work with disk units
5. Work with DST environment
6. Select DST console mode
7. Start a service tool
8. Perform automatic installation of the operating system
9. Work with save storage and restore storage
10. Work with remote service support

12. Work with system capacity
13. Work with system security
14. End batch restricted state

Selection
4

F3=Exit F12=Cancel
```

Figure 5-62 Work with disk units

12. Select **1. Work with disk configuration** as shown Figure 5-63.

```
Work with Disk Units

Select one of the following:

1. Work with disk configuration
2. Work with disk unit recovery

Selection
1

F3=Exit F12=Cancel
```

Figure 5-63 Work with disk units

13. Select **4. Work with mirrored protection** as shown in Figure 5-64.

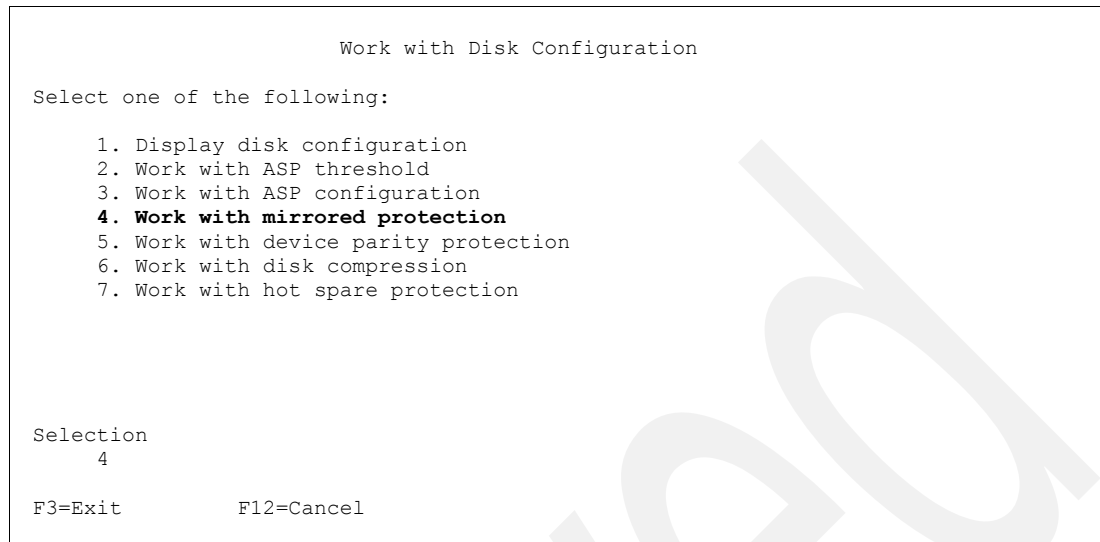


Figure 5-64 Work with mirrored protection

14. Select **3. Stop mirrored protection** as shown in Figure 5-65.




Figure 5-65 Stop mirrored protection

15. Enter 1 to select ASP 1, as shown in Figure 5-66.

```
                Select ASP to Stop Mirrored Protection

Select the ASPs to stop mirrored protection on.

Type options, press Enter.
  1=Select

Option  ASP  Protection
  1      1  Mirrored

F3=Exit          F12=Cancel
```

Figure 5-66 Selecting ASP to stop mirror

16. On the Confirm Stop Mirrored Protection panel, confirm that ASP 1 is selected, as shown in Figure 5-67, and then press Enter to proceed.

```
                Confirm Stop Mirrored Protection

Press Enter to confirm your choice to stop mirrored
protection. During this process the system will be IPLed.
You will return to the DST main menu after the IPL is
complete. The system will have the displayed protection.

Press F12 to return to change your choice.

ASP Unit  Serial          Resource
          Number         Type Model Name      Protection
  1
    1  50-105E951        2107  A85  DD001      Unprotected
    2  50-1061951        2107  A05  DMP003      RAID 5
    3  50-105F951        2107  A05  DMP001      RAID 5
```

Figure 5-67 Confirm to stop mirrored protection

17. When the stop for mirrored protection completes, a confirmation panel displays as shown in Figure 5-68.

```
                Disk Configuration Information Report

The following are informational messages about disk
configuration changes started in the previous IPL.

Information

Stop mirroring completed successfully

Press Enter to continue
```

Figure 5-68 Successful message to stop mirroring

18. The previously mirrored load source is now a *non-configured disk* unit, as highlighted in Figure 5-69.

Display Non-Configured Units					
Serial Number	Type	Model	Resource Name	Capacity	Status
50-1060951	2107	A85	DD002	35165	Non-configured

Press Enter to continue.

F3=Exit                      F5=Refresh                      F9=Display disk unit details  
F11=Display device parity status                      F12=Cancel

Figure 5-69 Non-configured disk

19. Now, you can exit from the DST panels to continue the manual mode IPL. At the Add All Disk Units to the System panel, select **1. Perform any disk configuration at SST** as shown in Figure 5-70.

Add All Disk Units to the System	
	System: RCHLTTN1
Non-configured device parity capable disk units are attached to the system. Disk units can not be added automatically. It is more efficient to device parity protect these units before adding them to the system. These disk units may be parity enabled and added at SST. Configured disk units must have parity enabled at DST.	
Select one of the following:	
1. Perform any disk configuration at SST	
2. Perform disk configuration using DST	
Selection	1

Figure 5-70 Message to add disks

20. You have stopped mirrored protection for the load source unit and re-IPLed the system successfully. Now, use the DS CLI to identify the volume groups that contain the two LUNs of your previously mirrored load source unit by entering the `showfbvol volumeID` command for the previously mirrored load source unit (for `volumeID` use the four digits from the disk unit serial number noted down in step 5) as shown in Figure 5-71.

```

dscli> showfbvol 1060
Date/Time: 9. November 2007 02:33:02 CET IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name          TN1mm
ID            1060
accstate      Online
datastate     Normal
configstate   Normal
deviceMTM     2107-A05
datatype      FB 520P
addrgrp       1
extpool       P4
exts          33
captype       iSeries
cap (2^30B)   32.8
cap (10^9B)   35.2
cap (blocks)  68681728
volgrp        V22
ranks         1
dbexts        -
sam           Standard
repcapalloc   -
eam           legacy
reqcap (blocks) 68681728

```

Figure 5-71 DS CLI: The `showfbvol` command

21. Enter `showvolgrp volumegroup_ID` for the two volume groups that contain the previously mirrored load source unit LUNs, as shown in Figure 5-72 and Figure 5-73.

```

dscli> showvolgrp v13
Date/Time: November 7, 2007 3:31:51 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name RedBookTN1LS_VG
ID V13
Type OS400 Mask
Vols 105E 105F 1061

```

Figure 5-72 DS CLI: The `showvolgroup` command

```

dscli> showvolgrp v22
Date/Time: November 7, 2007 3:31:54 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name RedBookTN1MM_VG
ID V22
Type OS400 Mask
Vols 105F 1060 1061

```

Figure 5-73 DS CLI: The `showvolgroup` command

22. To start using multipath for *all* volumes, including the load-source attached to the IOAs, add the previous load source mirror volume that has become the non-configured unit into the volume group of the load source IOA, as shown in Figure 5-74. At this point in the process, you have established two paths to the non-configured previous load source mirror LUN.

```

dscli> chvolgrp -action add -volume 1060 V13
Date/Time: November 7, 2007 3:36:34 AM IST IBM DSCLI Version: 5.3.0.991 DS:
IBM.2107-7589951
CMUC00031I chvolgrp: Volume group V13 successfully modified.
  
```

Figure 5-74 Adding a volume into a volume group

23. To finish the multipath setup, make sure that the current load source unit LUN (*LUN 105E* in our example) is also assigned to both System i IOAs. You assign the load source unit LUN to the second IOA by assigning the volume group (*V13* in our example) that now contains both previously mirrored load source unit LUNs to both IOAs. To obtain the IOAs host connection ID on the DS storage system for changing the volume group assignment, enter the **lshostconnect** command as shown in Figure 5-75. Note the ID for the lines that show the two load source IOA volume groups determined previously.

```

dscli> lshostconnect
Date/Time: November 7, 2007 3:30:36 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name ID WWPN HostType Profile portgrp volgrpID ESSIOpo
-----
RedBookTN1LS 0010 10000000C94C45CE iSeries IBM iSeries - OS/400 0 V13 all
RedBookTN1MM 001B 10000000C9509E12 iSeries IBM iSeries - OS/400 0 V22 all
  
```

Figure 5-75 DS SLI: The lshostconnect command

24. Change the volume group assignment of the IOA host connection that does not yet have access to the current load source. (In our example, volume group *V22* does not contain the current load source unit LUN, so we have to assign volume group *V13* that contains both previous load source units to host connection *001B*.) Use the **chhostconnect -volgrp volumegroupID hostconnecID** command as shown in Figure 5-76.

```

dscli> chhostconnect -volgrp V13 001B
Date/Time: November 7, 2007 3:40:25 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
CMUC00013I chhostconnect: Host connection 001B successfully modified.
dscli> lshostconnect
Date/Time: November 7, 2007 3:40:33 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name ID WWPN HostType Profile portgrp volgrpID ESSIOport
-----
RedBookTN1LS 0010 10000000C94C45CE iSeries IBM iSeries - OS/400 0 V13 all
RedBookTN1MM 001B 10000000C9509E12 iSeries IBM iSeries - OS/400 0 V13 all
  
```

Figure 5-76 DS CLI: Change the host connection volume group assignment

Now, we describe how to change two previously mirrored unprotected disk units to protected ones.

**Important:** It is not supported to change the LUN protection status of a LUN that is being configured, that is a LUN that is part of an ASP configuration. To convert the unprotected load source disk unit to a protected model follow, steps 12 to 18 in the process that follows.

Follow these steps:

1. Display the unprotected disk units by selecting **Display disk configuration status** and **Display non-configured disks** on System i SST or DST as shown in Figure 5-77.

```

Display Disk Configuration Status

ASP Unit      Serial          Type Model Name      Resource      Status
  1           50-105E951     2107 A85 DMP007           Configured
  2           50-1061951     2107 A05 DMP003           RAID 5/Active
  3           50-105F951     2107 A05 DMP001           RAID 5/Active

Press Enter to continue.

F3=Exit      F5=Refresh      F9=Display disk unit details
F11=Disk configuration capacity  F12=Cancel

Display Non-Configured Units

Serial          Type Model Name      Capacity Status
50-1060951     2107 A85 DMP005           35165 Non-configured

Press Enter to continue.

F3=Exit      F5=Refresh      F9=Display disk unit details
F11=Display device parity status  F12=Cancel
  
```

Figure 5-77 Displaying unprotected disks

2. On the storage system, use the DS CLI `lsfbvol` command output to display the unprotected, previously mirrored load source LUNs with a datatype of 520U that refer to unprotected volumes, as shown in Figure 5-78.

```

dscli> lsfbvol
Date/Time: November 7, 2007 3:25:51 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
Name      ID      accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)  cap (blocks)
-----
TN11s    105E  Online   Normal    Normal      2107-A85  FB 520U  P0      32.8       35.2       6868172
TN1Vol1  105F  Online   Normal    Normal      2107-A05  FB 520P  P0      32.8       35.2       6868172
TN1um    1060  Online   Normal    Normal      2107-A85  FB 520U  P4      32.8       35.2       6868172
TN1Vol2  1061  Online   Normal    Normal      2107-A05  FB 520P  P4      32.8       35.2       6868172
  
```

Figure 5-78 Listing unprotected disks



- Change only the *unconfigured* previous load source volume from *unprotected* to *protected* using the `chfbvol -os400 protected volumeID` command as shown in Figure 5-79.

```

dscli> chfbvol -os400 protected 1060
Date/Time: November 7, 2007 4:04:41 AM IST IBM DSCLI Version: 5.3.0.991 DS: IBM.2107-7589951
CMUC00026I chfbvol: FB volume 1060 successfully modified.

TN11s      105E Online Normal Normal 2107-A85 FB 520U P0 32.8 35.2 6868172
TN1Vol1    105F Online Normal Normal 2107-A05 FB 520P P0 32.8 35.2 6868172
TN1mm      1060 Online Normal Normal 2107-A05 FB 520P P4 32.8 35.2 6868172
TN1Vol2    1061 Online Normal Normal 2107-A05 FB 520P P4 32.8 35.2 6868172

```

Figure 5-79 Changing volume protection

- Perform an IOP reset for the IOA that is attached to the unconfigured previous load source volume on which you changed the protection mode on the storage system in the previous step.

**Note:** Note this IOP reset is required for System i to rediscover its devices for recognizing the changed LUN protection mode.

To reset the IOP from SST/DST select the following options:

- 1. Start a service tool
- 7. Hardware service manager
- 2. Logical hardware resources
- 1. System bus resources

Then, select the correct 2847 IOP (the one that is *not* the load source IOP), and choose **6. I/O debug**, as shown in Figure 5-80.

```

Logical Hardware Resources on System Bus

System bus(es) to work with . . . . . *ALL *ALL, *SPD, *PCI, 1-9999
Subset by . . . . . *ALL *ALL, *STG, *WS, *CMN, *CRP

Type options, press Enter.
2=Change detail      4=Remove      5=Display detail    6=I/O debug
7=Display system information
8=Associated packaging resource(s)    9=Resources associated with IOP

Opt Description          Type-Model  Status          Resource
Bus Expansion Adapter    28E7-      Operational     BCC10
System Bus               28B7-      Operational     LB09
Multi-adapter Bridge     28B7-      Operational     PCI11D
6 Combined Function IOP  2847-001   Operational     CMB03
HSL I/O Bridge           28E7-      Operational     BC05
Bus Expansion Adapter    28E7-      Operational     BCC05
System Bus               28B7-      Operational     LB04
More...

F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
F9=Failed resources    F10=Non-reporting resources
F11=Display serial/part numbers    F12=Cancel

```

Figure 5-80 Selecting IOP for reset

5. Select **3. Reset I/O processor** to reset the IOP as shown in Figure 5-81.

```
                                Select IOP Debug Function

Resource name . . . . . : CMB03
Dump type . . . . . : Normal

Select one of the following:

    1. Read/Write I/O processor data
    2. Dump I/O processor data
    3. Reset I/O processor
    4. IPL I/O processor
    5. Enable I/O processor trace
    6. Disable I/O processor trace

Selection
    3

F3=Exit      F12=Cancel
F8=Disable I/O processor reset    F9=Disable I/O processor IPL
```

Figure 5-81 Reset IOP option

6. Press Enter to confirm the IOP reset, as shown in Figure 5-82.

```
                                Confirm Reset Of IOP

You have requested that an I/O processor be reset.

Note: This will disturb active jobs running on
this IOP or on the devices attached to this IOP.

Press Enter to confirm your actions.
Press F12 to cancel this request.

F3=Exit      F12=Cancel
```

Figure 5-82 Confirming IOP reset

After a the IOP is reset successfully, a confirmation message displays, as shown in Figure 5-83.

```
                Select IOP Debug Function

Resource name . . . . . : CMB03
Dump type . . . . . : Normal

Select one of the following:

    1. Read/Write I/O processor data
    2. Dump I/O processor data
    3. Reset I/O processor
    4. IPL I/O processor
    5. Enable I/O processor trace
    6. Disable I/O processor trace

Selection

F3=Exit      F12=Cancel
F8=Disable I/O processor reset      F9=Disable I/O processor IPL
Reset of IOP was successful.
```

Figure 5-83 IOP reset confirmation message

7. Now, select **4. IPL I/O processor** in the Select IOP Debug Function menu to IPL the I/O as shown in Figure 5-84. Press Enter to confirm your selection.

```
                Select IOP Debug Function

Resource name . . . . . : CMB03
Dump type . . . . . : Normal

Select one of the following:

    1. Read/Write I/O processor data
    2. Dump I/O processor data
    3. Reset I/O processor
    4. IPL I/O processor
    5. Enable I/O processor trace
    6. Disable I/O processor trace

Selection
    4

F3=Exit      F12=Cancel
F8=Disable I/O processor reset      F9=Disable I/O processor IPL
```

Figure 5-84 IPL I/O

After a successful IPL, a confirmation message displays, as shown in Figure 5-85.

```

Select IOP Debug Function

Resource name . . . . . : CMB03
Dump type . . . . . : Normal

Select one of the following:

    1. Read/Write I/O processor data
    2. Dump I/O processor data
    3. Reset I/O processor
    4. IPL I/O processor
    5. Enable I/O processor trace
    6. Disable I/O processor trace

Selection

F3=Exit      F12=Cancel
F8=Disable I/O processor reset      F9=Disable I/O processor IPL
Re-IPL of IOP was successful.
  
```

Figure 5-85 I/O IPL confirmation message

8. Next, check the changed protection status for the unconfigured previous load source LUN in the SST Display non-configured units menu as shown in Figure 5-86.

```

Display Non-Configured Units

Serial          Resource
Number         Type Model Name   Capacity Status
50-1060951     2107 A05 DMP006   35165 Non-configured

Press Enter to continue.

F3=Exit      F5=Refresh      F9=Display disk unit details
F11=Display device parity status      F12=Cancel
  
```

Figure 5-86 SST - Display Non-Configured Units

Now, we explain the remaining steps to change the unprotected load source unit to a protected load source. To look at the current unprotected load source unit, we choose the DST menu function Display disk configuration status as shown in Figure 5-87.

```

Display Disk Configuration Status

ASP Unit      Serial          Resource
1             Number         Type Model Name   Status
1             50-105F951     2107 A85 DMP007   Unprotected
              50-1061951     2107 A05 DMP003   RAID 5/Active
              50-105F951     2107 A05 DMP001   RAID 5/Active

Press Enter to continue.

F3=Exit      F5=Refresh      F9=Display disk unit details
F11=Disk configuration capacity      F12=Cancel
  
```

Figure 5-87 Display Disk Configuration Status

9. Select **4. Work with disk units** in the DST main menu, as shown in Figure 5-88.

```
Use Dedicated Service Tools (DST)                               System:  RCHLTTN1
Select one of the following:
1. Perform an IPL
2. Install the operating system
3. Work with Licensed Internal Code
4. Work with disk units
5. Work with DST environment
6. Select DST console mode
7. Start a service tool
8. Perform automatic installation of the operating system
9. Work with save storage and restore storage
10. Work with remote service support
12. Work with system capacity
13. Work with system security
14. End batch restricted state
Selection
4
F3=Exit  F12=Cancel
```

Figure 5-88 DST: Main menu

10. Select **2. Work with disk unit recovery** as shown in Figure 5-89.

```
Work with Disk Units
Select one of the following:
1. Work with disk configuration
2. Work with disk unit recovery
Selection
2
F3=Exit  F12=Cancel
```

Figure 5-89 Work with Disk Units

11. Select **9. Copy disk unit data** as shown in Figure 5-90.

```
Work with Disk Unit Recovery

Select one of the following:

  1. Save disk unit data
  2. Restore disk unit data
  3. Replace configured unit
  4. Assign missing unit
  5. Recover configuration
  6. Disk unit problem recovery procedures
  7. Suspend mirrored protection
  8. Resume mirrored protection
  9. Copy disk unit data
 10. Delete disk unit data
 11. Upgrade load source utility
 12. Rebuild disk unit data
 13. Reclaim IOA cache storage

More...

Selection
  9

F3=Exit  F11=Display disk configuration status  F12=Cancel
```

Figure 5-90 DST: Copy disk unit data

12. Select the current unprotected load source unit 1 as the disk unit from which to copy, as shown in Figure 5-91.

```
Select Copy from Disk Unit

Type option, press Enter.
1=Select

      Serial
OPT Unit  ASP Number  Type Model Name  Status
-----
  1      1  50-105E951  2107  A85  DMP007  Configured
  2      1  50-1061951  2107  A05  DMP003  RAID 5/Active
  3      1  50-105F951  2107  A05  DMP001  RAID 5/Active

F3=Exit  F5=Refresh  F11=Display non-configured units  F12=Cancel
```

Figure 5-91 DST: Copy from Disk Unit

13. Select the unconfigured previous load source mirror as the copy-to-disk-unit, as shown in Figure 5-92.

```

Select Copy to Disk Unit Data

Disk being copied:

      Serial
Unit  ASP Number      Type Model Name      Resource      Status
  1    1 50-105E951    2107 A85 DMP007      Configured

1=Select

      Serial
Option Number      Type Model Name      Resource      Status
  1    50-1060951    2107 A05 DMP006      Non-configured

F3=Exit  F11=Display disk configuration status  F12=Cancel

```

Figure 5-92 DST: Select Copy to Disk Unit Data

14. Press Enter to confirm the choice, as shown in Figure 5-93.

```

Confirm Copy Disk Unit Data

Press Enter to confirm your choice for copy.
Press F12 to return to change your choice.

Disk being copied:

      Serial
Unit  ASP Number      Type Model Name      Resource      Status
  1    1 50-105E951    2107 A85 DMP007      Configured

Disk that is copied to:

      Serial
Number      Type Model Name      Resource      Status
  50-1060951  2107 A05 DMP006      Non-configured

F12=Cancel

```

Figure 5-93 Confirm Copy Disk Unit Data

During the copy process, the system displays the Copy Disk Unit Data Status panel, as shown in Figure 5-94.

```

Copy Disk Unit Data Status

The operation to copy a disk unit will be done in several phases.
The phases are listed here and the status will be indicated when
known.

Phase                                     Status

Stop compression (if needed) . . . . . : Completed
Prepare disk unit . . . . . : Completed
Start compression (if needed) . . . . . : Completed
Copy status . . . . . : 99 % Complete

Number of unreadable pages:

```

Figure 5-94 Copy status

15. After the copy process completes successfully, the system IPLs automatically. During the IPL, a message displays, as shown in Figure 5-95, because the system found an unconfigured unit as the previous load source IPL. You can continue by selecting **1. Keep the current disk configuration**, as shown in Figure 5-95.

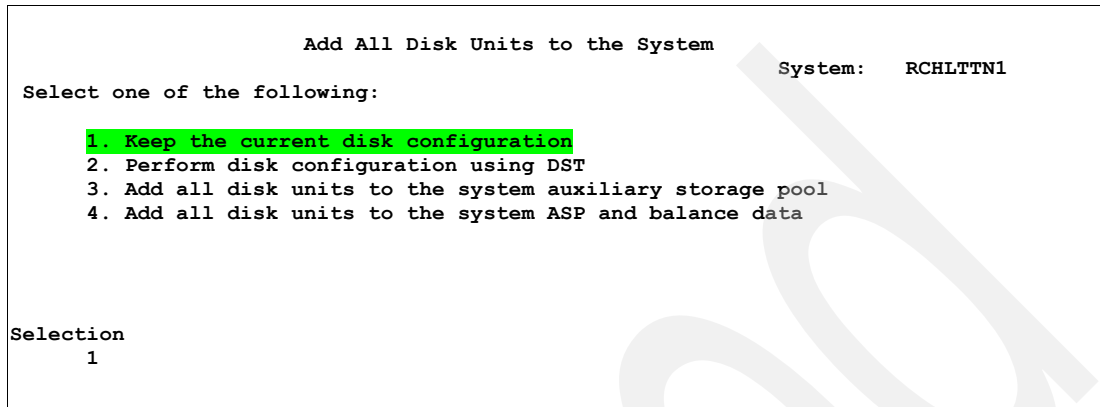


Figure 5-95 Add All Disk Units to the System

16. Next, look at the protected load source unit using the Display Disk Configuration Status menu, as shown in Figure 5-96.

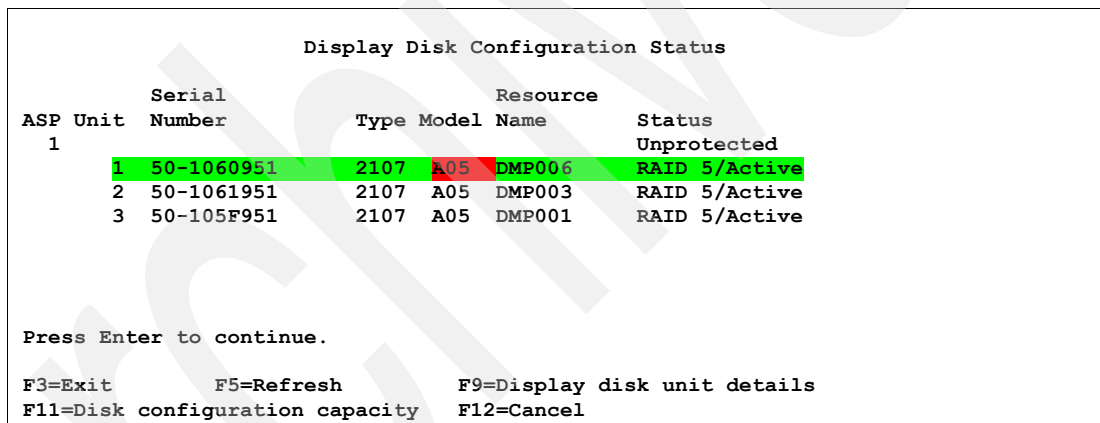


Figure 5-96 Display Disk Configuration Status

17. Then, look at the previous load source unit with its unprotected status using the Display Non-Configured Units menu as shown in Figure 5-97.

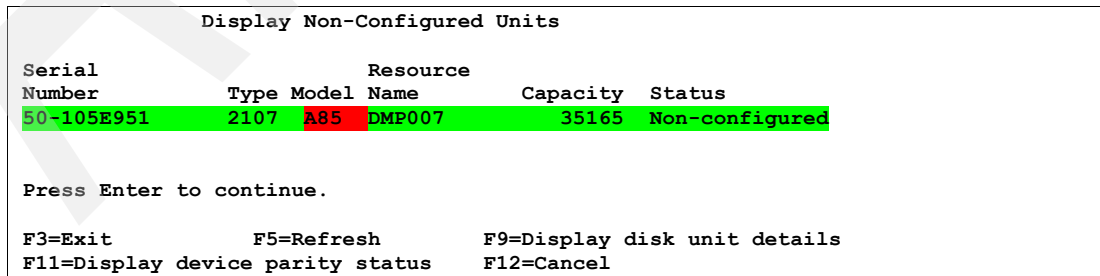


Figure 5-97 Display Non-Configured Units



18. If you want to change this non-configured unit that was the previous load source from which you migrated the data to a protected unit, then use the DS CLI `chfbvo1` command and an IOP reset or re-IPL as described in steps 1 to 4.

## 5.12 Migration considerations from IOP-based to IOP-less Fibre Channel

The migration from IOP-based to IOP-less Fibre Channel applies only to customers who continued using older #2787 or #5760 IOP-based Fibre Channel IOAs on a new System i POWER6 server (note that #2766 is not supported on System i POWER6) and now want to remove the old IOP-based technology to take advantage from the new IOP-less Fibre Channel performance and its higher integration.

**Note:** Carefully plan and size your IOP-less Fibre Channel adapter card placement in your System i server and its attachment to your storage system to avoid potential I/O loop or FC port performance bottlenecks with the increased IOP-less I/O performance. Refer to Chapter 3, “i5/OS planning for external storage” on page 51 and Chapter 4, “Sizing external storage for i5/OS” on page 89 for further information.

**Important:** Do *not* try to workaround the migration procedures that we discuss in this section by concurrently replacing the IOP/IOA pair for one mirror side or one path after the other. Concurrent hardware replacement is supported only for like-to-like replacement using the same feature codes.

Because the migration procedures are pretty much straightforward, we only outline the required steps for different configurations.

### 5.12.1 IOP-less migration in a multipath configuration

When using IOP-based i5/OS multipathing, you can perform the migration to IOP-less Fibre Channel concurrently without shutting down the system as follows:

1. Add an IOP-less IOA into another I/O slot.
2. Move the FC cable from the old IOP-based FC IOA to the new IOP-less IOA.
3. Change the host connection on the DS storage system to reflect the new WWPN.

Internally for each multipath group, this process creates a new multipath connection. Some time later, you need to remove the obsolete connection using the multipath reset function (see 5.13, “Resetting a lost multipath configuration” on page 242).

### 5.12.2 IOP-less migration in a mirroring configuration

When using IOP-based i5/OS mirroring for external storage, you can migrate to IOP-less Fibre Channel as follows:

1. Turn off the System i server.
2. Replace the Fibre Channel IOP/IOA cards with IOP-less cards.
3. Change the host connections on the DS storage system to reflect the new WWPNs.

### 5.12.3 IOP-less migration in a configuration without path redundancy

If you do not use multipath or mirroring, you need to follow these steps to migrate to IOP-less Fibre Channel:

1. Turn off the System i server.
2. Replace the Fibre Channel IOP/IOA cards with IOP-less cards.
3. Change the host connections on the DS storage system to reflect the new WWPNs.

## 5.13 Resetting a lost multipath configuration

If after changing your System i storage attachment configuration, any unknown paths are reported, which can happen typically when you have reduced the number of Fibre Channel paths to the System i host, follow the procedure that we describe in this section. A reset of the multipath configuration on the System i server frees up orphan path resource information after a multipath configuration change.

**Note:** An IPL might be required so that the System i recognizes the missing paths.

### 5.13.1 Resetting the lost multipath configuration for V6R1

To reset the lost multipath configuration for V6R1:

1. Log in to i5/OS System Service Tools using the STRSST command.
2. Select **1. Start a service tool**.
3. Select **7. Hardware service manager**.
4. Select **1. Packaging hardware resources**.
5. Select **9** for Disk Unit System as shown in Figure 5-98.

```

Packaging Hardware Resources
Local system type . . . . . : 9406
Local system serial number:  XX-XXXXX

Type options, press Enter.
2=Change detail      3=Concurrent maintenance      4=Remove      5=Display detail
8=Associated logical resource(s)      9=Hardware contained within package

Type-
Opt Description      Resource      Model      Unit ID      Name
Optical Storage Unit      =      6333-002 U787B.001.DNW5A3B      SD001
Tape Unit              6380-001      SD003
9 Disk Unit System      +      DE01
```

Figure 5-98 Packaging Hardware Resources

6. Select **7=Paths to multiple path disk** on the disks that you want to reset as shown in Figure 5-99.

```

Disk Units Contained Within Package
Resource name:  DE01

Type options, press Enter.
2=Change detail  4=Remove  5=Display detail  7=Paths to multiple path
disk
8=Associated logical resource(s)

Type-      Serial      Resource      Multiple
Opt Model  Number      Name          Path Disk
-----
 2107-A82  50-10000B2  DMP025       Unknown     Yes
 7 2107-A82  50-10000B2  DMP024       Operational Yes
 2107-A82  50-10010B2  DMP027       Unknown     Yes
 2107-A82  50-10010B2  DMP026       Unknown     Yes
 2107-A82  50-10010B2  DMP023       Unknown     Yes
 2107-A82  50-10010B2  DMP022       Operational Yes
 2107-A82  50-10020B2  DMP011       Unknown     Yes
 2107-A82  50-10020B2  DMP021       Operational Yes

F3=Exit      F5=Refresh    F6=Print      F12=Cancel    F14=Reset paths

```

Figure 5-99 Disk Units Contained Within Package

7. When prompted for confirmation, press F10, as shown in Figure 5-100.

```

Reset Paths to Multiple Path Disk Unit

WARNING: This service function should be run only under the direction of
the IBM Hardware Service.

You have selected to reset the number of paths on a multipath unit to
equal
the number of paths currently enlisted.

Press F10 to reset the paths to the following multipath disk units.

See help for more details

Type-      Serial      Resource      Logical Address
Model      Number      Name          /  /  /  /  /  /
2107-A82  50-10000B2  DMP024       2/ 34/0/ 32-2/ 6/ 0/ 3/ 1/ /
2107-A82  50-10000B2  DMP025       2/ 34/0/ 32-2/ 4/ 0/ 3/ 1/ /

F3=Exit      F10=Confirm    F12=Cancel

```

Figure 5-100 Reset Paths to Multiple Path Disk Unit

8. When the operation is complete, a confirmation panel displays, as shown in Figure 5-101.

```
Reset Paths to Multiple Path Disk Unit

WARNING: This service function should be run only under the direction of
the IBM Hardware Service.

You have selected to reset the number of paths on a multipath unit to
equal
the number of paths currently enlisted.

Press F10 to reset the paths to the following multipath disk units.

See help for more details

Type-      Serial      Resource
Model      Number      Name        Logical Address
2107-A82   50-10000B2  DMP024     2/ 34/0/ 32-2/ 6/ 0/ 3/ 1/
2107-A82   50-10000B2  DMP025     2/ 34/0/ 32-2/ 4/ 0/ 3/ 1/

F3=Exit    F10=Confirm  F12=Cancel
oval of the selected resources was successful.
```

Figure 5-101 SST: Multipath reset confirmation panel

**Note:** The DMPxxx resource name is *not* reset to DDxxx when multipathing is stopped.

### 5.13.2 Resetting a lost multipath configuration for versions prior to V6R1

To reset a lost multipath configuration for versions prior to V6R1:

1. Start i5/OS to DST or, if i5/OS is running, access SST and sign in. Select **1. Start a Service Tool**.
2. In the Start a Service Tool panel, select **1. Display/Alter/Dump**, as shown in Figure 5-102.

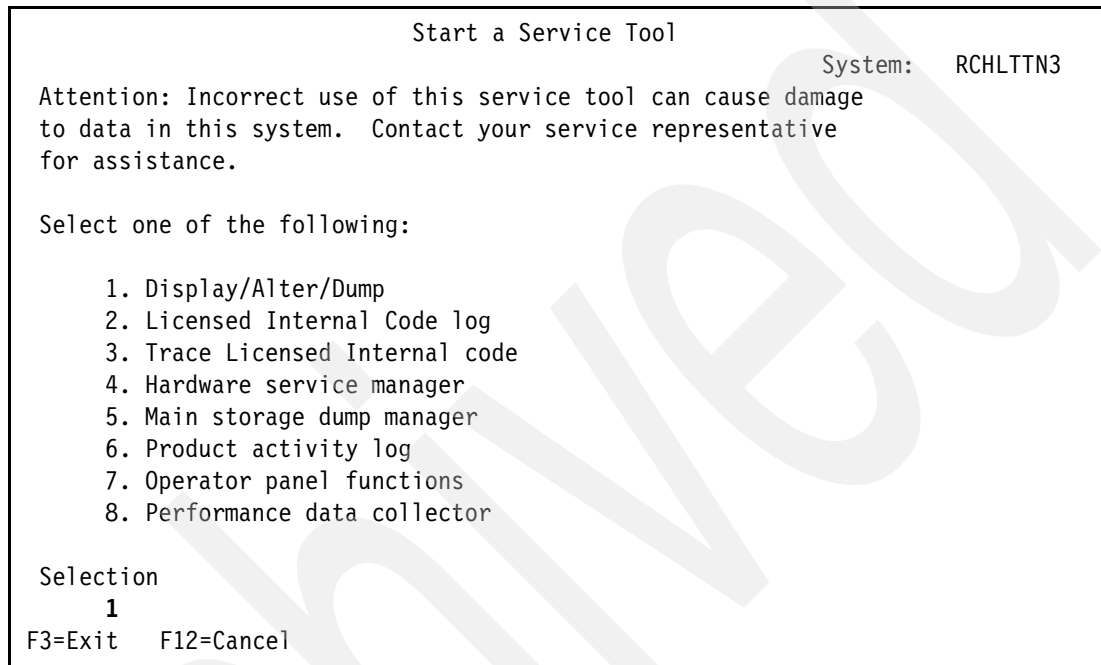


Figure 5-102 Starting a Service Tool panel

3. In the Display/Alter/Dump Output Device panel, select **1. Display/Alter/Dump**, as shown in Figure 5-103.

**Attention:** Use *extreme* caution when using the Display/Alter/Dump Output panel because you can end up damaging your system configuration. Ideally, when performing these tasks for the first time, do so after referring to IBM Support.

```
Display/Alter/Dump Output Device

Select one of the following:

    1. Display/Alter storage
    2. Dump to printer

    4. Dump to media

    6. Print dump from media
    7. Display dump status
Selection
    1
F3=Exit  F12=Cancel
```

Figure 5-103 Display/Alter/Dump Device Output panel

4. In the Select Data panel, select **2. Licensed Internal Code (LIC) data**, as shown in Figure 5-104.

```
Select Data

Output device . . . . . : Display

Select one of the following:

    1. Machine Interface (MI) object
    2. Licensed Internal Code (LIC) data
    3. LIC module
    4. Tasks/Processes
    5. Starting address
Selection
    2
F3=Exit  F12=Cancel
```

Figure 5-104 Selecting data for Display/Alter/Dump

5. In the Select LIC Data panel, scroll down the page, and select **14. Advanced analysis** (as shown in Figure 5-105), and press Enter.

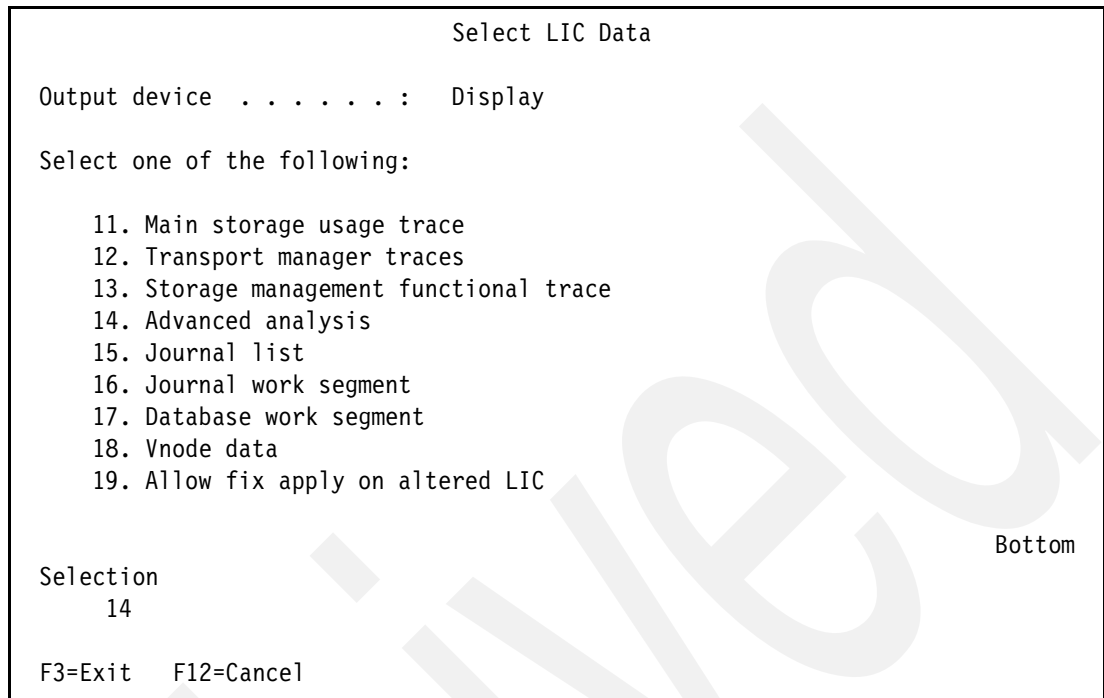


Figure 5-105 Selecting Advanced analysis

6. In the Select Advanced Analysis Command panel, scroll down the page and select **1** to run the MULTIPATHRESETTER macro, as shown in Figure 5-106.

```
                                Select Advanced Analysis Command

Output device . . . . . : Display

Type options, press Enter.
  1=Select

Option   Command

        JAVALOCKINFO
        LICLOG
        LLHISTORYLOG
        LOCKINFO
        MASOCONTROLINFO
        MASOWAITERINFO
        MESSAGEQUEUE
        MODINFO
        MPLINFO
  1     MULTIPATHRESETTER
        MUTEXDEADLOCKINFO
        MUTEXINFO

                                                More...

F3=Exit  F12=Cancel
```

Figure 5-106 Select Advanced Analysis Command panel

7. The multipath resetter macro has various options, which are displayed in the Specify Advanced Analysis Options panel (Figure 5-107). For Options, enter **-RESTMP -ALL**.

```
                                Specify Advanced Analysis Options

Output device . . . . . : Display

Type options, press Enter.

Command . . . . . : MULTIPATHRESETTER

Options . . . . . -RESTMP -ALL

F3=Exit  F4=Prompt  F12=Cancel
```

Figure 5-107 Multipath reset options



The Display Formatted Data panel displays as confirmation (Figure 5-108).

```
Display Formatted Data
Page/Line. . . . . 1 / 1
Columns. . . . . : 1 - 78
Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
DISPLAY/ALTER/DUMP
Running macro: MULTIPATHRESETTER -RESETMP -ALL
Reset the paths for Multiple Connections

***RESET MULTIPATH UNIT PATHS TO NUMBER CURRENTLY ENLISTED***

This service function should be run only under the direction of the
IBM Hardware Service Support. You have selected to reset the
number of paths on a multipath unit to equal the number of paths
that have currently enlisted.

To force the error, do the following:
  1. Press Enter now to return to the previous display.
  2. Change the '-RESETMP keyword on the Options line to
     '-CONFIRM' and press Enter

More...
F2=Find  F3=Exit  F4=Top  F5=Bottom  F10=Right  F12=Cancel
```

Figure 5-108 Multipath reset confirmation

8. Press Enter to return to the Specify Advanced Analysis Options panel (Figure 5-109). For Options, enter -CONFIRM -ALL.

```
Specify Advanced Analysis Options
Output device . . . . . : Display
Type options, press Enter.
Command . . . . . : MULTIPATHRESETTER
Options . . . . . -CONFIRM -ALL
F3=Exit  F4=Prompt  F12=Cancel
```

Figure 5-109 Confirming the multipath reset

- In the Display Formatted Data panel (Figure 5-110), press F3 to return to the Specify Advanced Analysis Options panel (Figure 5-107 on page 248).

```

                                Display Formatted Data
                                Page/Line. . .      1 /   1
                                Columns. . . : 1 - 78

Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
  DISPLAY/ALTER/DUMP
Running macro: MULTIPATHRESETTER          -CONFIRM -ALL
Reset the paths for Multiple Connections

*****
***CONFIRM RESET MULTIPATH UNIT PATHS TO NUMBER CURRENTLY ENLISTED***
*****

This service function should be run only under the direction of the
IBM Hardware Service Support.

You have selected to reset the number of paths on a multipath unit
to equal the number of paths that have currently enlisted.

Attempting to reset path for resource name: DMP003

More...

F2=Find  F3=Exit  F4=Top  F5=Bottom  F10=Right  F12=Cancel

```

Figure 5-110 Multipath reset results

- In the Specify Advanced Analysis Options panel (Figure 5-109 on page 249), repeat the confirmation process to ensure that the path reset is performed. Retain the setting for the Option parameter as -CONFIRM -ALL, and press Enter again.

11. The Display Formatted Data panel shows the results (Figure 5-111). In our example, it indicates that no disk unit paths have to be reset.

```
Display Formatted Data
Page/Line. . .      1 /  1
Columns. . . :  1 - 78
Find . . . . .
.....1.....2.....3.....4.....5.....6.....7.....
DISPLAY/ALTER/DUMP
Running macro: MULTIPATHRESETTER          -CONFIRM -ALL
Reset the paths for Multiple Connections

*****
***CONFIRM RESET MULTIPATH UNIT PATHS TO NUMBER CURRENTLY ENLISTED***
*****

This service function should be run only under the direction of the
IBM Hardware Service Support.

You have selected to reset the number of paths on a multipath unit
to equal the number of paths that have currently enlisted.

Could not find any disk units with paths which need to be reset.

F2=Find  F3=Exit  F4=Top  F5=Bottom  F10=Right  F12=Cancel

Bottom
```

Figure 5-111 No disks have to be reset

**Note:** The DMPxxx resource name is *not* reset to DDxxx when multipathing is stopped.

Archived



## Implementing FlashCopy using the DS CLI

In this chapter, we show how to implement IBM System Storage Copy Services with i5/OS using the DS command-line interface (CLI).

The examples in this book provide complete details about the IBM System i setup, the IBM System Storage DS8000 setup, and the IBM System Storage DS6000 setup. The examples are simple scenarios that System i users can work on in a test environment before creating them in a production environment.

In addition to this book, we recommend that you consult the list of books that we provide in “Related publications” on page 479.

## 6.1 Overview of IBM System Storage DS CLI

The IBM System Storage DS CLI enables open system hosts to invoke and manage FlashCopy and Remote Mirror and Copy functions through batch processes and scripts. The CLI provides a full-function command set that allows you to check the status of Copy Services and perform specific Copy Services functions when necessary.

Use the DS CLI to implement and manage the following Copy Services functions:

- ▶ FlashCopy
- ▶ Metro Mirror, formerly called *synchronous Peer-to-Peer Remote Copy* (PPRC)
- ▶ Global Copy, formerly called *PPRC Extended Distance* (PPRC-XD)
- ▶ Global Mirror, formerly called *asynchronous PPRC*

### 6.1.1 Installing and setting up the DS CLI

The DS CLI is supplied and installed through a CD that ships with the IBM System Storage platform. You can install and use the DS CLI on any supported operating system, including i5/OS. Although this chapter provides an example of DS CLI on Windows, you can manage Copy Services using DS CLI in i5/OS. However, you must ensure that the source server and the target server are shut down when using Copy Services for the entire direct access storage device (DASD) space of i5/OS.

To manage Copy Services, you can use the DS CLI functions on both a PC and i5/OS, depending on the operation that you are performing. Starting and stopping the Copy Services environment can, for example, be managed from i5/OS. The switchover of Metro Mirror or Global Mirror, or copying an entire i5/OS DASD space, requires the DS CLI on a Windows PC or other system or a partition running i5/OS. For more information about installing, setting up, and starting DS CLI, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

Before you start DS CLI, we recommend that you create a DS CLI profile or adjust the default DS CLI profile with values that are specific to your DS CLI environment. Examples of these values include the IP address of the Hardware Management Console (HMC) or the Storage Management Console (SMC), password file, and storage image ID. This way, you archive the values that you do not have to insert every time you invoke the DS CLI or use the DS CLI command.

## 6.2 Implementing traditional FlashCopy

In this section, we describe how to implement traditional FlashCopy, that is not DS8000 R3 space efficient FlashCopy, with the entire DASD space of i5/OS. Figure 6-1 provides an overview of the test environment that we used for this implementation example.

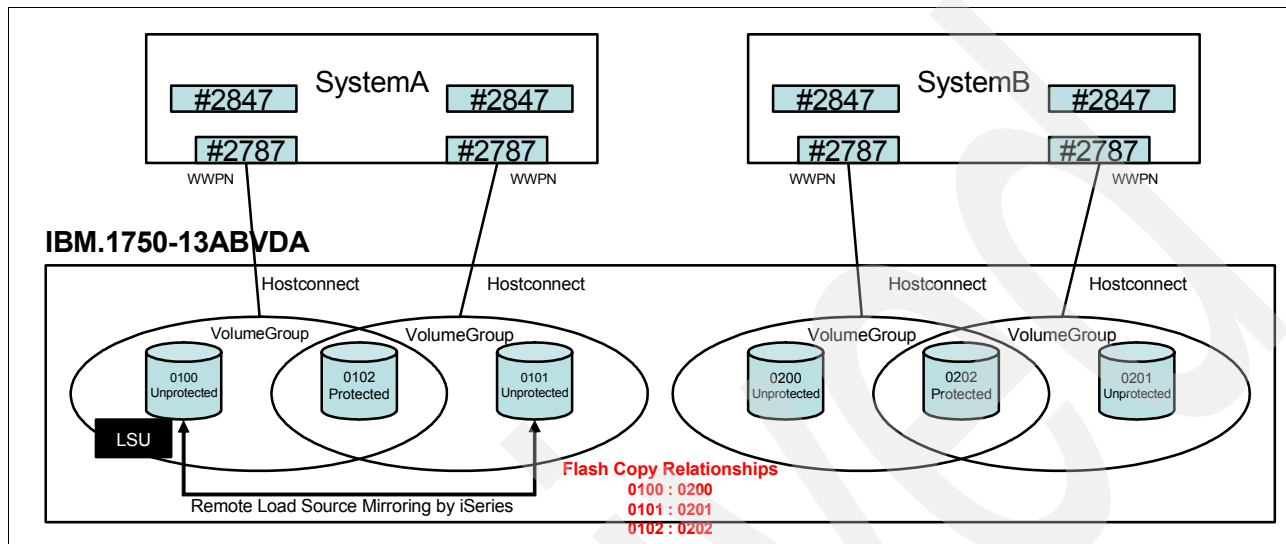


Figure 6-1 Test environment

System A is the source server, and system B is the target server for FlashCopy. Both systems are connected to a local external storage subsystem. System A is booted from the storage area network (SAN) with 2847 I/O processor (IOP). It has a load source unit in volume 0100 and mirrored load source unit in volume 0101. The first 2847 feature card is tagged as the load source in the HMC partition configuration. In order to maintain a simple scenario in this example, system A has only one volume 0102 for data.

**Note:** With the new i5/OS V6R1 multipath load source support, you do not need to mirror the load source to provide path protection. In our discussion, we continue to show the external load source mirroring setup for existing systems for demonstration purposes only.

The implementation example that we describe in this section uses FlashCopy to make a copy of the entire DASD space of system A and boot system B with the copied DASD space. This example assumes that the DASD environment is created and i5/OS V5R3M5 or later is installed. For more information about the setup and installation of the base disk space and i5/OS load, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

To implement FlashCopy for the entire DASD space of the System i environment, follow these steps:

1. Turn off or quiesce the source server.
2. Implement FlashCopy in IBM System Storage with DS CLI.
3. Perform an IPL or resume of the source server.
4. Perform an IPL of the target server.

**Important:** If you are considering using Metro Mirror, Global Mirror, or FlashCopy for the replication of the load source unit or other i5/OS disk units within the same DS6000 or DS8000 or between two or more DS6000 or DS8000 systems, the source volume and the target volume characteristics *must be identical*. The target and source must have matching capacities and matching protection types. For example, a *protected* 35 GB i5/OS volume must be replicated to another *protected* 35 GB volume if it is to be used in a replicated i5/OS configuration. It cannot be replicated to an *unprotected* 35 GB volume or to a volume of any other capacity. If you plan to migrate your load source from one size LUN to a larger size, use the *SLIC copy disk unit data utility* that is available in Dedicated Service Tools (DST).

In addition, the DS CLI offers the capability to change characteristics such as the protection type of a previously defined volume. After a volume is assigned to an i5/OS partition and added to that partition's configuration, its characteristics must *not* be changed. If there is a requirement to change the characteristic of a configured volume, you must first completely remove it from the i5/OS configuration. After you make the characteristic changes, for example to protection type, capacity, and so on, by destroying and recreating the volume or by using the DS CLI, you can then reassign the volume to the i5/OS configuration. To simplify the configuration, we recommend a symmetrical configuration between two IBM System Storage solutions, creating the same volumes with the same volume ID that determines the LSS\_ID.

For FlashCopy, we recommend that you place your target volumes on a different rank from where the source volumes are assigned for performance of the source server.

## 6.2.1 Turning off the source server

Before you initiate FlashCopy, turn off or quiesce your source server. Follow the procedures that pertain to your site or use the PWRDWNSYS command. Alternatively, to turn off your source server, you can use the i5/OS V6R1 quiesce for Copy Services CHGASPACT command to quiesce all database I/O for your system (see 15.1, "Using i5/OS quiesce for Copy Services" on page 432).

## 6.2.2 Setting up the FlashCopy environment

After you turn off the source server, you must create volumes for the target. You must also create volume groups and host connections to boot the target server. In order to allow the target server to boot from the copied load source unit on external storage, set the tagged load source unit to 2847 IOP or the Fibre Channel (FC) controller that is connected to the copied load source unit with HMC. Perform these tasks only the first time that you create the FlashCopy environment. Change the settings only if the source environment changes. In this section, these tasks have been performed already.

For more information about creating volumes, volume groups, and host connections, as well as tagging the load source IOP, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

**Important:** When you create the volumes for the target server, create the same number and type of volumes as for the source. You also need to match the protection type, either protected or unprotected. Target volumes must be planned in order to be assigned on a rank that is different from where the source volumes are assigned and in order to maintain the performance of the source server.



In this implementation example, we assume that you have a PC that is running Windows, has DS CLI installed, and is connected to the DS system and i5/OS. This example uses IBM System Storage DS6000. With DS CLI, implementing FlashCopy is a two-step process:

1. Check which fixed volumes are available as FlashCopy pairs on the DS system. Run the following command, where *storage\_image\_ID* relates to the DS6000:

```
dscli lsfbvol -dev <storage_image_ID>
```

Example 6-1 shows the results where the hexadecimal volume ID in the selected storage image. The second column shows the available fixed block volumes for the DS6000 system.

*Example 6-1 Output of the lsfbvol command*

---

```
dscli> lsfbvol -dev IBM.1750.13-ABVDA
Date/Time: October 26, 2005 6:25:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
```

Name	ID	accstate	datastate	configstate	deviceMTM	datatype	extpool	cap (2^30B)	cap (10^9B)
rchltn2-boot	0100	Online	Normal	Normal	1750-A85	FB 520U	P1	32.8	35.2
rchltn2-boot-mr	0101	Online	Normal	Normal	1750-A85	FB 520U	P1	32.8	35.2
rchltn2-disk2	0102	Online	Normal	Normal	1750-A05	FB 520P	P1	32.8	35.2
rchltn3-boot	0200	Online	Normal	Normal	1750-A85	FB 520U	P0	32.8	35.2
rchltn3-boot-mr	0201	Online	Normal	Normal	1750-A85	FB 520U	P0	32.8	35.2

---

2. Select volume pairs on the DS6000 system and create FlashCopy pairs using the following DS CLI command:

```
dscli mkflash -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID>
```

3. Specify the volume pairs by their ID. This implementation runs the following command with the **-nocp** parameter included:

```
dscli mkflash -dev <storage_image_ID> -nocp <source_volume_ID>:<target_volume_ID>
```

The **-nocp** parameter indicates that this example does *not* run a full-copy of the disk space but only creates the FlashCopy bitmap and copies tracks to the target which are going to be modified on the source system. For a full-copy of the disk space, omit the **-nocp** option.

Whether you do a full-copy or no-copy depends on how you want to use the FlashCopy targets. In a real environment where you might want to use the copy over a longer time and with high I/O workload, a full-copy is a better option to isolate your backup system I/O workload from your production workload when all data has been copied to the target. For the purpose of using FlashCopy for creating a temporary system image for save to tape during low production workload, the no-copy option is recommended.

Example 6-2 creates three FlashCopy pairs with source volumes 0100, 0101, and 0102 and target volumes 0200, 0201, and 0202 with the full-copy option on a DS6000 system.

*Example 6-2 Output of the mkflash command*

---

```
dscli> mkflash -dev IBM.1750-13ABVDA 0100-0102:0200-0202
Date/Time: October 15, 2005 4:09:52 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00137I mkflash: FlashCopy pair 0100:0200 successfully created.
CMUC00137I mkflash: FlashCopy pair 0101:0201 successfully created.
CMUC00137I mkflash: FlashCopy pair 0102:0202 successfully created.
```

---

## 6.2.3 Performing an IPL of the source server

After the FlashCopy pair relationship is established, you either perform an IPL of the source server if you turned off the server or, if you quiesced the system, resume it using the CHGASPACT DEVICE(*aspname*) OPTION(\*RESUME) command. If you perform an IPL, you most likely use a normal mode from the system B side. Follow the procedures that pertain to your site or use the HMC to activate the partition manually.

To display a FlashCopy relationship and its properties, enter the following command:

```
dscli lsflash -l -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID>..
```

To omit the current OutOfSyncTracks attribute, do not use the *target\_volume\_ID* parameter and the -l option.

Example 6-3 lists the FlashCopy sessions for volume pairs with source volumes 0100, 0101, and 0102 and target volumes 0200, 0201, and 0202. This example shows that the current number of tracks that are not synchronized as an *OutOfSyncTracks* attribute. By reviewing this attribute, you can determine the progress of the background copy. If you use a full copy option for initiating FlashCopy, when the number of OutOfSyncTracks becomes 0, all the data is copied from the source volume to the target volume through the background copy process.

*Example 6-3 Output of the lsflash command*

```
dscli> lsflash -l 0100-0102
Date/Time: October 15, 2005 4:12:02 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Reversible SourceWriteEnabled
TargetWriteEnabled BackgroundCopy OutOfSyncTracks DateCreated          DateSynced
=====
0100:0200 01      0          300    Disabled Disabled Disabled Disabled Enabled          Enabled
Enabled    536576          Fri Oct 14 22:58:24 JST 2005 Fri Oct 14 22:58:24 JST 2005
0101:0201 01      0          300    Disabled Disabled Disabled Disabled Enabled          Enabled
Enabled    536576          Fri Oct 14 22:58:24 JST 2005 Fri Oct 14 22:58:24 JST 2005
0102:0202 01      0          300    Disabled Disabled Disabled Disabled Enabled          Enabled
Enabled    496826          Fri Oct 14 22:58:24 JST 2005 Fri Oct 14 22:58:24 JST 2005
```

**Note:** Even if a background copy process is running, when the FlashCopy relationship is established and the bitmap of the source volume is created, you can access the source volume and the target volume for read and write. As soon as you receive a message stating that a FlashCopy pair is created successfully, move to the next step, which is to perform an IPL of the target server from these volumes.

To end or terminate a FlashCopy relationship, run the following command:

```
dscli rmflash -quiet -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID>
```

To suppress the confirmation message panel, omit the **-quiet** option.

Example 6-4 shows the termination of the FlashCopy sessions for volume pairs with source volumes 0100, 0101, and 0102 and target volumes 0200, 0201, and 0202 with the confirmation prompt.

*Example 6-4 Output of the rmflash command*

---

```
dscli> rmflash 0100-0102:0200-0202
Date/Time: November 1, 2005 1:43:11 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00144W rmflash: Are you sure you want to remove the FlashCopy pair 0100-0102:0200-0202:? [y/n]:y
CMUC00140I rmflash: FlashCopy pair 0100:0200 successfully removed.
CMUC00140I rmflash: FlashCopy pair 0101:0201 successfully removed.
CMUC00140I rmflash: FlashCopy pair 0102:0202 successfully removed.
```

---

## 6.2.4 Performing an IPL of the target server

After the FlashCopy pair relationship is established, you can perform an IPL of the target server. Keep in mind that the target server is a complete clone of the source server. Therefore, you must use *extreme care* when activating the target server. In particular, ensure that it does not connect to the network automatically because this automatic connection causes substantial problems within the target server and the source server.

Perform the IPL of the target server in restricted state, and do not attach the target server to your network until you resolve the potential conflicts the target server might have with the source server. For a CL scripting automation solution using a modified i5/OS start up script to prevent any such conflicts with a cloned system refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

The implementation example that we discuss in this section assumes that you have a 5250 console session open in the HMC. To perform the IPL of the target server in a restricted state, activate the partition with the manual mode set.

On the HMC:

1. In the HMC GUI, under Navigation Area, select **System Management** → **Servers** and select the server with which you want to work.
2. In the contents of server window, select the target partition that you want to activate.
3. Select **Tasks** → **Operations** → **Activate** from the drop-down menu (as shown in Figure 6-2).

**Note:** For below HMC V7, from the navigation area select **Management Environment** → **Server and Partition Server Management** → **Server management**. Then, in the Server and Partition: Server Management panel window, right-click the partition name of the target system, select **Activate**, and choose the profile that you want to use to activate the partition.

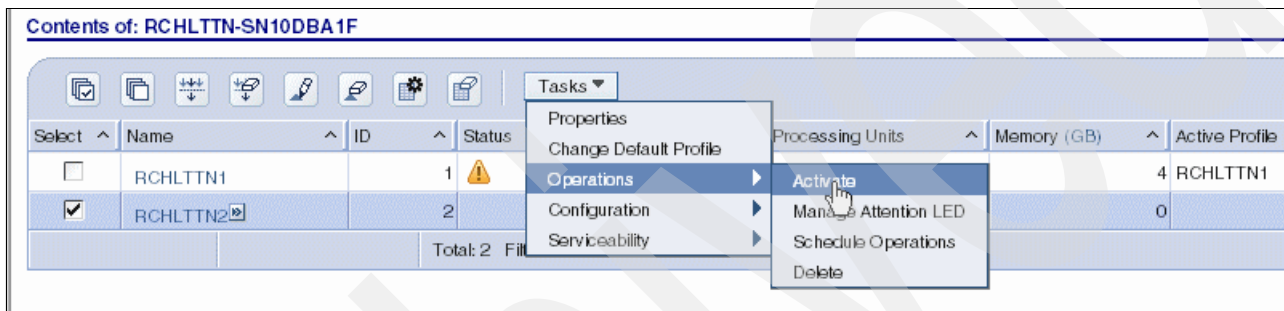


Figure 6-2 Activating the target partition

4. In the Activate Logical Partition window, select the partition profile to activate, and click **Advanced**, as shown in Figure 6-3.

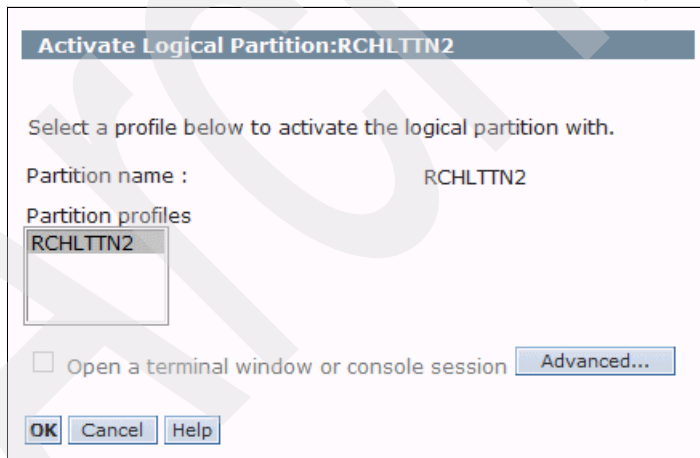


Figure 6-3 Selecting the profile to change

5. In the Activate Logical Partition - Advanced window, for the Keylock position, select **Manual**. For IPL type, select **B: IPL from the second side of the load source**. Then, click **OK** to set the activation setting. See Figure 6-4.

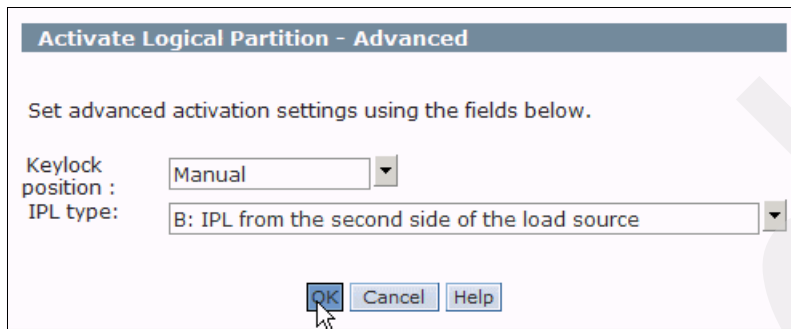


Figure 6-4 Changing the IPL type setting

6. The HMC shows a status dialog box that closes when the task is complete and the partition begins its activation. Wait for the login panel in the 5250 console session.

**Note:** A 5250 console opens only when working locally with the HMC or by Telnet to port 2300.

7. When the system has performed an IPL in manual mode, continue the IPL process by following the steps in the console panel.
8. If the source system has a mirrored load source volume and this is the *first* IPL of the target server, you might see a warning message during the manual IPL process stating that a Disk Configuration Error exists. If you do not see this message, proceed to the next step.

9. During the IPL process, the System Licensed Internal Code (SLIC) storage management task checks the mirror state of the load source. The *mirror state* information is maintained in three places on each of the load sources and in the vital product data (VPD) on the service processor on the central processor complex (CPC).

SLIC checks the mirror state in the following order (Figure 6-5):

- The load source
- The mirrored load source
- The CPC VPD

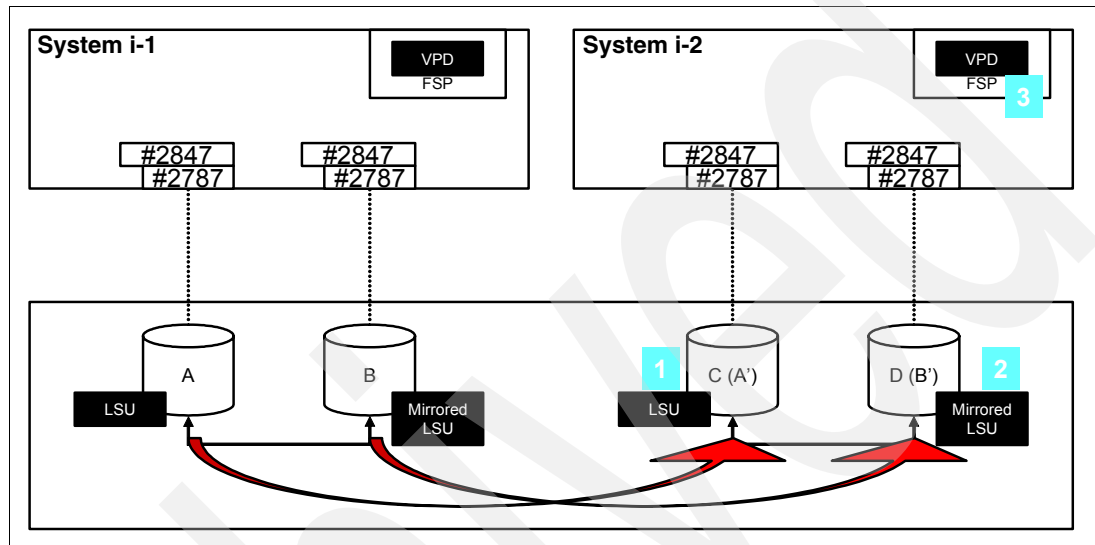


Figure 6-5 Mirror state check sequence

When you perform an IPL of the target server the *first* time, the VPD on the service processor on the target server does not have any information about the load source and the mirrored load source. Although SLIC knows the disk unit D has the mirrored load source volume, the SLIC storage management cannot determine whether disk unit D is the correct mirrored load source unit. Therefore, SLIC displays the report shown in Figure 6-6. After IPL is performed on the target system, the VPD has the mirror state information. The message cannot be seen unless the volume D or the connection is lost.

If you see the Disk Configuration Error Report (Figure 6-6), continue the IPL to bypass the warning message. Select **5=Display Detailed Report**, and press Enter.

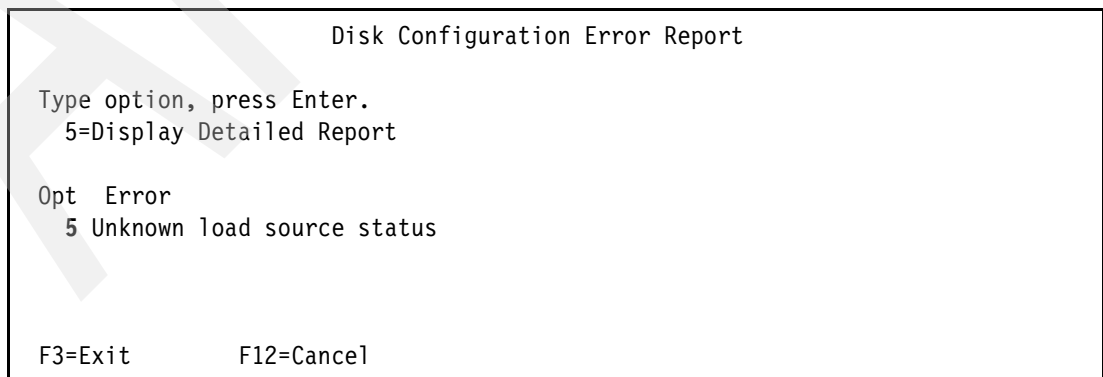


Figure 6-6 Disk Configuration Error Report

10. In the Display Unknown Mirrored Load Source Status panel, press Enter to continue and then press F3 to exit. See Figure 6-7.




Figure 6-7 Display Unknown Mirrored Load Source Status panel

11. If the source system has multipath volumes, you see a message that implies disk configuration attention during the manual IPL process. If you do not see this message, proceed to the next step.

When all the disk units of the system are copied, the path information of each disk unit is also copied. If the source system has disk units in a multipath arrangement, the multipath information is also available on the target server. Because multipath information is associated with physical hardware resource, the target system SLIC cannot detect the old path that is associated with the old physical hardware resource. Therefore, SLIC detects and configures the new paths through the new physical hardware resource, as shown in Figure 6-8.

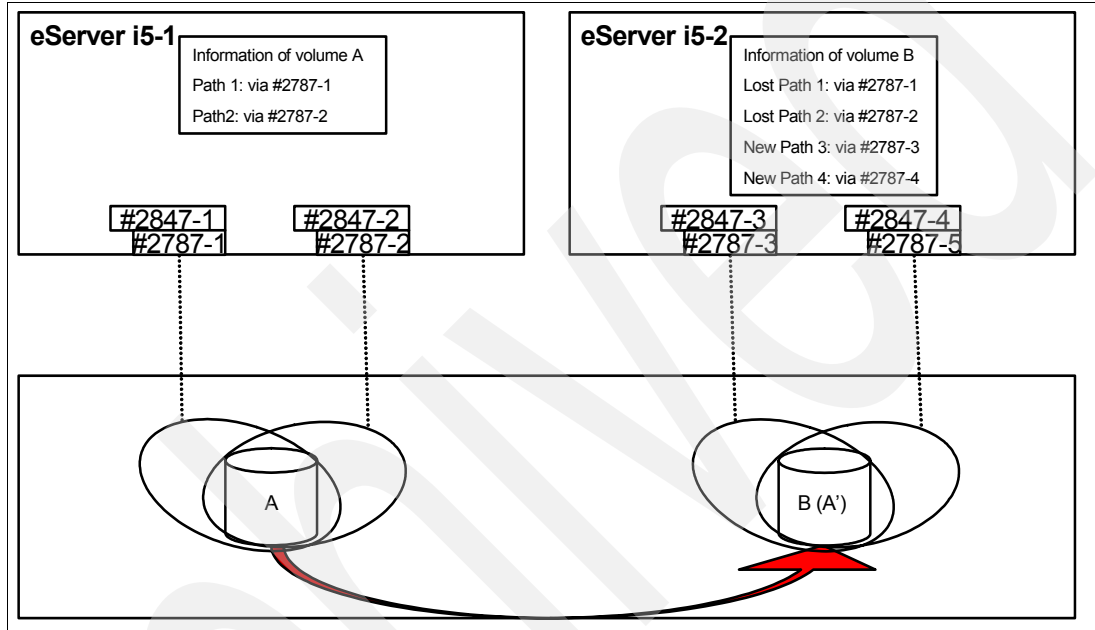


Figure 6-8 Multiple configuration on the source and target servers

In the Disk Configuration Attention Report display, press F10 as indicated by the message to accept the problem and continue (Figure 6-9). Later, you can reset the multipath information from Dedicated Service Tools (DST) or System Service Tools (SST). For more information, refer to 5.13, “Resetting a lost multipath configuration” on page 242.

Then, select **5=Display Detailed Report**.

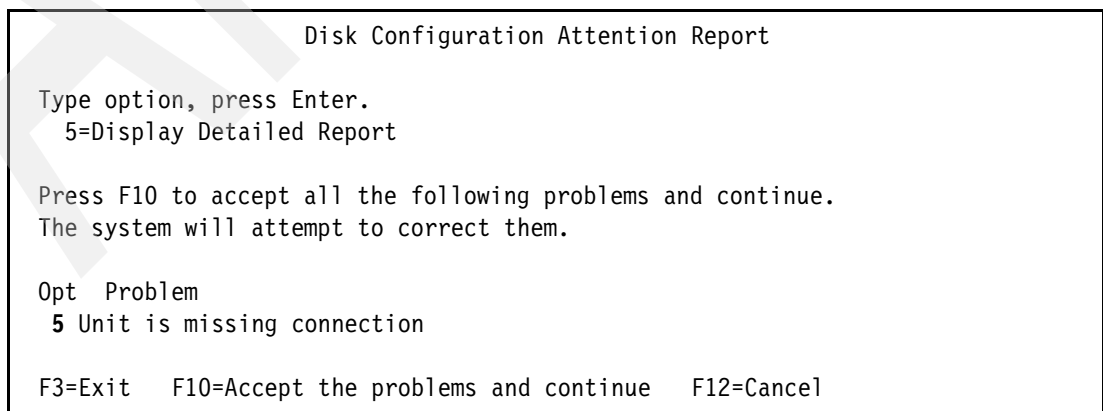


Figure 6-9 Disk Configuration Attention Report panel



12. The detailed Display Disk Units Causing Missing Connection report displays as shown in Figure 6-10. Press Enter to accept the configuration.

```
Display Disk Units Causing Missing Connection

One or more of the functional connections to a disk
unit have not been detected. The connections to the
disk unit were established by running ESS Specialist.

If you use the iSeries in this state, you may cause a
loss of data. The actual number of connections (paths)
detected and the number of connections expected are
listed below.

          Serial
ASP  Unit  Type  Model  Number      Actual  Expected
  1    2  1750  A05   30-0202C68    2        4

Press Enter to accept the configuration and continue.
```

Figure 6-10 Missing disk units report

13. In the IPL or Install the System panel, perform an IPL, install the operating system, or use DST by selecting one of the options that is provided (Figure 6-11). Under Selection, select **1. Perform an IPL**. Alternatively, you can perform an IPL to the operating system from the DST menu.

```
IPL or Install the System                                System:  RCHLTTN2

Select one of the following:

  1. Perform an IPL
  2. Install the operating system
  3. Use Dedicated Service Tools (DST)
  4. Perform automatic installation of the operating system
  5. Save Licensed Internal Code

Selection
  1
```

Figure 6-11 Performing an IPL of the target system

14. When the system is performing an IPL, the Licensed Internal Code IPL in Progress panel shows the status of this process. Wait for the next panel to open, as shown in Figure 6-12.

In this implementation example, the Previous system end parameter indicates a normal IPL. This parameter is set to *Normal* only if the source server was turned off before initiating the FlashCopy. If a FlashCopy of a system quiesced using the CHGASPACT command was taken, the indicator shows *Abnormal*; however, the IPL process does not take much longer because all database I/O activity was quiesced.

```

Licensed Internal Code IPL in Progress
10/14/05 14:40:36
IPL:
Type . . . . . : Attended
Start date and time . . . . . : 10/14/05 14:40:36
Previous system end . . . . . : Normal
Current step / total . . . . . : 1 16
Reference code detail . . . . . : C6004050

IPL step
>Storage Management Recovery          Time Elapsed  Time Remaining
Start LIC Log                          00:00:00
Main Storage Dump Recovery
Trace Table Initialization
Context Rebuild

Item:
Current / Total . . . . . :

Sub Item:
Identifier . . . . . :
Current / Total . . . . . :

```

Figure 6-12 IPL progress

15. Sign in to the operating system. Remember that the user profile and password that you use is the same as on the source system.

16. During the operating system IPL, the IPL Options panel displays, as shown in Figure 6-13. To boot the operating system in a restricted state, for the Start system to restricted state parameter, type Y.

```

                                IPL Options

Type choices, press Enter.

System date . . . . . 10 / 14 / 05  MM / DD / YY
System time . . . . . 14 : 43 : 26  HH : MM : SS
System time zone . . . . . Q0000UTC  F4 for list
Clear job queues . . . . . N        Y=Yes, N=No
Clear output queues . . . . . N     Y=Yes, N=No
Clear incomplete job logs . . . . . N Y=Yes, N=No
Start print writers . . . . . Y     Y=Yes, N=No
Start system to restricted state . . . . . Y  Y=Yes, N=No

Set major system options . . . . . N  Y=Yes, N=No
Define or change system at IPL . . . . . N Y=Yes, N=No

Last power-down operation was NORMAL

```

Figure 6-13 Setting Start system to restricted state

17. When the system has performed the IPL, change the following settings, depending on your requirements:

- System Name and Network Attributes
- TCP/IP settings
- Backup Recovery and Media Services (BRMS) Network information
- IBM eServer iSeries NetServer™ settings
- User profiles and passwords
- Job Schedule entries
- Relational Database entries

18. After you resolve the potential conflicts with the source server, attach the target to the network and restart the operating system.

**Note:** If you plan to use this clone to create backups using BRMS, refer to 15.2, “Using BRMS and FlashCopy” on page 436 to make sure the save information is reflected back to the source system.

## 6.3 Implementing space efficient FlashCopy

The new space efficient FlashCopy licensed function introduced with DS8000 R3 allows the amount of physical space allocated to the target volumes to be thinly-provisioned as proportional to the amount of write activity on the FlashCopy source and target. At the same time, it provides all of the functions available with the fully-provisioned traditional FlashCopy except that using it with background-copy makes no sense. For a detailed description of space efficient FlashCopy, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

In this section, we describe the steps to set up space efficient FlashCopy for system backup of a System i partition, including the monitoring actions that are required to prevent you from running out of space on the repository volume used as the backstore for physical storage capacity of the space efficient FlashCopy target volumes.

The space efficient FlashCopy target volumes on the DS8000 are defined of the same capacity and protection as System i production volumes. For our example a stand-by backup System i partition is connected to FlashCopy space efficient (SE) target volumes using boot from SAN (see Figure 6-14). The backup partition usually resides on the same System i server as the production partition but it can also be on a separate System i server.




Figure 6-14 Space efficient FlashCopy example setup

### 6.3.1 Configuring space efficient FlashCopy for a System i environment

Perform the following steps to configure space efficient FlashCopy between a System i production and backup partition:

1. Thoroughly plan the DS8000 storage configuration layout for your System i production volumes and the space efficient FlashCopy target volumes for your backup partition. To optimize performance locate your source and target volumes in different extent pools but within the same rankgroup and to prevent running out of space for the repository perform a

careful sizing for your required repository capacity (see 4.2.8, “Sizing for space efficient FlashCopy” on page 104).

In our example, we set up the production System i volumes (FlashCopy SE sources) and FlashCopy SE targets in 4 extent pools, each of them containing two RAID5 ranks in DS8000. Two extent pools belong to rankgroup 0, and two belong to rankgroup 1. We define both FlashCopy SE source and target LUNs in each extent pool. The source volumes have corresponding targets in another extent pool which for best performance belongs to the *same* rankgroup such as the extent pool with the source volumes, as shown in Figure 6-15.

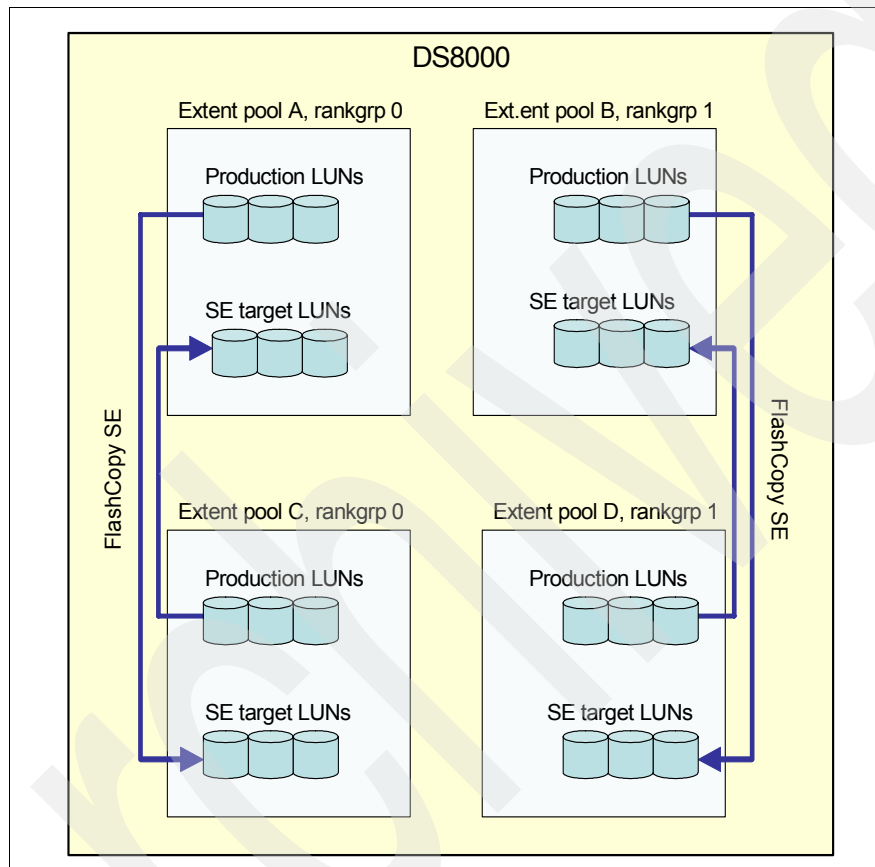


Figure 6-15 Layout of LUNs for FlashCopy SE

2. Define extent pools and LUNs for the production System i partition, create FlashCopy SE repository, and create FlashCopy SE LUNs. For information about how to perform the DS8000 logical storage configuration, for example for the extent pools and LUNs, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

To create the FlashCopy SE repository through DS CLI, use the **mksestg** command with the following parameters:

- |                  |                                                                           |
|------------------|---------------------------------------------------------------------------|
| <b>-extpool</b>  | Extent pool of repository                                                 |
| <b>-capytype</b> | (Optional) Denotes the type of specified capacity (GB, cylinders, blocks) |
| <b>-viricap</b>  | The amount of virtual capacity                                            |
| <b>-repcap</b>   | The physical capacity of the repository                                   |

Figure 6-16 shows an example of creating the repository storage.

```
dsccli> mksestg -extpool P14 -capytype gb -vircap 282 -repcap 70
Date/Time: October 30, 2007 2:17:46 PM CET IBM DSCLI Version: 5.3.0.977 DS:
IBM.2107-7520781
CMUC00342I mksestg:: The space-efficient storage for the extent pool P14 has
been created successfully.
dsccli> mksestg -extpool P15 -capytype gb -vircap 282 -repcap 70
Date/Time: October 30, 2007 2:18:21 PM CET IBM DSCLI Version: 5.3.0.977 DS:
IBM.2107-7520781
CMUC00342I mksestg:: The space-efficient storage for the extent pool P15 has
been created successfully.
dsccli> mksestg -extpool P34 -capytype gb -vircap 282 -repcap 70
Date/Time: October 30, 2007 2:19:03 PM CET IBM DSCLI Version: 5.3.0.977 DS:
IBM.2107-7520781
CMUC00342I mksestg:: The space-efficient storage for the extent pool P34 has
been created successfully.
dsccli> mksestg -extpool P47 -capytype gb -vircap 282 -repcap 70
Date/Time: October 30, 2007 2:19:41 PM CET IBM DSCLI Version: 5.3.0.977 DS:
IBM.2107-7520781
CMUC00342I mksestg:: The space-efficient storage for the extent pool P47 has
been created successfully.
```

Figure 6-16 Creating FlashCopy SE repository

To define space efficient logical volumes for System i using DS CLI, add the parameter **-sam tse** to the command **mkfbvol**, which you use to create System i LUNs.

**Note:** In this command, **sam** stands for *storage allocation method* and **tse** denotes *track space efficient*.

An example of this type of DS CLI command is as follows:

```
mkfbvol -extpool p14 -os400 A05 -sam tse -name Vol_SE_#h 1009-100f
```

In our example, we create four extent pools with 8 \* 35 GB System i LUNs each. We create a repository of 70 GB in each extent pool, which gives a total of 280 GB repository capacity for 1125 GB of production capacity. (For information about how to calculate the repository capacity, refer to 4.2.8, “Sizing for space efficient FlashCopy” on page 104.)

Next, we define 8 \* 35 GB space efficient LUNs in each extent pool to be used as FlashCopy SE targets. Figure 6-17 shows part of the display obtained by DS CLI command `lsfbvol` for one of our extent pools. Observe the `sam` column output showing standard and space efficient LUNs in our pool.

```
*dscli> lsfbvol -l -extpool p14
Date/Time: October 29, 2007 8:28:15 PM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
Name          ID   acstate  datastate  configstate  deviceMTM  datatype  extpool  sam      capttype
=====
ITSO_St_LS_1000 1000 Online   Normal    Normal      2107-A85   FB 520U   P14     Standard iSeries
ITSO_St_1001   1001 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1002   1002 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1003   1003 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1004   1004 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1005   1005 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1006   1006 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_St_1007   1007 Online   Normal    Normal      2107-A05   FB 520P   P14     Standard iSeries
ITSO_SE_LS_1008 1008 Online   Normal    Normal      2107-A85   FB 520U   P14     TSE      iSeries
ITSO_SE_1009   1009 Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100A   100A Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100B   100B Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100C   100C Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100D   100D Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100E   100E Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
ITSO_SE_100F   100F Online   Normal    Normal      2107-A05   FB 520P   P14     TSE      iSeries
```

Figure 6-17 System i standard and SE LUNs

3. Set up your System i production partition with all disk space on DS8000 and external load source connected using boot from SAN.

For more information about setting up a System i partition with external storage, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

In our example, we set up a System i partition using 32 \* 35 GB LUNs on DS8000 connected through four System i FC adapters. Two of the adapters are attached using boot from SAN #2847 IOPs. The external load source is connected using boot from SAN IOP and, because we do not use i5/OS V6R1 with multipath support for the load source, it is mirrored to another external LUN. All the other LUNs are connected using multipath.

Figure 6-18 shows some of the production LUNs as seen from System i System Service Tools (SST).

Display Disk Configuration Status						
ASP	Unit	Serial Number	Type	Model	Resource Name	Status
1						Mirrored
	1	50-1000781	2107	A85	DD019	Active
	1	50-1208781	2107	A85	DD020	Active
	14	50-120A781	2107	A05	DMP143	RAID-5/Active
	15	50-1304781	2107	A05	DMP195	RAID-5/Active
	16	50-1005781	2107	A05	DMP137	RAID-5/Active
	17	50-1508781	2107	A05	DMP191	RAID-5/Active
	18	50-150F781	2107	A05	DMP185	RAID-5/Active
	19	50-150B781	2107	A05	DMP183	RAID-5/Active
	20	50-1302781	2107	A05	DMP172	RAID-5/Active
	21	50-1004781	2107	A05	DMP159	RAID-5/Active
	22	50-1307781	2107	A05	DMP197	RAID-5/Active
	23	50-120B781	2107	A05	DMP109	RAID-5/Active
	24	50-150D781	2107	A05	DMP173	RAID-5/Active

More...

Press Enter to continue.

F3=Exit            F5=Refresh            F9=Display disk unit details  
 F11=Disk configuration capacity    F12=Cancel

Figure 6-18 System i production volumes

On the System i HMC, we tag as IPL device the FC adapter to which the external load source is connected and which is attached to the #2847 boot from SAN IOP.

### 6.3.2 Using space efficient FlashCopy with i5/OS

To use the space efficient FlashCopy volumes of the entire System i disk space for backup, perform the following steps:

1. If BRMS is used for backup, run the required commands to integrate BRMS with FlashCopy as described in 15.2, "Using BRMS and FlashCopy" on page 436.
2. Turn off the System i partition using the i5/OS command PWRDWNSYS or use the new i5/OS V6R1 quiesce for Copy Services function with the new CHGASPACT command (see 15.1, "Using i5/OS quiesce for Copy Services" on page 432).
3. Establish space efficient FlashCopy from fully-provisioned production LUNs to space efficient target LUNs, by using DS CLI `mkflash` command with the following parameters:

**-tgtse**            Denotes Space efficient target LUNs  
**-nocp**            Using no-copy to prevent background copy



Figure 6-19 shows the DS CLI commands that we used for space efficient FlashCopy of our production volumes to space efficient volumes.

```
dscli> mkflash -tgtse -nocp 1000-1007:1200-1207
Date/Time: October 30, 2007 10:13:40 AM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
CMUC00137I mkflash: FlashCopy pair 1000:1200 successfully created.
CMUC00137I mkflash: FlashCopy pair 1001:1201 successfully created.
CMUC00137I mkflash: FlashCopy pair 1002:1202 successfully created.
CMUC00137I mkflash: FlashCopy pair 1003:1203 successfully created.
CMUC00137I mkflash: FlashCopy pair 1004:1204 successfully created.
CMUC00137I mkflash: FlashCopy pair 1005:1205 successfully created.
CMUC00137I mkflash: FlashCopy pair 1006:1206 successfully created.
CMUC00137I mkflash: FlashCopy pair 1007:1207 successfully created.
dscli> mkflash -tgtse -nocp 1208-120f:1008-100f
Date/Time: October 30, 2007 10:14:03 AM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
CMUC00137I mkflash: FlashCopy pair 1208:1008 successfully created.
CMUC00137I mkflash: FlashCopy pair 1209:1009 successfully created.
CMUC00137I mkflash: FlashCopy pair 120A:100A successfully created.
CMUC00137I mkflash: FlashCopy pair 120B:100B successfully created.
CMUC00137I mkflash: FlashCopy pair 120C:100C successfully created.
CMUC00137I mkflash: FlashCopy pair 120D:100D successfully created.
CMUC00137I mkflash: FlashCopy pair 120E:100E successfully created.
CMUC00137I mkflash: FlashCopy pair 120F:100F successfully created.
dscli> mkflash -tgtse -nocp 1300-1307:1500-1507
Date/Time: October 30, 2007 10:14:20 AM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
CMUC00137I mkflash: FlashCopy pair 1300:1500 successfully created.
CMUC00137I mkflash: FlashCopy pair 1301:1501 successfully created.
CMUC00137I mkflash: FlashCopy pair 1302:1502 successfully created.
CMUC00137I mkflash: FlashCopy pair 1303:1503 successfully created.
CMUC00137I mkflash: FlashCopy pair 1304:1504 successfully created.
CMUC00137I mkflash: FlashCopy pair 1305:1505 successfully created.
CMUC00137I mkflash: FlashCopy pair 1306:1506 successfully created.
CMUC00137I mkflash: FlashCopy pair 1307:1507 successfully created.
dscli> mkflash -tgtse -nocp 1508-150f:1308-130f
Date/Time: October 30, 2007 10:14:32 AM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
CMUC00137I mkflash: FlashCopy pair 1508:1308 successfully created.
CMUC00137I mkflash: FlashCopy pair 1509:1309 successfully created.
CMUC00137I mkflash: FlashCopy pair 150A:130A successfully created.
CMUC00137I mkflash: FlashCopy pair 150B:130B successfully created.
CMUC00137I mkflash: FlashCopy pair 150C:130C successfully created.
CMUC00137I mkflash: FlashCopy pair 150D:130D successfully created.
CMUC00137I mkflash: FlashCopy pair 150E:130E successfully created.
```

Figure 6-19 Make FlashCopy SE

4. IPL the System i backup partition, which is connected to the space efficient FlashCopy target volumes, by activating the partition from the HMC. Make sure the boot from SAN I/O adapter (IOA) is tagged to which the external load source is connected.

The IPL of the System i backup partition brings up a *clone* of the System i production partition with disk space on space efficient FlashCopy target LUNs.

**Note:** Usually, you need to change the device descriptions and network attributes in the IPLed clone partition. If you use BRMS, then you also need to use commands for the integration of BRMS and FlashCopy. You can automate the procedure to do this. For more information, see *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

In our example the backup partition connects to the space efficient FlashCopy targets as shown in Figure 6-19 on page 273. In i5/OS these space efficient FlashCopy target LUNs are seen as regular System i disk units like shown in Figure 6-20. Notice the LUN ID, which is contained in the disk unit serial number characters 4 to 7.

Display Disk Configuration Status						
ASP	Unit	Serial Number	Type	Model	Resource Name	Status
1						Mirrored
	1	50-1200781	2107	A85	DD019	Active
	1	50-1008781	2107	A85	DD020	Active
	14	50-100A781	2107	A05	DMP143	RAID-5/Active
	15	50-1504781	2107	A05	DMP195	RAID-5/Active
	16	50-1205781	2107	A05	DMP137	RAID-5/Active
	17	50-1308781	2107	A05	DMP191	RAID-5/Active
	18	50-130F781	2107	A05	DMP185	RAID-5/Active
	19	50-130B781	2107	A05	DMP183	RAID-5/Active
	20	50-1502781	2107	A05	DMP172	RAID-5/Active
	21	50-1204781	2107	A05	DMP159	RAID-5/Active
	22	50-1507781	2107	A05	DMP198	RAID-5/Active
	23	50-100B781	2107	A05	DMP109	RAID-5/Active
	24	50-130D781	2107	A05	DMP173	RAID-5/Active

More...

Press Enter to continue.

F3=Exit            F5=Refresh            F9=Display disk unit details  
F11=Disk configuration capacity    F12=Cancel

Figure 6-20 FlashCopy SE targets in backup System i partition

5. To keep the occupation of the repository on a minimum level, remove the FlashCopy relationships after the backup completes, which releases all space that is allocated in the repository for the targets. Use the DS CLI `rmflash` command with the `-tgtreleasespace` parameter, as shown in Figure 6-21.

```
dsccli> rmflash -tgtreleasespace 1000-1007:1200-1207
Date/Time: November 1, 2007 11:48:52 AM CET IBM DSCLI Version: 5.3.0.977 DS:
IBM.2107-7520781
CMUC00144W rmflash: Are you sure you want to remove the FlashCopy pair
1000-1007:1200-1207:? [y/n]:y
CMUC00140I rmflash: FlashCopy pair 1000:1200 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1001:1201 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1002:1202 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1003:1203 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1004:1204 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1005:1205 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1006:1206 successfully removed.
CMUC00140I rmflash: FlashCopy pair 1007:1207 successfully removed.
```

Figure 6-21 Releasing repository space

### 6.3.3 Reactions of System i partitions to a full repository

If a proper sizing was done for your space efficient FlashCopy repository capacity (see 4.2.8, “Sizing for space efficient FlashCopy” on page 104) and if you maintain the occupation of your repository below the anticipated level with performing regular commands for cleaning its allocation, your repository probably will never be full.

However, if the used capacity from the repository reaches the default threshold of 85% and if you have set up Simple Network Management Protocol (SNMP) notifications on your DS8000, you receive a warning through SNMP trap 221. You can change this warning threshold to another value using the DS CLI `chsestg` command with the `-repcapthreshold` parameter.

For example, you can set the threshold value to 70% for extent pool P14 using the `chsestg -repcapthreshold 70 P14` command.

Figure 6-22 shows an example of SNMP trap 221 for reaching the default 85% threshold of used repository capacity. Notice that the corresponding extent pool is reported in hexadecimal notation and that the *Percentage full* information shows the *remaining* percentage of capacity until the repository is full. In our example, *Percentage full* shows 15%, meaning that 85% of the repository space is currently allocated.

```

0000 30 82 01 2D 02 01 00 04 06 70 75 62 6C 69 63 A4 0...public.
0010 82 01 1E 06 09 2B 06 01 04 01 02 06 81 02 40 04 .....+.....@.
0020 09 9B 3E 61 02 01 06 02 02 00 DD 43 04 06 37 E5 ...>a.....C..7.
0030 DB 30 81 FD 30 81 FA 06 0B 2B 06 01 04 01 02 06 .0..0....+.....
0040 81 02 03 02 04 81 EA 32 30 30 37 2F 31 31 2F 30 .....2007/11/0
0050 36 20 31 37 3A 32 31 3A 33 38 20 43 45 54 0A 53 6 17:21:38 CET.S
0060 70 61 63 65 20 45 66 66 69 63 69 65 6E 74 20 52 pace Efficient R
0070 65 70 6F 73 69 74 6F 72 79 20 6F 72 20 4F 76 65 epository or Ove
0080 72 2D 70 72 6F 76 69 73 69 6F 6E 65 64 20 56 6F r-provisioned Vo
0090 6C 75 6D 65 20 68 61 73 20 72 65 61 63 68 65 64 lume has reached
00A0 20 61 20 77 61 72 6E 69 6E 67 20 77 61 74 65 72 a warning water
00B0 6D 61 72 6B 0A 55 4E 49 54 3A 20 4D 6E 66 20 54 mark.UNIT: Mnf T
00C0 79 70 65 2D 4D 6F 64 20 53 65 72 69 61 6C 4E 6D ype-Mod SerialNm
00D0 0A 20 20 20 20 20 20 49 42 4D 20 32 31 30 37 2D . IBM 2107-
00E0 39 32 32 20 37 35 2D 32 30 37 38 31 0A 56 6F 6C 922 75-20781.Vol
00F0 75 6D 65 20 54 79 70 65 3A 20 30 0A 52 65 61 73 ume Type: 0.Reas
0100 6F 6E 20 63 6F 64 65 3A 20 30 0A 45 78 74 65 6E on code: 0.Extent
0110 74 20 50 6F 6F 6C 20 49 44 3A 20 32 32 0A 50 65 t Pool ID: 22.Pe
0120 72 63 65 6E 74 61 67 65 20 46 75 6C 6C 3A 20 31 rcentage Full: 1
0130 35 5

```

```

Frame Length: 305 bytes
Community: public
OID: .1.3.6.1.4.1.2.6.130
Address: 9.155.62.97
sysUpTime: 12 days, 01:47:46
Generic: 6 - Enterprise Specific
Specific: 221
OID: .1.3.6.1.4.1.2.6.130.3.2
ASN1 Type: Octet String 0x04 (4)
Value: 2007/11/06 17:21:38 CET|Space Efficient Repository or Over-provisioned
75-20781|Volume Type: 0|Reason code: 0|Extent Pool ID: 22|Percentage Full: 15

```

Figure 6-22 Repository threshold reached SNMP trap 221 warning

For more information about configuring and using SNMP notifications with DS8000 refer to *IBM System Storage DS8000 Series: Architecture and Implementation*, SG24-6786, which is available at:

<http://www.redbooks.ibm.com/abstracts/sg246786.html?Open>

If for any reason the repository runs out of space during the space efficient FlashCopy relationship, the FlashCopy relationship fails as shown by the DS CLI command outputs in Figure 6-23. In this case, the source volumes remain fully accessible for both reads and writes so the System i production partition continues to run without any failure. However, the System i backup partition that uses the disk space of the space efficient FlashCopy target volumes fails. In our case, it fails with SRC A6020266 entering a freeze state as soon as the FlashCopy relation fails at a fully occupied repository.

```

dscli> lssestg -l
Date/Time: November 6, 2007 5:31:05 PM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
extentpoolID stgtype datastate configstate repcapstatus %repcapthreshold repcap (2^30B) vircap repcapalloc vircapalloc
=====
P14          fb      Normal   Normal   below      0          70.0  282.0    70.0    264.0
P15          fb      Normal   Normal   below      0          70.0  282.0    70.0    264.0
P34          fb      Normal   Normal   below      0          70.0  282.0    70.0    264.0
P47          fb      Normal   Normal   below      0          70.0  282.0    70.0    264.0
dscli> lsflash 1000-15ff
Date/Time: November 6, 2007 5:31:16 PM CET IBM DSCLI Version: 5.3.0.977 DS: IBM.2107-7520781
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible SourceWriteEnabled TargetWriteEnabled
Backgrou
=====
1000:1200 - - - - - - - - - -
1001:1201 - - - - - - - - - -
1002:1202 - - - - - - - - - -
1003:1203 - - - - - - - - - -
1004:1204 - - - - - - - - - -
1005:1205 - - - - - - - - - -
1006:1206 - - - - - - - - - -
1209:1009 - - - - - - - - - -
120A:100A - - - - - - - - - -
120B:100B - - - - - - - - - -
120C:100C - - - - - - - - - -
120D:100D - - - - - - - - - -
120E:100E - - - - - - - - - -
1300:1500 - - - - - - - - - -
1301:1501 - - - - - - - - - -
1302:1502 - - - - - - - - - -
1303:1503 - - - - - - - - - -
1304:1504 - - - - - - - - - -
1305:1505 - - - - - - - - - -
1306:1506 - - - - - - - - - -
1508:1308 - - - - - - - - - -
1509:1309 - - - - - - - - - -
150A:130A - - - - - - - - - -
150B:130B - - - - - - - - - -
150C:130C - - - - - - - - - -
150D:130D - - - - - - - - - -
150E:130E - - - - - - - - - -

```

Figure 6-23 Space efficient FlashCopy relationships failed due to full repository

Archived



## Implementing Metro Mirror using the DS CLI

In this chapter, we explain how to implement the IBM System Storage Copy Services Metro Mirror function using the DS command-line interface (CLI).

## 7.1 Overview of the test environment

In this chapter, we describe how to implement Metro Mirror with the entire DASD space of the System i5 environment. Before we begin that discussion, we provide an overview of the test environment in this section. Figure 7-1 shows an overview of the test environment that we use for this implementation example.

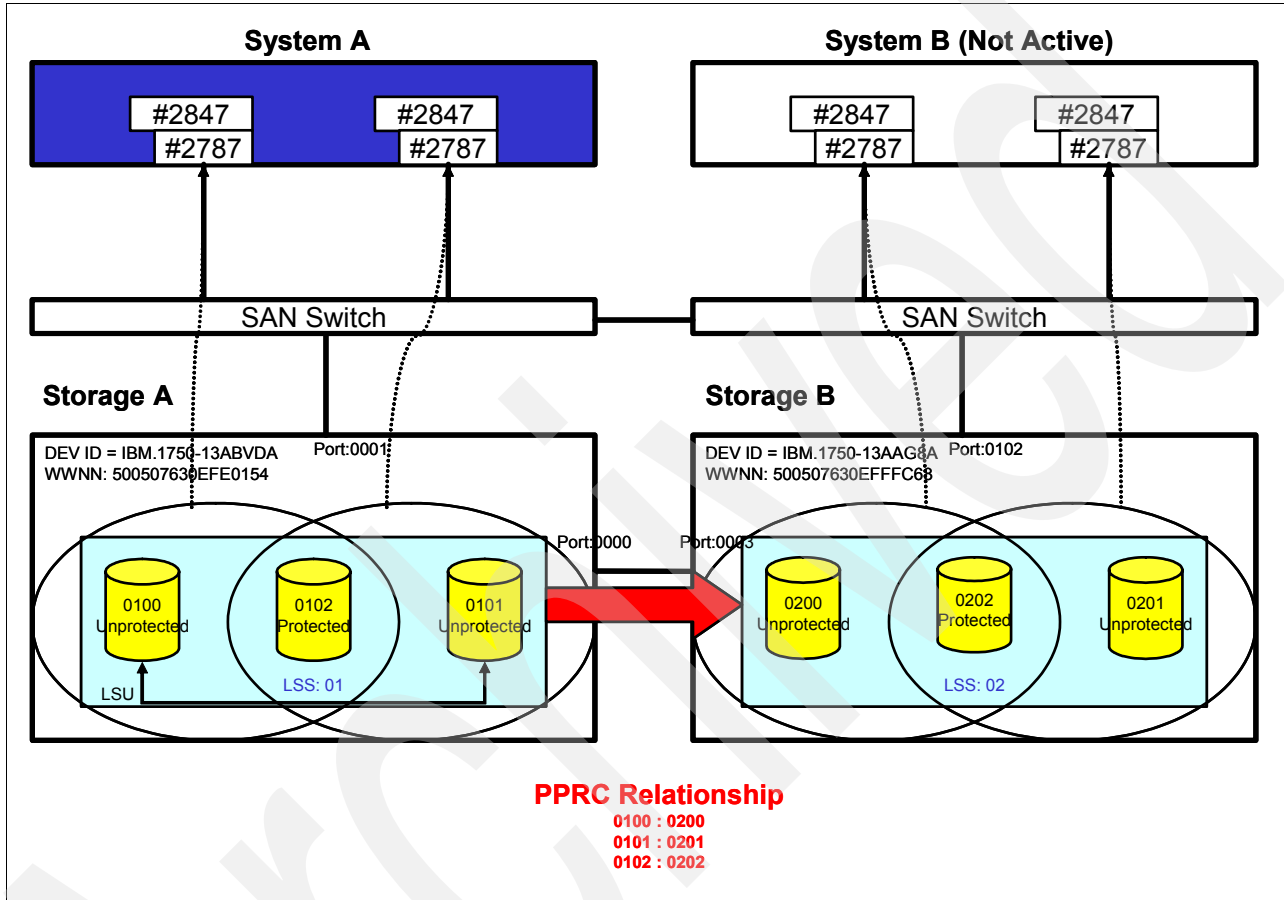


Figure 7-1 Test environment

System A is the production server, and system B is the backup server. System A is connected to storage A. An IPL is performed on System A from the storage area network (SAN) with 2847 I/O processor (IOP) having a load source unit in volume 0100 and a mirrored load source unit in volume 0101. The first 2847 feature card is tagged as the load source in the Hardware Management Console (HMC).

**Note:** If you are using i5/OS V6R1 or later, you need to use the new multipath load source support instead of mirroring the load source for providing path protection.

Storage A and storage B are connected with Fibre Channel (FC) cables. This implementation example assumes that system A and storage A are on local site A and that system B and storage B are on remote site B.

In this example, the Metro Mirror environment is created between storage A and storage B. The business application is then switched from local site A to remote site B. Finally, the business application is switched back from remote site B to local site A.



Implementation of Metro Mirror for the entire DASD space involves the following tasks:

1. Create a Metro Mirror environment:
  - a. Create the Peer-to-Peer Remote Copy (PPRC) paths.
  - b. Create the Metro Mirror relationships.
2. Switch the system from the local site to the remote site:
  - a. Make the volumes available on the remote site.
  - b. Perform the IPL of the backup server on the remote site.
3. Switch back the system from the remote site to the local site:
  - a. Start Metro Mirror in the reverse direction from the remote site to the local site.
  - b. Make the volumes available on the local site.
  - c. Perform the IPL of the production server on the local site.
  - d. Start Metro Mirror in the original direction from the local site to the remote site.

We describe these steps in detail in the remaining sections of this chapter.

## 7.2 Creating a Metro Mirror environment setup

Before you create a Metro Mirror environment, you must complete the following actions:

1. Create the volumes for the backup server.
2. Create volume groups and host connections to perform the IPL of the backup server.
3. To perform an IPL of the backup server from a copied load source unit on external storage, set the tagged load source unit to the 2847 IOP or the Fibre Channel I/O adapter (IOA) that is connected to the copied load source unit in the HMC.

You need to perform these tasks only the first time that you set up the Metro Mirror environment. Change the settings only if the source environment changes. Our example assumes that you have completed these tasks.

**Important:** If you are considering using Metro Mirror, Global Mirror, or FlashCopy for the replication of the load source unit or other i5/OS disk units within the same DS6000 or DS8000 or between two or more DS6000 or DS8000 systems, the source volume and the target volume characteristics must be identical. The target and source must have matching capacities and matching protection types. For example, a *protected* 35 GB i5/OS volume must be replicated to another *protected* 35 GB volume if it is to be used in a replicated i5/OS configuration. It cannot be replicated to an *unprotected* 35 GB volume or to a volume of any other capacity. If you plan to migrate your load source from one size LUN to a larger size, use the System i *copy disk unit data utility* that is available in Dedicated Service Tools (DST).

In addition, the DS CLI offers the capability to change characteristics, such as the protection type of a previously defined volume. After a volume is assigned to an i5/OS partition and added to that partition's configuration, its characteristics must *not* be changed. If there is a requirement to change a characteristic of a configured volume, you must first remove it completely from the System i ASP configuration. After you make the characteristic changes, for example protection type, capacity, and so on, by removing and recreating the volume or by using the DS CLI, you can reassign the volume into the System i configuration. To simplify the configuration, we recommend that you have a symmetrical configuration between two IBM System Storage solutions, creating the same volumes with the same volume IDs (LSSs and volume numbers).

For FlashCopy, we recommend that you plan the target volumes in order to be assigned on a different rank from where the source volumes are assigned for performance of the source server.

For more information about creating volumes, volume groups, and host connects, as well as tagging load source IOP, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

## 7.2.1 Creating Peer-to-Peer Remote Copy paths

For Metro Mirror, the PPRC paths must exist for every logical subsystem (LSS) between the source LSS, with which the source volumes are associated, and the target LSS, with which the target volumes are associated.

**Important:** When creating the PPRC paths, consider the following points during the planning phase:

- ▶ For performance, use dedicated I/O ports for PPRC; do *not* share them with the I/O of your servers. Use SAN switch zoning to restrict the server's storage system I/O port usage (see 3.2.5, "Planning for SAN connectivity" on page 67).
- ▶ For redundancy, create at least two PPRC paths between the same LSS. Use each I/O port of different controllers in case of failure or maintenance of one of the controllers. For example, one path can use port I00xx on controller 0, and the other path can use port I01xx on controller 1.

To create a PPRC path:

1. Check which source LSS and target LSS are available on each storage server. Which LSS is associated with which volume depends on the volume ID. Each volume has a four-digit hexadecimal volume ID, such as 0100. The first two digits to the left indicate the LSS ID.

If the volume that you want to define as the source has the volume ID 0100, its LSS ID is 01.

If the volume that you want to define as the target has the volume ID 0200, its LSS ID is 02.

Therefore, to see which source LSS and target LSS are available, look for the ID of the volume for which you want to create the PPRC relationship.

In this example, we run the following DS CLI command on an attached Windows PC:

```
dscli lsfbvol -dev <storage_image_ID>
```

The results of this command show the four-digit hexadecimal volume ID, where the first two digits indicate the LSS ID.

Example 7-1 lists the available fixed block volumes for the source in storage A.

*Example 7-1 Output of the lsfbvol command on system A*

```
dscli> lsfbvol -dev IBM.1750.13-ABVDA
Date/Time: October 26, 2005 6:25:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
rchltn2-boot  0100 Online    Normal    Normal      1750-A85  FB 520U  P1          32.8         35.2
rchltn2-boot-mr 0101 Online    Normal    Normal      1750-A85  FB 520U  P1          32.8         35.2
rchltn2-disk2  0102 Online    Normal    Normal      1750-A05  FB 520P  P1          32.8         35.2
```

Example 7-2 lists the available fixed block volumes for the target in storage B. In this example, the source LSS ID is identified as 01 in storage A, and the target LSS ID is identified as 02 in storage B.

*Example 7-2 Output of the lsfbvol command on system B*

```
dscli> lsfbvol -dev IBM.1750.13-AAG8A
Date/Time: October 26, 2005 6:26:22 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
rlttn2-boot-mmir 0200 Online    Normal    Normal      1750-A85  FB 520U  P0          32.8         35.2
rlttn2-btmmir 0201 Online    Normal    Normal      1750-A85  FB 520U  P0          32.8         35.2
rlttn2-dsk2-mmir 0202 Online    Normal    Normal      1750-A05  FB 520P  P0          32.8         35.2
```

**Symmetrical configuration:** In this example, the different volume IDs for the target volumes are configured from the source volumes to make it easier to understand the specified volume parameter for use in later commands. For your environment, we recommend that you use a symmetrical configuration, in which the target volumes have the same volume IDs as the source volumes, in order to simplify the configuration.

2. Check the worldwide node name (WWNN) of the target storage. The WWNN is unique in every IBM System Storage solution and is a required parameter for the command to create the PPRC paths. Use the following DS CLI command to display the list of storage images with their WWNNs configured in a storage complex:

```
dscli lssi
```

You need to run this command on both your source and target storage system, unless the source storage and the target storage are configured within the same storage complex on your System Management Console (SMC) or HMC so that both WWNNs are displayed in the result.

Example 7-3 lists the WWNN of source storage A and target storage B. In this example, the WWNN of target storage is identified as 500507630EFFF68 from its Storage Unit ID IBM.1750-13AAG8A. Take note of the WWNN of the source storage, 500507630EFE0154. This number is necessary for creating the PPRC path for the reverse direction.

*Example 7-3 Output of the lssi command*

```

dsccli> lssi
Date/Time: November 1, 2005 8:29:03 AM JST IBM DSCLI Version: 5.0.6.142
Name ID          Storage Unit      Model WWNN          State  ESSNet
-----
-   IBM.1750-13AAG8A  IBM.1750-13AAG8A  511   500507630EFFF68  Online Enabled
-   IBM.1750-13ABVDA  IBM.1750-13ABVDA  511   500507630EFE0154 Online Enabled

```

3. Check which I/O ports are available for the PPRC paths between the source LSS and the target LSS. Use the following command:

```

dsccli lsavailpprcport -dev <source storage_image_ID> -remotedev
<target storage_image_ID> -remotewwnn <WWNN of target Storage_image>
<Source_LSS_ID>:<Target_LSS_ID>

```

**Note:** For this command, you can specify any available LSS pair that you want for source and target.

The results of this command shows a list of FC I/O ports that can be defined as PPRC paths. Each row indicates the available I/O ports pair. The local port is a port on the local storage, and the attached port is a port on the remote storage.

Example 7-4 lists the available PPRC ports between storage A on site A and storage B on site B.

*Example 7-4 Output of the lsavailpprcport command on system A*

```

dsccli> lsavailpprcport -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwnn 500507630EFFF68 01:02
Date/Time: October 18, 2005 7:20:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Local Port Attached Port Type
-----
I0000      I0003      FCP
I0001      I0102      FCP

```

Example 7-5 lists the available PPRC ports from storage B on site B to storage A on site A for the reverse direction. This example shows that there are two available ports for PPRC.

*Example 7-5 Output of the lsavailpprcport command on system B*

```

dsccli> lsavailpprcport -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -remotewwnn 500507630EFE0154 02:01
Date/Time: November 2, 2005 1:31:33 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Local Port Attached Port Type
-----
I0003      I0000      FCP
I0102      I0001      FCP

```

An I/O port number has four digits to indicate its location (as shown in Figure 7-2):

- The first digit (*R*) is for the frame location.
- The second digit (*E*) is for the I/O enclosure.
- The third digit (*C*) is for the adapter.
- The fourth digit (*P*) is for the adapter's port.

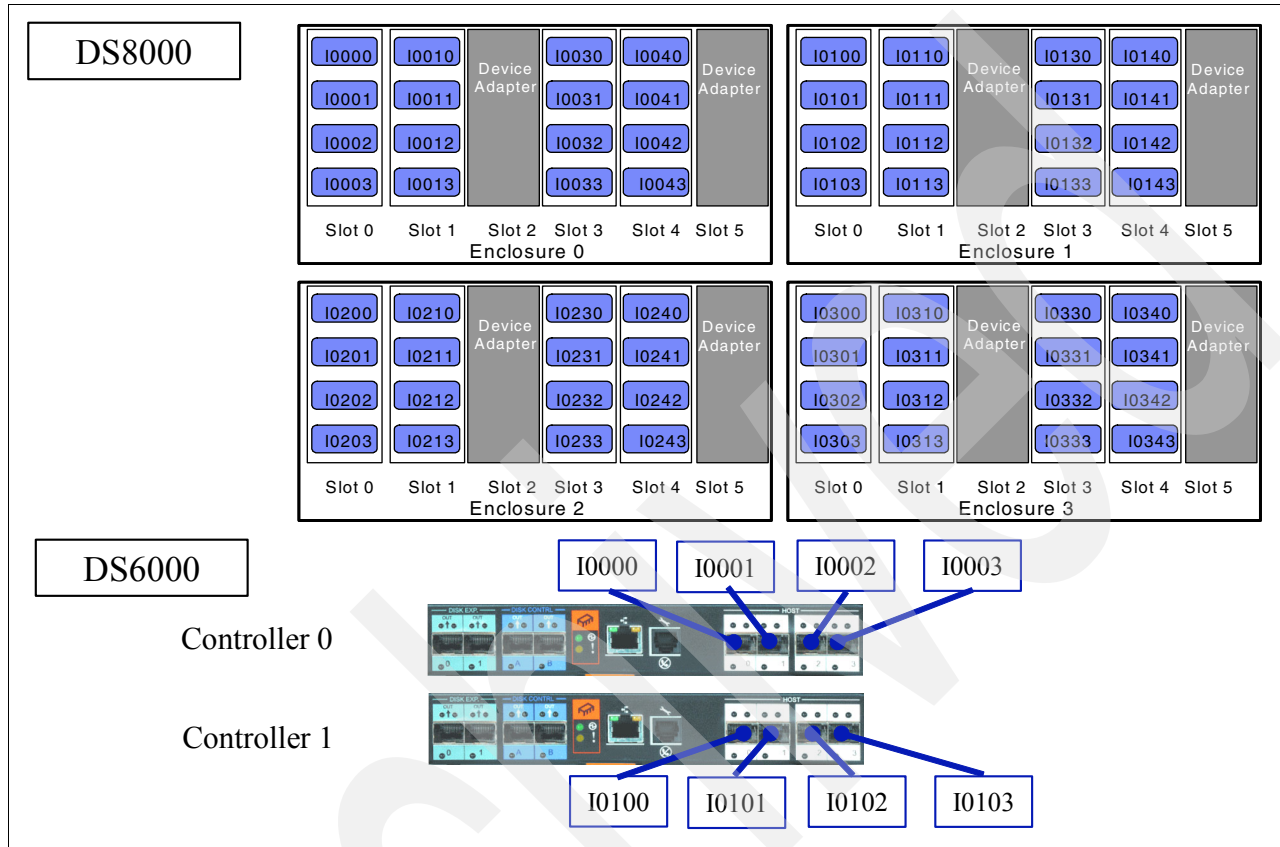


Figure 7-2 DS8000 and DS6000 port numbering

- For PPRC path redundancy, we recommend that you select two from the list of available I/O port pairs. Create the PPRC paths by running the following command:

```
dscli mkpprcpath -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -remotewwnn <WNN of target Storage_image> -src_lss <Source_LSS_ID> -tgt_lss <Target_LSS_ID> <Source_IO_Port>:<Target_IO_Port> <Source_IO_Port>:<Target_IO_Port> ...
```

The **mkpprcpath** command establishes or replaces PPRC paths between the source LSS and the target LSS over an FC connection. *Replaces* means that if you run this command again with different *source\_io\_port* and *target\_io\_port* parameters your previous paths are lost unless you also specify them with the command.

**Note:** This command creates a path or paths in one direction from the source LSS to the target LSS. If you want to run the mirror copy in the reverse direction to switch back the business application, create the path or paths also in the reverse direction.

Example 7-6 shows the creation of a PPRC path from source LSS 01 on storage A on site A on target LSS 02 on storage B on site B.

*Example 7-6 Output of the makepprcpath command: From system A to system B*

---

```
dscli> mkpprcpath -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwnn 500507630EFFF68
-srclass 01 -tgtlss 02 I0000:I0003
Date/Time: October 18, 2005 7:52:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00149I mkpprcpath: Remote Mirror and Copy path 01:02 successfully established.
```

---

Example 7-7 shows the creation of a PPRC path from source LSS 02 on storage B on site B on target LSS 01 on storage A on site A for the reverse direction.

*Example 7-7 Output of the makepprcpath command: From system B to system A*

---

```
dscli> mkpprcpath -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -remotewwnn 500507630EFE0154
-srclass 02 -tgtlss 01 I0003:I0000
Date/Time: October 18, 2005 7:52:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00149I mkpprcpath: Remote Mirror and Copy path 02:01 successfully established.
```

---

## Displaying a PPRC path

To display an established PPRC path, enter the following command:

```
dscli lspprcpath -dev <storage_image_ID> <source_LSS_ID>
```

Example 7-8 lists the PPRC path for LSS 01 on storage A.

*Example 7-8 Output of the lspprcpath command for system A*

---

```
dscli> lspprcpath -dev IBM.1750-13ABVDA 01
Date/Time: October 18, 2005 7:53:42 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Src Tgt State  SS  Port  Attached Port Tgt WNN
=====
01 02  Success FF02 I0000 I0003          500507630EFFF68
```

---

Example 7-9 lists the PPRC path for LSS 02 on storage B. This example shows that the current status of the path is Success.

*Example 7-9 Output of the lspprcpath command for system B*

---

```
dscli> lspprcpath -dev IBM.1750-13AAG8A 02
Date/Time: October 18, 2005 8:04:56 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Src Tgt State  SS  Port  Attached Port Tgt WNN
=====
02 01  Success FF01 I0003 I0000          500507630EFE0154
```

---

## Removing a PPRC path

To remove the established PPRC path, enter the following command:

```
dscli rmpprcpath -quiet -dev <storage_image_ID> -remotedev <storage_image_ID>
-remotewwnn <WNN of target Storage_image> <Source_LSS_ID>:<Target_LSS_ID>
```

To view the confirmation prompt, omit the **-quiet** option.

Example 7-10 shows the removal of the PPRC path from source LSS 01 on storage A on site A on target LSS 02 on storage B on site B.

*Example 7-10 Output of the rmpprcpath command for system A*

```
dsscli> rmpprcpath -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwnn 500507630EFFFFC68 01:02
Date/Time: November 2, 2005 3:00:23 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 01:02:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 01:02 successfully removed.
```

## 7.2.2 Creating a Metro Mirror relationship

After you create the PPRC paths, create the Metro Mirror relationships:

1. Check which fixed block volumes are available for Metro Mirror on the source LSS and the target LSS. Enter the following command to display the volumes:

```
dsscli lsfbvol -dev <storage_image_ID>
```

Example 7-11 lists the available fixed block volumes for the source in storage A.

*Example 7-11 Output of the lsfbvol command for system A*

```
dsscli> lsfbvol -dev IBM.1750.13-ABVDA
Date/Time: October 26, 2005 6:25:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
rchltn2-boot  0100 Online    Normal     Normal      1750-A85   FB 520U   P1         32.8        35.2
rchltn2-boot-mr 0101 Online    Normal     Normal      1750-A85   FB 520U   P1         32.8        35.2
rchltn2-disk2  0102 Online    Normal     Normal      1750-A05   FB 520P   P1         32.8        35.2
```

Example 7-12 lists the available fixed block volumes for the target in storage B.

*Example 7-12 Output of the lsfbvol command for system B*

```
dsscli> lsfbvol -dev IBM.1750.13-AAG8A
Date/Time: October 26, 2005 6:26:22 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
rlttn2-boot-mmir 0200 Online    Normal     Normal      1750-A85   FB 520U   P0         32.8        35.2
rlttn2-btmr-mmir 0201 Online    Normal     Normal      1750-A85   FB 520U   P0         32.8        35.2
rlttn2-dsk2-mmir 0202 Online    Normal     Normal      1750-A05   FB 520P   P0         32.8        35.2
```

2. Select volume pairs between both the sites and create the Metro Mirror pairs. This process is similar to creating FlashCopy pairs. To create Metro Mirror pairs, enter the following command:

```
dsscli mkpprc -dev <source storage_image_ID> -remotedev<target storage_image_ID>
-type mmir <Source_Volume>:<Target_Volume>
```

Example 7-13 shows the creation of three Metro Mirror pairs. The source volumes 0100, 0101, and 0102 on storage A are on site A, and the target volumes 0200, 0201, and 0202 on storage B are on site B.

*Example 7-13 Output of the mkpprc command for system B*

```
dsscli> mkpprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type mmir 0100-0102:0200-0202
Date/Time: October 18, 2005 9:27:44 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0100:0200 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0101:0201 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0102:0202 successfully created.
```

## Displaying the status and properties of Metro Mirror

To display a Metro Mirror relationship and its properties, run the following command:

```
dscli lsprrc -l -dev <source_storage_image_ID> -remotedev <target
storage_image_ID> <source_volume_ID>:<target_volume_ID>..
```

Leave out the *target\_volume\_ID* parameter and the *-l* option to omit the current *OutOfSyncTracks* attributes.

Example 7-14 lists the Metro Mirror relationship for volume pairs with source volumes 0100, 0101, and 0102 on storage A on site A and lists target volumes 0200, 0201, and 0202 on storage B on site B. This example shows that the current number of tracks that are not synchronized are displayed under *OutOfSyncTracks* with the status at *Copy Pending*. This attribute indicates the progress of the initial background copy. When the number of *OutOfSyncTracks* becomes 0, all the data is copied from the source volume on the source storage to the target volume on the target storage through the initial background copy process.

### Example 7-14 Output of the *lsprrc* command

```
dscli> lsprrc -l -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0100-0102
Date/Time: October 18, 2005 10:22:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode      First Pass          Status
=====
0100:0200 Copy Pending -      Metro Mirror 533124          Disabled Disabled  invalid
-      01 0      Disabled      Invalid
0101:0201 Copy Pending -      Metro Mirror 525783          Disabled Disabled  invalid
-      01 0      Disabled      Invalid
0102:0202 Copy Pending -      Metro Mirror 525805          Disabled Disabled  invalid
-      01 0      Disabled      Invalid
```

When the initial background copy is completed, the result shown in Example 7-15 displays. This example shows that the current number of tracks that are not synchronized is 0 and the status is Full Duplex. The initial asynchronous background copy from the source volume on storage A on site A to the target volume on site B is complete. After that, subsequent written data on the source volumes is copied to the target volumes synchronously.

### Example 7-15 Output of the *lsprrc* command: Copy complete

```
dscli> lsprrc -l -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0100-0102
Date/Time: October 18, 2005 11:03:57 PM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode      First Pass          Status
=====
0100:0200 Full Duplex -      Metro Mirror 0                Disabled Disabled  invalid
-      01 0      Disabled      Invalid
0101:0201 Full Duplex -      Metro Mirror 0                Disabled Disabled  invalid
-      01 0      Disabled      Invalid
0102:0202 Full Duplex -      Metro Mirror 0                Disabled Disabled  invalid
-      01 0      Disabled      Invalid
```

## Ending a Metro Mirror relationship

To end a Metro Mirror relationship, enter the following command:

```
dscli rmprrc -quiet -dev <source_storage_image_ID> -remotedev <target
storage_image_ID> <source_volume_ID>:<target_volume_ID>
```

To view the confirmation prompt, omit the *-quiet* option.



Example 7-16 shows the termination of Metro Mirror relationships for volume pairs. The source volumes 0100, 0101, and 0102 are on storage A on site A, and the target volumes 0200, 0201, and 0202 are on storage B on site B, with the confirmation prompt.

*Example 7-16 Output of the rmprrc command*

---

```

dscli> rmprrc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0100-0102:0200-0202
Date/Time: November 2, 2005 1:20:58 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00160W rmprrc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
0100-0102:0200-0202:? [y/n]:y
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0100:0200 relationship successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0101:0201 relationship successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0102:0202 relationship successfully withdrawn.

```

---

### Suspending Metro Mirror

Suspending a Metro Mirror relationship means that the background copy process is stopped but the storage system still keeps track of the changes to the source volume in its internal bitmap so that the relationship can be resumed later on without a full re-synchronization. To suspend the Metro Mirror synchronous copy, enter the following command:

```

dscli pausepprc -dev <source_storage_image_ID> -remotedev <target
storage_image_ID> <source_volume_ID>:<target_volume_ID>

```

Example 7-17 shows the suspension of Metro Mirror synchronous copy for volume pairs. The source volumes 0100, 0101, and 0102 are on storage A on site A, and the target volumes 0200, 0201, and 0202 are on storage B on site B with the confirmation prompt.

*Example 7-17 Output of the pausepprc command*

---

```

dscli> pausepprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0100:0200 0101:0201 0102:0202
Date/Time: October 21, 2005 7:11:25 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 0100:0200 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 0101:0201 relationship successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 0102:0202 relationship successfully paused.

```

---

If you look closely at the relationship and properties of Metro Mirror, you see that the status of the source volumes is *Suspended*, and the current number of tracks that are not synchronized to remote volume is growing. See Example 7-18.

*Example 7-18 Output of the lsprrc command: Showing OutOfSyncTracks*

---

```

dscli> lsprrc -l -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0100-0102
Date/Time: October 21, 2005 7:12:09 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State   Reason      Type          Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0100:0200 Suspended Host Source Metro Mirror 172          Disabled Disabled  invalid
-          01 0          Disabled      Invalid
0101:0201 Suspended Host Source Metro Mirror 171          Disabled Disabled  invalid
-          01 0          Disabled      Invalid
0102:0202 Suspended Host Source Metro Mirror 727          Disabled Disabled  invalid
-          01 0          Disabled      Invalid

```

---

## Resuming Metro Mirror

To resume the suspended Metro Mirror synchronous copy to bring the volumes back into a synchronous full-duplex state, enter the following command:

```
dscli resumepprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type mmir <source_volume_ID>:<target_volume_ID>
```

Example 7-19 shows that the suspended Metro Mirror synchronous copy for volume pairs has resumed with source volumes 0100, 0101, and 0102 on storage A on site A, and target volumes 0200, 0201, and 0202 on storage B on site B with the confirmation prompt.

### Example 7-19 Output of the resumepprc command

```
dscli> resumepprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type mmir 0100-0102:0200-0202
Date/Time: October 21, 2005 7:21:51 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 0100:0200 relationship successfully resumed. This
message is being returned before the copy completes.
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 0101:0201 relationship successfully resumed. This
message is being returned before the copy completes.
CMUC00158I resumepprc: Remote Mirror and Copy volume pair 0102:0202 relationship successfully resumed. This
message is being returned before the copy completes.
```

**Note:** When the Metro Mirror relationship is established, the target volumes are SCSI reserved and *not* accessible to the host. If you want to switch your production to the target site you need to failover PPRC so that the reserved target volumes become accessible to the host (see 7.3, “Switching over the system from the local site to remote site” on page 290).

If you try to access the target volumes while the relationship is established by performing an IPL of the backup server from the target volumes whose status is Full Duplex, the system IPL fails with the System Reference Code (SRC) B2003200 LP=002.

If you suspend the Metro Mirror with the **pausepprc** command and then try to access the target, the IPL fails. While the relationship is suspended, an IPL of the backup server from the target volumes whose its status is Target Suspended, causes the system IPL to fail with SRC B2003200 LP=002.

If you use the **-tgtread** option for the **mkpprc** command and then perform an IPL of the backup server from the target volumes for which read access is possible but not write access, the system IPL loops with SRC A60xxxxx and is not completed. This is because data is attempting to be written to the load source during the IPL.

## 7.3 Switching over the system from the local site to remote site

If you use the configuration provided in Figure 7-3 for your switchover environment and if a disaster occurs on storage A on site A, you must first check the status of the Metro Mirror environment as follows:

1. Check the PPRC path. Refer to “Displaying a PPRC path” on page 286 for more details.
2. Check the Metro Mirror relationships and properties. Refer to “Displaying the status and properties of Metro Mirror” on page 288 for more details.

### 7.3.1 Making the volumes available on the remote site

To perform an IPL of the target server from the copied volumes (target volumes) on the storage at the remote site, the copied volumes must be available to read and write. Making the volumes available on the remote site is a one-step process.

To make the earlier target volumes available to the host with DS CLI, enter the following command:

```
dscli failoverpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type mmir <source_volume_ID>:<target_volume_ID>
```

This command performs the following actions:

- ▶ Terminates the previous Metro Mirror relationship
- ▶ Establishes a new Metro Mirror relationship
- ▶ Suspends the new Metro Mirror relationship

The state of the target volumes that were previously source volumes is preserved by taking into account the fact that the previous source LSS might no longer be reachable.

This command changes the previous target volumes to new source volumes and its status to suspended, as shown in Figure 7-3. Thus, the server can access the new suspended source volumes to read and write.

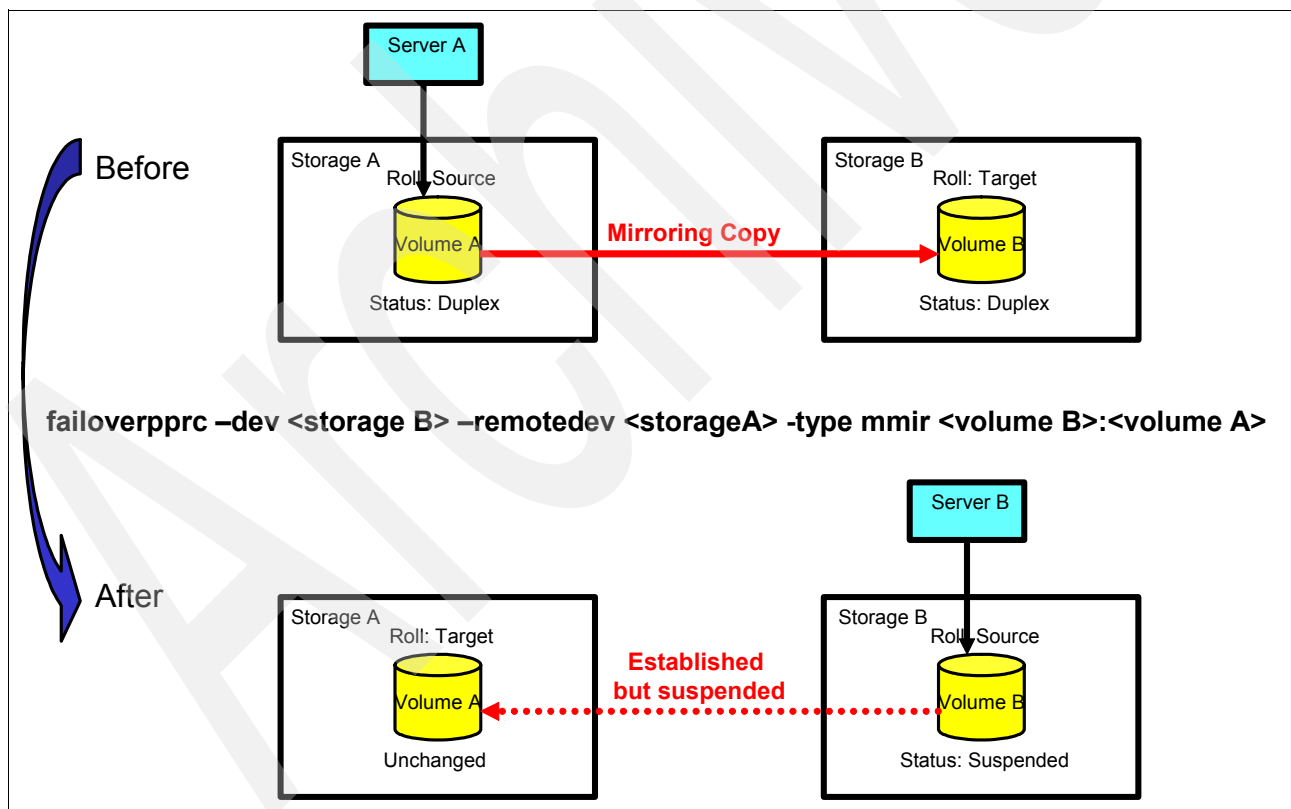


Figure 7-3 Failover of Metro Mirror volumes

Example 7-20 shows the failover of three Metro Mirror pairs. The source volumes 0100, 0101, and 0102 are on site A. The target volumes 0200, 0201 and 0202 are on site B. Also failover is to site B.

*Example 7-20 Output of the failoverpprc command*

```

dscli> failoverpprc -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -type mmir 0200-0202:0100-0102
Date/Time: October 19, 2005 3:58:34 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0200:0100 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0201:0101 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0202:0102 successfully reversed.

```

If you look closely at the relationship and properties of Metro Mirror, you see that the status of the new source volume is *Suspended*, and the reason is *Host Source*, as shown in Example 7-21.

*Example 7-21 Metro Mirror status after failover*

```

dscli> lsprrc -dev IBM.1750-13AAG8A -l 0200-0202
Date/Time: October 19, 2005 1:01:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID      State   Reason      Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
-----
0200:0100 Suspended Host Source Metro Mirror 0      Disabled Disabled invalid
-          02 0      Disabled Invalid
0201:0101 Suspended Host Source Metro Mirror 0      Disabled Disabled invalid
-          02 0      Disabled Invalid
0202:0102 Suspended Host Source Metro Mirror 0      Disabled Disabled invalid
-          02 0      Disabled Invalid

```

### 7.3.2 Performing an IPL of the backup server on the remote site

After running the `failoverpprc` command, you can perform an IPL of the backup server from the copied volumes (target volumes) on the remote site. When you activate the backup server, we recommend that you IPL the server manually with a restricted state because you might have to resolve the following issues in the backup server before you allow users to access the application on the backup server:

- ▶ Check the network and TCP/IP settings. They might need to be modified.
- ▶ Check the consistency of the application data because there is a possibility that some application data in the memory is *not* written to the disk. If necessary, apply the journal entries for the database.

To perform an IPL of the backup server with a restricted state, follow the steps in 6.2.4, “Performing an IPL of the target server” on page 259.

**Note:** You can see that this IPL is an abnormal IPL unless the operating system in the production server is in a state of shutdown when the Metro Mirror relationship is terminated. The abnormal IPL takes longer because processes are running, performing database recovery and journal recovery.

## 7.4 Switching back the system from the remote site to local site

If the production site is available again, schedule a switchover from the backup site to the production site. When the storage on the production site is available, check the condition of the previous configuration, for example, the volumes, the PPRC paths, and so on. In this section, we assume that these configuration components on the production site are not lost or have been recovered.

### 7.4.1 Starting Metro Mirror from the remote site to local site (reverse direction)

To switch back the system to the local site, resynchronize the data on the storage on the remote site to the storage on the local site, as shown in Figure 7-4. Resynchronization from the remote site to the local site is a one-step process.

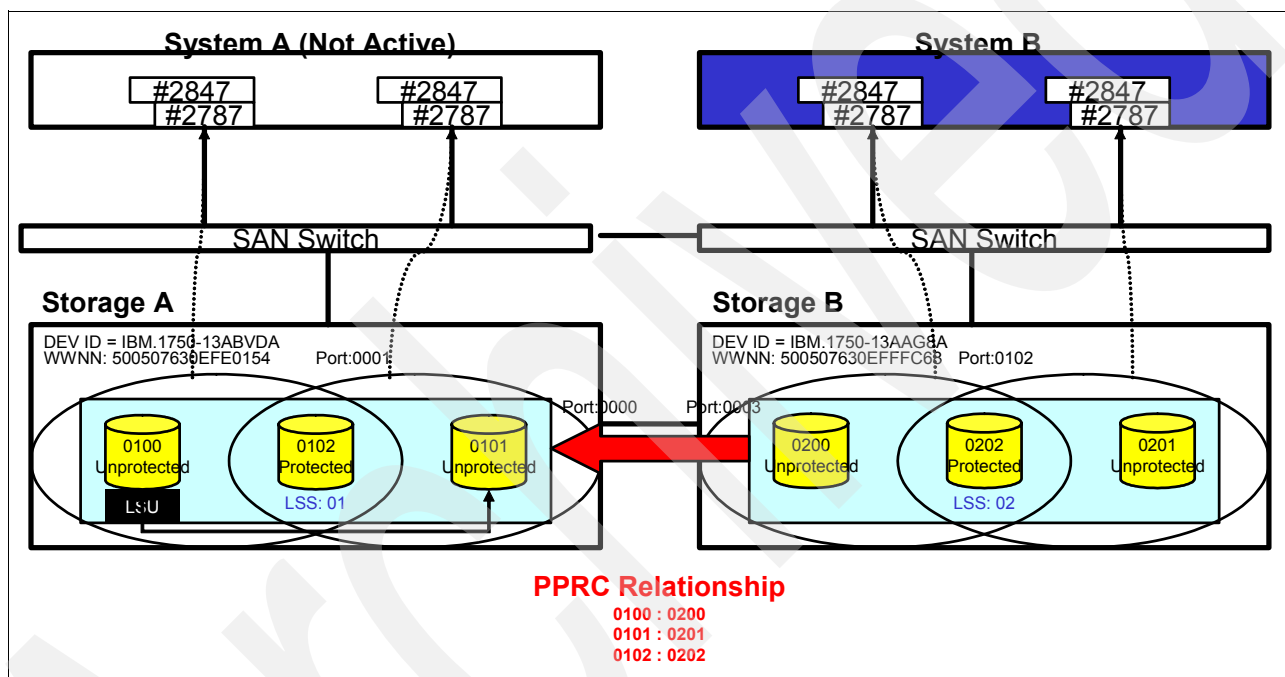


Figure 7-4 Starting Metro Mirror in reverse direction

When you start synchronization, the target volumes become unavailable.

**Important:** Before you start Metro Mirror, make sure the operating system of system A is in a state of shutdown. Otherwise, this operating system will hang.

Use the DS CLI to resynchronize from the new source volumes (old target) to the new target volumes (old source) by entering the following command:

```
dscli failbackpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type mmir <source_volume_ID>:<target_volume_ID>
```

The **failbackpprc** command performs the following actions:

- ▶ Checks the preserved state of the previous source volume to determine how much data to copy back.
- ▶ Copies either all the tracks or only the OutOfSyncTracks from the volume on the remote storage.
- ▶ Copies the subsequent written data on the source volumes to the target volumes synchronously after the initial copy, where the status of both the volumes becomes Full Duplex.

The DS CLI **failbackpprc** command changes the status of the target volume to Full Duplex, as shown in the *Before* step of Figure 7-5. Thus, the server cannot access the target volumes to read and write.

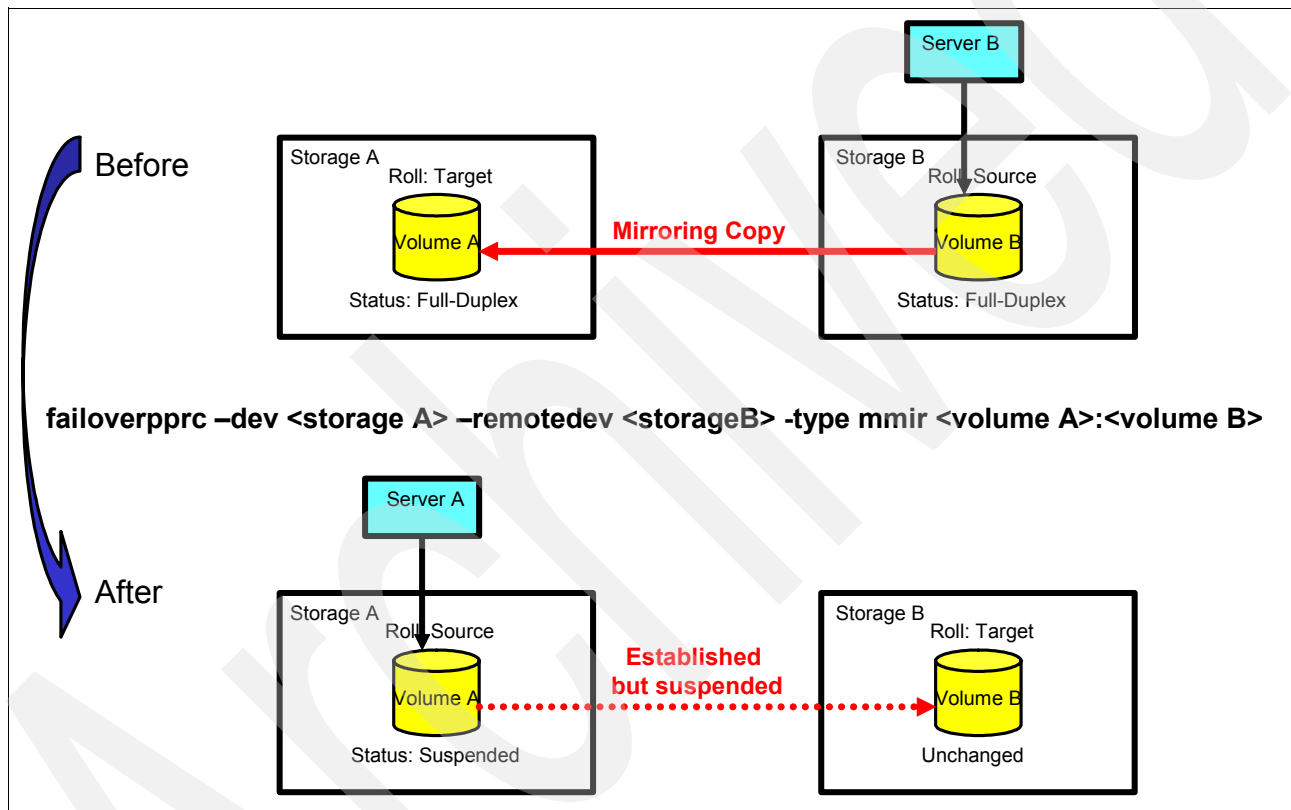


Figure 7-5 Running the failbackpprc command

Example 7-22 shows the failback of three former Metro Mirror pairs. The former source volumes 0100, 0101, and 0102 are on site A. The former target volumes 0200, 0201, and 0202 (in a suspended state) are on site B. The failback is to site B.

Example 7-22 Output of the failbackpprc command

```
dscli> failbackpprc -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -type mmir 0200-0202:0100-0102
Date/Time: October 19, 2005 3:54:12 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0200:0100 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0201:0101 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0202:0102 successfully failed back.
```

If you look closely at the relationship and properties of Metro Mirror, you can see that the status of the new source volume is Copy Pending and that the number of OutofSyncTracks is reduced, as shown in Example 7-23.

*Example 7-23 Output of the lsprrc command: Resync*

```

dscli> lsprrc -dev IBM.1750-13AAG8A -l 0200-0202
Date/Time: October 19, 2005 3:54:27 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0200:0100 Copy Pending -      Metro Mirror 22969      Disabled Disabled  invalid
-      02 0      Disabled Invalid
0201:0101 Copy Pending -      Metro Mirror 41251      Disabled Disabled  invalid
-      02 0      Disabled Invalid
0202:0102 Copy Pending -      Metro Mirror 23412      Disabled Disabled  invalid
-      02 0      Disabled Invalid
=====

```

After the initial copy is completed, you can see that the status of the new source volume is Full Duplex, and the number of OutofSyncTracks is 0, as shown in Example 7-24.

*Example 7-24 Output of the lsprrc command: Copy complete*

```

dscli> lsprrc -dev IBM.1750-13AAG8A -l 0200-0202
Date/Time: October 19, 2005 3:56:39 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0200:0200 Full Duplex -      Metro Mirror 0      Disabled Disabled  invalid
-      02 0      Disabled Invalid
0201:0201 Full Duplex -      Metro Mirror 0      Disabled Disabled  invalid
-      02 0      Disabled Invalid
0202:0202 Full Duplex -      Metro Mirror 0      Disabled Disabled  invalid
-      02 0      Disabled Invalid
=====

```

## 7.4.2 Making the volumes available on the local site

To perform an IPL of the server from the volumes on the storage on the local site, the volumes must be available for read and write. Making the volumes on the local site available is a two-step process. Follow these steps:

1. Turn off the server on the remote site. To switch back the system to the local site, the server on the remote site must be in a state of shutdown before the Metro Mirror relationship is suspended. Otherwise, the IPL of the server on the local site will be abnormal and might take longer to complete. If you do turn off the server, review the application data for consistency because there is a possibility that some application data in the memory is *not* written to the disk.

Follow the shutdown procedures for your site or use the PWRDWNSYS command. After you complete the shutdown process of the server on the remote site, ensure that the number of OutofSyncTracks is 0 and the status of Metro Mirror is Full Duplex with DS CLI.

2. Failover Metro Mirror to the local site with DS CLI in order to make the earlier source volumes available, by using the following command:

```

dscli failoverpprc -dev <source_storage_image_ID> -remotedev <target
storage_image_ID> -type mmir <source_volume_ID>:<target_volume_ID>

```

This command performs the following actions:

- Terminates the previous Metro Mirror relationship
- Establishes the new Metro Mirror relationship
- Suspends the new Metro Mirror relationship

The state of the previous source volume is preserved, taking into account the fact that the previous source LSS might no longer be reachable.

The **failoverpprc** command changes the previous target volumes to new source volumes and its status is changed to Suspended, as shown in Figure 7-6. Therefore, the server can access the new suspended volumes to read and write.

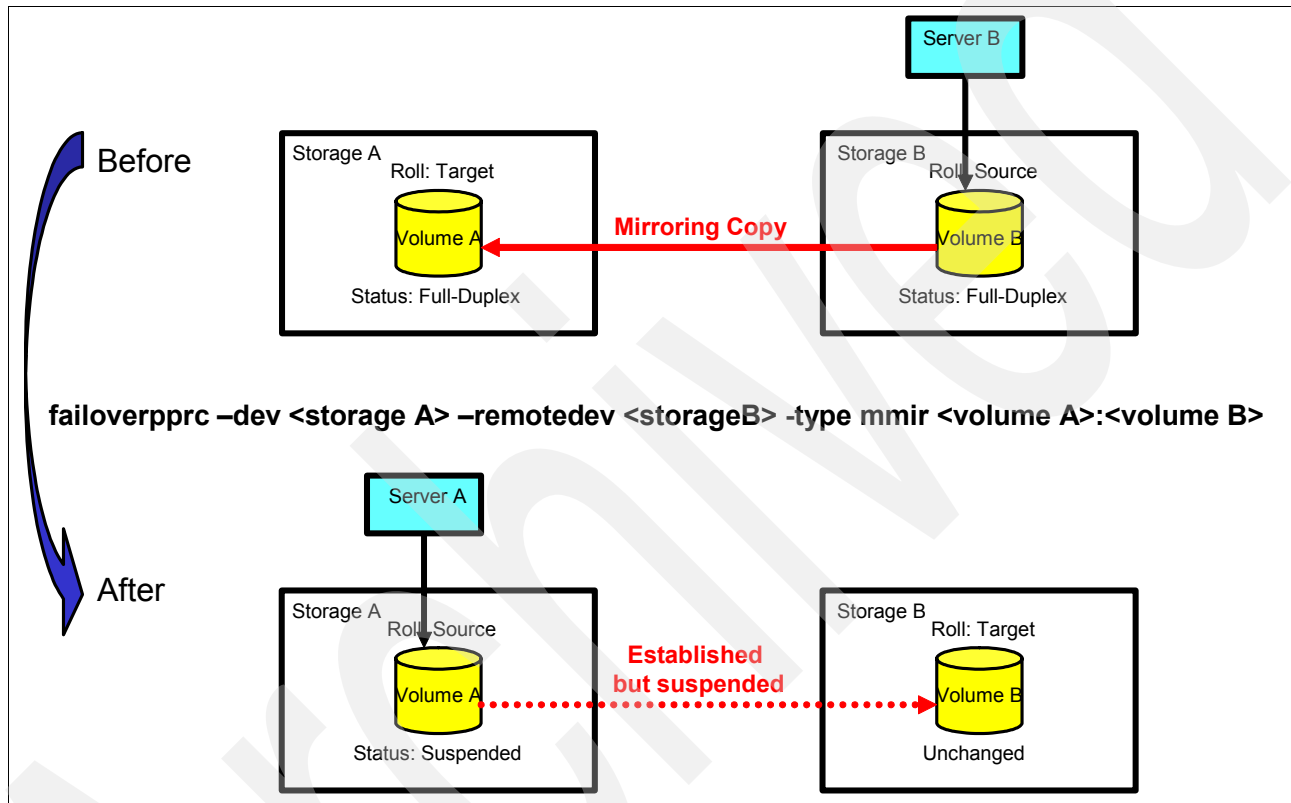


Figure 7-6 Running failoverpprc to local site

Example 7-25 shows the failover of three Metro Mirror pairs. The source volumes 0200, 0201, and 0202 (these were the original targets) are on site B. The target volumes 0100, 0101, and 0102 (these were the original source) are on site A. Failover is to site A.

*Example 7-25 Output of the failoverpprc command*

```

dscli> failoverpprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type mmir 0100-0102:0200-0202
Date/Time: October 19, 2005 3:58:34 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0100:0200 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0101:0201 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0102:0202 successfully reversed.

```



If you look closely at the relationship and properties of Metro Mirror, you can see that the status of the new source volume is Suspended and the reason is *Host Source*, as shown in Example 7-26.

*Example 7-26 Output of the lsprrc command*

```

dsccli> lsprrc -dev IBM.1750-13ABVDA -l 0100-0102
Date/Time: October 19, 2005 3:58:44 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State   Reason   Type           Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0100:0200 Suspended Host Source Metro Mirror 0           Disabled Disabled   invalid
-          01 0           Disabled      Invalid
0101:0201 Suspended Host Source Metro Mirror 0           Disabled Disabled   invalid
-          01 0           Disabled      Invalid
0102:0202 Suspended Host Source Metro Mirror 0           Disabled Disabled   invalid
-          01 0           Disabled      Invalid

```

### 7.4.3 Performing an IPL of the production server on the local site

After running the `failoverpprc` command successfully, you can perform the IPL of the production server from the volumes on the local site. If you did not change any settings on the system in the remote site, and shut down the operating system before failing over the Metro Mirror, you can perform an IPL on the production system with normal mode.

As long as the physical hardware resources of the production server, such as Ethernet adapters, expansion enclosures, and so on, have *not* changed on the local site, the server detects the hardware resources that are associated with the line descriptions again. The operating system then varies on LIND and starts the TCP/IP interface addresses automatically. However, the IPL of the recovered server should be performed carefully, regardless of whether you changed some settings on the system in the remote site. We recommend that you perform an IPL on the server manually to a restricted state before you allow users to access the application on the production server.

To perform an IPL on the production server with restricted state, follow the steps in 6.2.4, “Performing an IPL of the target server” on page 259.

## 7.4.4 Starting Metro Mirror from the local site to remote site (original direction)

The next stage of the recovery process is to start the synchronization from the volumes on the local storage to the volumes on the remote storage again. The resynchronization returns the system to a state of readiness and runs in normal disaster protection mode on the local site, as shown in Figure 7-7. The resynchronization from the local site to the remote site is a one-step process.

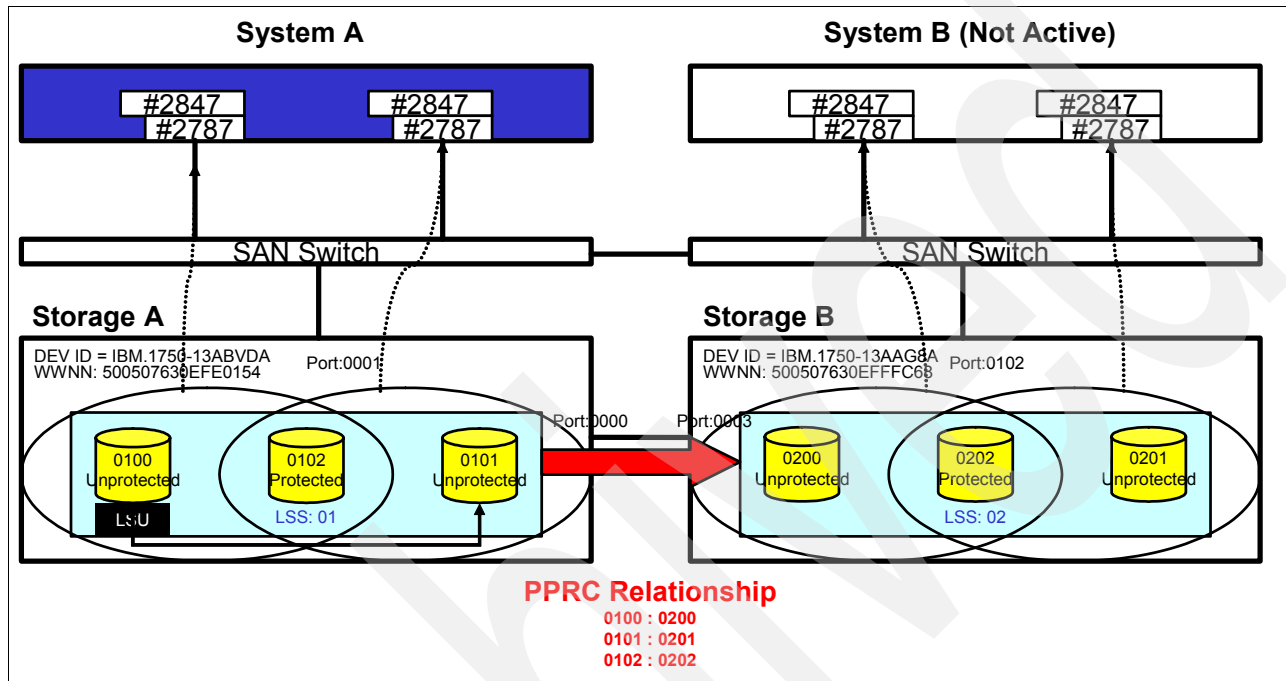


Figure 7-7 Resynchronization from the local site to the remote site

When you start synchronization, the target volumes become SCSI reserved again being unavailable to the host.

**Important:** Before you start Metro Mirror, the operating system of system B must be in a state of shutdown. Otherwise, this operating system will hang.

To resynchronize from the new source volumes to new target volumes, enter the following command:

```
dscli failbackpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type mmir <source_volume_ID>:<target_volume_ID>
```

This command performs the following actions:

- ▶ Checks the preserved state of the previous source volume to determine how much data to copy back.
- ▶ Copies either all the tracks or only the OutOfSyncTracks from the volume on the remote storage.
- ▶ Copies the subsequent written data on the source volumes to the target volumes synchronously after initial copy, where the status of both the volumes become Full Duplex.

The `failbackpprc` command changes the status of the target volume to Full Duplex, as shown in Figure 7-8. Therefore, the server cannot access the new target volumes to read and write.

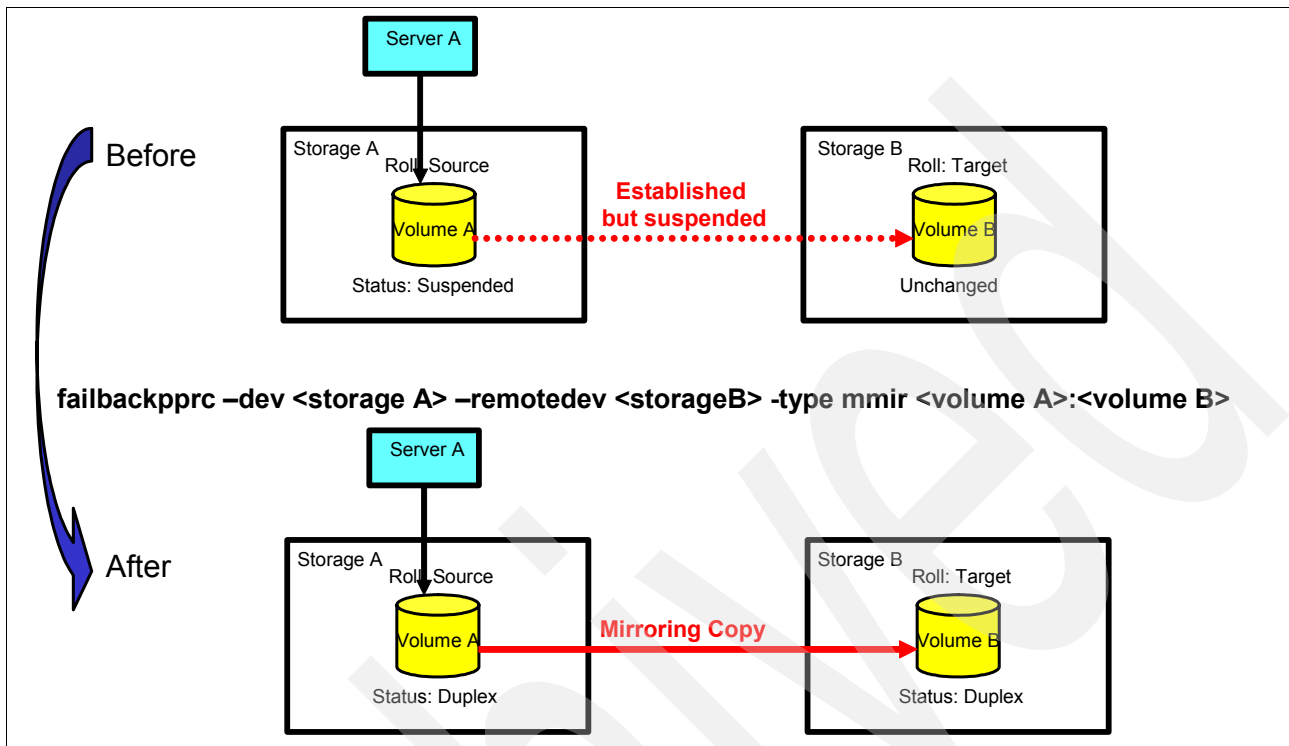


Figure 7-8 Running `failbackpprc`

Example 7-27 shows the failback of three former Metro Mirror pairs. The former source volumes 0200, 0201, and 0202 are on site B. The former target volumes 0100, 0101, and 0102 (in a suspended state) are on site A. Failback is to site A.

*Example 7-27 Output of the `failbackpprc` command*

```
dsccli> failbackpprc -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -type mmir 0200-0202:0100-0102
Date/Time: October 19, 2005 3:54:12 AM JST IBM DSCCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0200:0100 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0201:0101 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0202:0102 successfully failed back.
```

After finishing the initial copy, you can see that the status of the new source volume is Full Duplex and the number of OutofSyncTracks is 0, as shown in Example 7-28.

*Example 7-28 Output of the `lspprc` command*

```
dsccli> lspprc -l 0200-0202
Date/Time: October 19, 2005 3:59:21 AM JST IBM DSCCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks  Tgt Read Src Cascade  Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
-----
0200:0200 Full Duplex - Metro Mirror 0 Disabled Disabled invalid
- 02 0 Disabled Invalid
0201:0201 Full Duplex - Metro Mirror 0 Disabled Disabled invalid
- 02 0 Disabled Invalid
0202:0202 Full Duplex - Metro Mirror 0 Disabled Disabled invalid
- 02 0 Disabled Invalid
```

Archived



## Implementing Global Mirror using the DS CLI

In this chapter, we show how to implement the Global Mirror function of IBM System Storage Copy Services with i5/OS using the DS command-line interface (CLI).

## 8.1 Overview of the test environment

In this chapter, we describe how to implement Global Mirror with the entire direct access storage device (DASD) space of the System i5 environment. Before we begin that discussion, we provide an overview of the test environment in this section. Figure 8-1 provides an overview of the test environment used in this implementation example.

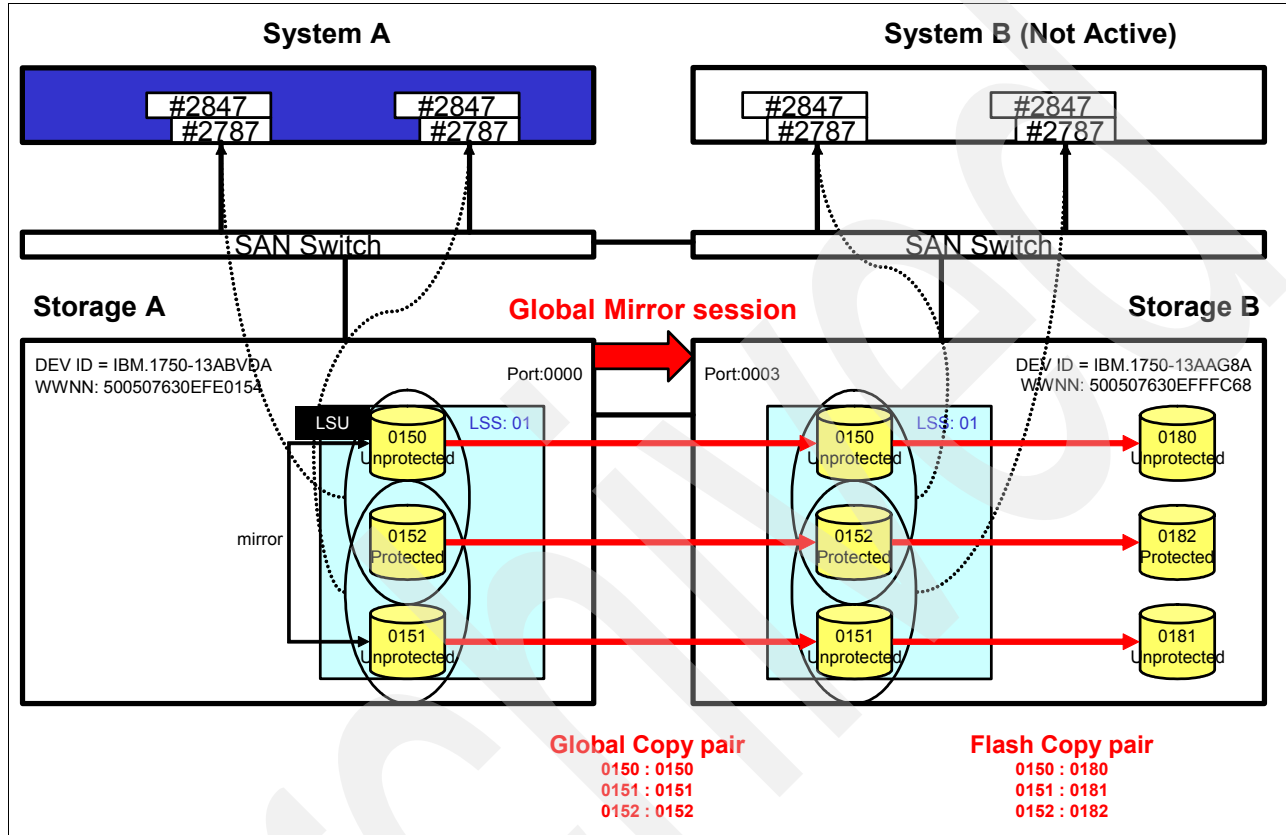


Figure 8-1 Test environment

System A is the production server, and system B is the backup server. System A is connected to storage A. An IPL is performed on System A from a storage area network (SAN) with 2847 I/O processor (IOP), having a load source unit in volume 0100 and a mirrored load source unit in volume 0101. The first 2847 feature card is tagged as the load source in the Hardware Management Console (HMC) partition profile.

**Note:** If you are using i5/OS V6R1 or later, you need to use the new multipath load source support instead of mirroring the load source for providing path protection.

Storage A and storage B are connected with Fibre Channel (FC) cables. This implementation example assumes that system A and storage A are on local site A and system B and storage B are on remote site B.

In this implementation example, a Global Mirror environment is created between storage A and storage B. The business application is then switched over from local site A to remote site B. Finally, the business application is switched back from remote site B to local site A.

Implementation of Global Mirror for the entire DASD space involves the following tasks:

1. Create a Global Mirror environment:
  - a. Create Peer-to-Peer Remote Copy (PPRC) paths.
  - b. Create Global Copy relationships.
  - c. Create FlashCopy relationships.
  - d. Create a Global Mirror session.
  - e. Start a Global Mirror session.
2. Switch over the system from the local site to the remote site:
  - a. Make the volumes available on the remote site.
  - b. Check and recover the consistency group of the FlashCopy target volumes.
  - c. Reverse the FlashCopy relationships.
  - d. Recreate the FlashCopy relationships.
  - e. Perform an IPL of the backup server on the remote site.
3. Switch back the system from the remote site to the local site:
  - a. Start Global Copy in reverse direction from the remote site to the local site.
  - b. Make the volumes available on the local site.
  - c. Start Global Copy in the original direction from the local site to the remote site.
  - d. Check or restart the Global Mirror session.
  - e. Perform an IPL of the production server on the local site.

We describe these steps in detail in the remaining sections of this chapter.

## 8.2 Creating a Global Mirror environment

Before you create a Global Mirror environment, you must:

1. Create the volumes for the backup server.
2. Create the volume groups and host connections for booting the backup server.
3. To perform an IPL from the copied load source unit on external storage, set the tagged load source unit to 2847 IOP or the Fibre Channel I/O adapter (IOA) that is connected to the copied load source unit in the HMC partition profile.

You need to perform these tasks only the first time that you set up the Global Mirror environment. Change the settings only if the source environment changes. Our example assumes that you have already completed these tasks.

**Important:** If you are considering using Metro Mirror, Global Mirror, or FlashCopy for the replication of the load source unit or other i5/OS disk units within the same DS6000 or DS8000 or between two or more DS6000 or DS8000 systems, the source volume and the target volume characteristics must be identical. The target and source must have matching capacities and matching protection types. For example, a *protected* 35 GB i5/OS volume must be replicated to another *protected* 35 GB volume if it is to be used in a replicated i5/OS configuration. It cannot be replicated to an *unprotected* 35 GB volume or to a volume of any other capacity. If you plan to migrate your load source from one size LUN to a larger size, use the System i *copy disk unit data utility* that is available in Dedicated Service Tools (DST).

In addition, the DS CLI offers the capability to change, such characteristics as the protection type of a previously defined volume. After a volume is assigned to an i5/OS partition and added to that partition's configuration, its characteristics must *not* be changed. If there is a requirement to change a characteristic of a configured volume, you must first completely remove it from the System i ASP configuration. After the characteristic changes, such as protection type and capacity, are made by destroying and recreating the volume or by using the DS CLI, you can reassign the volume into the System i configuration. To simplify the configuration, we recommend that you have a symmetrical configuration between two IBM System Storage solutions, creating the same volumes with the same volume IDs (LSSs and volume numbers).

For FlashCopy, we recommend that you plan the target volumes in order to be assigned on a different rank from where the source volumes are assigned for performance of the source server.

For more information about creating volumes, volume groups, and host connects, as well as tagging the load source IOP, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

## 8.2.1 Creating Peer-to-Peer Remote Copy paths

Global Mirror is a combination of Global Copy, which is asynchronous PPRC, and FlashCopy, as shown in Figure 8-2. Create the PPRC path for Global Copy between volume A and volume B.




Figure 8-2 Basic Global Mirror configuration

For Global Copy, the PPRC paths must exist for every logical subsystem (LSS) between the source LSS, with which the source volumes are associated, and the target LSS, with which the target volumes are associated.



If you have a subordinate storage server for consistency group, establish a PPRC path between each subordinate storage and the corresponding Global Copy target storage. Also, establish the PPRC path between the master storage and any subordinate storage, as shown in Figure 8-3.

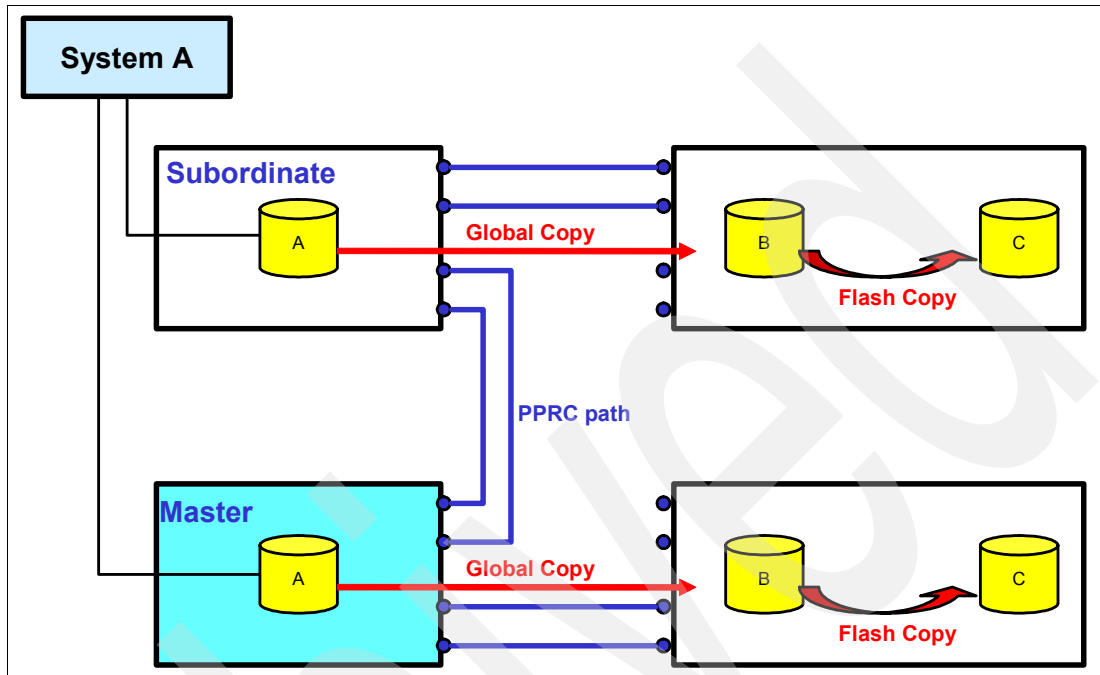


Figure 8-3 Global Mirror with subordinate storage

Creating the PPRC path for Global Copy is a four-step process.

**Important:** When creating the PPRC paths, consider the following points during the planning phase:

- ▶ For performance, use dedicated I/O ports for PPRC. Do *not* share them with the I/O of your servers. Use SAN switch zoning to restrict the server's storage system I/O port usage (see 3.2.5, "Planning for SAN connectivity" on page 67).
- ▶ For redundancy, create at least two paths between the same LSS. Use each I/O port of two or more different controllers in case of failure or maintenance of one of the controllers. For example, one path can use port I00xx on controller 0, and the other path can use port I01xx on controller 1.

Follow these steps:

1. Check which source LSS and target LSS are available on each System Storage solution.

Which LSS is associated with which volume depends on the LSS volume ID. Each volume has a four-digit hexadecimal volume ID, such as 0100. The first two digits indicate the LSS ID.

If the volume that you want to define as the source has the volume ID 0100, its LSS ID is *01*.

If the volume that you want to define as the target has the volume ID 0200, its LSS ID is *02*.

Therefore, to see which source LSS and target LSS are available, look at the volume ID of the volume you want to include in the Global Copy relationship. To obtain the volume ID, enter the following command:

```
dscli lsfvol -dev <storage_image_ID>
```

The results of the **lsfbvol** command show the hexadecimal volume ID, where the first two digits indicate the LSS ID. Example 8-1 lists the available fixed block volumes for the source in storage A.

*Example 8-1 Output of the lsfvol command on system A*

---

```
dscli> lsfvol -dev IBM.1750-13-ABVDA
Date/Time: October 26, 2005 6:42:15 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
ltn2-ls-prd   0150 Online   Normal    Normal     1750-A85   FB 520U   P1         32.8        35.2
ltn2-lsm-prd  0151 Online   Normal    Normal     1750-A85   FB 520U   P1         32.8        35.2
ltn2-dk2-prd  0152 Online   Normal    Normal     1750-A05   FB 520P   P1         32.8        35.2
```

---

Example 8-2 lists the available fixed block volumes for the target in storage B. Our example shows that the source LSS ID is 01 in storage A and that the target LSS ID is also 01 in storage B.

*Example 8-2 Output of the lsfvol command on system B*

---

```
dscli> lsfvol -dev IBM.1750-13AAG8A
Date/Time: October 26, 2005 6:42:27 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
ltn2-ls-gm    0150 Online   Normal    Normal     1750-A85   FB 520U   P3         32.8        35.2
ltn2-lsm-gm   0151 Online   Normal    Normal     1750-A85   FB 520U   P3         32.8        35.2
ltn2-dk2-gm   0152 Online   Normal    Normal     1750-A05   FB 520P   P3         32.8        35.2
ltn2-ls-gmf   0180 Online   Normal    Normal     1750-A85   FB 520U   P3         32.8        35.2
ltn2-lsm-gmf  0181 Online   Normal    Normal     1750-A85   FB 520U   P3         32.8        35.2
ltn2-dk2-gmf  0182 Online   Normal    Normal     1750-A05   FB 520P   P3         32.8        35.2
```

---

2. Check the worldwide node name (WWNN) of the target storage. The WWNN is unique in every IBM System Storage solution and is a required parameter for the command to create the PPRC paths. Use the following DS CLI command to display the list of storage images with their WWNNs configured in a storage complex:

```
dscli lssi
```

You need to run this command on both your source and target storage system, unless the source storage and the target storage are configured within the same storage complex on your System Management Console (SMC) or HMC so that both WWNNs are displayed in the result.

Example 8-3 lists the WWNN of source storage A and target storage B.

*Example 8-3 Output of the lssi command*

---

```
dscli> lssi
Date/Time: November 1, 2005 8:29:03 AM JST IBM DSCLI Version: 5.0.6.142
Name ID          Storage Unit      Model WWNN          State ESSNet
=====
-   IBM.1750-13AAG8A IBM.1750-13AAG8A  511   500507630EFFFFC68 Online Enabled
-   IBM.1750-13ABVDA IBM.1750-13ABVDA  511   500507630EFE0154 Online Enabled
```

---

This example shows that the WWNN of the target storage is 500507630EFFF68 from its Storage Unit ID IBM.1750-13AAG8A. Write down the WWNN of the source storage, 500507630EFE0154. This is essential to create the PPRC path for the reverse direction.

3. Check which I/O ports are available for the PPRC paths between the source LSS and the target LSS. Enter the following command to see the available ports:

```
dscli lsavailpprcport -dev <source storage_image_ID> -remotedev <target
storage_image_ID> -remotewwn <WWNN of target Storage_image>
<Source_LSS_ID>:<Target_LSS_ID>
```

**Note:** For this command, you can specify any available LSS pair that you want for source and target.

The result of the **lsavailpprcport** command displays a list of FC I/O ports that can be defined as PPRC paths, as shown in Example 8-4. Each row indicates the available I/O port pair. The local port is on local storage and the attached port is on remote storage.

Example 8-4 lists the available PPRC ports between storage A on site A and storage B on site B.

*Example 8-4 List of available ports*

---

```
dscli> lsavailpprcport -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwn 500507630EFFF68 01:01
Date/Time: October 26, 2005 9:36:35 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Local Port Attached Port Type
=====
I0000      I0003      FCP
I0001      I0102      FCP
```

---

Example 8-5 lists the available PPRC ports from storage B on site B to storage A on site A for the reverse direction. This example shows that there are two available ports for PPRC.

*Example 8-5 Output of the lsavailpprcport command*

---

```
dscli> lsavailpprcport -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -remotewwn 500507630EFE0154 01:01
Date/Time: November 2, 2005 1:31:33 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Local Port Attached Port Type
=====
I0003      I0000      FCP
I0102      I0001      FCP
```

---

An I/O port number has four digits to indicate its location, as shown in Figure 8-4:

- The first digit (*R*) is for the frame location.
- The second digit (*E*) is for the I/O enclosure.
- The third digit (*C*) is for the adapter.
- The fourth digit (*P*) is for the adapter's port.

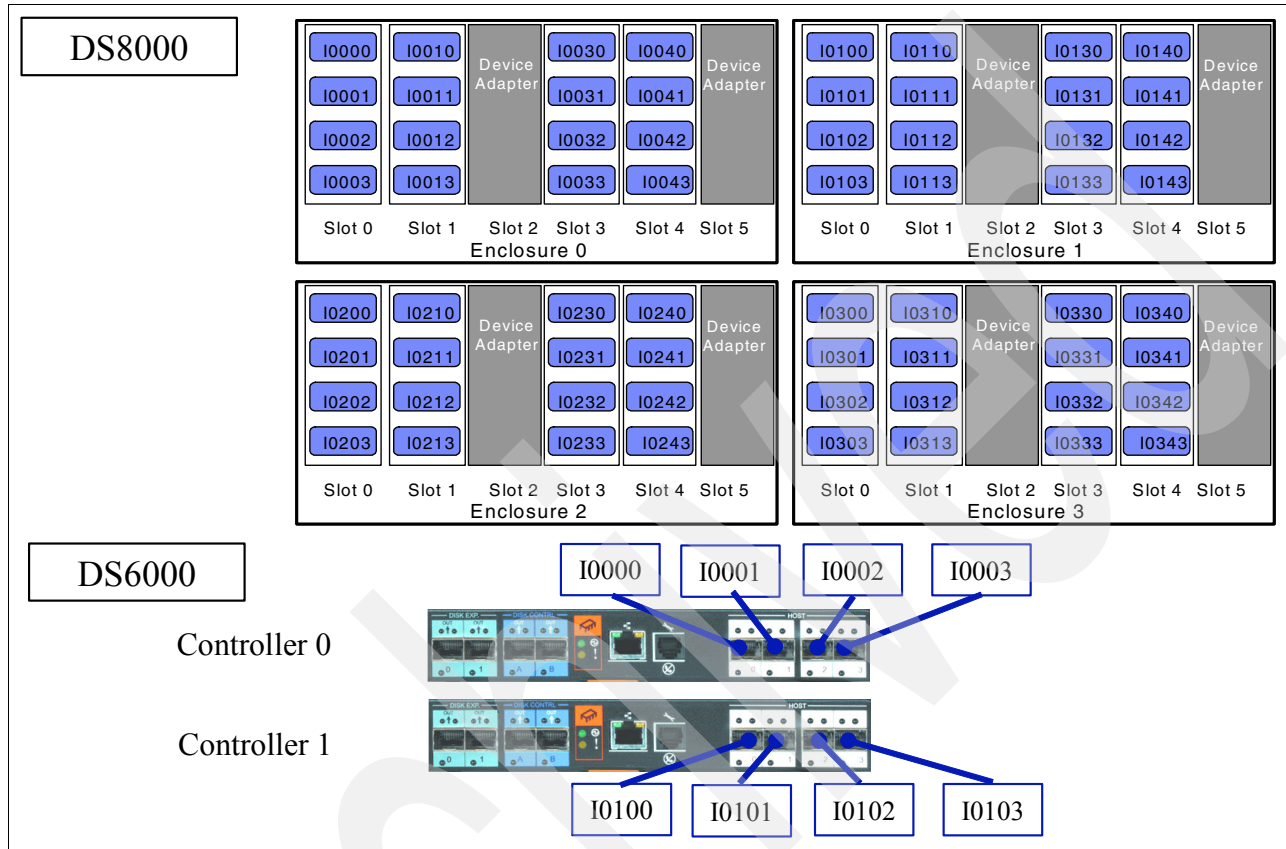


Figure 8-4 DS8000 and DS6000 port numbering

4. For PPRC path redundancy we recommend that you select two from the list of the available I/O port pairs. Create the PPRC paths by running the following command:

```
dsccli mkpprcpath -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -remotewwnn <WWNN of target Storage_image> -src_lss <Source_LSS_ID> -tgt_lss <Target_LSS_ID> <Source_IO_Port>:<Target_IO_Port> <Source_IO_Port>:<Target_IO_Port> ...
```

The `mkpprcpath` command establishes or replaces PPRC paths between the source LSS and the target LSS over an FC connection. *Replaces* means that if you run this command again with different `source_io_port` and `target_io_port` parameters your previous paths are lost unless you also specify them with the command.

**Note:** This command creates path or paths in one direction from the source LSS to the target LSS. If you want to run the mirror copy in the reverse direction to switch back the business application, create the path or paths also in the reverse direction.

Example 8-6 shows the creation of a PPRC path from source LSS 01 on storage A on site A on target LSS 01 on storage B on site B.

*Example 8-6 Output of the mkpprcpath command: From system A to system B*

---

```
dscli> mkpprcpath -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwnn 500507630EFFF68
-srclass 01 -tgtlss 01 I0000:I0003
Date/Time: October 26, 2005 9:40:52 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00149I mkpprcpath: Remote Mirror and Copy path 01:01 successfully established.
```

---

Example 8-7 shows the creation of a PPRC path from source LSS 01 on storage B on site B on target LSS 01 on storage A on site A for the reverse direction.

*Example 8-7 Output of the mkpprcpath command: From system B to system A*

---

```
dscli> mkpprcpath -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -remotewwnn 500507630EFE0154
-srclass 01 -tgtlss 01 I0003:I0000
Date/Time: October 26, 2005 9:52:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00149I mkpprcpath: Remote Mirror and Copy path 01:01 successfully established.
```

---

## Displaying a PPRC path

To display an established PPRC path, enter the following command:

```
dscli lspprcpath -dev <storage_image_ID> <source_LSS_ID>
```

Example 8-8 lists the PPRC path for LSS 01 on storage A.

*Example 8-8 Output of the lspprcpath command on system A*

---

```
dscli> lspprcpath -dev IBM.1750-13ABVDA 01
Date/Time: October 26, 2005 9:53:42 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Src Tgt State  SS  Port  Attached Port Tgt WNN
=====
01 01 Success FF01 I0000 I0003          500507630EFFF68
```

---

Example 8-9 lists the PPRC path for LSS 02 on storage B. This example shows that the current status of the path is *Success*.

*Example 8-9 Output of the lspprcpath command on system B*

---

```
dscli> lspprcpath -dev IBM.1750-13AAG8A 01
Date/Time: October 26, 2005 10:04:56 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Src Tgt State  SS  Port  Attached Port Tgt WNN
=====
01 01 Success FF01 I0003 I0000          500507630EFE0154
```

---

## Removing a PPRC path

To remove an established PPRC path, use the following command:

```
dscli rmprrcpath -quiet -dev <storage_image_ID> -remotedev <storage_image_ID>
-remotewwnn <WNN of target Storage_image> <Source_LSS_ID>:<Target_LSS_ID>
```

To show the confirmation prompt, omit the **-quiet** option.

Example 8-10 shows the removal of the PPRC path from source LSS 01, on storage A on site A, on target LSS 01 on storage B on site B.

*Example 8-10 Output of the rmpprcpath command*

```
dsscli> rmpprcpath -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -remotewwnn 500507630EFFFFC68 01:01
Date/Time: November 2, 2005 3:00:23 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00152W rmpprcpath: Are you sure you want to remove the Remote Mirror and Copy path 01:01:? [y/n]:y
CMUC00150I rmpprcpath: Remote Mirror and Copy path 01:01 successfully removed.
```

## 8.2.2 Creating a Global Copy relationship

**Note:** If Global Copy pairs are in several LSSs, select all of them during this process or run the process again on each LSS. If Global Copy pairs are spread over several storage images, run this process again on each of them.

Now that the PPRC path is ready, create the Global Copy relationship:

1. Check which fixed block volumes are available for Global Copy on the source LSS and the target LSS. Use the following command to display the volumes:

```
dsscli lsfbvol -dev <storage_image_ID>
```

Example 8-11 lists the available fixed block volumes for the source in storage A.

*Example 8-11 Output of the lsfbvol command on system A*

```
dsscli> lsfbvol -dev IBM.1750-13-ABVDA
Date/Time: October 26, 2005 6:42:15 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
ltnn2-ls-prd  0150 Online   Normal    Normal      1750-A85   FB 520U   P1        32.8        35.2
ltnn2-lsm-prd 0151 Online   Normal    Normal      1750-A85   FB 520U   P1        32.8        35.2
ltnn2-dk2-prd 0152 Online   Normal    Normal      1750-A05   FB 520P   P1        32.8        35.2
```

Example 8-12 lists the available fixed block volumes for the target in storage B. Our example shows that the source volumes for the Global Copy are 0150, 0151, and 0152 on storage A, and the target volumes for the Global Copy are 0150, 0151, and 0152 on storage B. In addition, the volumes 0180, 0181, and 0182 on storage B are the target volumes for FlashCopy.

*Example 8-12 Output of the lsfbvol command on system B*

```
dsscli> lsfbvol -dev IBM.1750-13AAG8A
Date/Time: October 26, 2005 6:42:27 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
ltnn2-ls-gm   0150 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8        35.2
ltnn2-lsm-gm  0151 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8        35.2
ltnn2-dk2-gm  0152 Online   Normal    Normal      1750-A05   FB 520P   P3        32.8        35.2
ltnn2-ls-gmf  0180 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8        35.2
ltnn2-lsm-gmf 0181 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8        35.2
ltnn2-dk2-gmf 0182 Online   Normal    Normal      1750-A05   FB 520P   P3        32.8        35.2
```

2. Select pairs of volumes between both the sites and create the Global Copy pairs. Use the following command to create the pairs:

```
dscli mkpprc -dev <source storage_image_ID> -remotedev<target storage_image_ID>
-type gcp <Source_Volume>:<Target_Volume> ...
```

If the Global Copy target volumes are already synchronized, use the **-mode nocp** option to omit the initial background synchronization.

Example 8-13 shows the creation of three Global Copy pairs. The source volumes 0150, 0151, and 0152 are on storage A on site A, and the target volumes 0150, 0151, and 0152 are on storage B on site B.

*Example 8-13 Output of the mkpprc command*

---

```
dscli> mkpprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type gcp -tgtread 0150:0150 0151:0151
0152:0152
Date/Time: October 26, 2005 9:45:48 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0150:0150 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0151:0151 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0152:0152 successfully created.
```

---

To view a Global Copy relationship and its properties, use the following command:

```
dscli lsprrc -l -dev <source storage_image_ID> -remotedev <target
storage_image_ID> <source_volume_ID>:<target_volume_ID> ...
```

To omit the current OutOfSyncTracks displayed attribute, remove the *target\_volume\_ID* parameter and the **-l** option.

Example 8-14 lists the Global Copy relationship for the volume pairs. The source volumes 0150, 0151, and 0152 are on storage A on site A, and the target volumes 0150, 0151, and 0152 are on storage B on site B.

*Example 8-14 Output of the lsprrc command*

---

```
dscli> lsprrc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 26, 2005 9:46:50 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Copy Pending - Global Copy 519038 Enabled Disabled invalid
- 01 300 Disabled False
0151:0151 Copy Pending - Global Copy 529669 Enabled Disabled invalid
- 01 300 Disabled False
0152:0152 Copy Pending - Global Copy 518732 Enabled Disabled invalid
- 01 300 Disabled False
```

---

This example shows that the current number of tracks that are *not* synchronized is displayed as the *OutOfSyncTracks* attribute. This attribute tells you how the initial asynchronous background copy is progressing. When the number of OutOfSyncTracks becomes 0, all the data is copied from the source volume on the source storage system, to the target volume on the target storage system through the initial asynchronous background copy process.

When the initial background copy is completed, the results shown in Example 8-15 are displayed. This example shows that the current number of tracks that are not synchronized is 0. The initial asynchronous background copy from the source volume on storage A on site A to the target volume on site B is completed. Then the subsequent data written on the source volumes is copied to the target volumes asynchronously. Therefore, the status continues to be Copy Pending.

*Example 8-15 Output of the lsprrc command: Progress*

```

dsccli> lsprrc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 26, 2005 11:30:01 PM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Copy Pending - Global Copy 0      Enabled Disabled invalid
-      01      300      Disabled      True
0151:0151 Copy Pending - Global Copy 0      Enabled Disabled invalid
-      01      300      Disabled      True
0152:0152 Copy Pending - Global Copy 0      Enabled Disabled invalid
-      01      300      Disabled      True

```

### Ending a Global Copy relationship

To end or terminate a Global Copy relationship, enter the following command:

```
dsccli rmprrc -quiet -dev <source storage_image_ID> -remotedev <target
storage_image_ID> <source_volume_ID>:<target_volume_ID> ...
```

To see the confirmation prompt, omit the `-quiet` option.

Example 8-16 shows the termination of Global Copy relationships for volume pairs. The source volumes 0150, 0151, and 0152 are on storage A on site A, and target volumes 0150, 0151, and 0152 are on storage B on site B with the confirmation prompt.

*Example 8-16 Output of the rmprrc command*

```

dsccli> rmprrc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A 0150-0152:0150-0152
Date/Time: November 2, 2005 1:20:58 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00160W rmprrc: Are you sure you want to delete the Remote Mirror and Copy volume pair relationship
0150-0152:0150-0152:? [y/n]:y
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0150:0150 relationship successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0151:0151 relationship successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0152:0152 relationship successfully withdrawn.

```

## 8.2.3 Creating a FlashCopy relationship

**Note:** If FlashCopy pairs are in several LSSs, select *all* of them during this process or run the process again on each LSS. If FlashCopy pairs are spread over several storage images, run this process again on each of them.

Now that the Global Copy relationship is established, create the FlashCopy relationships between source volume B and target volume C on storage B at site B. Follow these steps:

1. Check which fixed block volumes are available for FlashCopy pairs on the Global Mirror target storage image by using the following command:

```
dsccli lsfbvol -dev <storage_image_ID>
```



Example 8-17 lists the available fixed block volumes for the source in storage B. This example shows that the source volumes for the FlashCopy are 0150, 0151, and 0152 on storage B, and the target volumes for the FlashCopy are 0180, 0181, and 0182 on storage B.

*Example 8-17 Output of the lsfbvol command*

---

```

dscli> lsfbvol -dev IBM.1750-13AAG8A
Date/Time: October 26, 2005 6:42:27 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
Name          ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)
=====
ltnn2-ls-gm   0150 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8         35.2
ltnn2-lsm-gm  0151 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8         35.2
ltnn2-dk2-gm  0152 Online   Normal    Normal      1750-A05   FB 520P   P3        32.8         35.2
ltnn2-ls-gmf  0180 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8         35.2
ltnn2-lsm-gmf 0181 Online   Normal    Normal      1750-A85   FB 520U   P3        32.8         35.2
ltnn2-dk2-gmf 0182 Online   Normal    Normal      1750-A05   FB 520P   P3        32.8         35.2

```

---

2. Select volume pairs between both sites and create FlashCopy relationships between the Global Copy target volumes B becoming your FlashCopy source volumes and your designated FlashCopy target volumes C (see Figure 8-2 on page 304). When you create a Global Mirror environment, there is a FlashCopy relationship that requires certain attributes to be configured. These attributes are incremental, revertible, and nocopy FlashCopy functions. Use the following command to create the FlashCopy relationship:

```

dscli mkflash -dev <storage_image_ID> -record -persist -nocp
<Source_Volume>:<Target_Volume>...

```

In this command, the **-record** option is required to enable change recording because Global Mirror uses this FlashCopy relationship as *incremental* FlashCopy. The **-persist** option is required for incremental and revertible FlashCopy. The **-nocp** option is required for FlashCopy without background copy.

Example 8-18 shows the creation of three FlashCopy pairs with source volumes 0150, 0151, and 0152, and target volumes 0180, 0181, and 0182 on storage B on site B.

*Example 8-18 Output of the mkflash command*

---

```

dscli> mkflash -dev IBM.1750-13AAG8A -record -persist -nocp 0150:0180 0151:0181 0152:0182
Date/Time: October 26, 2005 9:56:40 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00137I mkflash: FlashCopy pair 0150:0180 successfully created.
CMUC00137I mkflash: FlashCopy pair 0151:0181 successfully created.
CMUC00137I mkflash: FlashCopy pair 0152:0182 successfully created.

```

---

### Displaying a FlashCopy relationship and properties

To display a FlashCopy relationship and its properties, enter the following command:

```

dscli lsflash -l -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID>..

```

To omit the current OutOfSyncTracks displayed attributes, remove the *target\_volume\_ID* parameter and the **-l** option.

Example 8-19 lists FlashCopy relationships for volume pairs with source volumes 0150, 0151, and 0152. This example shows that the SequenceNum of each pair is 0. This number changes each time there is an automatic incremental FlashCopy for saving a Global Mirror consistency group after starting the Global Mirror session.

*Example 8-19 Output of the lsflash command*

```

dscli> lsflash -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 26, 2005 9:57:11 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible SourceWriteEnabled
TargetWriteEnabled BackgroundCopy OutOfSyncTracks DateCreated          DateSynced
=====
0150:0180 01    0          300    Disabled Enabled Enabled Disabled Enabled
Enabled    Disabled 534475    Wed Oct 26 04:38:46 JST 2005 Wed Oct 26 04:38:46 JST 2005
0151:0181 01    0          300    Disabled Enabled Enabled Disabled Enabled
Enabled    Disabled 535788    Wed Oct 26 04:38:46 JST 2005 Wed Oct 26 04:38:46 JST 2005
0152:0182 01    0          300    Disabled Enabled Enabled Disabled Enabled
Enabled    Disabled 534353    Wed Oct 26 04:38:46 JST 2005 Wed Oct 26 04:38:46 JST 2005

```

**Ending a FlashCopy relationship**

To end or terminate a FlashCopy relationship, use the following command:

```
dscli rmflash -quiet -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID>
```

To view the confirmation prompt, omit the **-quiet** option.

Example 8-20 shows the termination of FlashCopy sessions for volume pairs with source volumes 0150, 0151, and 0152 and target volumes 0180, 0181, and 0182 without the confirmation prompt.

*Example 8-20 Output of the rmflash command*

```

dscli> rmflash -quiet -dev IBM.1750-13AAG8A 0150-0152:0180-0182
Date/Time: November 2, 2005 1:19:38 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00140I rmflash: FlashCopy pair 0150:0180 successfully removed.
CMUC00140I rmflash: FlashCopy pair 0151:0181 successfully removed.
CMUC00140I rmflash: FlashCopy pair 0152:0182 successfully removed.

```

**8.2.4 Creating a Global Mirror session**

After establishing the FlashCopy relationships, define the Global Mirror session and create a session ID between 1 and 255. Define this session number to all its LSSs that are participating in the session. Creating a Global Mirror session with DS CLI is a one-step process.

To create a Global Mirror session for an LSS by specifying the volume ID and session ID, enter the following command:

```
dscli mksession -dev <storage_image_ID> -lss <LSS ID> -volume
<volume_ID>,<volume_ID> <session ID>
```

Example 8-21 shows the creation of session number 01 for LSS 01 on storage A on site A.

*Example 8-21 Output of the mksession command*

---

```
dscli> mksession -dev IBM.1750-13ABVDA -lss 01 -volume 0150-0152 01
Date/Time: October 26, 2005 10:01:08 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00145I mksession: Session 01 opened successfully.
```

---

**Note:** Repeat this process for each LSS on each master storage server and subordinate storage server. Define the same session number to all the LSSs in the subordinate storage servers participating in the session.

### Displaying the volumes assigned to Global Mirror

To display a list of the volumes assigned to a Global Mirror session and their properties, use the following command:

```
dscli lssession -l -dev <storage_image_ID> <LSS ID>
```

Example 8-22 lists the volumes that are assigned to a Global Mirror session and its properties on LSS 01 on storage A. This example shows that the status of the volume is Join Pending because this Global Mirror session is not yet started.

*Example 8-22 Output of the lssession command*

---

```
dscli> lssession -dev IBM.1750-13ABVDA 01
Date/Time: October 26, 2005 10:01:46 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
LSS ID Session Status Volume VolumeStatus PrimaryStatus SecondaryStatus FirstPassComplete
AllowCascading
=====
01 01 Normal 0150 Join Pending Primary Copy Pending Secondary Simplex False
Disable
01 01 Normal 0151 Join Pending Primary Copy Pending Secondary Simplex False
Disable
01 01 Normal 0152 Join Pending Primary Copy Pending Secondary Simplex False
Disable
```

---

To change the volumes that are participating in the Global Mirror session for each LSS, use the following command:

```
dscli chsession -dev <storage_image_ID> -lss <LSS ID> -action <add or remove>
-volume <volume_ID> <session ID>
```

### Removing a Global Mirror session

To remove a Global Mirror session for each LSS, enter the following command:

```
dscli rmssession -quiet -dev <storage_image_ID> -lss <LSS ID> <session number>
```

To view the confirmation prompt, omit the **-quiet** option.

Example 8-23 shows the removal of session number 01 for LSS 01 on storage A on site A.

*Example 8-23 Output of the rmssession command*

---

```
dscli> rmssession -dev IBM.1750-13ABVDA -lss 01 01
Date/Time: November 2, 2005 2:39:46 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00148W rmssession: Are you sure you want to close session 01? [y/n]:y
CMUC00146I rmssession: Session 01 closed successfully.
```

---

## 8.2.5 Starting a Global Mirror session

Now that a Global Mirror session is created, start the Global Mirror session. We recommend that you start the session after you complete the initial copy of Global Copy. Starting a Global Mirror session with DS CLI is a single-step process. Enter the following command:

```
dscli mkgmir -dev <storage_image_ID> -lss <Master LSS ID> -cginterval <Number of seconds> -coordinate <Number of milliseconds> -drain <Number of seconds> -session <Session_ID> <Master_Control_Path_LSS_ID>:<Subordinate_Control_Path_LSS_ID>
```

This command helps you define the master LSS with the `-dev` and `-lss` parameters. You can optionally specify Global Mirror tuning parameters such as maximum drain time, maximum coordination time, and consistency group interval time. If you have subordinate storage servers, you can specify those servers as well.

**Note:** We recommend that you *not* specify a consistency group interval so that the Global Mirror code dynamically adjusts the consistency group interval creating consistency groups as often as the available bandwidth and I/O workload allow.

Example 8-24 shows the creation and start of a Global Mirror session with only one storage image A on site A.

### Example 8-24 Output of the mkgmir command

---

```
dscli> mkgmir -dev IBM.1750-13ABVDA -lss 01 -session 01
Date/Time: October 27, 2005 6:50:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00162I mkgmir: Global Mirror for session 01 successfully started.
```

---

## Displaying a Global Mirror relationship

To display a Global Mirror session relationship, use the following command:

```
dscli showgmir -dev <storage_image_ID> -metrics <Master_LSS_ID>
```

To omit performance statistics, remove the `-metrics` option.

Example 8-25 shows the Global Mirror relationship of master LSS 01 on storage A. This example shows that the Copy State is Running and the Current Time and the consistency group (CG) Time are the same. To understand this better, the properties of each component of Global Mirror are displayed after you start the Global Mirror session.

### Example 8-25 Output of the showgmir command

---

```
dscli> showgmir 01
Date/Time: October 27, 2005 7:02:19 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID IBM.1750-13ABVDA/01
Master Count 1
Master Session ID 0x01
Copy State Running
Fatal Reason Not Fatal
CG Interval (seconds) 0
XDC Interval(milliseconds) 50
CG Drain Time (seconds) 30
Current Time 10/27/2005 01:48:57 JST
CG Time 10/27/2005 01:48:57 JST
Successful CG Percentage 100
FlashCopy Sequence Number 0x435FB379
Master ID IBM.1750-13ABVDA
Subordinate Count 0
Master/Subordinate Assoc -
```

---

Example 8-26 shows the relationship, properties, and status of Global Copy in the implementation example environment. This example shows that Global Copy is working and that the number of OutOfSyncTracks attribute have become 0 in a few seconds. Thus, the consistency group is created and copied to the target volumes of the Global Copy in a few seconds.

*Example 8-26 Output of the lspprc command*

```

dscli> lspprc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 27, 2005 7:10:21 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Copy Pending -      Global Copy 23      Enabled Disabled invalid
-      01      300      Disabled      True
0151:0151 Copy Pending -      Global Copy 22      Enabled Disabled invalid
-      01      300      Disabled      True
0152:0152 Copy Pending -      Global Copy 172     Enabled Disabled invalid
-      01      300      Disabled      True
dscli>
dscli>
dscli> lspprc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 27, 2005 7:10:36 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Copy Pending -      Global Copy 0       Enabled Disabled invalid
-      01      300      Disabled      True
0151:0151 Copy Pending -      Global Copy 0       Enabled Disabled invalid
-      01      3       Disabled      True
00
0152:0152 Copy Pending -      Global Copy 0       Enabled Disabled invalid
-      01      3       Disabled      True
00

```

Example 8-27 shows the relationship, properties, and status of FlashCopy in the example environment. This example shows that FlashCopy is working and the number of SequenceNum changes every few seconds. Thus, a consistency group is created and copied to the target volume of the FlashCopy every few seconds.

*Example 8-27 Output of the lsflash command*

```

dscli> lsflash -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 27, 2005 7:11:15 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible SourceWriteEnabled
TargetWriteEnabled BackgroundCopy OutOfSyncTracks DateCreated          DateSynced
=====
0150:0180 01      435FB58C   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536463    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:09 JST 2005
0151:0181 01      435FB58C   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536489    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:09 JST 2005
0152:0182 01      435FB58C   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536486    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:09 JST 2005
dscli>
dscli>
dscli> lsflash -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 27, 2005 7:11:26 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible SourceWriteEnabled
TargetWriteEnabled BackgroundCopy OutOfSyncTracks DateCreated          DateSynced
=====
0150:0180 01      435FB598   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536572    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:21 JST 2005
0151:0181 01      435FB598   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536570    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:21 JST 2005
0152:0182 01      435FB598   300    Disabled Enabled   Enabled   Disabled Enabled
Disabled   Disabled   536539    Wed Oct 26 04:38:46 JST 2005 Thu Oct 27 01:53:21 JST 2005

```

Example 8-28 shows the relationship, properties, and status of the Global Mirror session for each LSS in this implementation example environment. This example shows that the status of consistency group (CG) is *In Progress* and VolumeStatus is *Active*.

*Example 8-28 Output of the ls session command*

```

dscli> ls session -l -dev IBM.1750-13ABVDA 01
Date/Time: October 27, 2005 7:12:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
LSS ID Session Status      Volume VolumeStatus PrimaryStatus      SecondaryStatus  FirstPassComplete
AllowCascading
=====
01    01 CG In Progress 0150  Active      Primary Copy Pending Secondary Simplex True          Disable
01    01 CG In Progress 0151  Active      Primary Copy Pending Secondary Simplex True          Disable
01    01 CG In Progress 0152  Active      Primary Copy Pending Secondary Simplex True          Disable

```

**Suspending a Global Mirror session**

To suspend a Global Mirror session, use the following command:

```

dscli pausegmir -dev <storage_image_ID> -lss <Master LSS ID> -session <session ID>
<Master_Control_Path_LSS_ID>:<Subordinate_Control_Path_LSS_ID>

```

**Note:** Suspending a Global Mirror only suspends building consistency groups but leaves Global Copy running.

Example 8-29 shows the suspension of a Global Mirror session of LSS 01 on storage A, and later, its status.

*Example 8-29 Output of the pausegmir command*

---

```
dscli> pausegmir -dev IBM.1750-13ABVDA -lss 01 -session 01
Date/Time: October 27, 2005 11:25:45 PM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00163I pausegmir: Global Mirror for session 01 successfully paused.
dscli>
dscli>
dscli> showgmir 01
Date/Time: October 27, 2005 11:26:05 PM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID                IBM.1750-13ABVDA/01
Master Count      1
Master Session ID 0x01
Copy State        Paused
Fatal Reason      Not Fatal
CG Interval (seconds) 0
XDC Interval(millisecons) 50
CG Drain Time (seconds) 30
Current Time      10/27/2005 18:12:37 JST
CG Time           10/27/2005 18:12:19 JST
Successful CG Percentage 99
FlashCopy Sequence Number 0x436099F3
Master ID         IBM.1750-13ABVDA
Subordinate Count 0
Master/Subordinate Assoc -
```

---

### Resuming a Global Mirror session

To resume a suspended Global Mirror session, enter the following command:

```
dscli resumegmir -dev <storage_image_ID> -lss <Master LSS ID> -cginterval <Number of seconds> -coordinate <Number of milliseconds> -drain <Number of seconds> -session <session_ID> <Master_Control_Path_LSS_ID>:<Subordinate_Control_Path_LSS_ID>
```

Use the **resumegmir** command to change the Global Mirror tuning parameters, for example, the maximum drain time, the maximum coordination time, and the consistency group interval time. You can also change the relationship between the master storage server and the subordinate storage server.

Example 8-30 shows that the suspended Global Mirror session has resumed and the status.

*Example 8-30 Output of the resumegmir command*

---

```
dscli> resumegmir -dev IBM.1750-13ABVDA -session 01 -lss 01
Date/Time: October 28, 2005 1:17:30 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00164I resumegmir: Global Mirror for session 01 successfully resumed.
dscli>
dscli>
dscli> showgmir 01
Date/Time: October 28, 2005 1:18:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID                IBM.1750-13ABVDA/01
Master Count      1
Master Session ID 0x01
Copy State        Running
Fatal Reason      Not Fatal
CG Interval (seconds) 0
XDC Interval(millisecons) 50
CG Drain Time (seconds) 30
Current Time      10/27/2005 20:04:44 JST
CG Time           10/27/2005 18:12:19 JST
Successful CG Percentage 99
FlashCopy Sequence Number 0x436099F3
Master ID         IBM.1750-13ABVDA
Subordinate Count 0
Master/Subordinate Assoc -
```

---

### Removing a Global Mirror relationship

To remove a Global Mirror session relationship, enter the following command:

```
dscli rmgmir -quiet -dev <storage_image_ID> -lss <Master LSS ID> -session <session ID> <Master_Control_Path_LSS_ID>:<Subordinate_Control_Path_LSS_ID>
```

Example 8-31 shows the removal of a Global Mirror session with only one storage A on site A.

*Example 8-31 Output of the rmgmir command*

---

```
dscli> rmgmir -dev IBM.1750-13ABVDA -lss 01 -session 01
Date/Time: November 2, 2005 2:36:59 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00166W rmgmir: Are you sure you want to stop the Global Mirror session 01:? [y/n]:y
CMUC00165I rmgmir: Global Mirror for session 01 successfully stopped.
```

---

To clean up the Global Mirror environment with DS CLI, perform the following tasks:

- ▶ Remove the Global Mirror session relationship between the Master Global Mirror session manager and its subordinates, that is, all the Global Mirror sessions with the same session number interconnected through the PPRC control paths.
- ▶ Remove the common Global Mirror session for each LSS on the source storage images.
- ▶ Remove the FlashCopy for each pair of the source volume, that is, read-only target volume of the Global Copy session, and the target volumes, which are also called *journal volumes*.



- ▶ Remove the Global Copy for each pair of source volume on site A and the target volumes on site B.
- ▶ Remove the PPRC paths between the source LSS and the target LSS for Global Copy. If you have a subordinate storage server, remove the PPRC path between the master storage server and the subordinate storage server.

## 8.3 Switching over the system from the local site to remote site

To better understand this concept, consider a situation where a disaster occurs on storage A on site A. In such a situation, you need to check the status of the Global Mirror environment and perform the following checks:

- ▶ The PPRC paths.
- ▶ The Global Copy relationships and properties.
- ▶ The FlashCopy relationships and properties.
- ▶ The Global Mirror session for each LSS.
- ▶ The Global Mirror session relationships.

### 8.3.1 Making the volumes available on the remote site

To perform an IPL from the backup server using the copied volumes on the remote site, make the Global Copy target volume available. With the DS CLI, make the previous Global Copy target volumes available by using the following command:

```
dscli failoverpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type gcp <source_volume_ID>:<target_volume_ID>
```

This command performs the following actions:

- ▶ Terminates the previous Global Copy relationship
- ▶ Establishes a new Global Copy relationship in the reverse direction
- ▶ Suspends the new Global Copy relationship

The state of the previous source volume is preserved by taking into account the fact that the previous source LSS might no longer be reachable.

This command changes the previous target volumes to new source volumes, and its status is changed to *Suspended*, as shown in Figure 8-5. Therefore, the new suspended source volume becomes accessible for read and write.

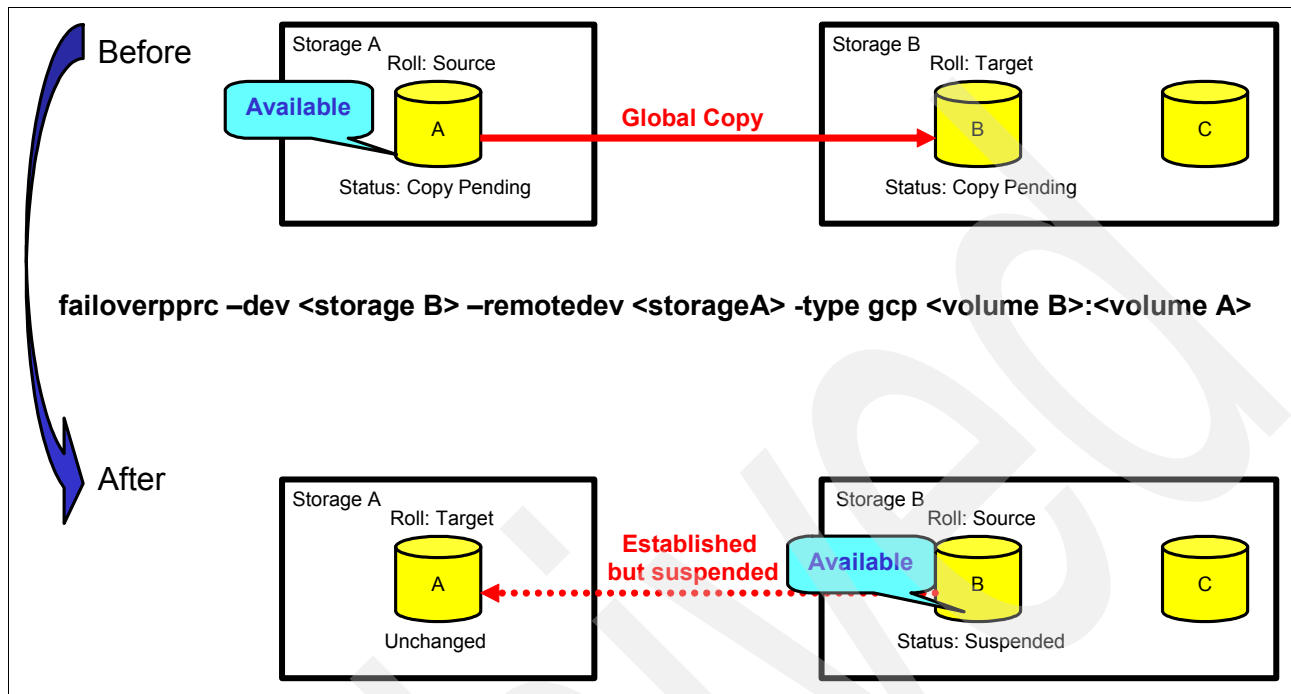


Figure 8-5 Failover of Global Copy to the remote site

Example 8-32 shows the failover of three Global Copy pairs. The source volumes 0150, 0151, and 0152 are on site A. The target volumes 0150, 0151, and 0152 are on site B. Failover is to site B.

Example 8-32 Output of the failoverpprc command: From system A to system B

```
dscli> failoverpprc -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -type gcp 0150:0150 0151:0151 0152:0152
```

```
Date/Time: October 28, 2005 8:02:09 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
```

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0150:0150 successfully reversed.
```

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0151:0151 successfully reversed.
```

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0152:0152 successfully reversed.
```

```
dscli>
```

```
dscli>
```

```
dscli> lspprc -l -dev IBM.1750-13AAG8A 0150-0152
```

```
Date/Time: October 28, 2005 8:02:38 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
```

ID	State	Reason	Type	Out Of Sync	Tracks	Tgt Read	Src Cascade	Tgt Cascade
Date Suspended	SourceLSS	Timeout (secs)	Critical Mode	First Pass	Status			

0150:0150	Suspended	Host	Source	Global Copy	0	Disabled	Disabled	invalid
-	01	300	Disabled	True				
0151:0151	Suspended	Host	Source	Global Copy	0	Disabled	Disabled	invalid
-	01	300	Disabled	True				
0152:0152	Suspended	Host	Source	Global Copy	0	Disabled	Disabled	invalid
-	01	300	Disabled	True				

### 8.3.2 Checking and recovering the consistency group of the FlashCopy target volume

Global Copy is an asynchronous copy. The data in the target volume of Global Copy is *not* used by the server. The consistent data must be on the target volume of FlashCopy. If a consistency group is being processed when a failure occurs, there is a possibility that volume C is not consistent, as shown in Figure 8-6.

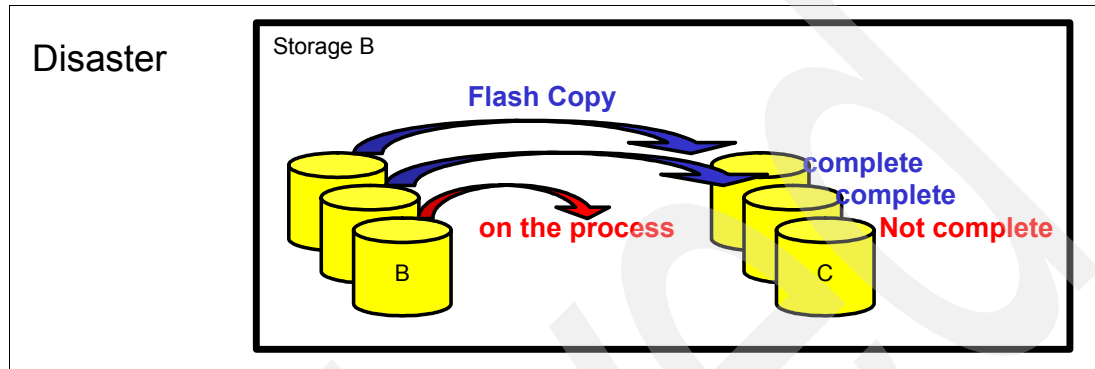


Figure 8-6 Invalid consistency

To check the consistency group, view its FlashCopy relationship using the following command:

```
dsccli lsflash -dev <storage_image_ID> <source_volume_ID>:<target_volume_ID> ...
```

Look for the Sequence Number and Revertible State attributes. Depending on whether the FlashCopy pairs are successfully set to a revertible state, and depending on whether their sequence numbers are equal or otherwise, enter either the **commitflash** command or the **revertflash** command.

#### Notes:

- ▶ Global Mirror is a distributed solution.
- ▶ When a consistency group is processed, the master Global Mirror session manager issues an incremental revertible FlashCopy on its own recovery site, and asks its subordinates to also perform this task on their recovery site.
- ▶ When a consistency group is in progress, several incremental revertible copies using FlashCopy might be running. The FlashCopy process looks at the change recording bitmap on the B volumes and compares it with the target bitmap on the C volumes. When the incremental FlashCopy is completed, the change recording bitmap is cleared. Therefore, all the changes are committed on the C volumes. The corresponding FlashCopy pairs are set to a nonrevertible state by the master Global Mirror session manager.

Because Global Mirror is a distributed solution, one of the following five situations can occur:

- ▶ If all the FlashCopy pairs are nonrevertible and their sequence numbers are equal, then the consistency group process is completed. No action is required because the consistency group is intact on the C volumes.
- ▶ If some FlashCopy pairs are revertible and their sequence numbers are equal and others are not revertible and their sequence numbers are equal but do *not* match the revertible FlashCopy sequence numbers, then some FlashCopy pairs are running in a consistency group process and some have not yet started their incremental process. To preserve the consistency, overwrite new data with the data saved at the last consistency formation for the FlashCopy pairs that are already in the new consistency group process. To accomplish this using DS CLI use the **revertflash** command.
- ▶ If all the FlashCopy pairs are revertible and their sequence numbers are equal, then all the FlashCopy pairs are running in a consistency group process and none have finished their incremental process. Overwrite the new data with the data saved at the last consistency formation for all the FlashCopy pairs. To accomplish this using DS CLI use the **revertflash** command.
- ▶ If some FlashCopy pairs are revertible and at least one is not revertible, but all their sequence numbers are equal, then some FlashCopy pairs are running in a consistency group process and some have already finished their incremental process. Commit data to a target volume in order to form a consistency between the source and the target. Only those FlashCopy pairs that have not already finished their incremental process must commit the changes because the nonrevertible pairs have already committed theirs. To accomplish this using DS CLI, use the **commitflash** command. Using this command on FlashCopy pairs that are not revertible only displays error messages.
- ▶ If none of these situations are possible, then the consistency group is corrupted.

Table 8-1 summarizes the consistency group status and the required action.

Table 8-1 Consistency group and FlashCopy validation decision table

<b>All FlashCopy relationships are revertible</b>	No	Yes	All but at least one are nonrevertible	Some FlashCopy pairs are revertible and others are not revertible
<b>All FlashCopy sequence numbers are equal</b>	Yes	Yes	Yes	<ul style="list-style-type: none"> <li>▶ Revertible FlashCopy pair sequence numbers are equal and not revertible.</li> <li>▶ FlashCopy pair sequence numbers are equal, but do not match the revertible FlashCopy sequence numbers.</li> </ul>
<b>Action</b>	No action required. All C volumes are consistent.	Withdraw all FlashCopy relations with the revert action.	Withdraw all FlashCopy relations with the commit action.	Withdraw all FlashCopy relations with the revert action.
<b>Comment</b>	Consistency group information ended.	All FlashCopy pairs are in a new consistency group process, and none have finished their incremental process.	Some FlashCopy pairs are running in a consistency group process, and some have already finished their incremental process.	Some FlashCopy pairs are running in a consistency group process, and some have not yet started their incremental process.

Usually all FlashCopy pairs are nonrevertible, and all sequence numbers are equal. This is a good condition, and you can proceed further. If not, perform the following recovery steps:

1. Revert to the previous consistent state using the following command:

```
dscli revertflash -dev <storage_image_ID> -seqnum <FlashCopy_Sequence_NB>
<Source_Volume> <Source_Volume>
```

2. Commit the data to the target volume in order to form a consistency group by running the following command:

```
dscli commitflash -dev <storage_image_ID> -seqnum <FlashCopy_Sequence_NB>
<Source_Volume> <Source_Volume>
```

Example 8-33 shows the consistency group of FlashCopy pairs, where the source volumes are 0150, 0151, and 0152, and the target volumes are 0180, 0181, and 0182. This example shows that all the FlashCopy relationships are nonrevertible and all the FlashCopy sequence numbers are equal. The consistent data is on the target volumes 0180, 0181, and 0182. However these FlashCopy target volumes contain only the changes from saving the consistency groups so they are not directly usable for the host.

Example 8-33 Output of the `lsflash` command

```

dsccli> lsflash -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 28, 2005 8:04:45 AM JST IBM DSCCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible SourceWriteEnabled
TargetWriteEnabled BackgroundCopy OutOfSyncTracks DateCreated          DateSynced
=====
0150:0180 01      436110CD    300    Disabled Enabled Enabled Disabled Enabled
Disabled Disabled 536570      Thu Oct 27 18:38:29 JST 2005 Fri Oct 28 02:34:43 JST 2005
0151:0181 01      436110CD    300    Disabled Enabled Enabled Disabled Enabled
Disabled Disabled 536552      Thu Oct 27 18:38:29 JST 2005 Fri Oct 28 02:34:43 JST 2005
0152:0182 01      436110CD    300    Disabled Enabled Enabled Disabled Enabled
Disabled Disabled 536571      Thu Oct 27 18:38:29 JST 2005 Fri Oct 28 02:34:43 JST 2005

```

### 8.3.3 Reversing a FlashCopy relationship

Now that you have consistent data on the FlashCopy target C volumes, reverse the FlashCopy relationship by copying the consistency group data from target C volumes to the B source volumes as shown in Figure 8-7. This reverse FlashCopy process ensures that the B volumes contain usable data for the host and not only the changes from saving the consistency groups after the initial Global Copy synchronization.




Figure 8-7 Reversing a FlashCopy

Reverse a FlashCopy relationship with the Fast Reverse Restore process by using the following command:

```

dsccli reverseflash -dev <storage_image_ID> -fast -tgtpprc -seqnum
<FlashCopy_Sequence_NB> <Source_Volume>:<Target_Volume> ...

```

Example 8-34 shows the result of using the `revertflash` command on the FlashCopy pairs with source volumes 0150, 0151, and 0152 on storage B on site B, and the FlashCopy original relationship. This example shows that after you enter the `reverseflash` command, the original FlashCopy relationship is terminated.

Example 8-34 Output of the `reverseflash` command

```

dsccli> reverseflash -dev IBM.1750-13AAG8A -fast -tgtpprc 0150:0180 0151:0181 0152:0182
Date/Time: October 28, 2005 8:07:17 AM JST IBM DSCCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00169I reverseflash: FlashCopy volume pair 0150:0180 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 0151:0181 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 0152:0182 successfully reversed.

```

```

dscli>
dscli>
dscli> lsflash -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 28, 2005 8:07:32 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMMCI9006E No Flash Copy instances named 0150-0152 found that match criteria: dev =
IBM.1750-13AAG8A.

```

### 8.3.4 Recreating a FlashCopy relationship

Consistent and usable data is now on volume B. However, the FlashCopy relationship is terminated and you have an isolated set of consistency group saves on volume C. Therefore, to have meaningful data on volume C and to prepare for re-enabling Global Mirror, you must re-establish the FlashCopy relationship between volume B and volume C. Use the same FlashCopy command that you used to establish FlashCopy when you created the Global Mirror environment. Refer to 8.2.3, “Creating a FlashCopy relationship” on page 312.

### 8.3.5 Performing an IPL of the backup server on the remote site

Perform an IPL on the backup server from volume B on the remote site, as shown in Figure 8-8.

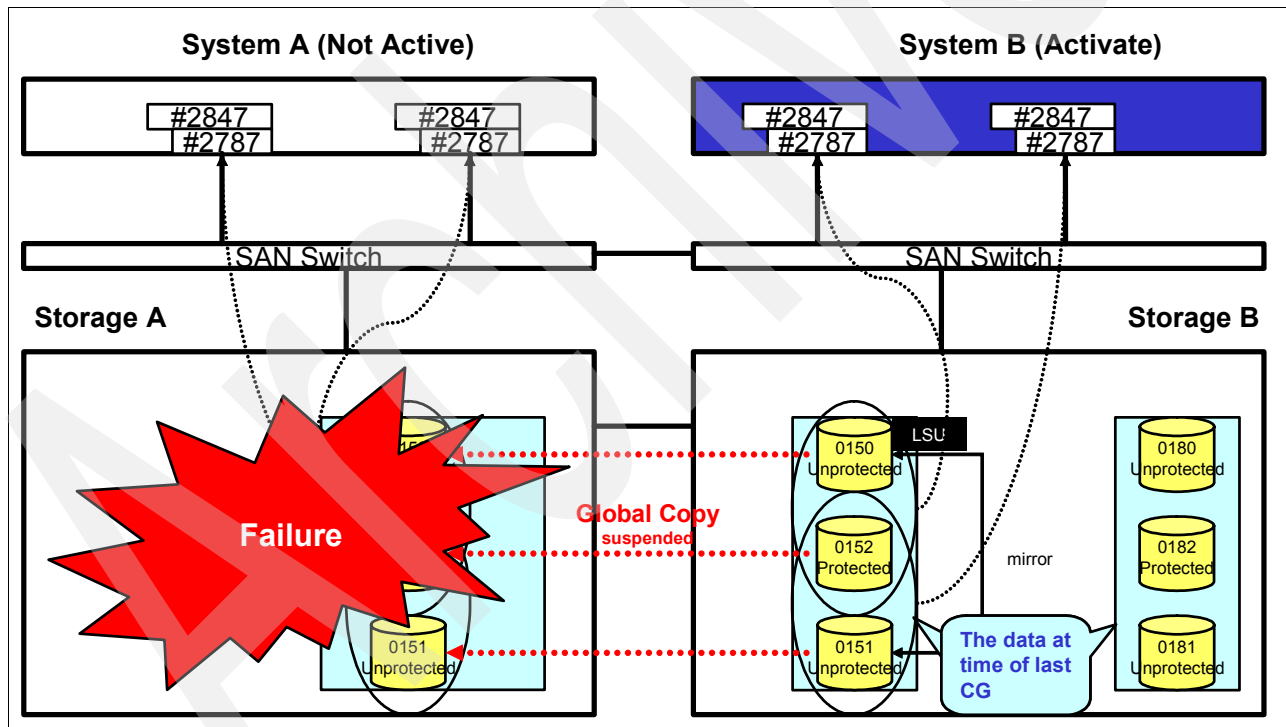


Figure 8-8 Performing an IPL of the backup server after failure

When you activate the backup server, we recommend that you perform an IPL on the server manually in a restricted state. We recommend this IPL because you might have to resolve the following issues in the backup server before you allow users to access the application on the backup server:

- ▶ Check and modify the network and the TCP/IP settings.
- ▶ Check the consistency of the application data because the application data on the remote storage is the data from the time that the last consistency group was processed. The data written by the application from the last consistency group to the failure can be lost on remote storage. There can also be a possibility that some application data in the memory is not written to the disk. The journal entries for the database might have to be applied.

To perform an IPL on the backup server with a restricted state, follow the steps in 6.2.4, “Performing an IPL of the target server” on page 259.

**Abnormal IPL:** This IPL is an abnormal IPL unless the operating system in the production server is in a state of shutdown when the Global Mirror relationship is terminated. The abnormal IPL might take longer because of the database recovery and journal recovery that is occurring.

## 8.4 Switching back the system from the remote site to local site

If the production site is available again, schedule a switchback from the system on the backup site to the production site. When the storage on the production site is available, check the condition of the previous configuration, for example, the volumes, the PPRC paths, and so on. In this section, we discuss the configuration of the storage on the production site that is not lost or is recovered.



## 8.4.1 Starting Global Copy from the remote site to local site (reverse direction)

To switch back the system to the local site, resynchronize the data on the storage on the remote site to the storage on the local site, as shown in Figure 8-9. Resynchronization from the remote site to the local site is a one-step process.

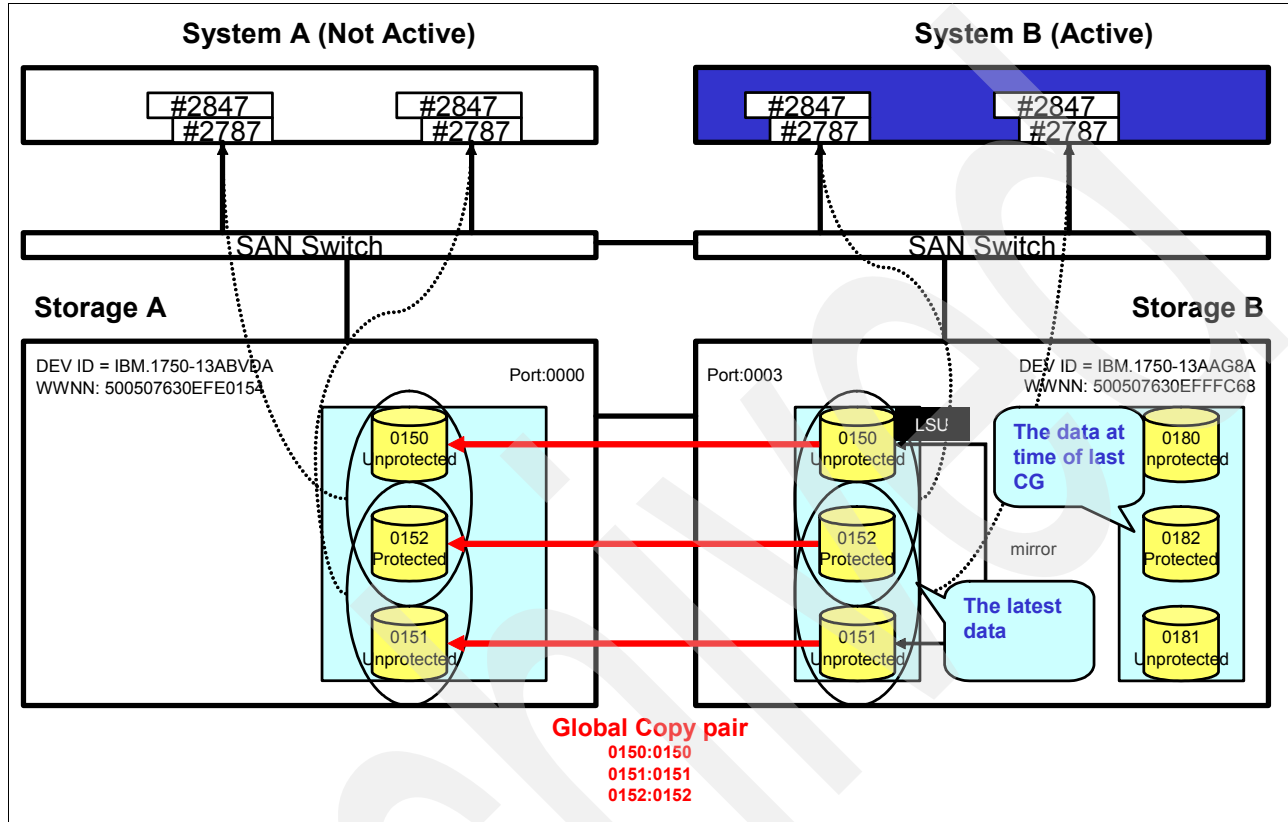


Figure 8-9 Reversing the direction of Global Copy

When you start Global Copy, the target volumes become unavailable.

**Important:** Before you start Global Copy, ensure that the operating system of system A is in a state of shutdown. Otherwise, this operating system will hang.

### Failback Global Copy

With DS CLI, resynchronize from the new source volumes at your recovery site to the new target volumes on your production site by using the following command:

```
dsccli failbackpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type gcp <source_volume_ID>:<target_volume_ID>...
```

This command performs the following tasks:

- ▶ Checks the preserved state of the previous source volume to determine how much data to copy back.
- ▶ Copies all the tracks or only OutOfSyncTracks from the volume on the remote storage.
- ▶ Copies the subsequent written data on the source volumes to the target volumes asynchronously, after the initial copy. The status of the volumes is Copy Pending because this copy process is asynchronous.

This command changes the previous source volumes to new target volumes, as shown in Figure 8-10. Thus, the server cannot access the new target volumes to read and write.

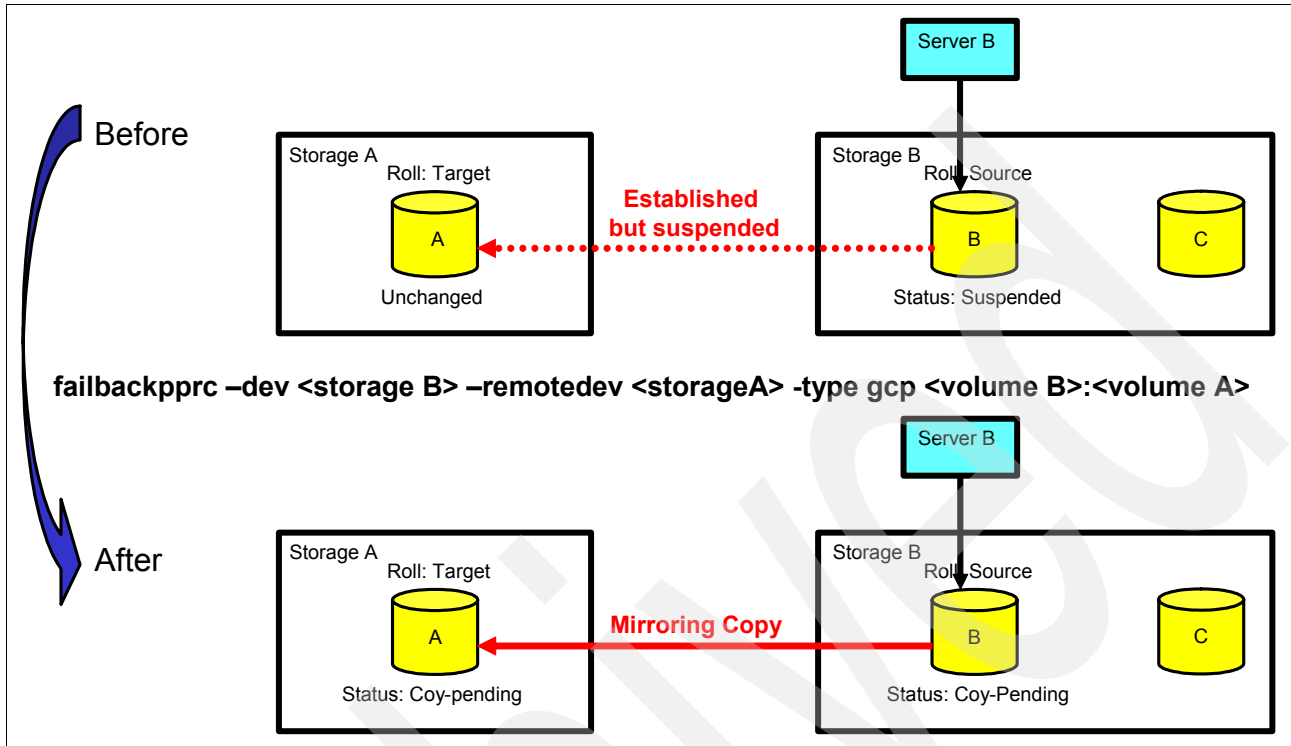


Figure 8-10 Failback Global Copy to a remote site

Example 8-35 shows the failback of three former Global Copy pairs. The former source volumes 0150, 0151, and 0152 are on site A. The former target volumes 0150, 0151, and 0152 (in a suspended state) are on site B. Failback is to site B. This example shows that the status of the Global Copy becomes Copy Pending, which means that the data is being copied asynchronously.

Example 8-35 Output of the failbackpprc command

```

dscli> failbackpprc -dev IBM.1750-13AAG8A -remotedev IBM.1750-13ABVDA -type gcp 0150:0150
0151:0151 0152:0152
Date/Time: October 28, 2005 9:20:34 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0150:0150 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0151:0151 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0152:0152 successfully failed back.
dscli>
dscli>
dscli> lspprc -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 28, 2005 9:21:08 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Copy Pending -      Global Copy 10597      Enabled  Disabled  invalid
-      01      300      Disabled      False
0151:0151 Copy Pending -      Global Copy 26189      Enabled  Disabled  invalid
-      01      300      Disabled      False
0152:0152 Copy Pending -      Global Copy 22163      Enabled  Disabled  invalid
-      01      300      Disabled      False

```

When the initial copy process is complete, the number of OutOfSyncTracks becomes almost 0.

### 8.4.2 Making the volumes available on the local site

To perform an IPL on the server from the volumes on the storage on the local site, the volumes must be available for read and write. To make the volumes on the local site available:

1. Turn off the server on the remote site. To switch the system to the local site, we recommend that you shut down the server on the remote site before the Global Copy relationship is suspended. Otherwise, the IPL of the server on the local site is abnormal, taking a longer time. If you do not shut down, review the application data for consistency. There is a possibility that some application data in the memory has *not* been written to the disk.

Follow the procedures for your site or use the PWRDWNSYS command. After completing the power down process in the server on the remote site, ensure with DS CLI that the number of OutOfSyncTracks is 0.

2. Failover the Global Copy to the local site. To make the previous source volumes available, enter the following command:

```

dscli failoverpprc -dev <source_storage_image_ID> -remotedev <target
storage_image_ID> -type gcp <source_volume_ID>:<target_volume_ID> ...

```

This command performs the following tasks:

- Terminates the previous Global Copy relationship
- Establishes the new Global Copy relationship
- Suspends the new Global Copy relationship

The `failoverpprc` command changes the previous target volumes to new source volumes and the status to Suspended, as shown in Figure 8-11. Thus, the server can access the new suspended source volumes for read and write.

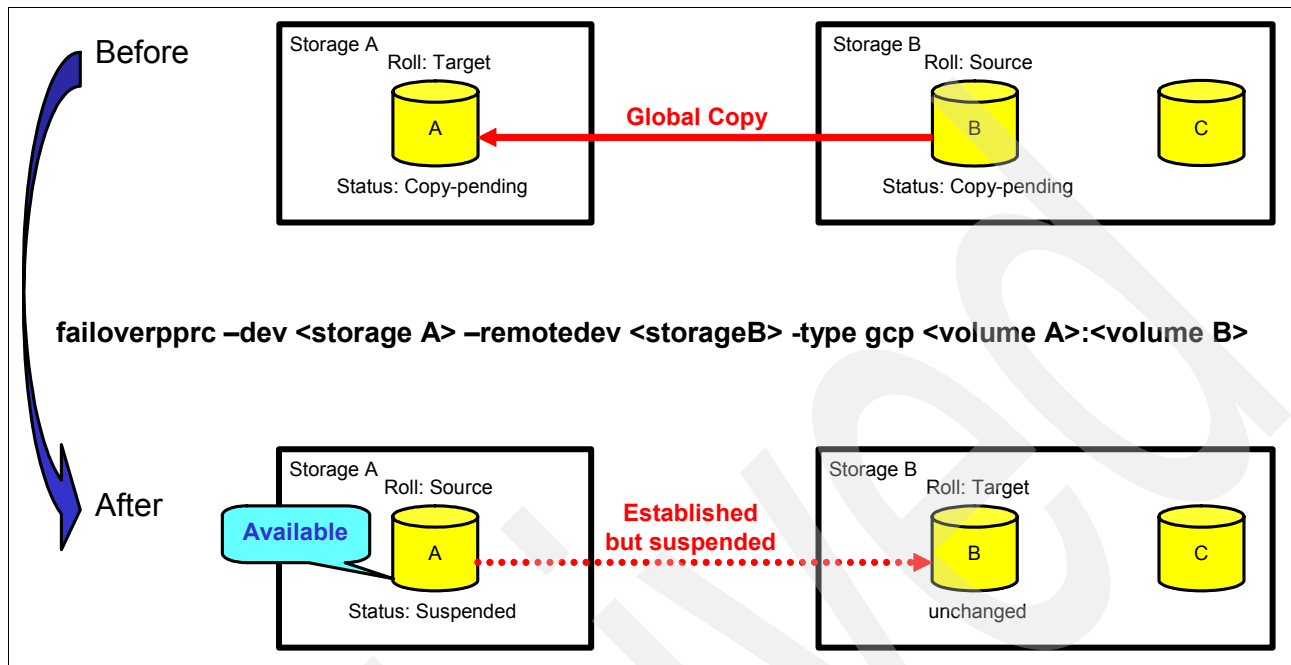


Figure 8-11 Failover Global Copy to a local site

Example 8-36 shows the failover of three Global Copy pairs. The source volumes 0150, 0151, and 0152 are on site B. The target volumes 0150, 0151, and 0152 are on site A. Failover is to site A.

*Example 8-36 Output of the failoverpprc command*

```

dscli> failoverpprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type gcp 0150:0150
0151:0151 0152:0152
Date/Time: October 28, 2005 9:38:48 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0150:0150 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0151:0151 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 0152:0152 successfully reversed.
dscli>
dscli>
dscli> lsprrc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 28, 2005 9:39:38 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State   Reason      Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
=====
0150:0150 Suspended Host Source Global Copy 0           Disabled Disabled  invalid
-          01 300      Disabled   True
0151:0151 Suspended Host Source Global Copy 0           Disabled Disabled  invalid
-          01 300      Disabled   True
0152:0152 Suspended Host Source Global Copy 0           Disabled Disabled  invalid
-          01 300      Disabled   True
dscli>
dscli>
dscli> lsprrc -l -dev IBM.1750-13AAG8A 0150-0152
Date/Time: October 28, 2005 9:39:50 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13AAG8A

```

ID	State	Reason	Type	Out Of Sync Tracks	Tgt Read Src Cascade	Tgt Cascade	
Date	Suspended	SourceLSS	Timeout (secs)	Critical Mode	First Pass	Status	
0150:0150	Copy	Pending	-	Global Copy 0	Enabled	Disabled	invalid
-	01	300		Disabled	True		
0151:0151	Copy	Pending	-	Global Copy 0	Enabled	Disabled	invalid
-	01	300		Disabled	True		
0152:0152	Copy	Pending	-	Global Copy 0	Enabled	Disabled	invalid
-	01	300		Disabled	True		

### 8.4.3 Starting Global Copy from the local site to remote site (original direction)

Start the Global Copy from the volumes on the local storage to the volumes on the remote storage and recreate the normal Global Mirror environment.

When you start the Global Copy, the target volumes become unavailable.

**Important:** Before you start Global Copy, ensure that the operating system of system B is in a state of shutdown. Otherwise, this operating system will hang.

With DS CLI, restart Global Copy from the new source volumes to the new target volumes by using the following command:

```
dsccli failbackpprc -dev <source_storage_image_ID> -remotedev <target_storage_image_ID> -type mmir -tgtread <source_volume_ID>:<target_volume_ID>...
```

**The -tgtread option:** In this command, the **-tgtread** option is required, because this Global Copy target volume is used as a FlashCopy source volume.

The **failbackpprc** command performs the following tasks:

- ▶ Checks the preserved state of the previous source volume to determine how much data to copy back.
- ▶ Copies either all the tracks or only OutOfSyncTracks from the volume on the remote storage.
- ▶ Copies the subsequent written data on the source volumes to the target volumes asynchronously, after the initial copy.

The `failbackpprc` command changes the previous source volumes to new target volumes, as shown in Figure 8-12. Therefore, the server *cannot* access the new target volumes for read and write.

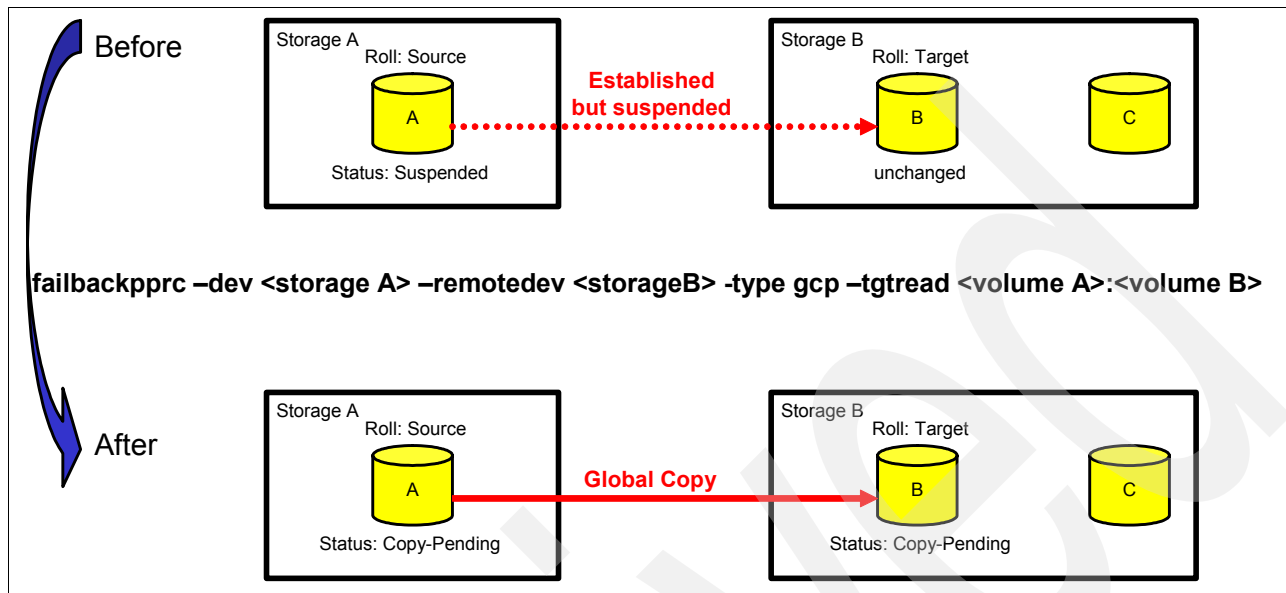


Figure 8-12 Failback Global Copy to a local site

Example 8-37 shows the failback of three former Global Copy pairs. The former source volumes 0150, 0151, and 0152 are on site B. The former target volumes 0150, 0151, and 0152 (in a suspended state) are on site A. Failback is to site A.

*Example 8-37 Output of the failbackpprc command*

```

dscli> failbackpprc -dev IBM.1750-13ABVDA -remotedev IBM.1750-13AAG8A -type gcp -tgtread
0150:0150 0151:0151 0152:0152
Date/Time: October 28, 2005 9:40:33 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0150:0150 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0151:0151 successfully failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 0152:0152 successfully failed back.
dscli>
dscli>
dscli> lspprc -l -dev IBM.1750-13ABVDA 0150-0152
Date/Time: October 28, 2005 9:41:14 AM JST IBM DSCLI Version: 5.0.6.142 DS: IBM.1750-13ABVDA
ID      State      Reason Type      Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade
Date Suspended SourceLSS Timeout (secs) Critical Mode First Pass Status
-----
0150:0150 Copy Pending - Global Copy 0          Enabled Disabled invalid
-          01         300          Disabled      True          Enabled Disabled invalid
0151:0151 Copy Pending - Global Copy 0          Enabled Disabled invalid
-          01         300          Disabled      True          Enabled Disabled invalid
0152:0152 Copy Pending - Global Copy 0          Enabled Disabled invalid
-          01         300          Disabled      True          Enabled Disabled invalid

```

## 8.4.4 Checking or restarting a Global Mirror session

Depending on whether the storage on the local site maintained the state of the Global Mirror session after the disaster, you might have to use either the `resumegmir` command or the `mkgmir` command. Alternatively, the state might be good, in which case, you do not have to enter any commands. To check and restart the Global Mirror session:

1. Check the Global Mirror session.
2. Resume the Global Mirror session.
3. Restart the Global Mirror session.

## 8.4.5 Performing an IPL of the production server on the local site

After you restart the Global Mirror session, perform an IPL on the production server from the volumes on the local site. If you did *not* change any settings on the system on the remote site and did *not* shut down the operating system before the failover of the Global Copy, perform an IPL on the system with normal mode.

As long as the physical hardware resources of the production server, such as Ethernet adapters, expansion enclosures, and so on, have *not* changed on the local site, the server detects the hardware resources that are associated with the line descriptions again. The operating system then varies on LIND and starts the TCP/IP interface addresses automatically. However, you must perform an IPL of the recovered server carefully, regardless of whether you changed settings on the system at the remote site. We recommend that you manually perform the IPL on the server to a restricted state before you allow users to access the application on the production server.

To perform an IPL of the production server to a restricted state, follow the steps in 6.2.4, “Performing an IPL of the target server” on page 259.

Archived





## Copy Services scenarios

In this chapter, we present the configurations and scenarios that we use in the examples in the remainder of this book.

## 9.1 Scenarios for System i using the DS GUI

In this section, we outline the scenarios that we use in this remainder of this book. These examples are designed to give you the information to set up simple test environments prior to configuring the production systems.

### 9.1.1 Scenario background

All of the scenarios in this book use the DS Storage Manager GUI. In some cases, DS command-line interface (CLI) commands are also shown because there is no GUI equivalent.

We used the DS Storage Manager GUI from DS8000 Release 3 which provides Javascript support for better response times and transparent page refresh. The release 3 GUI has a better look and feel and improved handling like the ability to scroll through larger tables instead of having to page through them. It also has some changes in functionality in its panels to cover new configuration related functions of DS8000 Release 3 such as space efficient FlashCopy and storage pool striping.

### 9.1.2 Test scenarios

We used the DS Storage Manager GUI to configure and manage the following environments:

- ▶ FlashCopy
- ▶ Metro Mirror, formerly known as synchronous Peer to Peer Remote Copy (PPRC)
- ▶ Global Mirror, formerly known as asynchronous PPRC Extended Distance (PPRC-XD)

### 9.1.3 Accessing the DS GUI interface

This section provides information about the DS GUI interface.

#### Accessing the DS6000 Storage Manager GUI

For the DS6000, you can install the DS Storage Manager GUI from the installation CD that comes with the storage system or download it from the Web at:

<http://www.ibm.com/servers/storage/support/disk/ds6800/downloading.html>

The DS6000 Storage Manager is installed on the Systems Management Console (SMC).

**Restriction:** Only one DS6000 Storage Management full management console can be installed per DS6000 storage unit. Additional consoles must be offline management consoles. This restriction does not apply to the DS8000 because the user accesses the DS Storage Manager application running on the Hardware Management Console (HMC) using a Web browser or the System Storage Productivity Center (SSPC).

#### Accessing the DS8000 Storage Manager GUI

The DS Storage Manager application is preloaded for the IBM System Storage DS8000 series. You can access it locally either using a Web Browser installed on the DS8000 Hardware Management Console (HMC) or using the System Storage Productivity Center (SSPC).

For DS8000 systems *without* SSPC installed, you access the DS Storage Manager GUI remotely using a Web browser pointing to the HMC as follows:

- ▶ For a *non-secure* HTTP connection to the HMC, enter the following URL:

`http://HMC_IP_address:8451/DS8000/Login`

- ▶ For a *secure* HTTPS connection to the HMC, enter the following URL:

`https://HMC_IP_address:8452/DS8000/Login`

For DS8000 systems *with* SSPC installed, you access the DS Storage Manager GUI remotely using a Web browser pointing to the SSPC using the following procedure:

1. Access the SSPC using your Web browser at the following URL (see Figure 9-1):

`http://SSPC_IP_address:9550/ITSRM/app/en_US/index.html`

2. Click **TPC GUI (Java™ Web Start)** to launch the TPC GUI.

**Note:** The TPC GUI requires an IBM 1.4.2 JRE™. Select one of the IBM 1.4.2 JRE links (shown in Figure 9-1) to download and install it based on your OS platform.

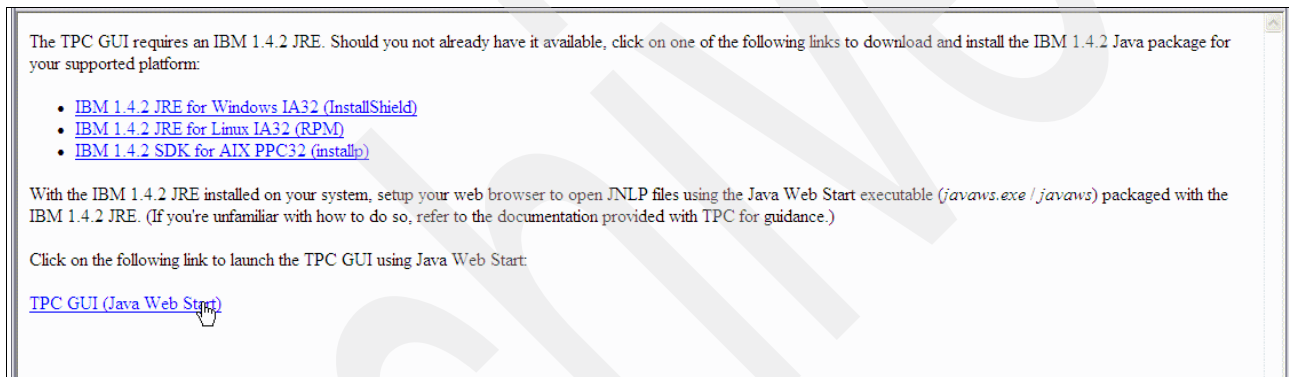


Figure 9-1 Index page of System Storage Productivity Center

3. The TPC GUI window display, as shown in Figure 9-2. Enter the information of user ID, password, and the SSPC server. Click **OK** to continue.

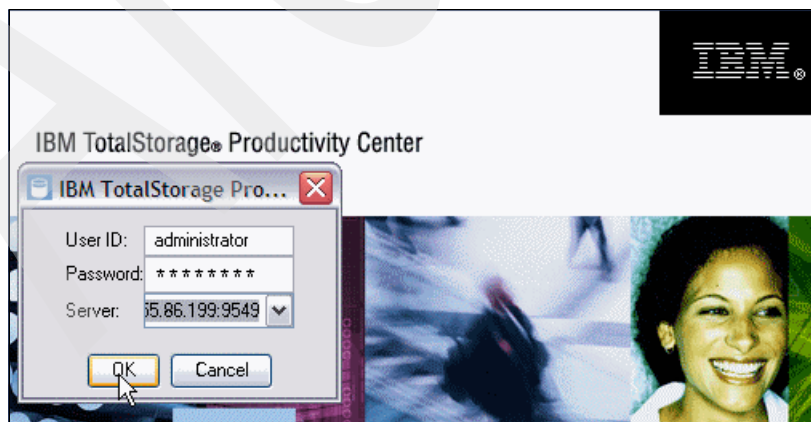


Figure 9-2 TCP GUI Sign On panel

- Click **Element Management** to get a list of element managers from DS8000 machines administrated by the SSPC as shown in Figure 9-3.

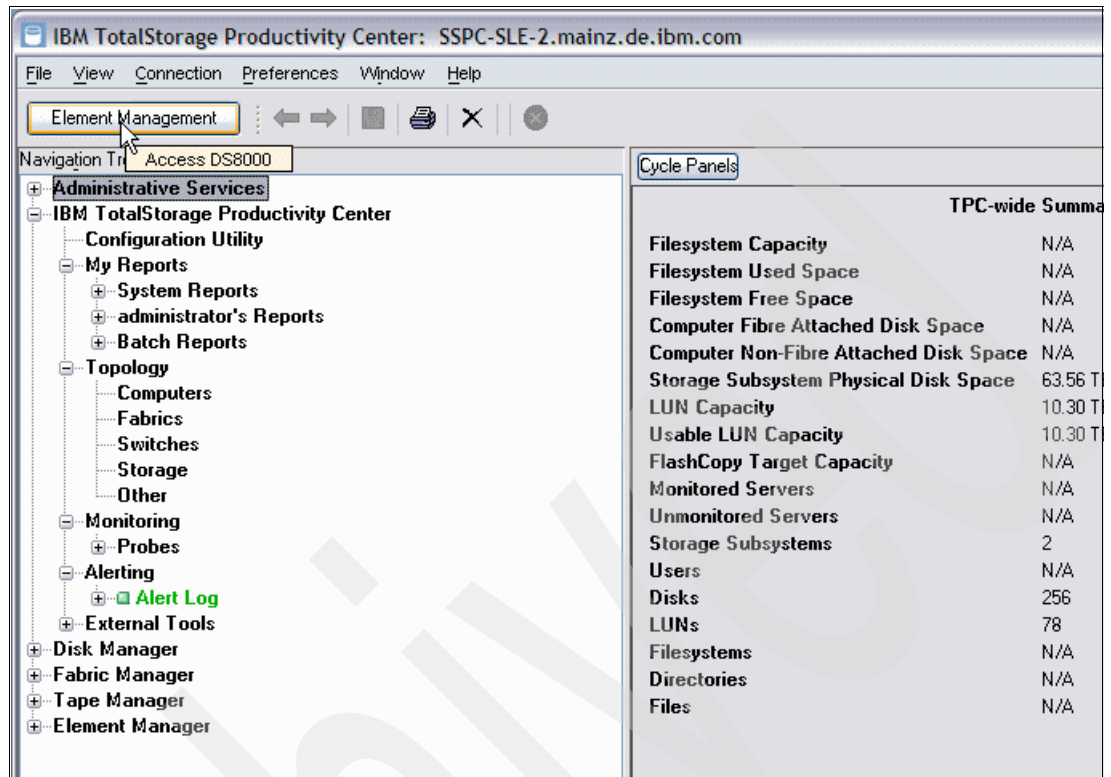


Figure 9-3 TPC GUI Enterprise Management window

- The Element Management of TPC GUI displays. Click one of the DS8000 machines to access its DS Storage Manager GUI (see Figure 9-4).

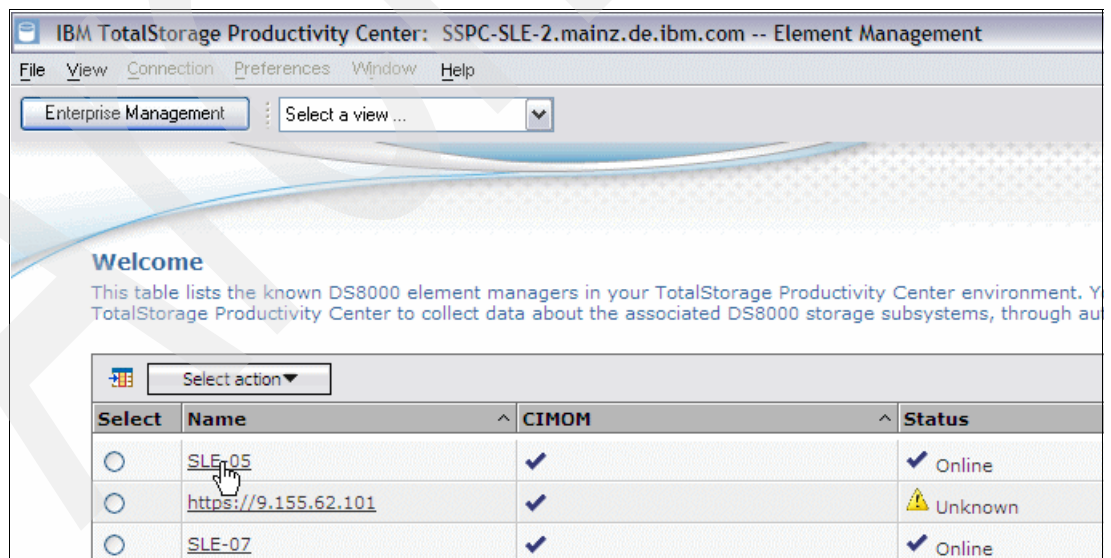


Figure 9-4 TPC GUI Element Management window

6. The DS8000 Storage Manager Welcome panel displays as shown in Figure 9-5.

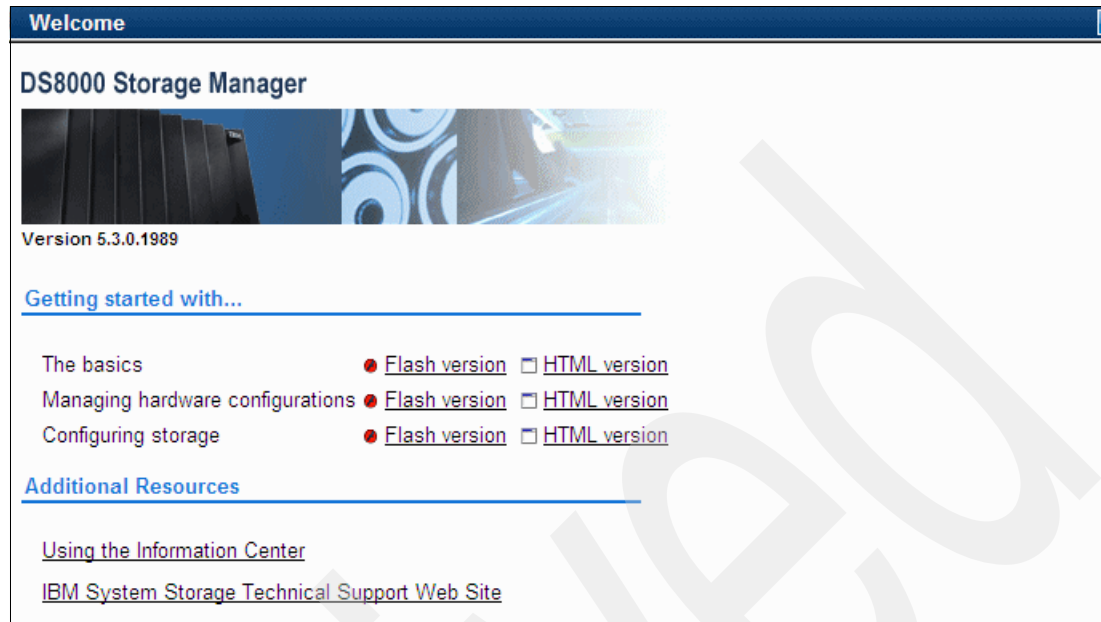


Figure 9-5 DS8000 Storage Manager Welcome panel

#### 9.1.4 System i5 models

We used two System i5 models for all of our testing. These models ran i5/OS V6R1. The first system, called *Reds*, is a System i5 model 9406 570. The second system, called *Expo*, is a System i5 model 9406 520. Both of these systems have two #2847 IOP-based Fibre Channel IOAs for boot from SAN and for i5/OS V6R1 load source multipath support.

#### 9.1.5 IBM System Storage server

For our testing, we used a DS8300 LPAR model with two independent storage facility images (SFIs)—storage image 51 and storage image 52. These two SFIs on a DS8000 LPAR model behave similar to two physically different DS8000 non-LPAR machines. We utilized this DS8000 LPAR concept for our examples to document the procedures for implementing Metro Mirror or Global Mirror using one DS8000 machine only as though it were two physically different DS8000 machines. Of course for a real production environment to provide disaster recovery protection Metro Mirror or Global Mirror should always be implemented between two physical IBM System Storage disk subsystems which are ideally at two different data center locations.

The System i5 model in the Reds environment used logical unit numbers (LUNs) on storage image 51 on the DS8300. The System i5 model in the Expo environment used LUNs on storage image 52. See Figure 9-6.




Figure 9-6 Configuration used in testing for this book

### 9.1.6 Test setup

Table 9-1 and Table 9-2 describe the configuration objects that we used to perform our tests. Both of the systems use multipath load source and have pair volumes in another LSS to support FlashCopy.

Table 9-1 Test configuration objects for Expo (storage image 52)

<b>Device adapter</b>	<ul style="list-style-type: none"> <li>▶ ExpoBS1</li> <li>▶ ExpoBS2</li> </ul>
<b>Volume group</b>	<ul style="list-style-type: none"> <li>▶ ExpoLS_VG</li> </ul>
<b>Logical subsystem (LSS)</b>	10
<b>Volumes (nicknames and IDs)</b>	<ul style="list-style-type: none"> <li>▶ ExpoLS0001 101A</li> <li>▶ ExpoLS0002 101B</li> <li>▶ ExpoLS0003 101C</li> <li>▶ ExpoLS0004 101D</li> </ul>
<b>Volume group</b>	<ul style="list-style-type: none"> <li>▶ ExpoFC_VG</li> </ul>
<b>Logical subsystem</b>	21
<b>Volumes (nicknames and IDs)</b>	<ul style="list-style-type: none"> <li>▶ ExpoFC0001 2100</li> <li>▶ ExpoFC0002 2101</li> <li>▶ ExpoFC0003 2102</li> <li>▶ ExpoFC0004 2103</li> </ul>

Table 9-2 Test configuration objects for Reds (storage image 51)

<b>Device adapter</b>	<ul style="list-style-type: none"> <li>▶ RedsBS1</li> <li>▶ RedsBS2</li> </ul>
<b>Volume group</b>	▶ RedsBookTN1LS_VG
<b>Logical subsystem</b>	10
<b>Volumes (nicknames and IDs)</b>	<ul style="list-style-type: none"> <li>▶ TN1ls 105E</li> <li>▶ TN1Vol1 105F</li> <li>▶ TN1Vol2 1060</li> <li>▶ TN1Vol3 1061</li> </ul>
<b>Volume group</b>	RedBookTN1MM_VG
<b>Logical subsystem</b>	21
<b>Volumes (nicknames and IDs)</b>	<ul style="list-style-type: none"> <li>▶ TN1FCls 2100</li> <li>▶ TN1FCVol1 2101</li> <li>▶ TN1FCVol2 2102</li> <li>▶ TN1FCVol3 2103</li> </ul>

## 9.2 FlashCopy scenario

In our implementation, we used FlashCopy to copy from the volumes (LUNs) in storage image 51 (volumes for the *Reds* system) to a similar set of volumes (LUNs) in the same storage system. See Figure 9-7.




Figure 9-7 FlashCopy environment for testing

## 9.3 Metro Mirror scenario

Implementation of Metro Mirror for the entire DASD space involves the following tasks:

1. Creating the PPRC paths between primary and secondary LSSs.
2. Creating the Metro Mirror volume relationships.
3. Switching over the system from the local site to the remote site.
  - Making the volumes available on the remote site.
  - Performing an IPL of the backup server on the remote site.

4. Switching back the system from the remote site to the local site.
  - Starting Metro Mirror in the reverse direction, from the remote site to the local site.
5. Making the volumes on the local site available.
  - Performing the IPL of the production server on the local site.
  - Starting Metro Mirror in the original direction, from the local site to the remote site.

In our test environment, the primary system Reds was connected to storage image 51, and the Metro Mirror targets for the Expo backup system were on storage image 52. See Figure 9-8.

The Reds system is the production server, and the Expo system is the backup server. The Reds system is connected to storage image 51, and an IPL is performed from external storage using boot from SAN. Storage image 51 and storage image 52 are connected with Fibre Channel cables as though they were two physically different DS8300 machines. This implementation example assumes that the Reds system and storage image 51 are on local site A. It also assumes that the Expo system and storage image 52 are on remote site B.

In this example, the Metro Mirror environment is created between storage image 51 and storage image 52. The business application is then switched over from local site A to remote site B. Finally, the business application is switched back from remote site B to local site A.




Figure 9-8 Metro Mirror environment for testing



## 9.4 Global Mirror scenario

In our scenario, we used the Reds system, with its LUNs in storage image 51, as the source system. We used Global Mirror to asynchronously copy the data to storage image 52. Global Mirror then uses FlashCopy to make sure a consistent copy of data is always on storage image 52. See Figure 9-9.

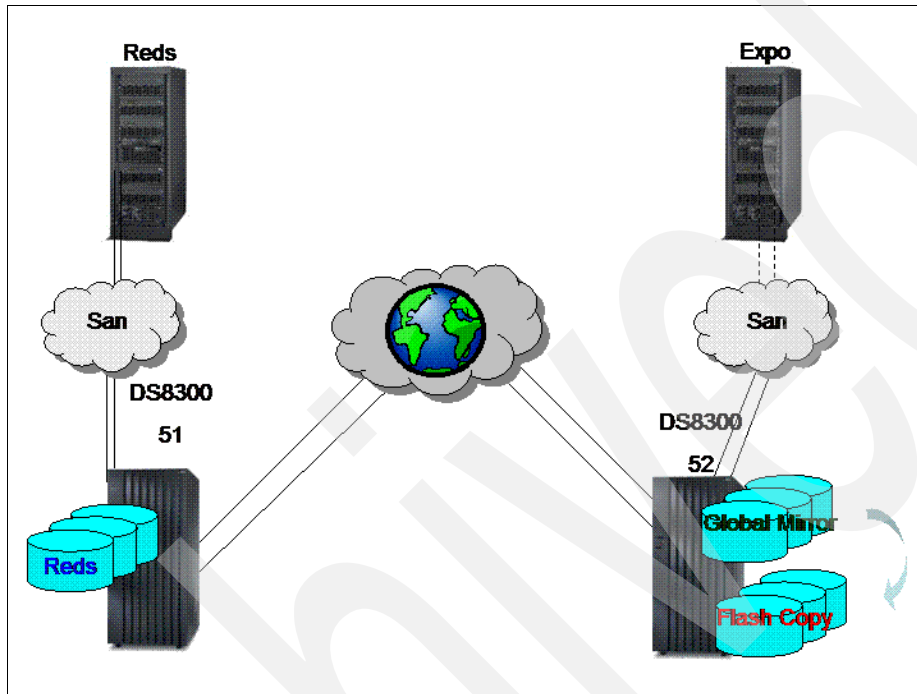


Figure 9-9 Global Mirror environment for testing

Archived

## Creating storage space for Copy Services using the DS GUI

Copy Services is an optional feature of the IBM System Storage DS6000 and DS8000. It brings powerful data copying and mirroring technologies to open system environments that were previously available only for the mainframe system.

In this chapter, we describe the steps that are required to configure the target LUNs using the Disk Storage (DS) GUI assuming an IBM System i server or LPAR with i5/OS is already attached to external storage.

For planning and implementing IBM System Storage disk subsystems for System i, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

Before you implement Copy Services, you must create the volumes that is used as the copy targets.

**Note:** In addition to creating the volumes, you also need to create arrays and ranks. However, we do not cover this topic in this book. Refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120.

This chapter describes the procedure to create volumes with the DS Storage Manager GUI using the following steps:

- ▶ Creating an extent pool
- ▶ Creating logical volumes
- ▶ Creating a volume group

## 10.1 Creating an extent pool

To create an extent pool, follow these steps:

1. Access the DS GUI as described in 9.1.3, “Accessing the DS GUI interface” on page 338 and sign on using an administrator user name and password.



Figure 10-1 DS8000 Sign On panel

- The main DS Storage Manager window is the starting point for all of your configuration, management, and monitoring needs for the DS disk and Copy Services tasks (Figure 10-2). The underlying hardware (two System p 570 models), I/O drawers, and I/O adapters are controlled by the storage HMC built into the DS8000 rack or a standalone desktop SMC.

From any of the GUI panels, you can access the Information Center by clicking the question mark (?) in the upper-right corner of the page.

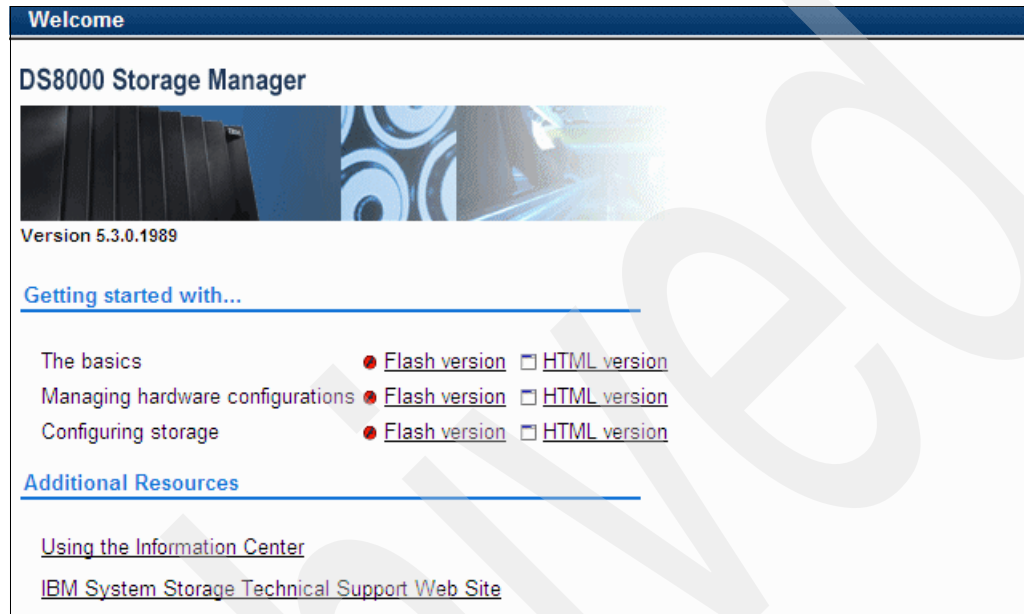


Figure 10-2 DS8000 Storage Manager: Getting started

- Next, create an extent pool to be assigned to a rank. You can check the availability of a rank with select **Real-time manager** → **Configure storage** → **Ranks**. In the panel on the right (see Figure 10-3), the rank R23 is listed as unassigned (that is no extent pool exists for this rank).

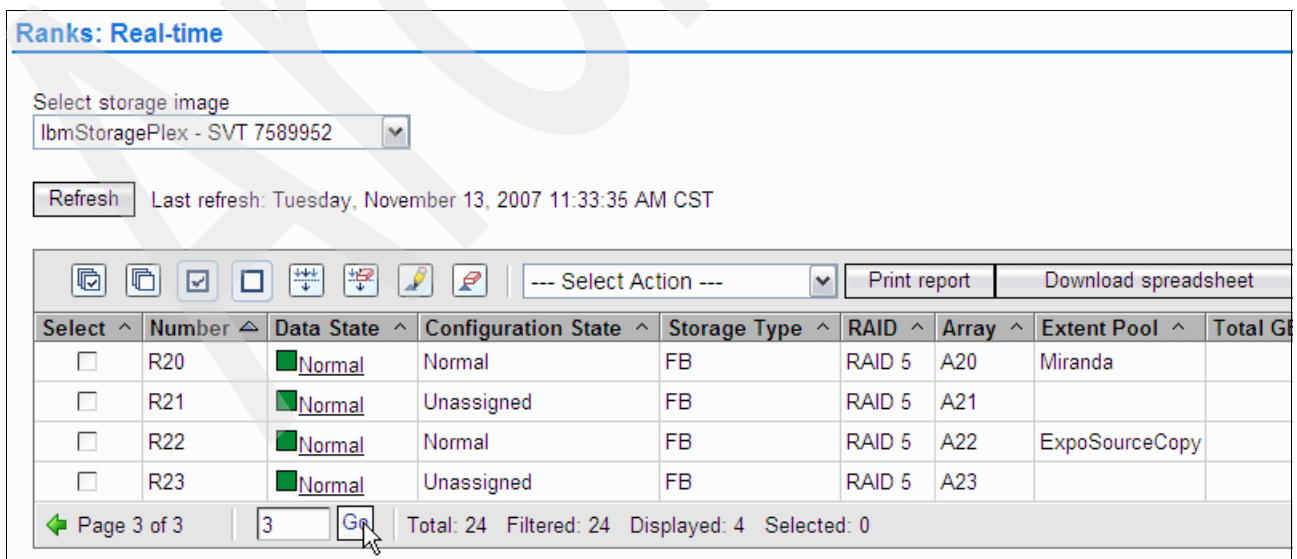


Figure 10-3 Check the availability of rank

4. Create a new extent pool for rank R23. Select **Real-time manager** → **Configure storage** → **Extent pools**, and select **Create New Extent Pools** from the Select action drop-down menu (see Figure 10-4).

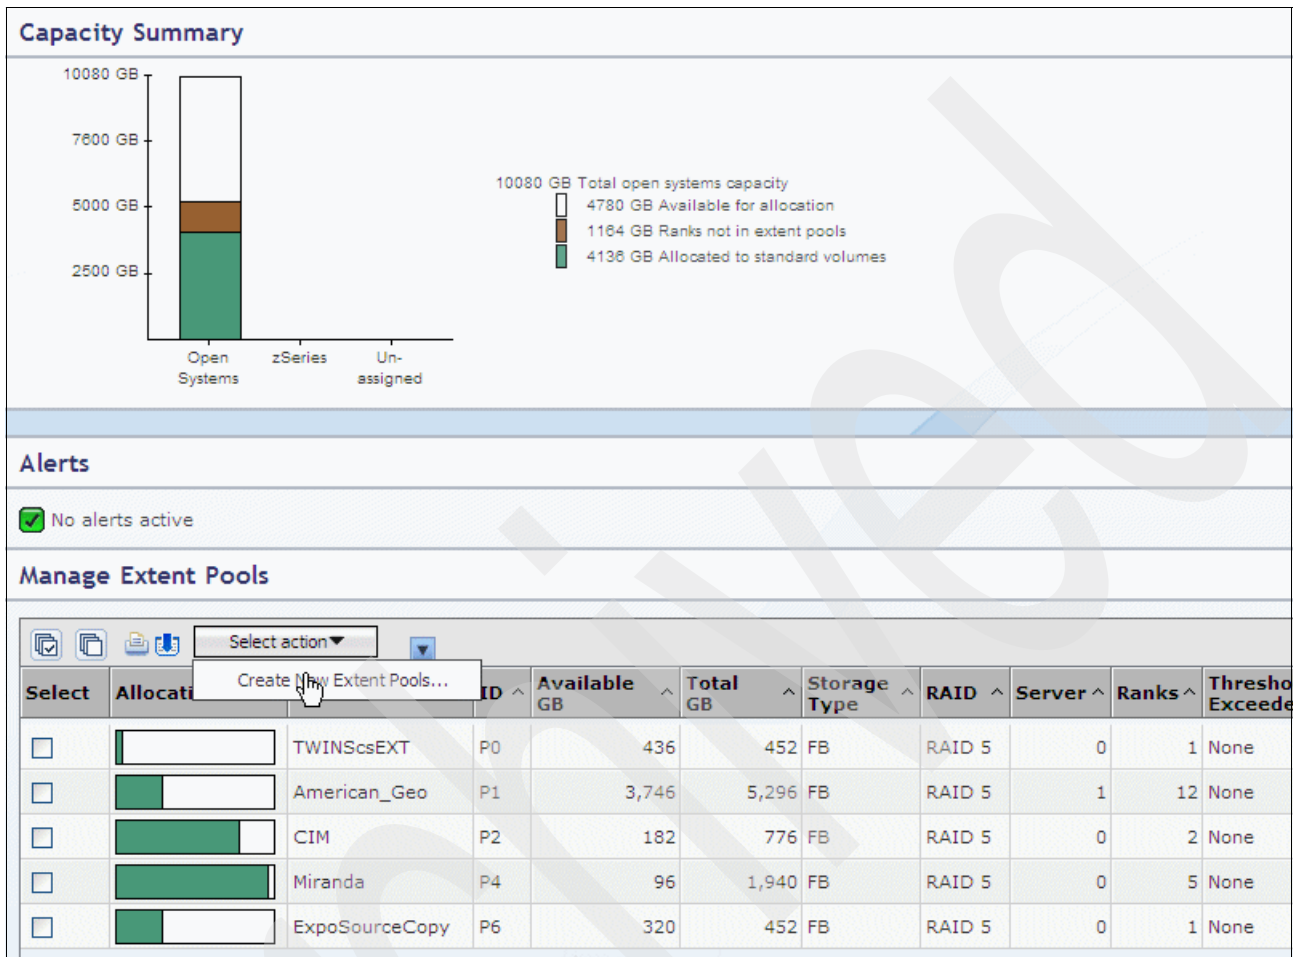


Figure 10-4 DS8000 Storage Manager Extent pools panel

5. The Create New Extent Pools panel displays (as shown in Figure 10-5 and Figure 10-6). In this panel:
  - a. Select FB for **Storage Type**, and select the **RAID Type** according to the type of RAID protection that is chosen for the arrays/ranks created before (see Figure 10-5).
  - b. Choose Manual for **Type of Configuration**. All of the ranks that have not been allocated to any extent pools will be displayed in the table (Figure 10-5). Choose only one of the ranks that are available in the table (R23 in our example).
  - c. Scroll down your window to see another option in the panel (see Figure 10-5 and Figure 10-6).

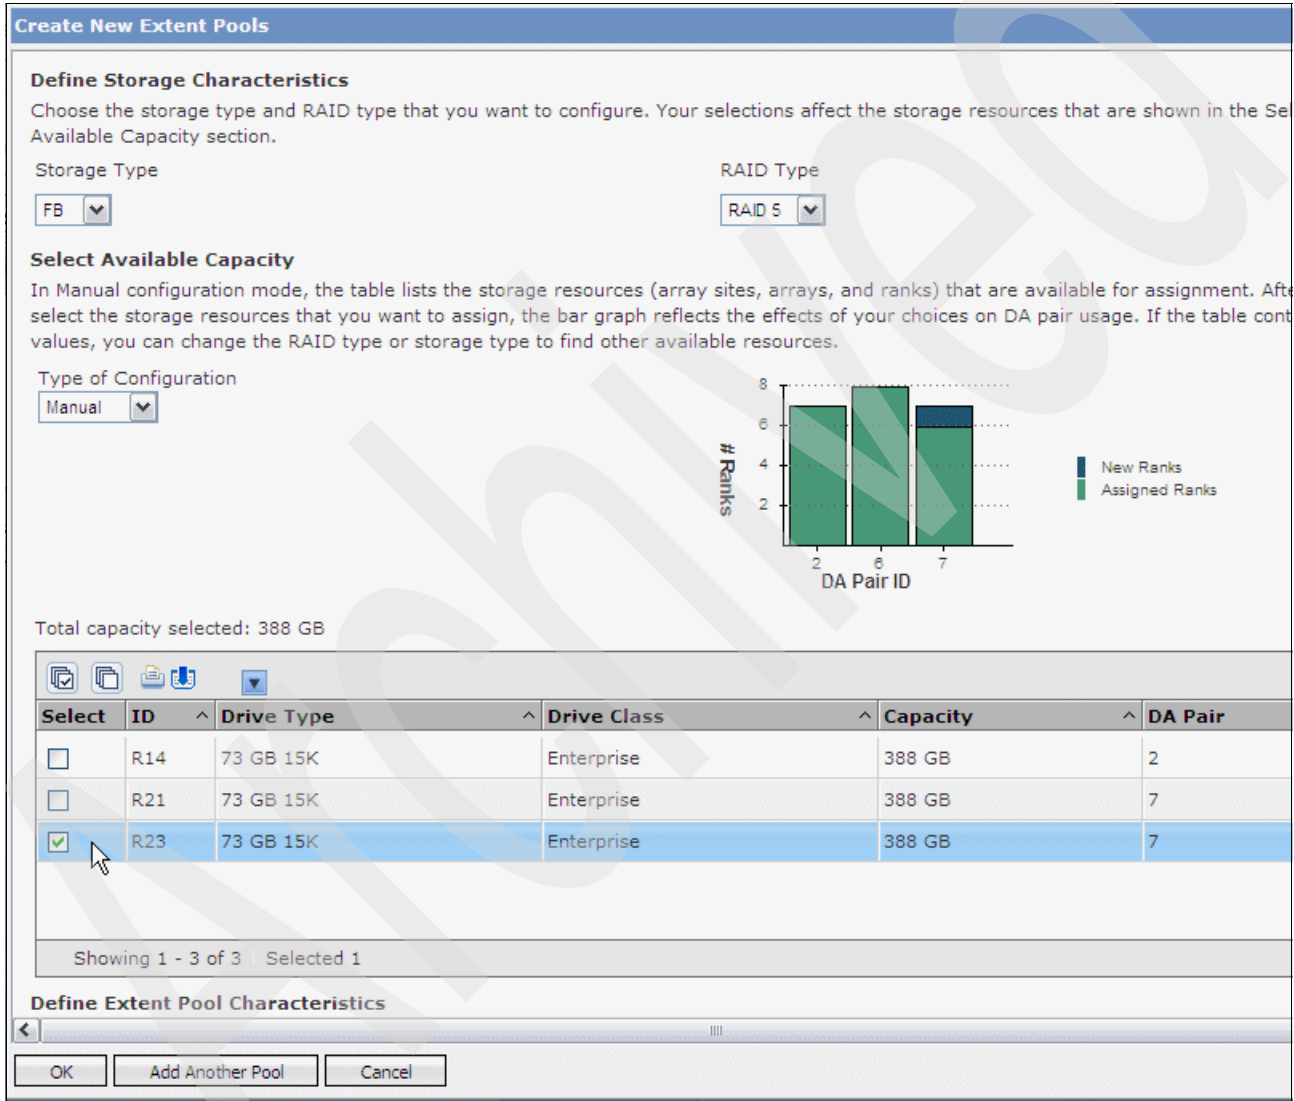


Figure 10-5 DS8000 Storage Manager Create New Extent Pools panel (upper)

- d. Select Single extent pool for **Number of extent pools**. Give a descriptive pool name prefix. Here we use *ExpoTargCopy* so that we can identify how this extent pool is used. Use **100** for **Storage Threshold** and **0** for **Storage Reserved** (see Figure 10-6).
- e. Select the server with which this extent pool to be associated for **Server Assignment**, and select **Add Another Pool** for continuing creating other extent pools or select **OK** to create only this extent pool (see Figure 10-6).

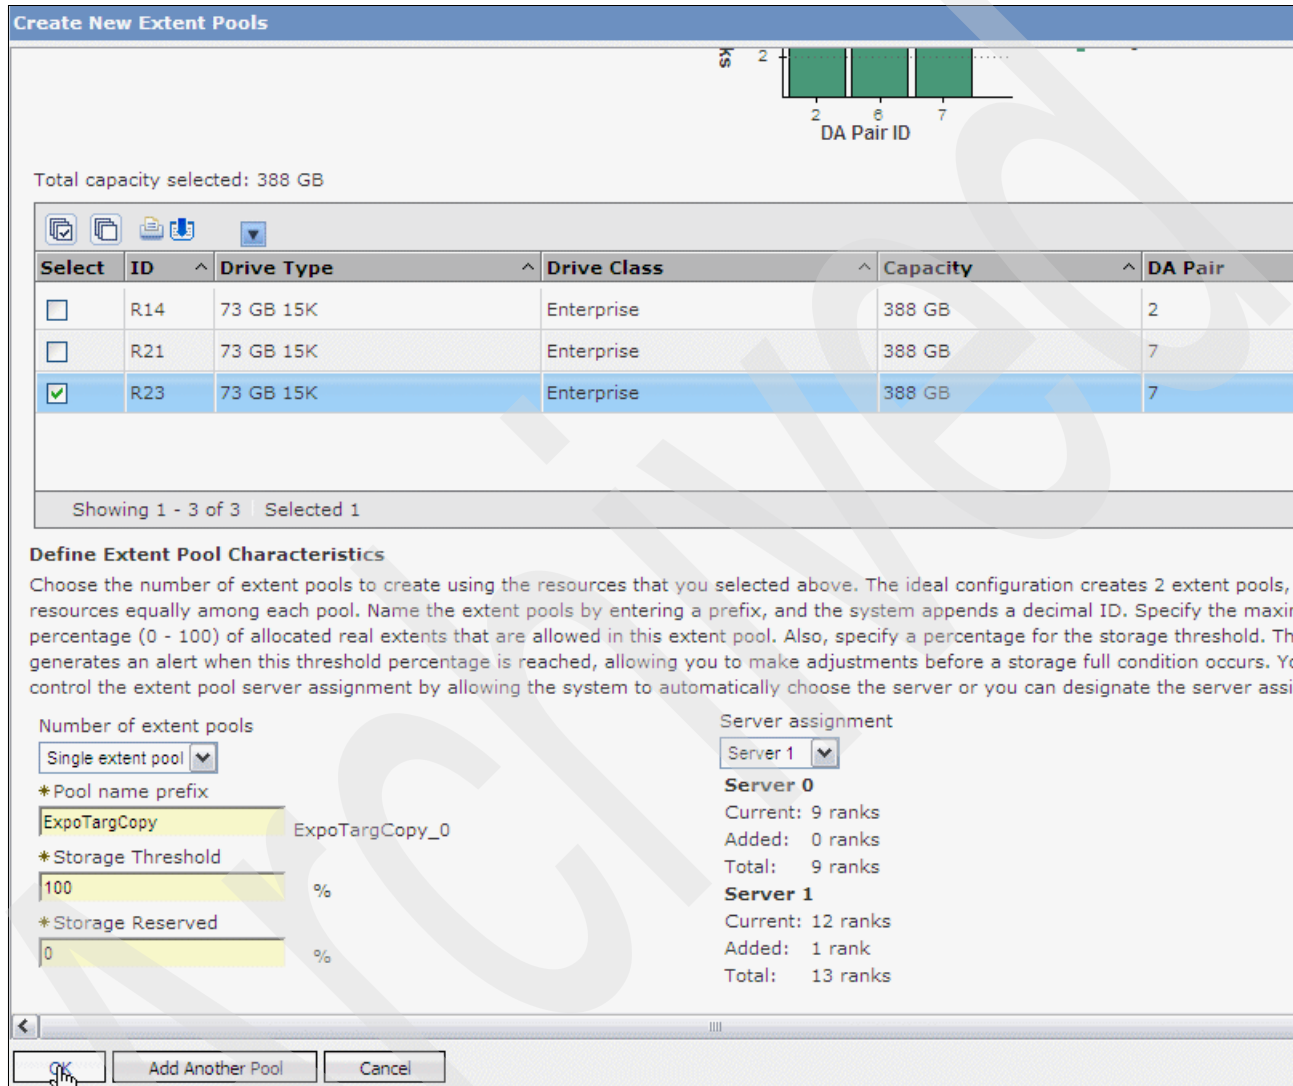


Figure 10-6 DS8000 Storage Manager Create New Extent Pools panel (lower)



6. In the Verification panel (Figure 10-7), review the information and verify whether everything is correct. If it is correct, click **Create All** to create the extent pool.



Figure 10-7 Verifying and confirming the creation of the extent pool

7. Depending on the size of the extent pool that you create, you might see a panel that shows a creating extent pools task for some time (shown in Figure 10-8).




Figure 10-8 Creating extent pools task message panel

8. Select **Real-time manager** → **Configure storage** → **Ranks**, and select the storage image 52 to see relationship between our newly created extent pool (*ExpoTargCopy\_0*) and R23 (see Figure 10-9).

**Ranks: Real-time**

Select storage image  
 IbmStoragePlex - SVT 7589952

Last refresh: Tuesday, November 13, 2007 11:45:42 AM CST

Select ^	Number ^	Data State ^	Configuration State ^	Storage Type ^	RAID ^	Array ^	Extent Pool ^
<input type="checkbox"/>	R20	<input checked="" type="checkbox"/> Normal	Normal	FB	RAID 5	A20	Miranda
<input type="checkbox"/>	R21	<input checked="" type="checkbox"/> Normal	Unassigned	FB	RAID 5	A21	
<input type="checkbox"/>	R22	<input checked="" type="checkbox"/> Normal	Normal	FB	RAID 5	A22	ExpoSourceCopy
<input type="checkbox"/>	R23	<input checked="" type="checkbox"/> Normal	Normal	FB	RAID 5	A23	ExpoTargCopy_0

Total: 24 Filtered: 24 Displayed: 4 Selected: 0

Figure 10-9 Viewing the ranks and associated extent pool

## 10.2 Creating logical volumes

To create logical volumes, follow these steps:

1. After creating the extent pool, we need to create volumes (LUNs) within our newly created extent pool. Select **Real-time manager** → **Configure storage** → **Open systems** → **Volumes - Open systems**, and select the appropriate storage image. Select **Create** from **Select Action** drop-down menu (see Figure 10-10).

**Volumes - Open systems: Real-time**

Select storage image:  Select secondary filter:

Last refresh: Tuesday, November 13, 2007 11:52:27 AM CST

Select	Nickname	ID	Status	Type	GB(2^30)	GB(10^9)	Blocks	RAID
<input type="checkbox"/>	TWINScs0001	1000	Normal	I-unpr	8.0	8.6	16,777,216	RAID 5
<input type="checkbox"/>	TWINScs0002	1001	Normal	I-unpr	8.0	8.6	16,777,216	RAID 5
<input type="checkbox"/>	CIM0001	1002	Normal	I-prot	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0002	1003	Normal	I-prot	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0003	1004	Normal	I-prot	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0004	1005	Normal	I-protected	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0005	1006	Normal	I-protected	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0006	1007	Normal	I-protected	32.8	35.2	68,681,728	RAID 5
<input type="checkbox"/>	CIM0007	1008	Normal	I-protected	65.7	70.6	137,822,208	RAID 5
<input type="checkbox"/>	CIM0008	1009	Normal	I-protected	65.7	70.6	137,822,208	RAID 5

Page 1 of 6   Total: 51 Filtered: 51 Displayed: 10 Selected: 0

Figure 10-10 Working with open systems volumes

- In the Select extent pool panel, select the newly created extent pool as shown in Figure 10-11.

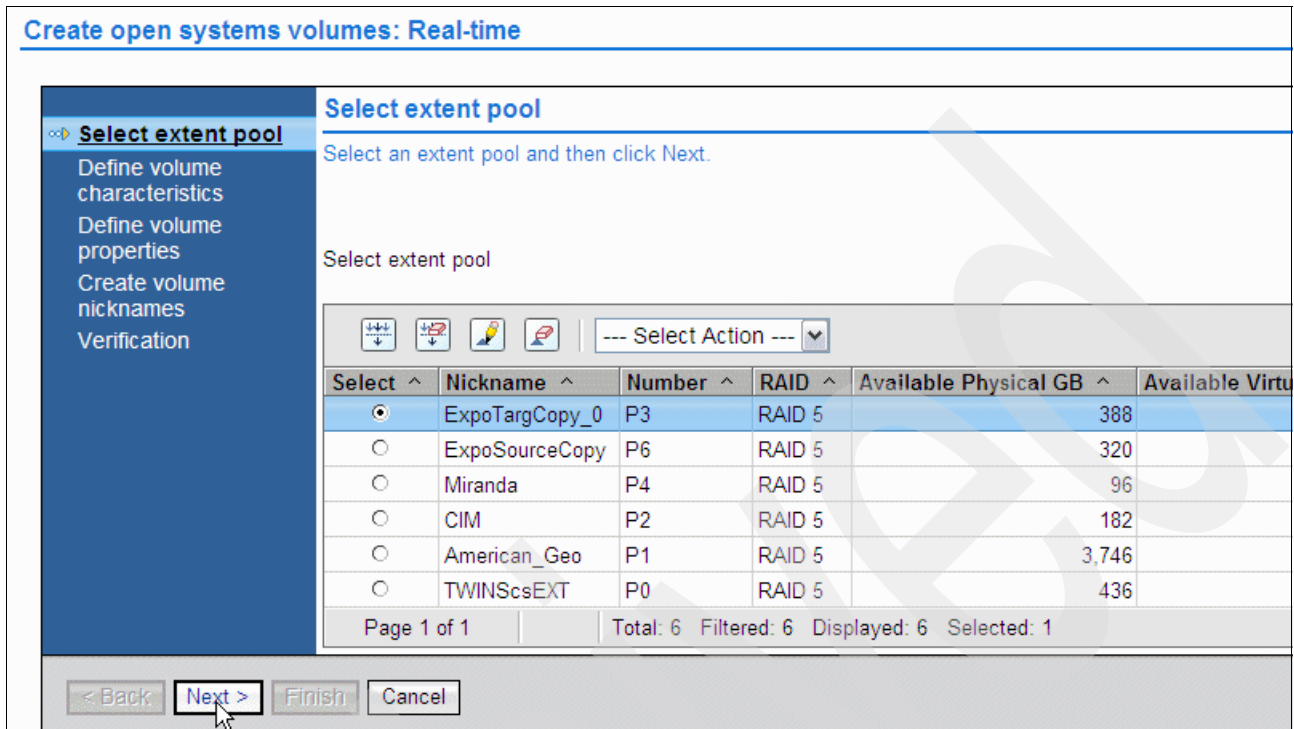


Figure 10-11 Selecting an extent pool

- Create the protected LUNs for the load source unit and all other LUNs by selecting **iSeries - Protected** for the Volume type.

**Note:** If your external load source is mirrored, for example, to provide path protection when using an older i5/OS version before V6R1, select **iSeries - Unprotected** only for creating the mirrored load source target volumes.

Because you have not created any volume groups, do not select any volume groups from the “Select volume groups” option. Select the default value for the “Extent allocation method” option, as shown in Figure 10-12. That is, do not use the rotate extents storage pool striping function (refer to 3.2.6, “Planning for capacity” on page 67).

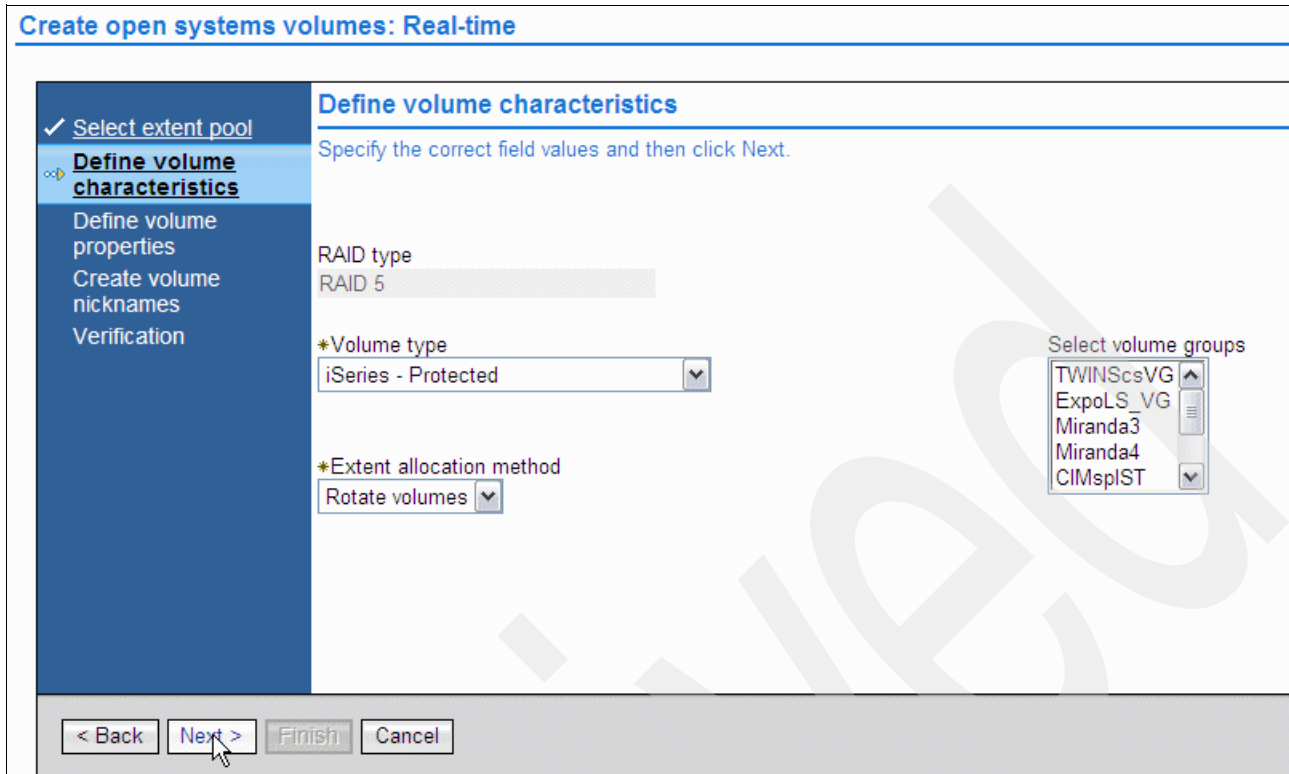


Figure 10-12 Define volume characteristics

- Specify to create the logical volumes (LUNs) by entering the information for Quantity, Size, and LSS, and then click **Next** to continue.

It is possible to create more than one volume at a time. In this example, we create four volumes. Because these LUNs are for an i5/OS environment only fixed LUN sizes are available. In our case, we are using 35.16 GB LUNs (see Figure 10-13). We also associate them with the logical subsystem (LSS) 0x21. Remember that the LSS is important when planning to use Metro Mirror or Global Mirror as it should preferably be the same for source and target volumes to help ease the administration.

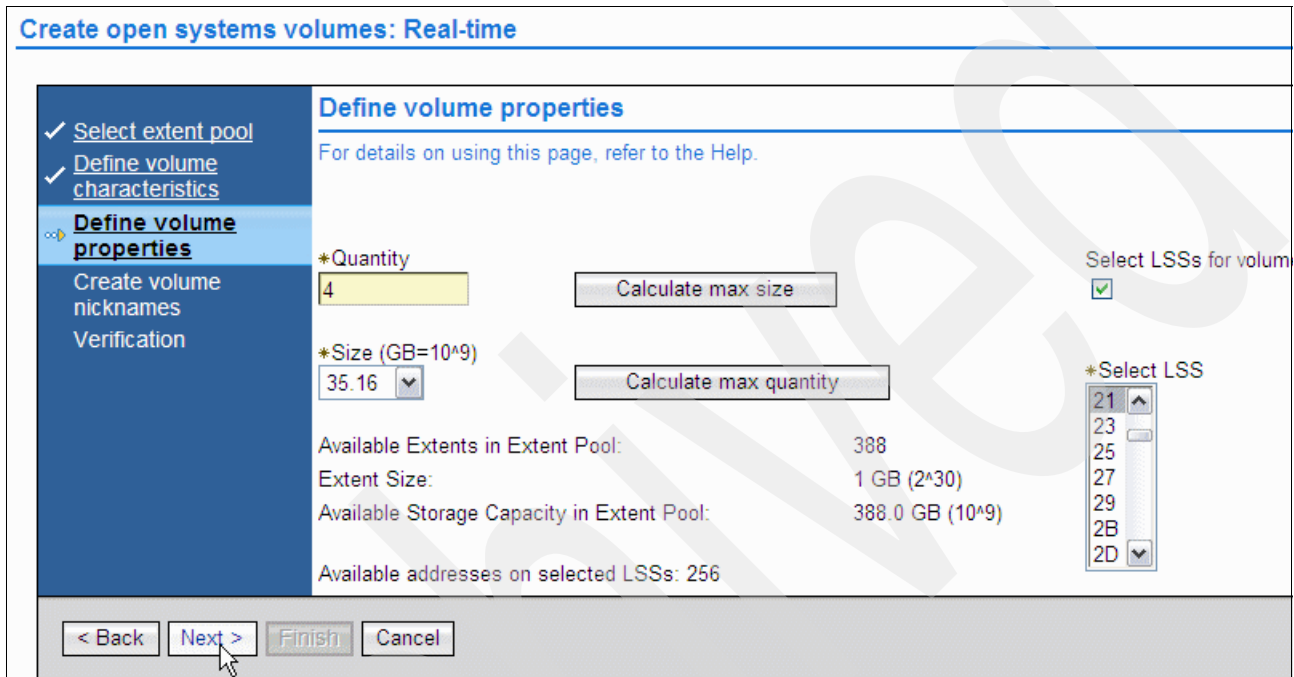


Figure 10-13 Define volume properties

- Define the naming convention to be used for the volumes. We use ExpoFC as the prefix because these LUNs are used for the FlashCopy of the Expo system (see Figure 10-14).

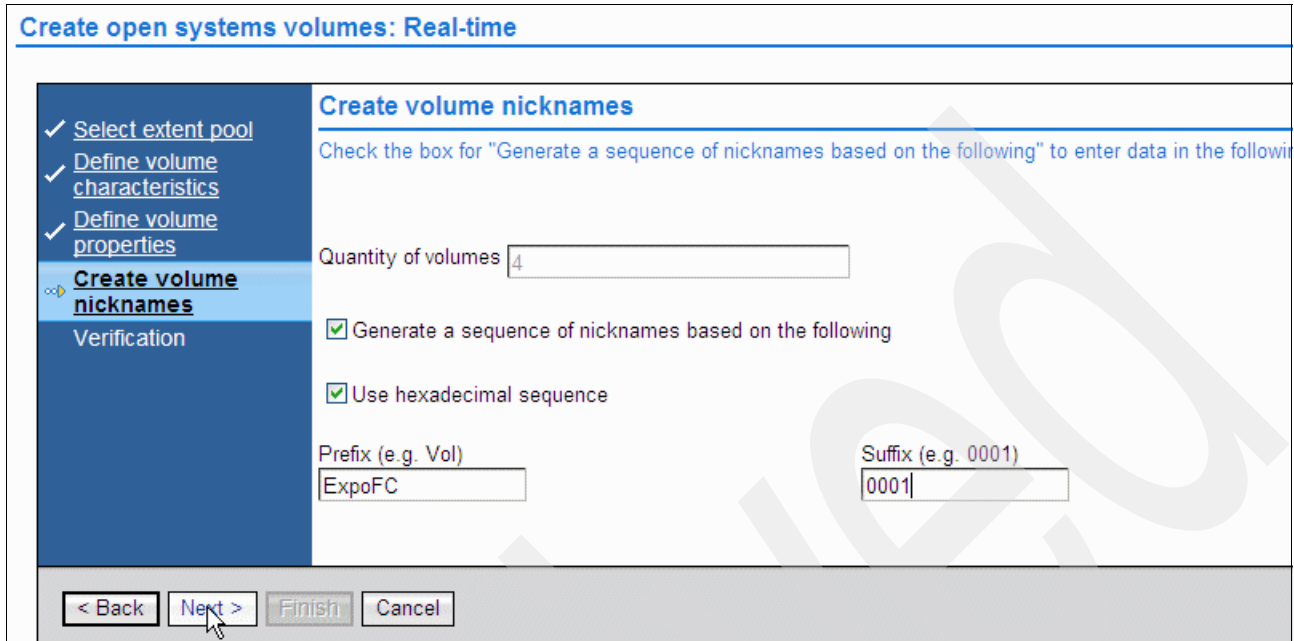


Figure 10-14 Creating volume nicknames

- The Verification panel displays as shown in Figure 10-15. Review the information, and click **Finish** to actually start the logical volume creation process.

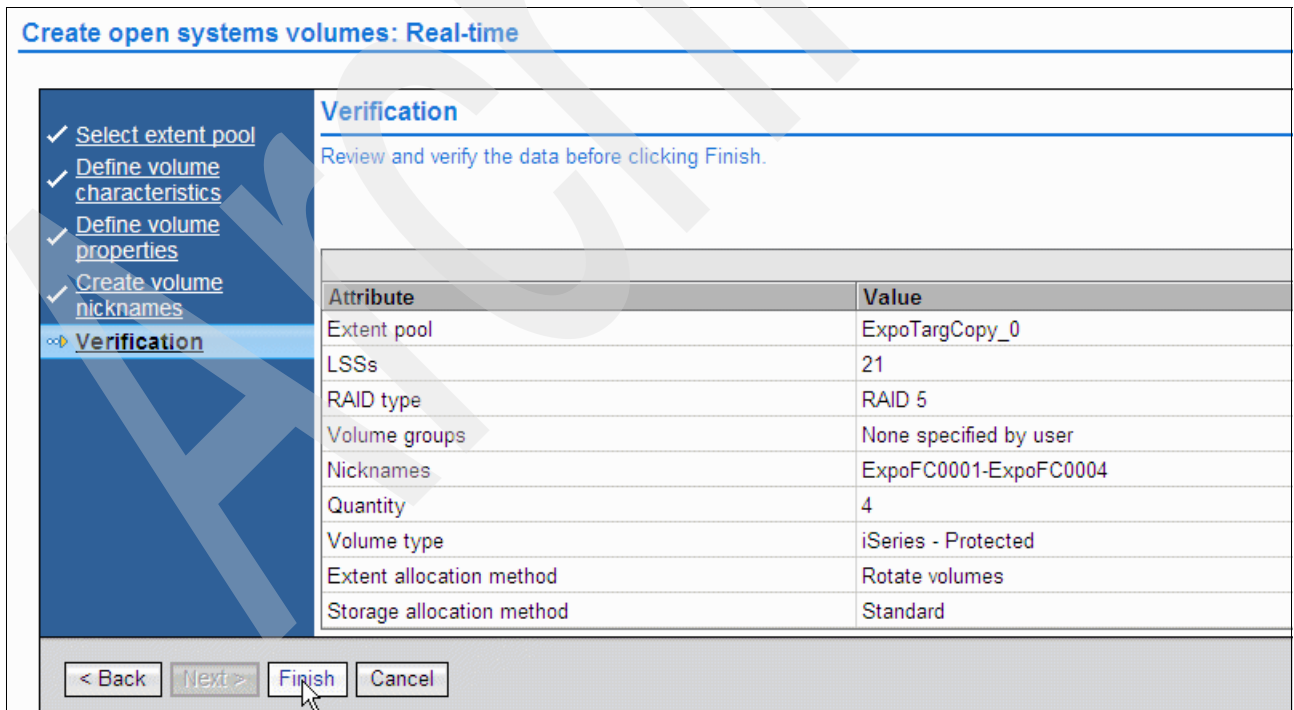


Figure 10-15 Verifying the creation of the open volume

- During the creation of the volumes, the Long Running Task Properties panel displays. You can close this panel by clicking **Close**. You can find all of the tasks detail by selecting **Real-time manager** → **Monitor system** → **Long running task summary**. You can also save the Long Running Task Properties to a file. See Figure 10-16.

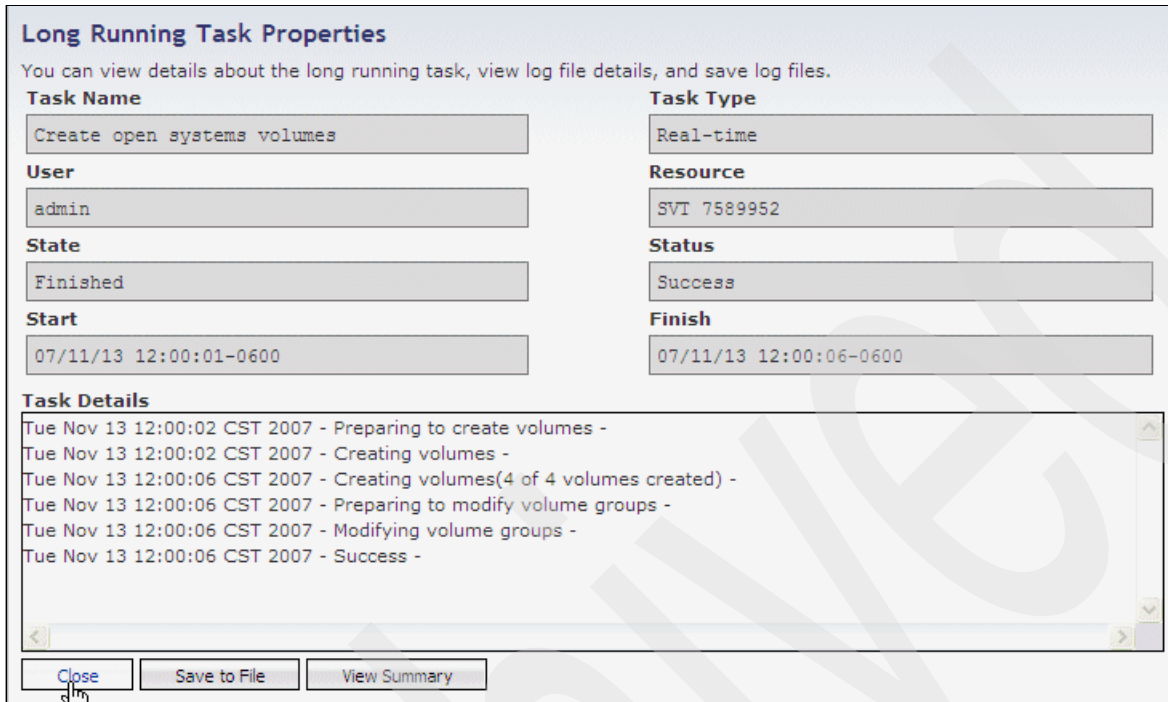


Figure 10-16 Long running task message for creating volumes

## 10.3 Creating a volume group

To create a volume group, follow these steps:

- Create a new volume group by selecting **Real-time manager** → **Configure storage** → **Open systems** → **Volume Groups**. Select **Create** from the Select action drop-down menu (see Figure 10-17).

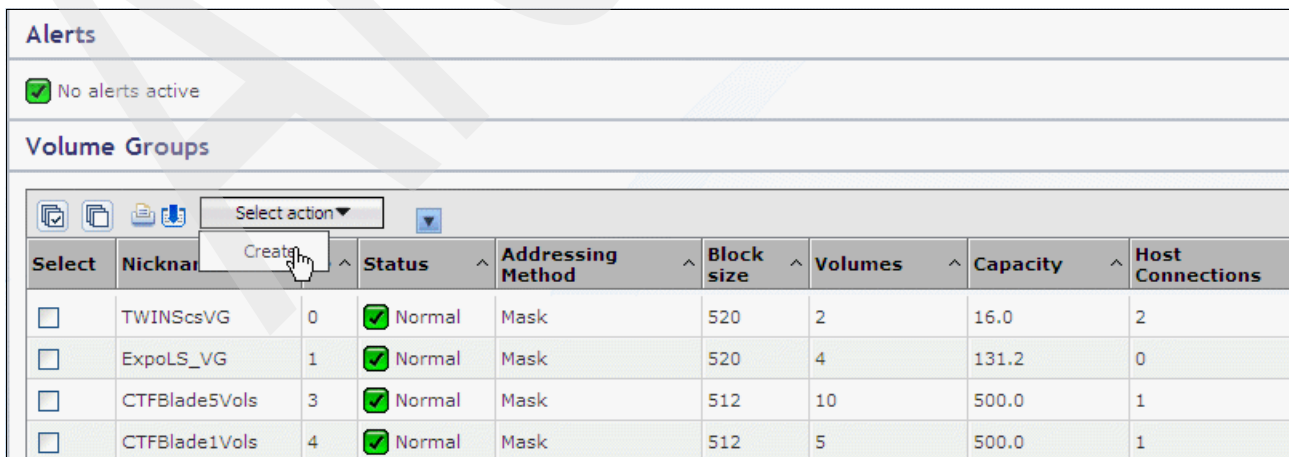


Figure 10-17 Working with volume groups



- The Create New Volume Group panel displays. Accept the default volume group nickname from **Volume Group Nickname** or enter a different nickname if desired. In our example, we use *ExpoFC\_VG* for our volume group nickname. Select **IBM iSeries and AS/400 Servers (OS/400)(iSeries)** for the **Host Type**, and select the volumes to be included in the group. In our example, we choose a filter for LSS 0x21 that we have defined before in step 4 on page 358). (see Figure 10-18).

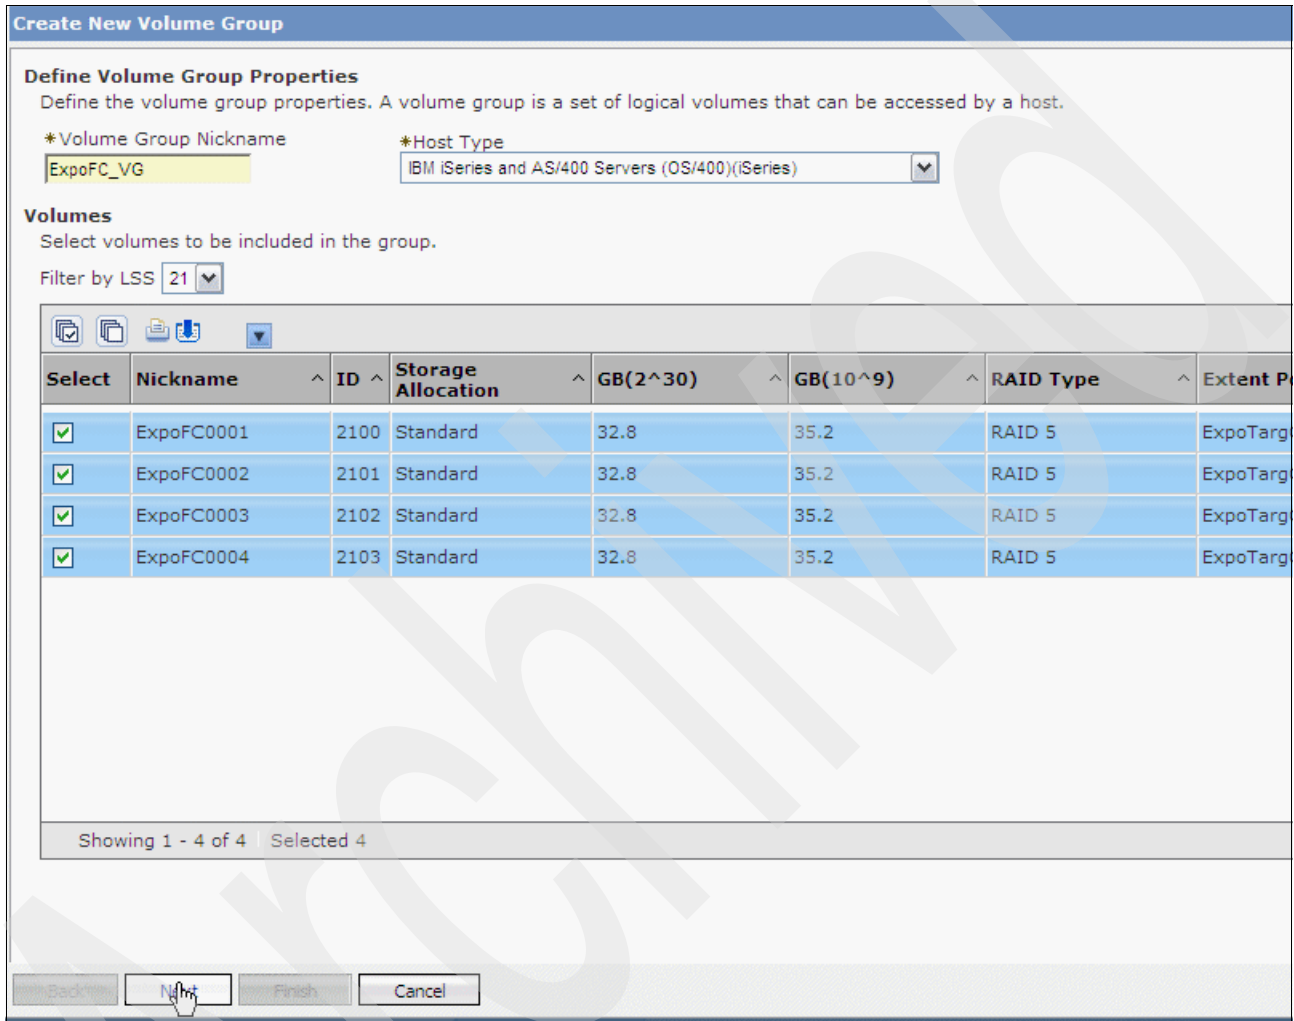


Figure 10-18 Define volume group properties

3. In the Verification panel, verify that the details are correct, and then click **Finish**. The volume group is now ready to be used for the FlashCopy (see Figure 10-19 and Figure 10-20).

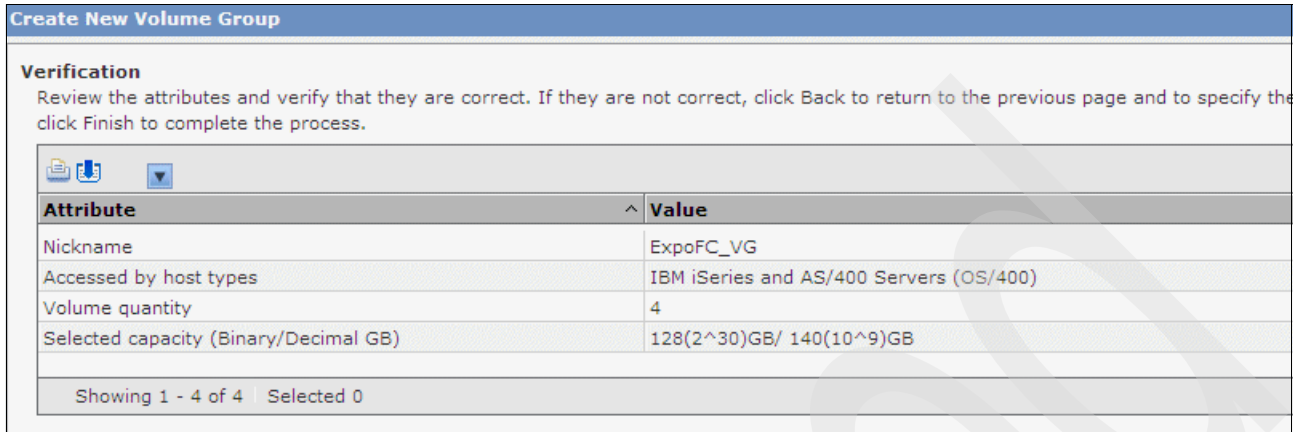


Figure 10-19 Volume group creation verification

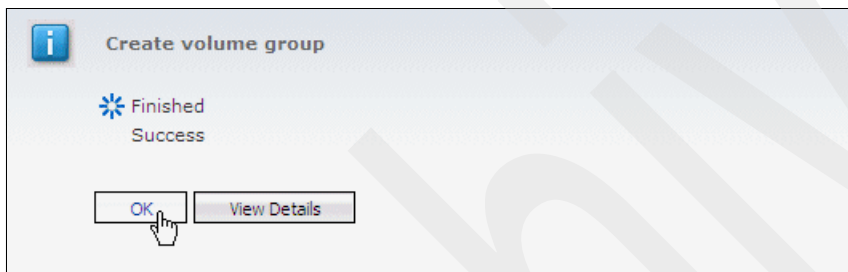


Figure 10-20 Volume group creation finished

## Implementing FlashCopy using the DS GUI

In this chapter, we explain how to implement FlashCopy using the GUI. You can also implement FlashCopy using the DS command-line interface (CLI) commands. For our environment, we copy from a single source to a single target, as shown in Figure 11-1.




Figure 11-1 FlashCopy environment

To implement FlashCopy using the GUI, follow these steps:

1. Access the DS GUI as described in 9.1.3, “Accessing the DS GUI interface” on page 338, and then sign on using an administrator user name and password, as shown in Figure 11-2.

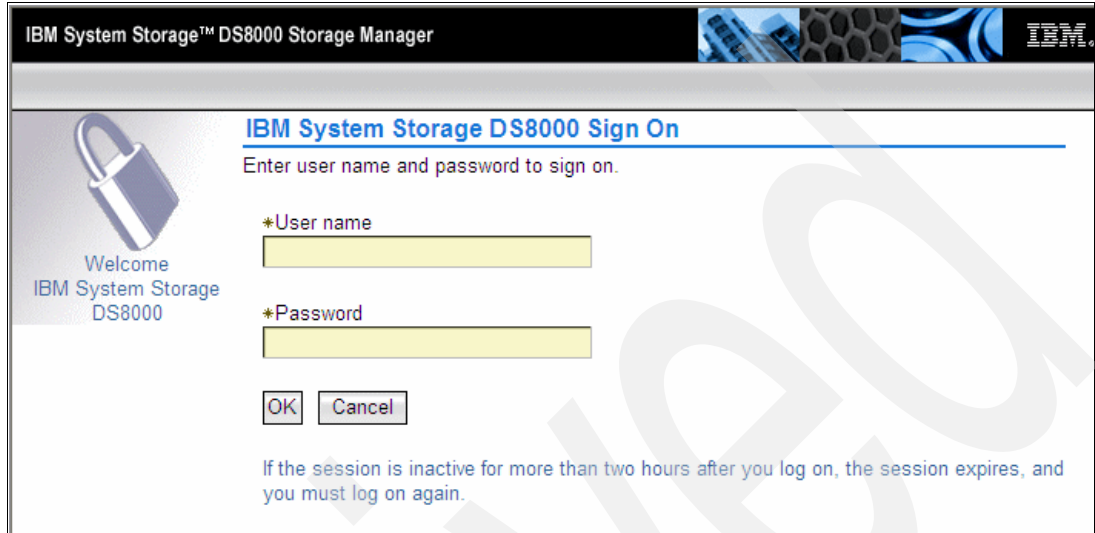


Figure 11-2 DS8000 Sign On panel

The main DS Storage Manager window (Figure 11-3) is the starting point for all your configuration, management, and monitoring needs for the DS disk and Copy Services tasks. The underlying hardware (two System p models), I/O drawers, and I/O adapters are controlled by the storage HMC built into the DS8000 rack or a standalone desktop SMC.

From any of the GUI panels, you can access the Information Center by clicking the question mark (?) in the upper-right corner of the page.

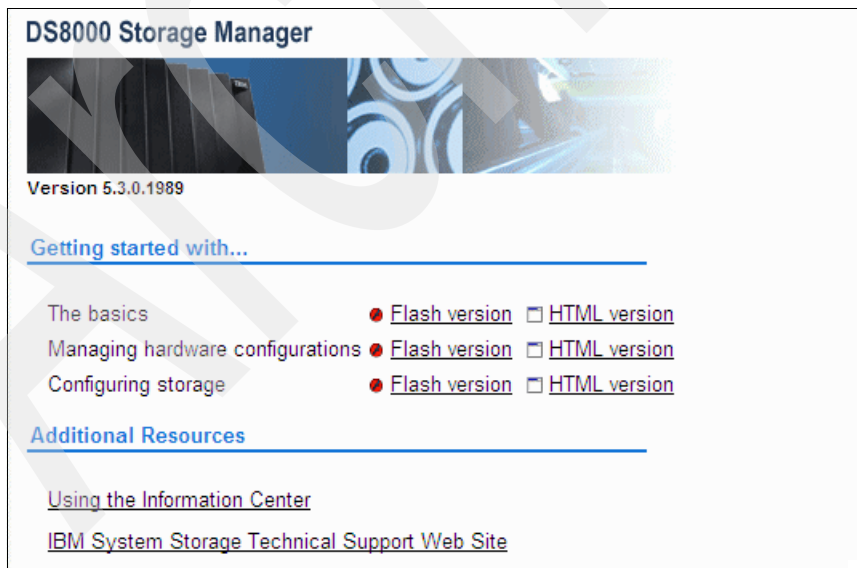


Figure 11-3 DS8000 Storage Manager - Getting started

2. Create the new FlashCopy implementation by selecting **Real-time manager** → **Copy services** → **FlashCopy**. Choose **Create** from the **Select action** drop-down box (see Figure 11-4).

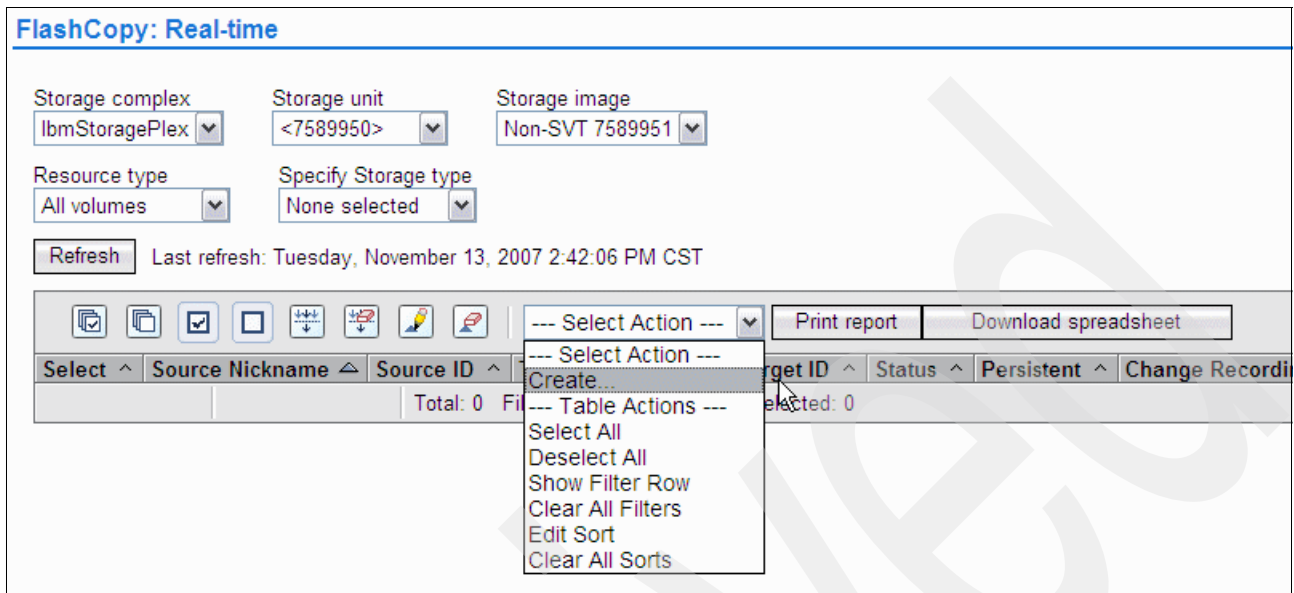


Figure 11-4 DS8000 Storage Manager FlashCopy window

3. In the right panel, select the type of relationship. In this example, we select **A single source with a single target**. Click **Next** to continue (see Figure 11-5).

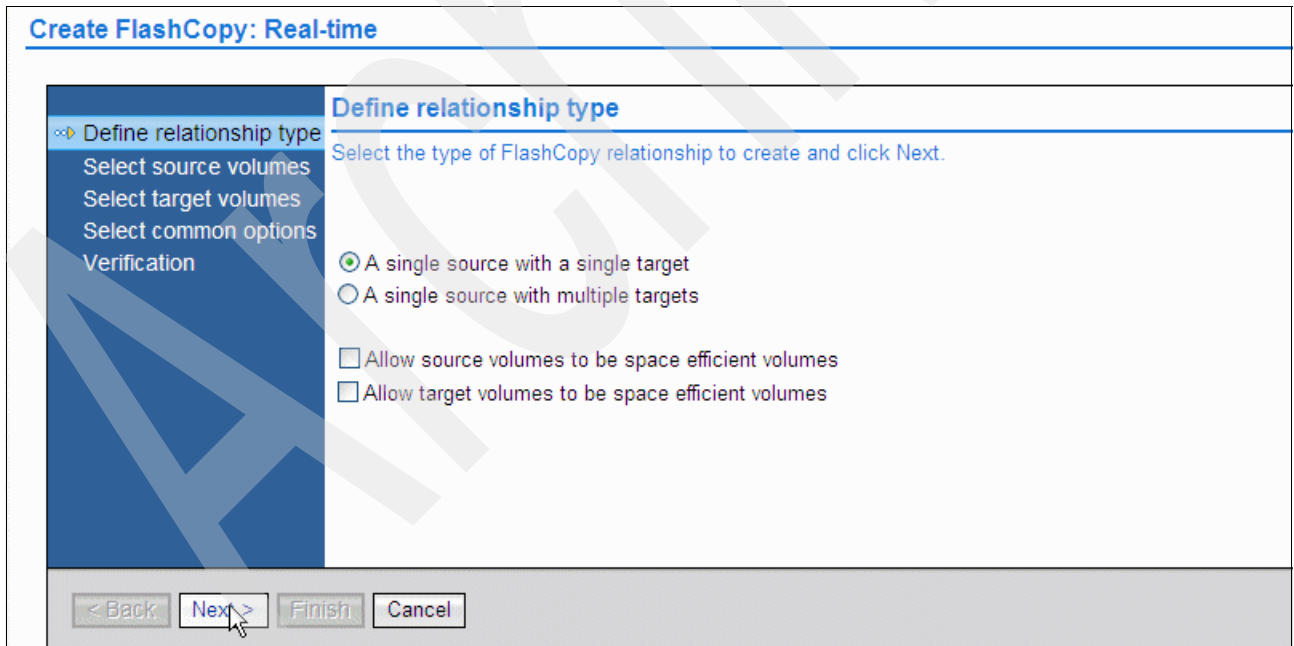


Figure 11-5 Defining the FlashCopy relationship

- In the Select the source volumes panel, specify the storage type for the FlashCopy. Because we are working with a System i environment, we are concerned only with fixed block (FB) volumes, so we select **All volumes** for **Resource type option** and **All FB volumes** for **Specify Storage type** option (see Figure 11-6).

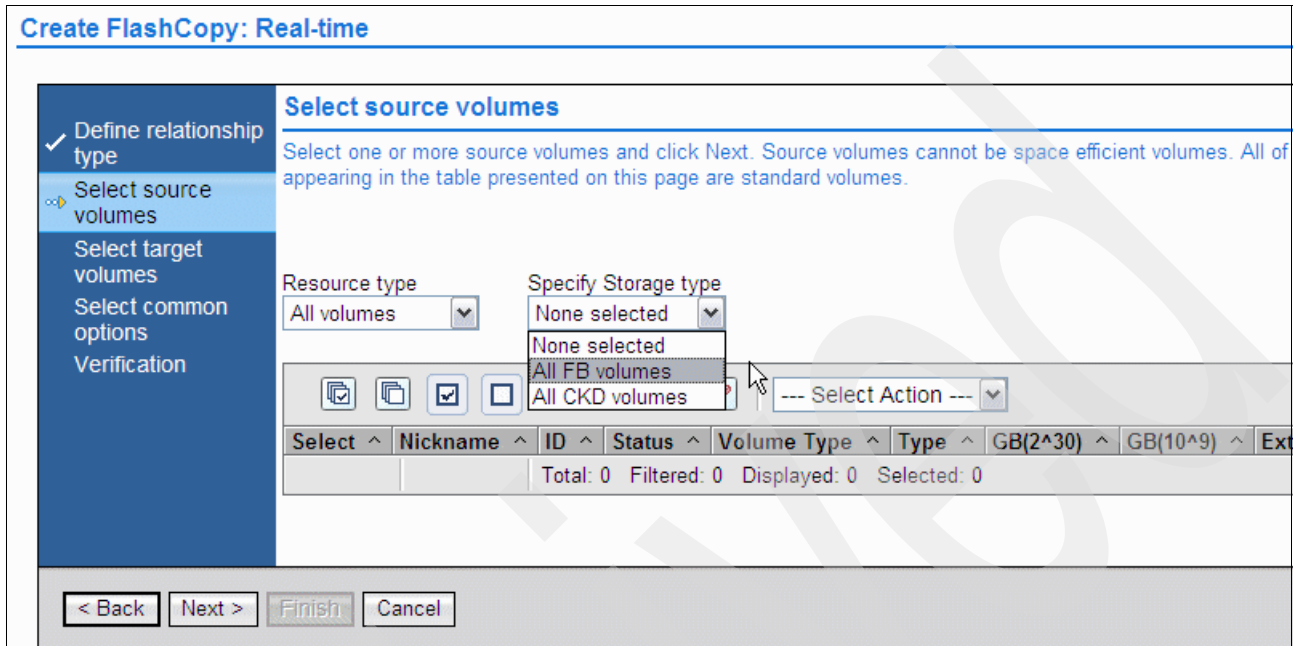


Figure 11-6 Specify Storage Type for source volumes

- In the next panel (Figure 11-7), select the volumes that are to be flashed. If the volumes that you want to select are on different pages, use the arrow key to go to the next page. Click **Next** to continue.

**Create FlashCopy: Real-time**

Define relationship type  
 **Select source volumes**  
 Select target volumes  
 Select common options  
 Verification

**Select source volumes**  
 Select one or more source volumes and click Next. Source volumes cannot be space efficient volumes in the table presented on this page are standard volumes.

Resource type: All volumes  
 Specify Storage type: All FB volumes

Select	Nickname	ID	Status	Volume Type	Type	GB(2^30)	GB
<input type="checkbox"/>	NDpool11	1056	Normal		I-protected	65.7	
<input type="checkbox"/>	NDpool12	1057	Normal		I-protected	65.7	
<input type="checkbox"/>	NDpool13	1058	Normal		I-protected	65.7	
<input type="checkbox"/>	NDpool14	1059	Normal		I-protected	65.7	
<input type="checkbox"/>	NDpool15	105A	Normal		I-protected	65.7	
<input type="checkbox"/>	jazMMtarget1	105B	Normal		I-protected	8.0	
<input type="checkbox"/>	jazMMtarget2	105C	Normal		I-protected	8.0	
<input type="checkbox"/>	jazMMtarget3	105D	Normal		I-protected	8.0	
<input checked="" type="checkbox"/>	TN1Is	105E	Normal		I-protected	32.8	
<input checked="" type="checkbox"/>	TN1Vol1	105F	Normal		I-protected	32.8	

Page 6 of 19 | 6 | Go | Total: 190 | Filtered: 190 | Displayed: 10 | Selected: 2

Figure 11-7 Selecting the FlashCopy source volumes

6. In the next panel (Figure 11-8), select the target volumes and click **Next** to continue.

**Note:** For System i environments always make sure the selected target volumes are the same System i volume model like the source volumes, i.e. they match in terms of volume capacity and protection mode.

**Create FlashCopy: Real-time**

- ✓ Define relationship type
- ✓ Select source volumes
- ⇨ Select target volumes
- Select common options
- Verification

### Select target volumes

Select a target volume for each selected source volume and click Next. Target volumes cannot be sp... volumes appearing in the table presented on this page are standard volumes.

Please select 4 targets

Resource type  
All volumes ▾

Specify Storage type  
All FB volumes ▾

☑ ☰ ☒ ☑ ☒ ☒ ☒ ☒ ☒ ☒
--- Select Action --- ▾

Select ^	Nickname ^	ID ^	Status ^	GB(2^30) ^	GB(10^9) ^	Extent Pool ^
<input type="checkbox"/>	RedsBS0048	1170	Normal	32.8	35.2	RedsBootSan
<input type="checkbox"/>	EndurNDmm	1171	Normal	65.7	70.6	ExpoBootSan
<input checked="" type="checkbox"/>	TN1FCIs	2100	Normal	32.8	35.2	ExpoBootSan
<input checked="" type="checkbox"/>	TN1FCVol1	2101	Normal	32.8	35.2	ExpoBootSan
<input checked="" type="checkbox"/>	TN1FCVol2	2102	Normal	32.8	35.2	ExpoBootSan

Figure 11-8 Selecting the target volumes

7. In the Select common options panel, select the parameters that you require (as shown in Figure 11-9). If you leave the default Initiate background copy option selected as in our example, a full copy of the data is forced from the source to the target.

When using FlashCopy to create a system or IASP image for backup to tape purposes, you typically should not use the background copy option to copy changed tracks only and thus limit the performance impact to the production system. In this case, clear the “Initiate background copy” option. If you are using DS CLI, use the `mkflash` command `-nocp` option. Click **Next** to continue.



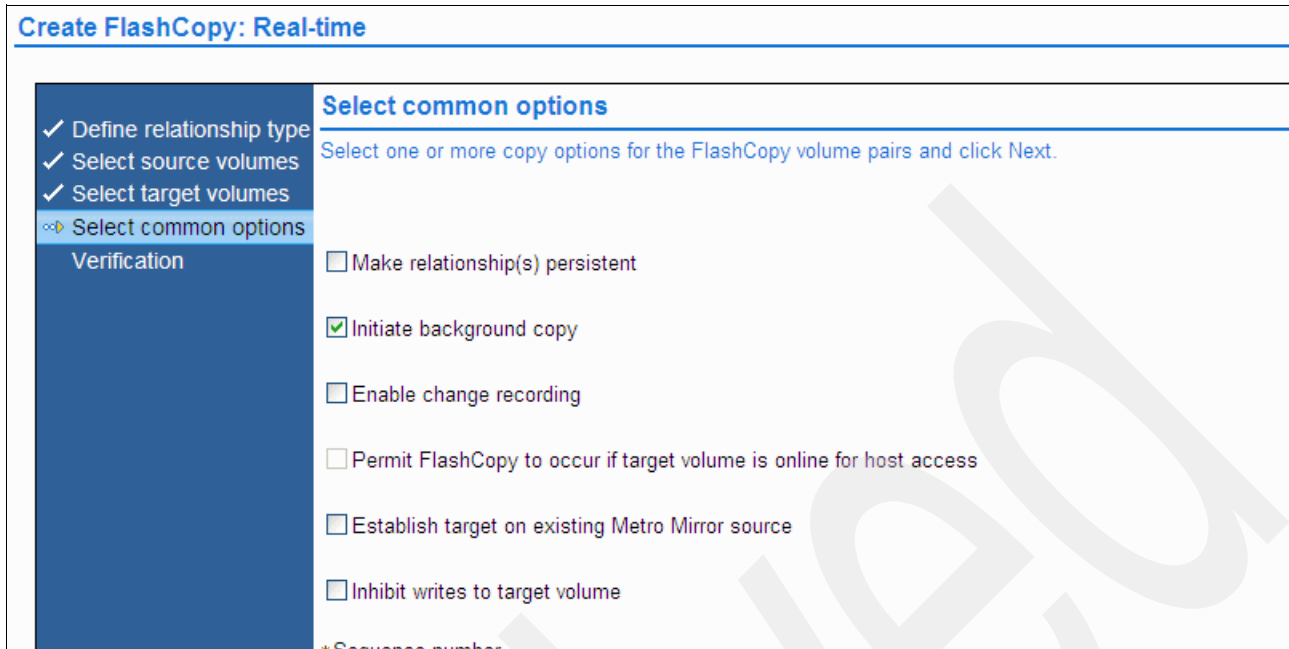


Figure 11-9 FlashCopy startup options

- In the Verification panel, verify that the source and target LUNs are as required, as shown in Figure 11-10. Click **Finish** to continue the FlashCopy implementation.

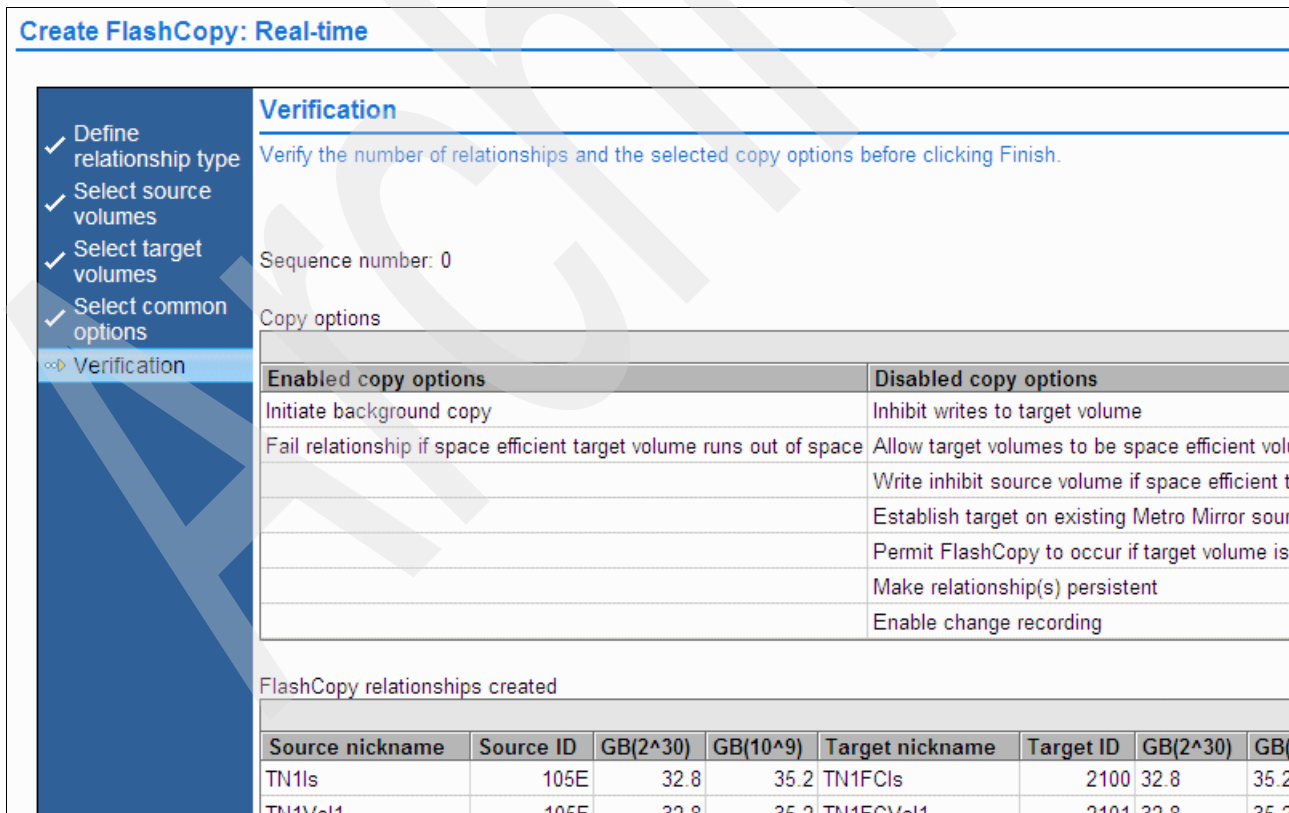


Figure 11-10 Verifying the FlashCopy options

- View the relationships between the source and target. Figure 11-11 shows that more than zero tracks are out of sync and that the copy process is still running in the background. Click one of the Source Nickname to see the FlashCopy properties.

**Note:** Independent from the FlashCopy completion state you can start using the FlashCopy target volumes without restriction for both host read and write access from another System i server or LPAR as soon as the FlashCopy relationship has been established, that is corresponding DS8000 internal track bitmaps are created.

**FlashCopy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Tuesday, November 13, 2007 3:04:40 PM CST

--- Select Action ---

Select ^	Source Nickname ^	Source ID ^	Target Nickname ^	Target ID ^	Status ^	Persistent ^	Ch
<input type="checkbox"/>	<a href="#">TN1Vol1</a>	105F	TN1FCVol1	2101	Out-of-sync tracks > 0	No	Dis
<input type="checkbox"/>	<a href="#">TN1Vol2</a>	1060	TN1FCVol2	2102	Out-of-sync tracks > 0	No	Dis
<input type="checkbox"/>	<a href="#">TN1Vol3</a>	1061	TN1FCVol3	2103	Background copy running	No	Dis
<input type="checkbox"/>	<a href="#">TN1Is</a>	105E	TN1FCIs	2100	Out-of-sync tracks > 0	No	Dis

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 11-11 FlashCopy status

10. The next panel lists the general properties of the FlashCopy (Figure 11-12). Verify the attributes, and click **Out of sync tracks** to see another properties of the FlashCopy.

**FlashCopy Properties: Real-time**

**General**
[Out-of-sync tracks](#)

**General**

Review the attributes of the current relationship.

Attribute	Value
Created	Nov 13, 2007 3:04:19 PM
Last Refresh	Nov 13, 2007 3:04:19 PM
Relationship will remain	No
Change recording	No
Revertible	No
Background copy initiated	No
Source is write inhibited	No
Target is write inhibited	No
Source ID	1060
Source Nickname	<a href="#">TN1Vol2</a>
Source Allocation Method	Standard
Target ID	2102

Figure 11-12 FlashCopy general properties

11. At any time you can look to see how many tracks are out of sync (see Figure 11-13). This is not an error condition but rather an indication of the number of tracks that have not been copied since the FlashCopy was initiated. When the FlashCopy has completed, the **Status** panel is changed to **Copy complete**. Click **Close** to exit the properties panel.

**Out-of-sync tracks**

Set a refresh interval to display the current out-of-sync tracks.

Source Nickname	Source ID	Target Nickname	Target ID
TN1Vol2	1060	TN1FCVol2	2102
Total: 1			

Refresh Interval  
 ▼

Out-of-sync tracks: 421221

Figure 11-13 FlashCopy Out-of-sync tracks properties

Archived



## Implementing Metro Mirror using the DS GUI

In this chapter, we describe the steps to configure the Metro Mirror function using the DS GUI, where the host is a System i server. We also reference the DS command-line interface (CLI) options when appropriate.

## 12.1 Metro Mirror arrangement

To implement Metro Mirror, you need a second storage system and a second System i partition to attach to that storage system in the event of a switchover.

In our example for sake of simplicity we used a DS8300 LPAR machine to set up Metro Mirror within one physical machine between both storage images 75-89951 and 75-89952 shown as separate machines 51 and 52 in Figure 12-1. For a real production environment disaster recovery solution, you need to set up Metro Mirror between different physical machines at different locations.




Figure 12-1 Metro Mirror arrangement with a System i5 environment

## 12.2 Implementing Metro Mirror volume relationships

**Important:** Before you can create Metro Mirror volume pairs, you must create PPRC paths between a source LSS in a specified storage unit and a target LSS in a specified storage unit. Either use the **DS Storage Manager GUI Realtime Manager** → **Copy Services** → **Paths** function or the DS CLI `mkpprcpath` command (see 7.2.1, “Creating Peer-to-Peer Remote Copy paths” on page 282).

To configure Metro Mirror volume relationships:

1. Access the DS GUI as described in 9.1.3, “Accessing the DS GUI interface” on page 338 and sign on using an administrator user name and password (see Figure 12-2).



Figure 12-2 DS8000 Sign On panel

2. On the main page of DS Storage Manager, in the left navigation panel, select **Real-time manager** → **Copy Services** → **Metro Mirror**, as shown in Figure 12-3.

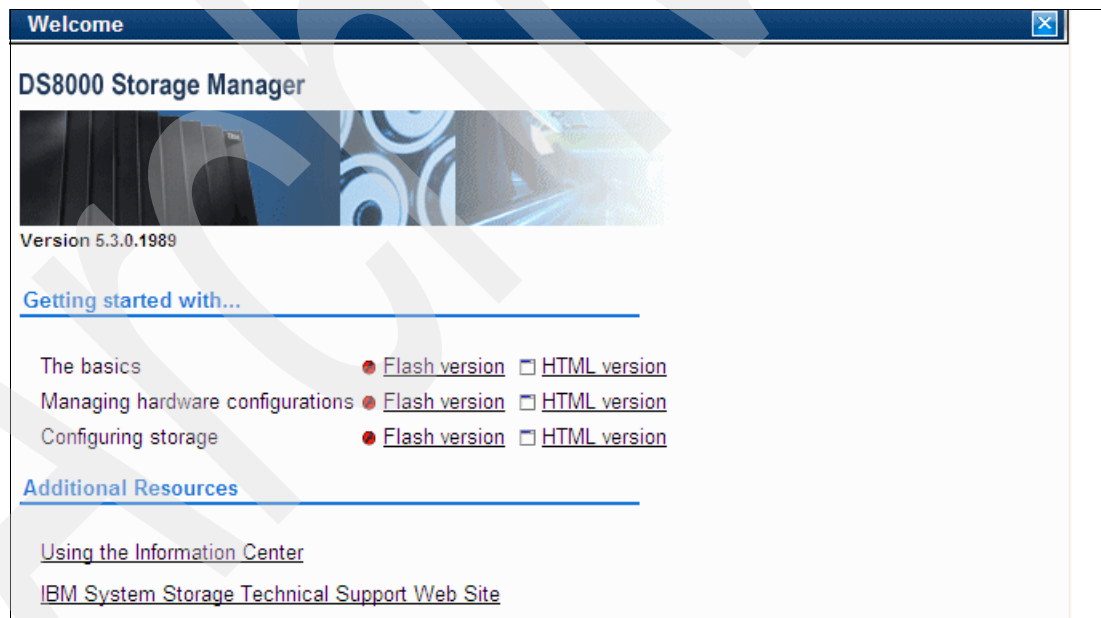


Figure 12-3 DS8000 Storage Manager: Getting started

**Note:** At any time you can access the online help for a description of the available functions, by clicking the question mark (?) in the upper-right corner.

3. In the next panel, connect to the storage image of the source DS system. Select **Create** from the Select Action drop-down menu.

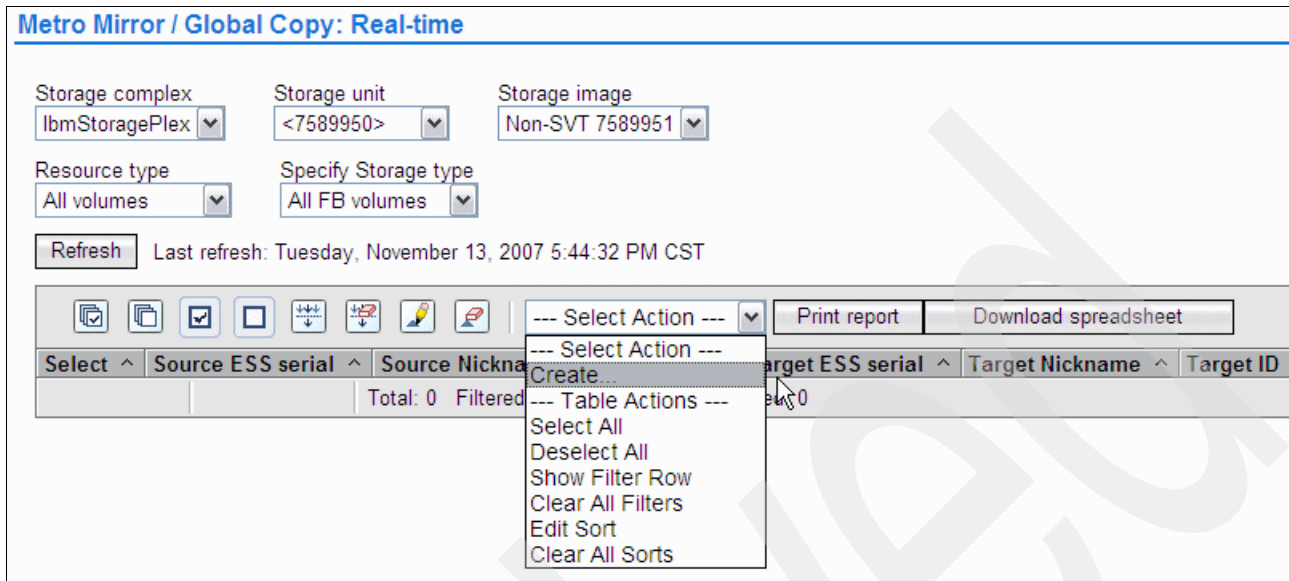


Figure 12-4 Creating a Metro Mirror relationship



4. In the Volume Pairing Method panel, you can choose to have the individual pairs linked automatically by the system (see Figure 12-5). If you select the “Automated volume pair assignment” option, the system pairs the first volume on the source with the first volume on the target, then the second, third, and so on until all volumes are paired. If your naming convention does not allow this, you must select **Manual volume pair assignment**. Click **Next** to continue.

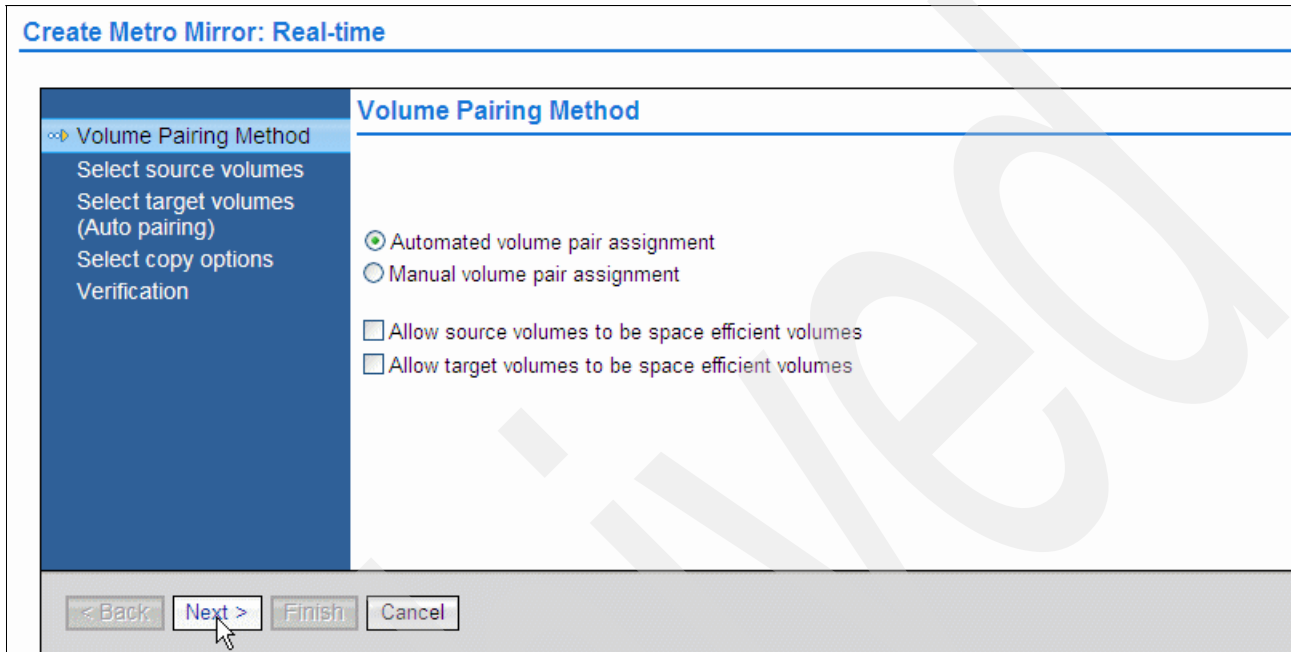


Figure 12-5 Volume Pairing Method for Metro Mirror

- In the Select source volumes panel, specify the source volumes that you want to include in this Metro Mirror implementation, as shown in Figure 12-6. If the volumes that you want to select are on the different pages, use the arrow key to go to the next page. Click **Next** to continue.

**Select source volumes**

Select one or more source volumes for the Metro Mirror relationships and click Next. Optionally, click the Create Paths button to create paths for the selected volumes. Source volumes cannot be space efficient volumes. All of the volumes appearing in the table presented on this page are standard volumes.

Resource type: LSS Specify LSS: 10

Select	Nickname	ID	Status	Volume Type	Type	GB(2^30)	GB(10^9)	Extent Pool
<input type="checkbox"/>	NDpool11	1056	Normal		I-protected	65.7	70.6	Non SVT 0
<input type="checkbox"/>	NDpool12	1057	Normal		I-protected	65.7	70.6	Non SVT 1
<input type="checkbox"/>	NDpool13	1058	Normal		I-protected	65.7	70.6	Non SVT 0
<input type="checkbox"/>	NDpool14	1059	Normal		I-protected	65.7	70.6	Non SVT 1
<input type="checkbox"/>	NDpool15	105A	Normal		I-protected	65.7	70.6	Non SVT 0
<input type="checkbox"/>	jazMMtarget1	105B	Normal		I-protected	8.0	8.6	Non SVT 0
<input type="checkbox"/>	jazMMtarget2	105C	Normal		I-protected	8.0	8.6	Non SVT 1
<input type="checkbox"/>	jazMMtarget3	105D	Normal		I-protected	8.0	8.6	ExpoFLASH
<input checked="" type="checkbox"/>	TN1Is	105E	Normal		I-protected	32.8	35.2	Non SVT 0
<input checked="" type="checkbox"/>	TN1Vol1	105F	Normal		I-protected	32.8	35.2	Non SVT 0

Page 6 of 9 | 6 | Go | Total: 83 | Filtered: 83 | Displayed: 10 | Selected: 4

Figure 12-6 Select source volumes for Metro Mirror

- For auto pairing, in the Select target volumes (Auto pairing) panel, select the target volumes (from another image of the DS8000 in this example), and let the system match them. See Figure 12-7. For additional volumes, use the arrow to go to the next panel or enter the number of the page that you require, and click **Go**.

**Select target volumes (Auto pairing)**

Select the same number of target volumes as selected source volumes and click Next. Target volumes cannot be space efficient volumes. All of the volumes appearing in the table presented on this page are standard volumes.

Please select 4 targets

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify LSS:

Select	Nickname	ID	Status	Type	GB(2^30)	GB(10^9)	Extent Pool	Host Connections
<input type="checkbox"/>	Miranda30701	1014	Normal	I-protected	65.7	70.6	Miranda	1
<input type="checkbox"/>	Miranda31401	1015	Normal	I-protected	131.4	141.1	Miranda	1
<input type="checkbox"/>	Miranda32801	1016	Normal	I-protected	262.9	282.3	Miranda	1
<input type="checkbox"/>	Miranda40701	1017	Normal	I-protected	65.7	70.6	Miranda	1
<input type="checkbox"/>	Miranda41401	1018	Normal	I-protected	131.4	141.1	Miranda	1
<input type="checkbox"/>	Miranda42801	1019	Normal	I-protected	262.9	282.3	Miranda	1
<input checked="" type="checkbox"/>	ExpoLS0001	101A	Normal	I-protected	32.8	35.2	ExpoSourceCopy	0
<input checked="" type="checkbox"/>	ExpoLS0002	101B	Normal	I-protected	32.8	35.2	ExpoSourceCopy	0
<input checked="" type="checkbox"/>	ExpoLS0003	101C	Normal	I-protected	32.8	35.2	ExpoSourceCopy	0

Figure 12-7 Select target volumes (Auto pairing) for Metro Mirror

7. The Select copy options panel offers additional options that you can select, as shown in Figure 12-8. Not all of the options are valid for the System i5 platform. For a detailed explanation of each option, see the online help text using the question mark (?). In our example, we select the **Perform initial copy** option. This option guarantees that the source and target volume contain the same data. When Metro Mirror relationship is created with this option, the entire source volume is copied to target volume. Click **Next** to continue.

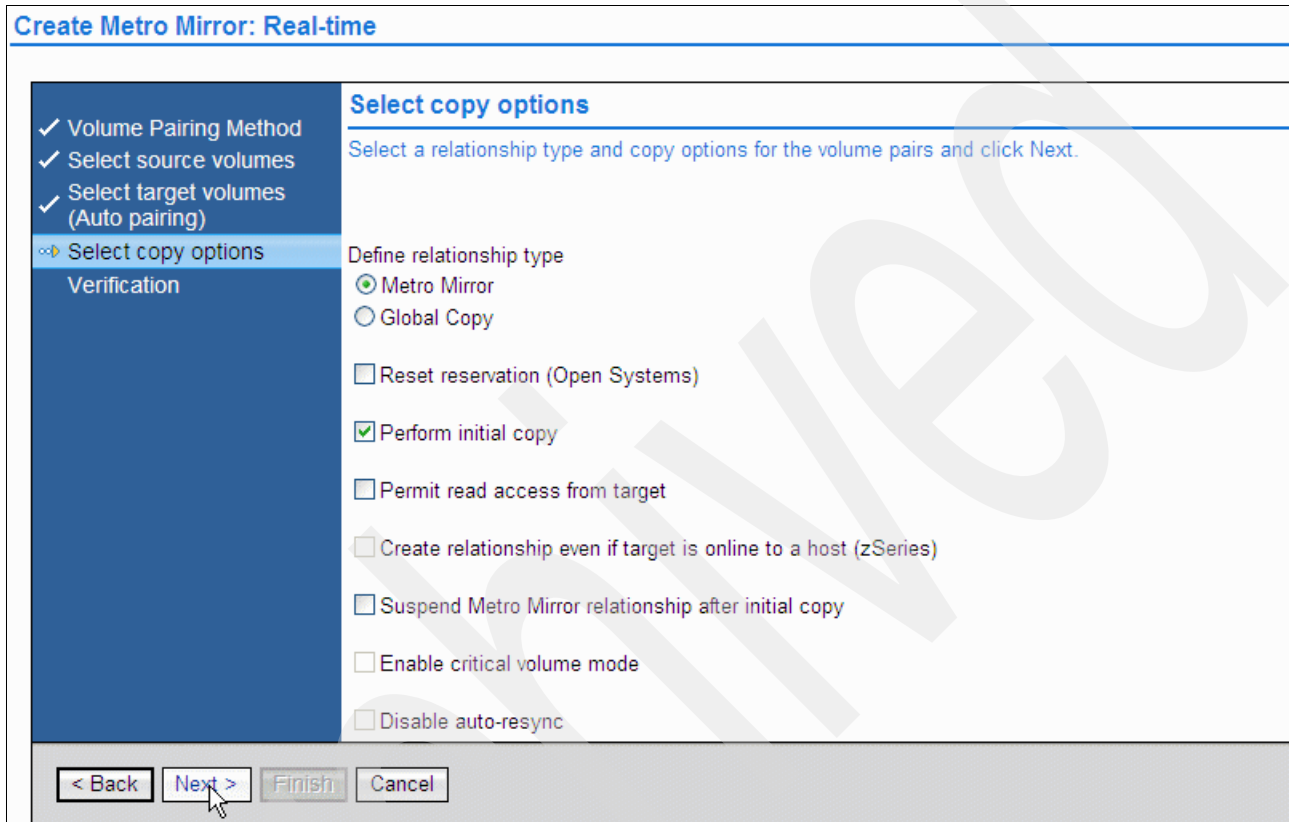


Figure 12-8 Select copy options for Metro Mirror

8. On the Verification panel, verify that the setup is correct, as shown in Figure 12-9.

**Create Metro Mirror: Real-time**

**Verification**

Verification

Source nickname	Source ID	Source Storage Complex	Source Binary GB	Source Decimal GB	Target nickname
TN1Is	105E	lbmStoragePlex	32.8	35.2	ExpoLS0001
TN1Vol1	105F	lbmStoragePlex	32.8	35.2	ExpoLS0002
TN1Vol2	1060	lbmStoragePlex	32.8	35.2	ExpoLS0003
TN1Vol3	1061	lbmStoragePlex	32.8	35.2	ExpoLS0004

Page 1 of 1      Total: 4    Displayed: 4

< Back    Next >    **Finish**    Cancel

Figure 12-9 Verifying the Metro Mirror relationship

Scroll to the right and verify the details there as well (Figure 12-10). If everything is correct, scroll back to the left, and click **Finish** to continue.

**Create Metro Mirror**

Complex	Source Binary GB	Source Decimal GB	Target nickname	Target ID	Target Storage Complex	Target Binary GB	Target Decimal GB
	32.8	35.2	ExpoLS0001	101A	lbmStoragePlex	32.8	35.2
	32.8	35.2	ExpoLS0002	101B	lbmStoragePlex	32.8	35.2
	32.8	35.2	ExpoLS0003	101C	lbmStoragePlex	32.8	35.2
	32.8	35.2	ExpoLS0004	101D	lbmStoragePlex	32.8	35.2

Figure 12-10 Additional information for verifying the Metro Mirror relationship

## 12.3 Displaying Metro Mirror volume properties

To display Metro Mirror volume properties, follow these steps:

1. The initial state of a newly created Metro Mirror relationship is *Copy pending*, as shown in Figure 12-11. Select one of the Metro Mirror relationships from the **Realtime Manager** → **Copy services** → **Metro Mirror / Global Copy** view, and select **Properties** from the Select Action drop-down menu to see the detailed properties.

**Metro Mirror / Global Copy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Wednesday, November 14, 2007 4:54:04 PM CST

---

--- Select Action ---
Print report
Download spreadsheet

Select ^	Source ESS serial ^	Source Nickname ^	Source ID ^	Target ESS serial ^	Target Nickname ^	Target ID
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:101B
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:101C
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:101D
<input type="checkbox"/>	7589951	TN1Is	105E	7589952	Not retrieved	0002:101A

Page 1 of 1      Total: 4   Filtered: 4   Displayed: 4   Selected: 0

Figure 12-11 Copy status of Metro Mirror relationship

- The Metro Mirror Properties panel, shown in Figure 12-12, displays similar status information (Copy pending) such as that shown in the previous panel (Figure 12-11). Click **Out-of-sync tracks** from the Metro Mirror Properties navigation panel to see another properties of the Metro Mirror relationship.

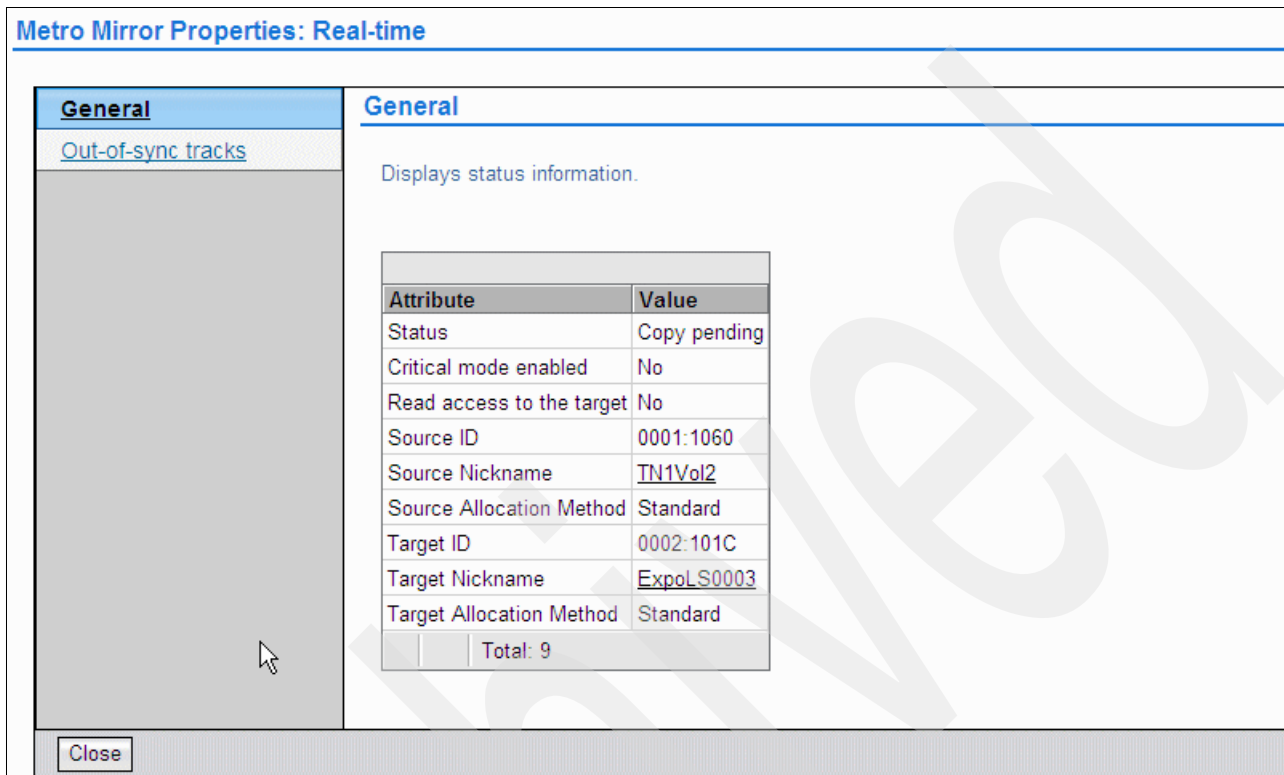


Figure 12-12 Metro Mirror relationship general properties

3. From the Out-of-sync tracks properties panel, verify the number of tracks that are not synchronized (see Figure 12-13).

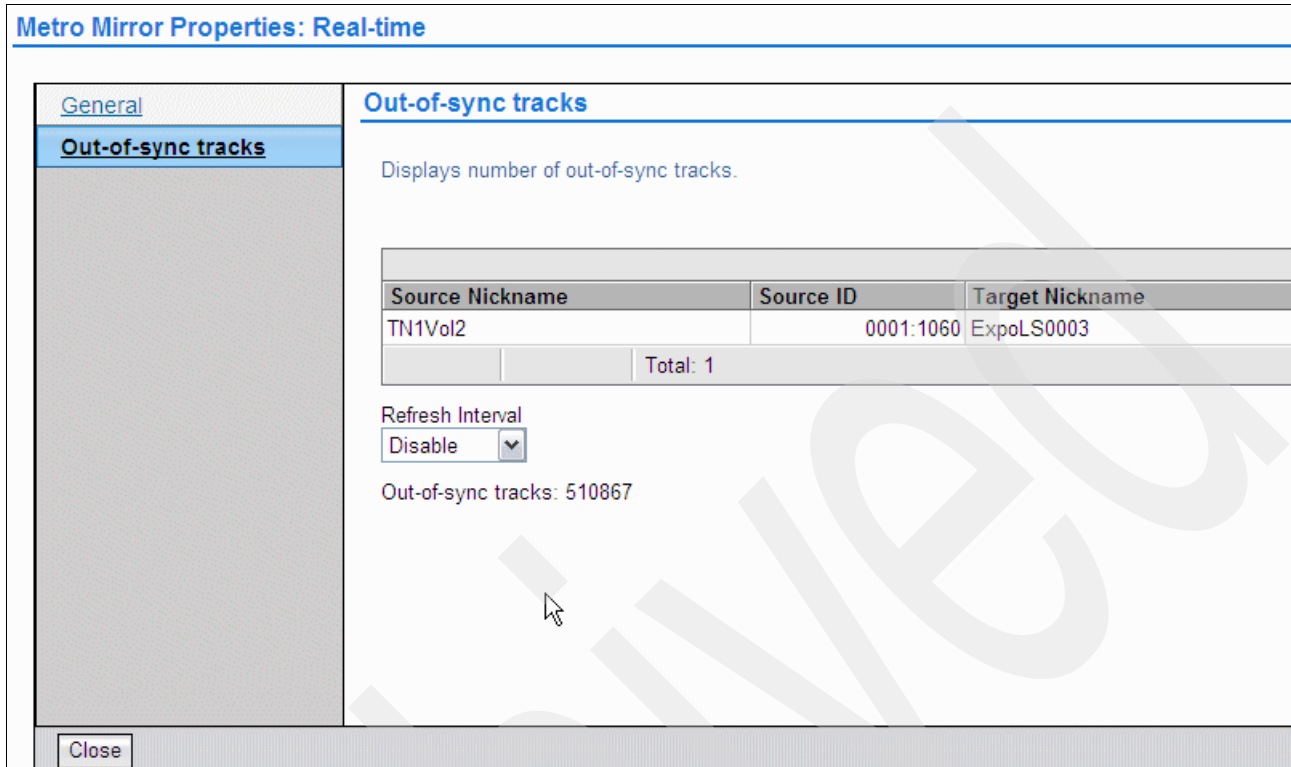


Figure 12-13 Out-of-sync tracks



Depending on the number and size of the volumes involved, it takes some time for the number of out-of-sync tracks to reach zero. Select one of the **Refresh Interval** options to refresh the Out-of-sync tracks automatically information (see Figure 12-14).

**Metro Mirror Properties: Real-time**

General

**Out-of-sync tracks**

Displays number of out-of-sync tracks.

Source Nickname	Source ID	Target Nickname
TN1Vol2	0001:1060	ExpoLS0003
Total: 1		

Refresh Interval  
1 Minute

Out-of-sync tracks: 505946

Close

Figure 12-14 Reduced out-of-sync trackers

As each volume completely synchronizes, it changes to a state of *Full duplex* (Figure 12-15).

**Metro Mirror / Global Copy: Real-time**

Storage complex: ibmStoragePlex  
Storage unit: <7589950>  
Storage image: Non-SVT 7589951

Resource type: All volumes  
Specify Storage type: All FB volumes

Refresh Last refresh: Wednesday, November 14, 2007 5:02:50 PM CST

--- Select Action --- Print report Download spreadsheet

Select	Source ESS serial	Source Nickname	Source ID	Target ESS serial	Target Nickname	Target ID
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:101B
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:101C
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:101D
<input type="checkbox"/>	7589951	TN1Is	105E	7589952	Not retrieved	0002:101A

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 12-15 Volumes state changed to Full duplex

After all volumes are full duplex, you know that the target is a true copy of the source (see Figure 12-16).

**Metro Mirror / Global Copy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Wednesday, November 14, 2007 5:19:12 PM CST

Select ^	Source ESS serial ^	Source Nickname ^	Source ID ^	Target ESS serial ^	Target Nickname ^	Target I
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1ls	105E	7589952	Not retrieved	0002:101

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 12-16 Metro Mirror relationship in full duplex

Looking at Metro Mirror relationship from the target system (Storage image 7589952), you can also see the same volume state information (see Figure 12-17).

**Metro Mirror / Global Copy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Wednesday, November 14, 2007 5:19:12 PM CST

Select ^	Source ESS serial ^	Source Nickname ^	Source ID ^	Target ESS serial ^	Target Nickname ^	Target I
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1ls	105E	7589952	Not retrieved	0002:101

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 12-17 Metro Mirror Relationship from the target system view

After the full duplex state is achieved, the Metro Mirror relationship is maintained until another action is undertaken. This can be a failover through a disaster or a planned outage of the source system.



## Managing Copy Services in i5/OS environments using the DS GUI

In this chapter, we discuss the tasks that are necessary to manage FlashCopy, Metro Mirror, and Global Mirror using the DS6000 and DS8000 GUI. You can manage the Copy Services functions from a DS command-line interface (CLI) or the GUI. In this chapter, we discuss the options that are available with the GUI for the three Copy Services functions.

## 13.1 FlashCopy options

FlashCopy has several options that you can use when you setup the FlashCopy relationship. The options that you select dictate what can be done after the flash is established. The following options are available:

- ▶ Make relationships persistent
- ▶ Initiate background copy
- ▶ Enable change recording
- ▶ Permit FlashCopy to occur if target volume is online for host access
- ▶ Establish target on existing Metro Mirror source
- ▶ Inhibit writes to target volume
- ▶ Fail relationship if space-efficient target volume becomes out of space
- ▶ Write inhibit the source volume if space-efficient target volume becomes out of space
- ▶ Sequence number for these relationships

You can also manage FlashCopy in an i5/OS environment using the iSeries Copy Services Toolkit.

In the following sections, we explain each of these options in detail.

**Note:** You can only attach FlashCopy LUNs to a System i i5/OS system or partition if they represent a full system or an IASP image. Then you can use this system image or IASP database to:

- ▶ Perform a backup
- ▶ Run reports
- ▶ Serve as a test environment
- ▶ Test an application update
- ▶ Test an operating system upgrade

### 13.1.1 Make relationships persistent

The Make relationships persistent option dictates whether the relationship continues after the copy is complete. If this option is not selected, the relationship ends after the copy is complete, that is after all tracks are copied from the source to the target volume. A persistent relationship remains even after the copy is complete. You can use this option for incremental or revertible FlashCopy.

### 13.1.2 Initiate background copy

With the Initiate background copy option, all data from the source volume is copied physically to the target volume or volumes. When the copy process is complete, the FlashCopy relationship ends unless the relationship is persistent. This option is the only option that is selected by default.

Clearing this option copies a track from the source to the target only if a track on the source is modified that is not copied yet or a background copy is initiated later.

### 13.1.3 Enable change recording

Selecting the Enable change recording option makes the relationship persistent automatically. It also monitors the writes and records changes on the volume pair in the FlashCopy

relationship. This option is required for *incremental FlashCopy*, that is if you plan to refresh the copy at a later date.

#### 13.1.4 Permit FlashCopy to occur if target volume is online for host access

This option is not available for the System i5 platform. It is used on the IBM eServer zSeries® platform.

#### 13.1.5 Establish target on existing Metro Mirror source

This option creates a point-in-time copy of a volume. With Metro Mirror a copy of that point-in-time copy is propagated to a remote site. This option creates a local point-in-time backup *and* a remote point-in-time backup.

If you do not select this option and the FlashCopy target volume is a Metro Mirror source volume, the create FlashCopy relationship task fails. This option defaults to not selected and displays on the Verification page as *disabled*.

#### 13.1.6 Inhibit writes to target volume

The Inhibit writes to target volume option prevents write operations from the host system (source volume) to the target volume while the FlashCopy relationship exists. It is used in context of Global Mirror to prevent host access to the FlashCopy target volumes containing the consistency group saves.

#### 13.1.7 Fail relationship if space-efficient target volume becomes out of space

This option if the target volume is a space-efficient volume. If the space-efficient target volume is full, this option will fail the FlashCopy relationship without impacting the production system.

#### 13.1.8 Write inhibit the source volume if space-efficient target volume becomes out of space

This option for space-efficient FlashCopy is not available if the target volume is not a space-efficient volume. It prevents write operations to the source volume if the space-efficient target volume is full.

**Important:** Using this option is not supported for i5/OS.

#### 13.1.9 Sequence number for these relationships

The Sequence number for these relationships option defines a number that can be used to group FlashCopy relationships. The sequence number is a maximum of eight hexadecimal digits in length. When defined during FlashCopy establish or resync, it can be used within subsequent commands to refer to multiple FlashCopy relationships.

If the FlashCopy sequence number that is specified does *not match* the sequence number of a current relationship or if a sequence number is not specified, the selected operation is

performed. If the FlashCopy sequence number that is specified *matches* the sequence number of a current relationship, the operation is not performed. The default value is zero.

## 13.2 FlashCopy GUI

In this section, we explain the options that are available to manage FlashCopy.

### 13.2.1 Delete

To delete the FlashCopy relationship, select **Real-time manager** → **Copy services** → **FlashCopy**. Select the relationship that you want to delete, and click **Delete** from the Select Action drop-down menu as shown in Figure 13-1.

The screenshot shows the 'FlashCopy: Real-time' interface. At the top, there are filters for 'Storage complex' (IbmStoragePlex), 'Storage unit' (<7589950>), and 'Storage image' (Non-SVT 7589951). Below these are 'Resource type' (LSS) and 'Specify LSS' (21). A 'Refresh' button and a timestamp 'Last refresh: Tuesday, November 13, 2007 3:25:45 PM CST' are also present.

The main area contains a table with columns: 'Select', 'Source Nickname', and 'Source ID'. The table lists four entries: TN1Vol1 (105F), TN1Vol2 (1060), TN1Vol3 (1061), and TN1ls (105E). The 'TN1Vol2' row is selected. A context menu is open over the table, showing options like 'Create...', 'Delete', 'Initiate Background Copy', 'Properties', and 'Reset Target Write Inhibit'. The 'Delete' option is highlighted.

At the bottom of the table, it says 'Page 1 of 1' and 'Total: 4'. To the right of the table are buttons for 'Print report' and 'Download spreadsheet'.

Select	Source Nickname	Source ID
<input type="checkbox"/>	TN1Vol1	105F
<input checked="" type="checkbox"/>	TN1Vol2	1060
<input type="checkbox"/>	TN1Vol3	1061
<input type="checkbox"/>	TN1ls	105E

Figure 13-1 Real-time manager in FlashCopy

The next panel displays a table that contains the FlashCopy relationship that you want to delete (Figure 13-2). Click **OK** to confirm the delete operation.

**Note:** Select the “Eliminate data and release allocated target space on space efficient target volumes” option to release the storage space that is allocated for the space-efficient target volume in the repository volume.

**Delete FlashCopy: Real-time**

Selected FlashCopy pairs to delete.

Source nickname ^	Source ID ^	Target nickname ^	Target ID ^
TN1Vol2	1060	TN1FCVol2	2102
Total: 1			

Eliminate data and release allocated target space on space efficient target volumes

Figure 13-2 Delete confirmation in FlashCopy

Deleting the FlashCopy relationship does not change the data on the target volume.

**Note:** You should reformat any previous FlashCopy target volumes that are configured to a System i host before using them on another System i server or partition.

## 13.2.2 Initiate Background Copy

After a FlashCopy relationship is established, it is possible to initiate a background copy as shown in Figure 13-3. This option ensures that all data is copied physically from the source to the target volume, that is all data is available on the target even after the FlashCopy relationship is removed. Select **Real-time manager** → **Copy services** → **FlashCopy**, and then select the FlashCopy relationship that you want to initiate. Choose **Initiate Background Copy** from **Select Action** drop-down menu.

The screenshot shows the 'FlashCopy: Real-time' interface. At the top, there are dropdown menus for 'Storage complex' (IbmStoragePlex), 'Storage unit' (<7589950>), and 'Storage image' (Non-SVT 7589951). Below these are 'Resource type' (LSS) and 'Specify LSS' (21). A 'Refresh' button and 'Last refresh: Tuesday, November 13, 2007 3:35:36 PM CST' are visible. The main area contains a table with columns 'Select', 'Source Nickname', and 'Source ID'. The table lists four source volumes: TN1Vol1 (105F), TN1Vol2 (1060), TN1Vol3 (1061), and TN1ls (105E). The 'Initiate Background Copy' option is selected in the 'Select Action' dropdown menu. To the right, there are buttons for 'Print report' and 'Download spreadsheet', and a table with columns 'Status', 'Persistent', and 'Ch'.

Select	Source Nickname	Source ID
<input type="checkbox"/>	TN1Vol1	105F
<input checked="" type="checkbox"/>	TN1Vol2	1060
<input type="checkbox"/>	TN1Vol3	1061
<input type="checkbox"/>	TN1ls	105E

Figure 13-3 Initiate Background Copy option for FlashCopy

Confirm the option to complete the background copy as shown in Figure 13-4, and click **OK** to continue.

The screenshot shows the 'Initiate background copy: Real-time' dialog box. It contains a message: 'Confirm the volume pairs from Non-SVT 7589951 on which to initiate a background copy and click Ok.' Below the message is a table with columns 'Select', 'Source nickname', 'Source ID', 'Target nickname', 'Target ID', and 'Relationship will remain'. The table lists one source volume pair: TN1Vol2 (1060) and TN1FCVol2 (2102) with a 'YES' relationship. Below the table are checkboxes for 'Reset target writes inhibit' and 'OK' and 'Cancel' buttons.

Select	Source nickname	Source ID	Target nickname	Target ID	Relationship will remain
<input checked="" type="checkbox"/>	TN1Vol2	1060	TN1FCVol2	2102	YES

Figure 13-4 Confirming the background copy



### 13.2.3 Resync Target

The Resync Target action is used to refresh the target volume of a selected FlashCopy relationship. Only data that has changed in the source volume since the initial FlashCopy or the last resynchronization operation is copied to the target volume.

**Note:** You must enable the “Make relationships persistent” and the “Enable change recording” options for the FlashCopy relationship before you can use the Resync Target feature.

To resynchronize the target volume of a FlashCopy relationship:

1. select **Real-time manager** → **Copy services** → **FlashCopy**, and select the FlashCopy relationship. Choose **Resync Target** from the Select Action drop-down menu as shown in Figure 13-5.

The screenshot shows the 'FlashCopy: Real-time' interface. At the top, there are dropdown menus for 'Storage complex' (IbmStoragePlex), 'Storage unit' (<7589950>), and 'Storage image' (Non-SVT 7589951). Below these are 'Resource type' (LSS) and 'Specify LSS' (21). A 'Refresh' button and the text 'Last refresh: Tuesday, November 13, 2007 3:48:28 PM CST' are also present.

The main area contains a table with columns: 'Select', 'Source Nickname', and 'Source ID'. The table lists four source volumes: TN1Vol1, TN1Vol2, TN1Vol3 (selected), and TN1Is. A context menu is open over the table, showing various actions. The 'Resync Target' action is highlighted. To the right of the table, there are buttons for 'Print report' and 'Download spreadsheet', and a table with columns 'Status', 'Persistent', and 'Char'.

Select	Source Nickname	Source ID
<input type="checkbox"/>	TN1Vol1	105F
<input type="checkbox"/>	TN1Vol2	1060
<input checked="" type="checkbox"/>	TN1Vol3	1061
<input type="checkbox"/>	TN1Is	105E

Page 1 of 1 | Total: 4

Status	Persistent	Char
Copy complete	Yes	Disab
Copy complete	Yes	Disab
Out-of-sync tracks > 0	Yes	Enab
Copy complete	Yes	Disab

Figure 13-5 Selecting Resync Target for FlashCopy

2. In the next panel (Figure 13-6), select the options for the resync.

**Resync FlashCopy: Real-time**

Select the copy options to use for re-synchronizing the target volumes and click OK.

Enable change recording  
Enable all ▼

Permit FlashCopy to occur if target volume is online for host access  
Select one ▼

Inhibit writes to target volume  
Disable all ▼

Sequence number  
00000000

Establish target on existing Metro Mirror source

Make relationship(s) persistent

OK Cancel

Figure 13-6 Selecting the resync options for FlashCopy

3. In this example, for Inhibit writes to target volume, select **Enable all** (see Figure 13-7), and then click **OK** to start resynchronization process.

**Resync FlashCopy: Real-time**

Select the copy options to use for re-synchronizing the target volumes and click OK.

Enable change recording  
Enable all ▼

Permit FlashCopy to occur if target volume is online for host access  
Disable all ▼

Inhibit writes to target volume  
Enable all ▼

Sequence number  
00000006

Establish target on existing Metro Mirror source

Make relationship(s) persistent

OK Cancel

Figure 13-7 Enabling writes on the target volume

- You can see the status of the FlashCopy, the options that you selected, when the copy was created, and when the copy was last refreshed from the properties panel (Figure 13-8). To access the properties panel, refer to the Figure 13-5 on page 393, and select **Properties** from Select Action drop-down menu.

**FlashCopy Properties: Real-time**

General	General																																
<a href="#">Out-of-sync tracks</a>	<p>Review the attributes of the current relationship.</p> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 10px;"> <thead> <tr style="background-color: #e0e0e0;"> <th style="text-align: left; padding: 2px;">Attribute</th> <th style="text-align: left; padding: 2px;">Value</th> </tr> </thead> <tbody> <tr><td>Created</td><td>Nov 13, 2007 3:47:40 PM</td></tr> <tr><td>Last Refresh</td><td>Nov 13, 2007 4:39:54 PM</td></tr> <tr><td>Relationship will remain</td><td>Yes</td></tr> <tr><td>Change recording</td><td>Yes</td></tr> <tr><td>Revertible</td><td>No</td></tr> <tr><td>Background copy initiated</td><td>Yes</td></tr> <tr><td>Source is write inhibited</td><td>No</td></tr> <tr><td>Target is write inhibited</td><td>Yes</td></tr> <tr><td>Source ID</td><td>1061</td></tr> <tr><td>Source Nickname</td><td><a href="#">TN1Vol3</a></td></tr> <tr><td>Source Allocation Method</td><td>Standard</td></tr> <tr><td>Target ID</td><td>2103</td></tr> <tr><td>Target Nickname</td><td><a href="#">TN1FCVol3</a></td></tr> <tr><td>Target Allocation Method</td><td>Standard</td></tr> <tr><td colspan="2" style="text-align: right;">Total: 14</td></tr> </tbody> </table>	Attribute	Value	Created	Nov 13, 2007 3:47:40 PM	Last Refresh	Nov 13, 2007 4:39:54 PM	Relationship will remain	Yes	Change recording	Yes	Revertible	No	Background copy initiated	Yes	Source is write inhibited	No	Target is write inhibited	Yes	Source ID	1061	Source Nickname	<a href="#">TN1Vol3</a>	Source Allocation Method	Standard	Target ID	2103	Target Nickname	<a href="#">TN1FCVol3</a>	Target Allocation Method	Standard	Total: 14	
Attribute	Value																																
Created	Nov 13, 2007 3:47:40 PM																																
Last Refresh	Nov 13, 2007 4:39:54 PM																																
Relationship will remain	Yes																																
Change recording	Yes																																
Revertible	No																																
Background copy initiated	Yes																																
Source is write inhibited	No																																
Target is write inhibited	Yes																																
Source ID	1061																																
Source Nickname	<a href="#">TN1Vol3</a>																																
Source Allocation Method	Standard																																
Target ID	2103																																
Target Nickname	<a href="#">TN1FCVol3</a>																																
Target Allocation Method	Standard																																
Total: 14																																	

Figure 13-8 FlashCopy properties

### 13.2.4 FlashCopy Revertible

The FlashCopy Revertible option is used to correct an inconsistency in the FlashCopy relationship by discarding or committing the changes to a target volume. This option is disabled after commit or discard change tasks is performed.

**Note:** The FlashCopy revertible option is valid for a FlashCopy relationship with the persistent, change recording, target write inhibit, and no copy options enabled and with the revertible option disabled.

To enable this option:

1. From the navigation panel, select **Real-time manager** → **Copy services** → **FlashCopy**. Select one of the FlashCopy relationship and choose **FlashCopy Revertible** from **Select Action** drop-down box as shown in Figure 13-9.

**FlashCopy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify LSS:

Last refresh: Tuesday, November 13, 2007 4:54:46 PM CST

Select	Source Nickname	Source ID	--- Select Action ---	Status	Persistent	Char
<input type="checkbox"/>	TN1Vol1	105F	FlashCopy Revertible	Out-of-sync tracks > 0	Yes	Disab
<input type="checkbox"/>	TN1Vol2	1060	Delete	Out-of-sync tracks > 0	Yes	Enab
<input checked="" type="checkbox"/>	TN1Vol3	1061	Initiate Background Copy	Copy complete	Yes	Enab
<input type="checkbox"/>	TN1Is	105E	Resync Target	Copy complete	Yes	Disab
Page 1 of 1		Total: 4	Reverse FlashCopy	1		

Properties:  --- Table Actions ---  
 Select All, Deselect All, Select All On Current Page, Deselect All On Current Page, Show Filter Row, Clear All Filters, Edit Sort, Clear All Sorts

Figure 13-9 FlashCopy Revertible

2. In the Select common options panel, select the necessary option, and click **Next** to continue, as shown in Figure 13-10.

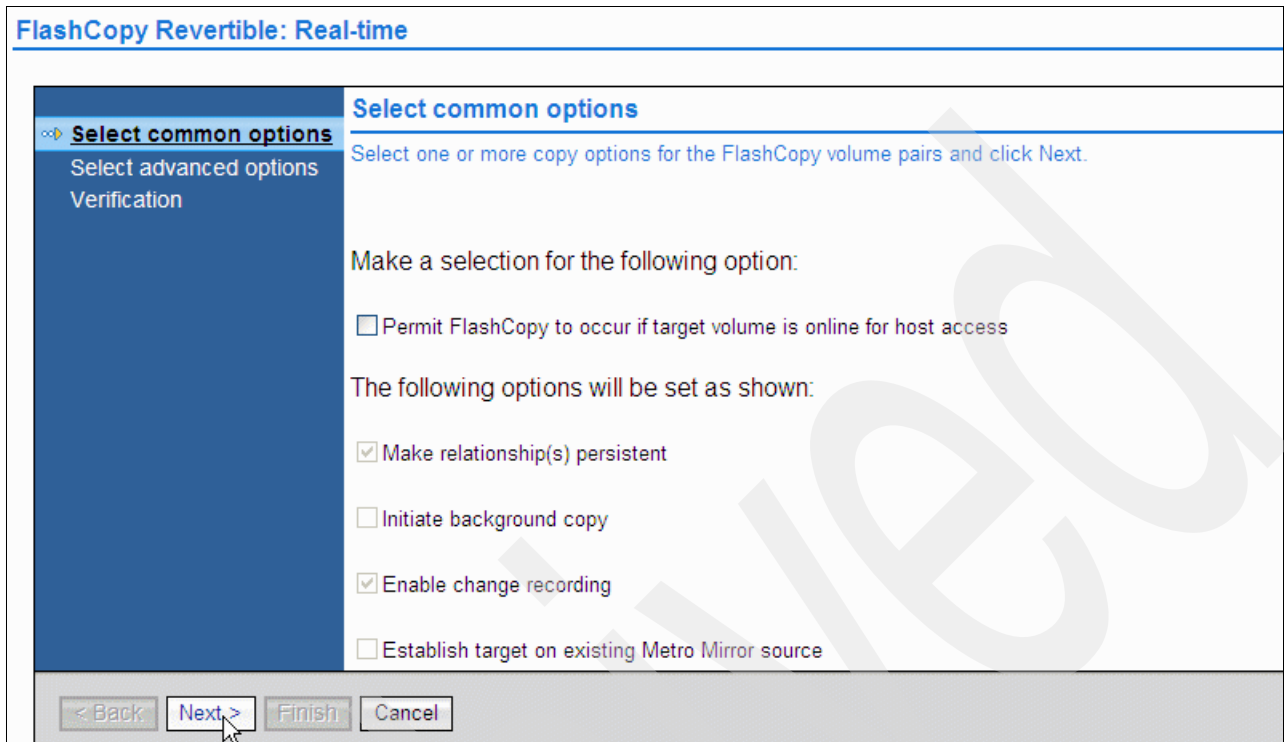


Figure 13-10 Revertible FlashCopy options

3. Enabling the FlashCopy Revertible option can impact the ability to use more advanced options. In the Select advanced options panel (Figure 13-11), you can see that because of previous selections, no advanced options are available. You can enter only the sequence number. Click **Next** to continue.

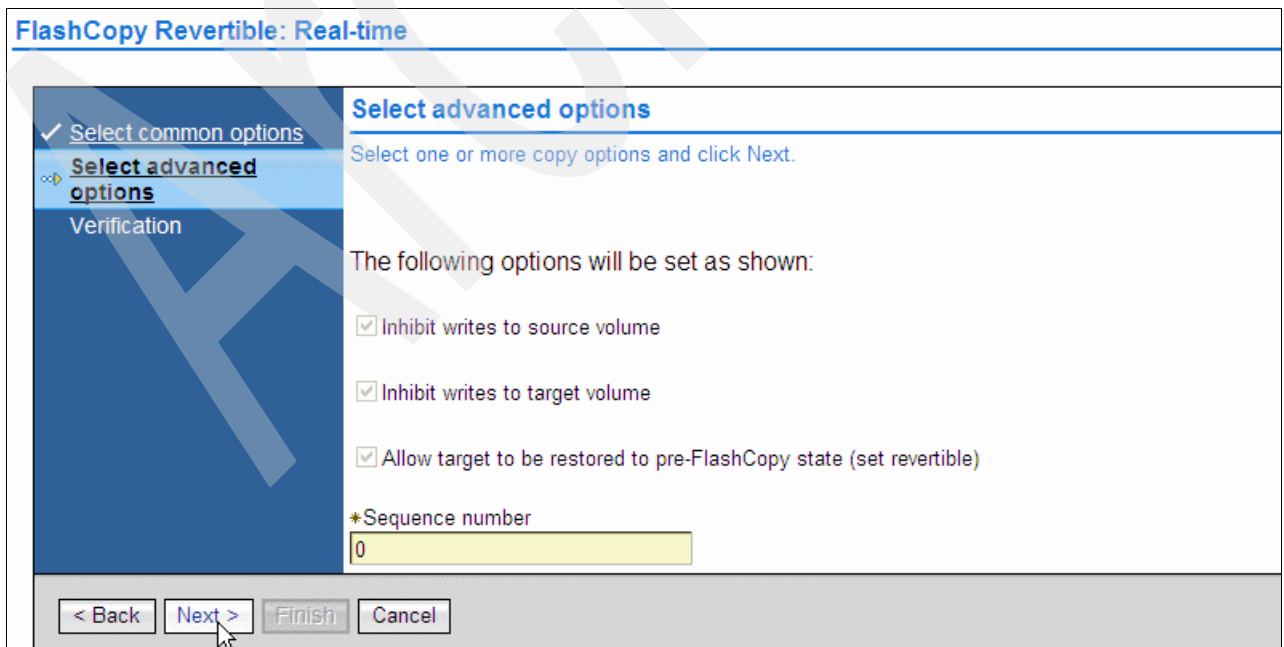


Figure 13-11 Advanced functions of FlashCopy Revertible

- On the Verification panel, verify the options, and click **Finish** to continue the revertible operation (see Figure 13-12).

- ✓ Select common options
- ✓ Select advanced options
- Verification

### Verification

Verify the number of relationships and the selected copy options before clicking Finish.

Sequence number: 0

Copy options

Enabled copy options	Disabled copy options
Fail relationship if space efficient target volume runs out of space	Write inhibit source volume if space efficient targ
Make relationship(s) persistent	Allow target volumes to be space efficient volum
Inhibit writes to source volume	Initiate background copy
Enable change recording	Permit FlashCopy to occur if target volume is on
Allow target to be restored to pre-FlashCopy state (set revertible)	Establish target on existing Metro Mirror source
Inhibit writes to target volume	

FlashCopy to set revertible

Source nickname	Source ID	Target nickname
TN1Vol3	1061	TN1FCVol3

< Back
Next >
Finish
Cancel

Figure 13-12 Verification of Revertible options

### 13.2.5 Reverse FlashCopy

From the navigation panel, select **Real-time manager** → **Copy services** → **FlashCopy**. Select one of the FlashCopy relationship, and click **Reverse FlashCopy** from **Select Action** drop-down box (see Figure 13-13).

**FlashCopy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify LSS:

Last refresh: Tuesday, November 13, 2007 4:54:46 PM CST

Select ^	Source Nickname ^	Source ID ^	--- Select Action ---	Status ^	Persistent ^	Char
<input type="checkbox"/>	TN1Vol1	105F	--- Select Action ---	Out-of-sync tracks > 0	Yes	Disab
<input type="checkbox"/>	TN1Vol2	1060	Create ...	Out-of-sync tracks > 0	Yes	Enab
<input checked="" type="checkbox"/>	TN1Vol3	1061	FlashCopy Revertible	Copy complete	Yes	Enab
<input type="checkbox"/>	TN1ls	105E	Delete	Copy complete	Yes	Disab
Page 1 of 1		Total: 4	Initiate Background Copy	1		
			Resync Target			
			Reverse FlashCopy			
			Properties			
			Reset Target Write Inhibit			
			--- Table Actions ---			
			Select All			
			Deselect All			
			Select All On Current Page			
			Deselect All On Current Page			
			Show Filter Row			
			Clear All Filters			
			Edit Sort			
			Clear All Sorts			

Figure 13-13 Reverse FlashCopy

On the panel shown in Figure 13-14, you can select one or more copy options to reverse the FlashCopy relationship. That is, the original source volume is now the target, whereas the original target volume becomes the source of the FlashCopy relationship.

When a relationship is reversed, only the data that is required to bring the target current to the source's point-in-time is copied. If no updates were made to the target since the last refresh, the direction change can be used to restore the source to the previous point-in-time state.

**Reverse FlashCopy: Real-time**

Select one or more copy options to use when reversing the volume pairs and click OK.

Make relationship(s) persistent : Yes

Enable change recording  
Disable all

Permit FlashCopy to occur if target volume is online for host access  
Disable all

Inhibit writes to target volume  
Disable all

Fast Reverse  
Disable all

Sequence number  
00000006

OK Cancel

Figure 13-14 Options to reverse the FlashCopy relationship



## 13.3 Metro Mirror GUI

In this section, we explain the options that are available to manage Metro Mirror using the GUI panels.

### 13.3.1 Recovery Failover

The Recovery Failover option is used to confirm which volume pairs to use during a failover operation to the recovery site. This option allows the target volumes at the recovery site to be used to restart the production environment during planned or unplanned outage.

From the navigation panel of the target system, select **Real-time manager** → **Copy services** → **Metro Mirror/ Global Copy**. Select one of the Metro Mirror relationship and select Recovery Failover from the Select Action drop-down menu (see Figure 13-15).

The screenshot displays the 'Metro Mirror / Global Copy: Real-time' interface. At the top, there are three dropdown menus for 'Storage complex' (set to 'IbmStoragePlex'), 'Storage unit' (set to '<7589950>'), and 'Storage image' (set to 'SVT 7589952'). Below these are 'Resource type' (set to 'All volumes') and 'Specify Storage type' (set to 'All FB volumes'). A 'Refresh' button and a timestamp 'Last refresh: Wednesday, November 14, 2007 5:19:12 PM CST' are also present.

The main area contains a table with columns: 'Select', 'Source ESS serial', 'Source Nickname', 'Target serial', 'Target Nickname', and 'Target I'. The table lists four source volumes (TN1Vol1, TN1Vol2, TN1Vol3, TN1ls) with serial number 7589951. The 'TN1ls' row is selected. A context menu is open over the table, listing actions such as 'Create...', 'Delete', 'Suspend...', 'Convert to synchronous', 'Resume', 'Modify', 'Recovery Failover', 'Recovery Failback', and 'Properties'. The 'Recovery Failover' option is highlighted.

Select	Source ESS serial	Source Nickname	Target serial	Target Nickname	Target I
<input type="checkbox"/>	7589951	TN1Vol1		Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol2		Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol3		Not retrieved	0002:101
<input checked="" type="checkbox"/>	7589951	TN1ls		Not retrieved	0002:101

Figure 13-15 Selecting the Recovery Failover option

Click **OK** to confirm the action to failover for the selected source. See Figure 13-16.

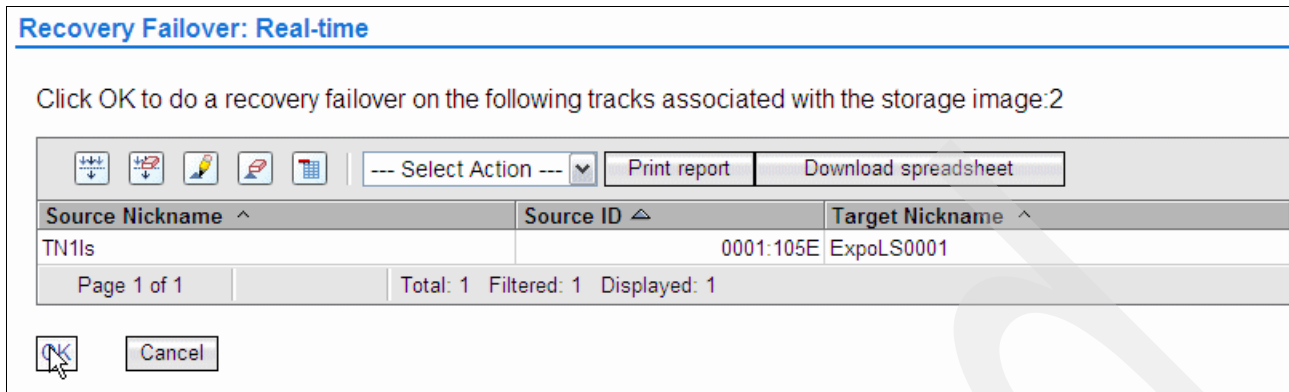


Figure 13-16 Confirming the failover

When the failover is initiated, the mirrored LUNs are in a *Suspended* state as shown in Figure 13-17. The previous source volume becomes the target volume and the previous target volume becomes the source volume.

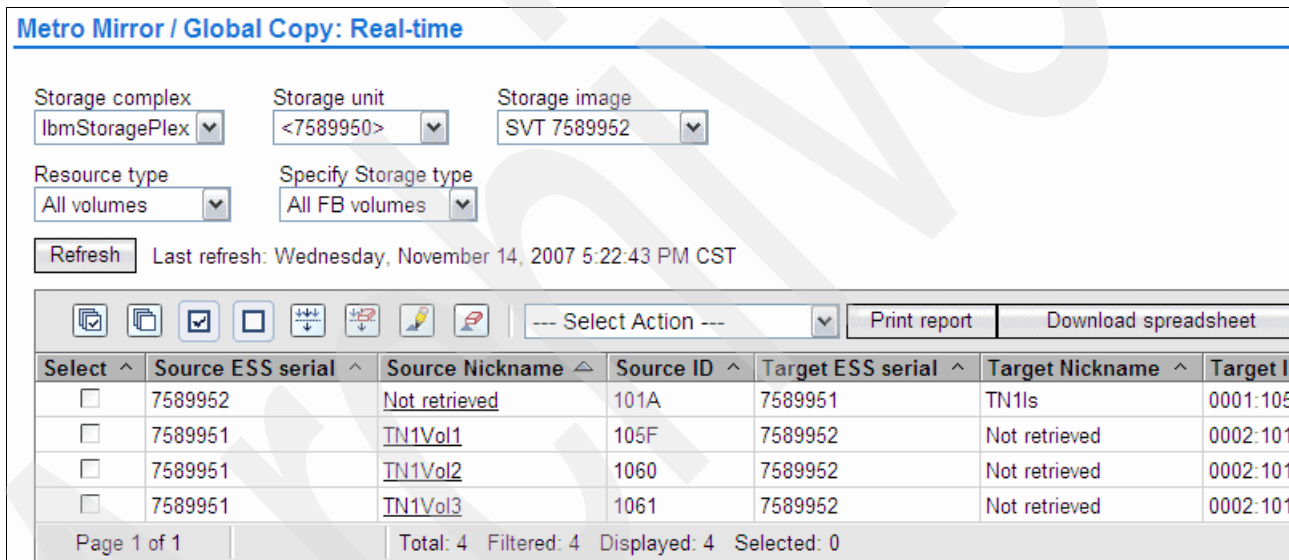


Figure 13-17 Failover initiated, mirrored LUNs in Suspended state

### 13.3.2 Recovery Failback

The Recovery Failback option is used to send the changed data from the recovery site back to the production site to synchronize the volume pairs. It changes the direction of the Metro Mirror data flow from the original target to the original source.

From the navigation panel of the target system, select **Real-time manager** → **Copy services** → **Metro Mirror/ Global Copy**. Because the failback process is done after the failover process, select the Metro Mirror relationship that has a data flow direction from the original target to the original source and is in *Suspended* state as shown in Figure 13-18. Select **Recovery Failback** from **Select Action** drop-down menu.

The screenshot shows the 'Metro Mirror / Global Copy: Real-time' interface. At the top, there are filter controls for 'Storage complex' (set to 'IbmStoragePlex'), 'Storage unit' (set to '<7589950>'), and 'Storage image' (set to 'SVT 7589952'). Below these are 'Resource type' (set to 'All volumes') and 'Specify Storage type' (set to 'All FB volumes'). A 'Refresh' button and the text 'Last refresh: Wednesday, November 14, 2007 5:25:42 PM CST' are also present.

The main area contains a table with columns: 'Select', 'Source ESS serial', 'Source Nickname', 'Target serial', 'Target Nickname', and 'Target I'. The first row is selected, and a context menu is open over it, showing options like 'Create...', 'Delete', 'Suspend...', 'Convert to synchronous', 'Resume', 'Modify', 'Recovery Failover', and 'Recovery Failback' (which is highlighted). Other options include 'Properties', 'Table Actions', 'Select All', 'Deselect All', 'Select All On Current Page', 'Deselect All On Current Page', 'Show Filter Row', 'Clear All Filters', 'Edit Sort', and 'Clear All Sorts'.

Select	Source ESS serial	Source Nickname	Target serial	Target Nickname	Target I
<input checked="" type="checkbox"/>	7589952	Not retrieved		TN1Is	0001:105
<input type="checkbox"/>	7589951	TN1Vol1		Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol2		Not retrieved	0002:101
<input type="checkbox"/>	7589951	TN1Vol3		Not retrieved	0002:101

Figure 13-18 Selecting the Recovery Failback option

In the next panel (Figure 13-19), confirm your action to complete the failback, and click **OK** to switch the direction of the data flow.

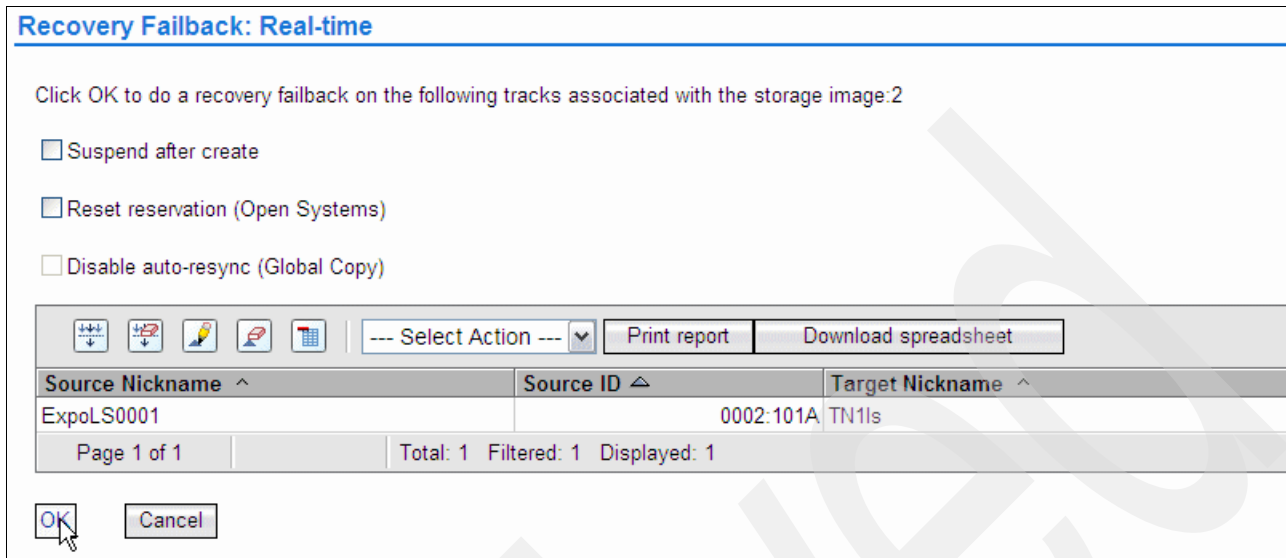


Figure 13-19 Confirming the recovery failback

When you refresh the panel, as shown in Figure 13-20, you see that the data flow is now from source 52 to target 51. Having fully synchronized, the state changes to *Full duplex*. The direction is still from source 52 to target 51.

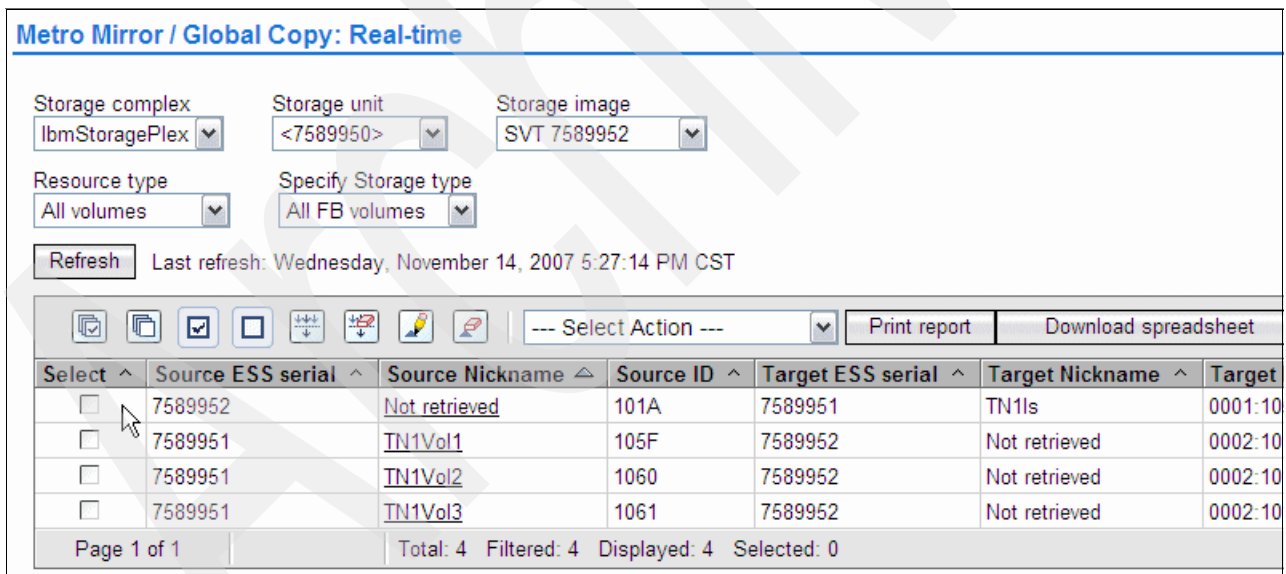


Figure 13-20 Fully synchronized failback

### 13.3.3 Suspend

The Suspend option is used to suspend the copy operation from the source volume to the target volume. Any host write updates after a suspend will result in unsynchronized mirror pairs. Use the Suspend option (Figure 13-21) for short outages and planned maintenance where it is not necessary to switch to the backup system.

From the navigation panel, select **Real-time manager** → **Copy services** → **Metro Mirror/ Global Copy**. Select one of the Metro Mirror relationship that will be suspended, and choose **Suspend** from the Select Action drop-down menu (see Figure 13-21).

The screenshot shows the 'Metro Mirror / Global Copy: Real-time' interface. At the top, there are filters for 'Storage complex' (IbmStoragePlex), 'Storage unit' (<7589950>), and 'Storage image' (SVT 7589952). Below these are 'Resource type' (All volumes) and 'Specify Storage type' (All FB volumes). A 'Refresh' button and 'Last refresh: Wednesday, November 14, 2007 5:28:57 PM CST' are also visible.

The main area contains a table with columns: 'Select', 'Source ESS serial', 'Source Nickname', 'S serial', 'Target Nickname', and 'Target'. The first row is selected, and a context menu is open over it, with 'Suspend...' highlighted. Other menu items include 'Create...', 'Delete', 'Convert to synchronous', 'Resume', 'Modify', 'Recovery Failover', 'Recovery Failback', 'Properties', and 'Table Actions'.

Select	Source ESS serial	Source Nickname	S serial	Target Nickname	Target
<input checked="" type="checkbox"/>	7589952	Not retrieved		TN1Is	0001:10
<input type="checkbox"/>	7589951	TN1Vol1		Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol2		Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol3		Not retrieved	0002:10

Figure 13-21 Selecting the Suspend option to suspend the Metro Mirror relationship

As in all previous examples, for Suspend, you also have the option to confirm your action (Figure 13-22). You can suspend on either the source or target system. If this is a planned outage, then suspend from the source system. You can suspend from the target if the source is no longer available.

The screenshot shows a dialog box titled 'Suspend Metro Mirror: Real-time'. It contains the text 'Select source or target for volumes to be suspended.' and two radio button options: 'Suspend at source' (which is selected) and 'Suspend at target'. At the bottom, there are 'OK' and 'Cancel' buttons.

Figure 13-22 Select volumes to be suspended

As shown in Figure 13-23, Metro Mirror is now in a *Suspended* state.

**Metro Mirror / Global Copy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Wednesday, November 14, 2007 5:31:41 PM CST

--- Select Action ---

Select ^	Source ESS serial ^	Source Nickname ^	Source ID ^	Target ESS serial ^	Target Nickname ^	Target I
<input type="checkbox"/>	7589952	Not retrieved	101A	7589951	TN1s	0001:10
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:10

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 13-23 Metro Mirror suspended

### 13.3.4 Resume

The Resume option is used to start a background copy and copy unsynchronized tracks from suspended Metro Mirror pairs.

From the navigation panel, select **Real-time manager** → **Copy services** → **Metro Mirror/ Global Copy**. Select the Metro Mirror relationship that you want to resume, and select **Resume** from the Select Action drop-down menu (see Figure 13-24).

**Metro Mirror / Global Copy: Real-time**

Storage complex:  Storage unit:  Storage image:

Resource type:  Specify Storage type:

Last refresh: Wednesday, November 14, 2007 5:31:41 PM CST

--- Select Action ---

Select ^	Source ESS serial ^	Source Nickname ^	Source ID ^	Target ESS serial ^	Target Nickname ^	Target I
<input checked="" type="checkbox"/>	7589952	Not retrieved	101A	7589951	TN1s	0001:10
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:10
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:10

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 1

--- Select Action ---  
 Create...  
 Delete  
 Suspend...  
 Convert to synchronous  
**Resume**  
 Modify  
 Recovery Failover  
 Recovery Failback  
 Properties  
 --- Table Actions ---  
 Select All  
 Deselect All

Figure 13-24 Selecting the Resume option for Metro Mirror

On the next panel, confirm the option to resume the Metro Mirror pair, and click **OK** to continue, as shown in Figure 13-25.

**Resume Metro Mirror: Real-time**

Click OK to resume the tracks associated the storage image:2

Permit read access from target  
 Create relationship even if target is online to a host (zSeries)  
 Reset reservation (Open Systems)  
 Suspend Metro Mirror relationship after initial copy  
 Enable critical volume mode

Figure 13-25 Confirming the resume option

The time during which the mirror is suspended and the amount of changes that occur determine the time that it takes for the mirror to return to a fully synchronized full duplex state (see Figure 13-26).

**Metro Mirror / Global Copy: Real-time**

Storage complex:   
 Storage unit:   
 Storage image:   
 Resource type:   
 Specify Storage type:   
 Last refresh: Wednesday, November 14, 2007 5:33:25 PM CST

Select	Source ESS serial	Source Nickname	Source ID	Target ESS serial	Target Nickname	Target ID
<input type="checkbox"/>	7589952	Not retrieved	101A	7589951	TN1Is	0001:104
<input type="checkbox"/>	7589951	TN1Vol1	105F	7589952	Not retrieved	0002:107
<input type="checkbox"/>	7589951	TN1Vol2	1060	7589952	Not retrieved	0002:107
<input type="checkbox"/>	7589951	TN1Vol3	1061	7589952	Not retrieved	0002:107

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4 Selected: 0

Figure 13-26 Metro Mirror Resume result

## 13.4 Global Mirror GUI

In this section, we explain the options that are available to manage Global Mirror using the GUI panels. Figure 13-27 shows the scenario that we use in this section.




Figure 13-27 Global Mirror relationship scenario



### 13.4.1 Create

Only one active Global Mirror session can exist between two storage systems. To create a new session, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the storage unit or image that will be the master for Global Mirror session, and choose **Create** from the Select Action drop-down menu (see Figure 13-28).

**Important:** Before we start to create the Global Mirror session, we need to set up the PPRC paths between the local site and the remote site. We also need to set up the Global Copy relationship and Flash Copy relationship for the Global Mirror session.

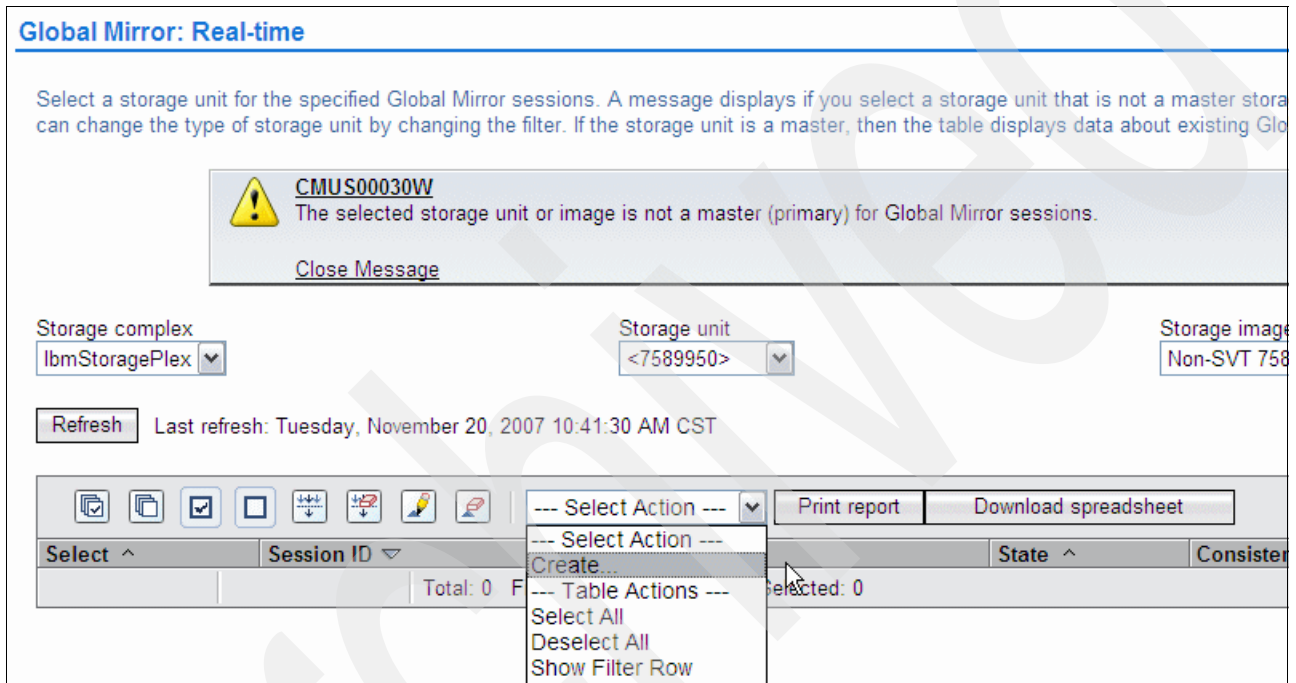


Figure 13-28 Selecting the Create option for Global Mirror

The select volumes panel display as shown in Figure 13-29. Select the source volumes of the Global Mirror session by expanding the required storage unit and the LSS. The selected volumes display in the Selected volumes table. Click **Next** to continue.

**Note:** If the Global Copy and FlashCopy relationship that is needed in Global Mirror session are not created yet, click **Create Metro Mirror** to start creating Global Copy relationships and then click **Create FlashCopy** to start creating FlashCopy relationship.

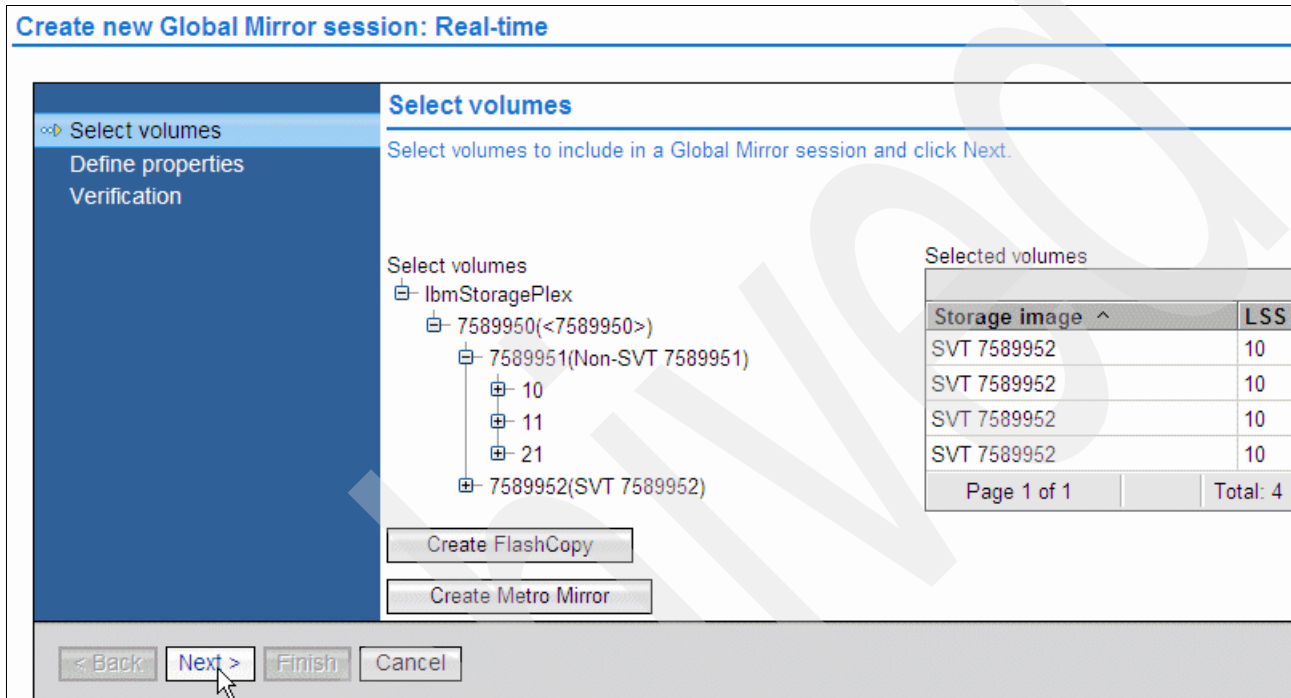


Figure 13-29 Selecting source volumes for Global Mirror

Define the properties of the Global Mirror session in the next panel as shown in Figure 13-30. Select the session ID that is available and select the LSS that will be used as the master LSS for the Global Mirror session. Click **Next** to continue.

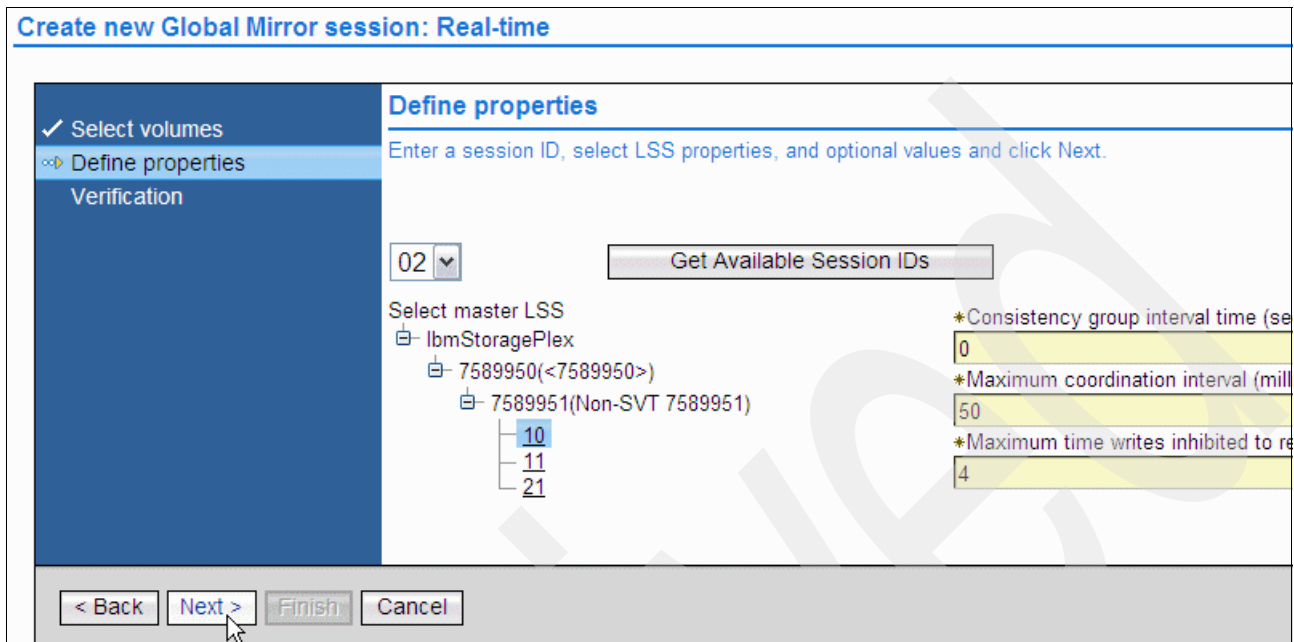


Figure 13-30 Define properties for Global Mirror

From the verification panel, shown in Figure 13-31, review all details, and click **Finish** to start Global Mirror session creation process.

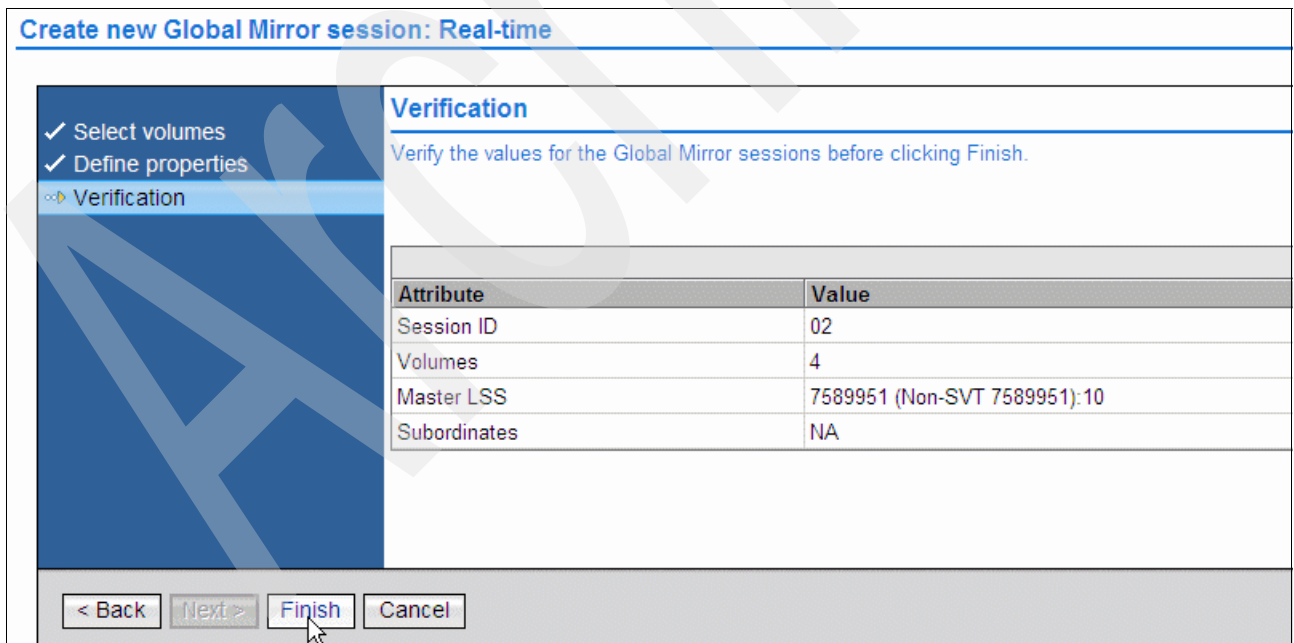


Figure 13-31 Verification panel of Global Mirror create process

To check the newly created Global Mirror session, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the storage unit or image that is configured as the master as shown in Figure 13-32.

**Global Mirror: Real-time**

Select a storage unit for the specified Global Mirror sessions. A message displays if you select a storage unit that is not a master storage unit. If the storage unit is a master, then the table displays data about existing Global Mirror sessions.

Storage complex:  Storage unit:  Storage image:

Last refresh: Wednesday, November 21, 2007 2:20:45 PM CST

--- Select Action ---

Select ^	Session ID	Storage images ^	State ^	Consistent
<input type="checkbox"/>		02	1 Running	Nov 21, 2007

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1 Selected: 0

Figure 13-32 New Global Mirror session

### 13.4.2 Delete

To delete a Global Mirror instance, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session, and choose **Delete** from the Select Action drop-down menu as shown in Figure 13-33.

**Global Mirror: Real-time**

Select a storage unit for the specified Global Mirror sessions. A message displays if you select a storage unit that is not a master storage unit. If the storage unit is a master, then the table displays data about existing Global Mirror sessions.

Storage complex:  Storage unit:  Storage image:

Last refresh: Wednesday, November 21, 2007 2:20:45 PM CST

--- Select Action ---

Select ^	Session ID	Storage images ^	State ^	Consistent
<input checked="" type="checkbox"/>		02	1 Running	Nov 21, 2007

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1 Selected: 1

--- Select Action ---  
 Create...  
 Delete  
 Modify...  
 Pause...  
 View session volumes  
 Properties  
 --- Table Actions ---  
 Select All  
 Deselect All  
 Select All On Current Page

Figure 13-33 Selecting the Delete option for Global Mirror

In the next panel (shown in Figure 13-34), click **OK** to confirm that you want to delete the Global Mirror session.




Figure 13-34 Confirming the Delete option for Global Mirror

### 13.4.3 Modify

To modify any of the properties of the Global Mirror, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session and click **Modify** from the Select Action drop-down menu as shown in Figure 13-35.

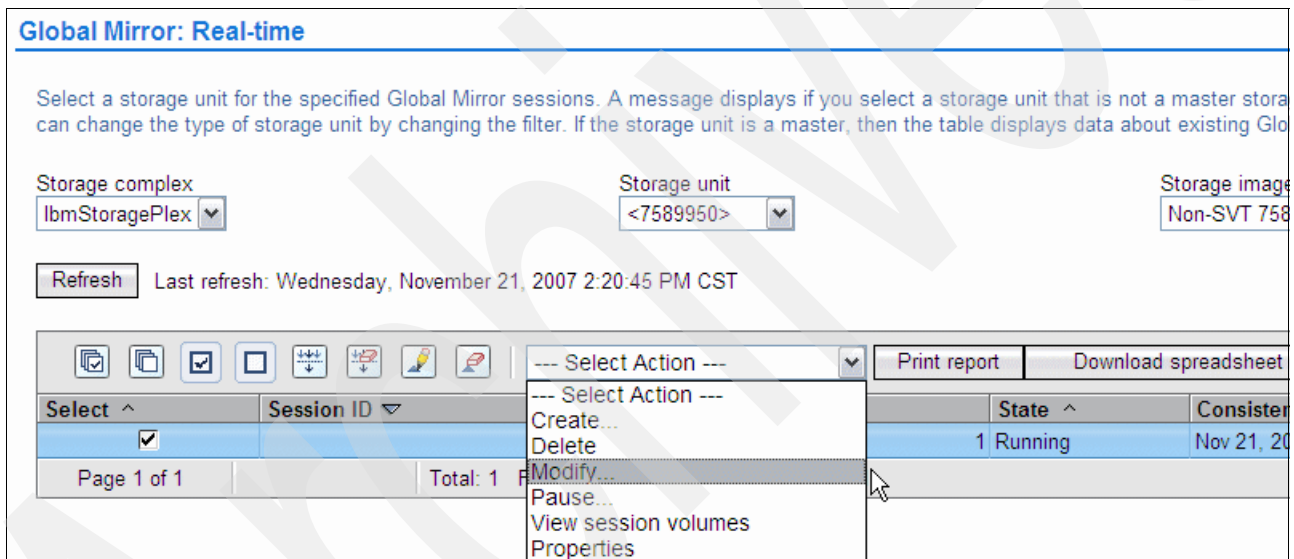


Figure 13-35 Selecting the Modify option for modify the Global Mirror properties

In the Select volumes panel, shown in Figure 13-36, select the volumes whose property you want to modify. The volumes that are selected can be removed from the session. You can also add new volume to the Global Mirror session.

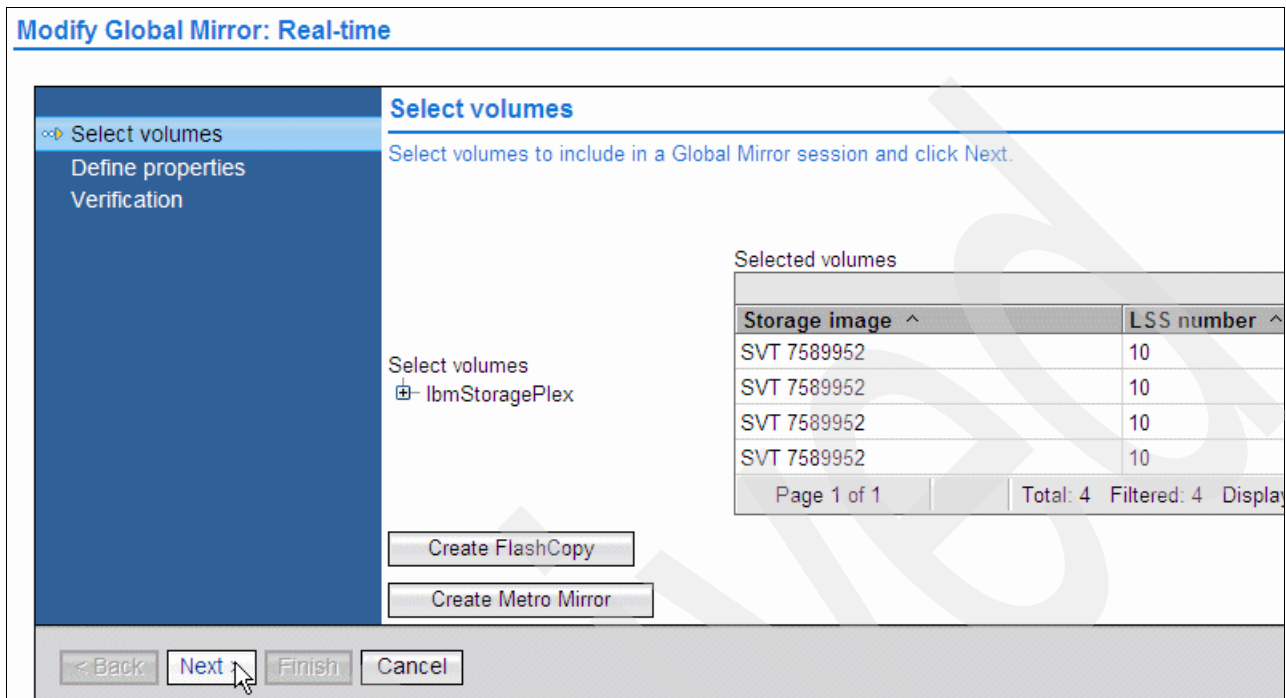


Figure 13-36 Selecting the volume for which to modify the properties

In the next panel, you can modify the Global Mirror session properties (see Figure 13-37).

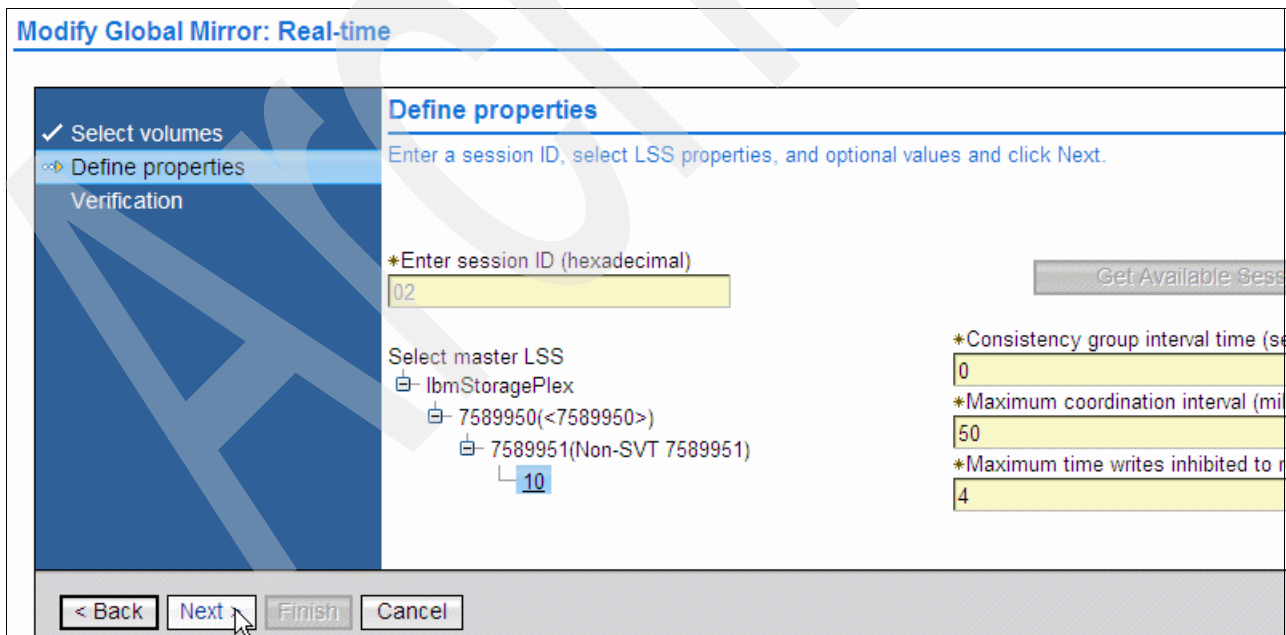


Figure 13-37 Modify Global Mirror properties

In the Verification panel, review the details, and if everything is correct, click **Finish** to confirm the modification, as shown in Figure 13-38.

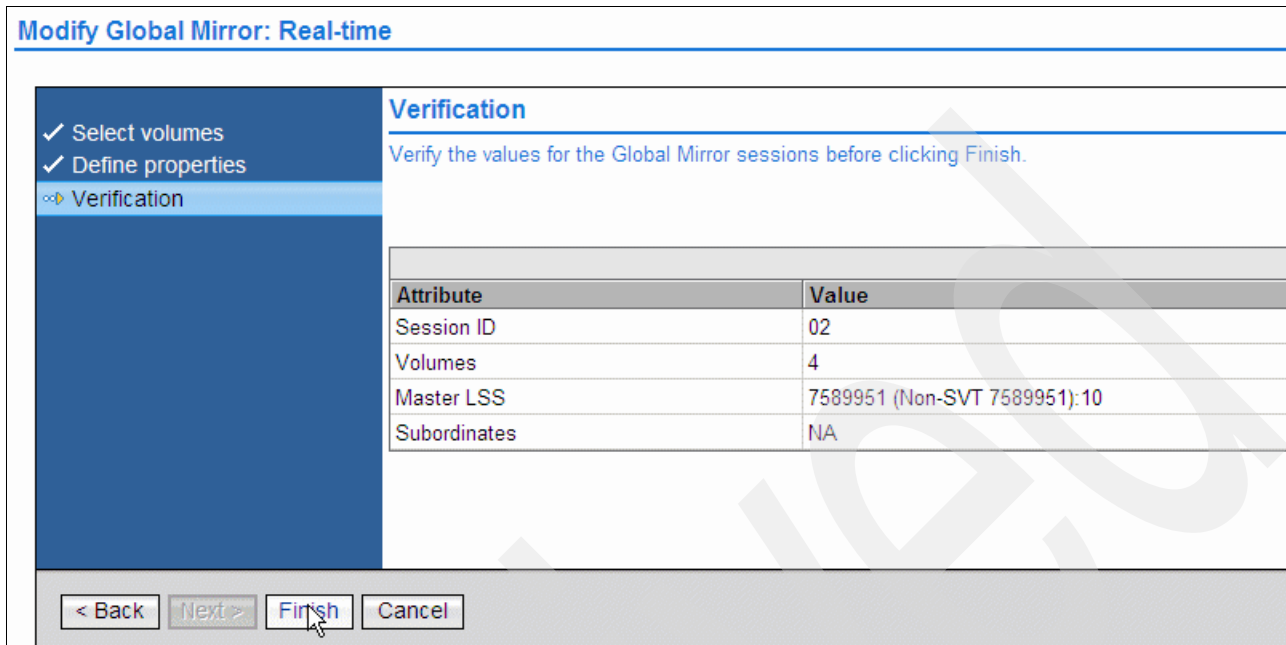


Figure 13-38 Verifying the details of the modified properties

### 13.4.4 Pause

To pause Global Mirror, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session and choose **Pause** from the Select Action drop-down menu as shown in Figure 13-39.

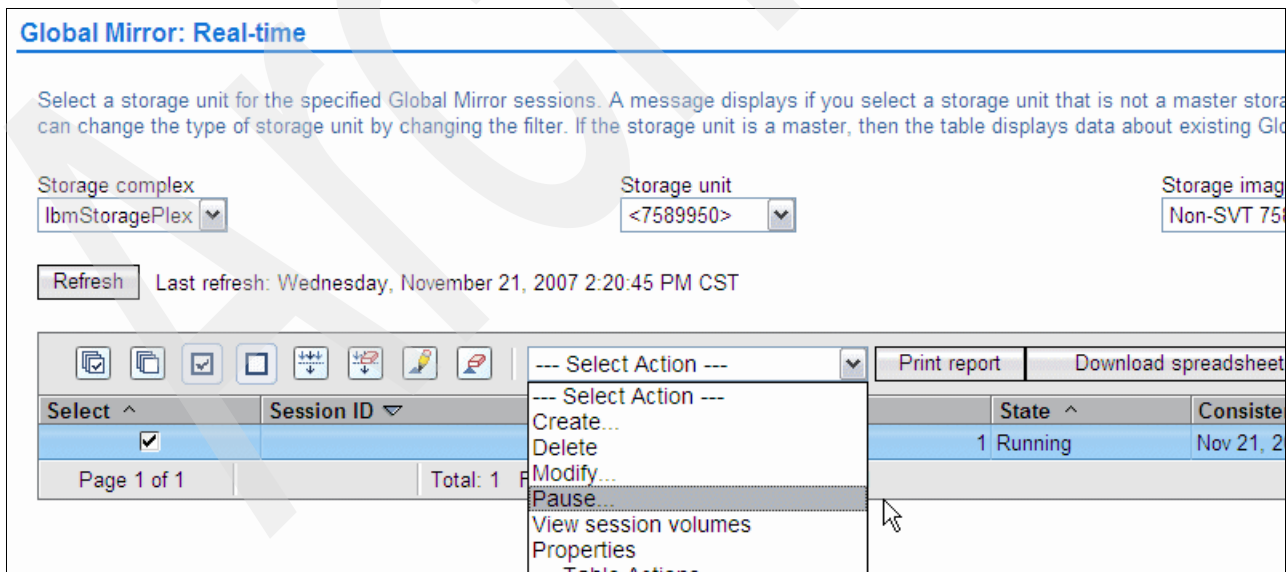


Figure 13-39 Selecting the Pause option to pause the Global Mirror relationship

In the next panel, shown in Figure 13-40, click **OK** to confirm the pause action for the Global Mirror session.




Figure 13-40 Confirming the Global Mirror Pause action

As shown in Figure 13-41, the Global Mirror instance is now in the *Paused* state.

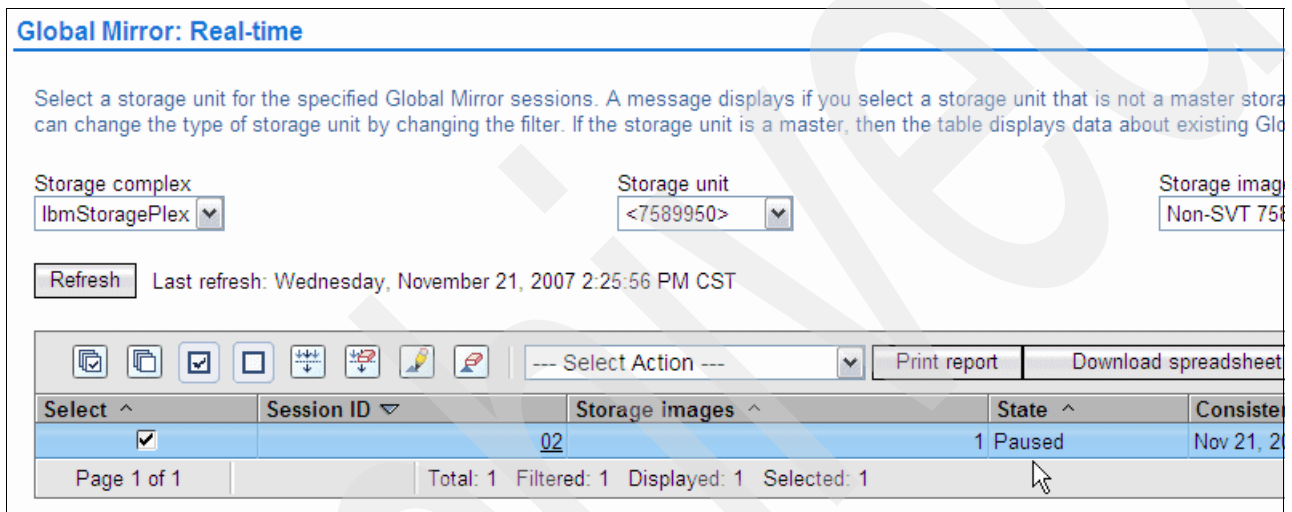


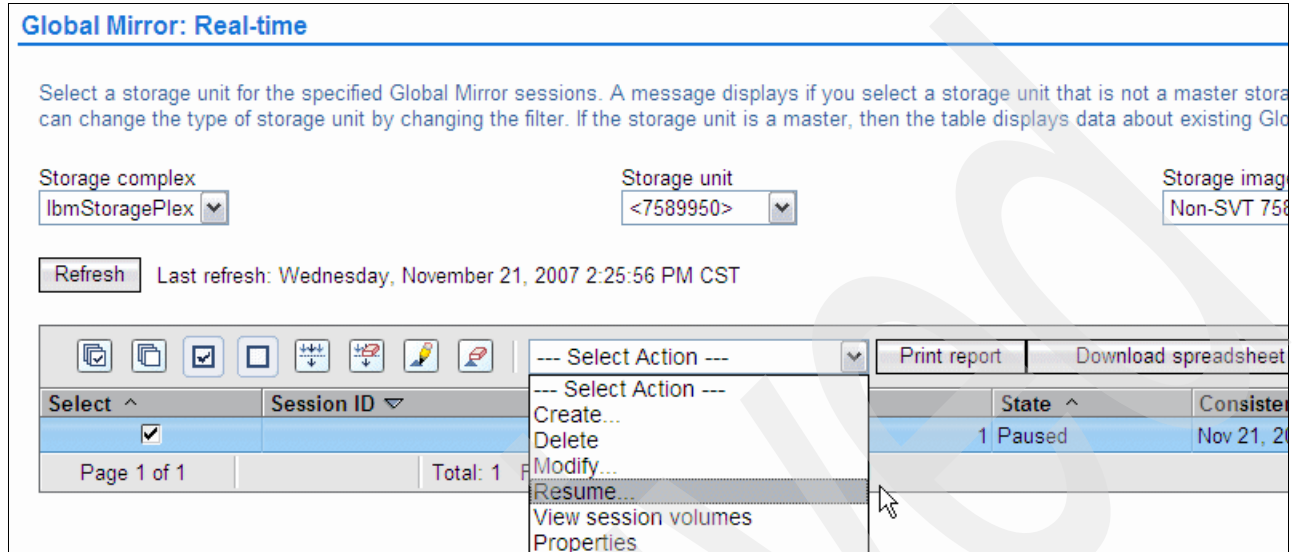
Figure 13-41 Global Mirror instance paused

**Note:** Pausing a Global Mirror session only pauses Global Mirror consistency group processing but leaves Global Copy running.



## 13.4.5 Resume

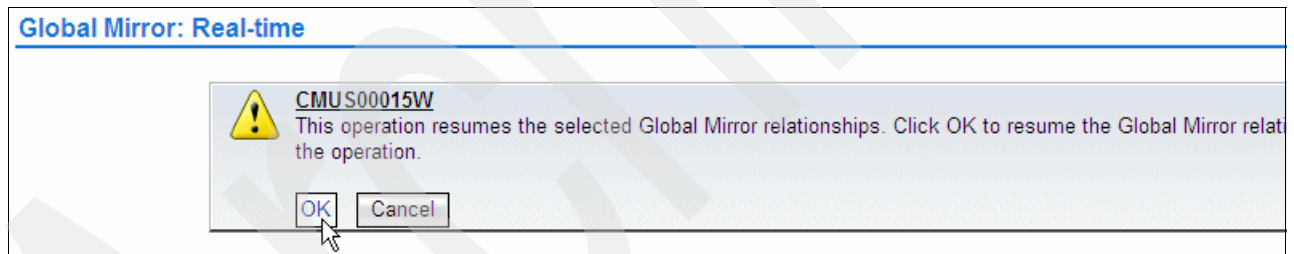
To resume a paused Global Mirror, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session and choose **Resume** from the Select Action drop-down menu as shown in Figure 13-42.



The screenshot shows the 'Global Mirror: Real-time' interface. At the top, there are filters for 'Storage complex' (IBM StoragePlex), 'Storage unit' (<7589950>), and 'Storage image' (Non-SVT 758). A 'Refresh' button is present, with the last refresh time being Wednesday, November 21, 2007 2:25:56 PM CST. Below the filters is a toolbar with various icons and a 'Select Action' dropdown menu. The dropdown menu is open, showing options: 'Create...', 'Delete', 'Modify...', 'Resume...', 'View session volumes', and 'Properties'. The 'Resume...' option is highlighted. To the right of the dropdown are 'Print report' and 'Download spreadsheet' buttons. Below the toolbar is a table with columns 'Select', 'Session ID', 'State', and 'Consistent'. The table has one row with '1 Paused' and 'Nov 21, 2007'. The 'Resume...' option is also visible in the table's context menu.

Figure 13-42 Global Mirror resume

In the next panel, shown in Figure 13-43, click **OK** to confirm the option to resume Global Mirror session.



The screenshot shows a confirmation dialog box titled 'Global Mirror: Real-time'. The dialog box has a yellow warning icon and the text: 'CMUS00015W This operation resumes the selected Global Mirror relationships. Click OK to resume the Global Mirror relationships. Click Cancel to cancel the operation.' Below the text are 'OK' and 'Cancel' buttons. A mouse cursor is pointing at the 'OK' button.

Figure 13-43 Confirm Global Mirror resume action

As shown in Figure 13-44, the state changes to *Running* to reflect the fact that Global Mirror has resumed.

**Global Mirror: Real-time**

Select a storage unit for the specified Global Mirror sessions. A message displays if you select a storage unit that is not a master storage unit. If the storage unit is a master, then the table displays data about existing Global Mirror sessions.

Storage complex:  Storage unit:  Storage image:

Last refresh: Wednesday, November 21, 2007 2:25:56 PM CST

--- Select Action ---

Select ^	Session ID	Storage images ^	State ^	Consistent
<input checked="" type="checkbox"/>		02	1 Running	Nov 21, 2007

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1 Selected: 1

Figure 13-44 Global Mirror pause status

### 13.4.6 View session volumes

You can view the session volumes of the Global Mirror by selecting **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session and choose **View session volumes** from the Select Action drop-down menu as shown in Figure 13-45.

**Global Mirror: Real-time**

Select a storage unit for the specified Global Mirror sessions. A message displays if you select a storage unit that is not a master storage unit. If the storage unit is a master, then the table displays data about existing Global Mirror sessions.

Storage complex:  Storage unit:  Storage image:

Last refresh: Wednesday, November 21, 2007 2:25:56 PM CST

--- Select Action ---

Select ^	Session ID	Storage images ^	State ^	Consistent
<input checked="" type="checkbox"/>		02	1 Running	Nov 21, 2007

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1 Selected: 1

--- Select Action ---  
 Create...  
 Delete  
 Modify...  
 Pause...  
 View session volumes  
 Properties  
 --- Table Actions ---  
 Select All

Figure 13-45 Selecting the View session volumes action for Global Mirror

The status of the volumes in Global Mirror displays in the next panel (Figure 13-46). Click **OK** to return to the previous panel.

**Global Mirror session volumes: Real-time**

Review the volumes that are associated with the session on Non-SVT 7589951 and click Ok.

Print report    Download spreadsheet

Nickname ^	ID ^	Status ^	Type ^	Capacity ^	RAID ^	Extent pool ^
TN1Is	105E	Normal	FB	35.2	RAID 5	Non SVT 0
TN1Vol1	105F	Normal	FB	35.2	RAID 5	Non SVT 0
TN1Vol2	1060	Normal	FB	35.2	RAID 5	ExpoFLASH
TN1Vol3	1061	Normal	FB	35.2	RAID 5	ExpoFLASH

Page 1 of 1    Total: 4    Filtered: 4    Displayed: 4

OK

Figure 13-46 Status of the Global Mirror session volumes

### 13.4.7 Properties

To view the properties of the Global Mirror or to view any errors, select **Real-time manager** → **Copy services** → **Global Mirror** from the left navigation panel. Select the Global Mirror session and choose **Properties** from the Select Action drop-down menu as shown in Figure 13-47.

**Global Mirror: Real-time**

Select a storage unit for the specified Global Mirror sessions. A message displays if you select a storage unit that is not a master storage unit. You can change the type of storage unit by changing the filter. If the storage unit is a master, then the table displays data about existing Global Mirror sessions.

Storage complex:     Storage unit:     Storage image:

   Last refresh: Wednesday, November 21, 2007 2:25:56 PM CST

Select ^	Session ID ^	State ^	Consistent ^
<input checked="" type="checkbox"/>		1 Running	Nov 21, 2007

Page 1 of 1    Total: 1

--- Select Action ---

- Select Action ---
- Create...
- Delete
- Modify...
- Pause...
- View session volumes
- Properties**
- Table Actions ---

Print report    Download spreadsheet

Figure 13-47 Selecting the Properties action to view the properties of the Global Mirror

The General properties panel displays, as shown in Figure 13-48. Choose **Failures** to view errors.

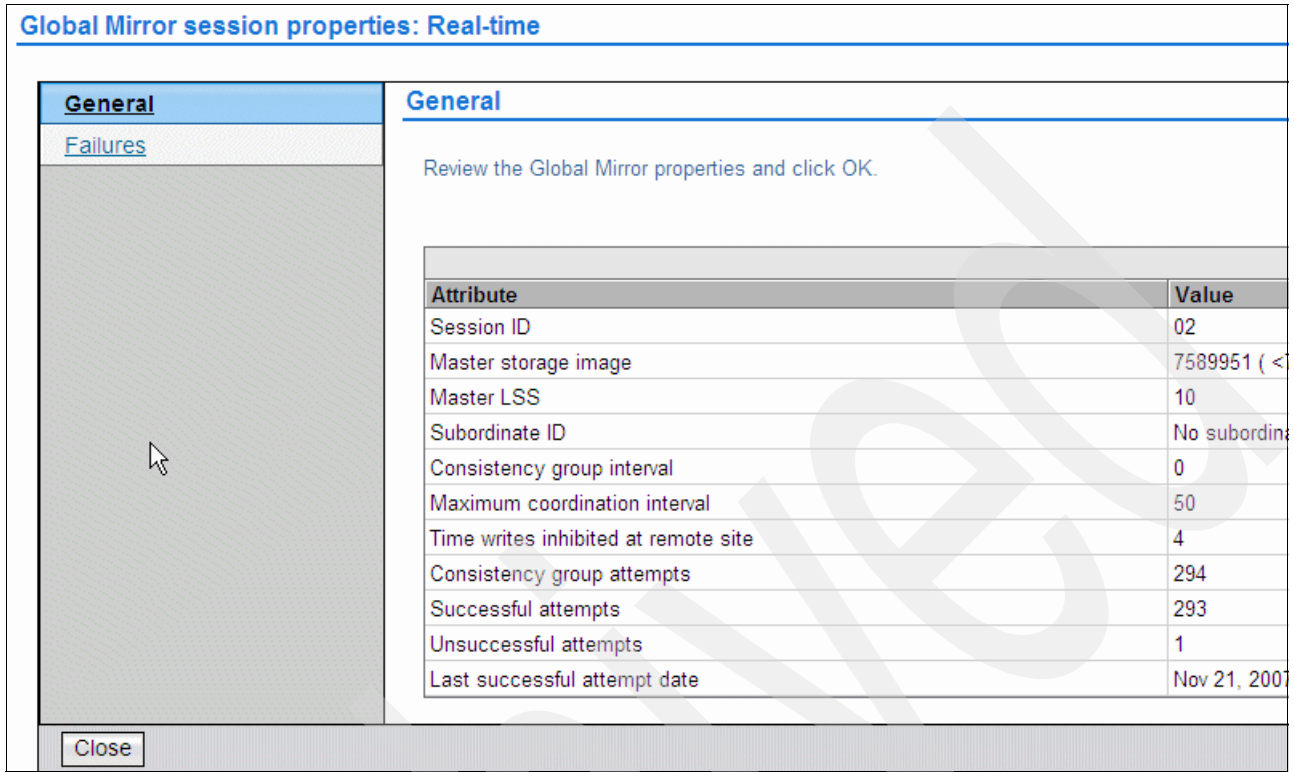


Figure 13-48 General properties of the Global Mirror

The failures properties display, as shown in Figure 13-49. Select the type of failure that you want to see. In our example, we select **Most recent failure**. Click **Close** to return to the previous panel.

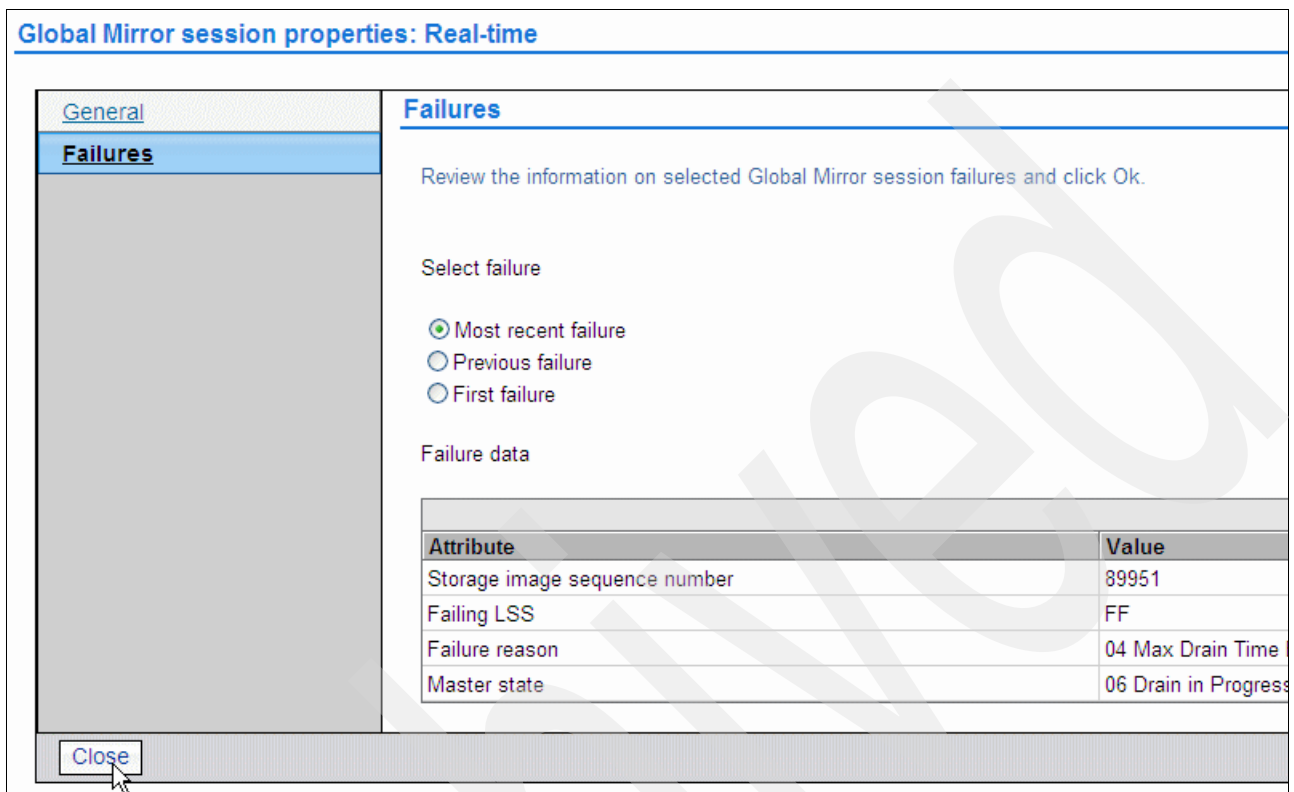


Figure 13-49 Viewing the most recent failure

Archived

## Performance considerations

In general, an external storage system has no knowledge of any file structures, library structures, and other host system specific characteristics. It has been designed only to handle the disks and their I/O. Thus, anything that you can see at the DS level is at a track level. When doing initial copies of Peer-to-Peer Remote Copy (PPRC) pairs, the system copies the entire disk, regardless of whether there is any data from the System i perspective.

When setting up the external storage, expert knowledge about the storage configuration is required to optimize its performance. You must have a clear view of what you want to do with the external storage both now and in the future. Based on experience, good planning of the final configuration of the system offers immediate pay-off during the configuration stages and for the entire configuration in the future.

When using IBM System Storage solutions in combination with the System i platform, four key areas require attention regarding performance:

- ▶ Configuration of the DS system
- ▶ Connectivity between the DS systems and System i environment
- ▶ Connectivity between the DS systems in case of Metro Mirror and Global Mirror solutions (physical and logical)
- ▶ I/O performance of the System i platform on the DS system

In this chapter, we look at each of these areas in detail, because the final solution depends on them to be tuned for maximum performance. We also focus on additional issues in relation to the various Copy Services solutions.

## 14.1 Configuration of the DS system

Although the System i platform is classified under Open Systems from a DS perspective, some specific considerations are related to the way that the System i platform handles I/O and logical unit number (LUN) sizes and properties.

The System i platform has a single-level storage architecture. This means that physical writes are spread across the available disks within the auxiliary storage pool (ASP) where the object is located. By doing this, you use as many disk resources (especially disk arms) as possible. In order to obtain the same effect on the DS system, you must follow these guidelines:

- ▶ Use separate ranks for System i disks.
- ▶ Try to get single-sized LUNs as per ASP or independent ASP (IASP), or a maximum of two adjacent sizes (for example, 17.5 and 35.2) with the majority of LUNs being of the larger size.
- ▶ Create the individual LUNs and extent pools on single ranks and not across ranks.
- ▶ Make sure that you balance the ranks and LUNs across both processors, associated with rankgroup 0 and 1, of the DS system, making maximum use of the full redundant setup of the DS system.
- ▶ Optimize the use of logical subsystems (LSSs; the first two digits of the LUN number).
- ▶ Place source and target LUNs on different ranks within the same DS processor (same rankgroup) for FlashCopy.

These guidelines might not make maximum use of the overall disk space, but they help to obtain maximum performance. Refer to 3.2.6, “Planning for capacity” on page 67 for detailed capacity planning considerations.

## 14.2 Connectivity between the DS systems and System i environment

The connectivity between the DS system and the System i environment is not directly related to the performance of Copy Services. However, problems or incorrect sizing of the connection can have a severe impact and must be considered when looking at the total solution.



## 14.2.1 Physical connections

To connect the two systems, we use optical connections, called *Fibre Channel (FC) adapters*, on the System i model and a host bus adapter (HBA) on the DS system. For various reasons, such as a limited number of HBAs on the DS system, you can place a storage area network (SAN) switch (Figure 14-1) between the systems to facilitate, manage, and share the connections.




Figure 14-1 SAN switch to facilitate, manage, and share connections between the two systems

A SAN switch can route the entering signal to the correct destination port, behind which is the destination worldwide port name (WWPN) but at the cost of some overhead. Both from performance perspective to prevent link utilization problems and because these installations hardly ever change after installation, it is best to create static paths from the source to the target, which is known as *zoning*. Refer to 3.2.5, “Planning for SAN connectivity” on page 67 for planning your SAN switch zoning. After you create the zoning, most SAN switches require this zoning definition to be *activated* before it becomes effective.

The HBA ports on the DS systems can be *zoned* similar to a SAN switch to restrict host system FC adapters’ login to preselected DS storage HBA ports only. We do not recommend to restrict the host logins to certain DS ports when creating the host connection definitions on the DS system because defining the zoning at the SAN switch proves to be much more flexible.

## 14.3 Connectivity between the DS systems

When using PPRC for your Metro Mirror or Global Mirror solution, you use two DS systems. You can connect these systems using any of the following methods:

- ▶ Dedicated fibre connections
- ▶ Shared fibre connections
- ▶ Multi-protocol routers, which transform the optical signal to TCP/IP-packets and back using a LAN or WAN connection in between the routers

Apart from this physical connection, we look briefly at the logical connection, which is how the DS system is handling the inter-DS I/Os.

### 14.3.1 Physical connections

The combination of the type of connectivity and the bandwidth used for the PPRC data flow is key to the performance. Three components to the connection must be sized:

- ▶ The effect of the backup DS model on the production site using Disk Magic
- ▶ The bandwidth between the sites
- ▶ The effect of the production site DS model on the backup site using Disk Magic

It is important to try and reduce the amount of data transiting between the DS systems and the amount of bandwidth that is allocated to this flow, dedicated or shared and with or without quality of service (QoS). The amount of data transiting between the DS systems is especially important for Metro Mirror.

Metro Mirror is based on synchronous updates (Figure 14-2). The write command as initiated from a System i (1 + 4) is not done until the remote DS system has confirmed the write (2 + 3) to the local DS system.

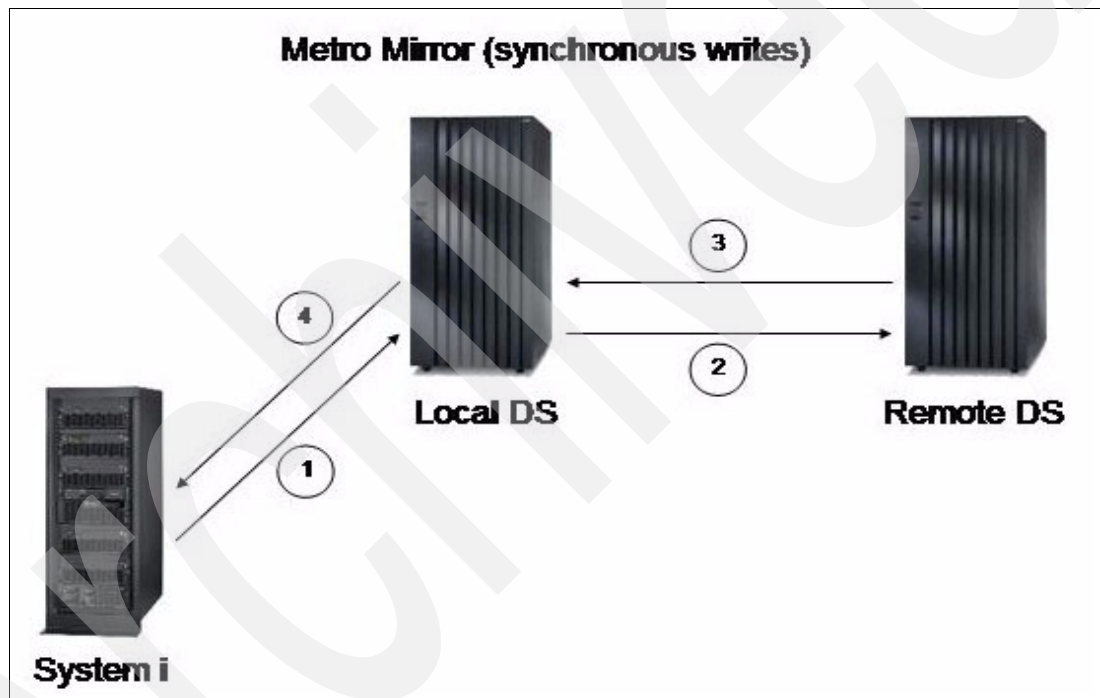


Figure 14-2 Metro Mirror: Synchronous updates

Global Mirror is asynchronous. The write on the local DS system (1 + 4) is the only write on which I/O responsiveness depends. The write to the remote DS-system (2 + 3) has no bearing on this.

Which solution is taken depends on the distance between the machines, the availability of the lines, and the costs involved. The most optimal solution of the connection methods described has a dedicated fibre connection between the systems. However, this solution is expensive and might even be unobtainable. The next best option is guaranteed bandwidth on either a fibre connection or WAN connection.

In order to avoid unpleasant surprises, we highly recommend that you do an accurate study of the amount of I/Os that will pass on from one DS system to the other. Given the level of initial investment for the external storage and the running costs involving data communications, it might be worth the effort to do a good benchmark to see the bandwidth that is needed.

As a rule of thumb, the I/O reports from the System i environment can be taken to see how much traffic will go across the connection between the two DS systems. This is not a one-to-one relation because it doesn't account for write efficiency by the cache but it is close enough for a first estimate.

### 14.3.2 Logical connections

The inter-DS I/Os are handled by the LSS. For each connection from one local LSS to a remote LSS, you must establish a PPRC path to the remote DS system and a path back. However, there are limitations:

- ▶ A DS system HBA port can handle both host I/O and PPRC traffic but for performance reasons dedicated HBA ports for PPRC are recommended.
- ▶ A primary LSS can have paths up to four secondary LSS.
- ▶ The maximum number of PPRC paths per LSS is eight.
- ▶ A physical PPRC link between the primary and secondary DS system supports up to 256 logical PPRC paths
- ▶ For each Metro Mirror path, there is a bandwidth limitation of approximately three LUNs in parallel, as you can see during the initial copy by simply looking at the number of Out-Of-Sync Tracks; the other values are the default for the setup.

Given these limitations, you must try to strike a balance between performance and what is possible. You must also keep in mind the possible evolutions of the systems to avoid creating bottlenecks in the future.

### 14.3.3 Using independent ASPs with Metro Mirror

One of the most rewarding methods to reduce the use of replication link bandwidth is the use of IASPs. The System i architecture attributes a temporary storage or memory space to each job (QTEMP). Applications use this space and other space for temporary indexes and files. IASPs separate the essential application I/O from the workfile type of I/O that remains in SYSBASE.

Figure 14-3 shows how the number of writes on the IASP remains almost flat, where the number of writes to SYSBASE creates continuous and considerable overhead, especially when the interactive users are working in the system.




Figure 14-3 IASP remains flat, where SYSBASE creates continuous and considerable overhead

Because the writes on SYSBASE are mainly for temporary use, they are of no importance when switching over from one DS system to the other one. When bringing up the system on the remote site after a crash, the System i platform first tries to repair the likely damaged object to determine whether it is only a temporary file and is of no use because the related user or job session is no longer available. Therefore, all the effort that has gone into replicating the files from the local to the remote site is of no use when switching over.

The use of IASPs has other major advantages. A switchover of a crashed system under a full system PPRC does not react any differently than trying to perform an IPL on the failed system. When performing an IPL after an abnormal end, the time that this IPL takes is unpredictable. Performing it on the remote site is not going to make it any faster. Therefore, a full system PPRC solution (Metro Mirror and Global Mirror alike) is the disaster recovery solution. Migrating to an IASP allows you to have both local and remote System i partitions running and switch over the IASP only.

To share this IASP as a resource between partitions, you must create a cluster between the two partitions (otherwise known as *nodes*) and make sure that all IASP-related information that is in SYSBASE (user profiles, job descriptions, and so on) is synchronized between these two nodes.

Figure 14-4 illustrates the IASP connectivity schema.




Figure 14-4 IASP connections in Metro Mirror

For further information about IASPs and clustering, refer to the following resources:

- ▶ *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs*, SG24-6802
- ▶ i5/OS Information Center, section System Management Clustering at:  
<http://publib.boulder.ibm.com/infocenter/series/v5r4/index.jsp>

There is one object that needs special attention. When creating a network server storage space (NWSSTG) in an IASP, the information regarding this NWSSTG has to be transferred separately to the second node. This is a one-time action after the creation of the NWSSTG.

As you can see in Figure 14-5, the directory NWSSTG is in \root\QFPNWSSTG. This directory contains the information that allows the System i platform to connect the NWSSTG correctly to the network server. The easiest way to copy this information is to save it to a save file and then restore it from this save file on the target system. Failing to do this prevents you from attaching the NWSSTG to the network server.

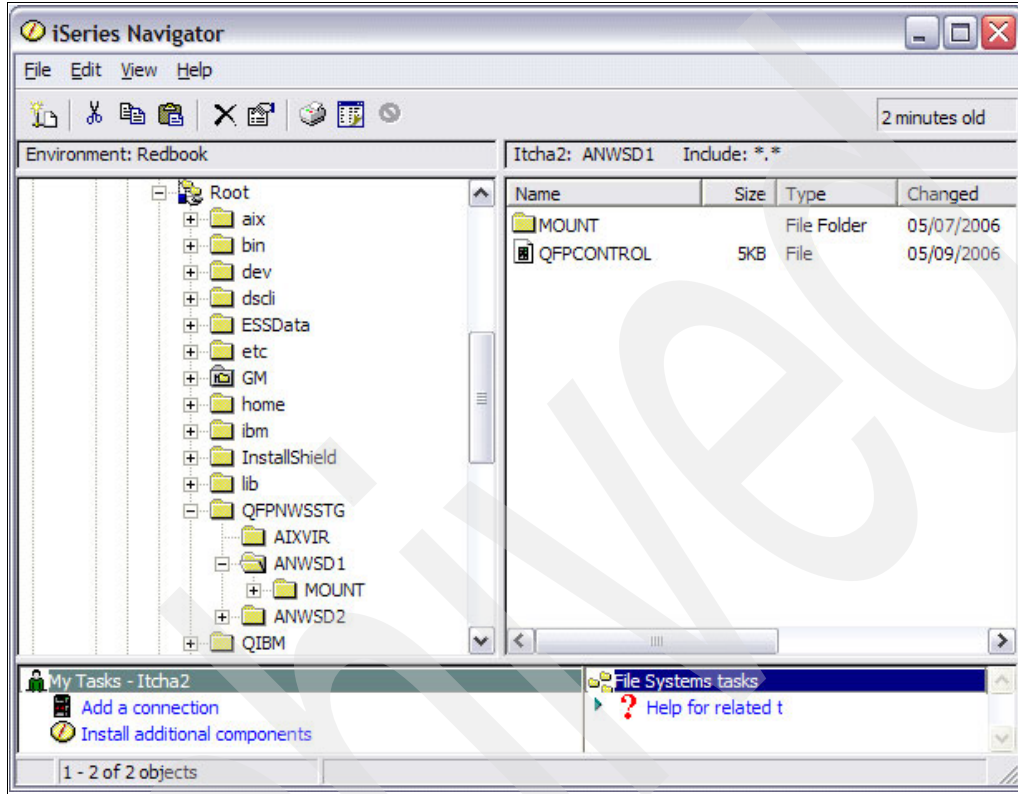


Figure 14-5 Contents of QFPNWSSTG



## FlashCopy usage considerations

In this chapter, we describe the functionality and usage of the new i5/OS V6R1 quiesce for Copy Services function that provides a new level of i5/OS availability for usage with FlashCopy.

We also discuss the special considerations that are required for using Backup Recovery and Media Services (BRMS) with IBM System Storage FlashCopy to make sure changes to the BRMS library by the backup system are rolled back properly to the original production system.

## 15.1 Using i5/OS quiesce for Copy Services

Prior to V6R1, to ensure all updates from main storage (memory) are flushed to disk storage, an i5/OS IASP had to be varied off or the system or LPAR had to be powered down before taking an IASP image or full system image using FlashCopy, for example for daily backups. Now the new i5/OS V6R1 function quiesce for Copy Services provides a higher level of availability for i5/OS for usage with IBM System Storage FlashCopy Copy Services.

**Note:** The i5/OS V6R1 quiesce for Copy Services function eliminates the IASP vary-off or power-down requirements before taking a FlashCopy by writing as much modified data from System i main memory to disk as possible allowing for nondisruptive use of FlashCopy with i5/OS.

When invoking the quiesce for Copy Services function, the flush of modified main memory content is internally performed within a two-phase flush to make this function very efficient with limiting the time required for suspending the I/O while paging out (destaging) the data modified since the first flush (see Figure 15-1).

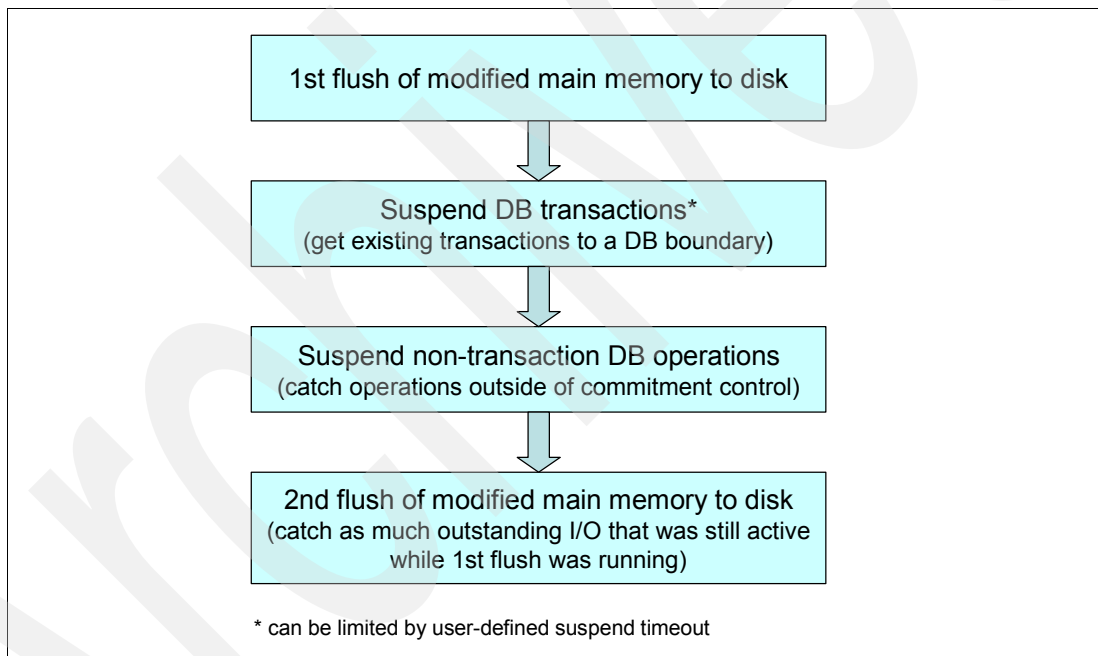


Figure 15-1 Quiesce for Copy Services two-phase flush process flow

Quiesce for Copy Services tries to flush as much modified data to disk as possible and then pauses (suspends) future *database* transactions and operations. Non-database write operations, like changing a message file, creating a library or IFS streamfile changes, are allowed to continue. Only database transactions and operations are suspended.

**Important:** For the IASP or system image created by FlashCopy after the quiesce for Copy Services completed, it is always an *abnormal vary-on* or *abnormal IPL*.



A new i5/OS V6R1 CL command CHGASPACT (change ASP activity) is used to invoke the quiesce for Copy Services function. Figure 15-2 shows the CL command user interface with with selected parameters to suspend \*SYSBAS I/O activity and end, that is *abort*, the suspend operation if the specified suspend timeout of 30 seconds would be succeeded.

```

Change ASP Activity (CHGASPACT)

Type choices, press Enter.

ASP device . . . . . *SYSBAS      Name, *SYSBAS
Option . . . . . *SUSPEND      *SUSPEND, *RESUME, *FRCWRT
Suspend timeout . . . . . 30      Number
Suspend timeout action . . . . . *END      *CONT, *END

F9=All parameters  F11=Keywords  F14=Command string  F24=More keys      Bottom
Parameter ASPDEV required.                                     +

```

Figure 15-2 CHGASPACT CL command user interface

The CHGASPACT parameters with their keywords denoted in brackets have the following meaning:

- ▶ **ASP device** (ASPDEV): Mandatory parameter to specify either the IASP device description name or \*SYSBAS comprising the system ASP 1 and any existing userASPs 2 to 31.
- ▶ **Option** (OPTION): Mandatory parameter to specify either to suspend, resume or force writes to the selected ASP device. The force writes option only triggers the first flush operation of the four phase quiesce function shown in Figure 15-1 without any suspend actions.

**Note:** The resume option should be run after using the suspend option. Otherwise, it takes 20 minutes until an automatic resume of a suspended ASP device is started.

- ▶ **Suspend timeout** (SSPTIMO): Mandatory parameter to specify the suspend timeout in seconds. If a parameter value of “0” is specified it is translated to 300s. The suspend timeout value is used as follows: *SSPTIMO* - 10s is the timeout used for suspending the DB transactions and the remaining fixed 10s are allotted for suspending the non-transaction operations. If SSPTIMO is specified for less than 11s, 11s are used. If there is a suspend timeout it is likely in the DB transaction part rather than in the suspend of the non-transaction operations.

- ▶ **Suspend timeout action (SSPTIMOACN):** Optional parameter to define the behavior of the quiesce function if the suspend timeout popped. The default value \*CONT means the timeout is ignored and the suspend operation continues. A message in the joblog indicates a succeeded suspend timeout.

**Note:** Use the \*END option if you do not accept taking a FlashCopy image from an unsuccessful DB transaction suspend. The \*END option will automatically invoke a resume of the ASP after a timeout.

The CHGASPACT command uses the following messages:

- ▶ CPCB717 Access to ASP &1 is suspended.
- ▶ CPCB718 Access to ASP &1 successfully resumed.
- ▶ CPDB717 SSPTIMO and SSPTIMOACN are only allowed with OPTION(\*SUSPEND)
- ▶ CPDB718 Suspend Timeout (SSPTIMO) is required with OPTION(\*SUSPEND)
- ▶ CPF8717 Suspend Access timed out and did not complete successfully.

Figure 15-3 shows a successful completion of the CHGASPACT suspend operation indicated by i5/OS message CPCB717.

```

Additional Message Information
Message ID . . . . . : CPCB717      Severity . . . . . : 00
Message type . . . . . : Completion
Date sent . . . . . : 11/03/07      Time sent . . . . . : 13:35:06

Message . . . . . : Access to ASP *SYSBAS is suspended.
Cause . . . . . : Access to ASP *SYSBAS is suspended. New transactions will
not be allowed to start until access is resumed. Transaction quiescing used
0 of 30 available seconds. The reason code is 0.
Possible reason codes are:
0 -- Transaction quiescing completed successfully within the specified
time.
-1 -- Transaction quiescing did not complete in the specified time. A
larger timeout value is required to allow existing transactions to complete.
Recovery . . . . . : None.

Bottom

Press Enter to continue.

F3=Exit  F6=Print  F9=Display message details  F12=Cancel
F21=Select assistance level

```

Figure 15-3 Quiesce for Copy Services successful suspend completion message

Figure 15-4 shows a successful completion of the quiesce for Copy Services resume operation, indicated through i5/OS message CPCB717 for the example of resuming a suspended \*SYSBAS using CHGASPACT ASPDEV(\*SYSBAS) OPTION(\*RESUME).

```

Additional Message Information

Message ID . . . . . : CPCB717      Severity . . . . . : 00
Message type . . . . . : Completion  Time sent . . . . . : 13:35:06
Date sent . . . . . : 11/03/07

Message . . . . . : Access to ASP *SYSBAS is suspended.
Cause . . . . . : Access to ASP *SYSBAS is suspended. New transactions will
not be allowed to start until access is resumed. Transaction quiescing used
0 of 1 available seconds. The reason code is 0.
Possible reason codes are:
0 -- Transaction quiescing completed successfully within the specified
time.
-1 -- Transaction quiescing did not complete in the specified time. A
larger timeout value is required to allow existing transactions to complete.
Recovery . . . . . : None.

Bottom

Press Enter to continue.

F3=Exit  F6=Print  F9=Display message details  F12=Cancel
F21=Select assistance level

```

Figure 15-4 Quiesce for Copy Services successful resume completion message

If a suspend transaction times out, message CPEB717 is posted and a spoolfile named QPCMTCTL is created, under the job that ran the suspend, that identifies the transactions that were unable to be suspended (see Figure 15-5). This spool-file allows to determine which files/transactions were outstanding and get a better idea of how long it will take to quiesce and whether those particular files are important for the FlashCopy image.

Record Level Status									
----- Changes ----- Lock									
Commit Cycle	File	Library	Member	Commit	Rollback	Pending	Level	Status	Journal
F	QUIESCE	F	F	0	0	1	*ALL	OPEN	QUIESCE/J 6

Figure 15-5 QPCMTCTL spool file example

For further information about the new i5/OS V6R1 CHGASPACT CL command and its QYASPCHGAA API functions allowing you to code up your own functionality refer to the i5/OS V6R1 Information Center, at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp>

The flowchart in Figure 15-6 shows the recommended process for using quiesce for Copy Services with FlashCopy. We recommend using an initial suspend timeout (SSPTIMO) value of 30s and reviewing the completion message of the CHGASPACT suspend operation which tells how many seconds of the provided timeout were used for quiescing the transactions.

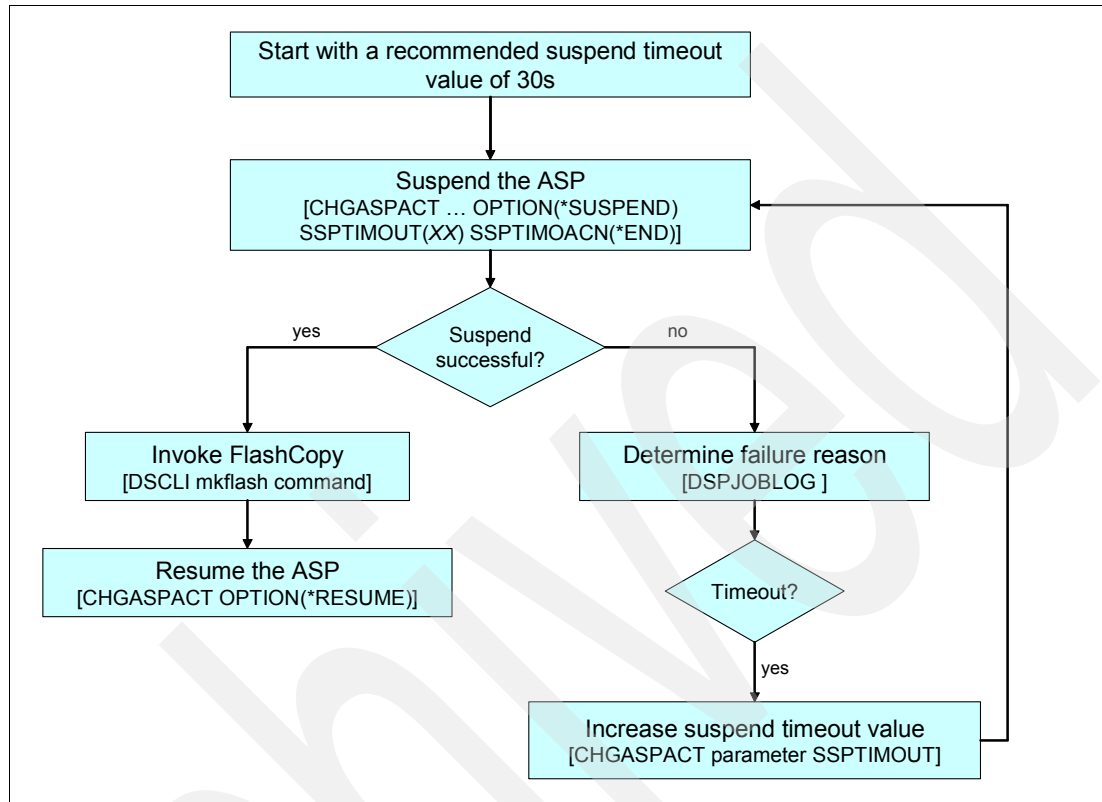


Figure 15-6 Example for using the quiesce for Copy Services function

If the suspend operation completes successfully (reason code 0) all database transactions have been successfully quiesced and a FlashCopy can be initiated which would be up to date to the last database transaction requiring no database recovery at a vary-on or IPL from the FlashCopy image. If the suspend operation times out, not all database transactions could be quiesced and the timeout value needs to be increased.

Note our usage of the CHGASPACT command with the non-default \*END option for the suspend timeout action (SSPTIMOACN) because we assume many customers will not accept a FlashCopy from an unsuccessful quiesce of their database operations. However, for those user-interactive scenarios that have no limit for a database transaction, specifying a timeout value for the database suspend makes no sense so that the default \*CONT option should be used for the suspend timeout action.

## 15.2 Using BRMS and FlashCopy

Backup Recovery and Media Services (BRMS) is the IBM strategic solution for performing backups and recovering System i5 environments. BRMS has a wealth of features, including the ability to work in a network with other systems to maintain a common inventory of tape volumes.

FlashCopy creates a copy of the source system onto a second set of disk drives, which are then attached and used by another system or logical partition (LPAR). The BRMS implementation of FlashCopy provides a way to perform a backup on a system that has been copied by FlashCopy and a BRMS history appears, as the backup is performed, on the production system.

In this chapter, we explore how you can use BRMS to perform backups and recoveries from a secondary LPAR. This can also be a separate stand-alone system. However, using the dynamic resource movement introduced in V5R1 and later of OS/400, the LPAR solution is the best way to use FlashCopy when attached to a System i platform.

**Attention:** If you plan to use online Domino backup, you must do the backup on the production system. You must save all journal receivers on the production system to avoid journal receiver conflict and to enable point-in-time recovery.

## 15.3 BRMS architecture

BRMS stores backup history and media information in a library called QUSRBRM. The files in this library define both the setup of the BRMS environment and the dynamic information gathered as a result of doing BRMS operations such as saves and restore tasks. This information is critical to the recovery of the system. When using FlashCopy to create a full system image, QUSRBRM is also copied from the production system to the backup system.

Figure 15-7 shows two partitions:

- ▶ A production partition for normal day-to-day processing
- ▶ A backup partition for taking offline backups

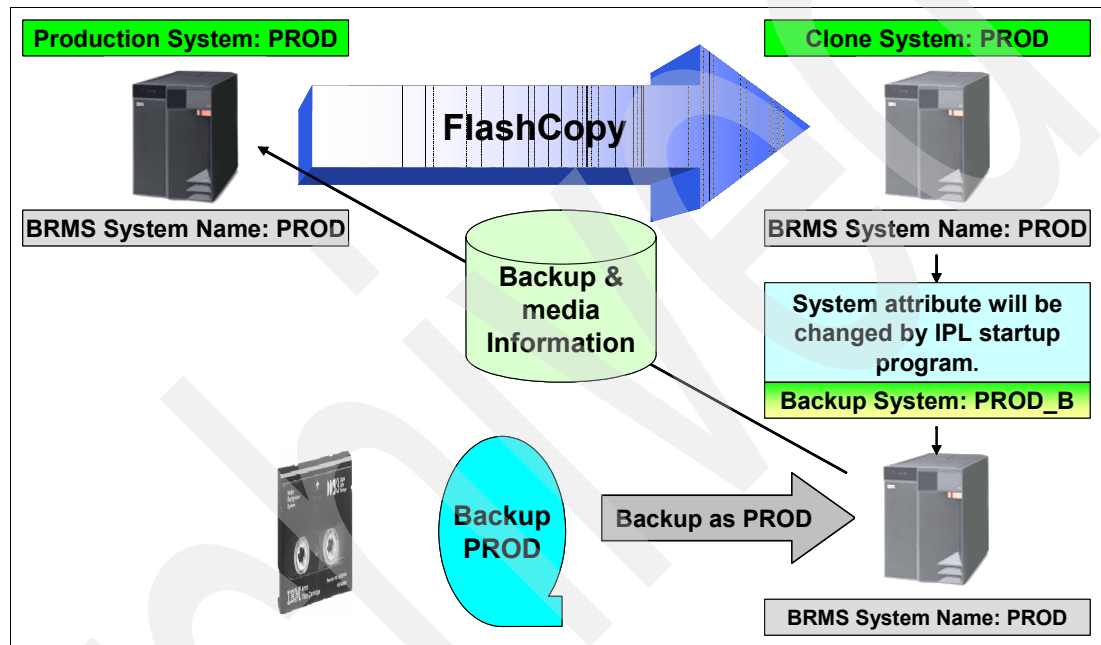


Figure 15-7 BRMS and FlashCopy

## 15.4 Enabling BRMS to use FlashCopy

The BRMS FlashCopy function requires the BRMS Network Feature product 5722-BR1. In order to use BRMS to perform a backup of the copy system, FlashCopy function must be enabled on the production system. After you enable the BRMS FlashCopy function, all backups that are performed on the backup system look like they were performed on the production system.

To enable the FlashCopy function for BRMS, enter the following command:

- ▶ For BRMS V6R1 and later:

```
WRKPCYBRM *SYS
```

Then, choose **1. Display or Change system policy** and select to enable FlashCopy using:

```
Enable FlashCopy . . . . . *YES
```

- ▶ Prior to BRMS V6R1:

```
QSYS/CALL PGM(QBRM/Q1AOLD) PARM('FLASHSYS ' '*YES')
```

**Note:** For all Q1AOLD program call commands in this section, you need to use all uppercase letters for all parameters.

By using this interface, BRMS can perform a backup of the backup system as though it were the production system. The backup history looks like a backup was performed on the production system.

### 15.4.1 Preliminary notification of FlashCopy mode

You must notify BRMS that the system's data is being copied using FlashCopy and the backup is performed on the backup system. This step is required prior to performing the FlashCopy function.

Enter the following command to set the BRMS system state to FlashCopy mode:

- ▶ For BRMS V6R1 and later:  
QSYS/INZBRM OPTION(\*FLASHCOPY) STATE(\*STRPRC)
- ▶ Prior to BRMS V6R1:  
QSYS/CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '\*BEGIN')

Enter the following command to display the BRMS FlashCopy state:

- ▶ For BRMS V6R1 and later:  
WRKPCYBRM \*SYS  
Then, choose **4. Change network group** and look for the FlashCopy state information.
- ▶ Prior to BRMS V6R1:  
QSYS/CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '\*DISPLAY')

When the system is in FlashCopy mode, the BRMS synchronization job does not run on the production system.

**Important:** Do not perform BRMS activity on the production system until all post FlashCopy steps are complete.

Any updates to the BRMS database on the production system using any BRMS activity, such as save, restore, BRMS maintenance, and so on, will be lost. When the system is in FlashCopy state, all incoming BRMS communication from the BRMS networked system is blocked. BRMS backup information about the current system might be outdated when a backup is performed on the backup system.

You should verify that this production system owns enough media for the backup in order to complete a successful backup. If a copy system can perform communication in a restricted state by using specified TCP/IP interface, then BRMS can use media owned by another system in the BRMS network.

### 15.4.2 Pre-backup step on backup system

To prevent a system name conflict in the network, in many situations, a default local location name and system name value in the Display Network Attribute (DSPNETA) command cannot be the same on the production system and backup system. To resolve a name conflict, a user might need to change these attributes through an IPL startup program on the backup system.

Because the production system is enabled for FlashCopy, any backup performed on the backup system uses the Display Network Attribute (DSPNETA) of the production system at the time of enabling the BRMS FlashCopy function.

### 15.4.3 Setting the BRMS system state to backup system

The status of the backup system is also in FlashCopy mode after its IPL. This prevents a BRMS synchronization job from sending an update to other systems in the BRMS network. From a BRMS perspective, at this time, the backup system is the production system, and all updates of the BRMS information should be sent to all systems in the network. In order to allow an update to another system, the state should be changed to *backup FlashCopy system*.

Enter the following command on the backup system to set the BRMS system state to backup FlashCopy system:

- ▶ For BRMS V6R1 and later:  
`QSYS/INZBRM OPTION(*FLASHCOPY) STATE(*STRBKU)`
- ▶ Prior to BRMS V6R1:  
`QSYS/CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '*BACKUPSYS')`

Enter the following command to display the BRMS system state:

- ▶ For BRMS V6R1 and later:  
`WRKPCYBRM *SYS`  
Then, choose **4. Change network group** and look for the FlashCopy state information.
- ▶ Prior to BRMS V6R1:  
`QSYS/CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '*DISPLAY')`

### 15.4.4 Setting the backup system to restricted state TCP/IP

When you are running SAVSYS backup procedures, the operating system must be in a restricted state. In a shared media inventory, if the current system does not have any volumes available, then BRMS needs to communicate with the remote systems for volume selection. In order to do this while in a restricted state, BRMS needs to start the TCP/IP interfaces that will be used to communicate with the remote systems. You need to specify those TCP/IP interfaces to BRMS. A restricted state TCP/IP interface specified on the production system might not be the same for the copy system.

On i5/OS V5R3 or later systems, enter the following command to specify the TCP/IP interfaces that BRMS should use during the restricted state:

```
QSYS/CALL QBRM/Q1AOLD PARM('TCPIIFC' '*ADD' 'interface')
```



Alternatively also the BRMS GUI from iSeries Navigator or Web support can be used to modify the TCP/IP restricted state interfaces by right-clicking **Backup, Recovery and Media Services**, selecting **Global Policy Properties**, choosing the **Network** tab from the dialog window, and selecting **Manage Interfaces to Start** as shown in Figure 15-8.

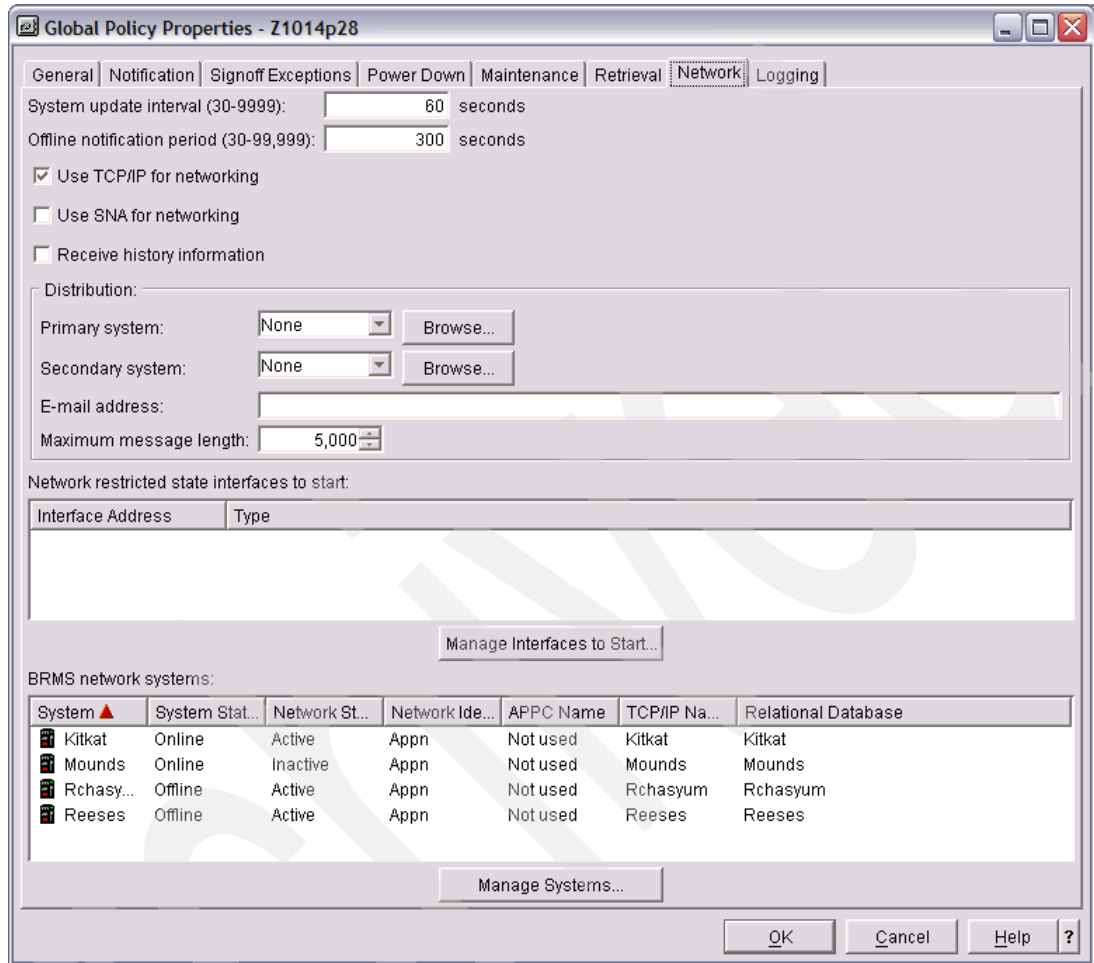


Figure 15-8 BRMS GUI - Global Policy

For more information about the restricted state TCP/IP interface, refer to:

<http://www-03.ibm.com/servers/eserver/iseries/service/brms/brmstcpip.html>

### 15.4.5 Changing hardware resource names on the backup system

It is highly unlikely that the hardware resource names associated with the tape drives in the production partition will match those on the backup partition. In this case, either change the device descriptions on the backup system after you have done the FlashCopy or create a CL program to perform this task automatically.

## 15.5 Performing the backup from the backup system

Because the BRMS system name is set on the production system and is stored in the BRMS database, after a FlashCopy, the BRMS system name on the backup system is the same as the BRMS system name on the production system.

Simply follow the backup procedure on the backup system as usual as you would on your production system.

## 15.6 Post FlashCopy steps

The following sections describe the required post FlashCopy steps to ensure that the BRMS database on the production system is updated with the BRMS backup information created by the performed backup from the backup system.

### 15.6.1 Indicating that the BRMS backup activity is complete

During the post FlashCopy step, do not perform BRMS activity on the production system and on the backup system. Enter the following command on the backup system to set the BRMS system state to end backup mode:

- ▶ For BRMS V6R1 and later:  
`QSYS/INZBRM OPTION(*FLASHCOPY) STATE(*ENDBKU)`
- ▶ Prior to BRMS V6R1:  
`QSYS/CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '*ENDBACKUP')`

This command prevents any incoming communication and feature BRMS synchronization updates to other systems in the BRMS network from the backup system. The Q1ABRMNET subsystem is ended during this step.

Do not use the backup system for any BRMS activity because all BRMS backup history information is sent to production, and all BRMS controls are sent back.

### 15.6.2 Sending QUSRBRM to the production system

You must save the QUSRBRM library to allow the BRMS management information to be transferred to the production partition. To save the QUSRBRM library, enter the following command on the copy system:

```
SAVLIBBRM LIB(QUSRBRM) DEV(tape-media-library-device-name) MEDPCY(media-policy)
OBJDTL(*OBJ) SAVTYPE(*FULL) SEQNBR(1) ENDOPT(*REWIND)
```

The final step is to restore QUSRBRM, which you saved from the backup system. This provides an accurate picture of the BRMS environment on the production partition, which reflects the backups that were just performed on the backup system. To restore QUSRBRM, use the media that was used to perform the backup of the QUSRBRM library and enter the following command on the production system:

```
QSYS/RSTLIB SAVLIB(QUSRBRM) DEV(tape-media-library-device-name)
VOL(volume-identifier) SEQNBR(1) OMITOBJ((QUSRBRM/*ALL *JRN)) ALWOBJDIF(*FILELVL
*AUTL *OWNER *AUTL) MBROPT(*ALL)
```

### 15.6.3 Indicating that the FlashCopy function is complete on the production system

Enter the following command on the production system to indicate that the FlashCopy function is complete:

- ▶ For BRMS V6R1 and later:

```
QSYS/INZBRM OPTION(*FLASHCOPY) STATE(*ENDPRC)
```

- ▶ Prior to BRMS V6R1:

```
CALL QBRM/Q1AOLD PARM('FLSSYSSTS' '*END')
```

This command starts the Q1ABRMNET subsystem if the system is not in a restricted state. It also starts all BRMS synchronization jobs.

## 15.7 Daily maintenance in BRMS

At this point, you should run maintenance on the production system. The BRMS maintenance function regularly and automatically cleans and updates media records. Regular removal of expired records from media and media content information files allows you to make more efficient use of your media.

The center of the BRMS maintenance function is the Start Maintenance for BRM (STRMNTBRM) command. This command processes the daily maintenance requirements that keep your system running efficiently. BRMS detects and records new and deleted libraries. By default, deleted libraries are not included in the "Recovering Your Entire System Report". This is important if you are saving libraries on auxiliary storage pool devices, i.e. independent ASPs. The auxiliary storage pool devices must be available when you run maintenance. Otherwise, BRMS is unable to locate the libraries and considers the libraries on unavailable auxiliary storage pool devices as having been deleted from the system.

For additional information about how to use the daily BRMS maintenance job, refer to IBM System - iSeries Backup, Recovery, and Media Services for iSeries, SC41-5345, which is available at:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r4/topic/books/sc415345.pdf>

## 15.8 Printing recovery reports

BRMS can generate a series of comprehensive recovery reports for use in recovering your entire system. If BRMS is offline due to system failure or other disaster, the recovery reports provide instructions on how to perform the first few steps manually. For example, the recovery reports tell you where to locate the volumes that are necessary to restore your system. In addition, they identify the manual steps that you must take to install the Licensed Internal Code and perform a restore of the operating system and the BRMS product.

After you complete the manual steps, you can use BRMS to assist in recovering the remainder of your system. Perform the following steps to print the recovery reports that you need to recover your system:

1. On any command line, enter the STRRCYBRM command. Then press F4.
2. On the Start Recovery using BRM display (Figure 15-9), press Enter.




Figure 15-9 BRMS - Start Recovery using BRM display

- As shown in Figure 15-10, in the Option field, type \*SYSTEM, and in the Action field, type \*REPORT. Press Enter.

```

Start Recovery using BRM (STRRCYBRM)

Type choices, press Enter.

Option . . . . . *SYSTEM      *SYSTEM, *ALLDLO, *ALLUSR...
Action . . . . . *REPORT      *REPORT, *RESTORE
Time period for recovery:
  Start time and date:
    Beginning time . . . . . *AVAIL      Time, *AVAIL
    Beginning date . . . . . *BEGIN      Date, *CURRENT, *BEGIN
  End time and date:
    Ending time . . . . . *AVAIL      Time, *AVAIL
    Ending date . . . . . *END        Date, *CURRENT, *END
  Use save files . . . . . *YES        *YES, *NO
  Use TSM . . . . . *YES          *YES, *NO
ASP device:
  From system . . . . . *LCL
  Auxiliary storage pool . . . . *ALL      Name, *ALL
  Objects . . . . . *ALL          *ALL, *LIB, *LNK
                                + for more values
More...

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

```

Figure 15-10 Start Recovery with BRM: Parameters view

4. The spooled files are generated, as shown in Figure 15-11, from which you can print the following reports:
  - **QP1ARCY**: Recovering Your Entire System (features the actual recovery steps)
  - **QP1A2RCY**: Recovery Volume Summary Report (tells you where to find the necessary volumes)
  - **QP1AASP**: Display ASP Information

Enter the Work with Spooled Files (WRKSPLF) command to print the reports.

Work with All Spooled Files								
Type options, press Enter.								
1=Send 2=Change 3=Hold 4=Delete 5=Display 6=Release 7=Messages								
8=Attributes 9=Work with printing status								
Opt	File	User	Device or Queue	User Data	Sts	Total Pages	Cur Page	Copy
	QP1ARCY	REDBOOK	QPRINT	STRRCYBRM	RDY	11		1
	QP1A2RCY	REDBOOK	QPRINT	STRRCYBRM	RDY	1		1
	QP1AASP	REDBOOK	QPRINT	STRRCYBRM	RDY	1		1

Bottom

Parameters for options 1, 2, 3 or command  
 ===>  
 F3=Exit F10=View 4 F11=View 2 F12=Cancel F22=Printers F24=More keys

Figure 15-11 Working with BRMS spooled files

To use BRMS to perform a recovery, you must have a copy of these reports available. Each time you complete a backup, print a new series of recovery reports. Be sure to keep a copy of these reports with each set of tapes at all locations where media is stored.

## 15.9 Recovering your entire system

BRMS recovery reports guide you, in a step-by-step manner, through the process of recovering your entire system. You can also use these reports to guide you through the recovery of selected aspects of your system. In the case of a total system failure, the reports guide you through the first manual steps of the recovery process. These initial, manual steps include recovery of the Licensed Internal Code and the operating system. For information about how to recover Licensed Internal Code and the operating system from failures refer to *IBM Systems - iSeries Backup and Recovery*, SC41-5304, which is available at:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r4/topic/books/sc415304.pdf>

After you complete the manual steps, you can use BRMS and the reports to help you restore the rest of your system. There are a variety of ways in which you can recover data. For example, you can restore information by control group, object, library, and document library

objects (DLOs). For more information about recovering your entire system, see Chapter 4, “Recovering Your Entire System” in *Backup, Recovery, and Media Services for iSeries*, SC41-5345.

**Important:** Because the backup on the backup system is done by BRMS as though it were for the production system, you do not need to update the system name in BRMS Media Information when you recover the production system.

Archived

Archived



## Troubleshooting

When the System i platform is used as the host server in an IBM System Storage environment, it can indicate performance, network, and hardware issues that are experienced on the IBM System Storage DS6000 or DS8000 system. You can use several tools to generate reports to help determine the cause of such issues. Two of these tools are System i Collection Services and System i Performance Explorer.

In this appendix, we discuss troubleshooting methodologies that you can use to determine the cause of I/O-related issues that are encountered when using FlashCopy, Metro Mirror, and Global Mirror functions for external storage in such an environment. We also describe various System i Performance Tools reports and Performance Explorer (PEX) reports that are used in the PD/PSI process.

**Important:** When using the various tools and utilities, remember to collect data for the *same time period*. Then, you can relate the data from one collection to the data in the other collections.

## Collection Services

You can use the System i Performance Tools licensed program product (LPP) 5722PT1 to generate reports from the System i Collection Services data. If this product is not installed on the system, you can find it on the CDs that are shipped with the system. The CDs are labeled Lxxxx, where xxxx is the system language code (for example, L2924 is U.S. English).

After you install the Performance Tools, there is a 70-day trial period for the Performance Tools LPP. Installing this product allows you to generate Performance Tools reports as well as to manage Collection Services from a series of menus. If Performance Tools are installed and the 70-day grace period expires, then you can manage Collection Services using a set of native system CL commands or a set of system APIs. However, report generation is not available without the Performance Tools LPP.

### Starting a performance collection

To manage Collection Services using the Performance Tools LPP:

1. Enter **go perform** on the command line to access the main panel, as shown in Figure A-1.

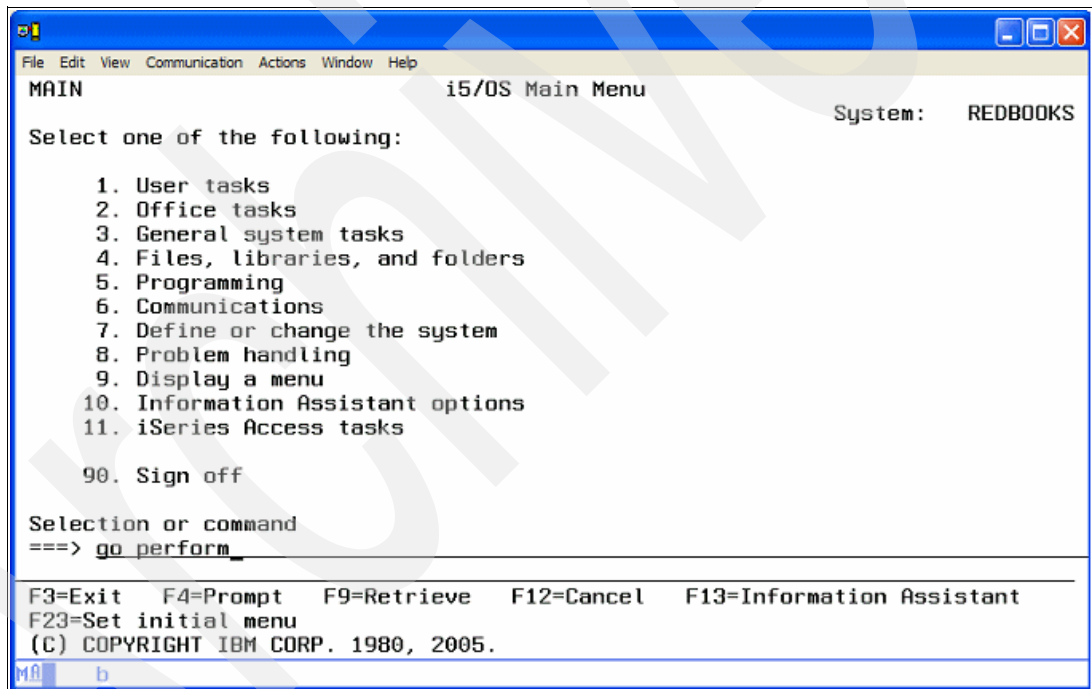


Figure A-1 Performance tools invocation

2. Set the desired collection attributes, and then start Collection Services (if it is not currently running) or cycle the currently running collection (if it is already started). Select **2. Collect performance data** (Figure A-2).

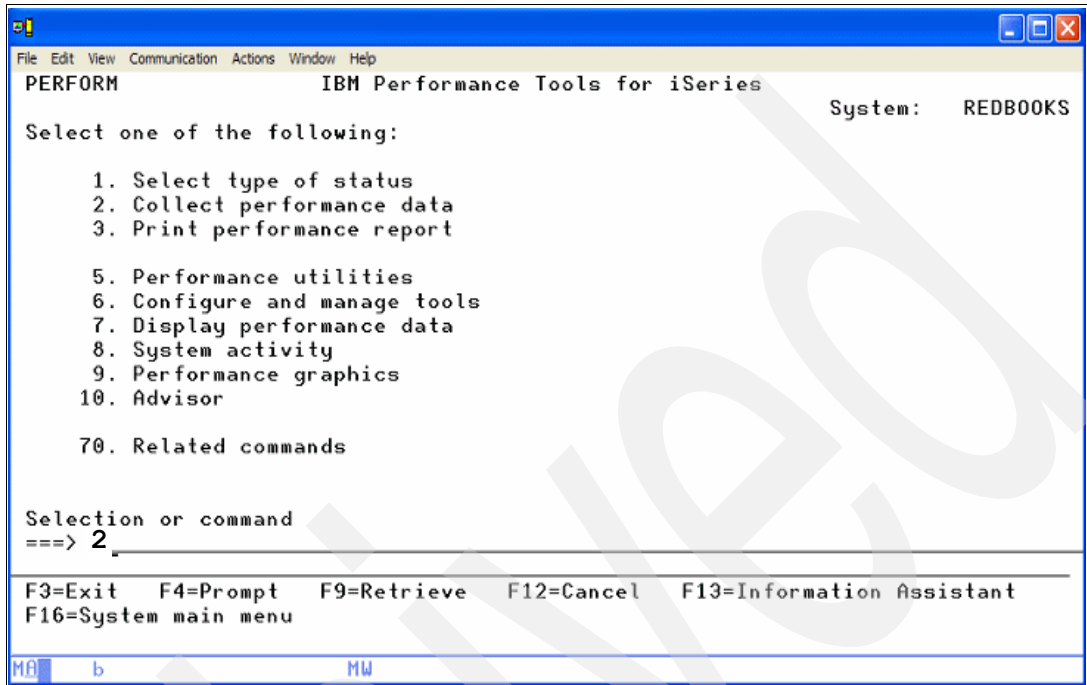


Figure A-2 Performance Tools main menu

3. Set the desired collection attributes. Select **2. Configure Performance Collection**. See Figure A-3.

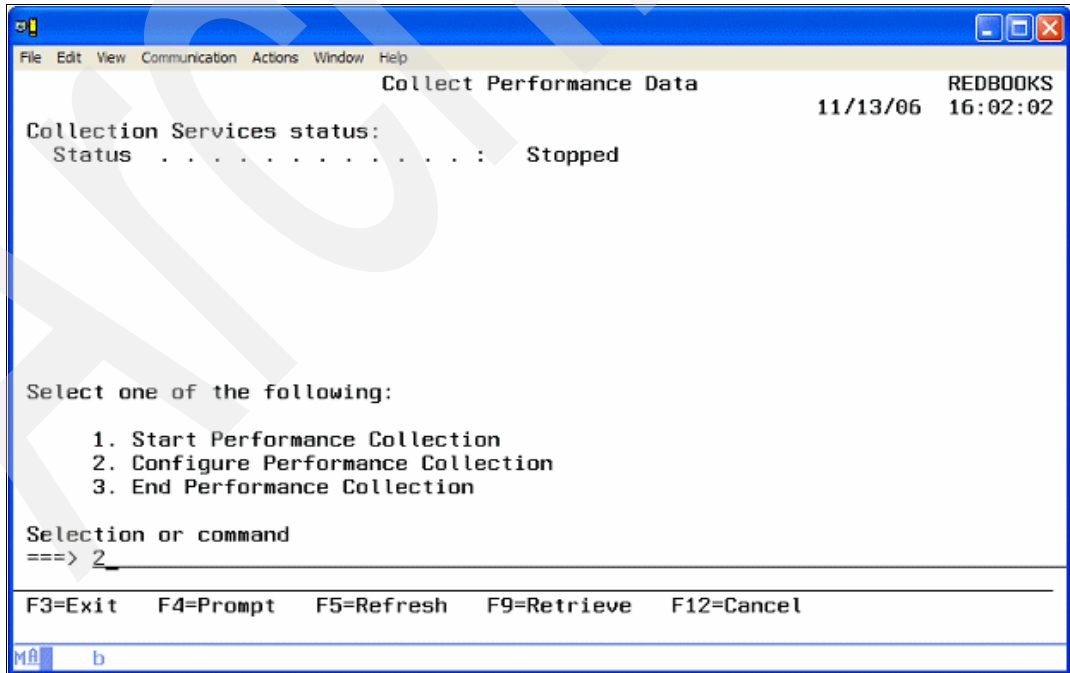


Figure A-3 Collect Performance Data options

4. Collection Services data is collected at intervals that you set using the Configure Perf Collection menu (Figure A-4).
  - a. Set the collection interval. Generally a 1 minute or 5 minute interval is desired for disk issues. Keep in mind that some of the data that is collected is averaged over the interval time period. The longer the collection interval is, the more diluted the data can become.
  - b. Ensure that Create database files is set to \*YES. This setting ensures that database files are created in the chosen library. The database files are necessary for report generation.

The native system CL command CFGPFCOL also accomplishes the functions on this panel, as follows:

```
CFGPFCOL INTERVAL(01.00) LIB(QMPGDATA) DFTCOLPRF(*STANDARDP) CYCTIME(000000)
CYCITV(24) RETPERIOD(00024 *HOURS) CRTDBF(*YES) CHGMLIB(*NO)
```

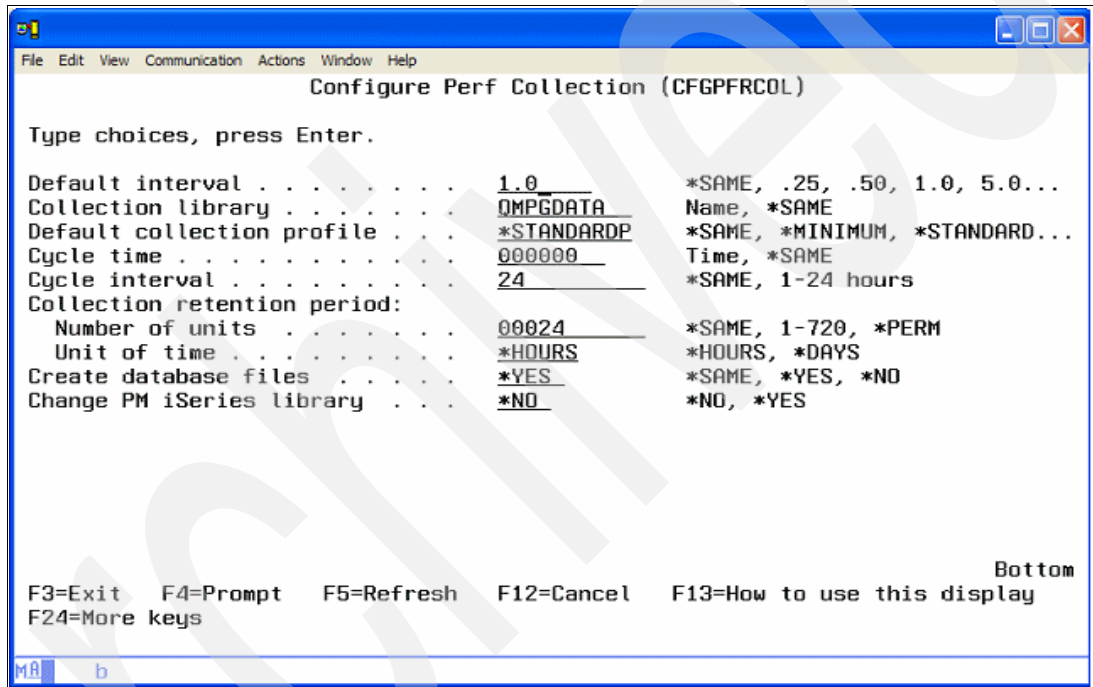


Figure A-4 Configure Perf Collection panel

- To start or cycle Collection Services, on the Collect Performance Data panel (Figure A-5), select **1. Start Performance Collection**. If Collection Services is currently running, the Status is *Started* and any additional attribute values are indicated.

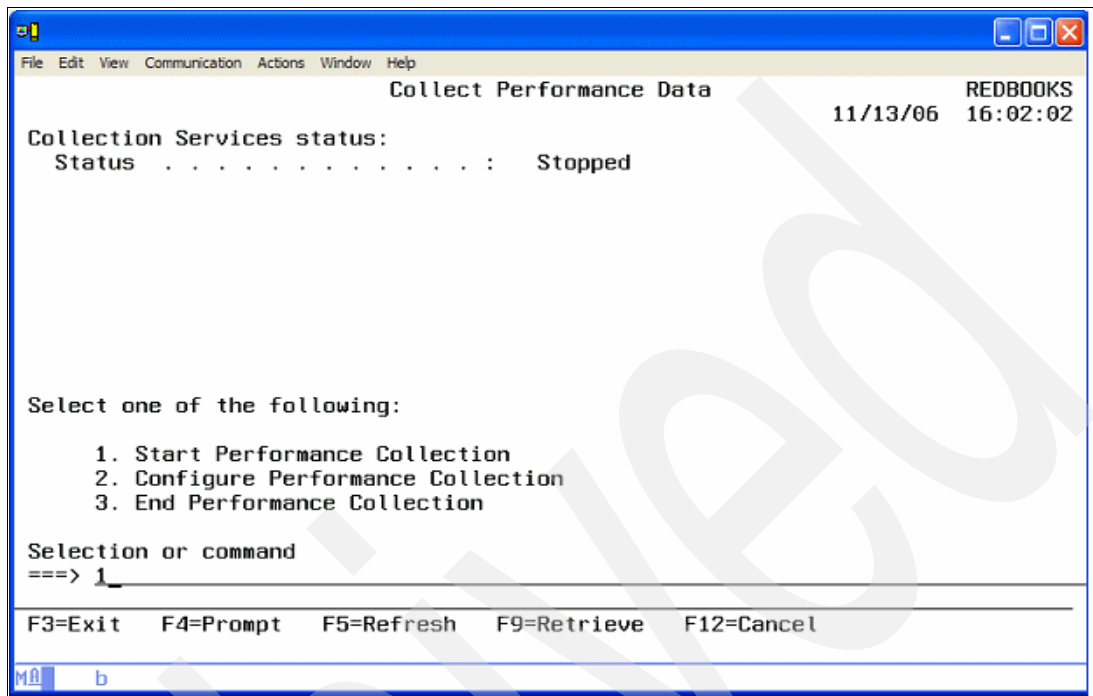


Figure A-5 Collect Performance Data panel

- 6. In the Start Performance Collection panel, the Collection profile setting of \*CFG uses the attributes that were set up previously. The Cycle collection setting of \*YES forces a currently running collection to be cycled. The new collection is started using the attributes that were set earlier.

Press Enter to start or cycle Collection Services. You might need to refresh the panel using F5.

The native system CL command STRPFRCOL also accomplishes this function as follows:

```
STRPFRCOL COLPRF(*CFG) CYCCOL(*NO)
```

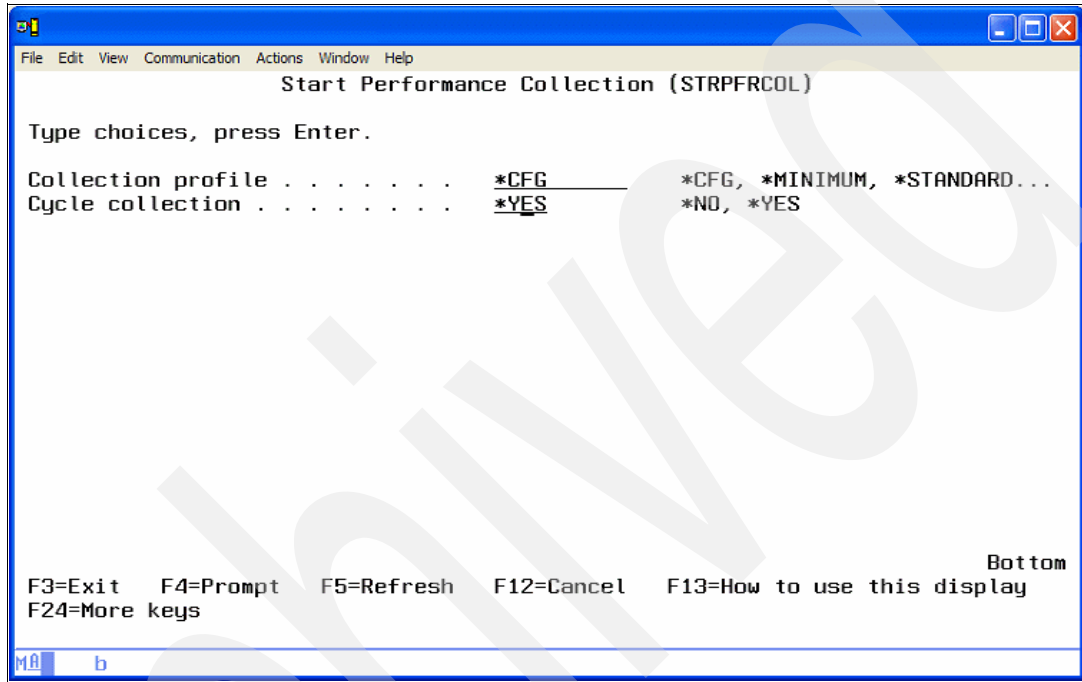


Figure A-6 Start Performance Collection attributes

7. Collection Services should now be started.

To end Collection Services, on the Collect Performance Data panel (Figure A-7), select **3. End Performance Collection.**

You can also end Collection Services using the ENDPFRCOL native system CL command as follows:

```
ENDPFRCOL FRCCOLEND(*NO)
```

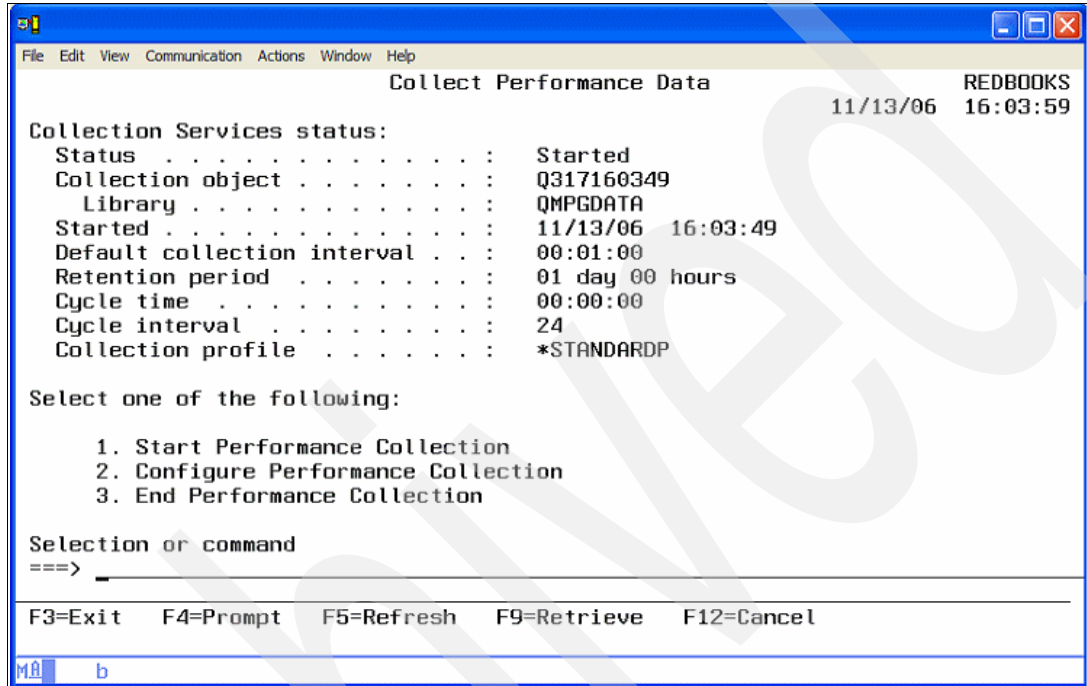


Figure A-7 Collection Services status

**Important:** After Collection Services is started, the data is collected and placed into a set of files that are found in the library that is chosen. The file data is retained for a period of time that is determined by the settings that are associated with the system Performance Monitor. You must check the period of time to ensure that the data collected is retained for the length of time that is desired, usually at least five days. If Performance Monitor is disabled, the retention period of the Collection Services data is permanent and the user must manage it.

## Checking the status of Performance Monitor

To check the status of Performance Monitor, you must review the attributes.

1. On a command line, enter `go pm400` (see Figure A-8).




Figure A-8 Starting Performance Monitor



2. On the Performance Monitor main menu, select **3. Work with PM eServer iSeries customization** (Figure A-9).

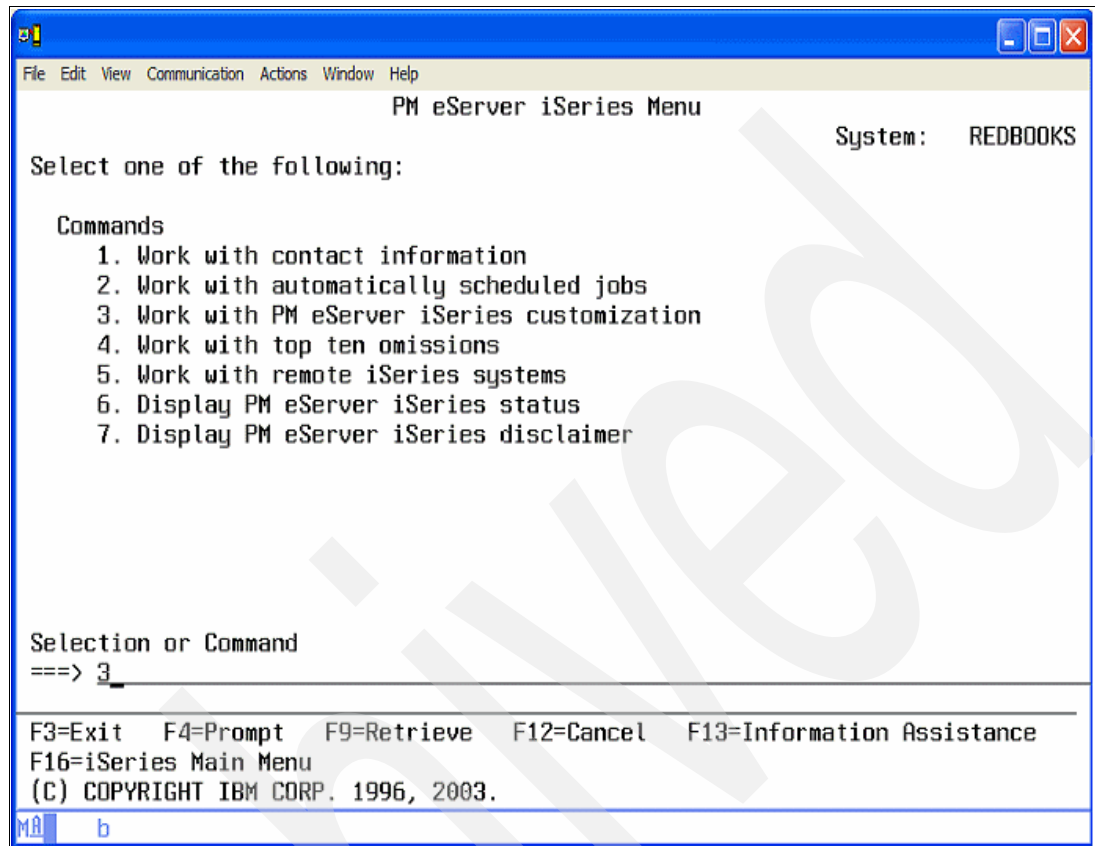


Figure A-9 Performance Monitor main menu

On the Work with PM eServer iSeries Customization panel (Figure A-10), the value of the Performance data purge days parameter determines how long to keep the database file data that is created by Collection Services.

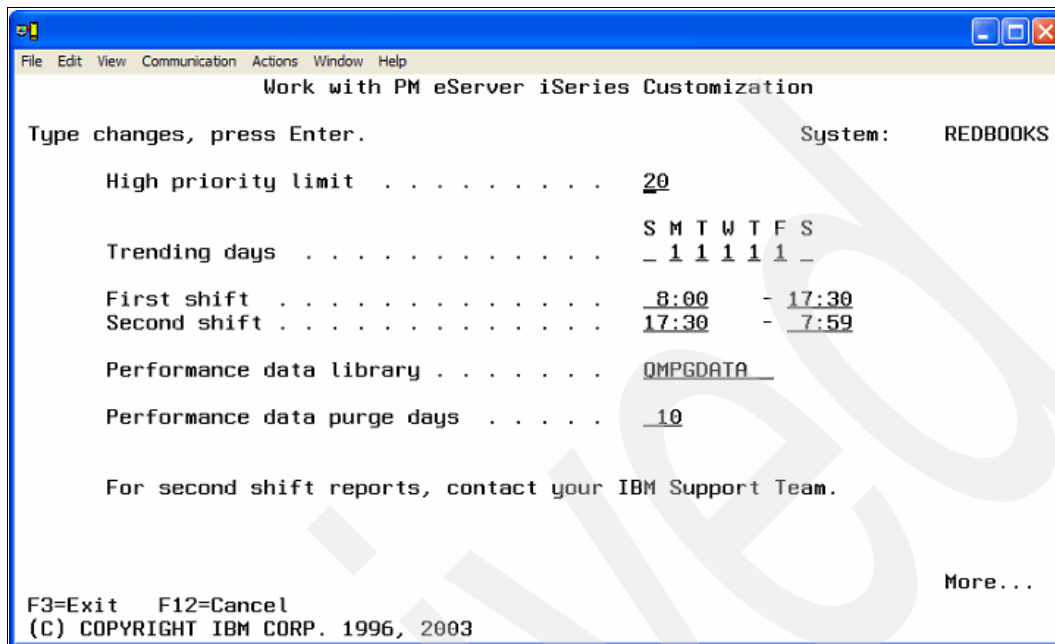


Figure A-10 Performance Monitor customization options

## Management of Collection Services

You can also manage Collection Services through calls to system APIs. The parameters for API calls are precise and must be entered in the correct order. Prior to the release of V5R3M0 i5/OS, these APIs were used to manage Collection Services on the iSeries system. Since V5R3M0, the native system CL commands STRPFCOL, CHGPFCOL, and ENDPFCOL have been made available. These commands virtually eliminate the need to use the APIs.

### Change Collection Services

To change Collection Service, use the following command:

```
CALL PGM(QYPSCSCA) PARM(*PFR ' X'00000384' 'QMPGDATA' X'000000A8'
X'0000000F' X'00000018' X'00000001' '*STANDARDP' X'00000000')
```

The command uses the following program:

```
QYPSCSCA /* API program
```

The command uses the following parameters:

```
*PFR' /* collection attribute, 10 characters
X'0000003C' /* interval 60 sec (1 minutes), 8 hexadecimal
'QMPGDATA' /* library name, 10 characters
'000000A8' /* retention 168 hrs (7 days), 8 hexadecimal
X'0000000F' /* cycle time 00:15:00 (12:15:00 AM), 8 hexadecimal
X'00000018' /* cycle interval 24 hrs, 8 hexadecimal
X'00000001' /* create database files, (1=Yes, 0=No) 8 hexadecimal
'*STANDARDP' /* default profile, 10 characters
X'00000000' /* return code, 8 hexadecimal
```

**Note:** Character values must be placed in single quotation marks. Hexidecimal values must be preceded by the *X* character with the value placed in single quotation marks.

### Start Collection Services

To start Collection Services, use the following command:

```
CALL QYPSSTRC PARM(*PFR ' *STANDARDP' X'00000000')
```

The command uses the following program:

```
QYPSSTRC /* API program
```

The command uses the following parameters:

```
*PFR' /* collection attribute, 10 characters  
*STANDARDP' /* default profile, 10 characters  
X'00000000' /* return code, 8 hexidecimal
```

### Cycle Collection Services

To cycle Collection Services, use the following command:

```
CALL QYPSCYCC PARM(*PFR ' X'00000000')
```

The command uses the following program:

```
QYPSCYCC /* API program
```

The command uses the following parameters:

```
*PFR' /* collection attribute, 10 characters  
X'00000000' /* return code, 8 hexidecimal
```

### End the collection

To end the collection, use the following command:

```
CALL QYPSENDC PARM(*PFR ' X'00000000')
```

The command uses the following program:

```
QYPSENDC /* API program
```

The command uses the following parameters:

```
X'00000000' /* return code, 8 hexidecimal
```

## Performance Tools reports

The Performance Tools product 5722PT1 is required to generate reports from Collection Services data. You can move the Collection Services collection object to any system that has the Performance Tools LPP installed at the same or later release level of the System i operating system. It is only necessary to restore the desired Collection Services management collection object to the target system. This object has a type of \*MGTCOL and an attribute of \*PFR.

To generate collection file data:

1. Locate the desired collection object. On the IBM Performance Tools for iSeries panel (Figure A-11), enter the WRKLIB CL command, and specify the collection library.




Figure A-11 WRKLIB command-line invocation

2. In the Work with Libraries panel, select **12=Work with objects** to work with objects in the collection library (Figure A-12).

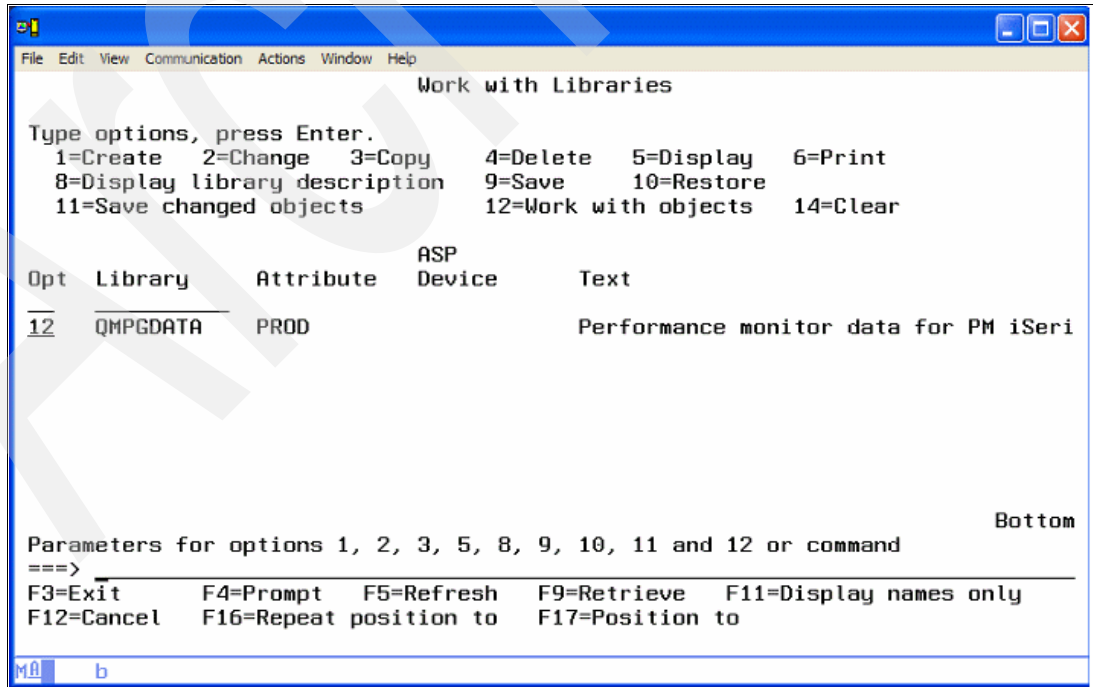


Figure A-12 Work with Libraries panel

- After you locate the desired object, save and then either restore or send the object to the target system using FTP (Figure A-13).

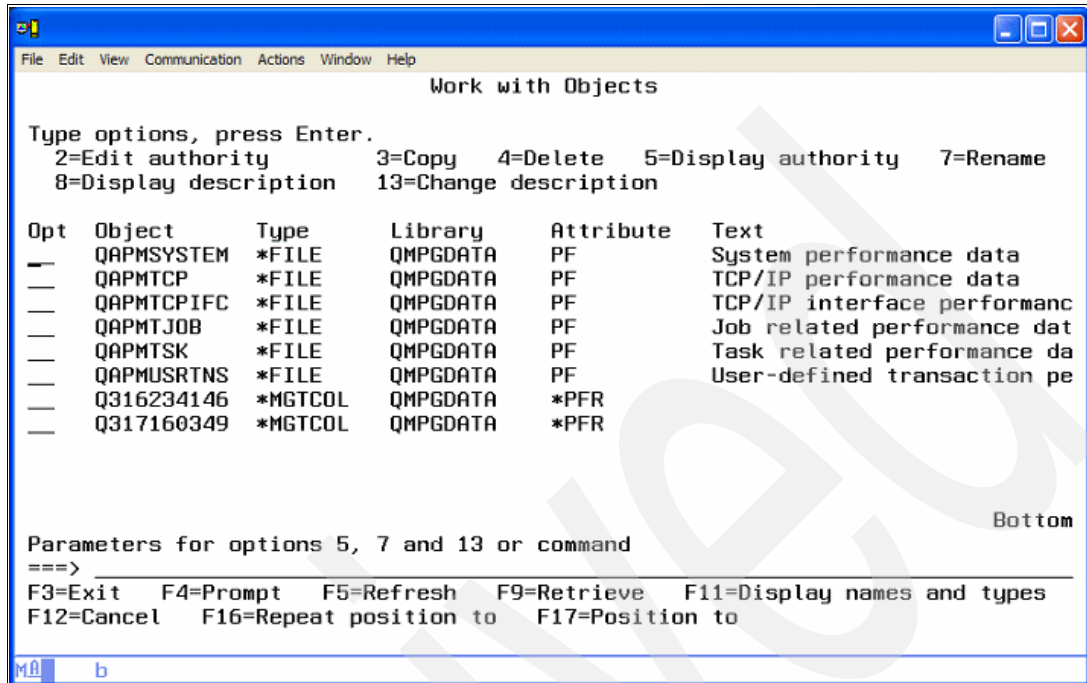


Figure A-13 Work with Objects panel

After the collection object is restored or received on the target system, generate the necessary database files using the Performance Tools menus.

- Enter **go perform** on the system command line to open the Performance Tools main menu (Figure A-14).

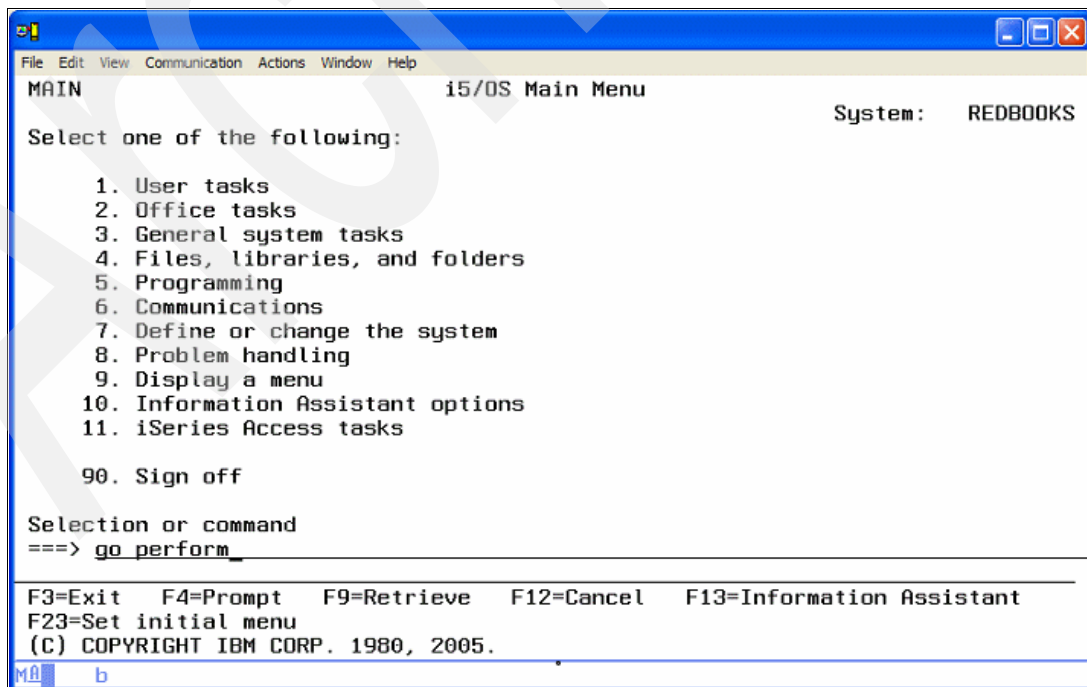


Figure A-14 Launching Performance Tools

5. In the IBM Performance Tools for iSeries panel (Figure A-15), select **6. Configure and manage tools**. The option to create performance data is contained in this option.

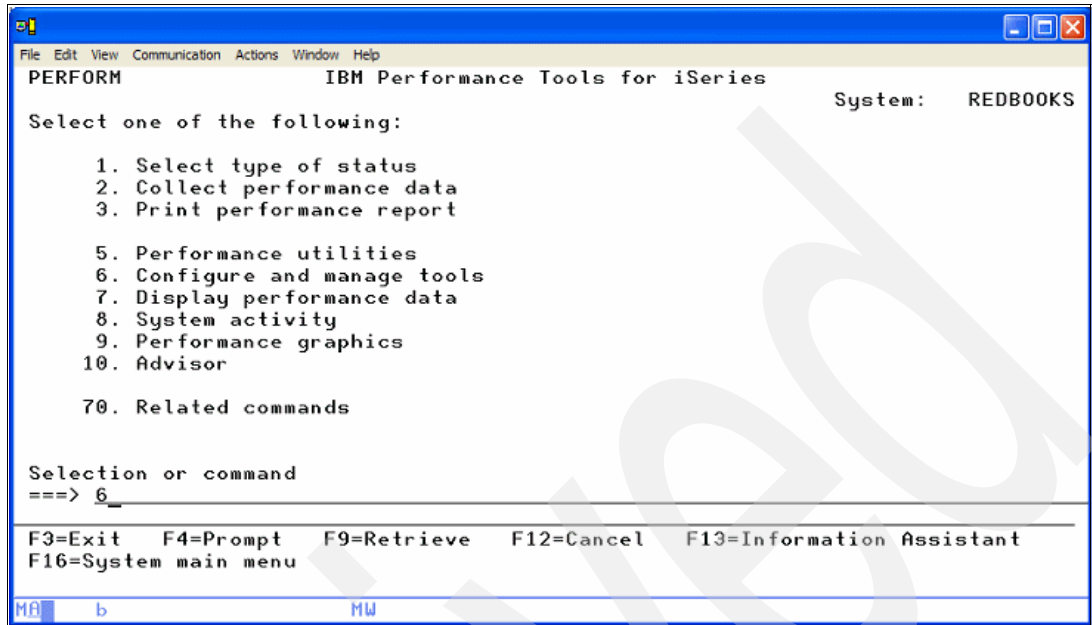


Figure A-15 Performance Tools main menu

6. On the Configure and Manage Tools panel (Figure A-16), select **5. Create performance data**.

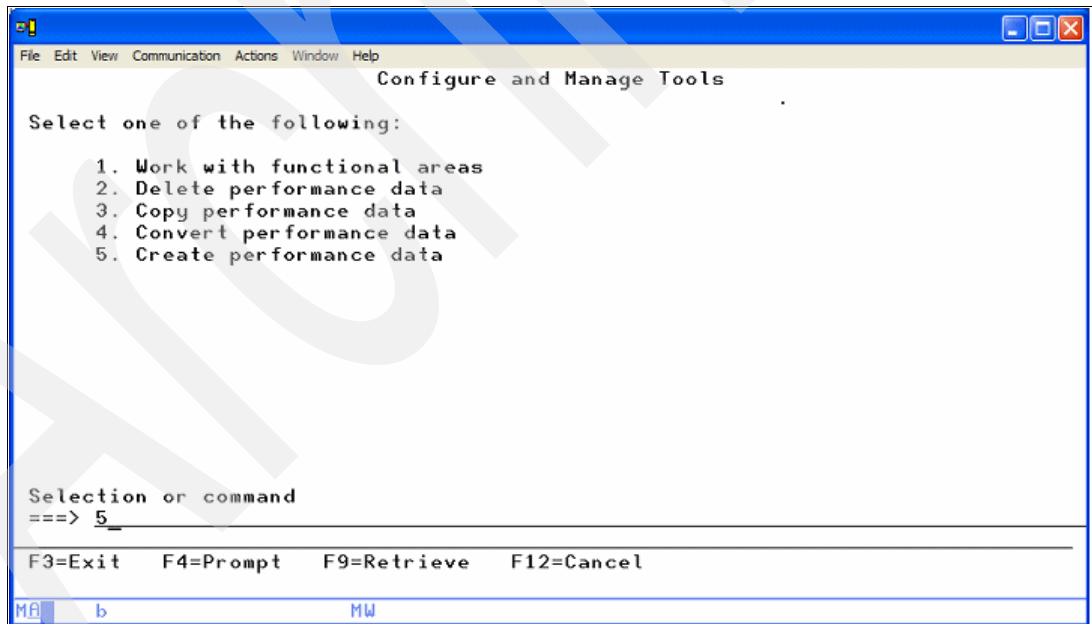


Figure A-16 Configure and Manage Tools panel

- On the Create Performance Data panel (Figure A-17), type the name of the collection and the name of the collection library. Then press Enter to create the collection file data.

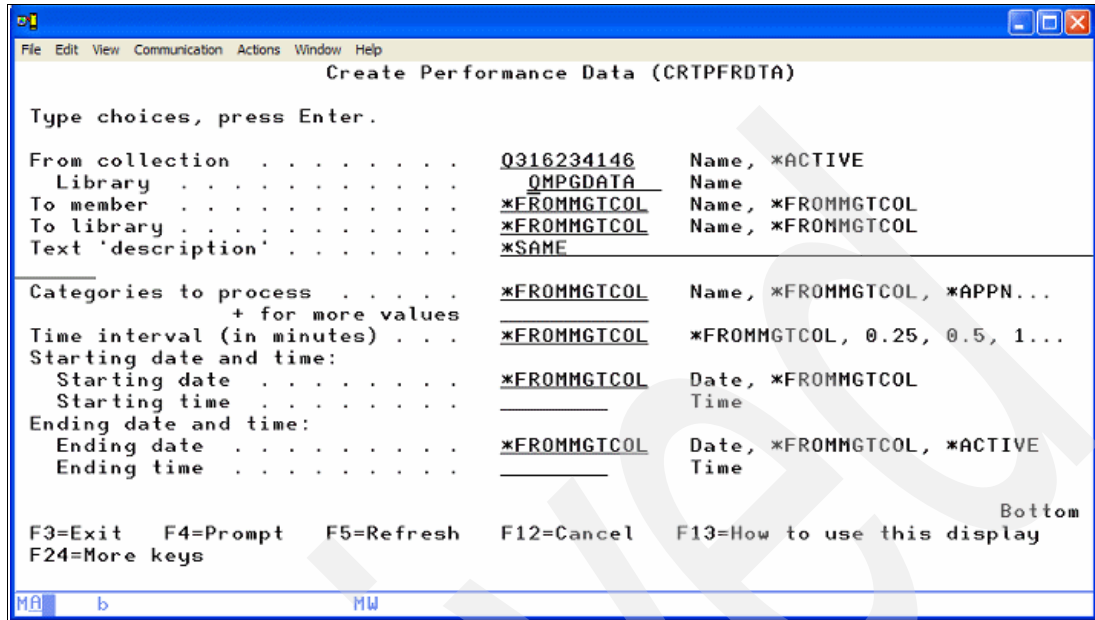


Figure A-17 Create Performance Data panel

## Generating a performance report

To generate a performance report using Performance Tools, choose the collection from which you want the report generated and then choose the type of report to produce. There are two reports of interest when working on disk issues:

- ▶ The Disk activity section of the System report
- ▶ The Disk utilization section of the Resource report

To generate a performance report:

1. Enter the **go perform** command to open the Performance Tools main menu (Figure A-18).

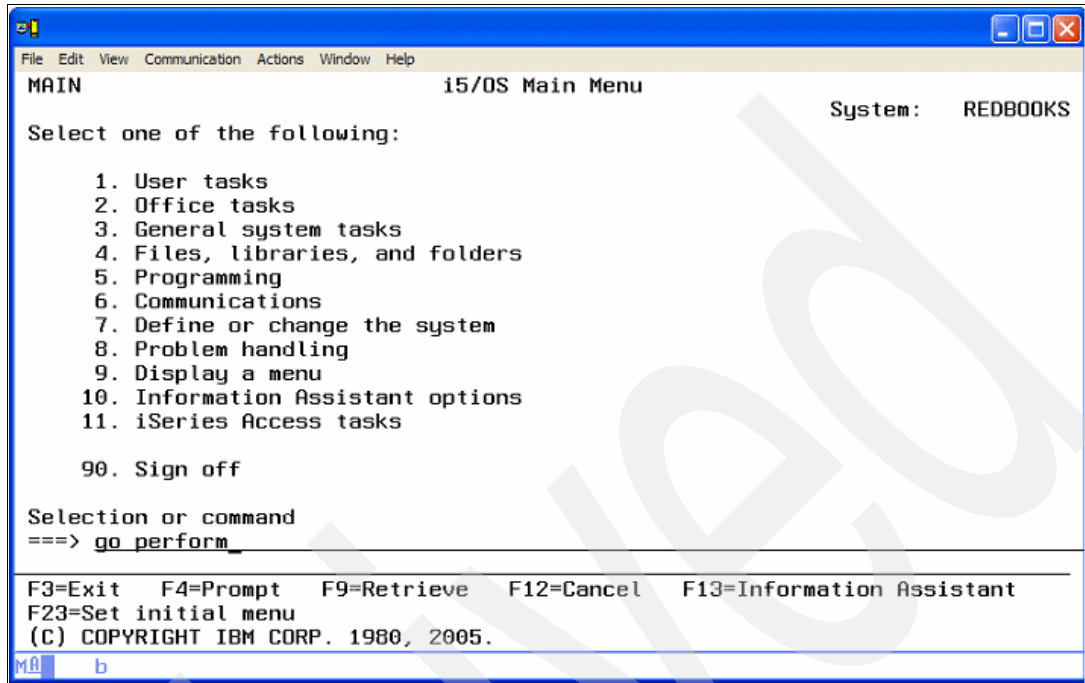


Figure A-18 Performance tools invocation

2. On the IBM Performance Tools for iSeries panel (Figure A-19), select **3. Print performance report**.

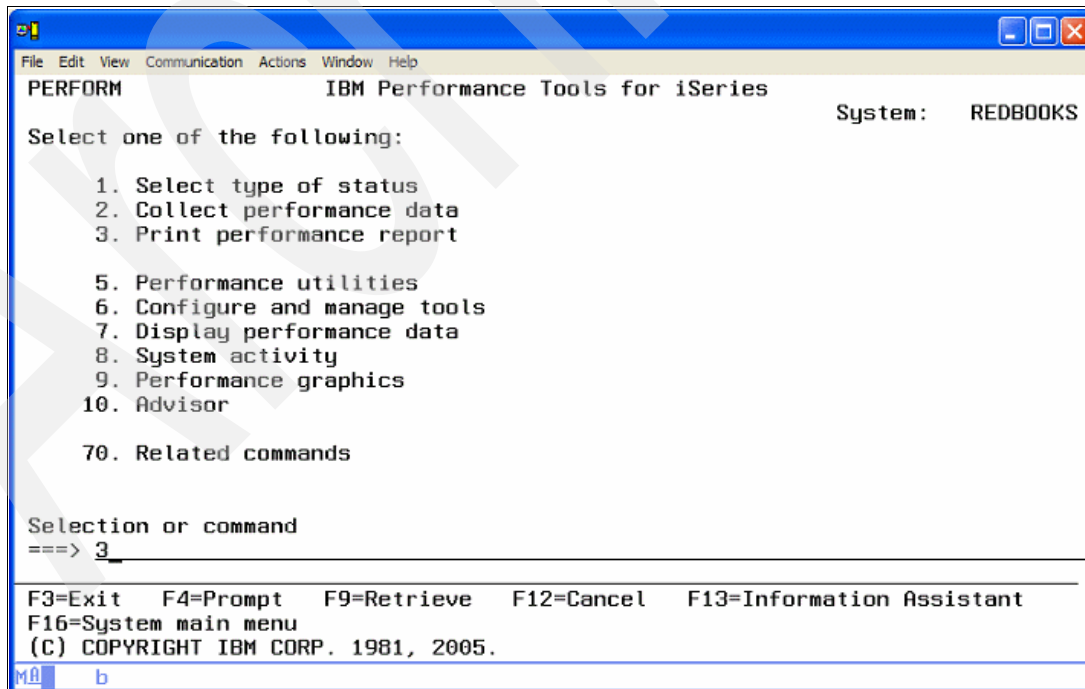


Figure A-19 Performance Tools main menu



- On the Print Performance Report panel (Figure A-20), choose the library that contains the data from which to generate the performance reports, and press Enter.

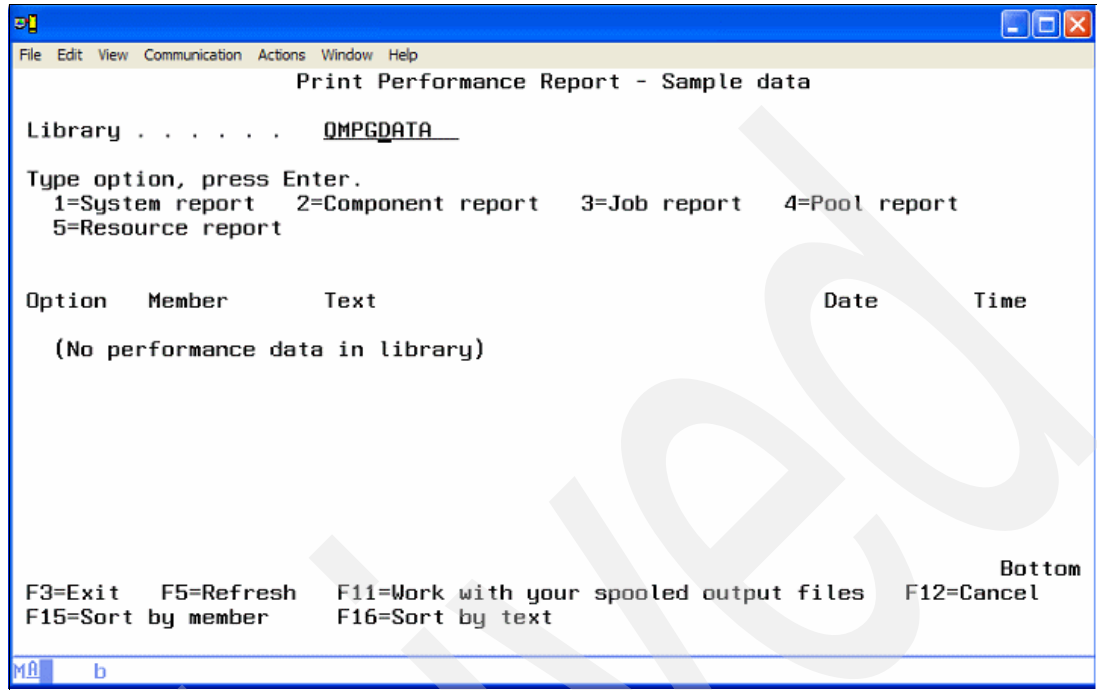


Figure A-20 Specifying the data library to generate the Performance reports

You now see all of the collections that are available to be used within the chosen library (Figure A-21).

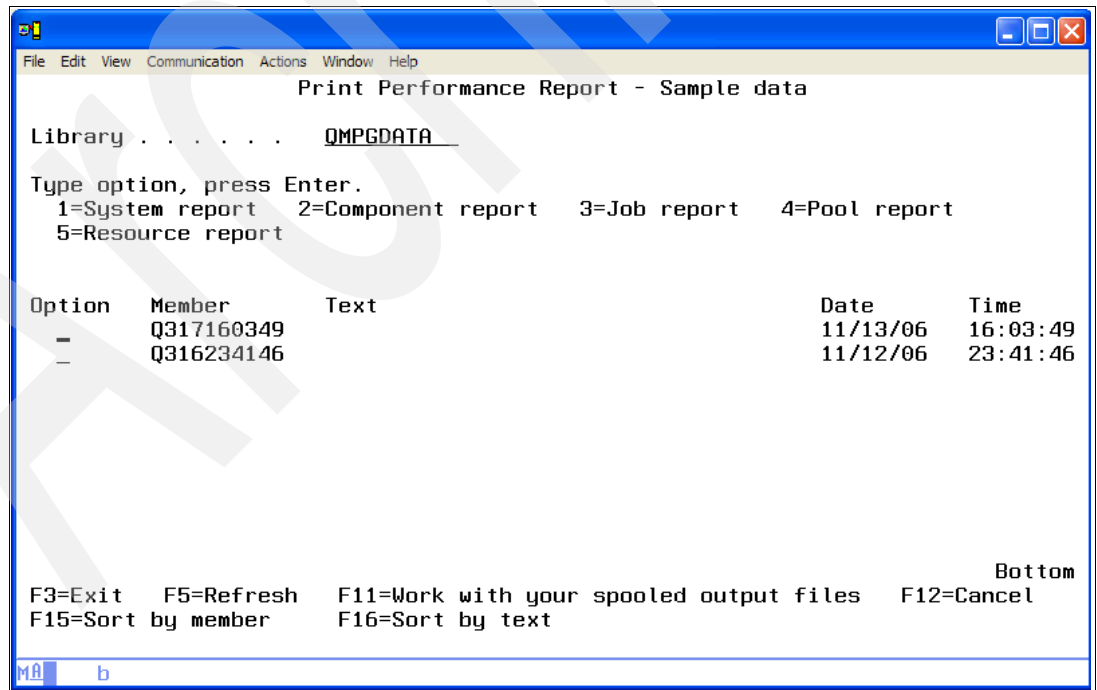


Figure A-21 Print Performance Report panel

## Generating a system report

To generate a system report:

1. Choose the type of report to generate and the collection to use. On the Print Performance panel (Figure A-22), choose **1=System report**.




Figure A-22 Choosing to create a system report

2. On the Select Sections for the Report panel (Figure A-23), choose the section of the report or press F6 to select all sections. In this example, we examine the Disk utilization section.

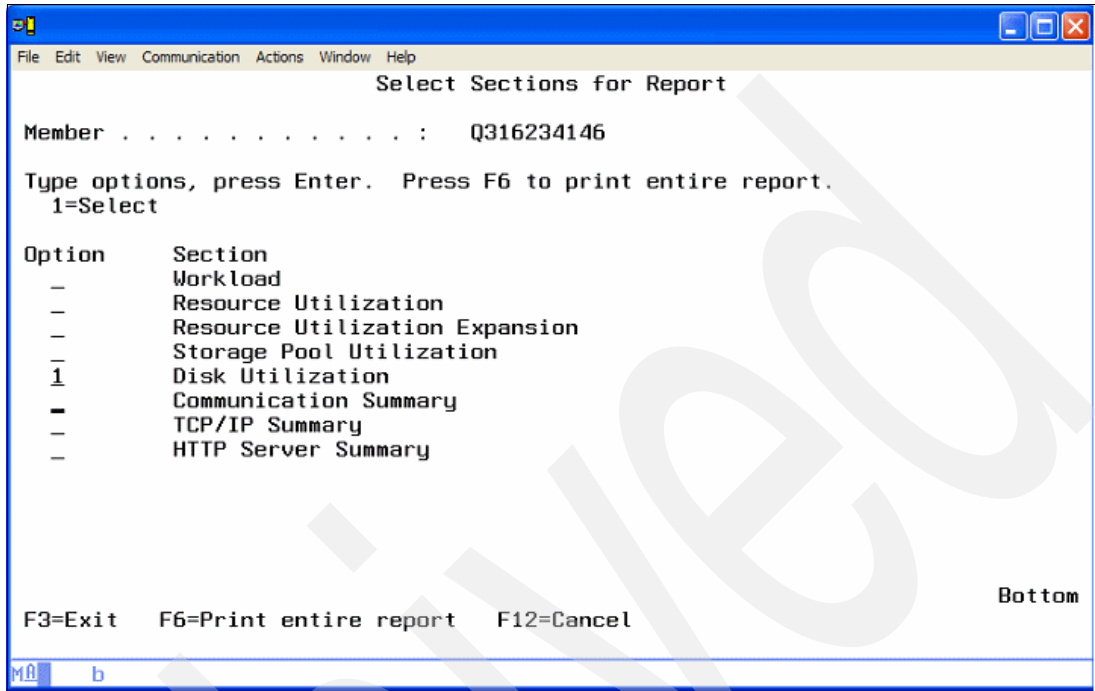


Figure A-23 Select Sections for Report panel

3. On the Select Categories for Report panel (Figure A-24), choose the data filter to use. We use the Time interval category to generate the report.

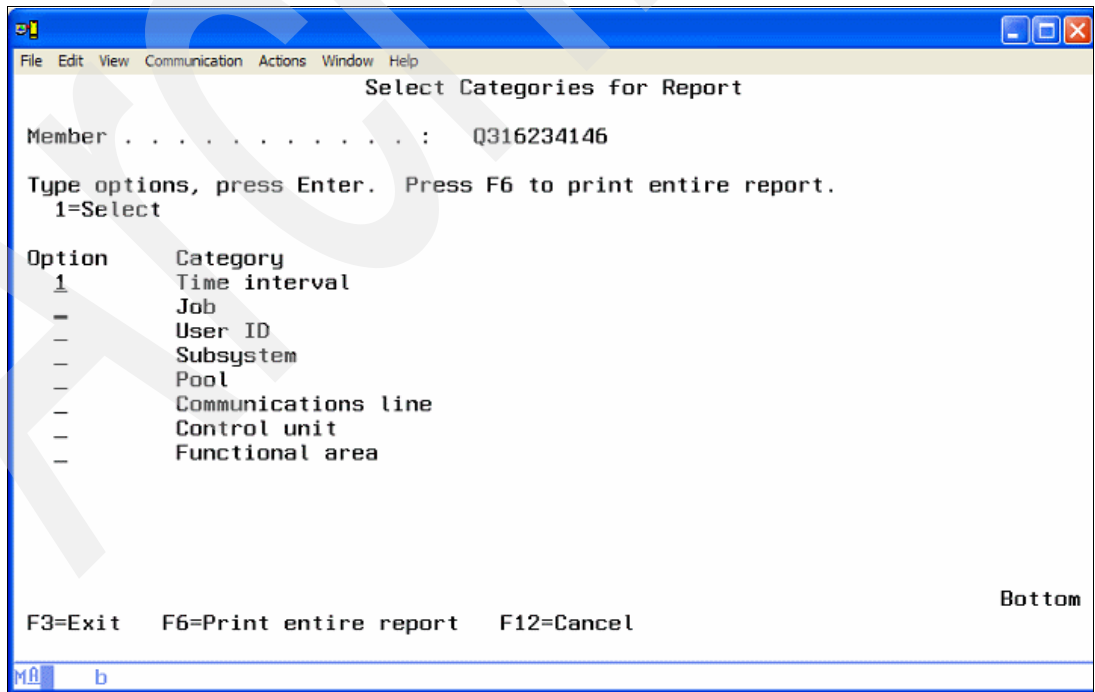


Figure A-24 Select Categories for Report panel

- On the Select Time Intervals panel (Figure A-25), select one or more time intervals.

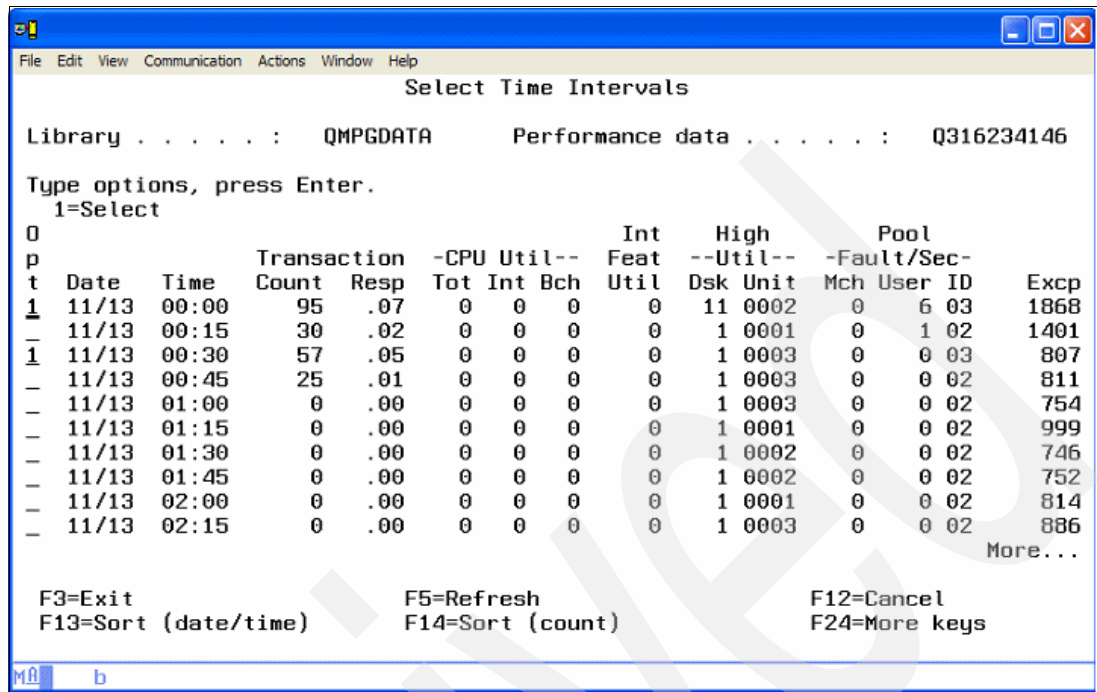


Figure A-25 Select Time Intervals panel

- On the Specify Report Options panel (Figure A-26), give the report a title. In this example, we specify a title of *System - Disk utilization*.




Figure A-26 Specify Report Options panel

- The report generation request is submitted to batch. You return to the Print Performance Report panel, which shows the information about the job that was submitted (see Figure A-27).

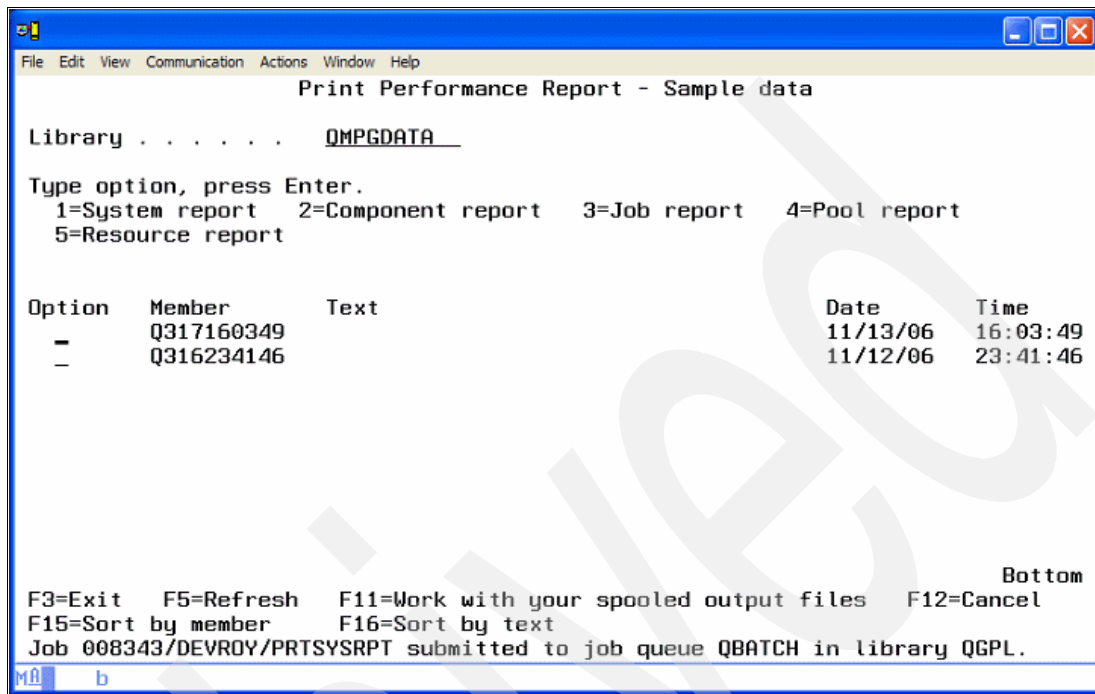


Figure A-27 Submitted job information

- The report is created as a spooled file found in the output queue for the user submitting the request. On the Work with Job Spooled Files panel (Figure A-28), select **5=Display** to display the system report that was generated.

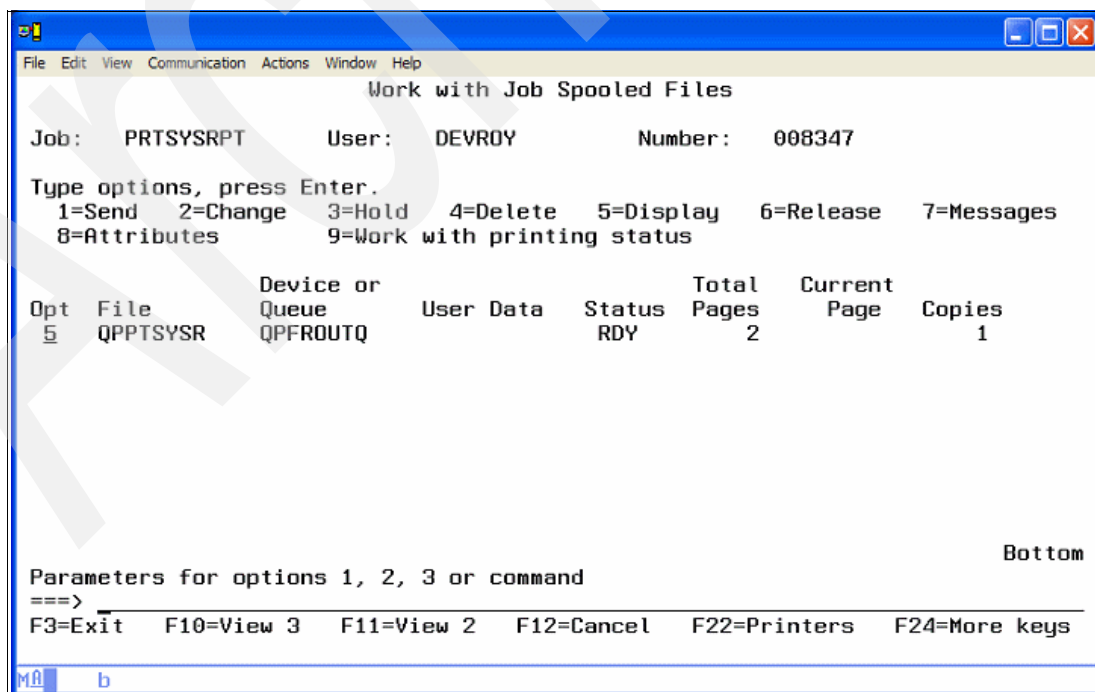


Figure A-28 Work with Job Spooled Files panel

The Disk utilization section of the system report (Figure A-29) contains much useful information:

- ▶ The *unit ID* and *unit name* along with the disk type and size.
- ▶ *I/O processor (IOP) utilization and name.*
- ▶ *Disk CPU utilization* refers to the processor utilization on the physical disk unit. This value is of no meaning with regard to external disk units.
- ▶ *Percent full and utilized* indicate how full and busy the disk units are.
- ▶ *I/O operations per second* are shown along with their average size.

The last three columns relate to the average I/O time per unit.

- ▶ *Service time* is the time spent outside of the System i environment.
- ▶ *Wait time* is the time spent on the System i environment.
- ▶ *Response time* is the sum of the service and wait times.

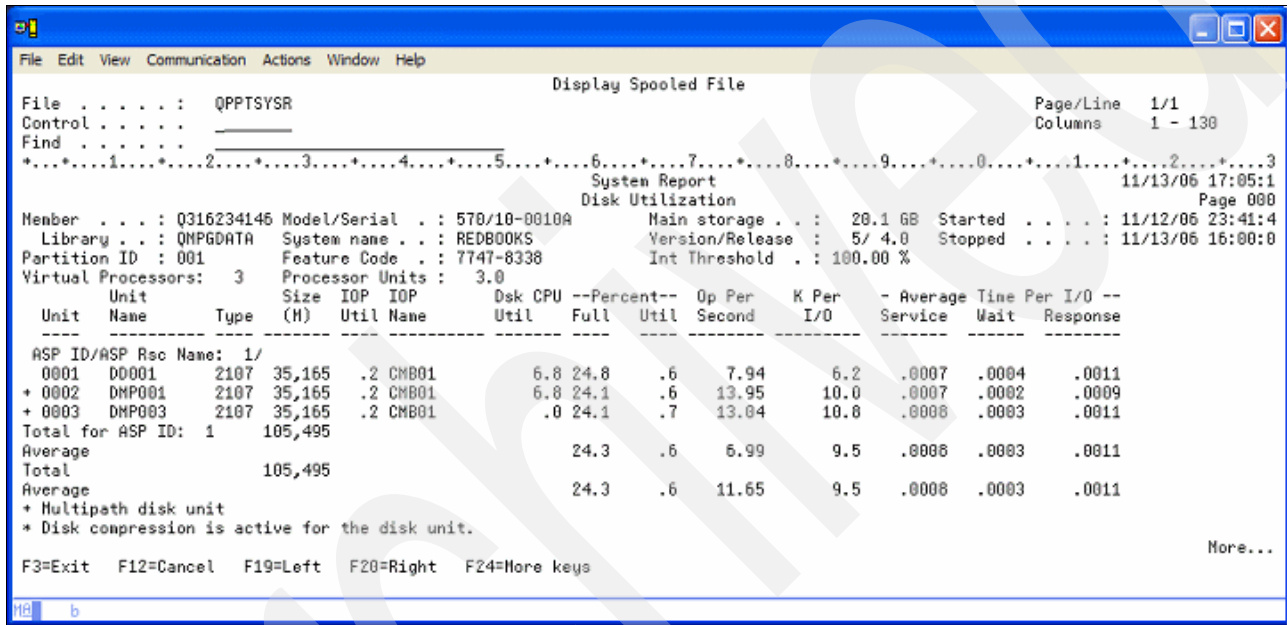


Figure A-29 System data report

## Generating a resource report

To generate a Resource report, use the following steps. In this section, we refer to some of the report panels from “Generating a system report” on page 466, because the panels are similar.

1. Choose the type of report to generate and the collection to use. On the Print Performance Report panel (Figure A-22 on page 466), choose **5=Resource report** to generate a resource report.
2. In the Select Section for Report panel (Figure A-23 on page 467), choose the section of the report or use F6 to select all sections. We generate the Disk utilization section.
3. In the Select Categories for Report panel (Figure A-24 on page 467), choose the data filter to use. We use the Time interval category to generate the report.
4. In the Select Time Intervals panel (Figure A-25 on page 468), select one or more time intervals. We choose the time period or periods of interest.

There are blank entries under the Disk high utilization column. No data is seen for disk resources for those time periods. This will only be seen when a one minute time interval is selected. Also notice that only those time periods with values in that column have been selected.

5. In the Specify Report Options panel (Figure A-26 on page 468), give the report a title. The report generation request is submitted to batch.
6. The report is created as a spooled file found in the output queue for the user who submitted the request. Use **5=Display** to display the system report generated.

This report provides summary and detail information regarding the disk units. The information is provided for each time interval that is selected. The detail information is provided for each disk unit for each time interval that was selected. This report shows the disk unit ID information along with I/O rates, disk utilization, service time, wait time, and queue length. The Disk CPU utilization values have no meaning for external disk units.

## Performance Explorer

Performance Explorer is an internal trace utility that can collect detailed information about the System i environment. It collects a large amount of data in a short time.

Prior to running this trace tool, you must apply several required PTFs. Failure to apply these PTFs can cause the system to terminate abnormally. It is always best to check with IBM software support before you run PEX traces.

For more information about Performance Explorer, review the information in the i5/OS Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v5r4/index.jsp?topic=/rzahx/rzahxcollectinfoappperf.htm>

## DS8000 troubleshooting

Use the following tips to help investigate any issues that might arise when dealing with Copy Services with a System i host:

- ▶ Verify that you have and installed the activation keys for your storage images. See Appendix B, “Installing the storage unit activation key using a DS GUI” on page 473.
- ▶ Verify that the bandwidth can handle the Copy Services solution that you have implemented.
- ▶ Check the Service log and verify that there are no bad host bus adapters or other errors:  
To locate the error log on the DS8000:
  - a. Log in to the DS8000 Hardware Management Console (HMC).
  - b. Click the **Service Applications** icon.
  - c. Click **Service Focal Point**.
  - d. Click **Manageable Service Events**.
  - e. Click **OK**.
  - f. The next window lists any issues that the DS8000 is having in your environment. To view details about the issue, select the issue, and then click **Selected** → **View Details** to learn more about the issue.

If any hardware problems are present, call your IBM Customer Service Representative.

- ▶ In a Copy Services solution, verify that the ports that are used in the primary or source system are the same ports that are used in the secondary or target system.

- ▶ If you are using a switch, verify that the ports that are used in the Copy Services solution are zoned.
- ▶ Verify that a sufficient number of FCP paths is assigned to your source and target sites.
- ▶ If you plan to use both Metro Mirror and Global Copy between a pair of storage units, we recommend that you use separate logical and physical paths for Metro Mirror and another set of logical and physical path for Global Copy.
- ▶ Verify that the source and target LUNs are of the same size.
- ▶ On the System i host, verify that there are no more than six FC adapter IOP pairs per one high-speed link (HSL) loop.
- ▶ On the System i host, verify that the 64-bit slots are being used for the Fibre Channel I/O adapter (IOA) and IOP card connections. The 64-bit slots are C01-05, C08-09, and C14-15. For more information, see *iSeries in Storage Area Networks A Guide to Implementing FC Disk and Tape with iSeries*, SG24-6220.
- ▶ Consider one FC adapter-IOP pair per multi-adapter bridge on a System i host.
- ▶ For FlashCopy, the source and target volumes or LUNs must be in the same DS8000 image (logical partition).
- ▶ If you are using Global Mirror, you must have Point-in-Time Copy function authorization for the secondary storage unit.
- ▶ If you will use Global Mirror during failback on the secondary storage unit, you must also purchase a Point-in-Time Copy function authorization for the primary storage system.
- ▶ Verify that multipath has been set up on the storage system.



## Installing the storage unit activation key using a DS GUI

In this appendix, we explain how to install the Licensed Internal Code feature activation keys for the IBM System Storage DS8000 products. These activation keys are essential to Copy Services on the DS8000 products. Without the installation of the activation keys, Copy Services does not work on the DS8000 products.

For the DS8000 system, apply the Licensed Internal Code feature activation keys:

1. Use a Web browser to connect to the IBM Disk storage feature activation Web page (see Figure B-1):

<http://www.ibm.com/storage/dsfa>

2. Click **IBM System Storage DS8000 series**.



The screenshot shows the IBM Disk storage feature activation web page. The page has a dark header with the IBM logo and a navigation menu with links for Home, Solutions, Services, Products, Support & downloads, and My IBM. A 'Country/' dropdown menu is visible in the top right. Below the navigation, there is a breadcrumb trail '← IBM System Storage' and a 'Welcome [ IBM S' message. The main content area is titled 'Disk storage feature activation' and contains a sidebar with links for 'Disk storage feature activation', 'Select DS6000 series machine', 'Select DS8000 series machine', and 'Help'. The main text explains that the DSFA application provides feature activation codes for DS6000 and DS8000 series machines. It also includes a 'Related links' section with links to 'IBM System Storage support feedback' and 'Electronic service call'. At the bottom, it states 'To begin, select your machine type:'.

Figure B-1 Disk storage feature activation page

3. To complete the required information to be entered on the following Disk storage feature activation Web site (Figure B-4 on page 475), perform the following steps:
  - a. Access the DS8000 Storage Manager GUI (refer to 9.1.3, “Accessing the DS GUI interface” on page 338) and select **Real-time manager** → **Manage hardware** → **Storage units** from the left navigation panel.
  - b. Note the *Model and Serial Number information* from the storage unit whose licensed functions are to be activated (see Figure B-2).
  - c. Select this storage unit by clicking **Select** and selecting **Properties** from the Select Action drop-down menu.

**Storage images: Real-time**

Refresh Last refresh: Friday, November 2, 2007 10:19:30 AM CDT

--- Select Action --- Print report

Select ^	Nickname ^	Serial Number ^	Status ^	Storage unit ^	Storage Complex ^	Total Ra
<input checked="" type="checkbox"/>	Non-SVT 7589951	7589951	<input checked="" type="checkbox"/> Normal	<7589950>	IbmStoragePlex	
<input type="checkbox"/>	SVT 7589952	7589952	<input checked="" type="checkbox"/> Normal	<7589950>	IbmStoragePlex	

Page 1 of 1 Total: 2 Filtered: 2 Displayed: 2 Selected: 1

Figure B-2 DS8000 Storage units panel

- d. Get the required **Machine signature** information from the General properties panel (Figure B-3).

**Storage Unit Properties: Real-time**

**\* General**

Hardware

General

Accept or modify the nickname and enter a description. Then click OK or Apply.

Attribute	Value
MTMS	IBM.2107.9A2.7589950
Machine signature	00a4-d796-fafc-3e41
WWNN	5005076303FFFD3E
Emulated MTMS	IBM.2107.9A2.7589950
Power state	<input checked="" type="checkbox"/> On
Processor complex status	<input checked="" type="checkbox"/> Normal

\*Nickname  
<7589950>

Description

Figure B-3 DS8000 Storage Unit Properties panel

4. Open the browser to the Disk storage feature activation Web site to display the required information. In the Select DS8000 series machine panel (Figure B-4), select your machine type. Then specify your machine's serial number and signature. Click **Submit** to continue.

IBM

Country/r

Home Solutions Services Products Support & downloads My IBM

Welcome [ IBM Sig

← IBM System Storage

Disk storage feature activation

Select DS6000 series machine

Select DS8000 series machine

- View machine summary
- Manage activations
- Retrieve activation codes
- Assign function authorization
- View authorization details

Help

Related links

- IBM System Storage support feedback
- Electronic service call

## Select DS8000 series machine

Enter your machine information. You can find this information in the DS Storage Manager application under storage unit properties.

You must complete all fields marked with an asterisk (\*).

Click Submit when you are finished.

Machine type \* 2107

Serial number \* 75-89950 (12-34567)

Machine signature \* 00a4-d796-fafc-3e41 (1234-5678-9012-3456)

Submit

Figure B-4 Select DS8000 series machine panel

5. From the left navigation pane of the browser, select **Retrieve activation codes**. From the Retrieve activation codes window, either write down the codes for each product and storage image or export the codes to a PC file.

6. Access the DS8000 Storage Manager GUI (refer to section 9.1.3, “Accessing the DS GUI interface” on page 338) and select **Real-time manager** → **Manage hardware** → **Storage images** from the left navigation panel. Select the check box for the storage image whose LIC features are to be activated and for Select Action, click **Apply Activation Codes** (see Figure B-5).

**Note:** In a 2107 Model 9A2, a logical partition model, repeat this step for each storage image.

Storage images: Real-time

Refresh Last refresh: Friday, November 2, 2007 10:19:30 AM CDT

--- Select Action --- Print report Do

Select ^	Nickname ^	Serial Number ^	Status ^	Storage unit ^	Storage Complex ^	Total Ra
<input checked="" type="checkbox"/>	Non-SVT 7589951	7589951	Normal	<7589950>	IbmStoragePlex	
<input type="checkbox"/>	SVT 7589952	7589952	Normal	<7589950>	IbmStoragePlex	

Page 1 of 1 Total: 2 Filtered: 2 Displayed: 2 Selected: 1

Figure B-5 DS8000 Storage Images panel

7. Enter the DS8000 Licensed Internal Code feature activation keys that you retrieved from the Disk storage feature activation Web site. Either manually type the keys or import the key file from your PC. Then click **OK** to continue (see Figure B-6).

**Note:** In order to see the capacity and storage type that is associated with the successful application of the activation codes, repeat this step.

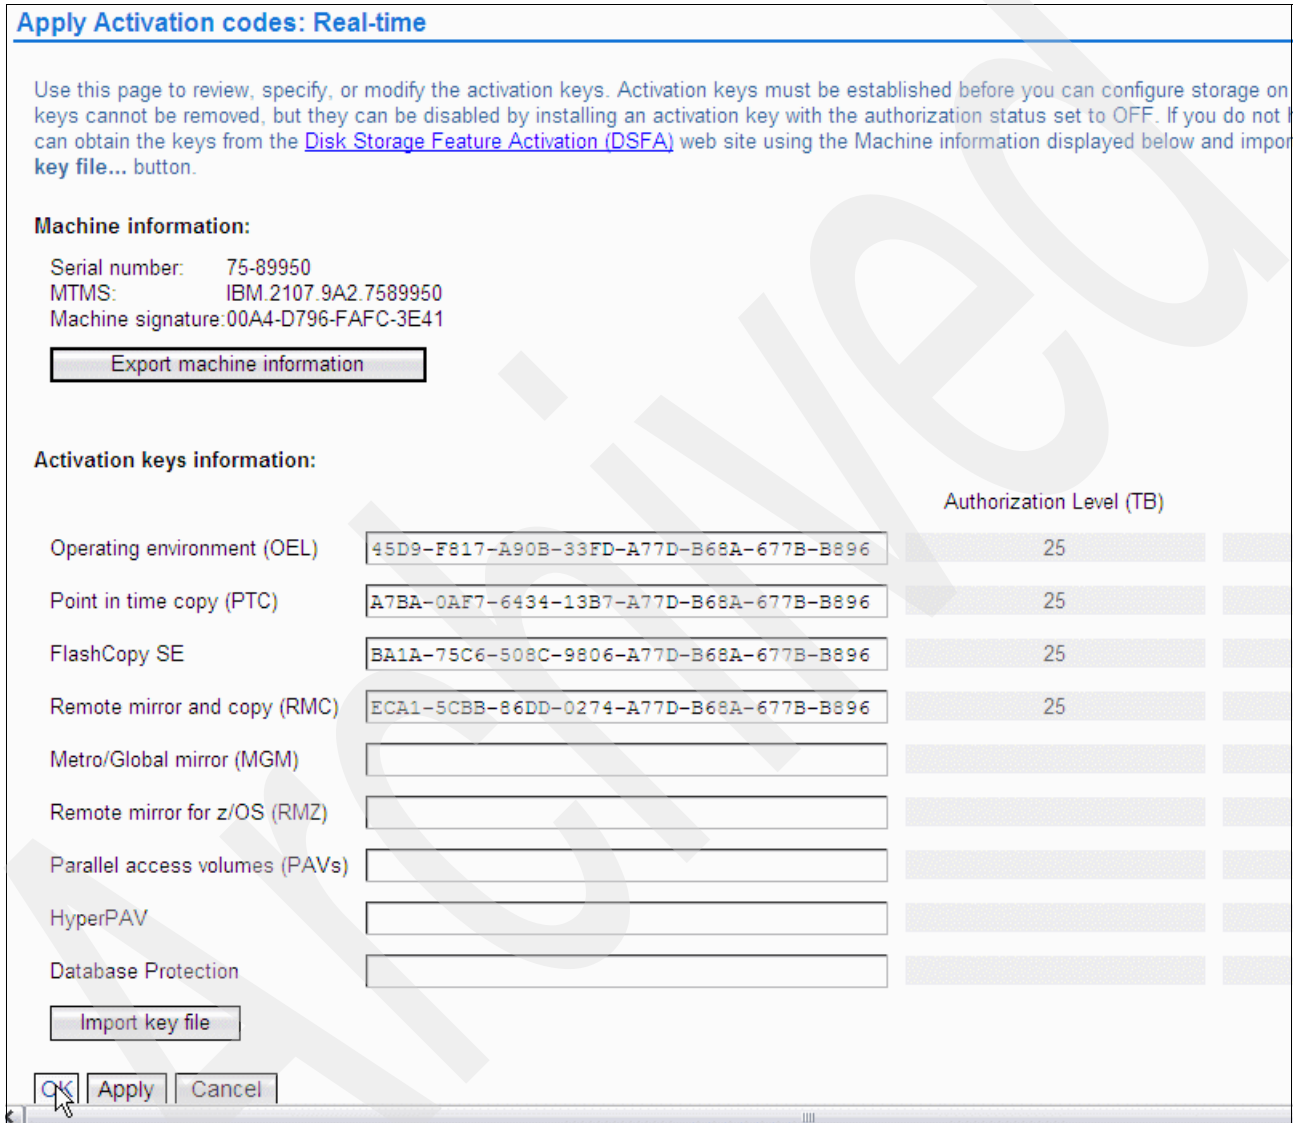


Figure B-6 Entering the activation codes

The following message displays:

CMUG00092W "This operation applies the activation codes to the storage image. Select OK to apply the activation codes. Select Cancel to cancel the operation."

8. Click **OK** to finish the application of the DS Licensed Internal Code activation codes.

Archived

# Related publications

We consider the publications that we list in this section particularly suitable for a more detailed discussion of the topics that we cover in this IBM Redbooks publication.

## IBM Redbooks publications

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 481. Note that some of the documents that we reference here might be available in softcopy only.

- ▶ *AS/400 Remote Journal Function for High Availability and Data Replication*, SG24-5189
- ▶ *Clustering and IASPs for Higher Availability on the IBM eServer iSeries Server*, SG24-5194
- ▶ *IBM eServer iSeries Migration: A Guide to Upgrades and Migrations to System i5*, SG24-7200
- ▶ *IBM eServer iSeries Migration: System Migration and Upgrades at V5R1 and V5R2*, SG24-6055
- ▶ *IBM eServer i5 and iSeries System Handbook i5/OS Version 5 Release 3 October 2005 -- Draft*, GA19-5486
- ▶ *IBM System i5, eServer i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006*, SG24-2155
- ▶ *IBM System Storage DS6000 Series: Architecture and Implementation*, SG24-6781
- ▶ *IBM System Storage DS6000 Series: Copy Services in Open Environments*, SG24-6783
- ▶ *IBM System Storage DS8000 Series: Architecture and Implementation*, SG24-6786
- ▶ *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788
- ▶ *PCI, PCI-X, PCI-X DDR, and PCIe Placement Rules for IBM System i Models*, REDP-4011
- ▶ *IBM TotalStorage DS6000 Series: Performance Monitoring and Tuning*, SG24-7145
- ▶ *IBM TotalStorage DS8000 Series: Performance Monitoring and Tuning*, SG24-7146

## Other publications

These publications are also relevant as further information sources:

- ▶ *Backup and Recovery*, SC41-5304
- ▶ *Backup, Recovery, and Media Services for iSeries*, SC41-5345

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM System Storage DS6000 Information Center  
<http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp>
- ▶ *IBM System Storage DS6000: Introduction and Planning Guide*  
[http://www-1.ibm.com/support/docview.wss?rs=1112&context=HW2A2&dc=DA400&q1=ssg1\\*&uid=ssg1S7001072&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=1112&context=HW2A2&dc=DA400&q1=ssg1*&uid=ssg1S7001072&loc=en_US&cs=utf-8&lang=en)
- ▶ IBM System Storage DS6000 Technical Notes  
[http://www-1.ibm.com/support/search.wss?q=ssg1\\*&tc=HW2A2&rs=1112&dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&dtm](http://www-1.ibm.com/support/search.wss?q=ssg1*&tc=HW2A2&rs=1112&dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&dtm)
- ▶ IBM System Storage DS8000 Information Center  
<http://publib.boulder.ibm.com/infocenter/ds8000ic/index.jsp>
- ▶ *IBM System Storage DS8000: Introduction and Planning Guide*  
[http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DA400&q1=ssg1\\*&uid=ssg1S7001073&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DA400&q1=ssg1*&uid=ssg1S7001073&loc=en_US&cs=utf-8&lang=en)
- ▶ IBM System Storage DS8000 Technical Notes  
<http://www-1.ibm.com/support/search.wss?dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&tc=HW2B2&rs=1113&dtm>
- ▶ *IBM System Storage DS8000 User's Guide*  
[http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DA400&q1=ssg1\\*&uid=ssg1S7001163&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DA400&q1=ssg1*&uid=ssg1S7001163&loc=en_US&cs=utf-8&lang=en)
- ▶ IBM Systems Information Centers  
<http://publib.boulder.ibm.com/eserver/>
- ▶ *IBM TotalStorage Enterprise Server Introduction and Planning Guide*  
[http://www-1.ibm.com/support/docview.wss?rs=503&context=HW26L&dc=DA400&q1=planning&uid=ssg1S7000003&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=503&context=HW26L&dc=DA400&q1=planning&uid=ssg1S7000003&loc=en_US&cs=utf-8&lang=en)
- ▶ Support for System Storage DS6800  
[http://www-1.ibm.com/support/search.wss?q=ssg1\\*&tc=HW2A2&rs=1112&dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&dtm](http://www-1.ibm.com/support/search.wss?q=ssg1*&tc=HW2A2&rs=1112&dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&dtm)
- ▶ Support for System Storage DS8100  
<http://www-1.ibm.com/support/search.wss?dc=DB500+D800+D900+DA900+DA800+DA600+DB400+D100&tc=HW2B2&rs=1113&dtm>
- ▶ *VPN Implementation (IBM System Storage DS8000)*  
[http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DB500&uid=ssg1S1002693&loc=en\\_US&cs=utf-8&lang=en](http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DB500&uid=ssg1S1002693&loc=en_US&cs=utf-8&lang=en)



## How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks, IBM Redpapers, Hints and Tips, draft publications, and Additional materials, as well as order a hardcopy of IBM Redbooks or CD-ROMs, in this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

Archived

Archived

# Index

## Numerics

12X loop 87

## A

abnormal IPL 292  
access density 99  
adapter 308  
    port 308  
ADDRMTJRN command 6  
administrative domain 11, 18  
application  
    data consistency 292  
    response time 93  
application resilient CRG 11  
ASP 13–14, 101, 104, 117, 199  
    group 14  
asynchronous background copy process 311  
asynchronous Peer-to-Peer Remote Copy 254  
auxiliary storage pool. *See* ASP

## B

backup node 10–11, 17  
Backup Recovery and Media Services (BRMS) 24, 436  
backup server 280, 302  
    issues 292  
backup system 8  
basic user ASP 14  
batch job, duration 93  
batch window 111  
blocksize 91, 121, 171, 176

## C

Capacity Magic 76  
    example 77  
change recording bitmap 323  
check or restart Global Mirror session 335  
clone 38  
    creating 21  
    i5/OS 38  
cluster 9  
    node 10–12  
    version 10, 12  
cluster component 10  
    backup node 11  
    cluster node 11  
    primary node 11  
    replicate node 11  
cluster resource group. *See* CRG  
Cluster Resource Services 10  
clustering  
    disaster recovery 9  
    high availability 9

Collection Services 450  
Commercial Processing Workload (CPW) 110  
commitflash command 324  
Compute Intensive Workload (CIW) 110  
consistency group 23, 305  
    FlashCopy target volume 323  
copied volumes 291  
Copy Services 253–254, 279, 301  
    Toolkit 27  
copy sets 25  
creating  
    Global Copy relationship 310, 321  
    Global Mirror environment 303  
    Global Mirror session 314  
    Metro Mirror relationship 287, 291, 296  
    Peer-to-Peer Remote Copy path 282, 304  
CRG 10–11  
    application resilient 11  
    data resilient 11  
    device resilient 11  
    exit program 11  
    peer 11  
cross-site mirroring 14  
    benefits 15, 17  
    concept 14  
    limitations 18

## D

DASD 18, 20, 24, 254–255, 280, 302  
    Copy Services 254  
    FlashCopy 255  
    Global Mirror 302  
    Metro Mirror 280  
    space 18, 20, 24  
        copying 18  
        regular copies 20  
data mining 15, 18  
data port services, LIC 16  
data resilient CRG 11  
Dedicated Service Tool. *See* DST  
device  
    CRG 11, 14  
    description 12  
    domain 10, 12  
    resilient CRG 11  
direct access storage device. *See* DASD  
direct attach of external storage, example 32  
disaster recovery 9, 40  
    using a copy 23  
disk arm 89, 95, 102  
    i5/OS workload 95  
disk configuration 196–197, 206, 221–223  
    error 261  
disk CPU utilization 470

- disk drive 96, 100, 170, 177
  - disk operations per second 96
- disk drive modules 96
- disk encryption 14
- Disk Magic 89, 94, 104, 106, 111, 113–117, 121–122, 125–126, 128–129, 132–133, 137, 143, 150–152, 156–157, 161, 163, 172–173, 175–176, 426
  - cache values 163
  - size DS8000 114
- disk pool 12–13, 62, 201–202, 205, 208–210
- disk response time 93
- disk service time 121, 130, 148–149, 162–163, 177
- disk space 90, 96, 102, 104, 106, 112–113, 122–125, 128–129, 131, 133–134, 141–142, 147, 158, 162–163, 174, 177
- disk subsystem 93, 110, 118, 120–124, 129, 135, 138, 151, 154–157, 173–174
  - ESS1 158, 162
  - icon 119
  - iSeries icon 138
  - model 141
- disk unit 57, 61–62, 170, 195–196, 200, 206, 211, 220
  - connection 62
  - recovery 196, 220
- display
  - disk configuration
    - capacity 197, 207
    - status 223
  - Global Mirror relationship 316
  - Metro Mirror status and properties 288
  - Peer-to-Peer Remote Copy path 286, 309
  - volumes assigned to Global Mirror 315
- D-mode IPL 191
- DS CLI 24, 28, 47, 185, 254, 256–257, 269–270, 277, 282, 293, 295, 301, 304, 314, 321, 329, 331, 333, 338
  - installing 254
  - setting up 254
- DS command-line interface. *See* DS CLI
- DS Storage Manager 24
- DS8000 352, 365
- DS8000 Storage Manager 341
- DST 62, 191, 193, 195, 220, 265
  - environment 220
  - main menu 194
- duration of batch job 93

## E

- end relationship
  - FlashCopy 314
  - Global Copy 312
  - Metro Mirror 288
- ESS 800 150, 154–155, 161
  - correct values 155
- expansion unit 17
- extent 73
- extent pool 95, 101, 128, 143–144, 183
  - needed capacity 143
- external load source unit 185–186, 215
  - path protection 215
- external storage 94–95, 181–182, 185, 215

- 7xx and prior only supported SCSI attachment 182
- main memory 97
- sizing iSeries workload 94
- supported environment 181
- System i5 customer plans 99

## F

- failback, Global Copy 329
- failover 14–15, 17
  - to backup 46
- failoverpprc command 46, 292, 296
- Fast Reverse Restore process 326
- FB 351
- FC adapter 89, 95, 99, 101, 103, 122, 125, 133, 136, 141, 150, 154–155, 159, 174, 425
- FC port 95
- Fibre Channel (FC) 95, 97, 182–183, 185
  - attachment 182
  - Disk Controller 182, 186
  - I/O adapter 214
  - request proceeds 92
- fixed block storage 183
- fixed block volumes 283
- FlashCopy 254, 259, 266, 268–269, 272, 282, 304, 343, 363, 388
  - i5/OS cloning 38
  - pair 258, 287, 312
    - consistency formation 324
    - consistency group 326
  - relationship 258, 313–314, 323, 326–327
  - system backup 20
  - target volume 303, 326
  - warm 20

## G

- geographic mirror 15, 17
- geographic mirroring 15
  - with external disk 47
- Global Copy 254, 303
  - in reverse direction 329
  - target storage 305
- Global Mirror 254, 282, 304, 314, 335, 345
  - asymmetrical replication 45
  - examples 43
  - full system replication 44
  - session 303, 316
  - switchable IASP replication 44
- Global Mirror session 22
- graphical user interface. *See* GUI
- GUI 12

## H

- HABP 5
- hard disk drive (HDD) 131
- hardware management 185
- Hardware Management Console. *See* HMC
- Hardware Service Manager 62
- HDD (hard disk drive) 131

- utilization 131–132, 135
- high availability 40
- high availability (HA) 9, 14, 17
  - software-based solution 5
- High Availability Business Partner. *See* HABP
- High Availability Solutions Manager 39, 66
- HMC 185–186, 191, 254
- host adapter 58, 92
  - PCI space 92
  - planned number 123
- host port 124, 141–142, 155, 158–159, 174
- HSL loop 58

## I

- I/O
  - enclosure 308
  - latency 41
  - operation 91–93, 99
  - operations 90
  - operations per second 470
  - port 307
  - property 188
  - rate 103
  - request 91–92
  - tower 57, 86, 214
- I/O adapter (IOA) 57, 61, 170, 185, 189
- I/O per second 98–99, 103, 133, 136, 163
- I/O processor 182
- I/O processor (IOP) 17, 40, 54, 57, 62, 92, 182, 216–217, 280, 302, 470
  - maximal I/O per second 98
- i5/OS 111, 137–140, 143, 146
  - cloning 38
  - current cache usage 140
  - current configuration 140
  - extent pool 143
  - mirroring 57
  - performance reports 111
  - separate extent pool 143
  - workload 90, 137
    - following configurations 141
    - sizing DS 90
- i5/OS Performance Tools 106
- i5/OS workload
  - disk operations 96
  - disk operations per second 96
- IASP 12, 14–15, 17, 27, 104, 199
- IBM Redbooks publications Web site 481
- IBM System Storage DS CLI. *See* DS CLI
- IBM Systems Director Navigator for i5/OS 12
- implementation
  - FlashCopy 3, 253, 337, 347, 363, 373, 387, 423, 449
  - Global Mirror 301
  - Metro Mirror 279
- incremental FlashCopy 313, 389
- independent auxiliary storage pool. *See* IASP
- independent user ASP 14
- initial copy 316
- initial program load. *See* IPL
- input/output processor. *See* I/O processor

- input/output. *See* I/O
- Integrated File System 4
- internal disk 113, 184
  - current cache usage 129
  - current read cache percentage 131
  - current workload 119
  - host workload 113
  - I/O rate 153
- internal load source unit 185
  - RAID disk controller 185
- inter-switch link (ISL) 58
- IOA 5760 103
- IOP-based Fibre Channel 103
- IOP-less Fibre Channel 54, 98–100
- IP connection 11
- IPL 62, 184, 189–190, 195, 220, 222–223, 261, 281
- iSeries 11–12, 14, 57–59, 61, 181–184, 195, 199–200, 205–206, 208–210
  - Copy Services toolkit 27
  - Model 825 150
- iSeries Copy Services Toolkit 66
- iSeries Navigator 12, 195, 199, 205, 208–209, 211
  - cluster management GUI 14
- ISL (inter-switch link) 58

## J

- journal entry 6–7
  - efficient transport 7
  - replication performance 7
- journal receiver 5–7, 23
- journal volumes 320
- journaling
  - local 6
  - remote 6

## K

- keylock position 190

## L

- LIC 6, 193
- Licensed Internal Code. *See* LIC
- load source 185, 188–189, 191, 193–195, 197, 220–223, 261, 290, 302–303, 356
  - IOP 185, 191, 215–216, 218
  - mirror state 262
- load source unit 182–185, 189, 191, 193, 215, 217, 219–220, 222–223, 255–256, 280, 302–303
- local journaling 5–6
  - high availability solution 6
- local port 307
- local site 290, 303
  - business application 280, 302
  - normal disaster protection mode 298
- logical unit 61, 195, 205, 211, 214
- logical unit number. *See* LUN
- logical volume 58, 92, 100, 181, 183–184, 195, 206
- LUN 38, 61, 92, 95, 99–100, 126, 132, 161, 177, 184, 215–216, 218–219

## M

- Machine Interface (MI) 8
- main memory 90–91
  - block of data 97
- manual IPL 189
- master Global Mirror session manager 323
- master storage 305
  - server 321
- master storage unit 22
- memory page 90
- Metro Copy in the original direction 333
- Metro Mirror 21, 254, 282, 287–290, 293, 295, 297, 304, 343
  - examples 41
  - full system replication 41
  - in original direction 298
  - in reverse direction 293
  - relationship 290
  - switchable IASP replication 42
- MI (Machine Interface) 8
- migration
  - external mirrored load source to boot load source 37
  - internal drives to external storage including load source 36
- mirror copy 16
  - data integrity 18
- mirror state 262
- mirrored load source unit 262
  - unprotected LUNs 356
- mirrored load source volume 262
- Model 825 150
- monitored resource entries 12
- multifunction IOP 222
- multipath
  - connection 61
  - I/O 57
  - information 264
  - volume 206, 208, 211, 264

## N

- needed number 96
- node 9, 428
- non-configured unit 195, 197, 203, 220

## O

- object type 8, 17
  - authorization lists (\*AUTL) 8
  - job descriptions (\*JOBID) 8
  - libraries (\*LIB) 8
  - non-journaled 8
    - replicating 8
  - programs (PGM) 8
  - user spaces (\*USRSPC) 8
- operating system 220, 265–266, 292–293, 297–298, 328–329, 333, 335
  - automatic installation 220, 265
- OS/400
  - mirroring 184
  - V5R3 59

## P

- page fault 91
- partition
  - name 186
  - profile name 186
  - property 187, 189–190
- PCI I/O card placement rules 86
- PCI-X IOP 185
- peak period 111
- peer CRG 11
- Peer-to-Peer Remote Copy Extended Distance (PPRC-XD) 254
- Peer-to-Peer-Remote Copy. *See* PPRC
- percent full and utilized 470
- performance expectations 85
- Performance Explorer 471
- performance report 106, 112
  - reported values 161
- Performance Tools 163
- Performance Tools reports 459
- permanent object 5
- physical capacity 67
- physical hardware resource 264
- planned number 124, 142
- PPRC 254, 281, 284, 286, 290, 303–307, 310
  - synchronous 254
- primary node 10–12, 17
- production copy 15
- production server 280–281, 292, 297, 302–303, 328, 335
  - physical hardware resource 297
- protected LUN 184
- protected mirror 74

## Q

- quiesce for CopyServices 432

## R

- RAID-5 rank 126
- rank group 101
- raw capacity 67
- read-only target volume 320
- Real-time manager 474
- recovery domain 11
- recovery point of objective 41
- recreate FlashCopy relationship 327
- remote journaling 5, 7
  - high availability solution 7
- remote site 281, 322
  - copied volumes 321
- removal
  - Global Mirror session 315
  - Peer-to-Peer Remote Copy path 286, 309
- replicate node 11
- report
  - Resource 470
  - System 466
- resilient applications 12
- resilient data 12
- resilient devices 12

- Resource report 470
- response time 470
  - application 93
  - disk 93
- resume
  - Global Mirror session 319
  - suspended Metro Mirror 290
- reverse direction
  - FlashCopy 326
  - Mirror Copy 285, 308
  - starting Metro Mirror 281, 293
- revertflash command 324, 326
- role swap 47
- RPO 41

**S**

- SAN LUNs 38
- Save While Active. *See* SWA
- SDD (Subsystem Device Driver) 57
- select Property 474
- service time 470
- service tool 62, 196, 220
- sessions 25
- setup
  - FlashCopy environment 256
  - System i5 253
  - System Storage DS6000 253
  - System Storage DS8000 253
- single point 58
- single-level storage 5, 90
- sizing 85
- SLIC (System Licensed Internal Code) 4, 262
- source
  - logical subsystem (LSS) 304, 321
  - server 255–256
  - storage system 311
  - volume 258, 320, 329
    - initial asynchronous background copy 288
- source site tracking 15
- source system 15–17, 24
  - DASD efficiency 7
  - journal entries 7
  - journal receivers 6
  - production copy 15
  - reader job 6
- space
  - total copy 24
- space efficient FlashCopy 19, 62
- spare disk drive 68
- sparing rule 68
- Storage Allocation Method 270
- Storage Management Console 254
- storage pool striping 74
- storage unit 474
  - Serial Number information 474
- StorWatch Expert 107
- subordinate storage server 305, 321
- Subsystem Device Driver (SDD) 57
- suspension
  - Global Mirror session 318
  - Metro Mirror 289
    - new Global Copy relationship 321
    - new Metro Mirror relationship 291, 296
- SWA 20
  - switchable IASP 14
  - switchback system 293, 328
  - swtchover 11, 14–15, 17
    - environment 290
    - from production to backup 45
  - swpprc command 45
  - synchronous delivery mode 7
  - synchronous mirroring 21–22
  - synchronous PPRC 254
  - system ASP 14, 18, 57, 184, 196–197
  - System i Copy Services Toolkit 39, 47
  - System i5
    - all disk in external storage 32
    - external storage HA environments 40
  - System Licensed Internal Code 4, 91
  - System Licensed Internal Code (SLIC) 262
  - System report 104, 106, 111, 115, 150, 163, 168–170, 175, 466
    - expert cache storage pools 175
    - Interactive CPU utilization 169
  - System Service Tools 196
  - System Storage Productivity Center 24
  - System Storage Productivity Center for Replication 24

**T**

- target
  - logical subsystem 304
  - server 255
  - storage system 311
  - volume 310, 320
- target site tracking 15, 18
- target system 6–8, 15, 17
  - journal receivers 6
  - main storage 7
  - mirror copy 15
  - receiver job 6
  - user profile 8
- TCP/IP setting 292
- temporary object 5
- terminate
  - previous Global Copy relationship 321
  - previous Metro Mirror relationship 291, 296
- test environment 302
- tgtread option 333
- TPC 340
- transfer size 91

**U**

- unprotected LUN 184
- user ASP 14
- user profile 18, 24

**V**

- virtual address 90

vital product data . *See* VPD  
volume group 101, 216, 361  
volumes  
    availability on local site 295, 331  
    availability on remote site 291, 321  
VPD 262

## **W**

wait time 470  
warm flash 21  
warm FlashCopy 20  
worldwide node name. *See* WWNN  
write penalty 96  
WWNN 284, 309

## **X**

XSM. *See* cross-site mirroring

## **Z**

zoning 425





**Redbooks**

# **IBM System Storage Copy Services and IBM i: A Guide to Planning and Implementation**

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages







# IBM System Storage Copy Services and IBM i:

## A Guide to Planning and Implementation



**Redbooks®**

**Discover IBM i 6.1 and DS8000 R3 Copy Services enhancements**

This IBM Redbooks publication describes the implementation of IBM System Storage Copy Services with the IBM System i platform using the IBM System Storage Disk Storage family and the Storage Management GUI and command-line interface. This book provides examples to create an IBM FlashCopy environment that you can use for offline backup or testing. This book also provides examples to set up the following Copy Services products for disaster recovery:

- ▶ Metro Mirror
- ▶ Global Mirror

**Learn how to implement Copy Services through a GUI and DS CLI**

**Understand the setup for Metro Mirror and Global Mirror**

The newest release of this book accounts for the following new functions of IBM System i POWER6, i5/OS V6R1, and IBM System Storage DS8000 Release 3:

- ▶ System i POWER6 IOP-less Fibre Channel
- ▶ i5/OS V6R1 multipath load source support
- ▶ i5/OS V6R1 quiesce for Copy Services
- ▶ i5/OS V6R1 High Availability Solutions Manager
- ▶ System i HMC V7
- ▶ DS8000 R3 space efficient FlashCopy
- ▶ DS8000 R3 storage pool striping
- ▶ DS8000 R3 System Storage Productivity Center
- ▶ DS8000 R3 Storage Manager GUI

**INTERNATIONAL  
TECHNICAL  
SUPPORT  
ORGANIZATION**

**BUILDING TECHNICAL  
INFORMATION BASED ON  
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)**