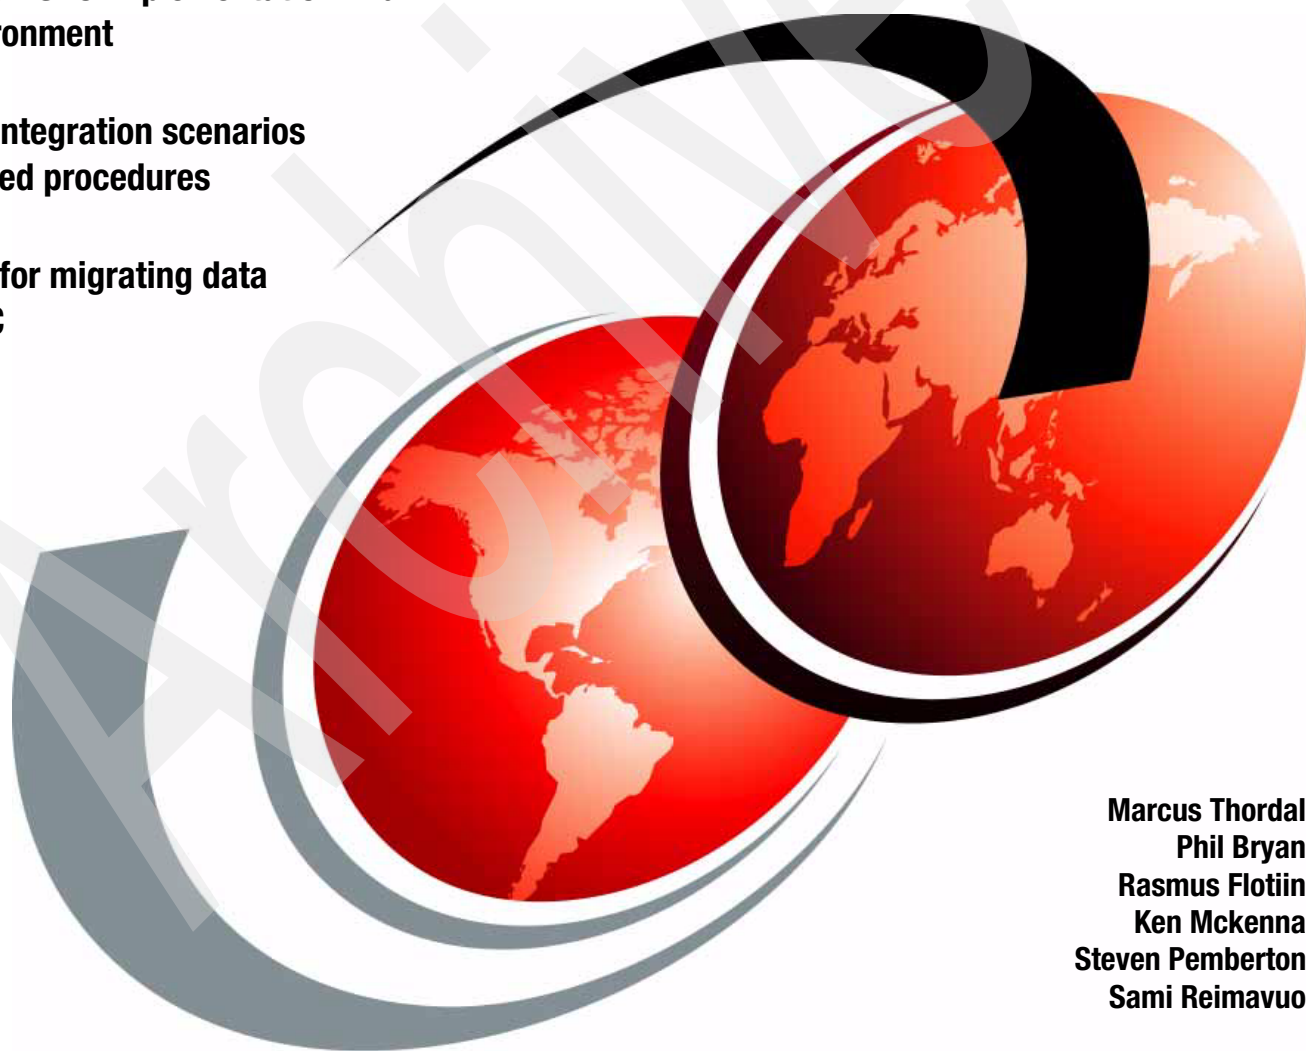


Implementing the SVC in an OEM Environment

Planning an SVC implementation in an
OEM environment

Practical integration scenarios
and detailed procedures

Guidance for migrating data
to the SVC



Marcus Thordal
Phil Bryan
Rasmus Flotiin
Ken Mckenna
Steven Pemberton
Sami Reimavuo

Redbooks



International Technical Support Organization

Implementing the SVC in an OEM Environment

March 2007

Archived

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

First Edition (March 2007)

This edition applies to IBM System Storage SAN Volume Controller 4.1.

© Copyright International Business Machines Corporation 2007. 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

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this book	ix
Become a published author	xi
Comments welcome	xi
Part 1. Planning for implementation in an OEM environment	1
Chapter 1. Introduction	3
1.1 Data storage challenges and SVC benefits	4
1.2 Supported environments	5
1.2.1 Storage controllers	6
1.2.2 Host operating systems	6
1.2.3 SAN Fabric	8
1.2.4 Solution assurance review	8
1.3 Unsupported environments	8
1.3.1 RPQ process	9
1.3.2 Multipath device driver compatibility	9
1.4 Planning and sizing considerations	10
1.4.1 Overview	11
1.4.2 Planning and sizing	11
1.4.3 Fabric	12
1.4.4 Zoning	14
1.4.5 Configuration of the SVC prior to actual migration	15
1.4.6 SVC physical installation	16
1.5 Using Disk Magic to plan an SVC implementation	16
Chapter 2. Methods for data migration	17
2.1 Overview of data migration methods	18
2.2 Storage controller-based	18
2.3 Host-based	20
2.4 Temporary storage-based	22
2.5 SVC Image Mode	23
2.6 SVC Managed Mode	25
2.7 Conclusion	27
2.8 More information	28
Chapter 3. Storage migration with the SVC	29
3.1 Generic SVC migration plan	30
3.1.1 Confirm support for your environment	30
3.1.2 Preparing the SAN for storage virtualization	32
3.1.3 Preparing the environment	33
3.1.4 Migration to SVC managed storage	34
3.1.5 Migration of SVC managed storage between controllers	34
3.1.6 Migration from SVC managed storage (optional)	34
3.2 Types of SVC volume migration	35
3.2.1 Migrating a VDisk between MDGs	35

3.2.2 Migrating a VDisk to image mode	37
3.2.3 Migrating a VDisk between I/O groups	39
3.2.4 Migrating VDisk extents (within an MDG)	41
3.2.5 Migrating extents off an MDisk that is being removed from the MDG.	43
3.2.6 Non-standard migration types	45
3.3 Monitoring the migration process	47
3.4 Summary of the SVC migration process	48
Part 2. Implementation in an OEM storage environment	51
Chapter 4. Implementing the SVC with DS4000	53
4.1 Scenario with DS4800.	54
4.1.1 Overview of integration steps	54
4.1.2 Scenario description	55
4.2 Preparing the SAN for storage virtualization	56
4.2.1 Confirm support for your environment.	56
4.2.2 Record the SVC FC port numbers	57
4.2.3 Prepare the SAN fabric	58
4.2.4 Prepare the DS8000 storage controllers.	59
4.2.5 Prepare the SAN Volume Controller	59
4.3 Storage virtualization.	61
4.4 Configure the OEM controller	71
4.5 Migration to the SAN Volume controller	74
4.5.1 Configure the SAN fabric	74
4.5.2 Configure the SAN Volume Controller	74
4.5.3 Configure the host system	76
4.6 Move SVC managed volumes between storage controllers	78
Chapter 5. Integrating the SVC with HDS	81
5.1 Integrating HDS Lightning with the SVC	82
5.1.1 Scenario description	82
5.1.2 Confirm support for your environment.	82
5.1.3 Preparing the SAN for storage virtualization.	84
5.1.4 Begin storage virtualization.	89
5.2 Integrating HDS TagmaStore USP with the SVC	93
5.3 Integrating HDS Thunder with the SVC	93
5.3.1 Scenario description	93
5.3.2 Summary of migration process	94
5.3.3 Preparing the SAN for storage virtualization.	94
5.3.4 Begin migration	98
Chapter 6. Integrating SVC with EMC	101
6.1 Integrating the EMC Symmetrix with the SVC	102
6.1.1 Scenario description	102
6.1.2 Summary of the SVC integration process.	103
6.1.3 Preparing the SAN for storage virtualization.	104
6.1.4 Begin integration	107
6.1.5 Integration into the SVC environment	117
6.2 Integrating the EMC FC4700 with the SVC.	120
6.2.1 Scenario description	120
6.2.2 Summary of integration process	121
6.2.3 Preparing the SAN for storage virtualization.	122
6.2.4 Integration of SVC managed storage	127
6.2.5 Migration to the SAN Volume controller	153

Chapter 7. Integrating SVC with HP.	161
7.1 Integrating the SVC in a SAN with a Hewlett-Packard EVA	162
7.2 Scenario description	162
7.3 Summary of integration process	163
7.4 Preparing the SAN for storage virtualization	164
7.5 Begin integration	168
7.6 Integration into the SVC environment	185
Chapter 8. Integrating SVC with NetApp.	189
8.1 Integration with NetApp V-Series	190
8.1.1 Scenario description	190
8.1.2 Summary of the SVC migration process	190
8.1.3 Confirm support for your environment	192
8.1.4 Preparing the SAN for storage virtualization	193
8.1.5 Migration to the SAN Volume controller	201
8.1.6 Move SVC managed volumes between storage controllers	208
8.2 Integration with NetApp FAS (back-end to SVC)	210
8.2.1 Scenario description	211
8.2.2 Summary of the SVC migration process	212
8.2.3 Confirm support for your environment	213
8.2.4 Preparing the SAN for storage virtualization	215
8.2.5 Migration to the SAN Volume controller	231
8.2.6 Move SVC managed volumes between storage controllers	234
Chapter 9. SAN boot migration to the SVC.	237
9.1 Scenario description	238
9.1.1 Confirm support for your environment	238
9.1.2 Preparing the SAN for storage virtualization	240
9.1.3 Move SVC managed volumes between storage controllers	245
9.2 Unsupported OEM controllers	247
9.3 Host cluster migration	248
Chapter 10. Troubleshooting	249
10.1 Overview	250
10.2 Solving errors	250
10.2.1 SAN Volume Controller	250
10.2.2 AIX	250
10.2.3 HP-UX	251
10.2.4 EMC CLARiiON	251
10.2.5 HDS Thunder	251

Appendix A. SVC administration commands	253
SAN Volume Controller commands	254
Appendix B. SAN zoning changes	257
B.1 Basic steps before zoning	258
B.2 Basic zoning steps	259
B.2.1 Using the CLI for the IBM b-type family	260
B.2.2 Using the CLI for the IBM m-type family	262
B.2.3 Using the CLI for Cisco	265
Related publications	269
IBM Redbooks	269
Other publications	269
Referenced Web sites	269
How to get IBM Redbooks	270
Help from IBM	270
Index	271

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

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

IBM Redbooks (logo) ™
xSeries®
zSeries®
AIX®
DS4000™
DS6000™
DS8000™

Enterprise Storage Server®
FlashCopy®
Freelance Graphics®
IBM®
Lotus®
Redbooks™
S/390®

System i™
System p™
System Storage™
Tivoli®
TotalStorage®
1-2-3®

The following terms are trademarks of other companies:

Network Appliance, FilerView, NetApp, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Solaris, StorageTek, Sun, Sun StorEdge, SunOS, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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™ document describes how to plan and perform integration of the IBM SAN Volume Controller (SVC) in existing OEM storage environments.

The first part of the book covers topics that apply to any OEM environment. It begins with a brief overview of the benefits of using the SVC to manage virtualized storage, and identifies the large number of currently supported environments. Planning and sizing considerations are discussed, along with data migration methodologies, generic SVC migration plans, and the types of SVC migration.

The second part of the book covers integration of the SVC in some of the most commonly used OEM storage environments from IBM, HDS, EMC, HP and NetApp. Practical scenarios are described for each environment, and step-by step details and code samples guide you through the planning, preparation, and performance of the migration.

The final section presents an overview of SVC administration commands and a detailed discussion of SAN zoning, along with sample code for performing the actual zoning changes.

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, San Jose Center.

Marcus Thordal is an IBM Certified IT Specialist in IBM Global Services, Denmark, and has been with IBM since 1998. He is BCFP, MCSE certified, and his areas of expertise include Open Systems storage solutions, in particular design and implementation of Storage Virtualization, Information Lifecycle Management and Disaster Recovery solutions. Marcus holds a Bachelor of Science in Engineering from The Technical University of Denmark. He has co-authored 8 previous IBM Redbooks on IBM Total Storage products and solutions.

Phil Bryan is a pre-sales Storage Specialist in the UK. He has 15 years of experience in the storage industry. He holds a degree in Electronic Engineering from Lancaster University. His areas of expertise include disk and storage development and testing, in particular with IBM 7133 (SSA) disk storage and SAN Volume Controller. Phil has been working in the storage pre-sales team for the last year, advising customers on storage solutions and strategies.

Rasmus Flotiin has been with IBM since 2001. He works for IBM as part of his education as a Data Technician at Technical Education Copenhagen. He has primarily worked with the architecture and implementation of Enterprise solutions, and he has great experience with implementing the SAN Volume Controller and IBM Disk systems from projects like Stockholm Road Pricing and the Carlsberg Transition.

Ken Mckenna is a pre-sales Storage Software consultant working in the IBM Software Group in the UK. He has been with IBM for 5 years and has worked in the SAN and Storage environment for 8 years. His areas of expertise include SAN and storage hardware and software solutions design and implementation.

Steven Pemberton is a storage pre-sales technical specialist with Synergy (an IBM business partner) and is based in Melbourne, Australia. Steven has considerable experience with the IBM Total Storage virtualization and Tivoli® Storage Manager products. He has consulted

and instructed throughout the Asia Pacific region. Steven is the co-author of two previous IBM Redbooks.

Sami Reimavuo is an IT Specialist in IBM Global Technology Services, Finland. He has 10 years of experience in IT, in a wide variety of infrastructure technologies ranging from networks and information security to system administration, storage systems, and virtualization concepts. His current focus is storage and host virtualization technologies and their application to infrastructure simplification and consolidation.



The team (left to right): Marcus, Ken, Phil, Steven, Rasmus, and Sami

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

Alex Ainscow, Iain Bethune, Dave Carr, Steve Chesney, Gregory Dalton, Mark Elliott, Robin Finlay, Huw Francis, Stephen Garraway, Geoff Lane, Richard Mawson, Rob Nicholson, Nick O'Rourke, Lucy Raw, Bob Sankey, Bill Scales, Dave Sinclair, Bruce Smith, Sukhi Sohal, Steve White, Barry Whyte, Tony Wrather
IBM Hursley

Thorsten Hoss, Torsten Knobloch, Chris Seiwert, Louise Schillig
IBM Mainz

Arun Batish, Sanjay Gandhi, Joseph Morabito
IBM Beaverton

Bill Wiegand
IBM Advanced Technical Support

Dorothy Faurot
IBM Raleigh

John Gressett, Mike Griesse
IBM Rochester

Chris Saul
IBM San Jose

Doris J Konieczny
IBM Tucson

Sharon Wang
IBM Chicago

Jens Toftgaard Nielsen, Anders Trier Rasmussen, Thomas M Steenholdt
IBM Denmark

Pirkka Palm, Juha Ekström, Pauli Rämö
IBM Finland

Charlotte Brooks, Tom Cady, Alison Chandler, Emma Jacobs, Leslie Parham, Deanna Polm,
Sangam Racherla, Jenifer Servais, Jon Tate, Erica Wazewski
IBM ITSO

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbooks document 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

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- Send your comments in an e-mail to:

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

Archived



Part 1

Planning for implementation in an OEM environment

This part describes the planning phase prior to implementation of the IBM System Storage™ SAN Volume Controller (SVC) in an OEM storage environment.

Archived

Introduction

This chapter presents an overview of the benefits of integrating the IBM System Storage SAN Volume Controller (SVC) in an existing storage environment. It also identifies the supported environments and how to work with currently unsupported environments, and describes how to approach the planning and sizing of a migration project.

An initial level of SVC skills is assumed. For a comprehensive discussion of the SVC, consult the latest SVC book: *IBM System Storage SAN Volume Controller*, SG24-6423.

<http://www.redbooks.ibm.com/abstracts/sg246423.html?Open>

1.1 Data storage challenges and SVC benefits

The motivation for implementing the IBM System Storage SAN Volume Controller (SVC) and thereby storage virtualization in your storage environment is usually based on the major challenges related to storage as seen from a business perspective.

From a business perspective, clients are faced with the following major storage challenges:

- Managing storage growth

Storage needs tend to grow at a very high rate. It is not uncommon for an organization to purchase storage subsystems based on an assumption that the capacity will satisfy their requirements for a period of three to five years, only to discover that their needs have outgrown available storage in just a year or two. The organization faced with such data growth will need to add additional storage capacity, often beyond the footprint of their existing storage controllers. They may try to meet their changing storage requirements by expanding current storage subsystems in chunks, if this is possible, or by buying different types of disk subsystems to match storage capacity needs. Either approach can strain IT budgets.

- Increased complexity

As storage capacity and the number of disk subsystems grow to meet an organization's needs, the cost of storage administration grows respectively. The organization's environment might include a variety of server platforms and operating systems, as well as Storage Area Networks (SAN) with multiple and diverse storage subsystems from different vendors, resulting in a complex storage infrastructure and consequently increased cost and risk with respect to storage administration.

- Maintaining availability

The requirements for data availability and data protection tend to be extremely stringent – not only from a business point of view, but for regulatory reasons as well, particularly in industries like banking and health care. These demands can be met by a wide range of advanced data replication and data protection features such as FlashCopy®, Metro/Global Mirror, and other copy services.

However, with the storage growth rate, and often no matching increase in storage budget, organizations are facing the challenge of managing more storage with minimal or no additional staff, while the requirements for availability, data protection, and maybe data mobility are increased. Changing the storage infrastructure in a traditional storage environment to accommodate the needed storage growth and evolving business requirements tends to be disruptive and leads to planned outages. Even if an outage is planned it is often very costly, either because the business is directly affected during the outage or because complex changes are performed outside normal office hours and require the involvement of expensive technical specialists and elaborate advance planning.

You can overcome these major business challenges by implementing the IBM System Storage SAN Volume Controller, thereby virtualizing storage.

From a high-level point of view, the SVC facilitates this by logically combining the storage capacity from multiple heterogeneous storage systems into a single reservoir of storage capacity and one central point of management with the following benefits:

- Single point of operation

After initial implementation and storage resources are given to the SVC, all management can be performed from one single interface. This improves efficiency and productivity, and minimizes the need to acquire and maintain skills on various storage controllers.

- Single point of access for the hosts

The multipath driver used to access storage devices on the SVC is available free of charge.

The host can have data placed on various storage controllers (storage tiers) without any concern, regardless of whether the multipath driver is supported on the different systems, since access is provided through the SVC.

- Application of Advanced Copy Services across heterogeneous storage controllers

The SVC Advanced Copy Services can be applied across the storage infrastructure independent of what copy services are available on the various storage controllers in the environment.

- Changes can be made non-disruptively

While storage is virtualized, and thereby the host applications are insulated from the physical infrastructure, changes to the back-end storage controllers can be performed without any disruption to the hosts.

New business needs, for example increased capacity needs, can be accommodated faster due to a simplified and flexible storage infrastructure.

Virtualization solutions can be implemented in the storage network, in the server, or in the storage device itself. The IBM storage virtualization solution is SAN-based, which helps to allow for a more open virtualization implementation. Locating virtualization in the SAN, and therefore in the path of input/output (I/O) activity, helps to provide a solid basis for policy-based management. The focus IBM places on open standards means its virtualization solution supports freedom of choice in storage device vendor selection.

1.2 Supported environments

This section provides information about the supported environments that the IBM System Storage SAN Volume Controller can be implemented in. It also covers the options available should you have hardware that is not currently supported by the SAN Volume Controller or that is at a different firmware or software level than that listed on the support list.

In the following sections we have identified the current supported hardware and recommended software levels for SAN Volume Controller V4.1. This information is based on the available information at the time of writing.

Refer to the SAN Volume Controller support pages on the IBM Web site to get the latest information about what is supported or has been tested.

<http://www.ibm.com/storage/support/2145>

The SAN Volume Controller covers host support including Windows NT®, 2000, and 2003; AIX®; SUSE SLES 8, SLES 9, and Red Hat EL AS Linux®; Novell NetWare 6.5 with clustering; VMWare ESX; Sun™ Solaris; and HP-UX.

The SAN switch support is also broad and includes the IBM TotalStorage® SAN Switches and members of the Brocade, McDATA, CNT, and Cisco families of SAN Switches and Directors.

IBM System Storage SAN Volume Controller Support of disk subsystems includes the IBM SystemStorage DS4000™, DS6000™, and DS8000™ series servers and the IBM TotalStorage Enterprise Storage Server®. It also supports several models of Hitachi Thunder, Lightning, and Tagma Store; EMC Clariion and Symmetrix; and Hewlett-Packard

StorageWorks Modular Arrays and EVA. Furthermore, the support includes the new IBM SystemStorage N series and NetApp® FAS.

Future releases of the IBM SystemStorage SAN Volume Controller will continue to grow the portfolio of supported disk subsystems, SAN switches, HBAs, hosts, and operating systems.

1.2.1 Storage controllers

Complete and current details regarding supported controllers and supported firmware releases can be found at the IBM System Storage SAN Volume Controller support site.

IBM

- ▶ DS4000
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_DS4000
- ▶ DS6000
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_DS6000
- ▶ DS8000
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_DS8000
- ▶ ESS
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_ESS
- ▶ N-Series
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_IBM_N-Series

Hewlett-Packard

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HP

Hitachi Data Systems

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HDS

EMC

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_EMC

SUN StorEdge

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_StorEdge

StorageTek

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_StorageTek

NetApp FAS

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_NetApp

1.2.2 Host operating systems

Begin by verifying the supported level for each OS and for multipath drivers. This information is available at the IBM System Storage SAN Volume Controller support site:

<http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002865>

Then consult the list of supported HBAs and firmware for each OS as follows:

IBM

- ▶ AIX on System p™
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_pSeries
- ▶ AIX on System i™
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_iSeries

Hewlett-Packard

- ▶ HP-UX 11.0
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HP-UX_11.0
- ▶ HOP_UX 11i_v1
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HP-UX_11i_v1
- ▶ HOP_UX 11i_v2
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HP-UX_11i_v2
- ▶ Tru64
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_Tru64
- ▶ OpenVMS
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_HP-UX_11i_v2

Microsoft

- ▶ NT 4.0
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_NT
- ▶ Windows® 2003 and Windows 2004 Server
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_Windows

Red Hat

- ▶ Red Hat EL Advanced Server 2.1 U7
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_RHEL2.1
- ▶ Red Hat EL Advanced Server 3.0 U5
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_RHEL3.0
- ▶ Red Hat EL Advanced Server 4.0 U3
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_RHEL4.0

SUN

- ▶ SunOS™ 5.8/Solaris™ 8
- ▶ SunOS 5.9/Solaris 9
- ▶ SunOS 5.10/Solaris 10
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_Solaris_8

Novell

- ▶ Novell NetWare 6.5
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_Novell

SuSE Linux Enterprise Server

- ▶ SLES 8 United Linux 1.0 SP4
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_SLES_8
- ▶ SLES 9 Novell/SuSE
http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_SLES_9

VMware

- ▶ VMware

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_VMware

1.2.3 SAN Fabric

Refer to the SAN Volume Controller V41. configuration requirements and guidelines to ensure that your SAN fabric environment is supported. See:

<http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002858>

Brocade

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_Brocade

Cisco MDS

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_MDS

CNT

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_CNT

McDATA

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864#_McData

1.2.4 Solution assurance review

For all IBM System Storage SAN Volume Controller solutions proposed by IBM or an IBM Business Partner, a solution assurance review will be carried out. The purpose of this is to ensure that the proposed configuration will meet the customer requirements and fit into the existing customer infrastructure. It includes a review of the existing customer SAN to confirm that the product can be installed successfully. In addition, all software and firmware levels for other hardware that the customer wants to connect to the SAN Volume Controller will be checked and validated against the IBM SAN Volume Controller supported hardware list and recommended software list.

This review will identify any hardware and software or firmware that is not currently supported. Ideally it is better to identify unsupported hardware as early as possible during the engagement with IBM or the business partner so that there is more time to resolve this.

1.3 Unsupported environments

The scope and breadth of the SAN Volume Controller supported hardware list is rapidly increasing. If you have a very unusual environment, however, it may be that your specific configuration is not currently supported. There are two distinct situations:

- ▶ The hardware you intend to use is unsupported.
- ▶ You are using supported hardware, but the firmware or software level it is using is not listed on the supported hardware list. This may be because the firmware or software is:
 - At a newer level than listed
 - At an earlier level than listed

The request for price quotation (RPQ) process enables you to request support for your specific environment and should be used in all of these situations.

If you have firmware or software at a level that is more recent than that listed on the SAN Volume Controller supported list, it is likely that testing will be underway already. You should still request that the RPQ process be initiated, as described in the next section.

1.3.1 RPQ process

In order to ensure that your environment will function properly at all times, even during error conditions, you should submit an RPQ to IBM.

Contact your IBM Sales person, pre-sales technical specialist, or business partner, providing as much information about your environment as possible – in particular the hardware you wish to use. For example, for a new host system, provide details about the processor type, memory, host bus adapter, and operating system. This should include information regarding the firmware and software levels of each component. For a new storage subsystem, you should provide all the details regarding the vendor and model type, together with the current firmware levels being used.

The sales team passes this information to the IBM SAN Volume Controller development team, which reviews the RPQ request to determine whether it can be supported and in what time frame. Please note that not all requests can be fulfilled.

Following the review, the development team will respond back to you through the sales group. Should your request be approved, you will be notified of the target date for the support.

1.3.2 Multipath device driver compatibility

The SAN Volume Controller can be used with a number of different vendor multipath device drivers. If your existing multipath driver is not compatible with SAN Volume Controller, you can download the IBM Subsystem Device Driver (SDD) free of charge from the IBM Web site. See:

<http://www.ibm.com/storage/support/2145>

Go to the recommended software levels page and use the link to the driver specific to your host platform.

What is a multipath driver?

Without a multipath driver, a host system will see many instances of the same SAN device. This is because through the SAN there are many routes to the same logical unit number (LUN). A multipath driver provides an interpretation and recognizes that a number of raw devices that the host sees are in fact the same LUN. Most multipath drivers provide:

- ▶ Enhanced data availability
- ▶ Dynamic I/O load balancing across multiple paths
- ▶ Automatic path failover protection
- ▶ Path selection policies for the host system

Supported vendor multipath drivers

Table 1-1 shows the supported vendor multipath drivers that can be used with SAN Volume Controller.

Table 1-1 Supported vendor multipath drivers

Operating system	Driver
Windows 2000/2003	Symantec VERITAS Storage Foundation v4.3 with Qlogic HBAs only
Solaris 5.8/5.9	Symantec/VERITAS Volume Manager 4.1
Solaris 5.9/5.10	MPxIO
HP-UX 11.0	Service Guard V11.14 -PVlinks
HP-UX 11i v1 and v2	Service Guard V11.16 -PVlinks
Novell NetWare	Native Multipath Driver
VMWare	Native Multipath Driver

Multipath driver coexistence with IBM Subsystem Device Driver (SDD)

Certain multipath drivers have been tested and can coexist on the same host running the IBM Subsystem Device Driver. Certain configurations for the host are required under these conditions. Refer to the Configuration Guide and Host Attachment Guide for this information.

Table 1-2 shows the operating systems and related drivers supported to coexist with Subsystem Device Driver.

Table 1-2 Multipath drivers tested for coexistence with IBM Subsystem Device Driver (SDD)

Operating System	Driver
Windows NT	RDAC
Windows 2000	RDAC
Windows 2003	RDAC
AIX 4.3.3/5.1/5.2/5.3	RDAC
Solaris 5.8	Symantec/VERITAS Volume Manager 4.1
Solaris 5.9	Symantec/VERITAS Volume Manager 4.1
HP-UX 11i v1	Service Guard V11.16, PVlinks

1.4 Planning and sizing considerations

In this section we discuss some of the items you should consider before you install an SVC solution. By creating a planned and scalable solution you will save time and money. When implementing an SVC you should always take into consideration how many hosts you have at the point of the introduction of the SVC and the number you will have in the future. You should also consider what features you will use and how your fabric is built.

1.4.1 Overview

The first step in any planning process should always be making sure that you have the most complete and up to date information. When planning to implement the IBM System Storage SAN Volume Controller, this means consulting the following documents for the latest updates:

V4.1.x Configuration Requirements and Guidelines

<http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002858>

IBM System Storage SAN Volume Controller Configuration Guide Version 4.10, SC26-7902

http://www-1.ibm.com/support/docview.wss?rs=591&context=STCFKTH&context=STCFKTW&dc=DA400&q1=english&uid=ssg1S7001413&loc=en_US&cs=utf-8&lang=en

IBM System Storage SAN Volume Controller V4.1.0 - Planning Guide, GA32-0551

http://www-1.ibm.com/support/docview.wss?rs=591&context=STCFKTH&context=STCFKTW&dc=DA400&q1=english&uid=ssg1S7001405&loc=en_US&cs=utf-8&lang=en

IBM System Storage SAN Volume Controller V4.1.0 - Installation Guide, GC26-7900

http://www-1.ibm.com/support/docview.wss?rs=591&context=STCFKTH&context=STCFKTW&dc=DA400&q1=english&uid=ssg1S7001411&loc=en_US&cs=utf-8&lang=en

IBM System Storage SAN Volume Controller V4.1.0 - Host Attachment Guide, SC26-7905

http://www-1.ibm.com/support/docview.wss?rs=591&context=STCFKTH&context=STCFKTW&dc=DA400&q1=english&uid=ssg1S7001416&loc=en_US&cs=utf-8&lang=en

1.4.2 Planning and sizing

This section describes some basic issues to consider when you plan a migration.

Supported environment

Make sure that your environment is supported. Carefully review the following items to ensure that they meet the requirements defined in the previous section and detailed on the referenced Web sites.

Applications

Determine whether any of your applications have dependency on a disk's unique identifier (UID). When migrating to or from an SVC the disk UID will change.

Operating system

Your operating system should be at a recommended level. If not, you should upgrade it prior to the migration.

Multipathing software

Make sure that your multipathing software is at a recommended level. If not, it should be upgraded.

Host Bus Adapter (HBA)

Check whether your Host Bus Adapter (HBA) firmware is at a supported level. Should the HBA not be at a level that is supported, you should upgrade it before a migration.

Switches

The switch should also be upgraded if needed. Make sure that the model is supported.

SAN Volume Controller

The SVC should also be upgraded prior to a migration. For information about supported hardware and recommended software levels:

<http://www-03.ibm.com/servers/storage/support/software/sanvc/>

Hardware

Before you start a migration you should make sure that all your hardware is in optimal state. Your environment should have redundant power, and there should be no physical errors and no errors pending in the error log.

1.4.3 Fabric

When creating a fabric it is always wise to consider the number of ports in your switch. Consider not only how many you need now, but also how many you will need in the future.

Here are a few recommendations to keep in mind when planning and sizing your fabrics:

- ▶ Use two fabrics for redundancy.
- ▶ Plan a flexible SAN. The cost of moving a SAN will usually surpass the initial cost of a switch with more ports. The majority of vendors offers switches that do not have to have all the ports activated. You can buy an extra license for them when they are needed.
- ▶ Avoid mixing switches from different vendors since not all switches are compatible and you may be at risk of not being able to implement special features across different vendor models.

The local or remote fabric must not contain more than three Inter-Switch Links (ISL) hops in each fabric. Any configuration that uses more than three ISL hops is not supported. When a local fabric is connected to a remote fabric for Metro Mirror, the ISL hop count between a local node and a remote node must not exceed seven. Therefore, some ISL hops can be used in a cascaded switch link between local and remote clusters if the internal ISL count of the local or remote cluster is less than three. If all three allowed ISL hops are used within the local and remote fabrics, the local remote fabric interconnect must be a single ISL hop between a switch in the local fabric and a switch in the remote fabric. The SAN Volume Controller supports the use of distance-extender technology, including DWDM (Dense Wavelength Division Multiplexing) and FCPIP extenders, to increase the overall distance between local and remote clusters. If this extender technology involves a protocol conversion, the local and remote fabrics should be regarded as independent fabrics, limited to three ISL hops each. The only restriction on the interconnection between the two fabrics is the maximum latency that is allowed in the distance extender technology.

For more information, refer to the “*Fibre-channel switches and interswitch link*” topic in *IBM System Storage SAN Volume Controller Configuration Guide Version 4.10. SC26-7902*.

There are some basic things to avoid when designing your SAN fabric. For instance, the nodes may not be separated by Inter-Switch Links (ISLs) as shown in Figure 1-1.




Figure 1-1 SVC divided from its Quorum disk by ISLs

The SVC nodes (in the same SVC cluster) should not be separated from each other by ISLs. In such a configuration the ISLs are considered as a single point of failure, while the loss of connection could mean that the SVC nodes in an SVC cluster lose connection with each other.

If your environment requires the SVC nodes to be divided by ISLs, you must ensure that a link failure does not cause nodes to fail when there are ISLs between nodes, thus making it necessary to use a redundant configuration. This is illustrated in Figure 1-2.

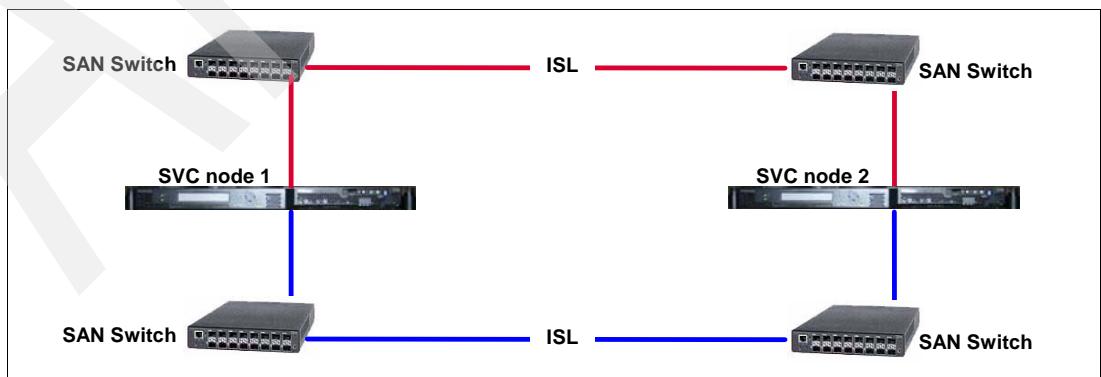


Figure 1-2 SVC nodes divided by ISLs

1.4.4 Zoning

The number of hosts in your fabric will influence how the zoning should be done.

When planning the zoning you should also take into consideration the number of paths from the host HBA to the SVC. The number of paths should not be greater than eight. This is because of the way SDD handles the connections to the SVC.

Zoning rules depend on the number of hosts

The configuration rules for host zones are different depending upon the number of hosts that will access the cluster. For smaller configurations of fewer than 64 hosts per cluster, the SAN Volume Controller supports a simple set of zoning rules that enable a small set of host zones to be created for different environments. For larger configurations of more than 64 hosts, the SAN Volume Controller supports a more restrictive set of host zoning rules. Zoning that contains host HBAs must not contain either host HBAs in dissimilar hosts, or dissimilar HBAs in the same host that are in separate zones. Dissimilar hosts means that the hosts are running different operating systems or are different hardware platforms; however, different levels of the same operating system are regarded as similar.

Clusters with fewer than 64 hosts

For clusters with fewer than 64 hosts attached, zones that contain host HBAs must contain no more than 40 initiators, including the SAN Volume Controller ports that act as initiators. A configuration that exceeds 40 initiators is not supported. A valid zone can be 32 host ports plus 8 SAN Volume Controller ports. Place each HBA port in a host that connects to a node into a separate zone. You should also include exactly one port from each node in the I/O groups that are associated with this host. This type of host zoning is not mandatory, but is preferred for smaller configurations.

Note: If the switch vendor recommends fewer ports per zone for a particular SAN, the more strict rules that are imposed by the vendor takes precedence over the SAN Volume Controller rules.

Clusters with 64 to 256 hosts

For clusters with 64 to 256 hosts attached, each HBA port in a host that connects to a node must be placed into a separate zone. In this separate zone, you must also include exactly one port from each node in the I/O groups that are associated with this host. The SAN Volume Controller does not specify the number of host ports or HBAs that a host or a partition of a host can have. The number of host ports or HBAs is specified by the host multipathing device driver. The SAN Volume Controller supports this number; however, it is subject to the other configuration rules that are specified here. To obtain the best performance from a host with multiple Fibre Channel ports, the zoning must ensure that each Fibre Channel port of a host is zoned with a different group of SAN Volume Controller ports. To obtain the best overall performance of the subsystem and to prevent overloading, the workload to each SAN Volume Controller port must be equal. This can typically involve zoning approximately the same number of host Fibre Channel ports to each SAN Volume Controller Fibre Channel port.

Clusters with 256 to 1024 hosts

For clusters with 256 to 1024 hosts attached, the SAN must be zoned so that each HBA port in a host that connects to a node can only see one SAN Volume Controller port for each node in the I/O group that is associated with the host. If you have 1024 hosts, each host must be associated with only one I/O group and each I/O group must only be associated with up to 256 hosts.

For more information see the “San Volume Controller Zones” topic in *IBM System Storage*

SAN Volume Controller Configuration Guide Version 4.10, SC26-7902, at:

http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp?topic=/com.ibm.svc.web.doc/svc_zoningguidelines_22n6zd.html

Note: You should always consult the IBM System Storage SAN Volume Controller Restrictions documentation for your particular level of SVC code:

<http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002866>

1.4.5 Configuration of the SVC prior to actual migration

MDisks

When creating mdisks and assigning them to the SVC, take into consideration the type of MDisk.

Depending on the numbers of underlying disks, we recommend that you do not mix MDisks with different attributes within one MDisk group.

If you have many physical disk within a single disk system, it is considered best practice to create one LUN within every single disk enclosure. This will reduce the time it takes to rebuild a parity, and it allows for load balancing on the controllers of the disk system.

When creating LUNs we strongly recommend that you use disk redundancy, either RAID-5 or RAID-10, to prevent a single disk failure from bringing down an MDisk group.

MDisk groups

Consider the following guidelines when you create MDisk groups:

- ▶ Allocate your image mode VDisks between your MDisk groups.
- ▶ All MDisks that are allocated to a single MDisk group should be of the same RAID type. This ensures that a single failure of a physical disk in the storage subsystem does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you must not mix RAID types. The performance of all VDisks is reduced to the lowest performer in the group.
- ▶ If you intend to keep the VDisk allocation within the storage exported by a storage subsystem, ensure that the MDisk group that corresponds with a single subsystem is presented by that subsystem. This also enables nondisruptive migration of data from one subsystem to another subsystem and simplifies the decommissioning process if you want to decommission a controller at a later time.
- ▶ Except when you migrate between groups, you must associate a VDisk with just one MDisk group.
- ▶ An MDisk can be associated with just one MDisk group.

For more information see the “MDisk groups and VDisks” topic in *IBM System Storage SAN Volume Controller Configuration Guide Version 4.10, SC26-7902*.

Extent size

When creating extent sizes, bear in mind that they cannot be easily changed. The size of the extent will also influence the total amount the SVC can handle.

Unless there is a special need for small Vdisks we recommend that all MDisk groups be made with 512 MB extents.

For more information see *IBM System Storage SAN Volume Controller Configuration Guide Version 4.10*, SC26-7902.

IO groups

There are some simple rules that you should bear in mind when planning the numbers of IO groups.

Note: One IO group is two SVC nodes in a cluster. The maximum numbers of IO groups is four, which would be eight SVC nodes in one cluster.

One IO group can handle 1024 VDisks. Table 1-3 shows how the maximum number of VDisks scales per SVC cluster with the numbers of IO groups.

Table 1-3 Total number of VDisks per cluster determined by number of IO groups

IO groups / VDisks	1	2	3	4
1024	X			
2048		X		
3072			X	
4096				X

1.4.6 SVC physical installation

Your SVC should be installed by certified personnel. This will ensure that the basic foundation is in order when you start using the SVC.

1.5 Using Disk Magic to plan an SVC implementation

IBM provides a tool that is able to give an estimate of the IO performance of your setup. This means that you will be able to plan your environment and give a good prediction of how your environment will perform. When you have an idea about how your environment will behave in its current setup, you will be able to predict how your environment will behave when you add more hosts or disk subsystems. Disk Magic will simulate your complete environment in detail, down to how your disk drawers are configured. Disk magic is able to take your data and put it into a graphical overview using Lotus® 1-2-3® 97. If you do not have the correct release of 1-2-3, you can still use Disk Magic but without graphics. Disk Magic can export data in a format that you can import into other graphic packages such as Freelance Graphics®, PowerPoint®, releases of 1-2-3 that are not supported by Disk Magic, or Excel®.

When you have a complete overview of your current setup you will be able to add more nodes to your setup to see how it will effect performance. By doing this you will be able to predict quite accurately when your system will start to loose performance, thus allowing you to intervene before you experience degraded performance. Another big advantage of this is that you will be able to plan your intervention for a time and place that will have minimum impact on your system.

For more information about Disk Magic contact your IBM representative or follow this link if you are able to:

http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s380&geoID=A11&prodID=IBM%20eServer%20And%20TotalStorage%20Products&docID=SSD5D00689DF4_moreinfo



Methods for data migration

In the previous chapter we explained how the fundamental demand for the ability to move data within a storage infrastructure can be achieved by performing infrastructure simplification. In this chapter we build on that, and on a conceptual level describe the different ways of migrating data from one environment to another.

2.1 Overview of data migration methods

With the traditional storage controllers we are limited to choosing from three basic migration methods:

- ▶ Storage controller-based
- ▶ Host-based
- ▶ Temporary storage-based

Note that these are generic descriptions, the actual implementation of these might differ depending on the vendor and the storage controllers in question.

With the SAN Volume Controller we have two new migration methods to use:

- ▶ SVC Image Mode
- ▶ SVC Managed Mode

Of these two, Image Mode is the transition phase, when moving to a storage infrastructure utilizing the SVC. This book is primarily focused on how to use Image Mode. Managed Mode is the easiest and most powerful method and can be used after the initial migration to the SVC.

In this chapter we describe each of the five data migration methods in detail and list the pros and cons of each.

2.2 Storage controller-based

In this method the old and new storage controllers are temporarily connected to each other either directly (port-to-port) or through a Storage Area Network (SAN). Data is copied using the storage controller's own copy services.

Support and technical limitations are different for every scenario since there are no standards for cross-vendor data movement or even for a single vendor's cross-model data movement.

To guarantee data integrity, the data must not change while the migration is happening, so the host systems have to be disconnected for the duration of the data copy. On the other hand, even though a maintenance window is required, host-side changes are limited to simple, repetitive tasks like uninstalling and installing drivers.

The configuration of the old storage controller limits the way the new one can be configured. For example, logical unit size has to be the same, otherwise they cannot be copied directly from the old storage controller. Issues like this might affect performance and ease of administration.

In most cases direct migrations from one storage controller to another are complex operations that are done by a vendor's experts using specialized software and hardware tools for the job (such as IBM Piper for Open Systems). Therefore, they usually are part of a delivery project that includes new hardware, migration, training, and so forth. Figure 2-1 illustrates a storage controller-based migration.




Figure 2-1 Storage controller-based migration

Example migration procedure

To migrate data in a production environment using the controller-based method the following steps are performed:

1. Assess support and compatibility of different components carefully.
2. Physically install the new storage controller and connect it to SAN.
3. Define disk layout and LUN masking based on the old storage controller.
4. Shut down host systems.
5. Copy all the data from the old storage controller to the new one.
6. Change host storage zoning so hosts only see the new storage controller.
7. Start up the hosts, uninstall old drivers, and install new drivers.

Pros

The advantages of this method are:

- If carried out as a service from the vendor, the responsibility and risk is on the vendor.
- If carried out as a part of purchase of a new storage controller the price can be negotiated.
- Apart from possibly changing drivers, this method does not require host-side operations.
- Copying is done by the storage controllers so host-side performance is not affected.

Cons

The disadvantages of this method are:

- Since there are countless different combinations, assessing risks and scheduling is difficult.

- ▶ There is no guarantee that this method is even possible for any particular situation; it depends completely on the brand and model of the storage controller in question.
- ▶ Migration has to be done while the hosts are offline and therefore requires a maintenance window.
- ▶ Normally the migration has to be done for the whole storage controller at one time to limit the complexity, so it causes a maintenance window to several hosts simultaneously.
- ▶ Layout of the new storage controller may be limited by the layout of the old one.
- ▶ It is very likely that the old and new storage controllers are incompatible, so migration will probably require uninstalling and installing software on every server.

2.3 Host-based

In this method both the old and new storage controllers are presented to the host simultaneously, and the host takes care of copying the data from old devices to new ones.

Since the host itself copies the data, migration can be done incrementally, one host at a time, when it best suits the users. On the other hand, this means that every different storage controller–operating system–volume manager combination requires its own migration plan and scripts that have to be implemented and tested.

To use this method, it must be technically possible to connect both storage controllers to the host at the same time. This requires the drivers to be able to coexist and might require separate HBAs for both storage controllers, which increases the risk of failure and affects performance. Figure 2-2 depicts a host-based migration.

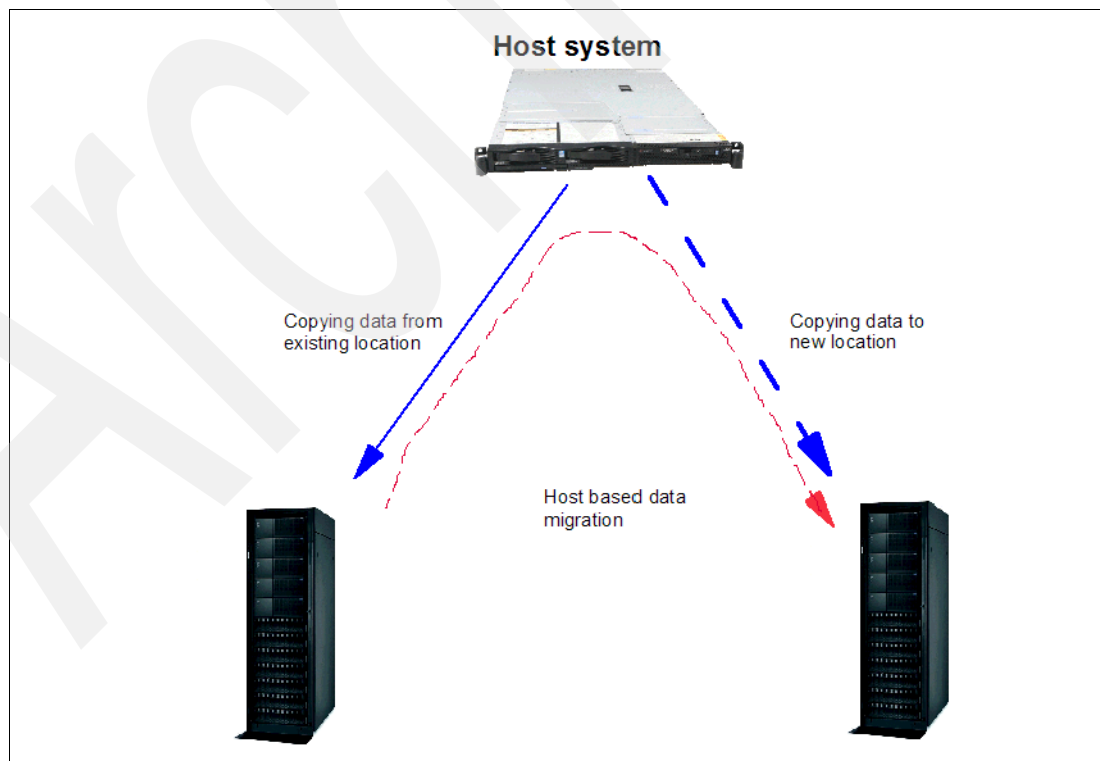


Figure 2-2 Host-based migration

Example migration procedure

To migrate data in a production environment using the host-based method the following steps are performed:

1. Physically install the new storage controller and connect it to SAN.
2. Define disk layout and LUN masking of the new storage controller.
3. Uninstall conflicting multipath drivers and install drivers for the new storage controller.
4. Change zoning to present the new storage controller to hosts.
5. Discover the new disks on hosts.
6. Migrate the data.
7. Undefine the old disks on the hosts.
8. Undefine zoning for the old storage controller.

Pros

The advantages of this method are:

- ▶ Migration can be done incrementally and one host at a time.
- ▶ Storage controllers do not have to know about each other so no deep vendor specialist knowledge is needed.
- ▶ The configuration of the new storage controller can be done freely and is not tied to the old storage controller.

Cons

The disadvantages of this method are:

- ▶ Switching drivers requires a maintenance window.
- ▶ Complexity. The actual implementation of this method depends on many details (host operating system, file system types, and so on), which results in so many different possibilities that custom migration scripts probably will have to be written and tested for every different host type.
- ▶ Slow speed. The host has to first read the data from the old storage controller and then write it to the new storage controller.
- ▶ Performance hit on the server. Since the host system moves the data, resulting I/O might use host resources like memory and CPU excessively.
- ▶ This method might require a long maintenance window if the copying has to be done offline (this depends on the host system).
- ▶ Possible reliability risk. Most hosts are connected with two HBAs, but the multipathing drivers cannot coexist, so during the migration both old and new environments are connected only with a single path.
- ▶ Possible compatibility problems. Drivers for the old and new environments need to be able to coexist on the host system.

2.4 Temporary storage-based

In this method the data is first copied off the host (for example, to a tape library), and copied back after the storage system has been replaced.

This method is simple because it often relies on the use of a data backup/restore solution that already exists and is tested. Therefore there are no complicated compatibility issues between the host and storage controller-based solutions.

Whether you can use this method or not depends on the amount of data to be migrated, the speed of the backup solution, and the maximum length of the maintenance window (downtime) acceptable to the business. For example, if the application only has a limited amount of data and is used only during normal business hours, the migration might be done at night or during the weekend. Figure 2-3 illustrates a temporary storage based migration.



Figure 2-3 Temporary storage based migration

Example migration procedure

To migrate data in a production environment using the temporary storage based method, the following steps are performed:

1. Shut down applications on the host.
2. Back up the data.
3. Remove old disk definitions and uninstall drivers.
4. Replace the storage system.
5. Create new disk definitions and install drivers.
6. Restore the data.
7. Start the applications.

Pros

The advantages of this method are:

- ▶ It is simple and easy to implement.
- ▶ The host does not need to see the old and new storage controller at the same time.
- ▶ Migration can be done at separate times for each host.

Cons

The disadvantages of this method are:

- ▶ It requires a long maintenance window.
- ▶ Slow migration. While the application is down the data has to be moved away and then moved back.
- ▶ It requires temporary storage from somewhere, typically from the backup tape library.

2.5 SVC Image Mode

This method involves moving hosts and storage controllers from stand-alone configurations to under SVC management. This is performed only once for each server and storage controller. Even though this method has many steps, most of them are nondisruptive to the host and can therefore be done while the host is online. This means the actual maintenance window will be as short as possible.

The idea is at first the disk the host accesses will be presented through the SVC as it is. This change can be performed very quickly. In general, the downtime equals the time needed for a restart of the host. Once the host is back online, the disk will be migrated to SVC Managed Mode (and to another storage controller, if needed) behind the scenes, while the host has data access.

This method requires operations on the storage controllers themselves, but it is fairly simple to carry out since advanced functions like copy services are not used. We only use basic features of storage controllers, like LUN mapping and LUN masking.

Host-side changes are limited to uninstalling old multipath device drivers and installing new ones. The maintenance window is short since data is not initially moved anywhere, it is just presented through the SVC.

Hosts can be moved one at a time because we can control the access with fabric zoning on the switches and LUN masking on the storage controllers. Figure 2-4 depicts an SVC Image Mode migration.

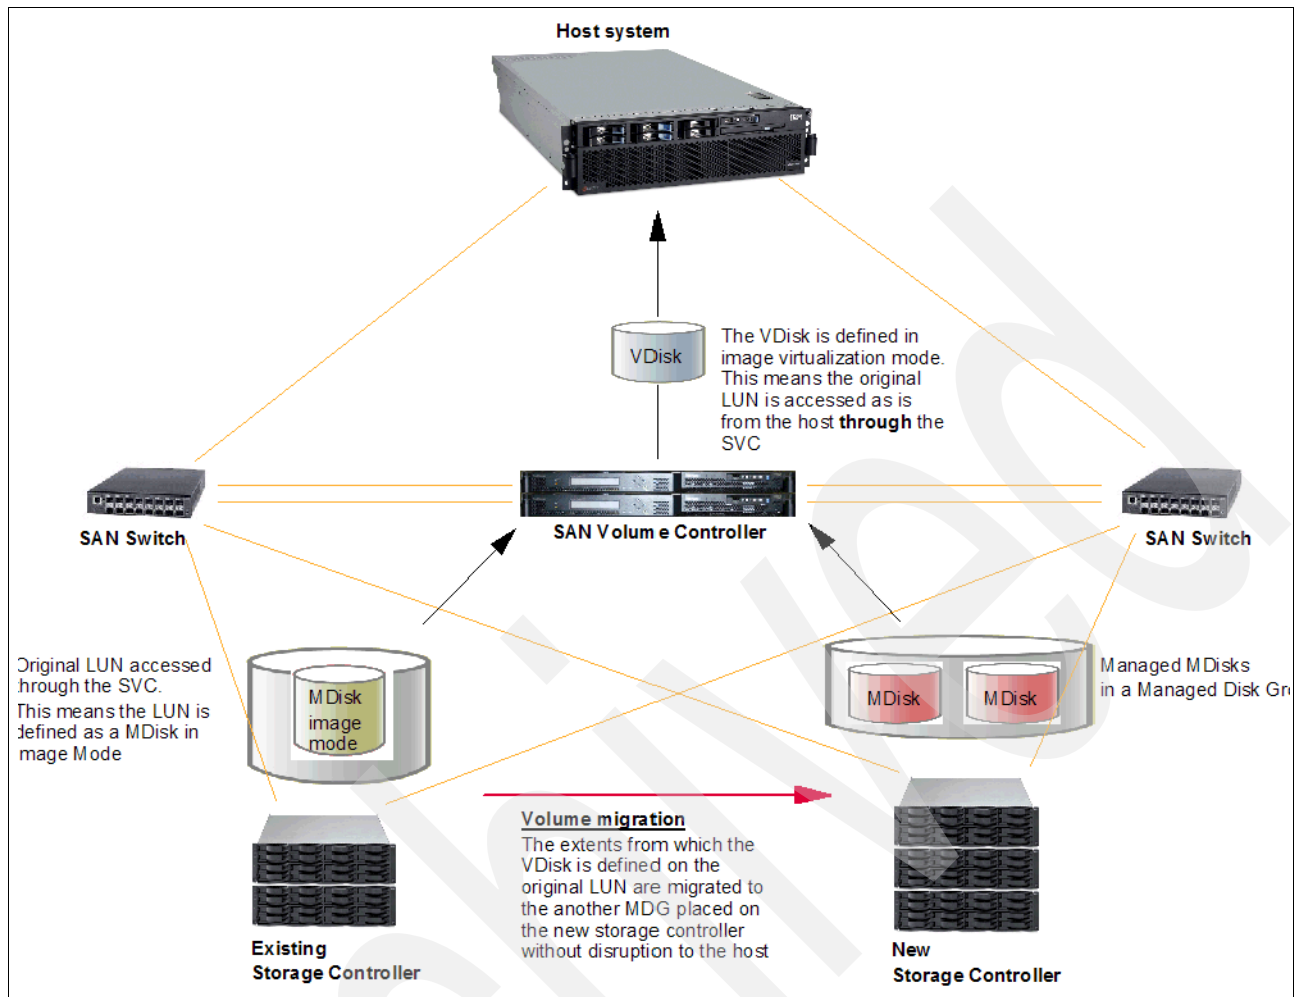


Figure 2-4 SVC Image Mode migration

Example migration procedure

To migrate data in a production environment using the SVC Image Mode based method the following steps are performed:

1. Host: Uninstall the old drivers and shut down. The maintenance window begins for the host and applications residing on it.
2. Switch: Change zoning so the host only sees SVC and SVC sees the storage controller.
3. Storage controller: Unmap disks from the host. Map disks to the SVC.
4. SVC: Define the host and the disks.
5. Host: Install SDD; disks should be visible as before. The maintenance window ends.
6. SVC: Migrate from Image Mode to Managed Mode.

Pros

The advantages of this method are:

- It is easy: No advanced features are needed from storage controllers.
- It is simple: Once the LUNs on the storage controller are presented to the host system through SVC, the rest of the procedure is always the same no matter which storage controller it is.

- ▶ It is generic: Only the handling of the storage controller differs from one migration to another migration.
- ▶ Migration is done by the SVC so host performance is not affected.
- ▶ The method is flexible: Migration can be done by one host at a time instead of a “big bang” where all hosts must be moved simultaneously.

Cons

The disadvantages of this method are:

- ▶ It requires a maintenance window to install drivers (and to uninstall existing ones).
- ▶ There are no other disadvantages, provided that the existing storage environment is supported by the SVC.

2.6 SVC Managed Mode

This method builds on the features of the SAN Volume Controller and storage virtualization. We can use this migration method once all the hosts have been connected to the SVC and all the storage controllers are managed by the SVC.

We have complete freedom to move, copy, or mirror data between any storage devices. We do not need to do any configuration changes to the host systems or storage systems, and everything happens while the hosts are online so there is no need for service downtime. All system administration actions needed are done from the single point of control – the SVC.

The ability to move data around freely brings new possibilities and advantages. For example, if performance requirements of the application change, the application data can be moved to a faster disk, or older data can be “retired” to slower and cheaper SATA-disk without downtime. Figure 2-5 illustrates an SVC Managed Mode migration.

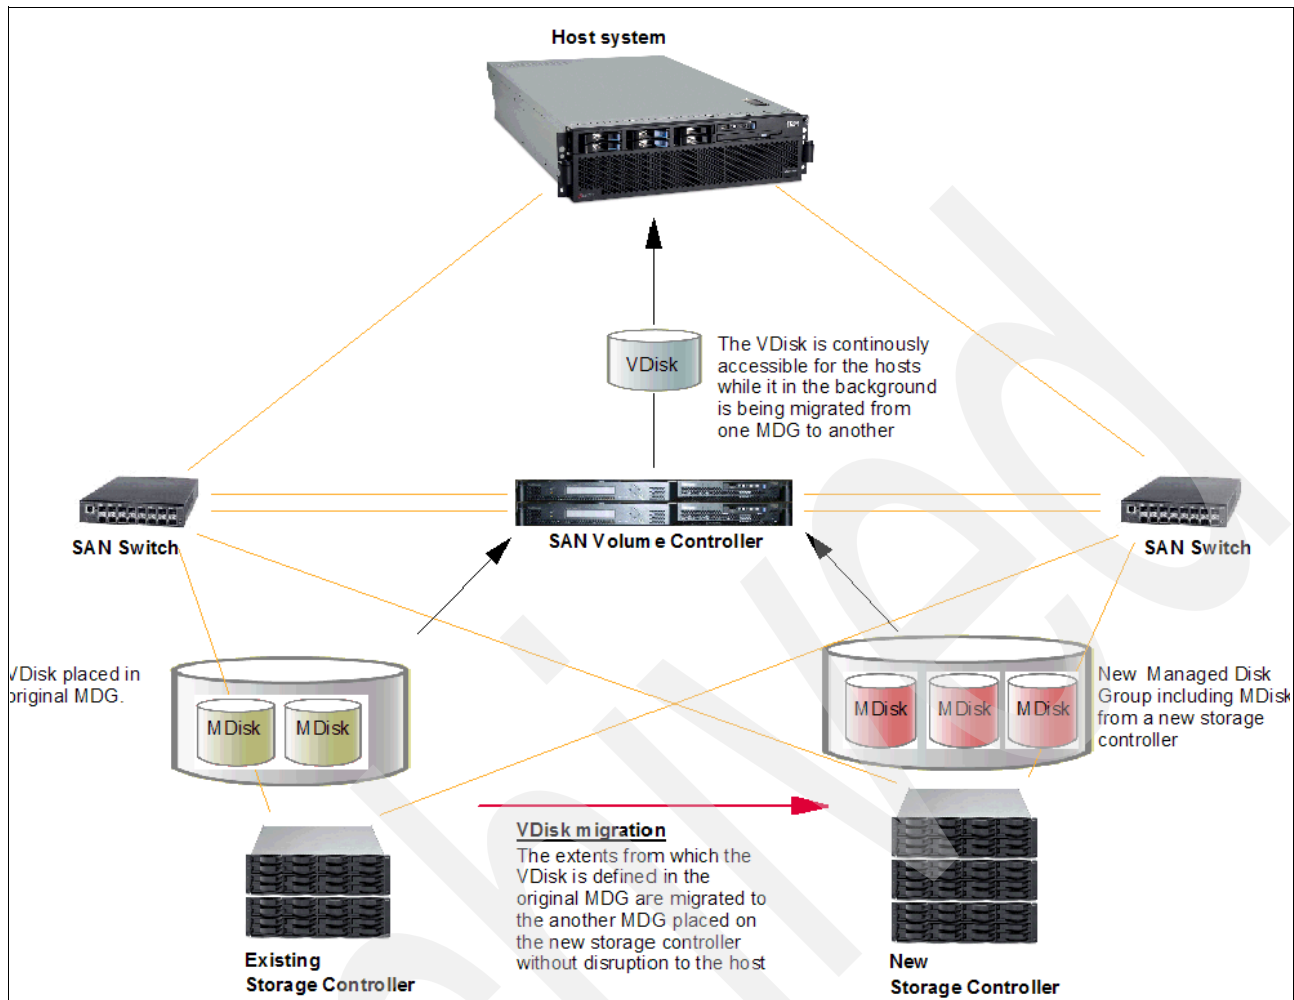


Figure 2-5 Managed Mode migration

Example migration procedure

Figure 2-5 shows the components that are involved in the different parts of the migration. To migrate data in a production environment using the SVC Managed Mode based method the following steps are performed.

1. Create the new MDG, including the MDisks to which you will migrate the existing VDisks (if they do not already exist).
2. Using the **SVC CLI** command **svctask migratevdisk** (or using the GUI), migrate the VDisk to the new MDG.
3. Monitor the progress of the VDisk migration using the SVC CLI command **svcinfolsmigrate**. Alternatively, you can monitor progress by selecting **Manage Progress** → **View Progress** from the GUI.

Note: The extent size must be the same for the MDG in order to facilitate VDisk migration between MDGs. In general it is recommended to use the same extent size for all MDGs on a SVC cluster.

4. Tell the SVC to migrate VDisks to the new storage controller.

Pros

The advantages of this method are:

- ▶ It is done online. No host-side changes are needed and no service downtime is incurred, so this method is nondisruptive in the production environment.
- ▶ It is easy. SVC handles data placement details (like striping) in the new storage controller.
- ▶ It is simple. Migration only requires knowledge of a few simple SVC commands or use of a simple GUI.
- ▶ It is generic. The same rules and commands apply every time, for every different host system and storage controller.
- ▶ It is scalable. Migrating the data of 100 host systems requires the same administrative effort as 10 or 1.
- ▶ It does not interfere with performance. Migration is done by the SVC so the host systems' performance is not affected.

Cons

If the support and features of the SVC itself are satisfactory, there are no disadvantages.

2.7 Conclusion

Migration and compatibility are not primary design features of storage controllers, and all three classic migration methods reflect this. They are either too complex, risky, or slow, and require lots of customized work, testing, and expert knowledge.

Storage controller-based

Due to complexity, the use of this method is limited to replacing decommissioned storage controllers with new ones. It can also be an alternative to host or temporary storage based methods, especially when the environment is so large (huge amounts of data, many different host types) or critical (strict usability requirements) that the other methods are not practical.

Host-based

Due to complexity and the customization needed, the practical use of this method is limited to replacing decommissioned storage controllers in smaller environments or as an alternative when storage controller-based migration is not possible due to compatibility issues.

Temporary storage-based

This method is only suitable for replacing decommissioned storage controllers in small non-critical environments due to the long service outage required.

SAN Volume Controller-based

SAN Volume Controller brings the same flexibility and ease of administration to Storage Area Networks that Logical Volume Management (LVM) brought to directly attached storage.

The aim of the SVC Image Mode is solely to get you to a situation where you can start using the SVC Managed Mode. If you use SVC merely as a one-time migration tool and remove it from the setup after the migration is done, you lose the use of SVC Managed Mode, which is the most powerful of all the presented methods.

2.8 More information

For a more detailed and technical description of migration with SVC, see Chapter 13, “Migrating to and from the SAN Volume Controller,” in *IBM System Storage SAN Volume Controller*, SG24-6423.

For more information about other IBM migration offerings, see the *IBM hardware-assisted data migration services* Web page found at:

http://www-03.ibm.com/servers/storage/services/featured/hardware_assist.html

Storage migration with the SVC

This chapter describes the process of implementing the SAN Volume Controller in an existing OEM storage environment. It presents a generic SVC migration plan, and other considerations when migrating storage with the SVC.

Topics covered in this chapter include:

- ▶ Generic overview of the SVC implementation process
- ▶ Different types of storage migration with the SVC
- ▶ Monitoring the storage migration process

3.1 Generic SVC migration plan

This section briefly describes the process of implementing an SVC in an existing OEM storage environment.

3.1.1 Confirm support for your environment

When implementing the SVC in an existing storage environment it is important to verify the SVC support for all of the included systems. The existing environment scenario can pose unique challenges for any implementation (not just for the SVC) because we do not have the opportunity to control the existing landscape. Therefore, it is important to validate the existing environment prior to the SVC implementation.

In the following discussion it is assumed that any new equipment in the proposed SVC solution is on the supported list. For details see 1.2, “Supported environments” on page 5.

Existing storage controller model

This section identifies some common factors to consider when confirming support for an existing vendor acquired storage controller.

Supported firmware levels

Generally, each of the supported storage controllers will be certified to a specific firmware level.

If the current firmware level is below the specified level, then you should upgrade to a supported level prior to implementation. If the current firmware level is above that specified, you should raise an RPQ to confirm support for that version prior to implementation. For details see 1.3.1, “RPQ process” on page 9.

Concurrent maintenance

Concurrent maintenance is the capability to perform I/O operations on the back-end storage controller while simultaneously performing maintenance on it.

Note: Not all storage controllers supported by the SVC are supported for concurrent maintenance. For details consult the IBM System Storage SAN Volume Controller Web site:

<http://www-03.ibm.com/servers/storage/support/software/sanvc/>

Sharing access to the storage controller

In most cases, access to the back-end storage controller can be shared between the SVC and other non-SVC hosts.

This configuration requires that some form of LUN security or controller partitioning is used to ensure that the SVC and other hosts do not share access to the same LUN on the back-end storage controller.

Quorum disk support

The SVC cluster identifies some disks as *quorum disks* and uses these disks to decide which cluster node takes control in certain failure situations.

Due to the limitations of some back-end storage controllers, not all storage controllers can be used to provide quorum-capable disks.

Advanced copy functions

Only a subset of the supported controllers are supported with their native copy services behind the SVC.

If the native copy services of your existing storage controller are supported with the SVC, then you will need to configure the host access using *cache-disabled* VDisks.

Note: With the current SVC code, there is no easy way to later change a cache-disabled VDisk to a VDisk in cache-enabled mode.

In order to take full advantage of the benefits of SVC storage virtualization, we strongly recommend that you consider using the SVC's own advanced copy functions, rather than those of the back-end storage controller.

Configuration settings

Some optional features and specific configuration options might need to be customized to work with the SVC. Each of the supported storage controllers has a list of configuration options that must be confirmed.

Existing SAN fabric switches

This section describes some common factors to consider when confirming support for any existing SAN switches.

Supported firmware levels

Generally, each of the supported SAN switches will be certified to a specific firmware level.

If the current firmware level is below that specified, then you should upgrade to a supported level prior to implementation. If the current firmware level is above that specified, raise an RPQ to confirm support for your version prior to implementation. For details see 1.3.1, "RPQ process" on page 9.

Switch zoning limitations

There are limitations in switch zoning for the SVC and the supported back-end storage controllers.

Typically, these limitations include:

- ▶ Two or more paths must be available to each LUN to avoid a single point of failure.
- ▶ Any LUNs presented by the storage controller to the SVC cluster must be accessible from all the SVC ports.

SAN topology limitations

The SVC is not supported on SANs that are composed of a *mesh* of switches. A mesh configuration is where four or more switches are connected together in a loop, but within that loop are paths that "short circuit" the loop.

Existing host operating systems

This section describes some common factors to be considered when confirming support for any existing hosts.

Multipath driver software

Host access to the SVC is supported with both the IBM-provided SDD multipath device driver, and a limited number of vendor acquired multipath device drivers.

Note: Not all combinations of supported hosts and multipath drivers are valid.

High availability clustering

There are some limitations on the type, configuration, and scale of the supported host clustering solutions. It is recommended that the host cluster be shut down when first migrating to the SVC managed storage.

Resolving supported environment issues

If some part of the existing environment is found to be non-compliant with the list of SVC-supported equipment, you need to rectify the situation before proceeding with the implementation.

Among the changes you might need to make are the following:

- ▶ Reconfigure the equipment to be compliant.
- ▶ Update the equipment microcode or firmware to a supported level.
- ▶ Update the host operating system to a supported level.
- ▶ Replace the non-compliant equipment prior to the SVC implementation.

In some cases it may be possible to obtain one-off support for non-compliant equipment through the RPQ process as described in 1.3.1, “RPQ process” on page 9.

If it proves impossible to create a supported environment then you should seriously reconsider whether to proceed with the SVC implementation.

Further information

See the following Web site for a list of supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Note: In some cases equipment that is not on the supported list may appear to function correctly with the SVC. However, this provides no expectation of future compatibility, and the equipment may prove incompatible in certain situations.

It is highly recommended that you remain within the bounds of the officially supported hardware and software combinations.

3.1.2 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Figure 3-1 shows the main steps involved in migrating from an existing storage environment to SVC managed storage on a new back-end storage controller.

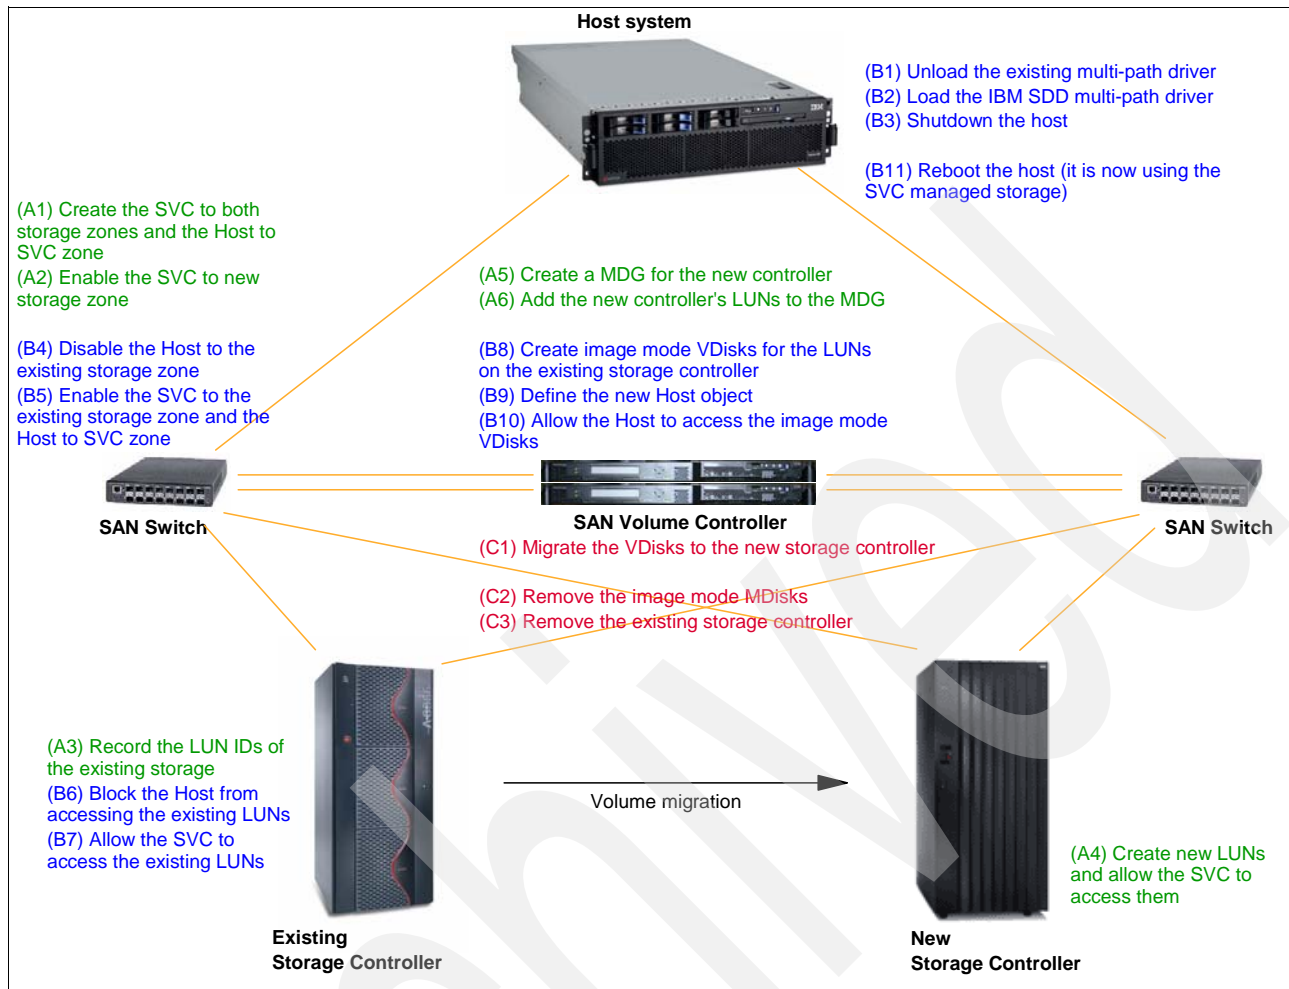


Figure 3-1 Using the SVC to virtualize existing storage and migrate to a new controller

The steps shown in this illustration describe a generic SVC implementation and migration project, and will vary slightly depending on the specific host platform and storage controller types.

3.1.3 Preparing the environment

The steps in this section are shown in Figure 3-1 prefixed with the letter A.

Preparing the SAN fabric (A1, A2)

- ▶ Create the Host to SVC zone (but do not activate it yet).
- ▶ Create the SVC to the existing storage controller zone (but do not activate it yet).
- ▶ Create the SVC to the new storage controller zone (activate it now).
- ▶ Repeat for both SAN fabrics.

Preparing the existing storage controller (A3)

- ▶ On the existing controller record LUN IDs and the assigned scsi ID for the hosts.

Preparing the new storage controller (A4)

- ▶ Create new LUNs for use by the SVC.
- ▶ Map the new LUNs to the SVC.

Preparing the SAN Volume Controller (A5, A6)

- ▶ Detect the LUNs on the new storage controller as MDisks.
- ▶ Create an MDG for the MDisks being presented by the new controller.
- ▶ Add the new controller's MDisks to the new MDG.

3.1.4 Migration to SVC managed storage

The steps in this section are shown in Figure 3-1 prefixed with the letter B.

Configure the host system (B1, B2, B3)

- ▶ Unload the existing multi-path device driver.
- ▶ Load the IBM SDD multi-path device driver.
- ▶ Shut down the host.

Configure the SAN fabric (B4, B5)

- ▶ Disable the Host to the existing storage controller zone.
- ▶ Enable the Host to the SVC zone.
- ▶ Enable the SVC to the existing storage controller zone.

Configure the existing storage controller (B6, B7)

- ▶ Unmap the LUNs from the Host.
- ▶ Remap the LUNs to the SVC.

Configure the SVC (B8, B9, B10)

- ▶ Create image mode VDisks for the LUNs on the existing storage controller.
- ▶ Define a new host object.
- ▶ Map the image mode VDisks to the host.

Configure the host system (B11)

- ▶ Restart the host.
- ▶ The host is now accessing the SVC managed storage (which is still physically on the original storage controller).

3.1.5 Migration of SVC managed storage between controllers

The steps in this section are shown in Figure 3-1 prefixed with the letter C.

Configure the SAN Volume Controller (C1, C2, C3)

- ▶ Migrate the image mode VDisks to the MDG on the new storage controller.
- ▶ The host is now accessing the SVC managed storage (which is now physically on the new storage controller).
- ▶ Remove the image mode MDisks (that represent LUNs on the existing storage controller).
- ▶ If desired, you can now remove the existing storage controller.

3.1.6 Migration from SVC managed storage (optional)

Typically, after implementing the SVC, it would remain in the storage environment. This allows the ongoing benefits of storage virtualization and administrative simplification that the SVC provides.

In some cases, however, you might wish to externalize volumes from the control of the SVC. This would allow the storage volumes to again be directly connected to the host systems.

You might want to perform this activity for any of the following reasons:

- ▶ You were only using SVC as a tool to migrate between storage systems.
- ▶ You were only using SVC to FlashCopy or Metro Mirror a VDisk onto another VDisk, and you no longer need that host connected to the SVC.
- ▶ You want to ship a host and its data that is currently connected to the SVC to a site where there is no SVC.
- ▶ Other changes to your environment no longer require this host to use the SVC.

3.2 Types of SVC volume migration

Once the existing storage has been virtualized with the SAN Volume Controller, you have various options to further manage the volume's physical location.

3.2.1 Migrating a VDisk between MDGs

A VDisk is composed of numerous *extents* that reside on one or more MDisks. Each MDisk represents a LUN presented by a back-end storage controller.

These MDisks are logically grouped into Managed Disk Groups (MDG). Each MDG should contain MDisks of similar performance, reliability, and capacity, and should not span across storage controllers.

A VDisk can be migrated between Managed Disk Groups within the SVC cluster. This allows you to migrate a VDisk (host accessible LUN) between physical MDisk locations to take advantage of differing characteristics of the underlying storage controllers.

Among the reasons you might wish to migrate to a VDisk are the following:

- ▶ Different performance characteristics of the back-end storage controllers.
- ▶ Different RAID levels (and performance) of the MDisks.
- ▶ To manage workload and capacity utilization of the back-end storage controller.
- ▶ When retiring a back-end storage controller.

The VDisk migration between MDGs is managed by the SVC and causes no disruption to the connected hosts.

Figure 3-2 on page 36 shows the migration of a VDisk's extent from one MDG to another. In this case the two MDGs contain MDisks from different storage controllers; therefore, the effect is to migrate the VDisk between storage controllers.

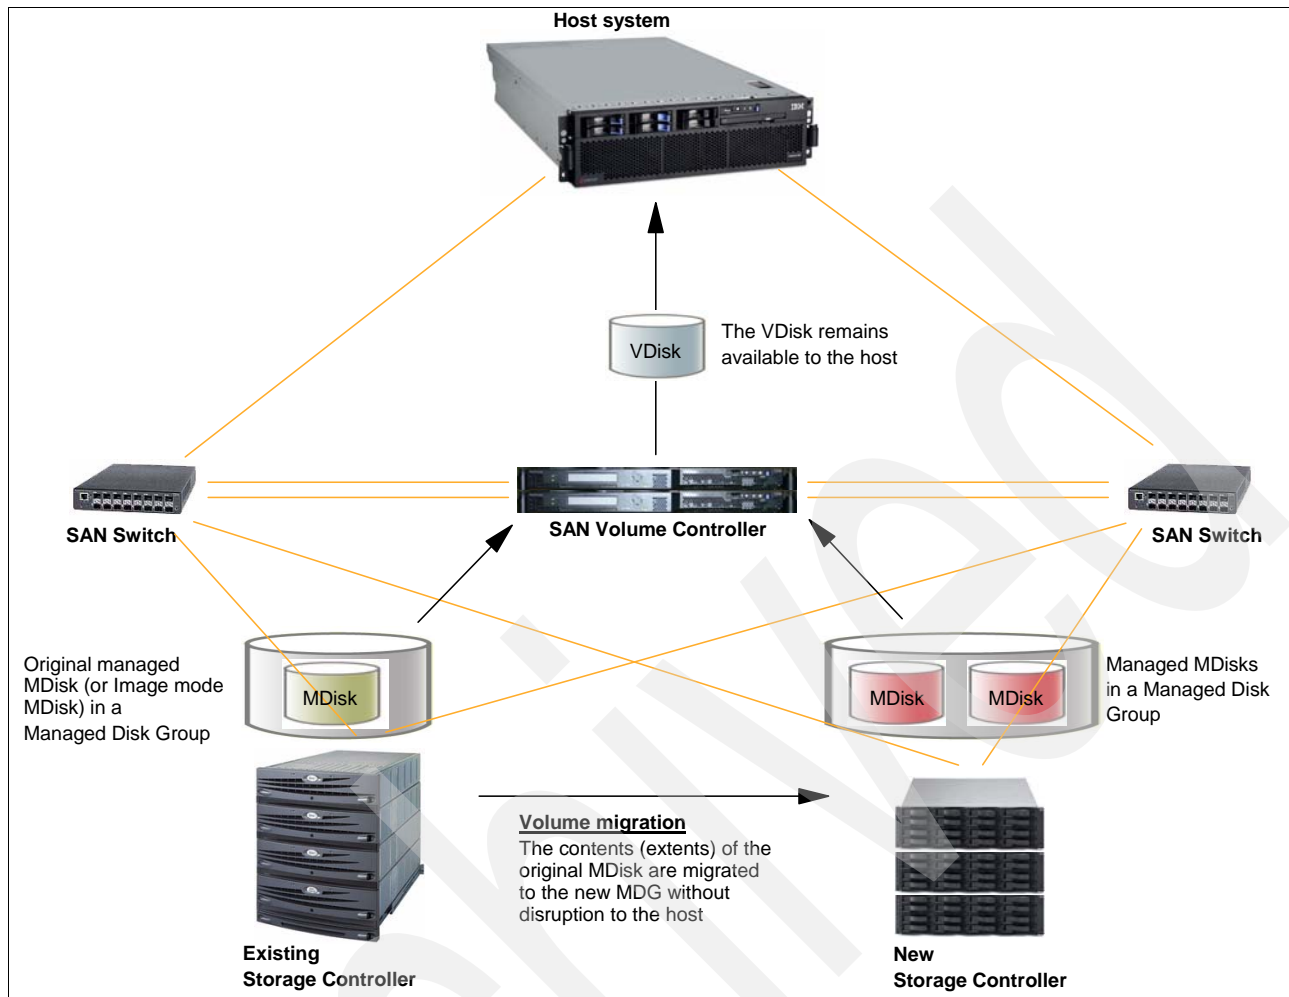


Figure 3-2 Migrating a VDisk between MDGs (and storage controllers)

Use the **svctask migratevdisk** command to start the volume migration.

The syntax of the command is shown in Example 3-1.

Example 3-1 **svctask migratevdisk**

```
>>> svctask -- -- migratevdisk -- ----->
>-- -mdiskgrp --+- mdisk_group_id ---+-- ----->
>--              '- mdisk_group_name -'
>--+-----+-----+----->
>--      '- -threads -- number_of_threads -'
>-- -vdisk --+- vdisk_id ---+-----><
>--              '- vdisk_name -'
```

It accepts the following parameters:

- ▶ **-mdiskgrp mdisk_group_id | mdisk_group_name**
Specifies the new managed disk group ID or name.

- ▶ **-threads number_of_threads**
 - Optionally specifies the number of threads to use while migrating these extents.
 - You can specify 1 - 4 threads. The default number of threads is 4.
- ▶ **-vdisk vdisk_id | vdisk_name**
 - Specifies the virtual disk ID or name to migrate into a new managed disk group.

This command runs asynchronously (in the background) and produces no feedback.

Restriction: To successfully migrate a VDisk between MDGs, the extent size of the two MDGs must be identical.

A VDisk can be migrated between MDGs regardless of the virtualization type (image, striped, or sequential); however, the migration will cause the transition to the virtualization type of striped.

The **migratevdisk** command will fail if there are insufficient free extents on the targeted managed disk group for the duration of the command.

3.2.2 Migrating a VDisk to image mode

Just as an existing storage controller's LUN can be virtualized using the SVC with all data intact, an SVC-managed VDisk can also be restored as a normal LUN with all data intact. This LUN can then be removed from the control of the SVC (*externalized*) and directly accessed by the host to the back-end storage controller.

Among the reasons you might wish to migrate a VDisk to image mode are the following:

- ▶ You are preparing to externalize the VDisk for direct connection to a host.
- ▶ As part of the process for toggling *cache-disabled* mode for the VDisk.

The initial VDisk migration to image mode is managed by the SVC and causes no disruption to the connected hosts. However, the host must un-mount the VDisk before it can be deleted and then re-mounted directly from the back-end storage controller.

Figure 3-3 on page 38 shows the migration of a VDisk from a managed mode on one storage controller, to image mode on an MDisk (which is placed in another MDG). The MDisk comes from another storage controller. The image mode VDisk is then deleted, which also removes the MDisk from the MDG. The unmanaged LUN on the back-end storage controller can then be accessed directly by the host.

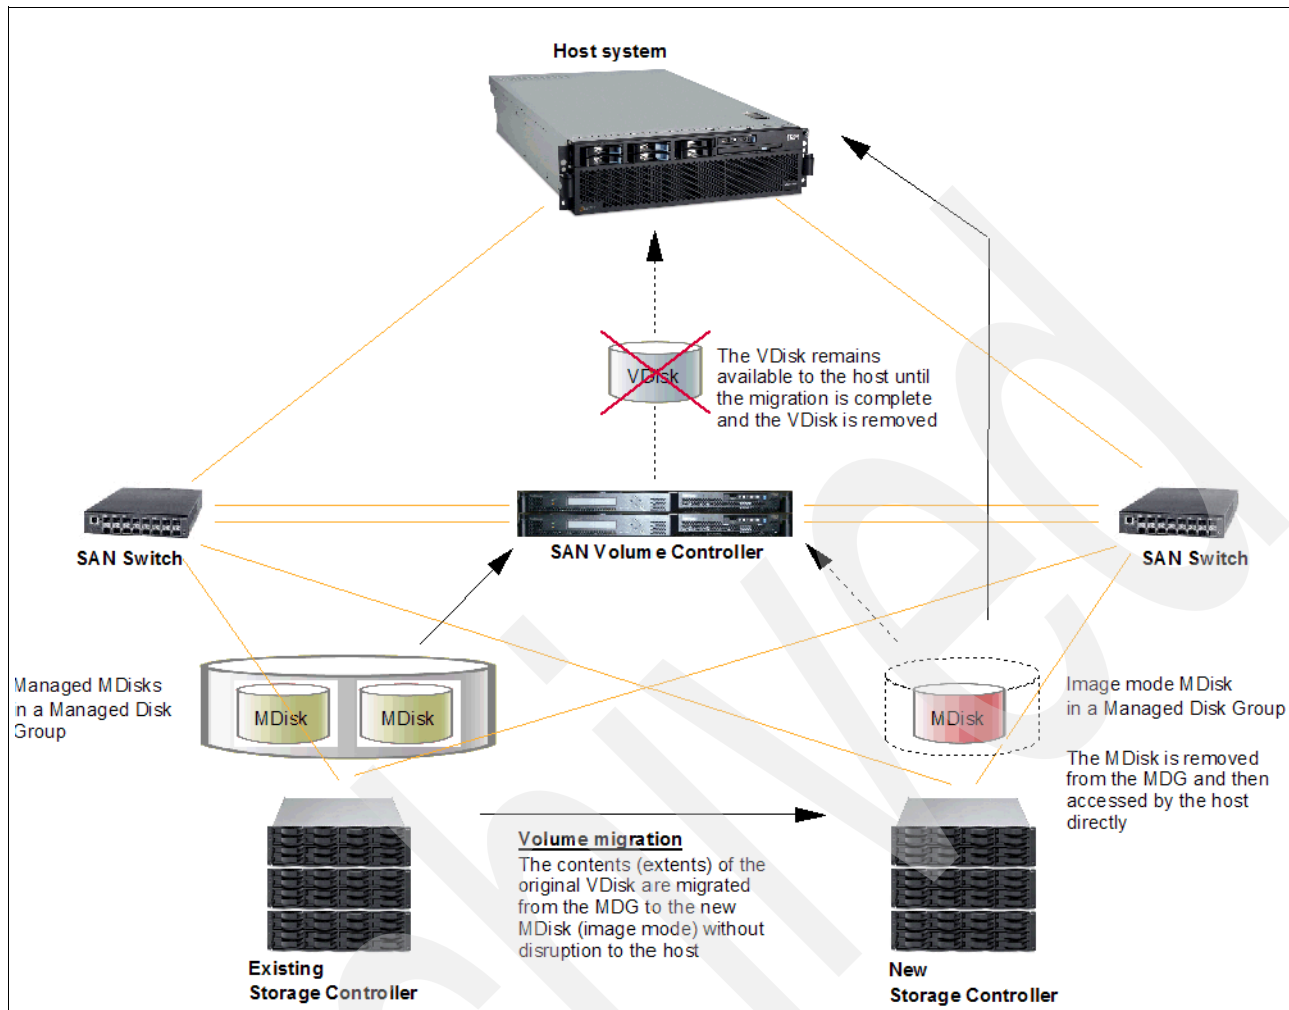


Figure 3-3 Migrating a VDisk to image mode (and between MDGs and storage controllers)

Use the **svctask migratetoimage** command to start the volume migration to image mode.

The syntax of the command is shown in Example 3-2.

Example 3-2

```
>>- svctask -- -- migratetoimage -- ----->
>-- -vdisk --+- source_vdisk_id ---+----->
      '- source_vdisk_name -'
>--+-----+-----+----->
      '- -threads -- number_of_threads -'
>-- -mdisk --+- unmanaged_target_mdisk_id ---+-- ----->
      '- unmanaged_target_mdisk_name -'
>-- -mdiskgrp --+- managed_disk_group_id ---+-----><
      '- managed_disk_group_name -'
```


It accepts the following parameters:

- ▶ **-vdisk source_vdisk_id | name**
Specifies the name or ID of the source VDisk to be migrated.
- ▶ **-threads number_of_threads**
 - Optionally specifies the number of threads to use while migrating these extents.
 - You can specify 1 - 4 threads. The default is 4.
- ▶ **-mdisk unmanaged_target_mdisk_id | name**
Specifies the name of the MDisk to which the data must be migrated.
- ▶ **-mdiskgrp managed_disk_group_id | name**
Specifies the MDisk group into which the MDisk must be placed after the migration has completed.

This command runs asynchronously (in the background) and produces no feedback.

Restriction: To successfully migrate a VDisk to image mode, the target MDisk must be unmanaged and large enough to contain the data being migrated.

The target MDisk must reside in an MDG that has the same extent size as the VDisk's current MDG.

Any striped or sequential VDisk can be migrated to image mode.

3.2.3 Migrating a VDisk between I/O groups

When you first create a VDisk, it is associated with one of the I/O groups in your SVC cluster. As your virtualized storage environment grows you may need to migrate to a VDisk between I/O groups in order to balance workload across the SVC nodes.

Each pair of nodes in your SVC cluster is one I/O group. For example, a six node SVC cluster is comprised of three I/O groups (pairs of nodes). If you only have one pair of SVC nodes then you only have a single I/O group, so this section does not apply to your environment unless you expand the SVC cluster with another I/O group.

Among the reasons you may wish to migrate to a VDisk between I/O groups are the following:

- ▶ To balance the workload between I/O groups.
- ▶ To remove an I/O group from the SVC cluster.

The VDisk migration between I/O groups is a disruptive procedure, meaning access to the VDisk will be lost during the VDisk migration.

In order to move a VDisk between I/O groups, the node's cache must be flushed. The SVC will attempt to de-stage all write data for the VDisk from the cache during the I/O group move. During the flush, the VDisk operates in cache write-through mode.

The cache flush will fail if data has been pinned in the cache for any reason (such as the MDG being offline). By default this will cause the migration between I/O groups to fail, but this behavior can be overridden using the **-force** flag. If the **-force** flag is used, and if the SVC is unable to de-stage all write data from the cache, the result is that the contents of the VDisk might be corrupted by the loss of the cached data.

You must quiesce host I/O before the migration in either of these situations:

1. If there is significant data in cache and it takes a long time to de-stage, the command will time out.
2. SDD vpaths associated with the VDisk are deleted before the VDisk move takes place in order to avoid data corruption. Data corruption could occur if I/O is still ongoing at a particular LUN ID when it is re-used for another VDisk.

Note: When migrating a VDisk between I/O Groups, the user does not have the ability to specify the preferred node. The preferred node is assigned by the SVC.

Figure 3-4 shows a single SVC cluster that contains two I/O groups (pairs of SVC nodes).

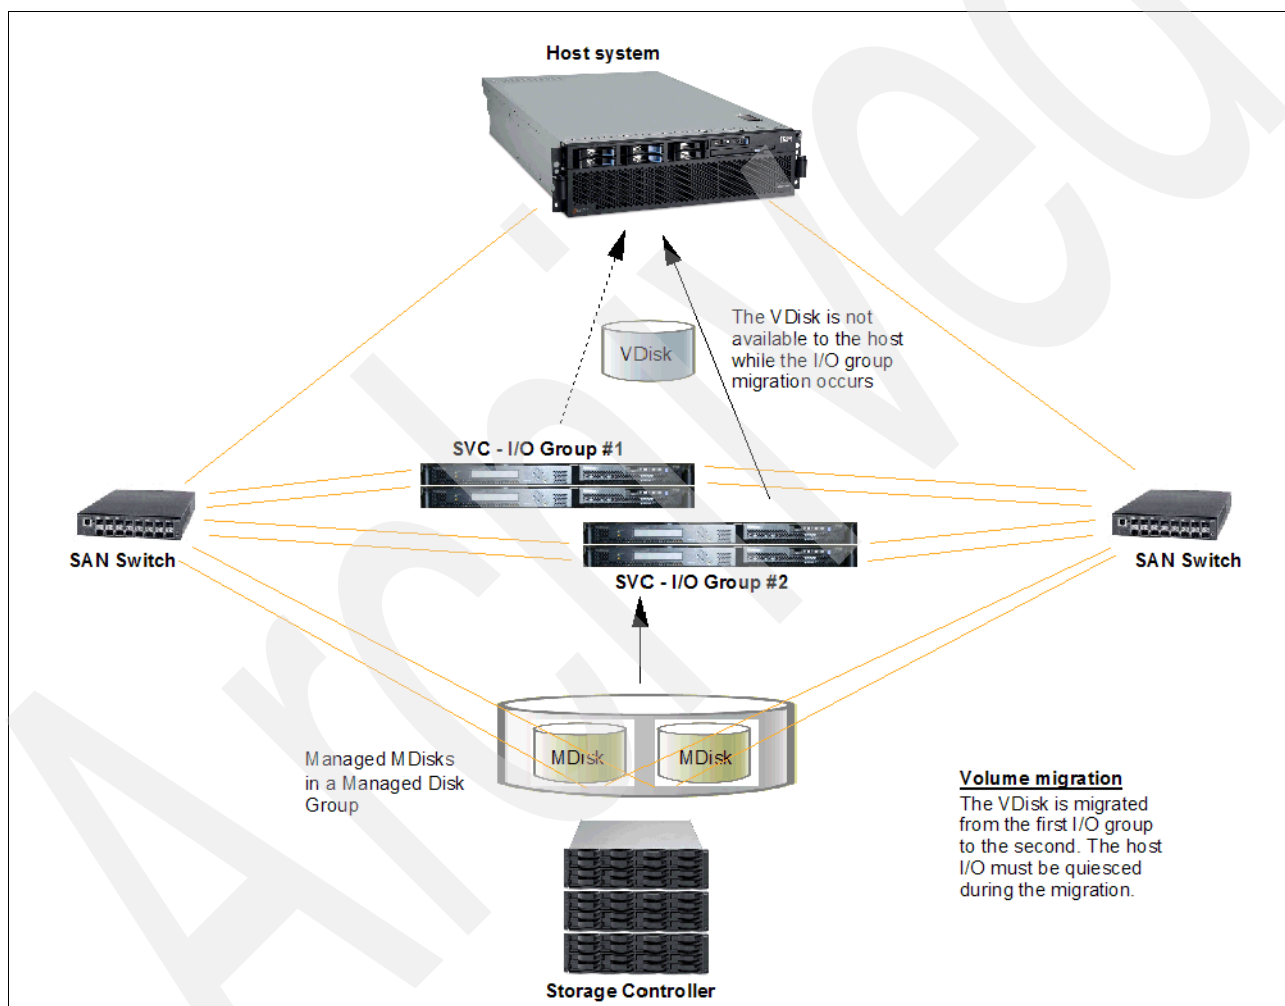


Figure 3-4 Migrating a VDisk between I/O groups

Use the **svctask chvdisk** command to start the VDisk migration between I/O groups.

The syntax of the command is shown in Example 3-3.

Example 3-3 *svctask chvdisk*

```
>>- svctask -- -- chvdisk -- ----->

>--+ -name -- new_name_arg -----+-->
```

```

+- -iogrp --+- io_group_id ---+---+-----+--+
|           '- io_group_name -' '- -force -' |
'- -rate -- throttle_rate --+-----+-----'
                               '- -unitmb -'

```

It accepts the following parameters:

► **-iogrp io_group_id | io_group_name**

- Optionally specifies a new I/O group to move the virtual disk to, either by ID or name. The **-force** flag can be used together with this parameter in order to force the removal of the VDisk to the I/O group.
- If the VDisk has a mapping to any hosts, it will not be possible to move the VDisk to an I/O group that does not include any of those hosts.

► **-force**

Specifies that you want to force the VDisk to be removed from an I/O group. This parameter can only be used together with **-iogrp**.

Attention: The **-force** flag can corrupt the contents of the VDisk. If the **-force** flag is used and if the cluster is unable to de-stage all write data from the cache, the result is that the contents of the VDisk might be corrupted by the loss of the cached data.

► **-rate throttle_rate [-unitmb]**

Optionally sets the I/O governing rates for the virtual disk. The default units are I/Os, but they can be used in conjunction with the **-unitmb** argument to specify in terms of MBps.

► **-name new_name_arg**

Optionally specifies a new name to assign to the virtual disk.

► **-udid vdisk_udid**

Optionally specifies the udid for the disk. Valid options are a decimal number from 0 to 32767, or a hex number from 0 to 0x7FFF. A hex number must be preceded by 0x (for example, 0x1234). If this parameter is omitted the default udid is 0.

► **vdisk_name | vdisk_id**

Specifies the virtual disk to modify, either by ID or by name.

The **-iogrp**, **-rate**, **-udid** and **-name** parameters are mutually exclusive. Only one of these parameters can be specified per command line.

This command runs asynchronously (in the background) and produces no feedback.

Attention: Under no circumstances should you move a VDisk to an offline I/O group. You must ensure the I/O group is online before moving the VDIs to avoid any data loss.

3.2.4 Migrating VDisk extents (within an MDG)

When you create a striped VDisk, its extents are evenly distributed within its MDG. This spreads the workload and capacity utilization evenly across all the MDIs in that MDG.

As your virtualized storage environment grows, the distribution of VDisk extents may no longer be even across all the MDisks in an MDG. This is because as you add new MDisks to the MDG, the existing MDisks already contain many extents, while the new MDisks are empty. The SVC does not automatically redistribute the extents across the new MDisks.

You may need to migrate VDisk extents within an MDG to balance workload across any new MDisks within that MDG.

Among the reasons you may wish to migrate the VDisk extents are the following:

- ▶ To balance workload across the MDisks with the MDG.
- ▶ You wish to free up the extents on a MDisk.

The migration of the VDisk extents is managed by the SVC and causes no disruption to the connect hosts.

Figure 3-5 shows the migration of VDisk extents within an MDG.

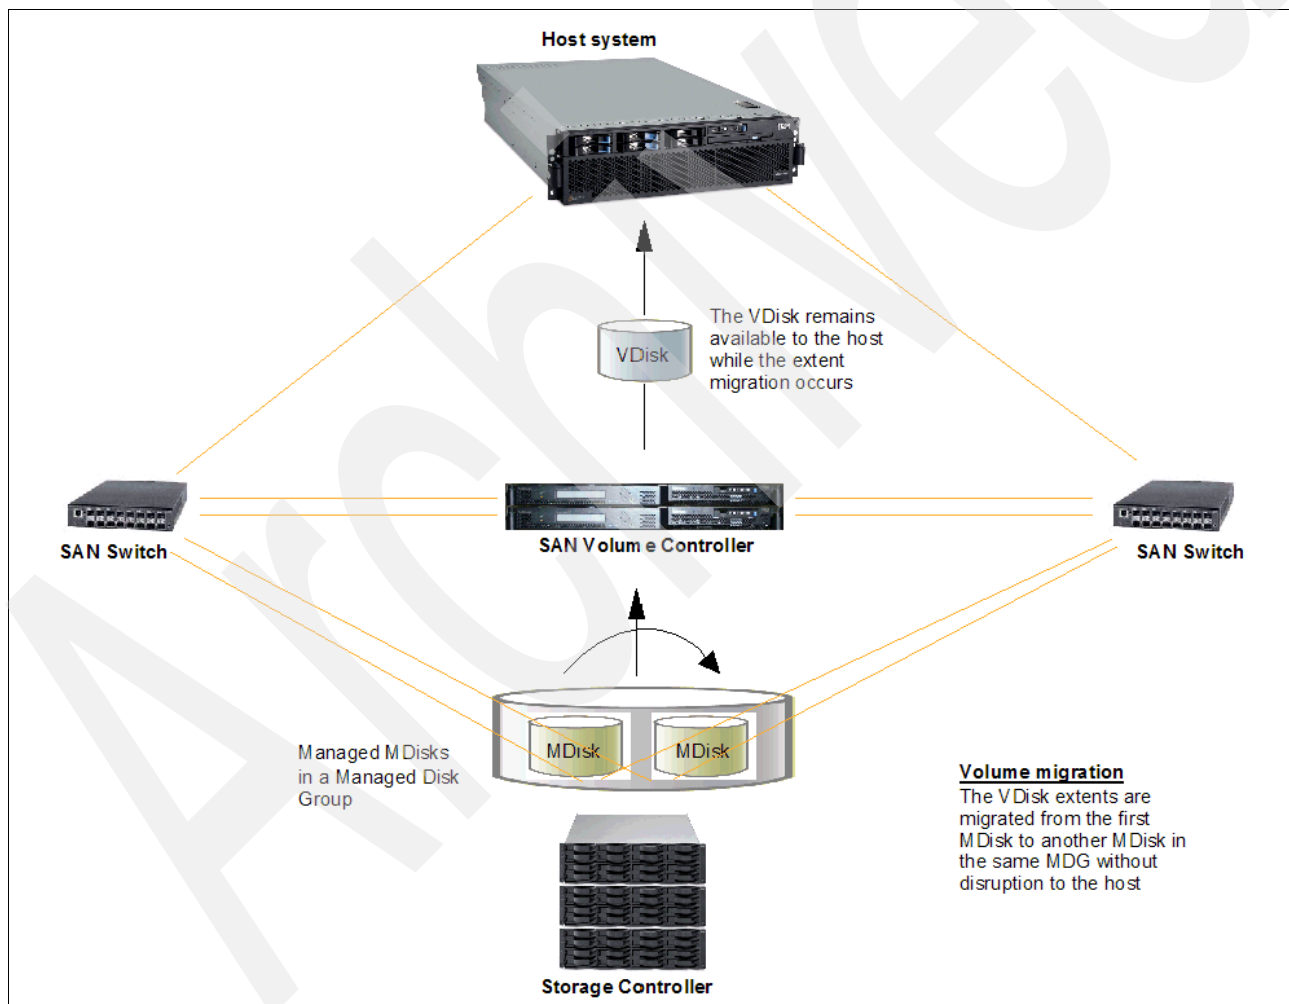


Figure 3-5 Migrating VDisk extents between MDisks in the same MDG

Use the **svctask migrateexts** command to start the volume migration to image mode.

The syntax of the command is shown in Example 3-4.

Example 3-4 *svctask migrateexts*

```
>>- svctask -- -- migrateexts -- ----->

>-- -source --+- source_mdisk_id ---+-- ----->
        '- source_mdisk_name -'

>-- -target --+- target_mdisk_id ---+-- ----->
        '- target_mdisk_name -'

>-- -exts -- number_of_extents -- ----->

>-- -vdisk --+- vdisk_id ---+-- ----->
        '- vdisk_name -'

>--+-----+----->
        '- -threads -- number_of_threads -'
```

It accepts the following parameters:

- ▶ **-source source_mdisk_id | source_mdisk_name**
Specifies the MDisk on which the extents currently reside.
- ▶ **-target target_mdisk_id | target_mdisk_name]**
Specifies the MDisk onto which the extents are to be migrated.
- ▶ **-exts number_of_extents**
Specifies the number of extents to migrate.
- ▶ **-vdisk vdisk_id | vdisk_na**
Specifies the VDisk to which the extents belong.
- ▶ **-threads number_of_threads**
Optionally specifies the number of threads to use while migrating these extents. You can specify 1 - 4 threads. The default is 4.

This command runs asynchronously (in the background) and produces no feedback.

For an example script to automatically balance the VDisks extents, consult the SVC scripting tools Web site:

<http://www.alphaworks.ibm.com/tech/svctools/download>

3.2.5 Migrating extents off an MDisk that is being removed from the MDG

An MDisk represents a LUN on the back-end storage controller that is presented to the SVC. These MDisks are logically grouped into MDGs and divided into *extents*. The extent size is fixed for each MDG upon creation; the available sizes are from 16 MB to 512 MB. A host-accessible VDisk is constructed from the extents in the defined MDG.

The removal of an MDisk from an MDG will fail if it contains any extents that are in use by a VDisk. This behavior can be overridden with the **-force** flag, which causes any occupied extents on the MDisk to be migrated to other MDisks in the same MDG prior to its deletion. The migration (and MDisk deletion) will fail if there are not enough free extents on other MDisks within the same MDG.

If a VDisk uses one or more extents that need to be moved as a result of a `remove mdisk` command, then the virtualization type for that VDisk is set to *striped* (if it previously was *sequential* or *image*).

If the MDisk is operating in *image* mode, and the MDisk is deleted, the MDisk transitions to *managed* mode while the extents are being migrated, and upon deletion it transitions to *unmanaged*. If there was a VDisk related to the image mode MDisk, then the VDisk transitions to *striped mode*.

Among the reasons you may wish to delete an MDisk is the following:

- You are preparing to delete the underlying LUN on the back-end storage controller.

The migration of the MDisk extents is managed by the SVC and causes no disruption to the connected hosts.

Attention: You must not destroy any controller LUNs until you have removed their MDisk representations from the MDG that they belong to. Best practice is to unmap the LUN on the storage controller prior to deletion, in order to verify that you are working with the intended MDisk. If the controller LUN has been deliberately removed, then the only method of removing the MDisk is to remove the entire MDG.

Figure 3-6 shows the automatic migration of extents from an MDisk when it is removed from the MDG.

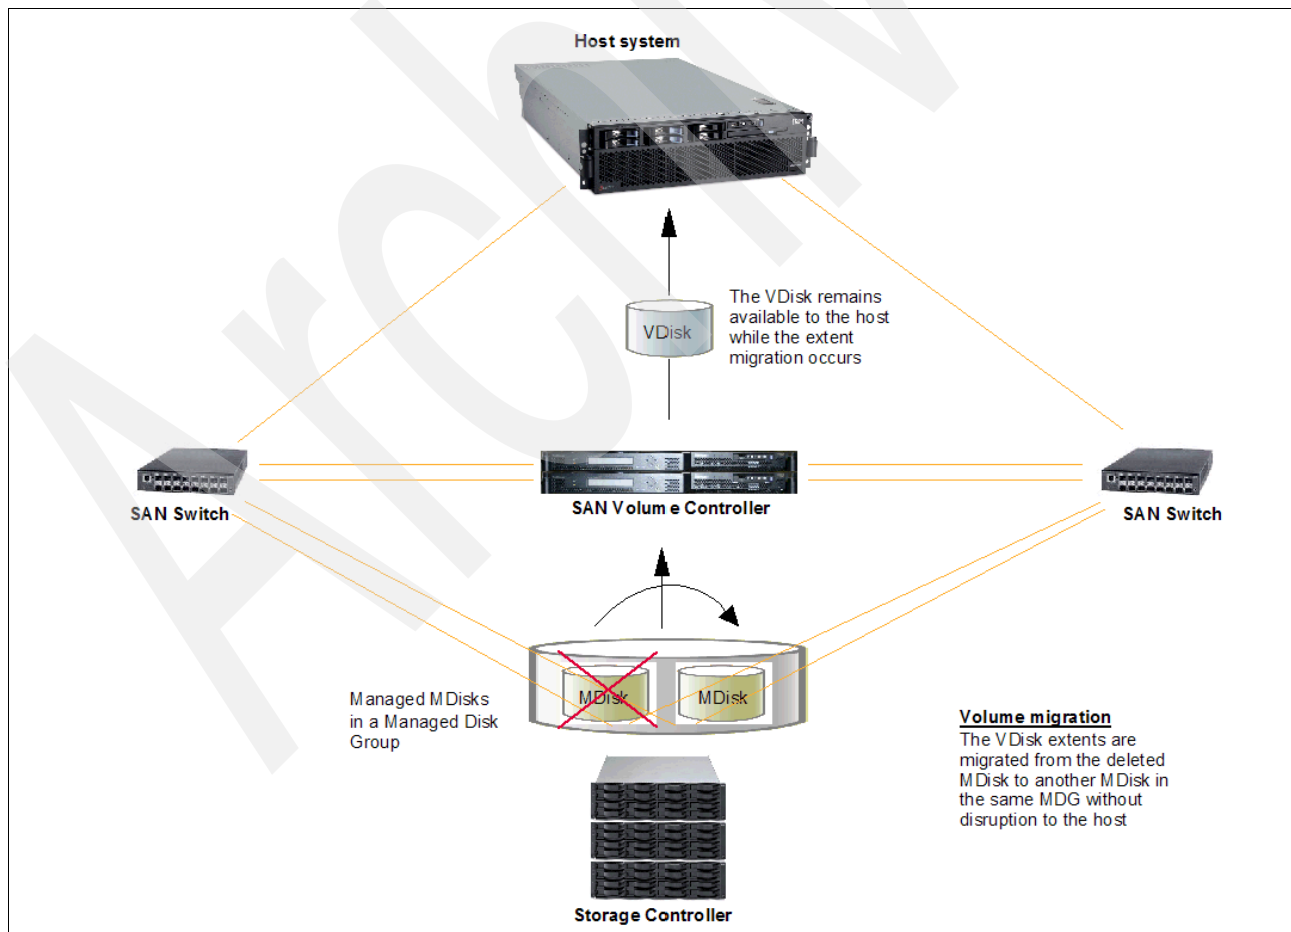


Figure 3-6 Migrating extents off an MDisk which is being deleted

Use the **svctask rmmdisk** command to delete an MDisk and migrate its contents within the MDG.

The syntax of the command is shown in Example 3-5.

Example 3-5 svctask rmmdisk

```
>>- svctask -- -- rmmdisk -- ----->

>-- -mdisk --+- mdisk_id_list ---+- --+-----+-- ----->
           '- mdisk_name_list -'      '- -force -'

>---+- mdisk_group_id ---+-----><
      '- mdisk_group_name -'
```

It accepts the following parameters:

► **-mdisk mdisk_id_list | mdisk_name_list**

Specifies one or more managed disk IDs or names to delete from the group.

► **-force**

Optionally specifies the force flag. If you do not supply the **-force** flag, and virtual disks exist that are made from extents on one or more of the managed disks specified, the command will fail. If you do supply the force flag, and virtual disks exist that are made from extents on one or more of the managed disks specified, any data on the disks will be migrated to other disks in the group if there are enough free extents in the group. This operation might take some time.

► **mdisk_group_id | mdisk_group_name**

Specifies the ID or name of the managed disk group from which to delete the disk or disks.

Note: The **rmmdisk** command is used to remove an MDisk from a MDG, it does not actually delete the MDisk itself.

This command runs asynchronously (in the background) and produces no feedback. The deletion of the MDisk is postponed until all extents are migrated, which may take some time.

3.2.6 Non-standard migration types

This section describes some advanced migration techniques that are not typically required in normal SVC operations.

Changing the preferred node of a VDisk

When a VDisk is first created it is associated with an SVC I/O group (pair of nodes), and is assigned a *preferred node* within that group. All host access to that VDisk is then provided through the associated I/O group, with all traffic going to the preferred node by default. If the preferred node is not available then the host can still access the Vdisk through the other node in the same I/O group.

Generally the I/O workload is spread reasonably evenly between the two nodes in an I/O group, but sometimes it may be desirable to rebalance the workload between the two nodes.

Although the current SVC release does not have a specific command to change a VDisk's preferred node, it is possible to do so with either of the following procedures.

Note: If you are unfamiliar with the following steps, then contact your IBM representative prior to carrying out the task.

Method 1

This method should not disrupt host access to the VDisk.

1. Remove one of the SVC nodes from the I/O group.

When the node is removed from the I/O group all access to its VDisks occurs through the remaining node.

Assuming that all host access is correctly configured across multiple paths, this should cause no interruption to the connected hosts.

2. Add the SVC node back into the I/O group.

When the node is rejoined to the I/O group the preferred node for all VDisks within that I/O group are evenly redistributed between the two nodes.

The hosts will then continue to access their VDisks through the new preferred nodes.

Note: This simple method does not guarantee an even distribution of the I/O workload between the two nodes in the I/O group.

Method 2

This method will briefly disrupt host access to the VDisk.

1. Create a new VDisk on the opposite node.

The new VDisk should be created on the other node in the same I/O group.

Note: If you wish to migrate the VDisk between I/O groups (not within the I/O group) then you should use the `chvdisk` command instead.

2. Use Metro Mirror to duplicate the original VDisk.

The contents of the existing VDisk will be mirrored to the new VDisk.

3. Remap the host access to the new VDisk.

- a. Stop the host from accessing the Vdisk. This may involve stopping host applications, un-mounting the volume, or briefly shutting down the host.
- b. Unmap the host access to the original VDisk.
- c. Remap the host access to the new VDisk.
- d. Restart the host accessing the VDisk.

Note: This method allows the migration of a single VDisk between nodes within the I/O group, and can therefore be used to rebalance the I/O workload.

It may also be possible to achieve this using host-level disk mirroring (LVM), and to do so with no disruption to the VDisk access.

Changing the cache mode of a VDisk

When a VDisk is first created it is by default made as a cache-enabled VDisk associated with an SVC I/O group (pair of nodes), and is assigned a preferred node within that group. All host access to that VDisk is then provided through the associated I/O group, with all traffic going to

the preferred node by default. If the preferred node is not available then the host can still access the Vdisk through the other node in the same I/O group.

Generally the I/O workload is by default spread reasonably evenly between the two nodes in an I/O group, but sometimes it may be desirable to rebalance the workload between the two nodes.

Although the current SVC release does not have a specific command to change a VDisk's cache mode, it is possible to do so with either of the following procedures.

Method 1

This method will briefly disrupt host access to the VDisk.

1. Create a new cache-enabled VDisk.
2. Use Metro Mirror to duplicate the original VDisk.
3. Remap the host access to the new VDisk.

Method 2

This method will briefly disrupt host access to the VDisk.

1. Migrate the cache-disabled VDisk to image mode.
2. Delete the cache-disabled image mode VDisk.
3. Create a new cache-enabled image mode VDisk.
4. Migrate the image mode VDisk to striped mode.

Note: It may also be possible to achieve the same result using host-level disk mirroring (LVM) and to do so with no disruption to the VDisk access.

3.3 Monitoring the migration process

You can monitor the progress of a running migration with the `svcinfolsmigrate` command.

This command will produce an output similar to that displayed in Example 3-6.

Example 3-6 Monitoring the VDisk migration process

```
IBM_2145:ldcluster:admin>svcinfolsmigrate
```

```
migrate_type MDisk_Group_Migration
progress 60
migrate_source_vdisk_index 0
migrate_target_mdisk_grp 1
max_thread_count 4
```

In Example 3-6, the progress field shows the completed percentage of the VDisk migration. This value may appear slow to increment because it only updates once per minute, and only when all threads have completed the migration of an extent. For large extent sizes with many threads, this can result in quite large increments in the percentage progress.

Another way to monitor the migration process is by listing the location of the MDisk extents that make up the VDisk that is being migrated. Initially, as an image mode VDisk, all the extents will be located on a single MDisk, but as the migration proceeds all the extents will be migrated to other MDisk.

You can view the MDisk extent distribution for a given VDisk with the **svcinfolsvdiskextent vdiskname** command.

This command will produce output similar to that as shown in Example 3-7.

Example 3-7 View the MDisk extents of a VDisk during migration

```
IBM_2145:ldcluster:admin>svcinfolsvdiskextent general01
```

id	number_extents
0	12974 (This was the original image mode VDisk/MDisk)
3	7
9	7
18	7
26	7
37	7
42	8
54	8

As the VDisk migration proceeds, the number of extents on the original MDisk will decrease. At the same time the number of extents on the MDisk in the target MDG will increase.

3.4 Summary of the SVC migration process

Figure 3-7 on page 49 provides a conceptual overview of the steps that are undertaken during the migration to SVC managed storage, and migration to a new back-end storage controller.

At a generic level, the key steps in the SVC migration process are as follows:

- ▶ On the SAN fabric switches
 - The SAN is re-zoned to present the existing LUN through the SVC.
- ▶ On the existing storage controller
 - The existing LUN is unmapped from the host and remapped to the SVC.
- ▶ On the new storage controller
 - The new LUNs are mapped to the SVC.
- ▶ On the SAN Volume Controller
 - An image mode VDisk is created for the MDisk from the existing storage controller.
 - The contents of the original LUN remain intact.
 - The new VDisk is presented to the host.
- ▶ On the host system
 - The previous multi-path device driver is unloaded.
 - The IBM SDD multipath device driver is loaded.
 - The host sees the VDisk as though it was the original LUN.
- ▶ On the SAN Volume Controller
 - The VDisk's extents are migrated to the new storage controller in the background.
 - There is no further disruption to the host during the extent migration.

Note: This overview does not attempt to describe the order of these steps or their interdependencies.

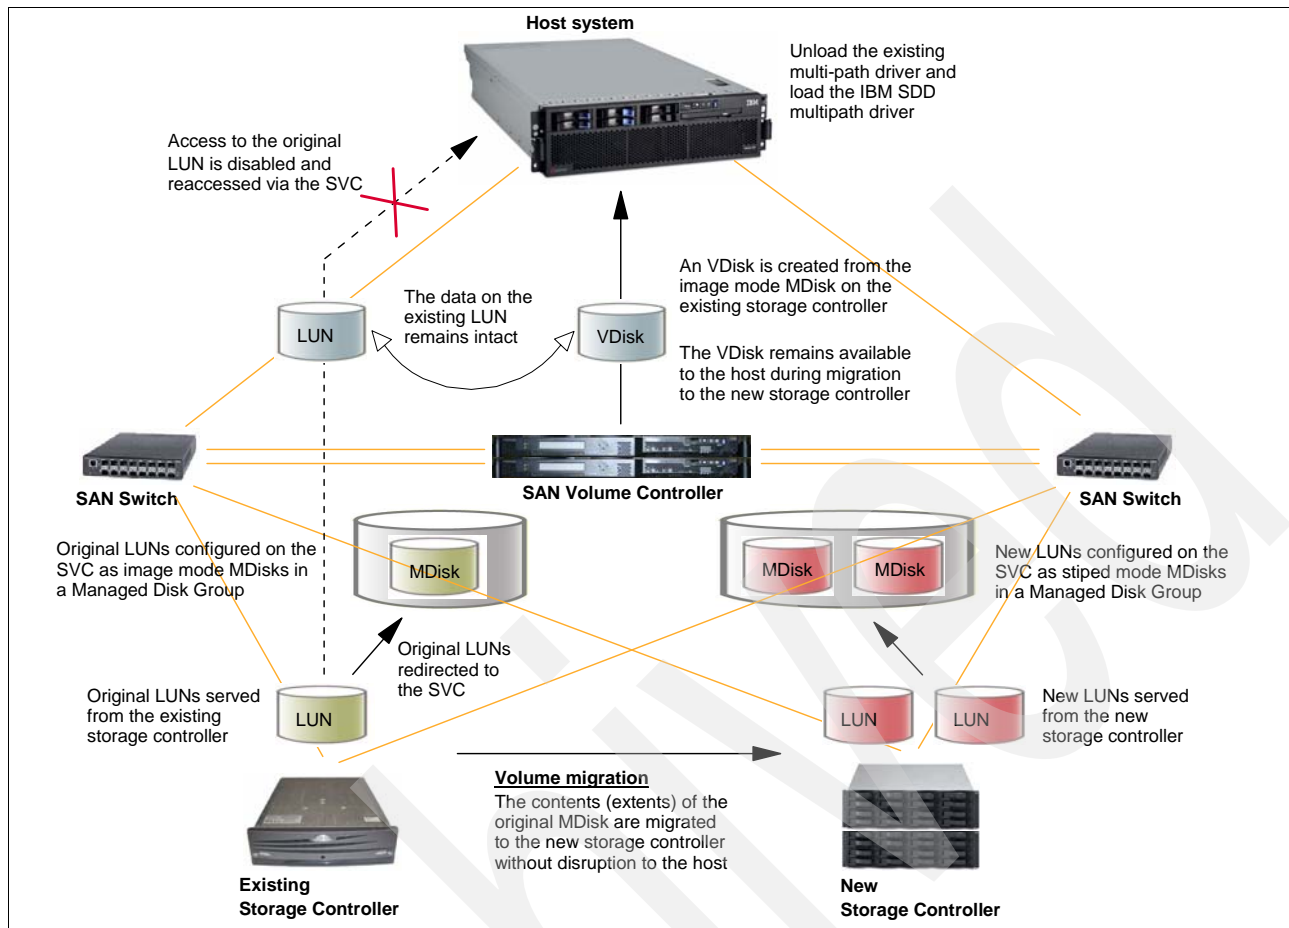


Figure 3-7 Overview of the SVC migration process

The second part of this book applies the general steps described here to actual implementations of the SVC on various vendor platforms.

Archived

Part 2

Implementation in an OEM storage environment

This part describes how to implement the SAN Volume Controller (SVC) in an existing vendor acquired storage environment. It covers several common vendor acquired storage controllers and presents practical, step by step instructions on how to migrate data to IBM storage systems.

The following table shows a cross-reference of the storage controllers and host operating systems covered in each of our SVC implementation examples.

Matrix of storage controller and host migration scenarios

	IBM AIX	HP HP-UX	Linux	SUN Solaris	Windows 2003	VMWare ESX	NetApp V-Series
HDS Thunder						X	
HDS Lightning					X		
HDS AMS							X
HDS UPS					See Lightning		
EMC Clariion	X	X					
EMC Symmetrix	X				X		
EMC DMX							
HP EVA	X				X		

	IBM AIX	HP HP-UX	Linux	SUN Solaris	Windows 2003	VMWare ESX	NetApp V-Series
HP XP (see HDS)							
STK D-Series							
STK FlexLine							
SUN 9900 (see HDS)							
NetApp FAS					X		
IBM DS4000			X		X		X
IBM DS8000			X		X		

This table only represents those combinations documented in our examples. For a full description of the SVC controller and host interoperability, refer to the following Web site:

<http://www-03.ibm.com/servers/storage/software/virtualization/svc/interop.html>



Implementing the SVC with DS4000

In this chapter we describe scenarios for implementing the SAN Volume Controller in an existing environment of DS4800 and DS8000 controllers.

4.1 Scenario with DS4800

In this chapter, we demonstrate the implementation of a SAN Volume Controller (SVC) in an existing environment with a DS4800 storage controller. The existing storage volumes on the DS4800 storage controller will be virtualized using the SAN Volume Controller and migrated to an IBM DS8000 storage controller.

Note: Be sure to read section 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization example.

4.1.1 Overview of integration steps

At a high level, the steps to integrate the SVC into an existing production environment are the following:

- ▶ General
 - Review SAN levels of software and firmware, supported hardware, and information from the configuration guide.
 - Determine SVC worldwide port names and host worldwide port names.
- ▶ SAN
 - Create host-to-SVC zone (Enable now or later).
 - Create OEM controller zone with SVC (Enable now or later).
 - Create DS8000 zone with the SVC (Enable now).
- ▶ DS8000 Controller
 - Configure some LUNs and present them to the SVC.
- ▶ SVC
 - Run discovery to detect new DS8000 managed disks (MDisk).
 - Create a managed disk group (MDG) with a suitable name.
 - Add DS8000 LUNs to this MDG.
- ▶ Host
 - Stop applications.
 - Remove existing multipath driver if not supported by the SVC.
 - Install subsystem device driver (SDD) or equivalent driver.
 - Shut down machine or unmount file systems and vary off volume groups.
- ▶ SAN
 - Enable new zone for host-to-SVC and the OEM controller-to-SVC.
 - Disable the OEM controller-to-host zone.
- ▶ OEM Controller
 - Unmap LUNs from host.
 - Remap LUNs to SVC.
- ▶ SVC
 - Discover OEM controller LUNs.
 - Create Image mode Virtual Disks (vDisk).
 - Create host with associated WWPNS.
 - Map vDisk to host.
- ▶ Host
 - Reboot or scan for new devices.

- Vary on volume groups.
- Remount file system.
- ▶ SVC
 - Migrate Image mode vDisks to the DS8000 MDG.
- ▶ OEM Controller
 - On completion of the migration you can remove the existing controller if desired.

4.1.2 Scenario description

In this scenario we start with a DS4800 storage controller and a single SAN-connected Linux or Windows host. We then introduce the SAN Volume Controller to virtualize the back-end storage. The virtualized storage can then be migrated to the new IBM DS8000 storage controller with a limited disruption to the host system.

The systems described in this section include:

- ▶ DS4800 storage controller
- ▶ IBM b-type Fibre Channel switches
- ▶ IBM DS8000 storage controller
- ▶ A Linux or Windows host

Figure 4-1 shows the relationships between the various systems used in this scenario.

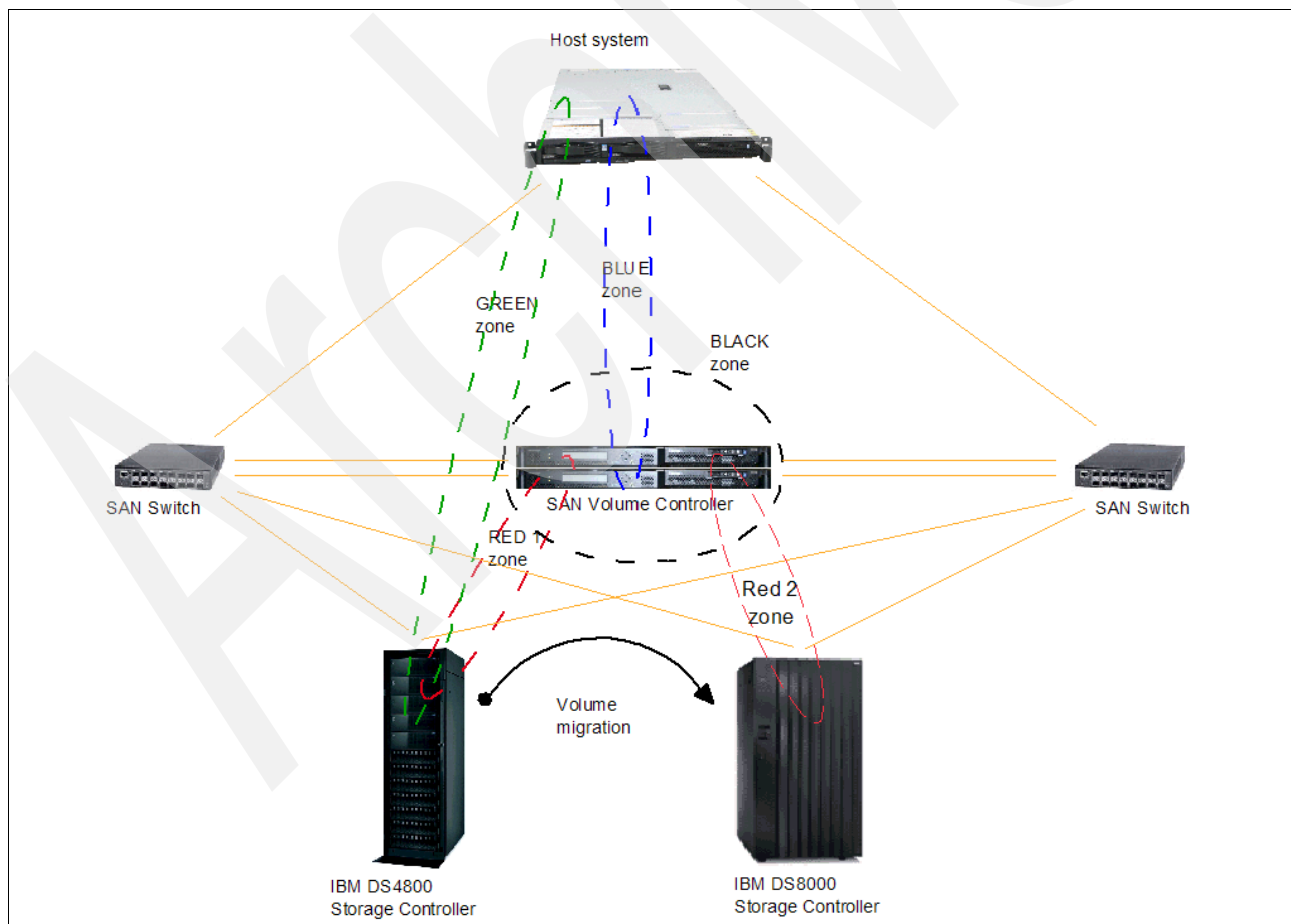


Figure 4-1 SVC implementation example with an existing DS4800 storage controller

After the storage volumes have been migrated to the IBM DS8000 storage controller the DS4800 storage controller could be removed from the environment.

4.2 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

4.2.1 Confirm support for your environment

This section provides information about configuring the SVC in an existing environment with a DS4800 storage controller.

Prior to implementing the SVC it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Concurrent host connection to the DS4800 and the SVC

This section provides information about configuring the IBM System Storage DS4000 series subsystem for attachment to a SAN Volume Controller. Certain models of the IBM DS4000 series of controllers are equivalent to StorageTek™ models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series, and the information in this chapter applies to the supported models as well. Table 4-1 shows the equivalents.

Table 4-1 IBM storage controllers and equivalent StorageTek models

Engenio Controller Number	IBM	StorageTek
2822	DS4100 (Previously FASTT100)	FLX210
2772	FastT200 (discontinued)	D173(discontinued)
2882 (Initial Release) (256MB)	DS4300 (Previously FASTT 600)	FLX 240 DL (discontinued)
2882 (Turbo) (1GB)	DS4300 Turbo (Previously FASTT 600 Turbo)	2882 with 1g memory FLX240
4884	DS4400 (discontinued)	D178
5884	DS4500	FLX280
6998 XBB (2Gb, 4Gb, 8Gb)	DS4800 Models 80, 82, 84, 88	FLX380
3992	DS4700	
3994	DS4700	

Consult the following Web site for specific DS4800 supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Configuring IBM DS4000 series disk controllers for the storage server

The IBM DS4000 series of disk controllers provide functionality that is compatible with the SAN Volume Controller.

Attention: The SAN Volume Controller does not concurrently support I/O operations with the download of ESM (Environmental Services Monitor) firmware. You must quiesce all I/O operations from the hosts that use storage provided by the IBM DS4000 series of controllers that you want to update before you install new ESM firmware.

The following limitations apply to partitions:

- ▶ Only one IBM DS4000 series storage partition that contains any of the ports of any of the nodes in a single SAN Volume Controller cluster can be created.
- ▶ Only map partitions to all of the ports on all of the nodes that are in the SAN Volume Controller cluster to avoid unexpected behavior. For example, you can lose access to your storage or you might not receive warning messages, even if there are errors logged in the SAN Volume Controller error log.

The following information applies to the IBM DS4000 series Copy Services:

- ▶ The IBM DS4000 series Copy Services must not be used when the SAN Volume Controller is attached to the IBM DS4000 series.
- ▶ You can use partitioning to allow IBM DS4000 series Copy Services to be used on other hosts.

The following information applies to the access LUN (also known as the Universal Transport Mechanism (UTM) LUN):

- ▶ The access/UTM LUN is a special LUN that allows the DS4000 to be configured through software over the Fibre Channel connection. The access/UTM LUN does not have to be in the partition that contains the SAN Volume Controller ports because the access/UTM LUN is not required by the SAN Volume Controller. No errors are generated if the access/UTM LUN is not in the partition.

The following information applies to the logical unit (LU):

- ▶ The SAN Volume Controller attempts to follow the preferred ownership that is specified by the IBM DS4000 series. You can specify which controller (A or B) is used for I/O operations to an LU.
- ▶ If the SAN Volume Controller can see the ports of the preferred controller and error conditions do not exist, the SAN Volume Controller accesses the LU through one of the ports on the preferred controller. If error conditions exist, the SAN Volume Controller ignores the preferred ownership of the IBM DS4000 series.

Further information

It is strongly recommended that you refer to the *IBM System Storage SAN Volume Controller 4.1.0 Configuration Guide*, SC26-7902, for further information regarding implementing the SVC with an existing storage controller. It is available at:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

4.2.2 Record the SVC FC port numbers

To obtain the needed information, connect to the SVC and record the WWPN for each of the SVC nodes.

1. Use the **svcinfo node** command to list the SVC nodes. Example 4-1 on page 58 shows the output of this command on our system.

Example 4-1 List the storage controllers

id	name	UPS_serial_number	WWNN	status
IO_group_id	IO_group_name	config_node	UPS_unique_id	hardware
1	node1	10005BL068	50050768010023EC	online 0
io_grp0	yes	2040000152700188	8F2	
2	node2	10005BL066	50050768010024E4	online 0
io_grp0	no	2040000152700186	8F2	

2. We need further information about each of the SVC nodes. Use the **svcinfn lsnod nomencl** command to list the properties of each of the SVC nodes.

Example 4-2 shows the output of this command on our system.

Example 4-2 List the storage controllers

```
IBM_2145:GREEN_SVC:admin>svcnfn lsnod 1
id 0
name node1
UPS_serial_number 10005BL068
50050768010023EC
status online
IO_group_id 0
IO_group_name io_grp0
partner_node_id 2
partner_node_name node2
config_node no
UPS_unique_id 2040000152700188
port_id 50050768014023EC (port 4)
port_status active
port_speed 4Gb
port_id 50050768013023EC (port 3)
port_status active
port_speed 2Gb
port_id 50050768012023EC (port 2)
port_status active
port_speed 2Gb
port_id 50050768011023EC (port 1)
port_status active
port_speed 2Gb
```

Repeat these steps for each of the SVC nodes in your SVC cluster.

4.2.3 Prepare the SAN fabric

We need to change the SAN zoning to allow the SVC to access the existing storage controller and the host whose data is to be virtualized.

Refer to *IBM System Storage SAN Volume Controller 4.1.0 Configuration Guide*, SC26-7902, for a description of the required SAN zones for implementing the SVC in an existing environment. In our example environment these are the RED1 and BLUE zones.

If you are planning to migrate the data to a new storage controller (as in our example), you should also zone the SAN to allow the SVC to access the new storage controller (RED2).

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendor’s SAN equipment.

Note: It is assumed that the BLACK zone (for SVC node to SVC node connection) already exists in the SAN. If not, create it now.

4.2.4 Prepare the DS8000 storage controllers

We need to configure access from the SVC to the LUN on the existing DS8000 storage controller. We also need to create some LUNs on the DS8000 storage controller to be our migration targets. The new LUNs can also be used as the SVC quorum disks.

Configure the new storage controller

We need to create one or more LUNs on the new storage controller. These LUNs will be used as the target volumes for our volume migration from the existing storage controller.

The general recommendation for SVC use is that each of these LUNs be configured to the full capacity of the containing RAID array. Therefore, if possible, create one LUN per RAID array.

Change the LUN masking on the storage controller to allow the SVC to access the LUNs that are to be virtualized.

Refer to the IBM DS8000 administration guide for detailed instructions on how to configure LUNs and LUN masking on the DS8000 storage controller. It is available at:

<http://www-1.ibm.com/support/docview.wss?rs=1113&uid=ssg1S7001073>

Since image mode volumes (from the existing storage controller) cannot be used as quorum disks, the new managed mode MDisks will also be used as the SVC quorum volumes.

4.2.5 Prepare the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized.

Use the **svcinfo lscontroller** command to list the storage controllers.

Example 4-3 shows the output of this command on our system.

Example 4-3 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfo lscontroller
```

id	controller_name	ctrl_s/n	vendor_id	product_id_low	product_id_high
0	controller0		IBM	1815	FAStT
1	ds8000		IBM	2107900	

Create an MDG for the new storage controller’s LUNs

Since we have already correctly zoned the SAN and configured LUN masking on the back-end storage controllers, the LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

Use the **svctask detectmdisk** command to rescan for newly available MDisks.

Note: This command produces no user feedback.

If the expected MDiskS are still not visible to the SVC, double check the SAN zoning and storage controller LUN masking.

Use the **svcinfolsmdisk** command to list the MDiskS on the new storage controller. Example 4-4 shows the output of this command on our system.

Example 4-4 List MDiskS on the new storage controller

```
IBM_2145:GREEN_SVC:admin>svcinfolsmdisk -delim , -filtervalue controller_name=ds8000

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
1,mdisk1,online,unmanaged,,,195.6GB,0000000000005000,ds8000,494...837
2,mdisk2,online,unmanaged,,,195.6GB,0000000000005100,ds8000,494...837
3,mdisk3,online,unmanaged,,,195.6GB,0000000000005200,ds8000,494...837
4,mdisk4,online,unmanaged,,,195.6GB,0000000000005300,ds8000,494...837
5,mdisk5,online,unmanaged,,,195.6GB,0000000000005400,ds8000,494...837
6,mdisk6,online,unmanaged,,,195.6GB,0000000000005500,ds8000,494...837
7,mdisk7,online,unmanaged,,,195.6GB,0000000000005700,ds8000,494...837
```

The command shown in this example has been edited for legibility. Note that the MDiskS for the new LUNs are in *unmanaged* mode.

Use the **svctask mkmdiskgrp** command to add the unmanaged MDiskS to a new MDG. Example 4-5 shows the output of this command on our system.

Example 4-5 Add the new controller's LUNs to a MDG

```
IBM_2145:GREEN_SVC:admin>svctask mkmdiskgrp -name ds8kgrp -ext 512 -mdisk
mdisk1:mdisk2:mdisk3:mdisk4:mdisk5:mdisk6:mdisk7

MDisk Group, id [1], successfully created
```

Use the **svctask lsmdisk** command to check that the MDG was created successfully. Example 4-6 shows the output of this command on our system.

Example 4-6 List the MDiskS that compose a given MDG

```
IBM_2145:GREEN_SVC:admin>svcinfolsmdisk -delim , -filtervalue mdisk_grp_name=ds8kgrp

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
1,mdisk1,online,managed,1,ds8kgrp,195.6GB,0000000000005000,ds8000,494...3837
2,mdisk2,online,managed,1,ds8kgrp,195.6GB,0000000000005100,ds8000,494...837
3,mdisk3,online,managed,1,ds8kgrp,195.6GB,0000000000005200,ds8000,494...837
4,mdisk4,online,managed,1,ds8kgrp,195.6GB,0000000000005300,ds8000,494...837
5,mdisk5,online,managed,1,ds8kgrp,195.6GB,0000000000005400,ds8000,494...837
6,mdisk6,online,managed,1,ds8kgrp,195.6GB,0000000000005500,ds8000,494...837
7,mdisk7,online,managed,1,ds8kgrp,195.6GB,0000000000005700,ds8000,494...837
```

The command shown in this example has been edited for legibility. Note that the MDiskS for the new LUNs are now in *managed* mode.

Define a host to the SVC

Before the host can access any VDiskS it needs to be defined to the SVC.

Use the **svcinfolshbaportcandidate** command to list any unassigned host WWPNS that are visible to the SVC. Example 4-7 shows the output of this command on our system.

Example 4-7 List unassigned host WWPNs that are visible to the SVC

```
IBM_2145:GREEN_SVC:admin>svcinfolshbaportcandidate

id
10000000C93618F7
210000E08B09E1FD    (This is the first HBA on our host)
210000E08B1A3FC6
10000000C937A664
210000E08B1A2EC6
210000E08B0BC71E    (This is the second HBA on our host)
```

Note: You will need to check on the host (or the SAN switch) to determine which WWPN belongs to the host with which you are working.

Use the **svctask mkhost** command to define the host to the SVC. Example 4-8 shows the output of this command on our system.

Example 4-8 Define a host to the SVC

```
IBM_2145:GREEN_SVC:admin>svctask mkhost -name Moe -hbawwpn
210000E08B09E1FD:210000E08B0BC71E -type generic
```

```
Host id [3] successfully created
```

Use the **svcinfolshost** command to check that the host was defined successfully. Example 4-9 shows the output of this command on our system.

Example 4-9 Check a host definition on the SVC

```
IBM_2145:GREEN_SVC:admin>svcinfolshost Moe

id 3
name Moe
port_count 2
type generic
mask 1111
iogrp_count 4
WWPN 210000E08B09E1FD
node_logged_in_count 8
state active
WWPN 210000E08B0BC71E
node_logged_in_count 8
state active
```

4.3 Storage virtualization

The preparation is now complete and we are ready to proceed with the storage virtualization.

Configure the existing storage controller

We need to change the LUN masking on the storage controller to allow the SVC to access the LUNs that are to be virtualized. There are several steps that need to be done first.

The administrative interface on the DS4800 can be accessed through the IBM System Storage DS4000/FastT Storage Manager.

For more information, see:

<http://www-03.ibm.com/servers/storage/support/disk/ds4800/stormgr1.html>

Figure 4-2 shows the DS4800 administrative graphical user interface (GUI).

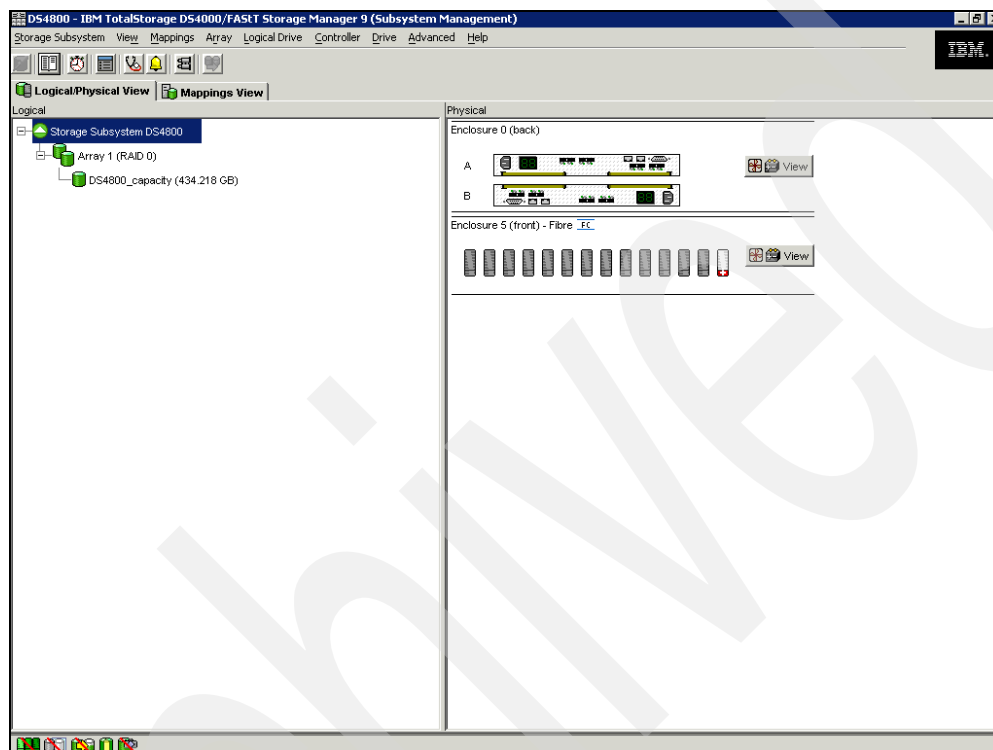


Figure 4-2 DS4800 administrative GUI

Figure 4-3 shows the existing LUN ID on the DS4800 for the host. Record this number.

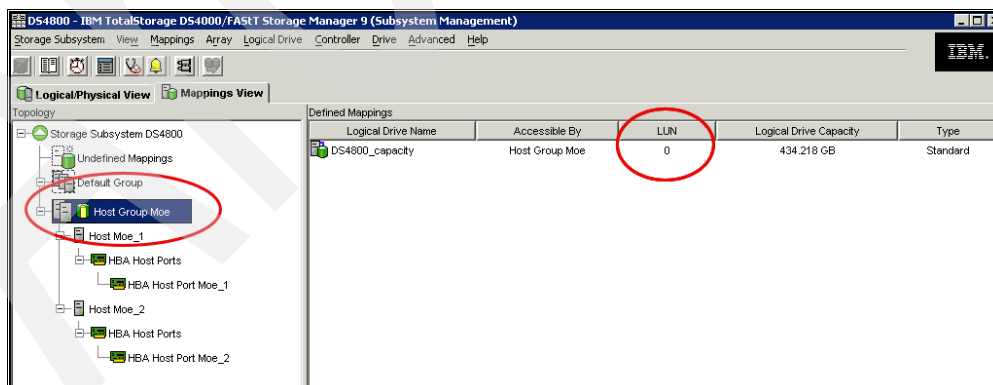


Figure 4-3 Checking the SCSI LUN ID for the existing volume

Record the SCSI LUN IDs that the existing storage controller presents to the host for the LUNs that are to be virtualized. We will need these ID numbers later to allow the SVC to present the same SCSI IDs to the host. This will allow the host to see the virtualized volumes as though they were the original, directly connected, volumes.

We will also need to identify the logical drive ID for the disk ID for later identification on the SVC.

Check the disk settings to see if they are compliant with the SVC. For more details, see:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

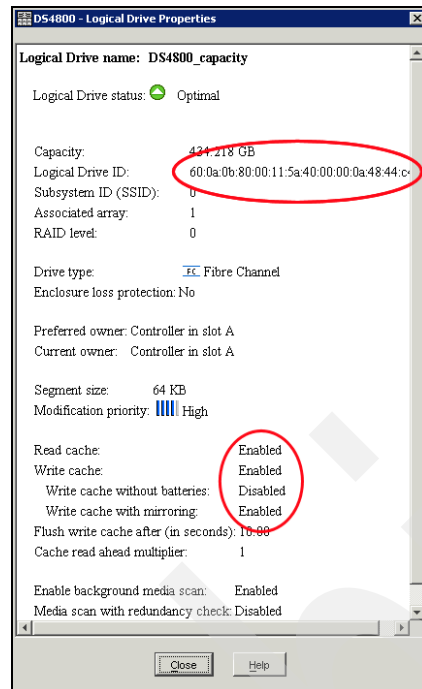


Figure 4-4 Recording the Logical Drive ID

Configure the host system

The SAN Volume Controller uses the IBM Subsystem Device Driver (SDD) multi-path driver to manage connectivity between the host and the virtualized storage.

In our environment it is necessary to unload the existing RDAC device driver (Redundant Disk Array Controller) before installing the IBM SDD driver.

Attention: All applications must be stopped before RDAC is removed. This ensures that there is no Read/Write activity to the disks.

Windows example

Start the Windows Control Panel to remove the RDAC multi-path device driver.

Figure 4-5 on page 64 shows the start of the uninstall process for RDAC on a Windows host.

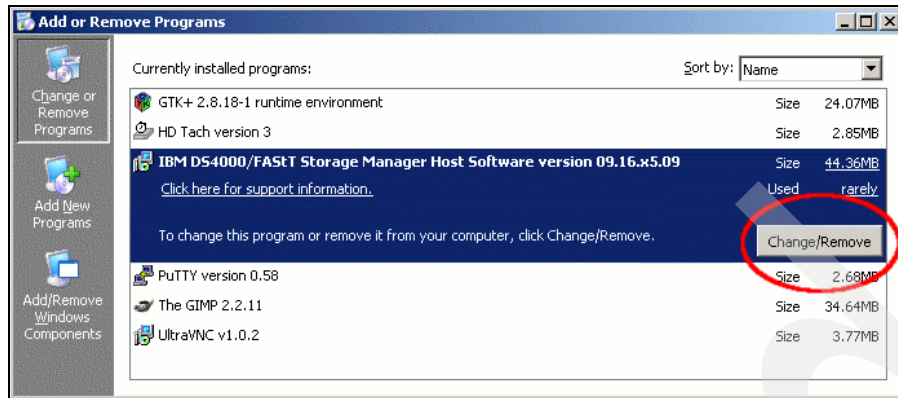


Figure 4-5 Uninstalling RDAC on Windows

When the uninstall application starts, choose to perform a **Complete Uninstall**, and follow the prompts to completion.

Reboot the host when prompted after uninstalling the RDAC multi-path driver.

Linux example

There are two steps to uninstalling the RDAC multipath driver on Linux:

1. Uninstall the RDAC multipath software.

Since RDAC is not installed using the Red Hat software package manager (RPM) it cannot be uninstalled using the normal RPM commands.

To uninstall the RDAC software you need access to the RDAC installation files, and to run the uninstall command.

Enter the **make uninstall** command in the installation directory (as root) to uninstall the RDAC package.

Example 4-10 shows the output of this command on our system.

Example 4-10 Uninstalling RDAC on a Linux host

```
[root@marie ~]# cd /home/ibm/rdac/linuxrdac-09.01.B5.30/

[root@marie linuxrdac-09.01.B5.30]# make uninstall

Preparing to uninstall MPP driver...
make[1]: Entering directory `/home/ibm/rdac/linuxrdac-09.01.B5.30'
make[1]: Leaving directory `/home/ibm/rdac/linuxrdac-09.01.B5.30'
MPP driver package has been sucessfully removed from your system.
```

Note that the path to the RDAC source will probably differ on your system.

2. Remove the RDAC entry from the boot menu.

When RDAC was originally installed a custom entry was made to the boot loader's configuration file. We need to remove this entry since we have uninstalled RDAC.

Edit the boot loader's configuration file (as root):

```
vi /etc/grub.conf
```

Comment out the RDAC entries as shown in Example 4-11.

Example 4-11 Removing the RDAC entry from grub.conf on a Linux host

```
#title Red Hat Enterprise Linux ES (2.6.9-34.ELsmp) with RDAC
#       root (hd0,0)
#       kernel /vmlinuz-2.6.9-34.ELsmp ro root=/dev/VolGroup00/LogVol100 rhgb quiet
#       initrd /mpp-2.6.9-34.ELsmp.img
```

Note that the RDAC entries in the boot loader menu will probably differ on your system.

You need not reboot the Linux host after uninstalling the RDAC multi-path driver. However, if you are using a different SAN storage controller and a different multipath driver, you should follow the other vendor's uninstallation instructions.

Install the SDD multipath driver

The IBM SDD multipath device driver provides support for failover and load balancing on the SVC managed volumes.

Windows example

Installing the SDD multipath driver on Windows is relatively straightforward. The steps are summarized as follows:

1. Download the SDD driver installation files.
2. Run the **setup.exe** (as Administrator) to start the installation.
 - a. Click **yes** to continue with the installation.
 - b. When prompted, choose **Continue Anyway** to install the SDD drivers.
 - c. Click **no** to reboot the host at the end of the installation.

The host should remain briefly offline while we reconfigure the SAN. While the host is offline we will remove direct access to the existing LUN and enable access to the virtualized LUN through the SVC.

It is advisable to configure the server to prevent the automatic starting of any applications on the next boot. This is to prevent any problems with applications trying to access the SAN volumes.

Linux example

Installing the SDD multipath driver on Linux is usually straightforward, but there are a few additional configuration steps required in a SAN boot environment.

First, install the SDD multipath driver on the Linux host using these steps:

1. Download the SDD driver installation files.
2. Run the following command (as root) to start the software installation:

```
rpm -Uph IBMsdd-1.6.1.0-3.i686.rhel4.rpm
```

Example 4-12 shows the output of this command on our system.

Example 4-12 Installing the SDD driver on a Linux system

```
[root@marie sdd]# rpm -Uph IBMsdd-1.6.1.0-3.i686.rhel4.rpm
##### [100%]
##### [100%]
Added following line to /etc/inittab:
srv:345:respawn:/opt/IBMsdd/bin/sdssrv > /dev/null 2>&1
```

Next, configure your system to be able to use the SDD driver at boot time.

Note: The following instructions are specifically for Red Hat Enterprise Linux 4.0 with LVM2 on the x86 platform. They do not apply for any other Linux version.

It is *strongly* recommended that you refer to Chapter 5 of the *SDD Multipath Subsystem Device Driver User's Guide* - SC30-4131-00 before attempting this procedure.

Perform the following steps to modify the boot/root and other devices for booting using the SDD driver:

1. Use **pvdisk** to get the physical volume for the root and swap lvm volume groups, as shown in Example 4-13.

Example 4-13 pvdisk

```
--- Physical volume ---
PV Name           /dev/sda2
VG Name           VolGroup00
PV Size           9.88 GB / not usable 0
Allocatable       yes
PE Size (KByte)   32768
Total PE          316
Free PE           2
Allocated PE      314
PV UUID           z4x552-xT4J-1sAx-JP4U-oY70-nJr3-JZ7P0e
```

In this procedure /dev/sda2 (sda) is the device that will be used for /dev/vpatha2 (vpatha).

2. Modify the /etc/fstab file, ensuring that:
 - LABEL= is not being used
 - /boot is mounted on a vpath device

Because Red Hat writes labels to the disk and uses labels in the /etc/fstab, the boot (/boot) device might be specified as a label; that is, LABEL=/boot. This can, however, be a different label other than LABEL=/boot. Check for a line in the /etc/fstab where /boot is mounted and change it to the correct vpath device. Also ensure that any other device specified with the LABEL= feature is changed to a /dev/sd or /dev/vpath device. Red Hat does not recognize LABEL= in a multipath environment.

There is a one-to-one correlation between SCSI disk and vpath minor devices; that is, sda1 and vpatha1. Major devices, however, might not correlate; that is, sdb1 could be vpathd1.

Because /boot was installed on /dev/sda1 and we have made vpatha1 correspond to sda in the /etc/vpath.conf file, then /dev/vpatha will be the mount device for /boot (see Example 4-14).

Example 4-14 Mount device for boot

```
Change:
/dev/VolGroup00/LogVol100    /      ext3    defaults    1 1
LABEL=/boot                  /boot  ext3    defaults    1 2
/dev/VolGroup00/LogVol101    swap   swap    defaults    0 0
To:
/dev/VolGroup00/LogVol100    /      ext3    defaults    1 1
/dev/vpatha1                 /boot  ext3    defaults    1 2
/dev/VolGroup00/LogVol101    swap   swap    defaults    0 0
```

3. Modify the `/boot/grub/menu.lst` file. Add an entry for the SDD/LVM boot using the new `initrd` file that we will create.

Example 4-15 Example file

```
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu

title Red Hat Enterprise Linux ES (2.6.9-34.ELsmp) with SDD
    root (hd0,0)
    kernel /vmlinuz-2.6.9-34.ELsmp ro root=/dev/VolGroup00/LogVol100 rhgb quiet
    initrd /initrd-2.6.9-34.ELsmp.sdd.img

title Red Hat Enterprise Linux ES (2.6.9-34.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.6.9-34.ELsmp ro root=/dev/VolGroup00/LogVol100 rhgb quiet
    initrd /initrd-2.6.9-34.ELsmp.img
```

4. Modify `/etc/lvm/lvm.conf`.

Change: `filter = ["a/*/"]`
To: `filter = ["a/vpath*/", "r/sd*/"]`
In the `types =` section, add an entry for `vpath`:
 `types = ["vpath", 16]`

5. Start SDD (as root)

```
# sdd start
```

This will create the `/etc/vpath.conf` file. Ensure that `vpatha` is the root device. We must get the root's physical device's LUN ID using the **`cfgvpath query`** command (in this procedure `sda` is the root device).

The **`cfgvpath query`** command will produce an output similar to that shown in Example 4-16.

Example 4-16 `cfgvpathquery` command

```
/dev/sdd (8,48) host=1 ch=0 id=1 lun=0 vid=IBM pid=2145 serial=600...006
lun_id=600...006 ctrl_flag=1 ctrl_nbr=0 df_ctrl=0
```

Note that some data from the output has been modified for ease of reading.

The `lun_id` for `/dev/sda` is `600507680183804668000000000000006` (shown as `600...006`).

Edit the `/etc/vpath.conf` file using the correct `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD).

For example:

```
vpatha 600507680183804668000000000000006
```

6. Prepare the `initrd` file.

The `initrd` file refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the method shown in Example 4-17.

Example 4-17 Determining initrd

```
# ls -lA /boot | grep initrd | grep $(uname -r)
# cd /boot
# cp [initrd file] to initrd.vp.gz
# gunzip initrd.vp.gz
# mkdir /boot/mnt
```

Note: For the rest of this procedure we will be working from /boot/mnt.

7. Change directory to /boot/mnt and unarchive the initrd image.

Example 4-18 Changing directory and unarchive

```
# cd /boot/mnt
# cpio -iv < ../initrd.vp
```

8. Make additional directories in /boot/mnt, as shown in Example 4-19.

Example 4-19 Making directories

```
# mkdir mnt
# mkdir -p opt/IBMsdd/bin
# chmod -R 640 opt/IBMsdd
# mkdir -p lib/tls
```

9. Copy the /etc files to /boot/mnt/etc, as shown in Example 4-20.

Example 4-20 Copying files

```
# cp /etc/vpath.conf /boot/mnt/etc/
# cp /etc/group /boot/mnt/etc/
# cp /etc/passwd /boot/mnt/etc/
# cp /etc/nsswitch.conf /boot/mnt/etc/
# cp /etc/fstab /boot/mnt/etc/
# cp /etc/lvm/lvm.conf /boot/mnt/etc/lvm/
```

10. Modify the /boot/mnt/etc/nsswitch.conf file (for rhel4, this might already be done).

```
Change:  passwd: compat
To:      passwd: files
```

```
Change:  group: compat
To:      group: files
```

11. Copy /opt/IBMsdd/bin/cfgvpath to /boot/mnt/opt/IBMsdd/bin/ and change the file permissions to 755, as shown in Example 4-21.

Example 4-21 Copy and change file permissions

```
# cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
# chmod 755 /boot/mnt/opt/IBMsdd/bin/*
```

12. Copy required library files for cfgvpath.

Use the **ldd** command to determine the library files and locations.

Example:

```
# ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}'
```

The **ldd** command will return a list similar to that shown in Example 4-22.

Example 4-22 Determine library files and locations

```
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

These files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

13. Copy the correct `sdd-mod` to the `initrd` file system.

Use the **uname -r** command to determine the correct `sdd-mod`, as shown in Example 4-23.

Example 4-23 Determine correct sdd-mod

```
# uname -r  
2.6.9-34.ELsmp  
  
# cp /opt/IBMsdd/sdd-mod.ko-2.6.9-34.ELsmp /boot/mnt/lib/sdd-mod.ko
```

14. Copy the following binaries and change their permissions to 755, as shown in Example 4-24.

Example 4-24

```
# cp /bin/mknod /boot/mnt/bin/  
# cp /bin/mount /boot/mnt/bin/  
# cp /bin/umount /boot/mnt/bin/  
# cp /bin/cp /boot/mnt/bin/  
# chmod 755 /boot/mnt/bin/*
```

15. Copy the required library files for each binary that you copied to the `initrd` file system.

Use the **ldd** command to determine the library files and locations, as shown in Example 4-25.

Note: There might be duplications during copying because many binaries use the same libraries.

Example 4-25 Determine files and locations

```
# ldd /bin/mknod | awk '{print $(NF-1)}'  
  
/lib/libselinux.so.1  
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

These files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

Also, copy the library files shown in Example 4-26 to `/boot/mnt/lib/`.

Example 4-26

```
# cp /lib/libproc-3.2.3.so /boot/mnt/lib/  
# cp /lib/libtermcap.so.2 /boot/mnt/lib/  
# cp /lib/libnss_files.so.2 /boot/mnt/lib/
```

16. Modify the `/boot/mnt/init` file.

Add the lines shown in Example 4-27 just before the statement: [echo "Loading SDD module"].

Example 4-27 Modify init file

```
echo "Loading SDD module"
insmod /lib/sdd-mod.ko echo
"Creating vpath devices"
/opt/IBMsdd/bin/cfgvpath makenodes
```

Ensure that an updated copy of vpath.conf is copied to the /root file system by using the following syntax to mount the root file system:

```
/bin/mount -o rw -t [fstype] [device] /mnt
```

Add the lines shown in Example 4-28 just after [insmod /lib/dm-snapshot.ko]. The values used for the [fstype] and [device] here are only examples. Use the correct values for the system that is being configured.

Example 4-28 Adding lines

```
/bin/mount -o rw -t ext3 /dev/rootVolGroup/rootVol /mnt
/bin/cp /etc/vpath.conf /mnt/etc/
/bin/umount /mnt
```

17. Repackage the updated initrd file.

Use **cpio** to archive the /boot/mnt directory and **gzip** to compress it.

Example 4-29 Archive and compress directory

```
# find . | cpio -H newc -vo > ../initrd-2.6.9-34.ELsmp.sdd.img
# cd /boot
# gzip initrd-2.6.9-34.ELsmp.sdd.img
# mv initrd-2.6.9-34.ELsmp.sdd.img.gz initrd-2.6.9-34.ELsmp.sdd.img
```

18. Shut down the Linux host.

```
# shutdown now
```

Attention: The SDD driver only supports a specific list of Linux kernel versions. If you use automatic system updates (for example, Red Hat **up2date**), you should disable the automatic updates for kernel levels.

Shut down the host

The host should remain briefly offline while we reconfigure the SAN. While the host is offline we will remove direct access to the existing LUN and enable access to the virtualized LUN through the SVC.

It is advisable to configure the server to prevent the automatic starting of any applications on the next boot. This is to prevent any problems with applications trying to access the SAN volumes.

Note: Although some parts of the storage virtualization process may be conducted with the host system online, we choose to shut down the host to minimize risk during the change.

4.4 Configure the OEM controller

The following steps are used to configure the existing controller to allow the SVC access:

1. Figure 4-6 shows the preliminary screen for adding a host group for the SVC to the DS4800.

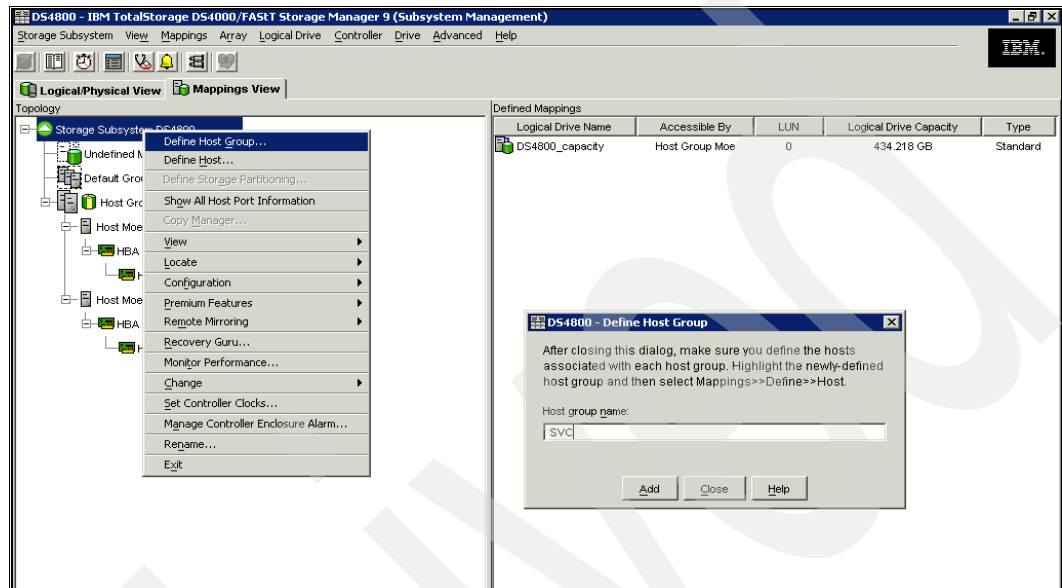


Figure 4-6 Adding an SVC host group to the DS4800

2. Create a host for each SVC node. (Figure 4-7 shows creating a host within a host group.) Each node will have 4 WWPNs.




Figure 4-7 Creating a host

3. Associate the node's WWPNs with the new host. Figure 4-8 on page 72 shows the naming of the host, in this case Node1. It also displays the WWPNs we have associated with the host.




Figure 4-8 Associating the node's WWPNs on the DS4800

4. Add the alias for the WWPN associated with the node (Figure 4-9). The alias must be unique for each WWPN.




Figure 4-9 Setting the alias for the WWPN

5. Choose which host type to use for the SVC (Figure 4-10 on page 73). Select **IBM TS SAN VCE**.




Figure 4-10 choosing the host type

6. Remap the LUN. Figure 4-11 shows changing the mapping of the LUN from the host to the SVC. Make sure that the LUN ID is the same as the directly connected LUN ID.




Figure 4-11 Assigning LUNs to the SVC on the DS4000

Important: Make sure that the host is powered down at this point. This action will remove the host's access to the LUN.

When prompted, confirm that you want to accept the changes.

4.5 Migration to the SAN Volume controller

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SVC.

4.5.1 Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SAN volume controller. In our example environment we needed to disable the GREEN zone.

At this point the SAN should already be configured to allow the host to access the SVC.

Refer to Appendix B, "SAN zoning changes" on page 257 for examples of the commands used to define the required SAN zones on the SAN equipment.

4.5.2 Configure the SAN Volume Controller

Now that the other systems have been prepared, we can proceed with the virtualization of the existing storage volume.

Create an MDG for the image mode volumes

Before we can configure the existing storage controller's LUN as an image mode VDisk, we should define a new MDG to logically house those volumes.

Use the `svctask mkmdiskgrp` command to create an empty MDG for your image mode volumes. Example 4-30 shows the output of this command on our system.

Example 4-30 Create an empty MDG for image mode volumes

```
IBM_2145:GREEN_SVC:admin>svctask mkmdiskgrp -name imagegrp -ext 512
```

```
MDisk Group, id [9], successfully created
```

Although it is not strictly necessary to create a separate MDG for image mode volumes, we believe that it is good practice.

Important: To be able to migrate volumes between MDGs, the extent size of the source and target MDGs must be the same.

Create image mode VDisks from existing storage controller's LUNs

Now we can create the image mode VDisk using the existing unmanaged MDisk, which was previously connected directly to the host.

Use the `svcinfo lsmdisk` command to list the MDisk. Example 4-31 on page 75 shows the output of this command on our system.

Example 4-31 List MDisk on the existing storage controller

```
IBM_2145:GREEN_SVC:admin>svcinfolsmdisk 9
id 9
name mdisk9
status online
mode unmanaged
mdisk_grp_id
mdisk_grp_name
capacity 434.2GB
quorum_index
block_size 512
controller_name controller0
ctrl_type 4
ctrl_WWNN 200400A0B80BDE24
controller_id 0
path_count 0
max_path_count 2
ctrl_LUN_# 0000000000000000
UID 600a0b8000115a4000000a4844c4b1100000000000000000000000000000000
preferred_WWPN 201400A0B80BDE24
active_WWPN 202500A0B80BDE24
```

The commands shown in Example 4-31 have been edited for legibility. Note that the MDisk for the existing LUN is in *unmanaged* mode.

Use the **svctask mkvdisk** command to create the image mode VDisk. Example 4-32 shows the output of this command on our system.

Example 4-32 Define a new image mode VDisk

```
IBM_2145:GREEN_SVC:admin>svctask mkvdisk -mdiskgrp imagegrp -iogrp 0 -vtype image
-name Moe01 -mdisk mdisk9
```

Virtual Disk, id [1029], successfully created

Attention: If you were to add the existing volumes to an MDG as normal *striped* MDisk then all data on those volumes would be lost. Since we intend to virtualize previously existing volumes with data intact, we *must* create *image mode virtual disks* instead.

Use the **svcinfolsvdisk** command to check that the VDisk was created successfully. Example 4-33 shows the output of this command on our system.

Example 4-33 Display the properties of the new image mode VDisk

```
IBM_2145:GREEN_SVC:admin>svcinfolsvdisk Moe01

id 1029
name Moe01
IO_group_id 0
IO_group_name io_grp0
status online
mdisk_grp_id 3
mdisk_grp_name ds8kgrp
capacity 434.2GB
type image
```

```
formatted no
mdisk_id
mdisk_name
FC_id
FC_name
RC_id
RC_name
vdisk_UID 600a0b8000115a4000000a4844c4b11
throttling 0
preferred_node_id 2
fast_write_state empty
cache readwrite
udid
```

Map the image mode VDisk to the Host

Next, we need to define a relationship between the new image mode VDisk and our host system. This is the equivalent of LUN masking (or storage partitioning) on a normal storage controller.

Use the **svctask mkvdiskhostmap** command to map the new VDisk to the host. Example 4-34 shows the output of this command on our system.

Example 4-34 Map a VDisk to a Host

```
IBM_2145:GREEN_SVC:admin> svctask mkvdiskhostmap -host Moe -scsi 0 Moe01
```

```
Virtual Disk to Host map, id [1], successfully created
```

Note: It is important to define the same SCSI ID number to the VDisk as the existing storage controller has presented to the host. This will allow the virtualized LUN to appear as the original LUN when later exposed to the host.

4.5.3 Configure the host system

Now that the SAN has been reconfigured, you can start the host system.

When the host reboots it will see multiple paths to the SAN volume.

After the host reboots (with SDD loaded), it should see a single path to the virtualized storage volume.

Note: If you disabled the automatic start of any applications on the host you should re-enable them now.

Windows

Figure 4-12 shows the Windows device manager view after SDD is installed.

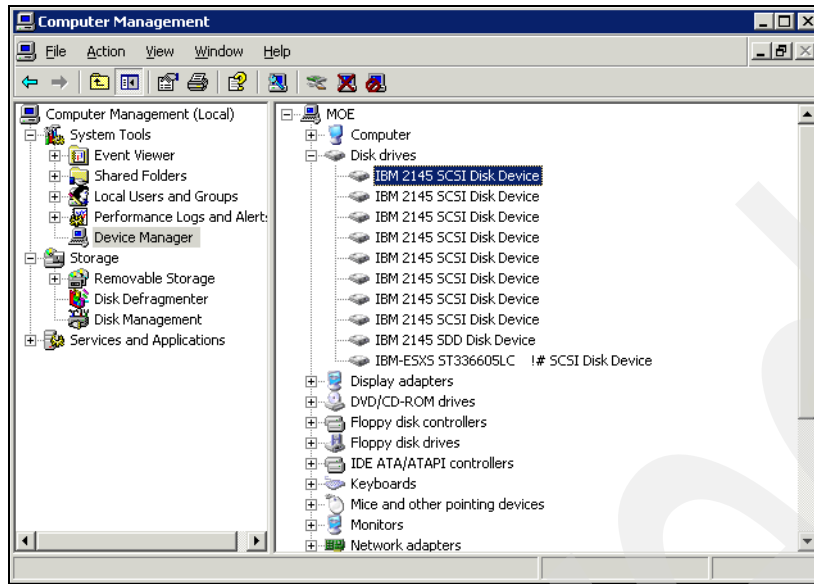


Figure 4-12 Windows device managed view with SDD installed

You can use the IBM SDD command line to check that the SAN volumes are available. Example 4-35 shows the output of this command on our system.

Example 4-35 Checking SAN devices with SDD

```
C:\Program Files\IBM\SDDDSM>datapath query device
PTotal Devices : 1
DEV#: 0 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 60050768019A011F6000000000000002
=====
```

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	3556	0
1	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
2	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	1194	0
3	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	1275	0
4	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	1243	0
5	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	1232	0
6	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
7	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	0	0

The data on the SAN disk is now available on the host for application and user access.

Linux example

On a Linux host (with SDD installed) you can use the commands shown in Example 4-36 to view the available disk drives.

Example 4-36

```
cat /proc/scsi/scsi
datapath query device
```

Example 4-37 shows the output of these commands on our system.

Example 4-37 List available SCSI LUNs on a Linux host

```
[ibm@marie ~]# cat /proc/scsi/scsi
```

Attached devices:

```
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 2145      Rev: 0000
  Type:   Direct-Access      ANSI SCSI revision: 04
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 2145      Rev: 0000
  Type:   Direct-Access      ANSI SCSI revision: 04
Host: scsi1 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 2145      Rev: 0000
  Type:   Direct-Access      ANSI SCSI revision: 04
Host: scsi1 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 2145      Rev: 0000
  Type:   Direct-Access      ANSI SCSI revision: 04
```

```
[root@marie ~]# datapath query adapter
```

Active Adapters :2

Adpt#	Name	State	Mode	Select	Errors	Paths	Active
0	Host0Channel0	NORMAL	ACTIVE	8477	0	2	2
1	Host1Channel0	NORMAL	ACTIVE	7038	0	2	2

```
[root@marie ~]# datapath query device
```

Total Devices : 1

```
DEV#: 0 DEVICE NAME: vpatha TYPE: 2145 POLICY: Optimized Sequential
SERIAL: 60050768018380466800000000000006
```

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Host0Channel0/sda	OPEN	NORMAL	1	0
1	Host0Channel0/sdb	OPEN	NORMAL	8476	0
2	Host1Channel0/sdc	OPEN	NORMAL	0	0
3	Host1Channel0/sdd	OPEN	NORMAL	7055	0

The data on the SAN disk is now available on the host for application and user access.

Note: The storage volume has now been virtualized, but still resides on the original storage controller. Access to the volumes is managed through the SAN Volume Controller.

In the next section we migrate the storage volume to the new storage controller.

4.6 Move SVC managed volumes between storage controllers

The SAN volume controller enables the migration of a storage volume between storage controllers with no disruption to the host systems.

Configure the SAN Volume Controller

In this example we migrate an image mode VDisk from an MDG on one storage controller to another MDG on the new storage controller. During the migration the image mode VDisk is automatically converted to striped mode.

Migrate the VDisk to the new storage controller

Use the **svctask migratevdisk** command to start the volume migration. Example 4-38 shows the output of this command on our system.

Example 4-38 Migrating a VDisk between MDGs

```
IBM_2145:GREEN_SVC:admin> svctask migratevdisk -vdisk Moe01 -mdiskgrp ds8kgrp
```

Note: This command produces no user feedback.

Remember that for volume migration to occur the extent sizes of the source and target MDGs must be the same.

Monitor the VDisk migration

Use the **svcinfolsmigrate** command to monitor the progress of the migration. Example 4-39 shows the output of this command on our system.

Example 4-39 Monitoring the VDisk migration process

```
IBM_2145:GREEN_SVC:admin>svcinfolsmigrate
```

```
migrate_type MDisk_Group_Migration
progress 60
migrate_source_vdisk_index 0
migrate_target_mdisk_grp 1
max_thread_count 4
```

You can also use the **svcinfolsvdiskextent** command to view the location of the MDisk extents that comprise the VDisk. Example 4-40 shows the output of this command on our system.

Example 4-40 View the MDisk extents of a VDisks during migration

```
IBM_2145:GREEN_SVC:admin>svcinfolsvdiskextent Moe01
id          number_extents
0           12974 (This was the original image mode VDisk/MDisk)
3           7
9           7
18          7
26          7
37          7
42          8
54          8
```

As the VDisk migration proceeds, the number of extents on the original MDisk will decrease; at the same time, the number of extents on the MDisk in the target MDG will increase.

When the volume migration is complete no extents will remain on the original LUN.

Note: The storage volume now resides on the new storage controller. Access to the virtualized volume is managed through the SAN Volume Controller.

Remove the original volume

After the volume migration is complete, the original image mode MDisk remains defined in the original MDG even though it contains no extents. This MDisk can now be removed from the MDG, which changes its status to unmanaged mode.

Use the **svctask rmmdisk** command to remove the original MDisk from the MDG. Example 4-41 shows the output of this command on our system.

Example 4-41 Remove an MDisk from a MDG

```
IBM_2145:GREEN_SVC:admin>svctask rmmdisk -mdisk mdisk9 imagegrp
```

From that point the MDisk could be reassigned to another MDG or deleted from the storage controller.

Note: The **rmmdisk** command does not delete the MDisk or the LUN that it represents, it only removes it from the MDG and returns it to *unmanaged* mode. If you wish to remove the MDisk entirely, you will need to do so on the storage controller.

Use the **svcinfo lsmdisk** command to view the status of the MDisk. Example 4-42 shows the output of this command on our system.

Example 4-42 Display the properties of an MDisk

```
IBM_2145:GREEN_SVC:admin>svcinfo lsmdisk 9
id 9
name mdisk9
status online
mode unmanaged
mdisk_grp_id
mdisk_grp_name
capacity 434.2GB
quorum_index
block_size 512
controller_name controller0
ctrl_type 4
ctrl_WWNN 200400A0B80BDE24
controller_id 0
path_count 0
max_path_count 2
ctrl_LUN_# 0000000000000000
UID 600a0b8000115a4000000a4844c4b1100000000000000000000000000000000
preferred_WWPN 201400A0B80BDE24
active_WWPN 202500A0B80BDE24
```

If you wish, you can now delete the original LUN from the original storage controller.

Note: You should ensure that you do not destroy any controller LUNs until you have deleted them from the MDisk group and returned them to *unmanaged* mode.



Integrating the SVC with HDS

In this chapter, we show how to integrate the SVC in an existing environment with various Hitachi Data Systems (HDS) storage controllers.

5.1 Integrating HDS Lightning with the SVC

In this section, we demonstrate the integration of the SVC in an existing environment with an HDS Lightning storage controller and a single SAN-connected Windows host. The existing storage volumes on the HDS controller will be virtualized using the SVC.

5.1.1 Scenario description

In this scenario, we start with an HDS Lightning storage controller and a SAN-attached Windows host. We then introduce the SAN Volume Controller in order to virtualize the back-end storage. The systems described in this section include:

- ▶ HDS Lightning 9970V storage controller
- ▶ IBM b-type family switches

Figure 5-1 shows the relationships between the systems used in this scenario.




Figure 5-1 SVC integration example with an existing HDS Lightning storage controller

5.1.2 Confirm support for your environment

This section provides information about configuring the SVC in an existing environment with an HDS Lightning storage controller.

Prior to integrating the SVC, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SAN Volume Controller supports most models of the HDS Lightning. Certain models of HDS Lightning are equivalent to Sun StorEdge™ and HP XP models; therefore, the SAN Volume Controller also supports models of Sun StorEdge and HP XP.

Table 5-1 lists the models of HDS Lightning, Sun StorEdge, and HP XP that are supported by the SAN Volume Controller.

Table 5-1 Supported models of the HDS Lightning, Sun StorEdge and HP XP

HDS Lightning models	Sun StorEdge models	HP XP models
Lightning 9910	StorEdge 9910	XP48
Lightning 9960	StorEdge 9960	XP512
Lightning 9970V	StorEdge 9970	XP128
Lightning 9980V	StorEdge 9980	XP1024

See the following Web site for specific HDS Lightning supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Note: Concurrent upgrade of the controller firmware is not supported with the SAN Volume Controller.

Sharing an HDS Lightning between a host and the SVC

You can share an HDS Lightning 99xxV between a host and the SVC, subject to the following limitations:

- ▶ A controller port cannot be shared between a host and a SAN Volume Controller. If a controller port is used by a SAN Volume Controller, it must not be present in a switch zone that allows a host to access the port.
- ▶ Logical units (LUNs) cannot be shared between a host and a SAN Volume Controller.

Note: A single host must not be connected to both the SVC and an HDS controller at the same time because the multipathing drivers IBM Subsystem Device Driver (SDD) and Hitachi HiCommand Dynamic Link Manager (HDL) cannot coexist on the host.

Supported topologies

The SAN Volume Controller supports connection to the HDS Lightning under the following conditions:

- ▶ The SAN Volume Controller resolves up to four worldwide node names (WWNNs) per subsystem and allows up to 512 LUNs per WWNN. The HDS Lightning assigns one WWNN per port; therefore, the SAN Volume Controller can be a limitation to both capacity (2048 LUNs) and bandwidth (4 ports).

You can use the following procedure for HDS Lightning subsystems with 8 ports if more capacity or bandwidth is required:

- Divide the set of ports into groups of between two and four.
- Assign a discreet set of LUNs to each group.

The SAN Volume Controller interprets each group as a separate subsystem.

- ▶ If an LU is mapped to the SAN Volume Controller port as LUNx, the LU must appear as LUNx to all the SAN Volume Controller ports in the cluster and must also appear as LUNx through all of the controller ports that it is mapped to.
- ▶ Command LUNs must not be mapped to the SAN Volume Controller.
- ▶ LUNs that are created using LUSE and Virtual LVI/LUN can be mapped to the SAN Volume Controller after they have been created. LUN Expansion (LUSE) and Virtual LVI/LUN operations cannot be run on a disk that is managed by the SAN Volume Controller.
- ▶ Only disks with open emulation can be mapped to the SAN Volume Controller. S/390® disks cannot be used with the SAN Volume Controller. Only Fibre Channel connections can be used to connect the SAN Volume Controller to the HDS Lightning.

Quorum disk support

The HDS Lightning 99xxV is not an approved host for quorum disks. Therefore, SVC configurations with only HDS Lightning are not possible.

Further information

Refer to the latest SVC Configuration Guide for further information regarding integrating the SVC with an existing storage controller. It is available at:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

5.1.3 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Prerequisite conditions

In this scenario we started with an HDS Lightning that was already configured to provide SAN storage over an existing SAN.

Since the installation of the SVC is outside the intended scope of this book, we have started our example at the point where all this preparatory work has been completed.

We have assumed the following initial conditions:

- ▶ The HDS Lightning is already configured to provide SAN storage, and LUN security is enabled.
- ▶ The SVC nodes have been installed and minimally configured as an SVC cluster with quorum disks on another storage controller.
- ▶ The switches have been installed, and appropriate zones have been created for the existing environment.

Preparing the SAN fabric

We need to change the SAN zoning to allow the SVC to access the existing storage controller and the host whose data is to be virtualized.

Refer to Figure 5-1 on page 82 for an overview of the required SAN zones for integrating the SVC in an existing environment. In our example environment these are the RED1 and BLUE zones.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendor’s SAN equipment.

Note: It is assumed that the BLACK zone (for SVC node to SVC node connection) already was defined when the SVC cluster was created.

Preparing the storage controllers

Next, we need to configure access from the SVC to the LUN on the existing storage controller.

Record the SVC fibre port numbers

We need to connect to the SVC and record the WWPN for each of the SVC nodes.

1. Connect to the SVC, then use the **svcinfn node** command to list the SVC nodes. Example 5-1 shows the output of this command on our system.

Example 5-1 List the storage controllers

IBM_2145:ldcluster:admin>svcinfn lsnode					
id	name	UPS_serial_number	WWNN	status	IO_group_id
0	node1	100062L102	500507680100055C	online	0
1	node2	100062L097	50050768010000D0	online	0

2. We need further information about each of the SVC nodes. Use the **svcinfn node nodename** command to list the properties of each of the SVC nodes. Example 5-2 shows the output of this command for one node on our system.

Example 5-2 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfn lsnode node1
id 0
name node1
UPS_serial_number 100062L102
WWNN 500507680100055C (WWNN for node1)
status online
IO_group_id 0
IO_group_name io_grp0
partner_node_id 1
partner_node_name node2
config_node no
UPS_unique_id 2040000182701002
port_id 500507680140055C (WWPN for port #4)
port_status active
port_speed 4Gb
port_id 500507680130055C (WWPN for port #3)
port_status active
port_speed 2Gb
port_id 500507680110055C (WWPN for port #1)
port_status active
port_speed 2Gb
port_id 500507680120055C (WWPN for port #2)
port_status active
port_speed 2Gb
hardware 8F4
```

Repeat these steps for each of the SVC nodes in your SVC cluster.

Configure the existing storage controller

Use the following steps to change the LUN masking on the storage controller to allow the SVC to access the LUNs that are to be virtualized.

1. Launch the Storage Navigator and log on with administrative privileges.
2. You need to change the Lightning administrative interface to “change mode” before you can make any changes to the controller’s configuration. To set the Storage Navigator in change mode:
 - a. Click the **pen** icon in the top right corner.
 - b. Click the **key & lock** icon in the left row of icons to access the **LU Path and Security** menu.
3. From the **LUN Manager** tab, expand the ports on which the host has access to the list of assigned LUNs.
4. Record the SCSI LUN IDs that the existing storage controller presents to the host for the LUNs that are to be virtualized. These ID numbers are needed later to allow the SVC to present the same SCSI IDs to the host. This will allow the host to see the virtualized volumes as though they were the original, directly connected, volumes.

Define the SVC host group on the HDS controller

1. Define a host group for the SVC on the HDS Lightning. Name the host group SVC, set the host mode to **Windows**, and click **OK**.
2. To associate the SVC’s WWPNs with the new host group SVC, right-click the **host group** and select **Add New WWN**. Add the WWPN and define a name, for example, Node1P1. Repeat this procedure for all SVC node ports.
3. Select the **Port** tab to verify the port settings for the ports we use for SVC access are set correctly. The port settings should be:

Host Speed:	Auto
Fibre Addr:	default assigned
Fabric:	ON
Connection:	P-to-P
4. Map the LUNs to the SVC with the following steps:
 - a. From the LUN Manager tab, highlight the **SVC** host group.
 - b. Select the LUNs in the **LDEV** pane.
 - c. Click **Add LU Path**.
 - d. Click **Apply** to commit the changes.

Refer to the HDS administration guide for detailed instructions on how to configure LUN masking on the Lightning storage controller.

Note: Though we recommend performing the integration process using the safest steps possible, it is safe to give the SVC access to the existing LUN (which is still being directly accessed by the host) as long as you do not also provide host access to the same LUN through the SVC at the same time.

Preparing the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized.

Use the **svcinfo lscontroller** command to list the storage controllers. Example 5-3 on page 87 shows the output of this command on our system.

Example 5-3 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfolsccontroller
```

id	controller_name	ctrl_s/n	vendor_id	product_id_low	product_id_high
0	controller0		HITACHI	DISK-SUB	SYSTEM
1	controller1		HITACHI	DISK-SUB	SYSTEM
2	sharkc11		IBM	2105800	

Use the **svcinfolsmdisk** command to list the MDiskS on the existing storage controller. Example 5-4 shows the output of this command on our system.

Example 5-4 List MDiskS on the existing storage controller

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue  
controller_name=controller*
```

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
0,mdisk0,online,unmanaged,,,203.5GB,0000000000000001,controller0,484...000

These commands have been edited for legibility. Note that the MDisk for the existing LUN is in *unmanaged* mode.

Create an MDG for the new storage controller's LUNs

Since we have already correctly zoned the SAN and configured LUN masking on the back-end storage controllers, the LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

Use the **svctask detectmdisk** command to rescan for newly available MDiskS.

Note: This command produces no user feedback.

If the expected MDiskS are still not visible to the SVC, double check the SAN zoning and storage controller LUN masking.

Use the **svctask mkmdiskgrp** command to create a new MDG. Example 5-5 shows the output of this command on our system.

Example 5-5 Add the new controller's LUNs to an MDG

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name MDG_HDS1 -ext 512
```

```
MDisk Group, id [1], successfully created
```

Note: You can define the MDiskS to be included in the MDG upon creation, but if you want to preserve any data on the MDiskS from the new storage controller, they must not be included in the MDG since this will destroy the data on the LUN. The only way to preserve data on a LUN is to create an image mode VDisk, and as part of this process the MDisk becomes managed.

Use the **svctask lsmdiskgrp** command to check that the MDG was created successfully. Example 5-6 on page 88 shows the output of this command on our system.

Example 5-6 List the MDisks that compose a given MDG

```
IBM_2145:ldcluster-65:admin>svcinfolsmdiskgrp -delim ,  
  
id,name,status,mdisk_count,vdisk_count,capacity,extent_size,free_capacity  
  
0,MDG_Sharkc,online,2,6,99.2GB,512,49.2GB  
1,MDG_HDS1,online,0,0,0GB,512,0GB
```

This command has been edited for legibility. Note that there are currently no MDisks defined in the MDisk group MDG_HDS1.

Define a host to the SVC

Before the host can access any VDisks, it needs to be defined to the SVC.

Use the **svcinfolshbaportcandidate** command to list any unassigned host WWPNs that are visible to the SVC. Example 5-7 shows the output of this command on our system.

Example 5-7 List unassigned host WWPNs that are visible to the SVC

```
IBM_2145:ldcluster:admin>svcinfolshbaportcandidate  
  
id  
10000000C93618F7  
210100E08B3A9AC6 (This is the first HBA on our host)  
210000E08B1A3FC6  
10000000C937A664  
210000E08B1A2EC6  
210000E08B1A9AC6 (This is the second HBA on our host)
```

Note: Check on the host (or the SAN switch) to determine which WWPN belongs to the host with which you are working.

Use the **svctask mkhost** command to define the host to the SVC. Example 5-8 shows the output of this command on our system.

Example 5-8 Define a host to the SVC

```
IBM_2145:ldcluster:admin>svctask mkhost -name general -hbawwn  
210100E08B3A9AC6:210000E08B1A9AC6 -type generic  
  
Host id [3] successfully created
```

Use the **svcinfolshost** command to check that the host was defined successfully. Example 5-9 shows the output of this command on our system.

Example 5-9 Check a host definition on the SVC

```
IBM_2145:ldcluster-65:admin>svcinfolshost general  
  
id 3  
name general  
port_count 2  
type generic  
mask 1111  
iogrp_count 4
```

```
WWPN 210100E08B3A9AC6
node_logged_in_count 8
state active
WWPN 210000E08B1A9AC6
node_logged_in_count 8
state active
```

The preparation is now complete and we are ready to proceed with the storage virtualization.

5.1.4 Begin storage virtualization

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SVC.

Configure the host system

The SAN Volume Controller uses the IBM Subsystem Device Driver (SDD) multi-path driver to manage connectivity between the host and the virtualized storage.

In our environment it was necessary to unload the existing HDS multi-path device driver (known as Hitachi Dynamic Link Manager or HDLM) before installing the IBM SDD driver.

Note: Keep the HDLM license in case you later want to stop using the SVC, and need to reinstall it.

For details on multipath driver coexistence with SDD consult the SVC support documentation:

<http://www-03.ibm.com/servers/storage/support/software/sanvc/>

Figure 5-2 shows the uninstall of the HDLM driver from our Windows host.

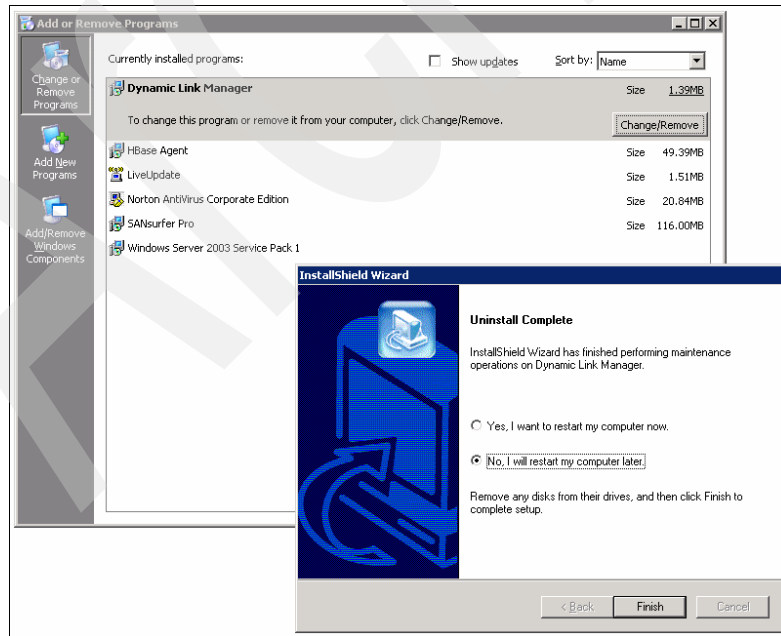


Figure 5-2 Uninstalling HDS HDLM for Windows

Be sure to follow the existing storage controller vendor's recommendations when removing the existing multipath device driver.

Shut down the host

Stop all applications and shut down the host.

The host should remain offline briefly while we reconfigure the SAN. While the host is offline, we remove direct access to the existing LUN and enable access to the virtualized LUN through the SVC.

It may be advisable to configure the server to prevent the automatic starting of any applications on the next boot. This is to prevent any problems with applications trying to access the SAN volumes before we can load the IBM SDD multi-path device driver.

Note: Although some parts of the storage virtualization process can be performed with the host system online, we chose to shut down the host to minimize risk during the change.

Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SVC.

Figure 5-1 on page 82 shows the zones. Remove the GREEN zone from the active zoning configuration and enable the BLUE zone so that the host being integrated can access the SVC.

Refer to Appendix B, "SAN zoning changes" on page 257 for examples of the commands used to define the required SAN zones on some common vendors' SAN equipment.

Configure the existing storage controller

Change the LUN masking on the existing storage controller to block the host from directly accessing the existing LUN.

To do this, right-click the host group on the **LUN Manager** tab and select **Clear Host Group**. When prompted, confirm the removal of the host to LUN access.

Refer to the HDS administration guide for detailed instructions on how to configure LUN masking on the Lightning storage controller.

Configure the SAN Volume Controller

Now that the other systems have been prepared, we can proceed with the virtualization of the existing storage volume.

Create an MDG for the image mode volumes

Before configuring the existing storage controller's LUNs as image mode VDisks, define a new MDG to logically house those volumes.

Use the **svctask mkmdiskgrp** command to create an empty MDG for the image mode volumes. Example 5-10 shows the output of this command on our system.

Example 5-10 Create an empty MDG for image mode volumes

```
IBM_2145:ldcluster-65:admin>svctask mkmdiskgrp -name imagegrp -ext 512
```

```
MDisk Group, id [9], successfully created
```

Although it is not strictly necessary to create a separate MDG for image mode volumes, we believe that it is good practice.

Important: To be able to migrate volumes between MDGs, the extent size of the source and target MDGs must be the same.

Create image mode VDisks from existing storage controller's LUNs

Now you can create the image mode VDisk using the existing unmanaged MDisk, which was previously connected directly to the host.

Use the **svctask mkvdisk** command to create the image mode VDisk. Example 5-11 shows the output of this command on our system.

Example 5-11 Define a new image mode VDisk

```
IBM_2145:ldcluster:admin>svctask mkvdisk -mdiskgrp imagegrp -iogrp 0 -vtype image  
-name general01 -mdisk mdisk0
```

```
Virtual Disk, id [1029], successfully created
```

Attention: If you were to add the existing LUNs to an MDG, then all data on those LUNs would be lost. Since we intend to virtualize previously existing LUNs with the data intact, we must create *image mode virtual disks*, which is done by defining a VDisk as being image mode and corresponding to an unmanaged MDisk.

Use the **svcinfn lsvdisk** command to check that the VDisk was created successfully. Example 5-12 shows the output of this command on our system.

Example 5-12 Display the properties of the new image mode VDisk

```
IBM_2145:ldcluster-65:admin>svcinfn lsvdisk general01
```

```
id 1029  
name general01  
IO_group_id 0  
IO_group_name io_grp0  
status online  
mdisk_grp_id 10  
mdisk_grp_name imagegrp  
capacity 203.5GB  
type image  
formatted no  
mdisk_id  
mdisk_name  
FC_id  
FC_name  
RC_id  
RC_name  
vdisk_UID 600507680191002AE0000000000000980  
throttling 0  
preferred_node_id 2  
fast_write_state empty  
cache readwrite  
udid
```

Map the image mode VDisk to the Host

Next, we need to define a relationship between the new image mode VDisk and our host system. This is the equivalent of LUN masking (or storage partitioning) on a normal storage controller.

Use the **svctask mkvdiskhostmap** command to map the new VDisk to the host. Example 5-13 shows the output of this command on our system.

Example 5-13 Map a VDisk to a Host

```
IBM_2145:ldcluster-65:admin>svctask mkvdiskhostmap -host general -scsi 1 general01
```

Virtual Disk to Host map, id [1], successfully created

Note: It is important to define the same SCSI ID number to the VDisk as the existing storage controller has presented to the host. This will allow the virtualized LUN to appear as the original LUN when later exposed to the host.

Configure the host system

Now that the SAN has been reconfigured, you can restart the host system.

When the host reboots, it will see multiple paths to each of the SAN volumes. Load the IBM SDD multipath device driver to provide correct access to the SAN volume, and reboot.

Note: If you have disabled the automatic start of any applications on the host you should re-enable them now.

After the host reboots (with SDD loaded), it should see a single path to the virtualized storage volume.

You can use the IBM SDD command line to check that the SAN volumes are available. Example 5-14 shows the output of this command on our system.

Example 5-14 Checking SAN devices with SDD

```
C:\Program Files\IBM\SDDDSM>datapath query device
```

Total Devices : 1

```
DEV#: 0  DEVICE NAME: Disk1 Part0  TYPE: 2145          POLICY: OPTIMIZED
SERIAL: 600507680191002AE0000000000000980
```

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
1	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	879	0
2	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
3	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	877	0

The data on the SAN disk is now available on the host for application and user access.

5.2 Integrating HDS TagmaStore USP with the SVC

When performing an integration with TagmaStore Universal Storage Platform (USP) and the SVC, the process is in practice identical to the one described previously with Lightning.

5.3 Integrating HDS Thunder with the SVC

In this section, we demonstrate the integration of the SVC in an existing environment with a Hitachi Data Systems Thunder storage controller. The existing storage volumes on the HDS Thunder storage controller will be virtualized using the SVC.

5.3.1 Scenario description

In this scenario, we start with an IBM xSeries® server, running VMware ESX Server 2.5, that is attached to the HDS Thunder 9200 storage controller through two SAN switches.

We then introduce the SVC to virtualize the back-end storage.

The systems described in this section include:

- ▶ HDS Thunder 9200 storage controller
- ▶ IBM b-type switches
- ▶ IBM xSeries server running VMware ESX Server v2.5.3
- ▶ IBM SAN Volume Controller v4.1

Figure 5-3 shows the relationships between the various systems used in this scenario.




Figure 5-3 SVC integration example with an existing HDS Thunder storage controller

5.3.2 Summary of migration process

At a high level, the steps to integrate the SVC into this an existing production environment are the following:

1. Preparation
 - a. General
 - i. Review SAN levels of software and firmware, supported hardware, and information from the configuration guide.
 - ii. Determine SAN Volume Controller worldwide port names and Host worldwide port names.
 - b. SAN
 - i. Create host zone with SVC (enable now or later).
 - ii. Create Thunder controller zone with SVC (enable now or later).
 - c. Thunder Controller
 - i. Identify the SCSI LUN IDs and Unique IDs (UID) for all those to be migrated under SAN Volume Controller.
2. Migration to SVC managed storage
 - a. Host
 - i. Shut down the virtual machines.
 - ii. Shut down the ESX Server.
 - b. SAN
 - i. Enable new zone for host-SVC and Thunder controller-SVC.
 - ii. Disable Thunder controller-host zone.
 - c. Thunder Controller
 - i. Unmap LUNs from host.
 - ii. Map LUNs to SVC.
 - d. SAN Volume Controller
 - i. Discover Thunder controller LUNs.
 - ii. Create Image Mode Vdisks.
 - iii. Create host with associated WWPNs.
 - iv. Map Vdisks to host.
 - e. Host
 - i. Start the ESX Server.
 - ii. Check the storage setup, compare to previous configuration.
 - iii. Start the virtual machines.

Note: Be sure to read 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization example.

5.3.3 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SAN Volume Controller is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment, it may require coordination of effort between various support teams and an understanding of several technical domains.

Confirm support for your environment

This section provides information about configuring the SVC in an existing environment with a Hitachi Data Systems Thunder storage controller.

Prior to integrating the SVC, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

Both the 9200 and 9500 series of Thunder systems are supported by the SVC. For details on specific microcode levels and such, refer to the SVC v4.1.0 Supported Hardware Web site.

Quorum disk support

A quorum disk is used to resolve tie-break situations when the nodes disagree on the current state of the cluster. The SVC cluster requires a quorum disk, but not all storage arrays are supported for quorum device. When you are planning the migration, make sure you have at least one disk system that is supported.

All HDS Thunder models are supported as sources of quorum disk, so any HDS Thunder model can be used as the only back-end controller of the SVC Cluster, but bear in mind that there are limitations for quorum disks. The most important limitation in our situation is that LUNs already in use by hosts that you are migrating over to the SVC cannot be used for quorum. Therefore, if you don't already have a quorum device, you must have some unused capacity on the Thunder that you can map to the SVC.

Further information

SVC version 4.1 Configuration Guide:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

SVC v4.1.0 Supported Hardware Web site:

<http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002864>

Preparing the SAN fabric

We need to prepare two additional zones: one to allow SVC to be visible to the host whose data is to be virtualized (the BLUE zone in Figure 5-3 on page 93) and the second to enable the SVC to be visible to the HDS Thunder (the RED1 zone in Figure 5-3). Depending upon your environment, you may choose to simply create these zones at this time and only enable them later when the host applications have been stopped.

Refer to Appendix B, "SAN zoning changes" on page 257 for examples of the commands used to define the required SAN zones on some common vendors' SAN equipment.

Note: It is assumed that the BLACK zone (for SAN Volume Controller node to SAN Volume Controller node connection) was already defined when the SVC cluster was created.

Record the SVC fibre adapter WWPNs

We will need the World Wide Port Names (WWPN) for each of the SVC nodes. This will enable us to configure the HDS Thunder to make LUNs visible to the SVC. There are several methods for collecting this information. You can use the SAN switch interface to identify the WWPNs, provided you know which ports the SAN Volume Controller nodes are connected to. Alternatively, you can use the SAN Volume controller GUI or CLI. The following example shows how to use the SVC CLI.

Establish a secure shell (ssh) connection to the SVC, then use the **svcinfolsnode** command to list the SVC nodes. Example 5-15 shows the output of this command on our system.

Example 5-15 List the SVC nodes

```
IBM_2145:GREEN_SVC:admin>svcinfolsnode
id name      UPS_serial_number  WWNN              status  IO_group_id
1  node1      10005BL068         50050768010023EC  online  0
2  node2      10005BL066         50050768010024E4  online  0
IBM_2145:GREEN_SVC:admin>
```

This shows the nodes in the SVC cluster (in this case node1 and node2) and the WWNN of each.

You need some further information about each of the SVC nodes. Use the **svcinfolsnodenodename** command to list the properties of each of the SVC nodes. Example 5-16 shows the output of this command on our system.

Example 5-16 List the SVC node details

```
IBM_2145:GREEN_SVC:admin>svcinfolsnodenode1
id 1
name node1
UPS_serial_number 10005BL068
WWNN 50050768010023EC          (WWNN for node1)
status online
IO_group_id 0
IO_group_name io_grp0
partner_node_id 2
partner_node_name node2
config_node yes
UPS_unique_id 2040000152700188
port_id 50050768014023EC      (WWPN for port 1)
port_status active
port_speed 2Gb
port_id 50050768013023EC      (WWPN for port 2)
port_status active
port_speed 2Gb
port_id 50050768011023EC      (WWPN for port 3)
port_status active
port_speed 2Gb
port_id 50050768012023E      (WWPN for port 4)
port_status active
port_speed 2Gb
hardware 8F2
IBM_2145:GREEN_SVC:admin>
```

Repeat this process for each SVC node in the cluster.

Record the host WWPNS

There are several methods of identifying the WWPNS of the host involved in the migration. The simplest method is to capture the information through the SAN switch. Alternatively, you can access the information about the host through the host bus adapter BIOS.

On VMware ESX Server you can also log in by secure shell (ssh) and enter the command **wwpn.pl**. Example 5-17 shows how the command is used.

Example 5-17 Output of the wwpn.pl command

```
[root@bryher root]# wwpn.pl
vmhba2: 10000000c94a0382 (Emulex) 7:3:0
vmhba3: 10000000c94a0383 (Emulex) 7:3:1
```

Record the HDS Thunder WWPNs

To display the WWPNs on HDS Thunder, use the supplied graphical management interface, Disk Array Management Program (DAMP). In DAMP select the **Configuration menu** → **Fibre Channel** tab, and the details for each port are displayed. Record these details for later use.

Prepare the host system

We need to check which LUNs are presented to the host. Even if there is documentation available on the environment, it is good practice to verify this before the migration.

To register the disk mappings on the VMware ESX, log on to the VMware ESX Server Web interface and select **Storage Management** settings, or telnet to the ESX Server and use the command line interface.

When using the Web interface, start by selecting the **Settings** tab from the main screen, then select the **Storage Management** link, then the **Failover Paths** tab.

From this screen you can see which LUNs the ESX Server has mapped. In our scenario the ESX Server has three disks mapped from the storage array, the ESX server has two path for each LUN, their Logical Unit Numbers (LUN or SCSI ID) are 0, 1, and 2, and the WWPNs of the storage array (SAN Target column) are, respectively, 50:06:0e:80:00:01:11:02 and 50:06:0e:80:00:01:11:12.

Preparing the HDS Thunder

Now that we have the host side details (LUNs, SAN Target WWPNs), we can check that the storage array side matches and that there is no misconfiguration. It is important to verify before the migration that everything really works and is configured properly.

In DAMP, select the **Target ID** tab on the **Configuration** menu. There are two different modes the storage array can be in: **S-TID** (Single Target ID), where every LU is presented to every host through every port; or **M-TID** (Multi Target ID), where you have to separately map the LUs to ports and assign them LUNs (comparable to enable/disable LUN masking). In our example, our Thunder is using M-TID since storage arrays seldom are dedicated to just one host.

Note: In Hitachi terminology, LU is sometimes called LUN and LUN is instead called H-LUN, Host LUN or LUN Number.

In the **Current Configuration** pane, verify the LUN numbers. In the H-LUN column (on the far right side) you should see the same LUN numbers you saw on the ESX Server, also the storage controller ports (for example 0A, 1A) in the Ports column. Ports should correspond to the SAN Target WWPNs that we identified on the ESX Server WWPNs. The left column (LUN) in the **Current Configuration** pane is the internal ID for the LU; record this value because you will need it a bit later when you check the LUN Security configuration.

Select the **Fibre Channel** tab from the DAMP Configuration menu; the LUN masking is done from this screen. Each of the storage array host ports has its own settings. First see if the **Security Information Enable** check box is checked for the port. If it is, only the specified

WWPNs have access to the specifies LUNs. If it is not checked, everyone connected to that port has access to anything that storage array is presenting on that port.

The buttons on the right (Add, Change, File, Delete) are used to add new host definitions for LUN Security. The button on the left is used for defining which host can access which LUN.

5.3.4 Begin migration

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SVC.

Configure the host system

On the host, in this example a VMWare ESX server, unmount the volumes on the LUNs that you are about to migrate, or alternatively, shut down the ESX server while the LUNs are integrated to the SVC (be sure to stop the guest OSs residing on these volumes prior to unmount). Since the ESX server handles multipath LUN access natively, we do not install SDD or any other multipath driver on the ESX server.

Note: Although some parts of the storage virtualization process can be performed with the host system online, we chose to shut down the host to minimize risk during the change.

Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SVC.

Figure 5-3 on page 93 shows the zones. Remove the GREEN zone from the active zoning configuration and enable the BLUE zone so that the host being integrated can access the SVC.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Configure the existing storage controller

Change the LUN masking on the existing storage controller to block the host from directly accessing the existing LUN.

This is done from the **Fibre Channel** tab from the DAMP configuration menu.

Refer to the HDS administration guide for detailed instructions on how to configure LUN masking on the storage controller.

Configure the SAN Volume Controller

Now that the other systems have been prepared, we can proceed with the virtualization of the existing storage volume.

Create an MDG for the image mode volumes

Before configuring the existing storage controller’s LUNs as image mode VDisks, define a new MDG to logically house those volumes.

Use the `svctask mkmdiskgrp` command to create an empty MDG for the image mode volumes. Figure 5-18 on page 99 shows the output of this command on our system.

Example 5-18 Create an empty MDG for image mode volumes

```
IBM_2145:ldcluster-65:admin>svctask mkmdiskgrp -name imagegrp -ext 512
```

```
MDisk Group, id [10], successfully created
```

Although it is not strictly necessary to create a separate MDG for image mode volumes, we believe that it is good practice.

Important: To be able to migrate volumes between MDGs, the extent size of the source and target MDGs must be the same.

Create image mode VDisks from existing storage controller's LUNs

Now we can create the image mode VDisk using the existing unmanaged MDisk, which was previously connected directly to the host.

Use the **svctask mkvdisk** command to create the image mode VDisk. Example 5-19 shows the output of this command on our system.

Example 5-19 Define a new image mode VDisk

```
IBM_2145:ldcluster:admin>svctask mkvdisk -mdiskgrp imagegrp -iogrp 0 -vtype image  
-name Bryher01 -mdisk mdisk01
```

```
Virtual Disk, id [1030], successfully created
```

Attention: If you were to add the existing volumes to an MDG as normal *striped* MDisks, then all data on those volumes would be lost. Since we intend to virtualize previously existing volumes with data intact, we must create *image mode virtual disks* instead.

Use the **svcinfa lsvdisk** command to check that the VDisk was created successfully. Example 5-20 shows the output of this command on our system.

Example 5-20 Display the properties of the new image mode VDisk

```
IBM_2145:ldcluster-65:admin>svcinfa lsvdisk Bryher01
```

```
id 1029  
name Bryher01  
IO_group_id 0  
IO_group_name io_grp0  
status online  
mdisk_grp_id 10  
mdisk_grp_name imagegrp  
capacity 203.5GB  
type image  
formatted no  
mdisk_id  
mdisk_name  
FC_id  
FC_name  
RC_id  
RC_name  
vdisk_UID 600507680191002AE000000000000981  
throttling 0
```

```
preferred_node_id 2
fast_write_state empty
cache readwrite
udid
```

Map the image mode VDisk to the Host

Define a relationship between the new image mode VDisk and your host system. This is the equivalent of LUN masking (or storage partitioning) on a normal storage controller.

Use the **svctask mkvdiskhostmap** command to map the new VDisk to the host. Example 5-21 shows the output of this command on our system.

Example 5-21 Map a VDisk to a Host

```
IBM_2145:ldcluster-65:admin>svctask mkvdiskhostmap -host Bryher -scsi 0 Bryher01
```

```
Virtual Disk to Host map, id [0], successfully created
```

Note: It is important to define the same SCSI ID number to the VDisk as the existing storage controller has presented to the host. This will allow the virtualized LUN to appear as the original LUN when later exposed to the host.

Configure the host system

Now that the SAN has been reconfigured you can restart the host system.

The data on the SAN disk is now available on the host for application and user access.



Integrating SVC with EMC

In this chapter we describe how to integrate the SVC in an existing environment with various EMC storage controllers.

6.1 Integrating the EMC Symmetrix with the SVC

In this section we demonstrate the integration of the SVC in an existing environment with an EMC Symmetrix storage controller. The existing storage volumes on the EMC Symmetrix (8730) storage controller will be virtualized using the SVC.

Note: The same steps can be used for any EMC Symmetrix storage controller because they all share a common management interface.

6.1.1 Scenario description

In this scenario, we start with an EMC Symmetrix storage controller and SAN-attached Windows and AIX hosts. We then introduce the SVC in order to virtualize the back-end storage.

The systems described in this section include:

- ▶ EMC Symmetrix storage controller
- ▶ IBM b-type family switches

Figure 6-1 shows the relationships between the systems used in this scenario as well as the SAN zones that will be used.




Figure 6-1 SAN zoning for SVC

6.1.2 Summary of the SVC integration process

At a high level, the steps to integrate the SVC into an existing production environment are the following:

1. Preparation

a. General

- i. Review SAN levels of software, firmware, and supported hardware using information from the configuration guide.
- ii. Determine SVC worldwide port names and Host worldwide port names.

b. SAN

- i. Create a host with SVC zone (enable now or later).
- ii. Create an EMC Symmetrix with SVC zone (enable now or later).

Note: In this example the SVC cluster is already set up and installed. For details on how to set up and install an SVC cluster consult *IBM System Storage SAN Volume Controller*, SG24-6423.

<http://www.redbooks.ibm.com/abstracts/sg246423.html?open>

c. EMC Symmetrix

- i. Identify the LUN unique IDs and SCSI IDs for all the LUNs to be integrated under SVC.

2. Integration

a. Host

- i. Host-specific actions.
- ii. Stop applications.
- iii. While the only supported multipath driver at this time is SDD, the existing multipath driver (PowerPath) should be uninstalled on the host for the integration to be successful.

Note: Keep the PowerPath license in case you later want to stop using the SVC, and need to reinstall it. For details on multipath driver coexistence with SDD consult the SVC support documentation:

<http://www-03.ibm.com/servers/storage/support/software/sanvc/>

- iv. Install IBM Subsystem Device Driver (SDD).
- v. Shut down machine or unmount file systems and vary off volume groups.

b. SAN

- i. Enable new zone for host-to-SVC and the OEM controller-to-SVC.
- ii. Remove the OEM controller-to-host zone from the active zoning configuration.

c. EMC Symmetrix

- i. Unmap LUNs from host (LUN masking).
- ii. Remap LUNs to SVC (LUN masking).

- d. SVC
 - i. Discover EMC Symmetrix LUNs.
 - ii. Create Image Mode Virtual Disks (VDisk).
 - iii. Create host with associated WWPNS.
 - iv. Map VDisks to host.
- e. Host
 - i. Host-specific actions.
 - ii. Reboot or scan for new devices.
 - iii. Vary on volume group.
 - iv. Remount file systems.

Note: We recommend that you review 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization example.

6.1.3 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Confirm support for your environment

This section provides information about configuring the SVC in an existing environment with an EMC Symmetrix storage controller.

Prior to integrating the SVC, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SVC supports most models of the EMC Symmetrix.

Table 6-1 lists the models of the EMC Symmetrix, including DMX, that the SVC supports.

Table 6-1 Supported models of EMC Symmetrix

Series	Models supported
Symmetrix	8130, 8230, 8430, 8530, 8730, 8830
DMX	800, 1000, 2000, 3000
DMX-2	800, 1000-M2, 1000-P2, 2000-M2, 2000-M2-3, 2000-P2, 2000-P2-3, 3000-M2-3
DMX-3	DMX-3

The list in Table 6-1 is current at the time of writing. See the following Web site for the latest supported EMC Symmetrix models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Sharing an EMC Symmetrix between a host and the SVC

You can directly connect hosts to an EMC Symmetrix at the same time as the SVC, provided the following conditions are met:

- ▶ You must zone the fabric such that other hosts cannot access the target ports that the SVC uses.
- ▶ You must configure the EMC Symmetrix such that other hosts cannot access the LUs that the SVC manages.

Note: A single host must not be connected to both the SVC and an EMC Symmetrix at the same time because the multipathing drivers IBM Subsystem Device Driver (SDD) and EMC PowerPath cannot coexist on the host.

Switch zoning considerations for the EMC Symmetrix and Symmetrix DMX

The following considerations apply to switch zoning for the SVC and the EMC Symmetrix subsystems:

- ▶ The SVC switch zone must include at least one target port on two or more Fibre Channel adapters to avoid a single point of failure.
- ▶ If Volume Logix is not enabled, this implies potential security risk if zoning is not performed correctly. Therefore, if Volume Logix is not enabled, ensure only SVC initiator ports that are LUN-masked on the EMC Symmetrix controller are present in the fabric zone.

Connecting the EMC Symmetrix to the SVC through the SAN

The EMC Symmetrix connects to the SAN through a Fibre Channel director. Directors are installed in pairs and each consists of two boards, one of which is a Fibre Channel adapter. The Fibre Channel adapter provides 2 - 12 target ports. The EMC Symmetrix assigns a worldwide node name (WWNN) per target port and the SVC can resolve up to four WWNs per subsystem. If you want to connect more than four target ports to the SVC, perform the following steps:

1. Divide the set of target ports into groups of two to four.
2. Define a discrete set of logical units (LUs) for each group.
3. Map the LUs to each target port in their group.

The SVC views each group of target ports as a separate subsystem. Ensure that no LUs are members of more than one group.

Note: You must configure the EMC Symmetrix to present logical units (LUs) to all SVC initiator ports that are in the fabric zone.

Quorum disk support

The SVC uses a small portion of three separate Managed Disks (MDisks) for the quorum disks. MDisks presented from the EMC Symmetrix are supported as quorum disks.

Further information

We strongly recommend that you refer to the SVC version 4.1 Configuration Guide for further information regarding integrating the SVC with an existing storage controller:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

Preparing the SAN fabric

Change the SAN zoning to allow the SVC to access the existing storage controller and the host whose data is to be virtualized.

Refer to Figure 6-1 on page 102 for an indication of the required SAN zones for integrating the SVC in an existing environment. In our example environment, these are:

- ▶ BLACK zone (SVC node-to-node connection). We assume that this has already been activated when the SVC cluster was created.
- ▶ RED1 zone (SVC- to-EMC Symmetrix)
- ▶ BLUE (SVC-to-host) zones

The GREEN zone is the original zone connecting the host to the EMC Symmetrix.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors SAN equipment.

Record the SAN Volume Controller Worldwide Port Names

Record the SVC Worldwide Port Names (WWPNs) for each of the SVC nodes. This will enable you to configure the EMC Symmetrix to make LUNs visible to the SVC (LUN masking). There are several methods of collecting this information. You can use the SAN switch interface to identify the WWPNs, provided you know which ports the SVC nodes are connected to. Alternatively, you can use the SVC GUI or CLI. Example 6-1 shows how to list the WWPN of the SVC nodes using the SVC CLI.

Record the SVC Fibre Channel port numbers

Connect to the SVC and record the WWPN for each of the SVC nodes.

Connect to the SVC; we did this using the CLI. Use the **svcinfo lsnode** command to list the SVC nodes. Example 6-1 shows the output of this command on our system.

Example 6-1 List the SVC nodes

```
IBM_2145:GREEN_SVC:admin>svcinfo lsnode
id name      UPS_serial_number  WWNN                status  IO_group_id
1  node1      10005BL068         50050768010023EC   online  0
2  node2      10005BL066         50050768010024E4   online  0
IBM_2145:GREEN_SVC:admin>
```

This shows the nodes in the SVC cluster (in this case node1 and node2) and the WWNN of each.

You now need to obtain further information about each of the SVC nodes.

Use the **svcinfo lsnode nodename** command to list the properties of each of the SVC nodes. Example 6-2 shows the output of this command on our system.

Example 6-2 List the SVC node details

```
IBM_2145:GREEN_SVC:admin>svcinfo lsnode node1
id 1
name node1
UPS_serial_number 10005BL068
WWNN 50050768010023EC           (WWNN for node1)
status online
IO_group_id 0
```

```

IO_group_name io_grp0
partner_node_id 2
partner_node_name node2
config_node yes
UPS_unique_id 2040000152700188
port_id 50050768014023EC          (WWPN for port 1)
port_status active
port_speed 2Gb
port_id 50050768013023EC          (WWPN for port 2)
port_status active
port_speed 2Gb
port_id 50050768011023EC          (WWPN for port 3)
port_status active
port_speed 2Gb
port_id 50050768012023E          (WWPN for port 4)
port_status active
port_speed 2Gb
hardware 8F2
IBM_2145:GREEN_SVC:admin>

```

Repeat this process for each SVC node in the cluster.

Record the host Worldwide Port Names (WWPN)

There are several methods of identifying the WWPNs of the host involved in the integration. The simplest method is to capture the information through the SAN switch. Alternatively, you can access the information through the host from the host bus adapter BIOS. On some platforms you can capture the information in a configuration file.

Configure the EMC Symmetrix

Change the LUN masking on the EMC Symmetrix controller to allow the SVC to access the LUNs that are to be virtualized.

We can access the administrative interface on the EMC Symmetrix through EMC ControlCenter or through the CLI. In this example we use EMC ControlCenter.

On the Modify Masking Configuration window we identify the LUNs that are masked to our Windows host *Krusty*. In our example this is Device 016, to which our Windows host is granted access through Symmetrix ports FA-5B and FA-14B (also called Storage Adapters).

We record the SCSI LUN IDs that the Symmetrix storage controller presents to the host for the LUNs that are to be virtualized. You will need these ID numbers later to allow the SVC to present the same SCSI IDs to the host.

6.1.4 Begin integration

This section describes the steps to be performed on the hosts in our scenario.

Host

The SVC is provided with a multipath driver called the IBM Subsystem Device Driver (SDD). The driver manages the resilience and connectivity between the host and the virtualized storage. Multipath drivers from other vendors are also tested and supported. For the most up to date support matrix, refer to the SAN Volume Controller supported hardware list at:

<http://www.ibm.com/storage/support/2145>

In the examples shown in this section the existing multipath driver (PowerPath), which is specific to the EMC storage controller, is uninstalled, and then SDD is installed on the host. The procedure is similar for other device drivers as well, but we recommend that you read the relevant documentation available for those drivers.

Microsoft Windows

Uninstalling PowerPath and installing SDD is a disruptive process that requires two reboots.

Figure 6-2 shows the current configuration of the Windows host, with a 25.29GB E: drive that is directly allocated from the EMC Symmetrix.

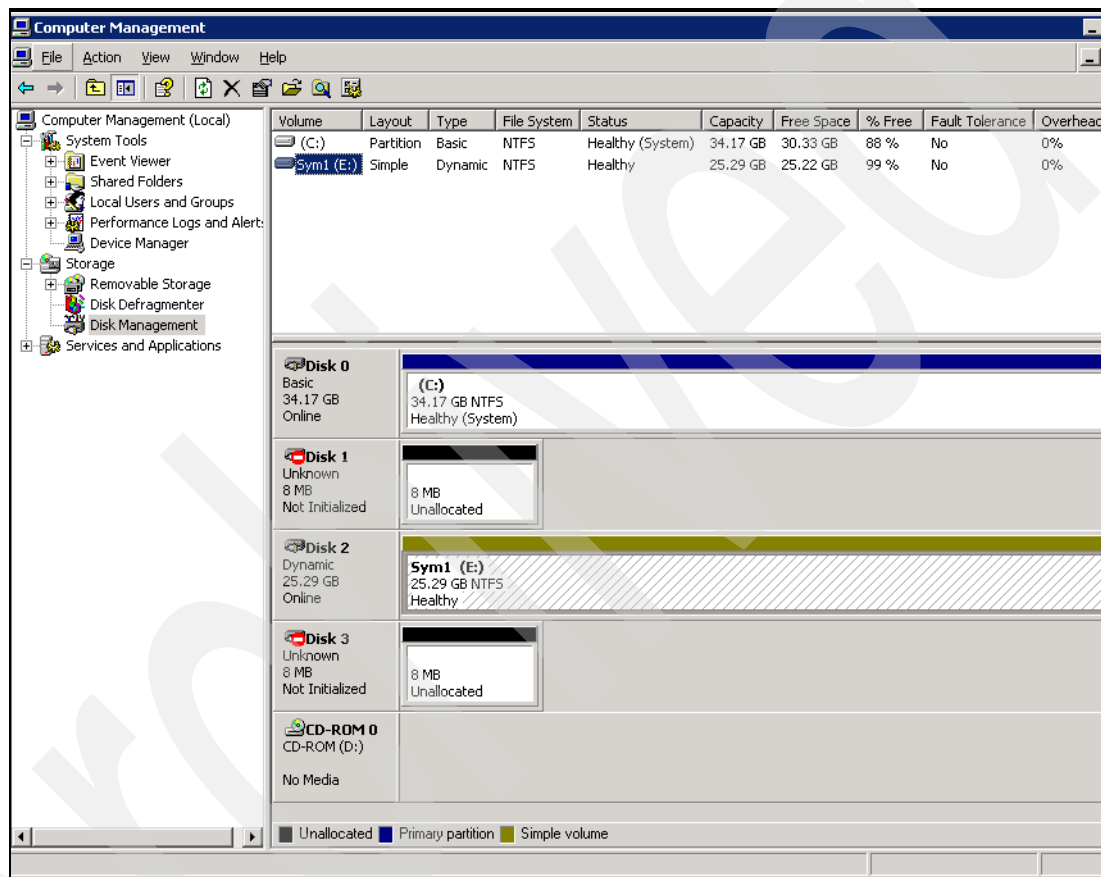


Figure 6-2 Windows host with EMC Symmetrix volume

Perform the following steps to uninstall Powerpath and install SDD:

1. Stop all applications from using the host.
2. Remove the drive letter from the volume to be integrated (in this case the E: drive) as shown in Figure 6-3.




Figure 6-3 Windows host with drive letter removed

3. Uninstall the EMC PowerPath multipathing software from the host before installing the IBM SDD multipathing software. Although it is possible to perform both of these tasks consecutively, it is recommended that you perform them in the following sequence:

Note: Be sure to consult the existing storage controller vendors uninstall recommendations when removing the existing multi-path device driver.

- a. Uninstall the EMC PowerPath using **Add or Remove Programs** (from the Control Panel), as shown in Figure 6-4.




Figure 6-4 Remove PowerPath from the Windows host

- b. Reboot the host to complete the removal of EMC PowerPath.

- c. Install the IBM multipath driver, SDD. The latest version of the filesets can be downloaded from the Web. Refer to the SVC recommended software levels document, which will direct you to the relevant download site:

<http://www.ibm.com/storage/support/2145>

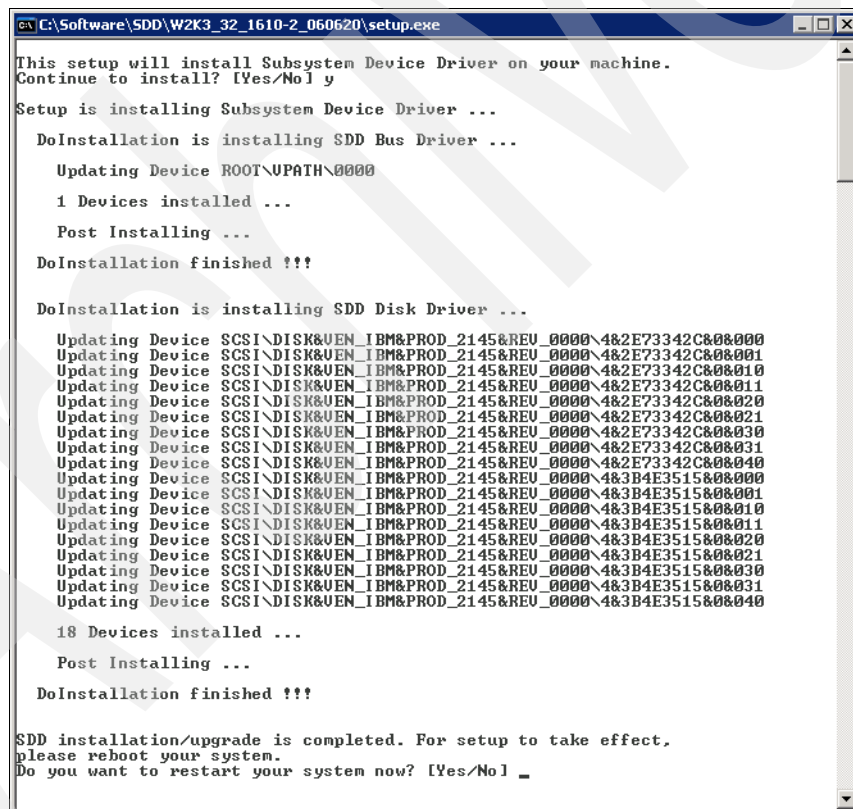
It is advisable to configure the server to prevent automatic startup of any applications dependant on access to volumes on the SAN during the next boot. This is to prevent any problems with applications trying to access the SAN volumes, which are currently not available.

The host should remain briefly offline while you reconfigure the SAN. Remove direct access to the existing LUN and enable access to the virtualized LUN through the SAN Volume Controller.

Note: Although some parts of the storage virtualization process may be conducted with the host system online, we choose to shutdown the host so as to minimize risk during the change.

Figure 6-5 shows the output when SDD installs.

- d. Reboot the host to complete the installation.



```
C:\Software\SDD\W2K3_32_1610-2_060620\setup.exe

This setup will install Subsystem Device Driver on your machine.
Continue to install? [Yes/No] y

Setup is installing Subsystem Device Driver ...

DoInstallation is installing SDD Bus Driver ...

  Updating Device ROOT\UPATH\0000
  1 Devices installed ...
  Post Installing ...
  DoInstallation finished !!!

DoInstallation is installing SDD Disk Driver ...

  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808000
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808001
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808010
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808011
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808020
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808021
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808030
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808031
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\482E73342C808040
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808000
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808001
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808010
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808020
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808021
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808030
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808031
  Updating Device SCSI\DISK&VEN_IBM&PROD_2145&REV_0000\483B4E3515808040

  18 Devices installed ...
  Post Installing ...
  DoInstallation finished !!!

SDD installation/upgrade is completed. For setup to take effect,
please reboot your system.
Do you want to restart your system now? [Yes/No] _
```

Figure 6-5 SDD installation output

AIX

On this platform, all the changes to the host system can be made without a full shutdown. In our example, the AIX system had the EMC PowerPath multipathing driver installed.

1. Use the **lsdev -Cc disk** command to determine what disks are being presented to the host system. Example 6-3 shows the output of this command on our system.

Example 6-3 List disks presented by EMC Symmetrix to AIX host

```

root@p615-tic-6[/usr/lpp] # lsdev -Cc disk
hdisk0      Available 1S-08-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1      Available 1S-08-00-5,0 16 Bit LVD SCSI Disk Drive
hdisk2      Available 1H-08-01      EMC Symmetrix FCP Raid1
hdisk3      Available 1H-08-01      EMC Symmetrix FCP Raid1
hdisk4      Available 1V-08-01      EMC Symmetrix FCP Raid1
hdisk5      Available 1V-08-01      EMC Symmetrix FCP Raid1
hdiskpower0 Available 1V-08-01      PowerPath Device
hdiskpower1 Available 1V-08-01      PowerPath Device
root@p615-tic-6[/usr/lpp] #

```

2. Use the **lspv** command to determine what volume group they are part of. Example 6-4 shows the output of this command on our system.

Example 6-4 List physical volume to volume group mapping

```

root@p615-tic-6[/] # lspv
hdisk0      0051a28abb619770 rootvg      active
hdisk1      0051a28a1f16b48c  None
hdisk2      none                None
hdisk3      none                None
hdisk4      none                None
hdisk5      none                None
hdiskpower0 0051a28a8b6742bf SymmVG      active
hdiskpower1 0051a28a8b67446c SymmVG      active

```

On the system we are using you will see that the EMC Symmetrix is presenting two disks, **hdiskpower0** and **hdiskpower1**, which are part of the **SymmVG** volume group.

You need to stop the applications that access this volume group. In your environment there may be more than one volume group involved, in which case you should execute the following set of commands with that in mind:

1. Use the **lsvg -o |lsvg -i -l** command to determine what mount points the volume group has. Example 6-5 shows the logical volumes associated with the host system volume groups and the relevant mount points for our system.

Example 6-5 Mapping between volume group and logical volumes

```

root@p615-tic-6[/] # lsvg -o |lsvg -i -l
SymmVG:
LV NAME      TYPE      LPs  PPp  PVs  LV STATE MOUNT POINT
fs1v00       jfs2      500  500  2    open/syncd /emc
log1v00      jfs2log   1    1    1    open/syncd N/A
rootvg:
LV NAME      TYPE      LPs  PPp  PVs  LV STATE MOUNT POINT
hd5          boot      1    1    1    closed/syncd N/A
hd6          paging    8    8    1    open/syncd  N/A
hd8          jfslog    1    1    1    open/syncd  N/A
hd4          jfs       1    1    1    open/syncd  /
hd2          jfs       18   18   1    open/syncd  /usr
hd9var       jfs       1    1    1    open/syncd  /var
hd3          jfs       33   33   1    open/syncd  /tmp

```

hd1	jfs	1	1	1	open/syncd	/home
hd10opt	jfs	4	4	1	open/syncd	/opt
lv00	jfs	4	4	1	open/syncd	/usr/vice/cache

Note that the SymmVG has two logical volumes associated. fslv00 is mounted to /emc.

2. Unmount the file system and vary off the volume group in readiness for integrating it under the SVC environment.

To do this, run **/usr/sbin/umount filesystem** followed by **varyoffvg Volume Group**.

Example 6-6 shows this set of commands being run on our system.

Example 6-6 Unmount of file system and vary off of volume group

```
root@p615-tic-6[/] # /usr/sbin/umount /emc
root@p615-tic-6[/] # /usr/sbin/varyoffvg SymmVG
```

Perform the following steps to uninstall the EMC PowerPath device driver and the existing LUN definitions:

1. Remove the EMC PowerPath driver by running the command **installp -u EMCpower**. Information about the progress of the removal is displayed, ending with the details shown in Example 6-7.

Example 6-7 PowerPath deinstall

Summaries:				
Installation Summary				
Name	Level	Part	Event	Result
EMCpower.consistency_grp	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.hr	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.mobility	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_aa	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_ap	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_clariio	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.ppvm	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.vcsagent	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.base	4.4.0.0	USR	DEINSTALL	SUCCESS

2. Use the following command to remove the disk definitions: **rmdev -d1 hdiskn**, where **n** represents the corresponding number of each hdisk presented by the EMC Symmetrix. Example 6-8 shows the command output.

Example 6-8 Removing the EMC Symmetrix related hdisks

```
root@p615-tic-6[/] # rmdev -d1 hdisk2
hdisk2 deleted
root@p615-tic-6[/] # rmdev -d1 hdisk3
hdisk3 deleted
root@p615-tic-6[/] # rmdev -d1 hdisk4
hdisk4 deleted
root@p615-tic-6[/] # rmdev -d1 hdisk5
hdisk5 deleted
```

You now need to install the IBM SDD and host attachment scripts onto the host system. The latest version of these filesets can be downloaded from the Web. Refer to the SAN Volume Controller recommended software levels document:

<http://www.ibm.com/storage/support/2145>

This will direct you to the relevant download site. Perform the following steps to install the IBM SDD:

1. If you are uncertain which level of AIX you have installed issue the command **oslevel**, which will return this information. This will determine which filesets you should download. The system in our example is running AIX v5.3, so the two filesets we downloaded are:

```
devices.fcp.disk.ibm.rte.tar
devices.sdd.53.rte.tar
```

Place the filesets in the same directory. We recommend creating a new directory called IBM.

2. Untar the files by running the commands as shown in Example 6-9.

Example 6-9

```
tar -xvf devices.fcp.disk.ibm.rte.tar
tar -xvf devices.sdd.53.rte.tar
```

This will generate two unpacked filesets. Example 6-10 shows the output from listing the directory in which they were untarred.

Example 6-10 List of files following untar command

```
root@p615-tic-6[/IBM] # ls -l
total 6168
-rwxrwx---1 155350 449628      31744 09 May 23:59 devices.fcp.disk.ibm.rte
-rw-r-----1 root   system    40960 21 Jul 16:15 devices.fcp.disk.ibm.rte.tar
-rw-r-----1 root   system  1536000 09 May 20:02 devices.sdd.53.rte
-rw-r-----1 root   system  1546240 21 Jul 16:15 devices.sdd.53.rte.tar
```

3. Install these filesets by using smit/smitty or through the command line. Example 6-11 shows how we set the install source directory as IBM and then installed the filesets in that directory.

Example 6-11 Commands to install through the command line

```
root@p615-tic-6[/] # cd /IBM
root@p615-tic-6[/IBM] # inutoc $PWD
root@p615-tic-6[/IBM] # installp -acd . -einstallp.log ALL
```

4. Check the output to ensure that the filesets were installed successfully. The output will show the cause of any failures.

Now that you have uninstalled the EMC drivers and installed the SDD, you can reconfigure the SAN fabric and map the existing EMC Symmetrix LUNs to the SVC.

Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SVC. (Figure 6-1 on page 102 shows the zones. You should remove the GREEN zone from the active zoning configuration).

You should now enable the BLUE zone so that the host being integrated can access the SVC.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Configure the EMC Symmetrix

Identify the EMC Symmetrix ports that are to be used by the SVC and create the appropriate SAN zones to allow access.

Use the **svcinfolsccontroller** command to confirm that the SVC can see the EMC Symmetrix once the SAN zones have been activated.

There are two controllers, *controller1* and *controller2*, with a vendor/product ID of EMC SYMMETRI X. Each port on an EMC Symmetrix uses a different WWNN, therefore the two ports appear as two separate controllers. Since the SVC can support up to 64 back-end storage controllers, we recommend that you change the generic name controllers to a friendly name in order to easily identify them in the future. We used the commands shown in Example 6-12 to rename the controllers (in this case to Sym8730-1 and Sym8730-2).

Example 6-12

```
svctask chcontroller -name Sym8730-1 controller1
svctask chcontroller -name Sym8730-2 controller2
```

Example 6-13 shows the output of these commands on our system.

Example 6-13 List of SVC back-end controllers

```
IBM_2145:GREEN_SVC:admin>svcinfolsccontroller
id controller_name ctrl_s/n vendor_id product_id_low product_id_high
0 DS8000 IBM 2107900
1 controller1 EMC SYMMETRI X
2 controller2 EMC SYMMETRI X
IBM_2145:GREEN_SVC:admin>svctask chcontroller -name Sym8730-1 controller1
IBM_2145:GREEN_SVC:admin>svctask chcontroller -name Sym8730-2 controller2
IBM_2145:GREEN_SVC:admin>svcinfolsccontroller
id controller_name ctrl_s/n vendor_id product_id_low product_id_high
0 DS8000 IBM 2107900
1 Sym8730-1 EMC SYMMETRI X
2 Sym8730-2 EMC SYMMETRI X
```

Create the mapping from the EMC Symmetrix ports to the SVC. In this example we are using ports FA-4A and FA-13A.

Prior to changing the mapping we recommend that you verify the EMC Symmetrix port flag settings for each port to be assigned to the SVC. In our example we verify from the **ECC Port Flags Settings** window that only **H**, **PP**, and **UWN** are set.

Note: Refer to the SVC version 4.1 Configuration Guide for details of flag settings for EMC Symmetrix ports:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

Use the ECC Help to obtain descriptions of the port flag bit settings.

Creating the SVC on the EMC Symmetrix

Create a new host on the EMC Symmetrix, in our case GREEN_SVC, defining all of the SVC WWPNs.

On the EMC Symmetrix remove the LUN masking to the original host.

If you are using different EMC Symmetrix ports to connect to the SVC, then move the mapping of the LUNs to the new ports. In our example we moved the devices 016 and 019 to the new ports FA-4A and FA-13A.

You can now grant the masking of the devices to the SVC to enable them to be seen as MDisks.

Configure the SVC

When access to a volume is changed from direct access on the existing storage controller (un-virtualized) to access through the SVC (virtualized), you must define an *Image Mode* Virtual Disk (VDisk) pointing to the LUN to preserve the existing data on the disk.

Once the LUN has been un-mapped from the host and mapped to the SVC, use the **svctask detectmdisk** command to discover the new MDisks.

The **svcinfolsmdisk** command produces no user feedback, so use the command to list all MDisks in the cluster.

Example 6-14 shows the result of these commands on our system.

Example 6-14 Discover the new disk from EMC Symmetrix

```
IBM_2145:GREEN_SVC:admin>svctask detectmdisk
IBM_2145:GREEN_SVC:admin>svcinfolsmdisk
```

id	name	status	mode	mdisk_grp_name	capacity	ctrl_LUN_#	controller
0	mdisk0	online	managed	RedBook_DS8K	25.0GB	4011400000000000	DS8000
1	mdisk1	online	managed	RedBook_DS8K	25.0GB	4011400100000000	DS8000
2	mdisk2	online	managed	RedBook_DS8K	25.0GB	4011400200000000	DS8000
3	mdisk3	online	managed	RedBook_DS8K	25.0GB	4011400300000000	DS8000
4	mdisk4	online	unmanaged		25.3GB	0000000000000001	Sym8730-1
5	mdisk5	online	unmanaged		25.3GB	0000000000000001	Sym8730-2

```
IBM_2145:GREEN_SVC:admin>
```

Two new MDisks called *mdisk4* and *mdisk5* with IDs of 4 and 5, and size 25.3GB, have been allocated from the EMC Symmetrix.

We now create a special type of Virtual Disk (VDisk) called *Image mode*. This creates a one-to-one mapping of the virtual-to-physical addresses to protect the integrity of the data on the volume.

Important: When importing a data volume into the SVC using the CLI, you must create the VDisk in Image Mode (there is a special **Create an Image Mode VDisk** option if you use the GUI). Creating it in any other mode will corrupt the data. It is good practice to create an MDG called IMAGE (or similar) as a default location for image mode disks. This helps to remind the administrator to select image mode.

Use the following command on the SVC to create an MDG called *IMAGE* and, because you will probably wish to integrate the data to another MDG later, define the appropriate extent size, equal to the extent size in the destination MDG, which in our example is 512Mb:

```
svctask mkmdiskgrp -name IMAGE -ext 512
```

We used the commands shown in Example 6-15 to create two new image mode VDisks called *p615-1* and *p615-2* using MDisks *mdisk4* and *mdisk5*, in MDG *IMAGE*.

Example 6-15

```
svctask mkvdisk -mdiskgrp IMAGE -iogrp 0 -vtype image -mdisk mdisk4 -name p615-1
svctask mkvdisk -mdiskgrp IMAGE -iogrp 0 -vtype image -mdisk mdisk5 -name p615-2
```

Upon successful completion of the commands, we used the **svcin** **lsvdisk** command to list the VDisks.

Example 6-16 Create the Image Mode VDisk

```
IBM_2145:GREEN_SVC:admin>svctask mkmdiskgrp -name IMAGE -ext 512
MDisk Group, id [1], successfully created
IBM_2145:GREEN_SVC:admin>svctask mkvdisk -mdiskgrp IMAGE -iogrp 0 -vtype image
-mdisk mdisk4 -name p615-1
Virtual Disk, id [0], successfully created
IBM_2145:GREEN_SVC:admin>svctask mkvdisk -mdiskgrp IMAGE -iogrp 0 -vtype image
-mdisk mdisk5 -name p615-2
Virtual Disk, id [1], successfully created
IBM_2145:GREEN_SVC:admin>svcin lsvdisk
```

id	name	IO_group_id	IO_group_name	status	mdisk_grp_id	mdisk_grp_name	capacity	type
1	p615-1	0	io_grp0	online	1	IMAGE	25.3GB	image
2	p615-2	0	io_grp0	online	1	IMAGE	25.3GB	image

```
IBM_2145:GREEN_SVC:admin>
```

You can see that there are two VDisks with IDs 1 and 2 called *p615-1* and *p615-2* in MDG *IMAGE*, each with size 25.3GB. The VDisk type is image.

Activate the BLUE zone (Figure 6-1 on page 102) that was created earlier to allow the SVC to map the VDisk to the Windows host.

Now create a host called *p615* on the SVC, specifying the WWPNs of the host HBAs, with the following command:

```
svctask mkhost -name p615 -hbawwpn 100000e08b1881c7:100000e08b0e1fd2
```

The **svcin** **lshost** command lists the hosts.

Example 6-17 Create the Host on the SVC

```
IBM_2145:GREEN_SVC:admin>svctask mkhost -name p615 -hbawwpn
10000000c93581c7:10000000c93a1fd2
Host id [0] successfully created
IBM_2145:GREEN_SVC:admin>svcin lshost
```

id	name	port_count	iogrp_count
0	p615	2	4

```
IBM_2145:GREEN_SVC:admin>
```

The next step is to map the VDisks *p615-1* and *p615-2* to the host *p615*. The commands to do this are shown in Example 6-18.

Example 6-18

```
svctask mkvdiskhostmap -host p615 p615-1
svctask mkvdiskhostmap -host p615 p615-2
```

Upon successful completion, use the **svcinfolsvdisk host map p615-1** command to check the host mappings.

Example 6-19 Map Image Mode VDisk to Windows host

```
IBM_2145:GREEN_SVC:admin>svctask mkvdiskhostmap -host p615 p615-1
Virtual Disk to Host map, id [1], successfully created
IBM_2145:GREEN_SVC:admin>svctask mkvdiskhostmap -host p615 p615-2
Virtual Disk to Host map, id [2], successfully created
IBM_2145:GREEN_SVC:admin>svcinfolsvdiskhostmap p615-1
id name          SCSI_id host_id host_name  wwpn
0 p615-1         1          1      p615      10000000c93581c7
0 p615-1         1          1      p615      10000000c93a1fd2
IBM_2145:GREEN_SVC:admin>
```

The preparation is now complete and we are ready to proceed with the storage virtualization.

6.1.5 Integration into the SVC environment

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SAN Volume Controller.

Configure the host system

Now that the SAN fabric has been rezoned and the SAN Volume Controller has been configured to present the host LUNs, you can now install the IBM SDD and access the LUNs.

Microsoft Windows

Restart the host system.

When the host reboots, it should now see the same LUNs as when directly attached, but now they are being presented through SAN Volume Controller.

Run Disk Management and Rescan Disks. Figure 6-6 shows the discovery of the 25.29GB dynamic disk that was formerly the E: drive.

Now, with Disk Management, add a drive letter and restart your applications.

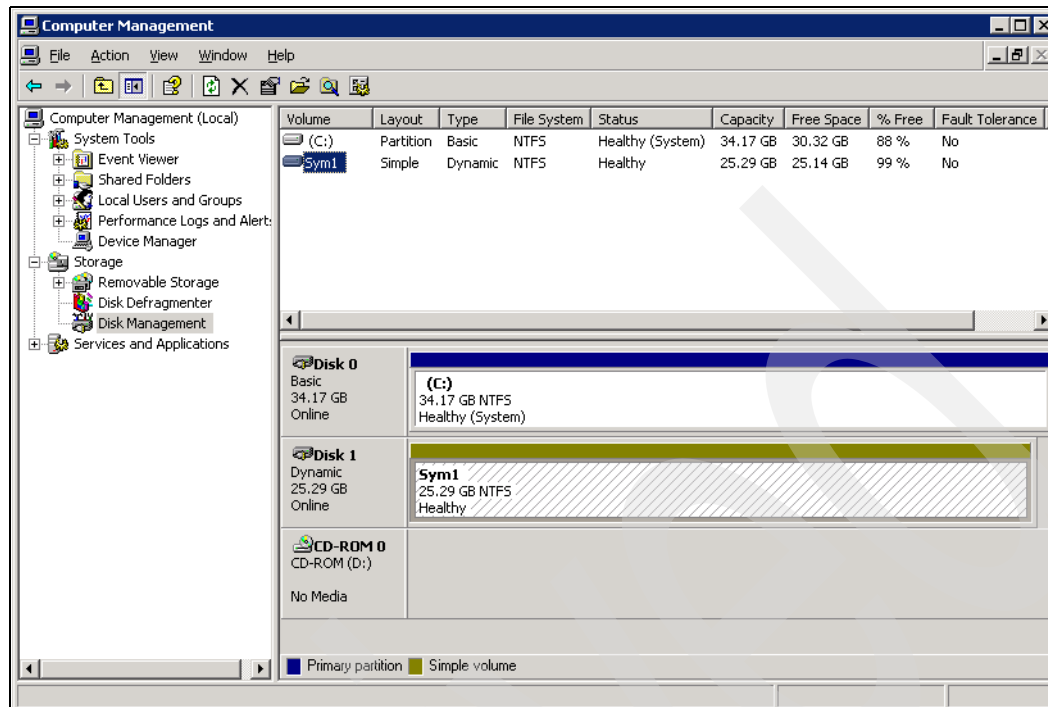


Figure 6-6 SVC VDisk detected by Windows host

Figure 6-7 shows the Windows Device Manager screen with the 8 paths to the SVC (IBM 2145 SCSI Disk Device). 2145 is the product number for the SVC. There is also a 2145 SDD Disk Device, which is the logical volume presented to Windows.

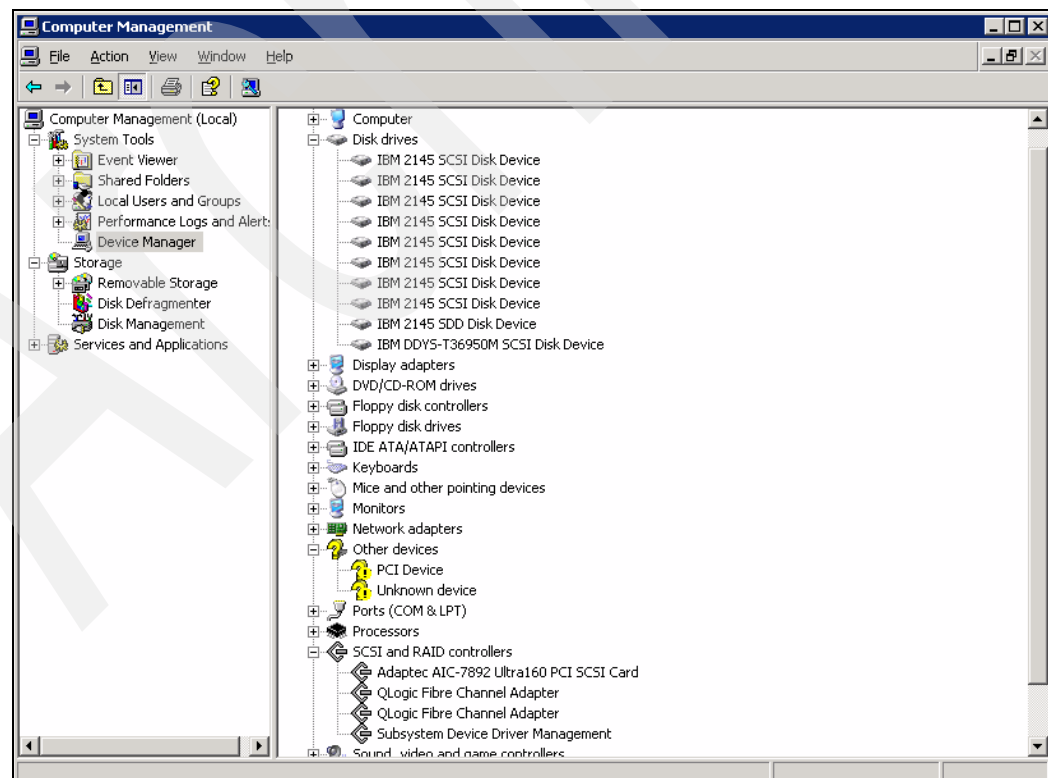


Figure 6-7 SDD devices on Windows Device Manager

AIX

Having successfully installed the SDD, run the **cfgmgr** command to discover the SVC disks.

This command produces no user feedback, so use the **lsdev -Cc disk** command to display the disk devices.

You should see a SAN Volume Controller Device for each path to the SVC VDisks and a Data Path Optimizer Pseudo Device Driver for each logical device. Example 6-20 shows the output of these commands.

Example 6-20 Discovery of SDD devices on AIX

```
root@p615-tic-6[/] # cfgmgr

root@p615-tic-6[/] # lsdev -Cc disk
hdisk0 Available 10-60-00-0,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-60-00-1,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 10-68-01      SAN Volume Controller Device
hdisk3 Available 10-68-01      SAN Volume Controller Device
hdisk4 Available 10-68-01      SAN Volume Controller Device
hdisk5 Available 10-68-01      SAN Volume Controller Device
hdisk6 Available 10-68-01      SAN Volume Controller Device
hdisk7 Available 10-68-01      SAN Volume Controller Device
hdisk8 Available 10-68-01      SAN Volume Controller Device
hdisk9 Available 10-68-01      SAN Volume Controller Device
vpath0 Available              Data Path Optimizer Pseudo Device Driver
vpath1 Available              Data Path Optimizer Pseudo Device Driver
root@p615-tic-6[/] #
```

If you do not see the devices as expected, check that the SDD has been installed correctly. If necessary, check that the SAN zoning is correct and enabled. Also verify that the SVC is presenting the LUNs to the host.

AIX writes a Physical Volume Identifier (PVID) to each LUN, so this means that the newly presented disks should be recognized as part of the SymmVG volume group.

Example 6-21 shows the output from the **lspv** command. You should see that the vpaths are assigned to their respective volume groups and you will also notice that they are not yet active.

Example 6-21 List physical volumes

```
root@p615-tic-6[/] # lspv
hdisk0      0051a28abb619770      rootvg      active
hdisk1      0051a28a1f16b48c      None
hdisk2      none                  None
hdisk3      none                  None
hdisk4      none                  None
hdisk5      none                  None
hdisk6      none                  None
hdisk7      none                  None
hdisk8      none                  None
hdisk9      none                  None
vpath0      0051a28a471f2ffb      SymmVG
vpath1      0051a28a472a5224      SymmVG
root@p615-tic-6[/] #
```

You now need to reactivate the volume group. To do this, issue the command **varyonvg SymmVG**.

Example 6-22 shows the output from this command.

Example 6-22 List physical volumes output following varying on the volume groups.

```

root@p615-tic-6[/] # varyonvg SymmVG
root@p615-tic-6[/] # lspv
hdisk0          0051a28abb619770          rootvg          active
hdisk1          0051a28a1f16b48c          None
hdisk2          none                     None
hdisk3          none                     None
hdisk4          none                     None
hdisk5          none                     None
hdisk6          none                     None
hdisk7          none                     None
hdisk8          none                     None
hdisk9          none                     None
vpath0          0051a28a471f2ffb          SymmVG          active
vpath1          0051a28a472a5224          SymmVG          active
root@p615-tic-6[/] #

```

You can now remount the logical volumes.

In our example we have one mount point, /emc, so the command is **mount /emc**.

The data stored on the EMC Symmetrix is now available to the host system again.

6.2 Integrating the EMC FC4700 with the SVC

In this section, we demonstrate the integration of the SVC in an existing environment with an EMC FC4700 storage controller. The existing storage volumes on the EMC FC4700 storage controller will be virtualized using the SAN Volume Controller. The procedure is very similar for all EMC CLARiiON models that can be managed using the Navisphere software supplied with the equipment by EMC. The example shows how to achieve integration where the EMC storage controller has Access Logix installed.

6.2.1 Scenario description

In this scenario we start with an EMC FC4700 storage controller and a dual redundant SAN connected AIX and HP-UX host. We then introduce the SAN Volume Controller to virtualize the back-end storage.

The systems described in this section include:

- ▶ EMC FC4700 storage controller
- ▶ IBM b-type family switches
- ▶ Host system running AIX or HP-UX

Figure 6-8 shows the relationships between the systems used in this scenario.




Figure 6-8 SAN Volume Controller integration example with an existing EMC FC4700 controller

6.2.2 Summary of integration process

At a high level, the steps to integrate the SVC into the existing production environment are the following:

1. Preparation
 - a. General
 - i. Review SAN levels of software, firmware, and supported hardware using information from the configuration guide.
 - ii. Determine SAN Volume Controller worldwide port names and Host worldwide port names.
 - b. SAN
 - i. Create host zone with SVC (enable now or later).
 - ii. Create EMC FC4700 controller zone with SVC (enable now or later).

Note: In this example the SVC Cluster is already set up and installed. For details on how to set up and install an SVC cluster consult *IBM System Storage SAN Volume Controller*, SG24-6423.

<http://www.redbooks.ibm.com/abstracts/sg246423.html?Open>

- c. FC4700 Controller
 - i. Identify the SCSI LUN IDs and Unique IDs (UID) for all those to be integrated under SAN Volume Controller.

2. Integration to SAN Volume Controller managed storage
 - a. Host (dependent on host type)
 - i. Stop applications.
 - ii. Remove existing multipath driver if not supported by SAN Volume Controller.
 - iii. Install SDD or equivalent driver.
 - iv. Shut down machine or unmount file systems and vary off volume groups.
 - b. SAN
 - i. Enable new zone for host-SAN Volume Controller and FC4700 controller-SAN Volume Controller.
 - ii. Disable FC4700 controller-host zone.
 - c. FC4700 Controller
 - i. Unmap LUNs from host.
 - ii. Remap LUNs to SVC.
 - d. SAN Volume Controller
 - i. Discover FC4700 controller LUNs.
 - ii. Create Image Mode VDisks.
 - iii. Create host with associated WWPNS.
 - iv. Map VDisks to host.
 - e. Host (dependent on host)
 - i. Reboot, or
 - ii. Scan for new devices, vary on volume group, and remount file systems.

Note: We recommend that you review 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization example.

6.2.3 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SAN Volume Controller is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Confirm support for your environment

This section provides information about configuring the SAN Volume Controller in an existing environment with an EMC CLARiiON storage controller.

Prior to integrating the SAN Volume Controller it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SAN Volume Controller supports most models in the EMC CLARiiON range.

The models of EMC CLARiiON that are supported by the SAN Volume Controller are the following:

- ▶ FC4700-1
- ▶ FC4700-2
- ▶ CX200
- ▶ CX300
- ▶ CX400
- ▶ CX500
- ▶ CX600
- ▶ CX700

This list is current at the time of writing. See the following Web site for the latest EMC CLARiiON supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Sharing the EMC CLARiiON between a host and the SAN Volume Controller

The EMC CLARiiON can be shared between a host and a SAN Volume Controller, subject to the following limitations:

- ▶ Split controller access is only supported when Access Logix is installed and enabled.
- ▶ LUs must not be shared between a host and a SAN Volume Controller.
- ▶ Partitions in a RAID group must not be shared between a host and a SAN Volume Controller.

Note: A single host must not be connected to both the SVC and an EMC Symmetrix at the same time because the multipathing drivers IBM Subsystem Device Driver (SDD) and EMC PowerPath cannot coexist on the host.

Switch zoning considerations for the EMC CLARiiON

Following are switch zoning limitations for the SAN Volume Controller and EMC CLARiiON that you need to be aware of:

- ▶ The EMC CLARiiON must be configured to present logical units to all SVC initiator ports in the fabric zone.
- ▶ Only the LUN-masked SVC initiator ports on the EMC CLARiiON controller should be present in the fabric zone.
- ▶ One of the considerations is the number of connections (as defined by process logins) consumed by the SVC cluster and the EMC CLARiiON. To determine the number of connections for a single fabric, use the following calculation:

Number of SVC nodes × number of initiator ports × number of target ports

If this exceeds the capabilities of the subsystem, take care to reduce the number of initiator or target ports in the configuration, but do not create a single point of failure.

- If you need to decrease the number of initiator ports, use two of the four ports on each SAN Volume Controller node (which equates to one per HBA) and configure two fabrics (fabric zones), ensuring these are the only initiator ports that are visible to the target ports.
- To decrease the number of target ports, utilize ports from more than one controller.
- ▶ The EMC CLARiiON CX200 presents two ports and will support 30 connections. Therefore, if you are utilizing a single SAN fabric, a 4-node cluster will require 32 connections (4 x 4 x 2). This is in excess of the EMC CLARiiON CX200's capability and compromises the SVC cluster integrity. Since only two target ports are available, you must

decrease the number of initiator ports. This will utilize only 16 of the 30 connections that are available.

Note: The EMC CLARiiON CX200 cannot be utilized in an 8-node cluster configuration. The reason is that the number of initiator ports cannot be fewer than 16 (2 on each node) and the number of target ports cannot be less than 2. This will require 32 connections and will exceed the subsystem limit.

- ▶ Both the EMC CLARiiON FC4700 and EMC CLARiiON CX400 systems present four target ports, and 64 connections are supported. Therefore, if you are utilizing a single SAN fabric, a 4-node cluster will require 64 connections (4 x 4 x 4). This will equal the capabilities of the EMC CLARiiON and may only be a problem if split support with other hosts is needed. Decreasing either the number of initiator ports or the number of target ports will require 32 of the available 64 connections.
- ▶ The EMC CLARiiON CX600 presents 8 target ports and the number of connections supported is 128. A 4-node cluster will require 128 connections (4 x 4 x 8). An 8-node cluster is in excess of the connection limit, and therefore none of the reduction schemes can be used.

Quorum disk support

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the EMC CLARiiON controllers as quorum disks.

Further information

It is strongly recommended that you refer to the SAN Volume Controller V4.1 Configuration Guide for further information regarding integrating the SAN Volume Controller with an existing storage controller:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

Preparing the SAN fabric

We need to prepare two additional zones. One zone allows SAN Volume Controller to be visible to the host whose data is to be virtualized (Figure 6-1 BLUE zone), and the second zone enables the SVC to be visible to the EMC FC4700 (Figure 6-1 RED1 zone). Depending upon your environment, you may choose to simply create these zones at this time and only enable them later when the host applications have been stopped.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendor’s SAN equipment.

Note: It is assumed that the BLACK zone, which is the SVC node to SVC node zone, was created at the same time as the SVC cluster and therefore already exists in the SAN.

Record the SAN Volume Controller WWPN

You will need the SAN Volume Controller worldwide port names (WWPNs) for each of the SAN Volume Controller nodes. This will enable you to configure the EMC FC4700 to make LUNs visible to the SAN Volume Controller. There are several methods of collecting this information. You can use the SAN switch interface to identify the WWPNs, provided you know which ports the SAN Volume Controller nodes are connected to. Alternatively, you can use the SAN Volume controller GUI or CLI.

The following example shows how to use the GUI:

1. Open a browser window and connect to the San Volume Controller GUI.
2. Select **Work With Nodes** → **Nodes**
Figure 6-9 shows the resulting GUI panel on our system.




Figure 6-9 Selecting the node to determine the associated WWPNs

3. Select each node in turn by clicking the node name link (Figure 6-9). The panel shown in Figure 6-10 is then presented, showing general data about that node.




Figure 6-10 Selecting the Port Information for the node

4. Select the Port information link (Figure 6-10) to show the four WWPNs associated with that node. You will see a panel similar to Figure 6-11. Record the WWPNs to enable you to map the EMC FC4700 LUNs to the SAN Volume Controller later on.




Figure 6-11 List of WWPNs for the node

Repeat these steps for each of the SAN Volume Controller nodes in your SAN Volume Controller cluster.

Record the host WWPNs

There are several methods of identifying the WWPNs of the host involved in the integration. The simplest method is to capture the information through the SAN switch. Alternatively, you can access the information through the host through the host bus adapter BIOS. On some platforms you can capture the information in a configuration file.

Preparing the EMC FC4700

We need to identify the LUNs to be integrated and capture their SCSI LUN identifiers and their Unique IDs. This is important when we present the LUNs through SAN Volume Controller because some host platforms will expect to see the LUNs presented with the same SCSI LUN identifier. The Unique ID helps to identify the LUNs, having presented them to SAN Volume Controller. If you have a large number of LUNs, this is the best method to confirm which they are.

The administrative interface on the FC4700 can be accessed through a Web browser using the EMC Navisphere software.

Usually a *Storage Group* is created for each host or host cluster. Find the storage group that includes the host you are planning to integrate. In our case the host name is *Frank*.

When you expand the view with the host *Frank*, the associated LUNs presented to that host are displayed.

For each LUN presented, capture the unique identifier (UID). This will help in confirming the LUNs when they are presented to the SAN Volume Controller later.

To collect this information, point to a LUN, right-click it, and select the **Properties** link. On the General tab the Unique ID for the LUN is displayed; record it for later use. Repeat these steps for each LUN assigned to the host to be integrated.

6.2.4 Integration of SVC managed storage

In this section we describe the steps performed on the hosts in our scenario.

Host

The SVC is provided with a multipath driver called the IBM Subsystem Device Driver (SDD). The driver manages the resiliency and connectivity between the host and the virtualized storage. Multipath drivers from other vendors are also tested and supported. For the most up-to-date support matrix, refer to the SAN Volume Controller supported hardware list at:

<http://www.ibm.com/storage/support/2145>

HP-UX

In our example, the HP-UX system was using the native PV Links multipath driver to access the FC4700 LUNs.

Use the **ioscan -fnC disk** command to determine what disks are being presented to the host system.

Example 6-23 shows the output of this command on our system.

Example 6-23 Output of ioscan command

```
bash-2.04# ioscan -fnC disk
Class      I  H/W Path          Driver  S/W State  H/W Type  Description
=====
disk       2  0/0/2/0.6.0       sdisk   CLAIMED    DEVICE    SEAGATE ST336704LC
              /dev/dsk/c2t6d0   /dev/rdisk/c2t6d0
disk       3  0/0/2/1.6.0       sdisk   CLAIMED    DEVICE    SEAGATE ST318404LC
              /dev/dsk/c3t6d0   /dev/rdisk/c3t6d0
disk       4  0/4/0/0.33.2.0.0.0.0 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c36t0d0   /dev/rdisk/c36t0d0
disk       9  0/4/0/0.33.2.0.0.0.1 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c36t0d1   /dev/rdisk/c36t0d1
disk      10  0/4/0/0.33.2.0.0.0.2 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c36t0d2   /dev/rdisk/c36t0d2
disk       5  1/10/0/0.97.6.19.0.0.0 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c35t0d0   /dev/rdisk/c35t0d0
disk       6  1/10/0/0.97.6.19.0.0.1 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c35t0d1   /dev/rdisk/c35t0d1
disk       7  1/10/0/0.97.6.19.0.0.2 sdisk   CLAIMED    DEVICE DGC      C4700WDR5
              /dev/dsk/c35t0d2   /dev/rdisk/c35t0d2
```

This shows the raw disk devices attached to the HP-UX system. It shows 2 different paths to the same LUNs. The number in bold is the SCSI LUN ID being used by the FC4700.

The raw disks of interest here are c36t0d0, d1, d2 and c35t0d0, d1 and d2. We now need to find which volume group these belong to and which logical volumes they have mounted on them.

Use the **vgdisplay -v** command to determine what volume group they are part of.

Example 6-24 shows the output of this command on our system. Information for /dev/vg00 has been removed for clarity.

Example 6-24 List physical volume to volume group mapping

```
bash-2.04# vgdisplay -v
--- Volume groups ---
VG Name                                /dev/vg00
.....
VG Name                                /dev/FC4700_PV
VG Write Access                         read/write
VG Status                              available
Max LV                                 255
Cur LV                                1
Open LV                                1
Max PV                                 16
Cur PV                                3
Act PV                                 3
Max PE per PV                         5120
VGDA                                   6
PE Size (Mbytes)                       4
Total PE                              8957
Alloc PE                              5000
Free PE                               3957
Total PVG                             0
Total Spare PVs                       0
Total Spare PVs in use                0

--- Logical volumes ---
LV Name                                /dev/FC4700_VG/FC4700_LV
LV Status                             available/syncd
LV Size (Mbytes)                      5000
Current LE                            1250
Allocated PE                          1250
Used PV                               1

--- Physical volumes ---
PV Name                                /dev/dsk/c36t0d0
PV Name                                /dev/dsk/c35t0d0 Alternate Link
PV Status                             available
Total PE                              1279
Free PE                               0
Autoswitch                            On

PV Name                                /dev/dsk/c36t0d1
PV Name                                /dev/dsk/c35t0d1 Alternate Link
PV Status                             available
Total PE                              2559
Free PE                               0
Autoswitch                            On

PV Name                                /dev/dsk/c36t0d2
PV Name                                /dev/dsk/c35t0d2 Alternate Link
PV Status                             available
Total PE                              5119
Free PE                               3957
Autoswitch                            On
```

On the system we are using, note that the FC4700 is presenting three LUNs that are part of the /dev/FC4700_PV volume group.

We now need to stop the applications that access this volume group. In your environment there may be more than one volume group involved, in which case you should stop all relevant applications.

Perform the following steps to install the IBM Subsystem Device Driver and host attachment scripts onto the host system. Do this before unmounting the file systems or the install process will fail. The latest version of these filesets can be downloaded from the Web. Refer to the SAN Volume Controller recommended software levels document, which will direct you to the relevant download site:

<http://www.ibm.com/storage/support/2145>

1. Determine the version of HP-UX you are running and download the required SDD fileset. The system in our example is running HP-UX 11i PA-RISC, so the fileset to be downloaded is:

IBMsddHP64_050906.depot.tar

Place the fileset in the same directory. We recommend creating a new directory called IBM.

2. Untar the file by running the command **tar -xvf IBMsddHP64_050906.depot.tar**

Note that you will need to change the file name to suit the fileset for your version of HP-UX.

This will generate an unpacked fileset. Example 6-25 shows the output from listing the directory in which it was untarred.

Example 6-25 List of files following untar command

```
bash-2.04# ls -l
total 24496
-rw-r--r--  1 500      513      6256640 Sep  8  2005 IBMsddHP64_050906.depot
-rw-r-----  1 root      sys      6266880 Aug  2 16:14 IBMsddHP64_050906.depot.tar
```

3. Install this fileset using SAM or through the command line. Example 6-26 shows how we set the install source directory as SDD and then installed the fileset in that directory.

Example 6-26 Command to install through the command line

```
bash-2.04# swinstall -s /SDD/IBMsddHP64_050906.depot IBMsdd_tag

===== 08/04/06 11:38:10 BST BEGIN swinstall SESSION
(non-interactive) (jobid=frank-0080)

* Session started for user "root@frank".

* Beginning Selection
* Target connection succeeded for "frank:/".
* Source:                /SDD/IBMsddHP64_050906.depot
* Targets:                frank:/
* Software selections:
    IBMsdd_tag.IBMsdd_commands,r=1.6.0.12,a=HP-UX_B.11.00_64,v=IBM
    IBMsdd_tag.IBMsdd_kernel,r=1.6.0.12,a=HP-UX_B.11.00_64,v=IBM
* Selection succeeded.

* Beginning Analysis and Execution
```

- * Session selections have been saved in the file
"/.sw/sessions/swinstall.last".
- * The analysis phase succeeded for "frank:/".
- * The execution phase succeeded for "frank:/".
- * Analysis and Execution succeeded.

NOTE: More information may be found in the agent logfile using the command "swjob -a log frank-0080 @ frank:/".

```
===== 08/04/06 11:39:00 BST  END swinstall SESSION (non-interactive)
(jobid=frank-0080)
```

4. Check the output to ensure that the fileset was installed successfully. The output will show the cause of any failures.

You can now unmount the file systems and prepare the volume groups for export. Do this for each mount point and volume group.

Use the **mount** command to determine what mount points the volume group has.

Example 6-27 shows the logical volumes associated with the host system volume groups and the relevant mount points.

Example 6-27 Mapping between volume group and logical volumes

```
bash-2.04# mount
/ on /dev/vg00/lvol3 log on Thu Jul 13 15:36:07 2006
/stand on /dev/vg00/lvol1 defaults on Thu Jul 13 15:36:10 2006
/var on /dev/vg00/lvol8 delaylog,nodatainlog on Thu Jul 13 15:36:32 2006
/usr on /dev/vg00/lvol7 delaylog,nodatainlog on Thu Jul 13 15:36:32 2006
/tmp on /dev/vg00/lvol4 delaylog,nodatainlog on Thu Jul 13 15:36:32 2006
/opt on /dev/vg00/lvol6 delaylog,nodatainlog on Thu Jul 13 15:36:33 2006
/home on /dev/vg00/lvol5 delaylog,nodatainlog on Thu Jul 13 15:36:33 2006
/EMC on /dev/FC4700_VG/FC4700_LV delaylog,nodatainlog,nolargefiles on Wed Aug 2 13:56:10 2006
```

You will see that the FC4700_VG volume has a single mount point, /EMC.

Before you go any further, make a configuration backup of the volume groups you plan to integrate.

To do this run the **vgcfgbackup Volume Group** command.

Unmount the file system and vary off the volume group in readiness for integrating it into the SVC environment.

To do this run the **/usr/sbin/umount Filesystem** command.

Deactivate the volume group using the **vgchange -a n Volume Group** command.

Export the volume group using the **vgexport Volume Group** command.

Example 6-28 shows this series of commands being run on our system.

Example 6-28 Unmount of file system and export volume group

```
bash-2.04# vgcfgbackup FC4700_VG
Volume Group configuration for /dev/FC4700_VG has been saved in
/etc/lvmconf/FC4700_VG.conf
bash-2.04# /usr/sbin/umount /EMC
```

```

bash-2.04#
bash-2.04# vgchange -a n FC4700_VG
Volume group "FC4700" has been successfully changed.
bash-2.04# vgexport FC4700_VG
bash-2.04#

```

Now that we have installed the IBM Subsystem Device Driver and exported the volume groups, we can reconfigure the SAN fabric and map the existing EMC FC4700 LUNs to the SAN Volume Controller.

AIX

On this platform, all the changes to the host system can be made without a full shutdown. In our example, the AIX system had the EMC PowerPath multipathing driver installed.

Use the **lsdev -Cc disk** command to determine what disks are being presented to the host system. Example 6-29 shows the output of this command on our system.

Example 6-29 List disks presented by EMC Clariion to AIX host

```

[root@wibbler]/
>lsdev -Cc disk
hdisk0      Available 10-60-00-0,0 16 Bit LVD SCSI Disk Drive
hdisk1      Available 10-60-00-1,0 16 Bit LVD SCSI Disk Drive
hdisk2      Available 10-68-01      EMC CLARiion FCP RAID 5 Disk
hdisk3      Available 10-68-01      EMC CLARiion FCP RAID 5 Disk
hdisk4      Available 20-58-01      EMC CLARiion FCP LUNZ Disk
hdisk5      Available 20-58-01      EMC CLARiion FCP LUNZ Disk
hdisk6      Available 20-58-01      EMC CLARiion FCP RAID 5 Disk
hdisk7      Available 20-58-01      EMC CLARiion FCP RAID 5 Disk
hdiskpower0 Available 10-68-01      PowerPath Device
hdiskpower1 Available 10-68-01      PowerPath Device
[root@wibbler]/
>

```

Use the **lspv** command to determine what volume group they are part of. Example 6-30 shows the output of this command on our system

Example 6-30 List physical volume to volume group mapping

```

[root@wibbler]/emc/tmp
>lspv
hdisk0      0051a28abb619770      rootvg      active
hdisk1      0051a28a1f16b48c      None
hdisk2      none      None
hdisk3      none      None
hdisk4      none      None
hdisk5      none      None
hdisk6      none      None
hdisk7      none      None
hdiskpower0 0051a28a471f2ffb      ClarVG      active
hdiskpower1 0051a28a472a5224      ClarVG      active
[root@wibbler]/emc/tmp
>

```

On the system we are using you will see that the EMC Clariion is presenting two disks, hdiskpower0 and hdiskpower1, which are part of the ClarVG volume group.

We now need to stop the applications that access this volume group. In your environment there may be more than one volume group involved, in which case you should execute the following commands with that in mind.

Use the **lsvg -o |lsvg -i -l** command to determine what mount points the volume group has.

Example 6-31 shows the logical volumes associated with the host system volume groups and the relevant mount points.

Example 6-31 Mapping between volume group and logical volumes

```
[root@wibbler]/emc/tmp
>lsvg -o |lsvg -i -l
ClarVG:
LV NAME          TYPE      LPs  PPs  PVs  LV STATE  MOUNT POINT
emclv             jfs       500  500   2  open/syncd  /emc
loglv01           jfslog    1    1    1  open/syncd  N/A
rootvg:
LV NAME          TYPE      LPs  PPs  PVs  LV STATE  MOUNT POINT
hd5              boot      1    1    1  closed/syncd N/A
hd6              paging    8    8    1  open/syncd  N/A
hd8              jfs2log   1    1    1  open/syncd  N/A
hd4              jfs2      1    1    1  open/syncd  /
hd2              jfs2      18   18    1  open/syncd  /usr
hd9var           jfs2      1    1    1  open/syncd  /var
hd3              jfs2      2    2    1  open/syncd  /tmp
hd1              jfs2      1    1    1  open/syncd  /home
hd10opt          jfs2      1    1    1  open/syncd  /opt
loglv00          jfslog    1    1    1  open/syncd  N/A
lv00             jfs       4    4    1  open/syncd  /usr/vice/cache
```

Note that the ClarVG has two logical volumes associated. fslv00 is mounted to /emc.

Unmount the file system and vary off the volume group in preparation for integrating it into the SVC environment.

To do this, run **/usr/sbin/umount Filesystem** followed by **varyoffvg Volume Group**.

Example 6-32 shows this set of commands being run on our system. Neither command produces any user feedback.

Example 6-32 Unmount of file system and vary off of volume group

```
[root@wibbler]/emc/tmp
>/usr/sbin/umount /emc
[root@wibbler]/emc/tmp
>/usr/sbin/varyoffvg ClarVG
[root@wibbler]/emc/tmp
>
```

Uninstall the EMC PowerPath device driver and the existing LUN definitions using the following steps:

1. Uninstall the EMC PowerPath driver by running the **installp -u EMCpower** command.

Information about the progress of the removal is displayed, ending with the details shown in Example 6-33.

Example 6-33 PowerPath deinstall

```

+-----+
|                               Summaries:                               |
+-----+
| Installation Summary |
+-----+

```

Name	Level	Part	Event	Result
EMCpower.consistency_grp	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.hr	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.mobility	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_aa	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_ap	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.multi_path_clariion	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.ppvmm	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.vcsagent	4.4.0.0	USR	DEINSTALL	SUCCESS
EMCpower.base	4.4.0.0	USR	DEINSTALL	SUCCESS

```

[root@wibbler]/
>

```

2. Remove the disk definitions by running the **rmdev -dl hdiskn** command, where **n** represents the corresponding number of each hdisk presented by the EMC Clariion.

Example 6-34 shows the command output.

Example 6-34 Removing the EMC Clariion related hdisks

```

[root@wibbler]/
>rmdev -dl hdisk2
hdisk2 deleted
[root@wibbler]/
>rmdev -dl hdisk3
hdisk3 deleted
[root@wibbler]/
>rmdev -dl hdisk4
hdisk4 deleted
[root@wibbler]/
>rmdev -dl hdisk5
hdisk5 deleted
[root@wibbler]/
>rmdev -dl hdisk6
hdisk6 deleted
[root@wibbler]/
>rmdev -dl hdisk7
hdisk7 deleted

```

You now need to install the IBM SDD and host attachment scripts onto the host system. The latest version of these filesets can be downloaded from the Web. Refer to the SAN Volume

Controller recommended software levels document, which will direct you to the relevant download site:

<http://www.ibm.com/storage/support/2145>

1. If you are uncertain which level of AIX you have installed, use the command **oslevel** to discover it. This will determine which filesets you should download. The system in our example is running AIX v5.3, so the two filesets to be downloaded are:

```
devices.fcp.disk.ibm.rte.tar
devices.sdd.53.rte.tar
```

Place the filesets in the same directory. We recommend creating a new directory called IBM.

2. Untar the files by running the commands shown in Example 6-35.

Example 6-35

```
tar -xvf devices.fcp.disk.ibm.rte.tar
tar -xvf devices.sdd.53.rte.tar
```

This will generate two unpacked filesets. Example 6-36 shows the output from listing the directory in which they were untarred.

Example 6-36 List of files following untar command

```
[root@wibbler]/
>ls -l
total 6168
-rwxrwx---1 155350 449628 31744 09 May 23:59 devices.fcp.disk.ibm.rte
-rw-r-----1 root system 40960 21 Jul 16:15 devices.fcp.disk.ibm.rte.tar
-rw-r-----1 root system 1536000 09 May 20:02 devices.sdd.53.rte
-rw-r-----1 root system 1546240 21 Jul 16:15 devices.sdd.53.rte.tar
```

3. Install these filesets using **smit/smitty** or through the command line. Example 6-37 shows how we set the install source directory as IBM and then installed the filesets in that directory.

Example 6-37 Commands to install through the command line

```
[root@wibbler]/
>cd /IBM
[root@wibbler]/
>inutoc $PWD
[root@wibbler]/
>installp -acd . -einstallp.log ALL
```

4. Check the output to ensure that the filesets were installed successfully. The output will show the cause of any failures.

Now that you have uninstalled the EMC drivers and installed the IBM SDD, you can reconfigure the SAN fabric and map the existing EMC Clariion LUNs to the SVC.

Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SAN Volume Controller. (Figure 6-8 on page 121 shows the zones. You should remove the GREEN zone from the active zoning configuration).

Enable the RED1 zone and the BLUE zone (see Figure 6-8) so that the SAN Volume Controller can access the FC4700 LUNs you are about to remap and so that the host under integration can access the SAN Volume Controller.

Refer to Appendix B, “SAN zoning changes” on page 257 for a description of the required SAN zones for integrating the SAN Volume Controller in an existing environment. This appendix also provides examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Configure the EMC FC4700

You need to unmap the LUNs that are currently presented directly to the host and remap them to the SAN Volume Controller nodes.

To unmap the LUNs perform the following steps:

1. Right-click the mouse while pointing to the name of the storage group that contains the host on which SAN volumes are to be virtualized and click **Select LUNs**.
2. In the Storage Group window select the LUNs to be removed from the host mapping. Move each LUN from the **Selected LUNs** pane to the **Available LUNs** pane by dragging and dropping it, then click **Apply** to perform the changes.

You have now removed the host’s access to the LUNs and can start the process of mapping them to SVC. Use the following steps to do the remapping:

1. Register the SAN Volume Controller WWPNs on the FC4700 before you create the new host group. Open the **Connectivity Status** panel to register the WWPNs. Place the pointer over the controller name and right click. Select **Connectivity Status** from the list.
2. The **Connectivity Status** panel is displayed (provided your SAN zoning allowing the SAN Volume Controller node ports to see the FC4700 is in place). This panel includes a list of node ports. Click the **Group Edit** button to assign all these WWPNs to a host. In our example, the host was called *SVC*.
3. In the **Group Edit** panel select all the SAN Volume Controller node ports from the **Available** pane and add them to the **Selected** pane.
4. Set the initiator information. Table 6-2 shows the required settings for SAN Volume Controller.

Table 6-2 Initiator settings for SAN Volume Controller

Parameter	Setting
Initiator Type	CLARiiON Open
Failover Mode	2
Array CommPath	Disabled
Unit Serial Number	Any
Host Name	Any

5. Click **OK** to submit the configuration.

You have now assigned the WWPNs to the SAN Volume Controller host.

Creating a storage group and adding the SVC

To create a new storage group and add the SVC, perform the following steps:

1. Right-click with the pointer over the storage controller name, select **Create Storage Group**, and give it a memorable name. We entered Storage Group SVC.
2. Add the host **SVC** to the newly created host group.
3. At the same time, you can also present the LUNs that were originally directly attached to the host. To do this, right-click with the pointer over the new storage group, in our case Storage Group SVC. Select **Connect Hosts**.
4. In the Storage Group Properties window select the *SVC* and move it to the **Hosts to be Connected** pane.
5. Click **Apply** to submit the changes.

You are now ready to map the LUNs to the SAN Volume Controller host group. Do this using the following steps:

1. Click the **LUN** tab and add the LUNs to the host group by moving them from the **Available LUNs** pane to the **Selected LUNs** pane. As you will see, the LUN IDs remain the same as when presented directly to the host. Ensure that you select the correct LUNs in the event you have other unassigned LUNs.
2. Click **Apply** to submit the changes. You have now assigned the host LUNs to the SAN Volume Controller.

Tip: If you are not familiar with configuring the FC4700 storage controller, refer to the installation and configuration documentation: *EMC Navisphere Manager, ADMINISTRATOR'S GUIDE*, P/N 300-003-511.

Configure the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized. Do this by accessing the SAN Volume Controller GUI and selecting **Work With Managed Disks** → **Disk Controller Systems**.

Figure 6-12 shows the output of this command on our system. Note that a new controller has been identified as *controller2*. We recommend that you rename the controller to something more memorable.

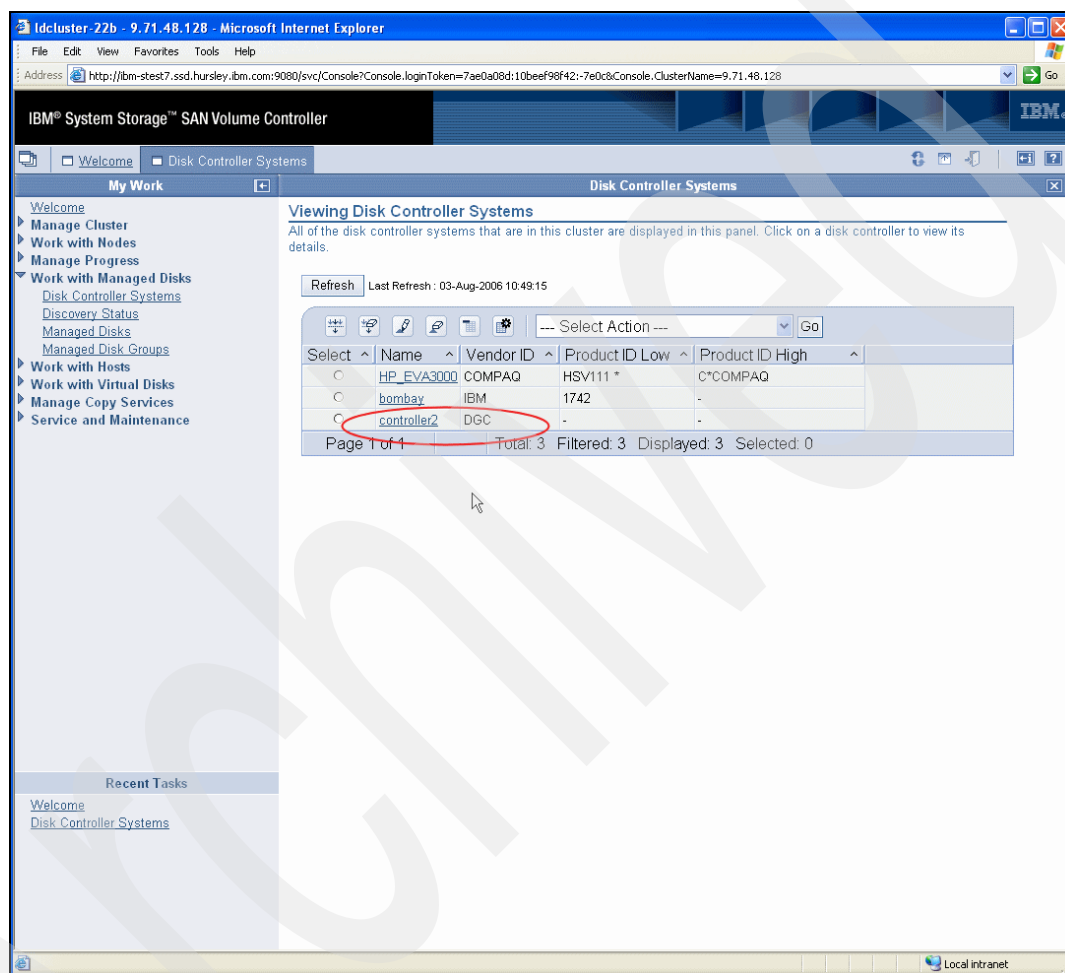


Figure 6-12 List of attached storage controllers

To do this check the **Select** button to the left of the controller name, click the drop-down menu button, and select **Rename a disk controller system**. Figure 6-13 shows this selection.

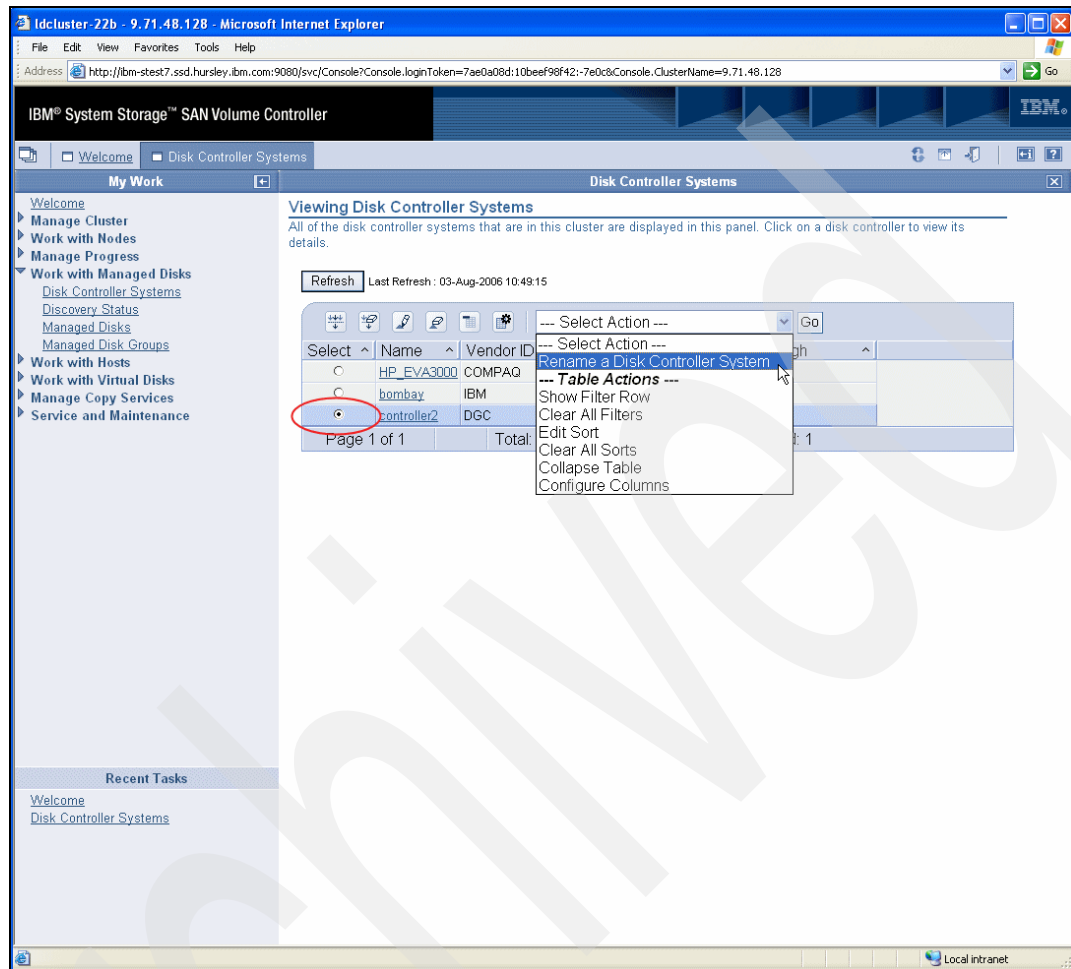


Figure 6-13 Renaming a storage controller

Check which LUNs the EMC FC4700 storage controller is presenting. To do this, select **Work With Managed Disks** → **Managed Disks**. A filter option page will be shown. Click the **Bypass Filter** box.

Figure 6-14 shows the Managed Disk view.

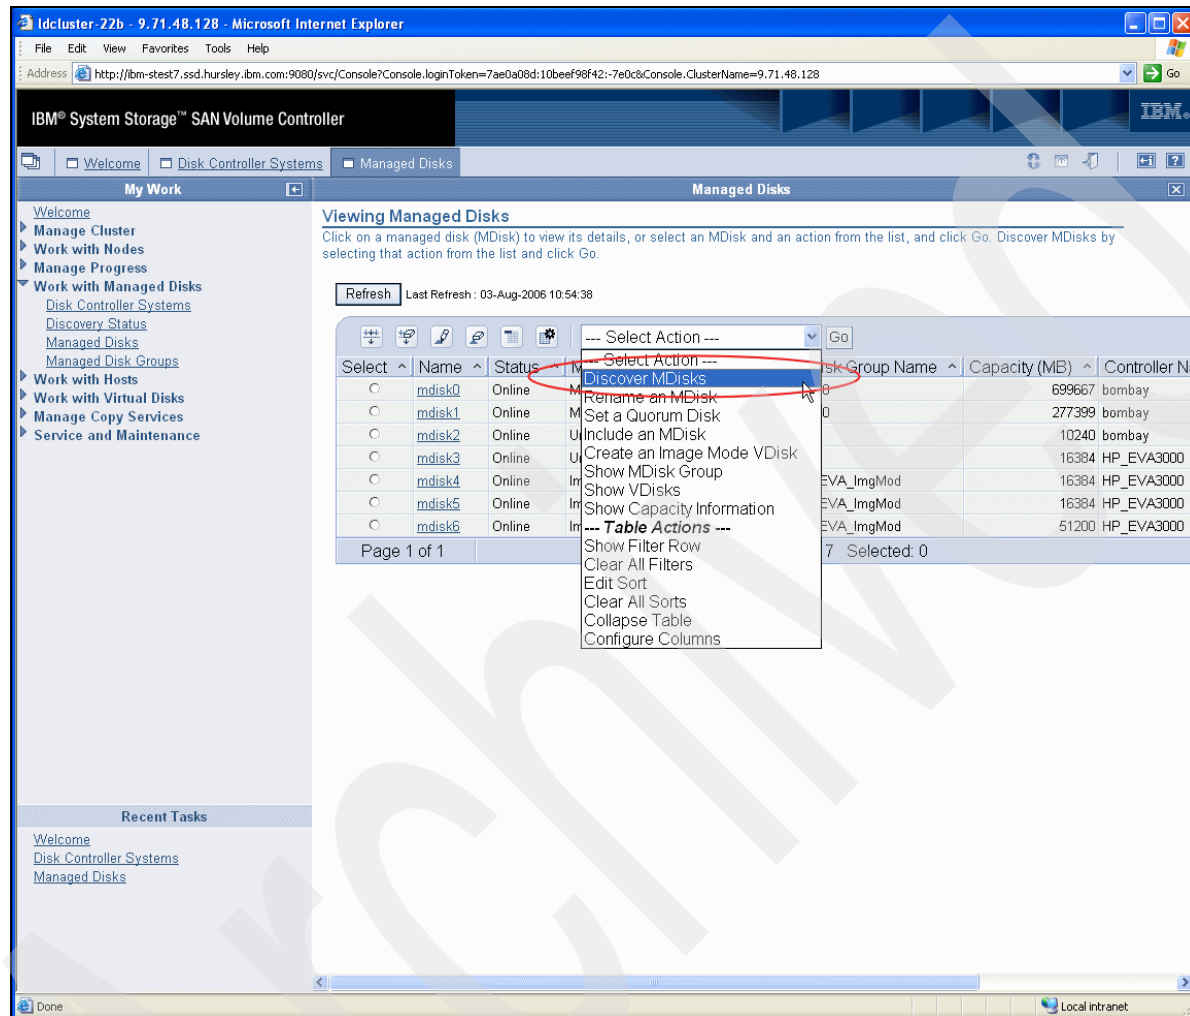


Figure 6-14 Managed Disk view

Since the EMC FC4700 controller has been recently added, it may be that the LUNs it is presenting to SAN Volume Controller have not yet been discovered. Run the **Discover MDisks** command to scan for new devices if this is the case.

Figure 6-15 shows the result for the command.

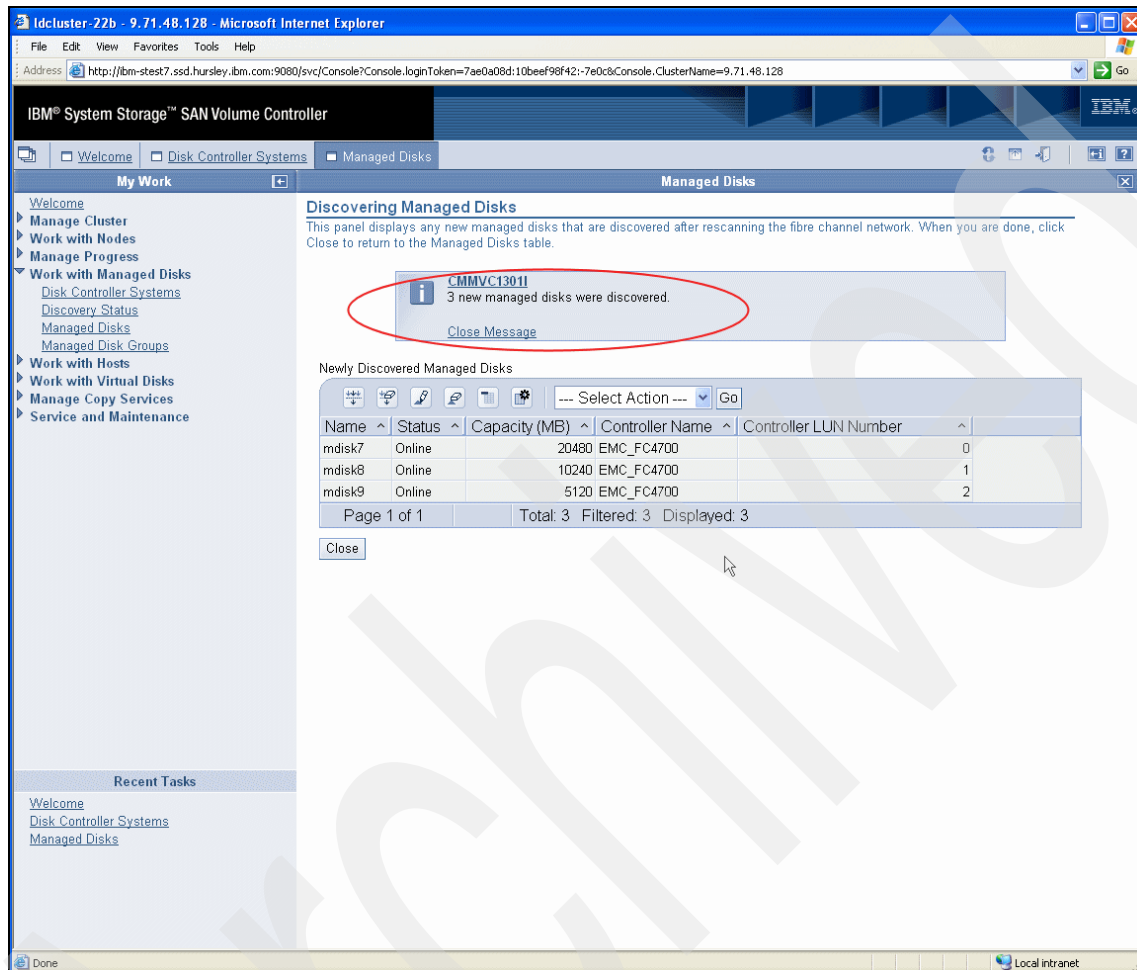


Figure 6-15 Result of discover MDisks command

On our system, three new MDisks were found. Click **Close** to return to the view screen.

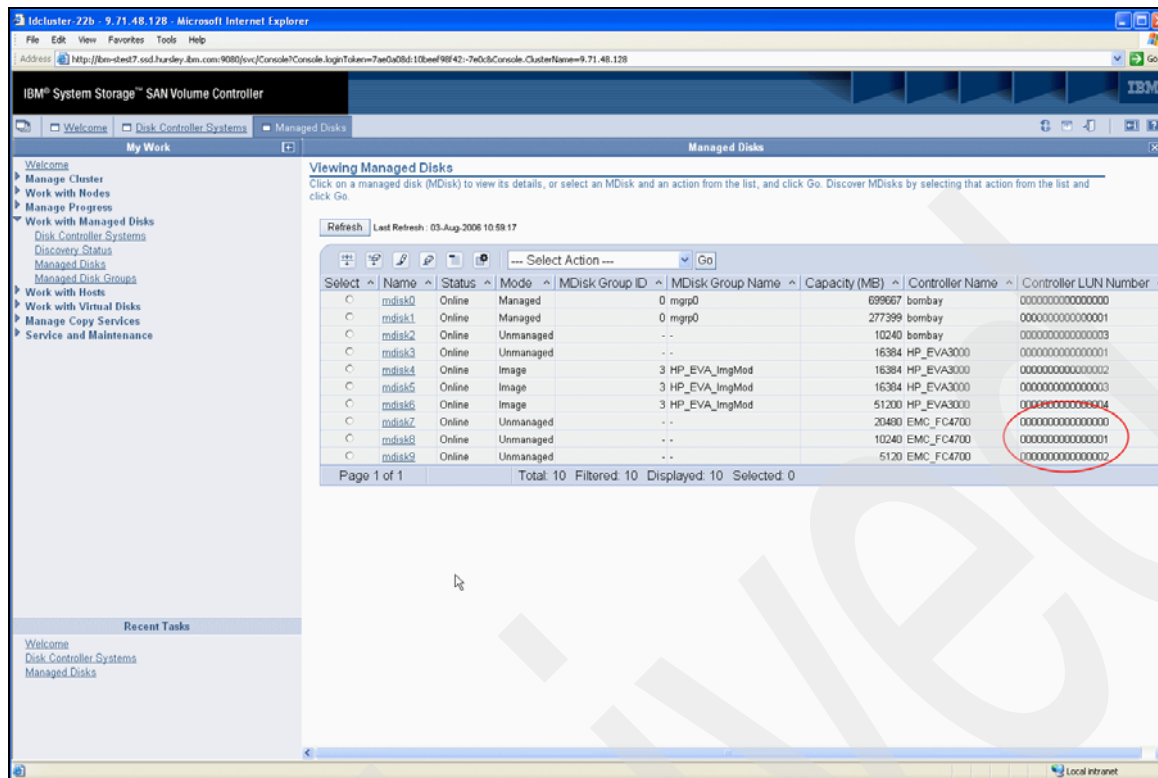


Figure 6-16 Managed Disk View after discovery

Figure 6-16 shows three disks being presented from the FC4700 controller. At present they are in the *Unmanaged* state. You will now create a managed disk group and define image mode VDisks which point to these managed disks. Unlike the SAN Volume Controller CLI, the GUI enables you to create a managed disk group as part of the process of creating an Image Mode VDisk.

Creating the image mode VDisks

Earlier in the process you collected information about the SCSI LUN ID and the Unique ID (UID) to differentiate between the LUNs presented by the FC4700. In Figure 6-16 you can determine the SCSI LUN ID that the FC4700 controller has assigned to each of the MDisks. However, you cannot determine the UID that each MDisk has. If you need to do this to identify each one, then you must use the command line interface, with the command **svcinfo lsmdisk *name***, where *name* is the MDisk in question. The UID is clearly identified. It is important to know the LUN numbers for each of the MDisks you are about to make into Image Mode Disks because you cannot access this information through the GUI once you begin the process.

Perform the following steps to create the image mode VDisks:

1. Select **Work With Virtual Disks** → **Virtual Disks**.
2. Bypass the filter.
3. From the drop-down menu select **Create Image Mode VDisk**. Figure 6-17 on page 142 illustrates selecting this option.

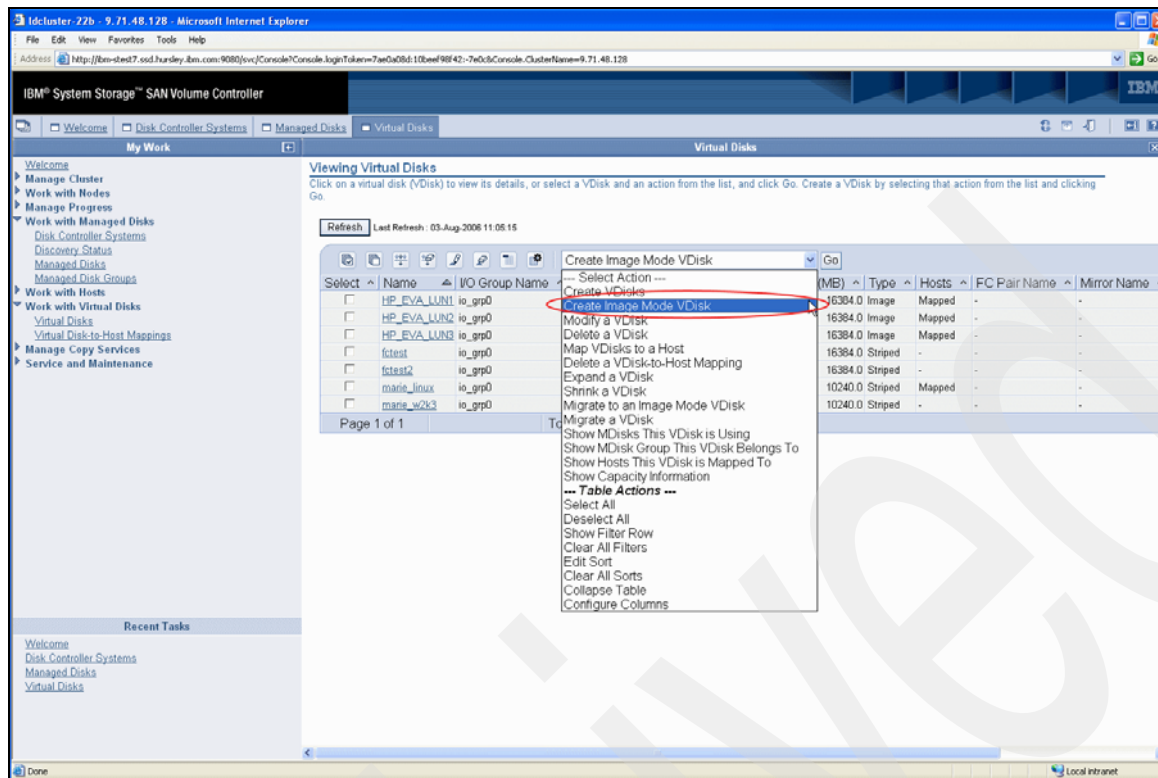


Figure 6-17 Create Image Mode VDisk

4. You are now presented with a wizard that will take you through the whole process. Figure 6-18 shows the overview of the steps. Click **Next**.



Figure 6-18 Create Image Mode Virtual Disk introduction

5. Figure 6-19 shows the second screen in this wizard, on which several attributes are set:
 - Give the Image Mode VDisk a name. We recommend that you give it a memorable name to identify it more easily in the future.
 - Select the cache mode as **Read/Write**.

- The ID is the **SCSI LUN ID** that SAN Volume Controller will use to present this to the host. You should retain the IDs that had been used by the EMC FC4700 when presenting the LUNs directly to the host.
 - If this is the first Image Mode VDisk you have created from the FC4700, check the box **Create an empty MDisk group**. If you need to repeat these steps for additional FC4700 LUNs you can choose not to check this box.
 - Select the EMC FC4700 LUN that you want to create an Image Mode VDisk from.
- Click **Next**.

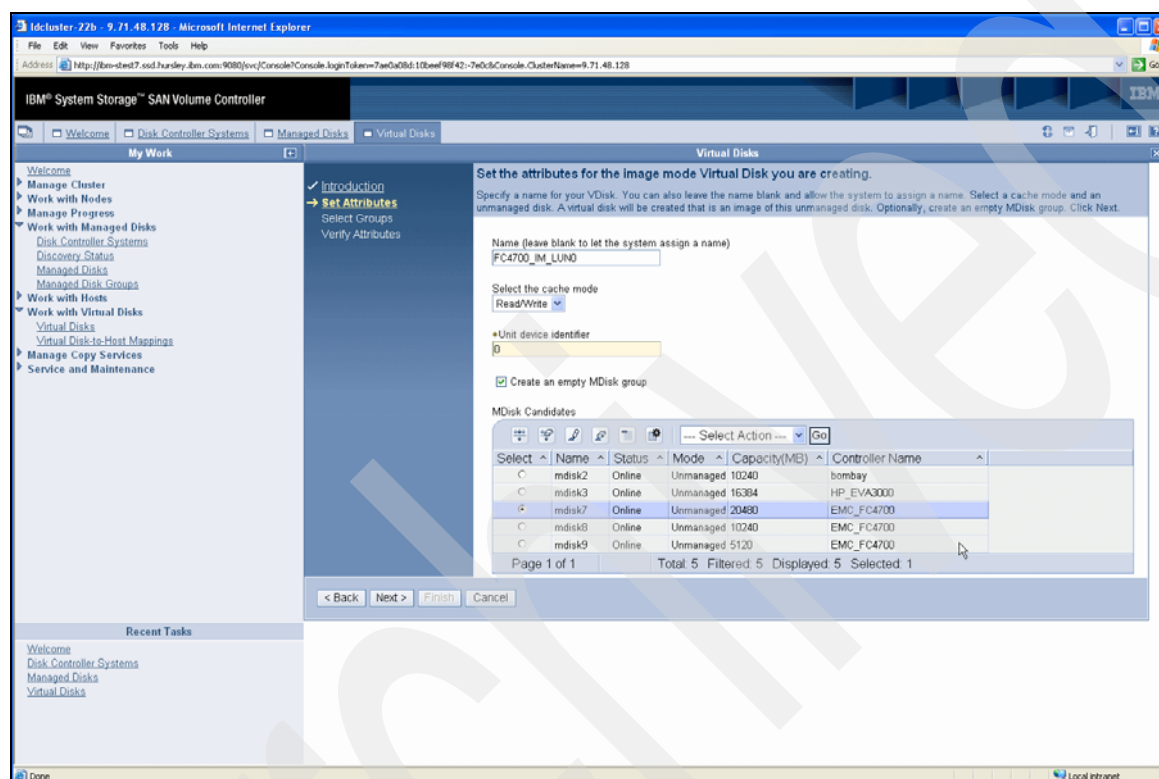


Figure 6-19 Create Image Mode VDisk set attributes

6. Figure 6-20 shows the next panel, where you must enter an MDisk group name. This panel will only appear if you checked the **Create an empty MDisk group box**. You should try to make the name memorable.

Click **Next**.




Figure 6-20 Create Image Mode VDisk name MDisk group

7. Figure 6-21 shows the panel where you define the extent size for the MDisk group. Once again, this will only appear if you are creating a new MDisk group as part of the process. Click **Next**.

Important: The extent size is not specifically an attribute associated with Image Mode VDisks. If you intend to integrate the disks to managed mode later on (to another MDG), the extent size of both managed disk groups has to be the same.

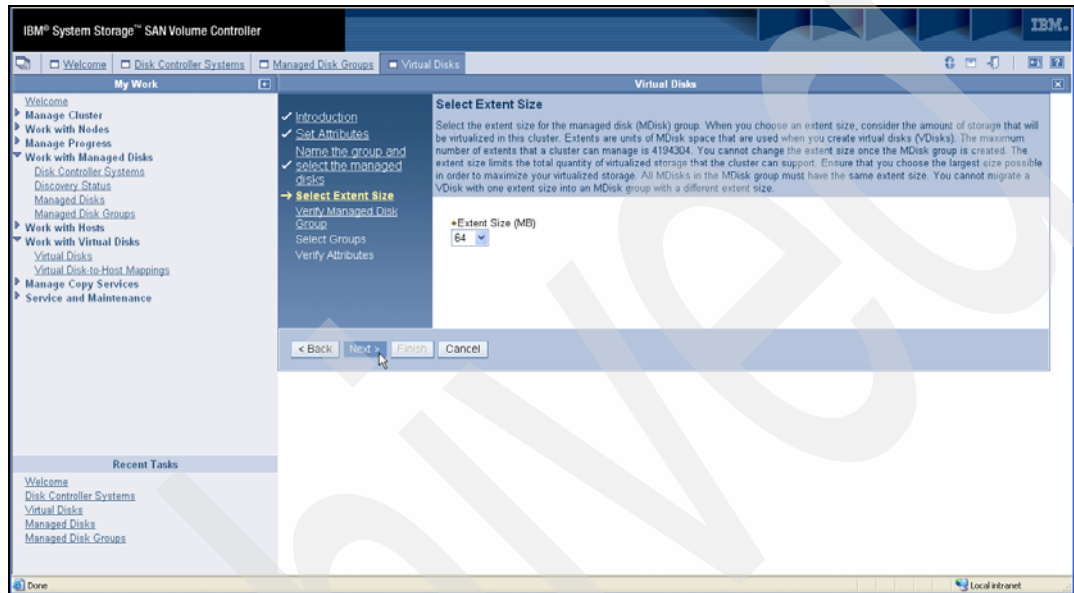


Figure 6-21 Create Image Mode VDisk select extent size

- On the next panel (Figure 6-22) verify the MDisk group attributes before it is created. If everything is correct, click **Next**.

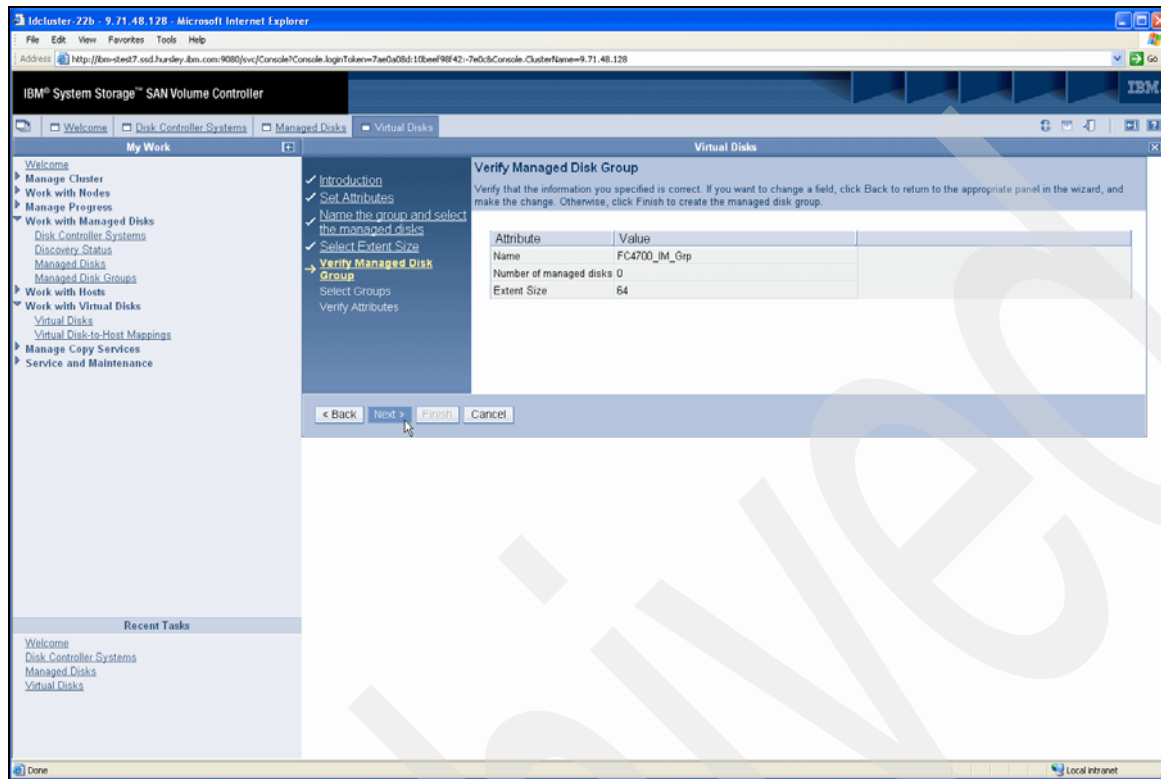


Figure 6-22 Create Image Mode VDisk verify managed disk group attributes

9. On the next panel (Figure 6-23) you have the choice of defining which I/O group the Image Mode VDisk will be part of, which node is preferred, and which MDisk group it will be assigned to. We recommend that you allow the system to choose the preferred node and I/O group. Select the MDisk group that you have created for this Image Mode VDisk. Click **Next**.

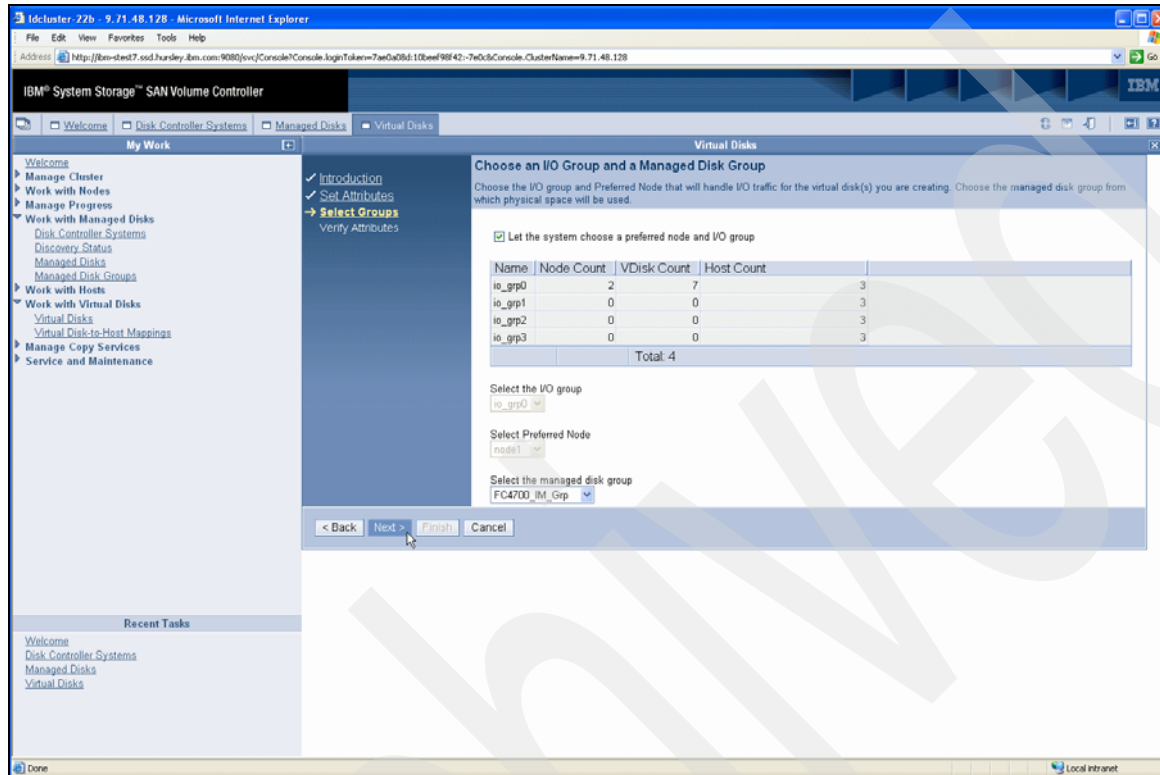


Figure 6-23 Create Image Mode VDisk choose I/O group and managed disk group

10. On the final panel (Figure 6-24) check carefully all the parameters you have entered. This is your last opportunity to change the configuration before the Image Mode VDisk is created. If everything is in order, click **Finish**.

Once you have completed the creation process, repeat the steps for each of the EMC FC4700 LUNs you want to map as Image Mode VDisks.

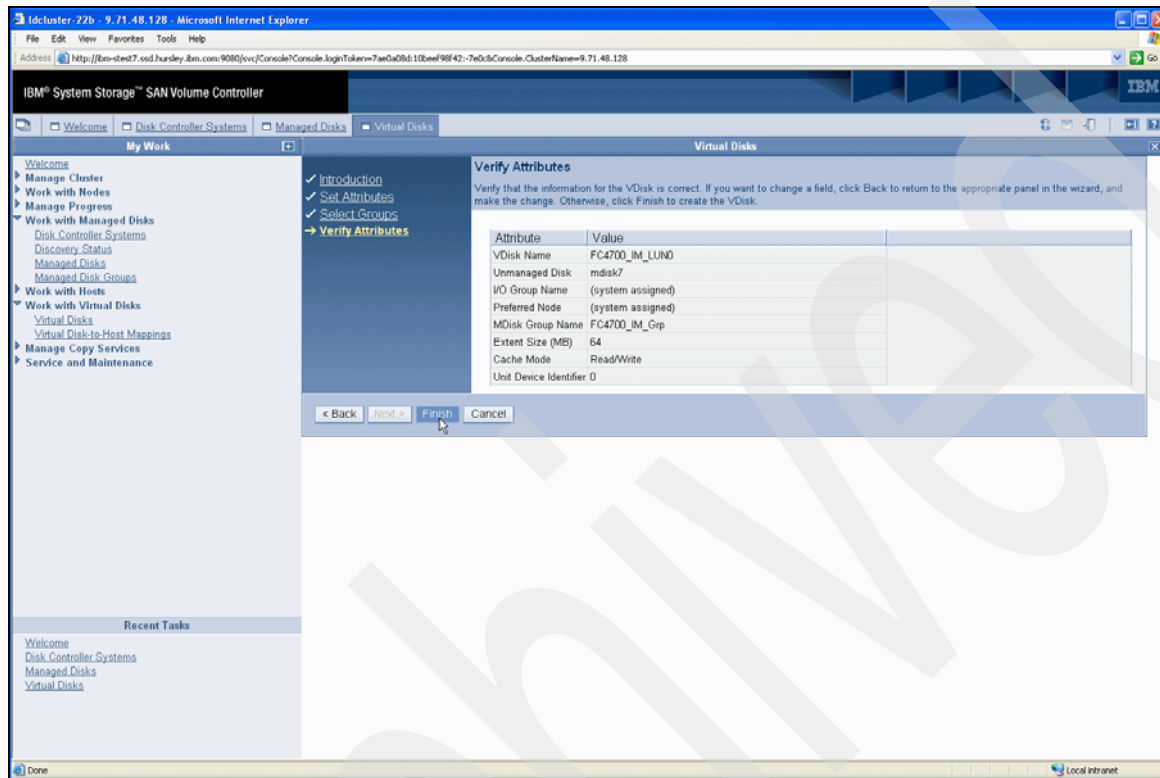


Figure 6-24 Create Image Mode VDisk verify attributes

Define a host to the SAN Volume Controller

Before the host can access any VDisks it needs to be defined to the SAN Volume Controller. Do this with the following steps:

1. Confirm that all the SAN rezoning work is complete and that the host WWPNs have been made visible to the SAN Volume Controller cluster before starting the process.
2. From the SAN Volume Controller GUI, select: **Work With Hosts** → **Hosts**.

3. Bypass the filter. The panel shown in Figure 6-25 is displayed.




Figure 6-25 Host View

4. From the drop-down menu select **Create a host**.

5. The panel shown in Figure 6-26 is displayed. Name the host and specify its attributes on this panel.
 - Enter a host name.
 - The host type can either be **HP-UX** or **Generic**.
 - Port Mask is a new feature in V4.1 of the SAN Volume Controller software. It allows you to mask which ports present devices to the defined host. Unless you specifically need to use this function, we recommend that you leave the default setting of **1111**.
 - Select the **WWPNs** of the host from the list. If you are unsure what the WWPNs are for your host, refer to “**Record the host WWPNs**” on page 126.

Note: If the WWPNs of the host you are attempting to define are not visible, this indicates that either the SAN zoning is not correct or that the host is not correctly logged into the switch.

You can optionally manually enter the WWPNs if they are not listed.

Click **OK** to complete the host creation. Figure 6-27 on page 151 shows the panel after the new host has been added.




Figure 6-26 Create a host

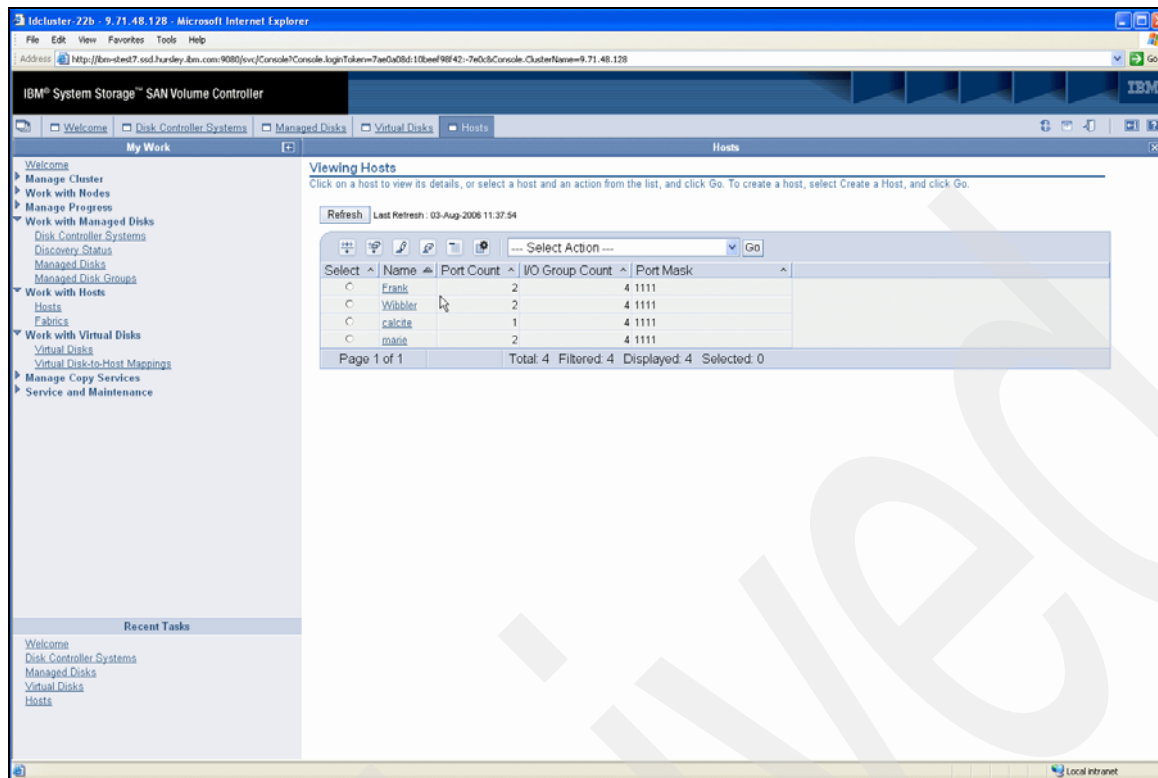


Figure 6-27 Host view after adding new host

Map the Image Mode VDisks

The final preparation procedure is to map the newly created Image Mode VDisks to the host. Use the following steps to do this:

1. Select **Work With Virtual Disks** → **Virtual Disks**.
2. Bypass the filter and continue.
3. You will be presented with a panel similar to that in Figure 6-28. Check the boxes in the **Select** column for all the VDisks you want to map to the host. In our example, this is the first three VDisks listed. Select **Map VDisks to a Host** from the drop-down menu as shown in the same figure.

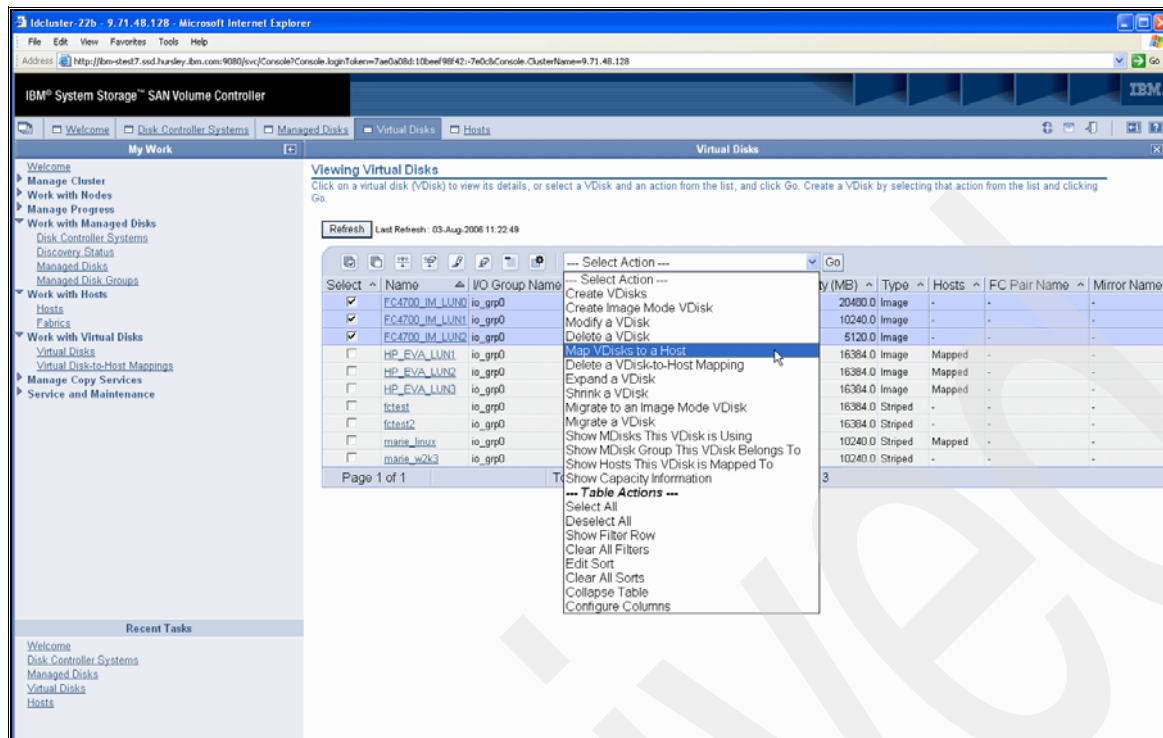


Figure 6-28 Map VDisks to host

- On the next panel (Figure 6-29) select the host that you want to map the VDisks to. You have already defined the host, so you should see the host name listed. Select the host and click **OK**.

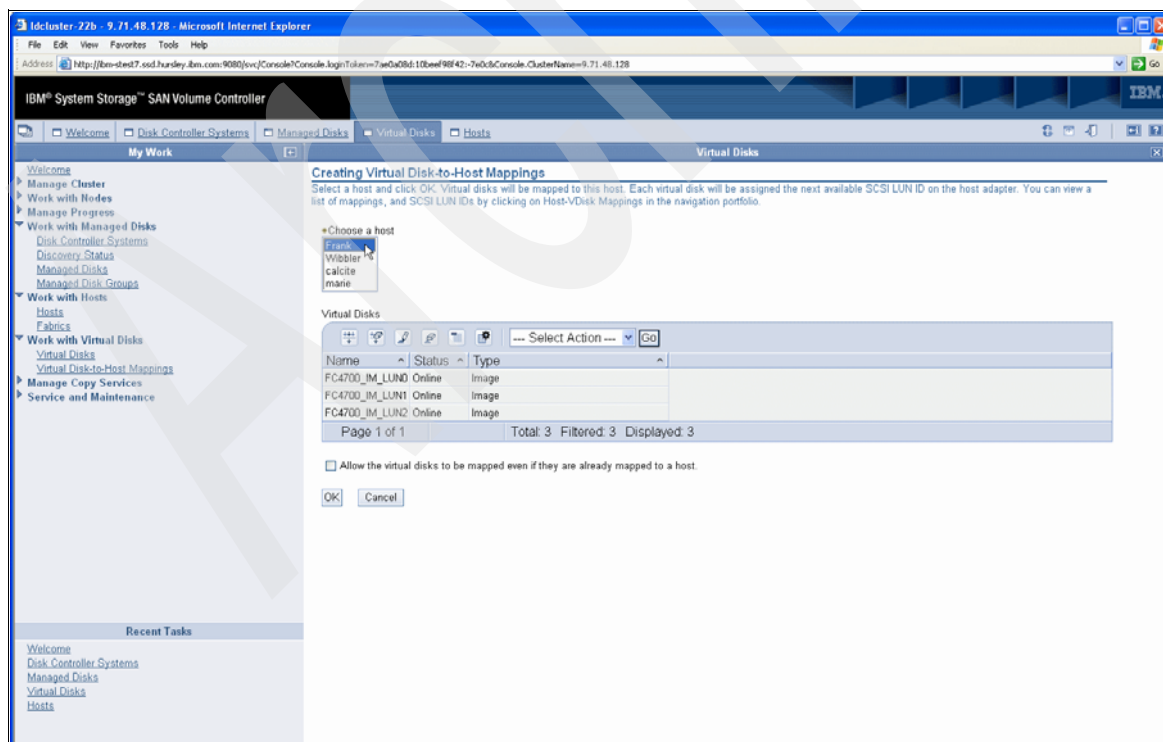


Figure 6-29 Selecting the host to map the VDisks to

5. Figure 6-30 shows the completion panel, which should indicate that your VDisks have been mapped successfully.

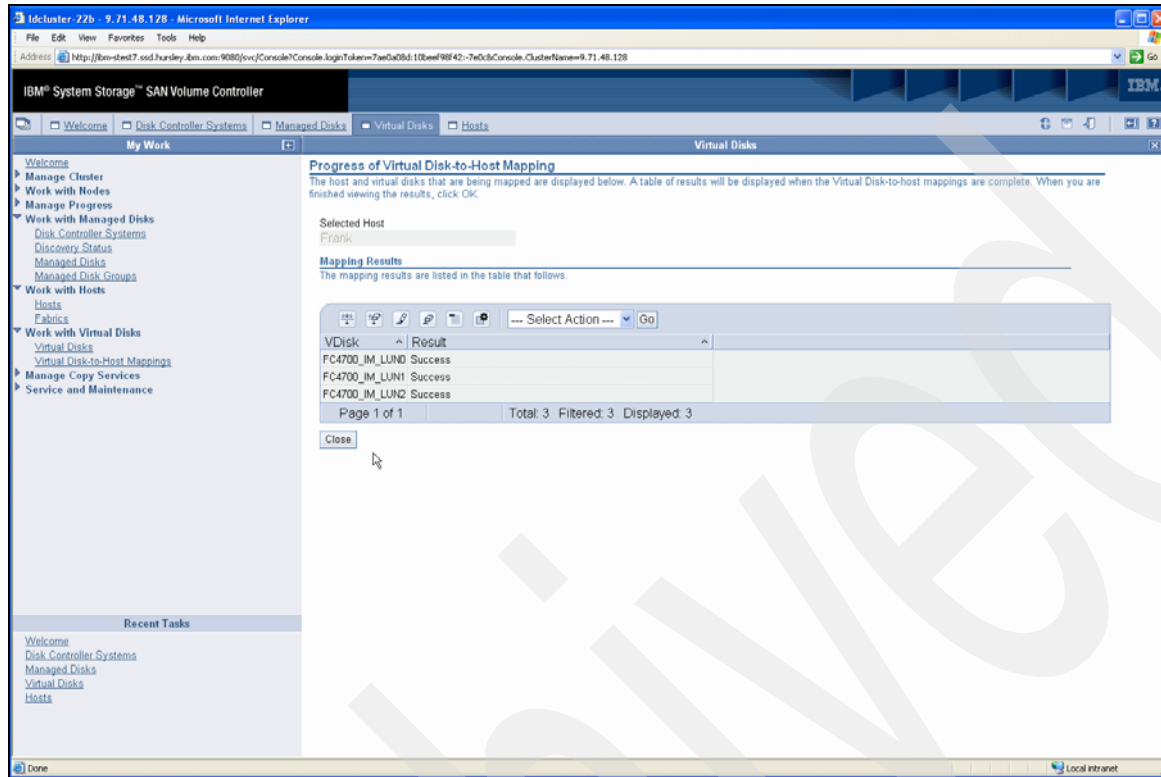


Figure 6-30 Completion panel from map VDisk to Host operation

Since image mode volumes (from the existing storage controller) cannot be used as Quorum disks, the new managed mode MDisks will also be used as the SAN Volume Controller quorum volumes.

The preparation is now complete and you are ready to proceed with the storage virtualization.

6.2.5 Migration to the SAN Volume controller

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SAN Volume Controller.

Configure the host system

Now that the SAN fabric has been rezoned and the SAN Volume Controller has been configured to present the host LUNs, you can install the IBM Subsystem Device Driver and restore the LUNs.

HP-UX

Earlier, we successfully installed the Subsystem Device Driver and the host attachment script earlier and we have now presented the LUNs through SAN Volume Controller.

Run the `ioscan -fnC disk` command. Example 6-38 shows the output of this command on our system. You should see IBM 2145 devices that are the FC4700 LUNs and indicate alternate routes to the LUNs.

Example 6-38 Checking the Controller LUNs have been identified

```
bash-2.04# ioscan -fnC disk
Class      I  H/W Path          Driver  S/W State  H/W Type  Description
=====
disk       2  0/0/2/0.6.0        sdisk   CLAIMED    DEVICE    SEAGATE ST336704LC
              /dev/dsk/c2t6d0    /dev/rdisk/c2t6d0
disk       3  0/0/2/1.6.0        sdisk   CLAIMED    DEVICE    SEAGATE ST318404LC
              /dev/dsk/c3t6d0    /dev/rdisk/c3t6d0
disk       4  0/4/0/0.33.2.0.0.0 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c36t0d0   /dev/rdisk/c36t0d0
disk       5  0/4/0/0.33.2.0.0.1 sdisk   NO_HW      svcinfo lsnode DEVICE    DGC
C4700WDR5
              /dev/dsk/c36t0d1   /dev/rdisk/c36t0d1
disk       6  0/4/0/0.33.2.0.0.2 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c36t0d2   /dev/rdisk/c36t0d2
disk      13  0/4/0/0.33.4.0.0.0 sdisk   CLAIMED    DEVICE    IBM      2145
disk      14  0/4/0/0.33.4.0.0.1 sdisk   CLAIMED    DEVICE    IBM      2145
disk      15  0/4/0/0.33.4.0.0.2 sdisk   CLAIMED    DEVICE    IBM      2145
disk      16  0/4/0/0.33.6.0.0.0 sdisk   CLAIMED    DEVICE    IBM      2145
disk      17  0/4/0/0.33.6.0.0.1 sdisk   CLAIMED    DEVICE    IBM      2145
disk      18  0/4/0/0.33.6.0.0.2 sdisk   CLAIMED    DEVICE    IBM      2145
disk       7  1/10/0/0.97.6.19.0.0 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c35t0d0   /dev/rdisk/c35t0d0
disk       8  1/10/0/0.97.6.19.0.1 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c35t0d1   /dev/rdisk/c35t0d1
disk       9  1/10/0/0.97.6.19.0.2 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c35t0d2   /dev/rdisk/c35t0d2
disk      10  1/10/0/0.97.7.19.0.0 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c37t0d0   /dev/rdisk/c37t0d0
disk      11  1/10/0/0.97.7.19.0.1 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c37t0d1   /dev/rdisk/c37t0d1
disk      12  1/10/0/0.97.7.19.0.2 sdisk   NO_HW      DEVICE    DGC      C4700WDR5
              /dev/dsk/c37t0d2   /dev/rdisk/c37t0d2
disk      19  1/10/0/0.97.9.19.0.0 sdisk   CLAIMED    DEVICE    IBM      2145
disk      20  1/10/0/0.97.9.19.0.1 sdisk   CLAIMED    DEVICE    IBM      2145
disk      21  1/10/0/0.97.9.19.0.2 sdisk   CLAIMED    DEVICE    IBM      2145
disk      22  1/10/0/0.97.11.19.0.0 sdisk   CLAIMED    DEVICE    IBM      2145
disk      23  1/10/0/0.97.11.19.0.1 sdisk   CLAIMED    DEVICE    IBM      2145
disk      24  1/10/0/0.97.11.19.0.2 sdisk   CLAIMED    DEVICE    IBM      2145
```

If you do not see the devices as expected, verify that the SAN zoning has been carried out correctly (and is enabled). Also check that the SAN Volume Controller is presenting the LUNs to the host.

Install the special files for the newly recognized paths by running the **insf** command. Example 6-39 shows the result of running this command on our system.

Example 6-39 Installing the HP-UX special files

```
bash-2.04# insf
insf: Installing special files for sdisk instance 13 address 0/4/0/0.33.4.0.0.0
insf: Installing special files for sdisk instance 14 address 0/4/0/0.33.4.0.0.1
insf: Installing special files for sdisk instance 15 address 0/4/0/0.33.4.0.0.2
insf: Installing special files for sctl instance 30 address 0/4/0/0.33.4.255.0.0.0
insf: Installing special files for sdisk instance 16 address 0/4/0/0.33.6.0.0.0.0
insf: Installing special files for sdisk instance 17 address 0/4/0/0.33.6.0.0.0.1
insf: Installing special files for sdisk instance 18 address 0/4/0/0.33.6.0.0.0.2
insf: Installing special files for sctl instance 31 address 0/4/0/0.33.6.255.0.0.0
insf: Installing special files for sdisk instance 19 address 1/10/0/0.97.9.19.0.0.0
```

```

insf: Installing special files for sdisk instance 20 address 1/10/0/0.97.9.19.0.0.1
insf: Installing special files for sdisk instance 21 address 1/10/0/0.97.9.19.0.0.2
insf: Installing special files for sctl instance 32 address 1/10/0/0.97.9.255.1.3.0
insf: Installing special files for sdisk instance 22 address 1/10/0/0.97.11.19.0.0.0
insf: Installing special files for sdisk instance 23 address 1/10/0/0.97.11.19.0.0.1
insf: Installing special files for sdisk instance 24 address 1/10/0/0.97.11.19.0.0.2
insf: Installing special files for sctl instance 33 address 1/10/0/0.97.11.255.1.3.0

```

Once the special files are installed, configure the IBM Subsystem Device Driver. Run the **cfgvpath** command to do this.

Example 6-40 shows the results of running this command on our system.

Example 6-40 Running the configure vpath command

```

bash-2.04# cfgvpath
Vpath: Configuring 3 devices (disks)

```

Re-start the system. Stop all applications using the host system at this point, if you have not done this already, by issuing the **shutdown -r 0** command.

When the system reboots, you should be able to see the virtual paths that are configured and available. To do this run the **datapath query device** command.

Example 6-41 shows the result of this command on our system.

Example 6-41 Using datapath query device

```

bash-2.04# datapath query device
Total Devices : 3
Dev#: 0 Device Name: vpath1 Type: 2145 Policy: Optimized
Serial: 6005076801AA00809800000000000010
=====
Path#      Adapter H/W Path      Hard Disk      State      Mode      Select      Error
0           1/10/0/0      c47t0d2      CLOSE     NORMAL          0          0
1           1/10/0/0      c46t0d2      CLOSE     NORMAL          0          0
2           0/4/0/0      c45t0d2      CLOSE     NORMAL          0          0
3           0/4/0/0      c44t0d2      CLOSE     NORMAL          0          0

Dev#: 1 Device Name: vpath2 Type: 2145 Policy: Optimized
Serial: 6005076801AA00809800000000000011
=====
Path#      Adapter H/W Path      Hard Disk      State      Mode      Select      Error
0           1/10/0/0      c47t0d1      CLOSE     NORMAL          0          0
1           1/10/0/0      c46t0d1      CLOSE     NORMAL          0          0
2           0/4/0/0      c45t0d1      CLOSE     NORMAL          0          0
3           0/4/0/0      c44t0d1      CLOSE     NORMAL          0          0

Dev#: 2 Device Name: vpath3 Type: 2145 Policy: Optimized
Serial: 6005076801AA00809800000000000013
=====
Path#      Adapter H/W Path      Hard Disk      State      Mode      Select      Error
0           1/10/0/0      c47t0d0      CLOSE     NORMAL          0          0
1           1/10/0/0      c46t0d0      CLOSE     NORMAL          0          0
2           0/4/0/0      c45t0d0      CLOSE     NORMAL          0          0
3           0/4/0/0      c44t0d0      CLOSE     NORMAL          0          0

```

Re-create the volume groups you exported earlier. To do this execute the commands shown in Example 6-42.

Example 6-42

```
mkdir /dev/Volume Group Name
mknod /dev/Volume Group Name/group c 64 0x0n0000
```

In this example **n** is the volume group number. This number is usually assigned in sequence each time a volume group is created. The value assigned does not matter provided that it is unique in all other volume group IDs. The command will fail if you choose an existing ID. Example 6-43 shows the commands as run on our system.

Example 6-43 Re-creating the volume group

```
ash-2.04# cd /
bash-2.04# mkdir /dev/FC4700_VG
bash-2.04# cd /dev/FC4700_VG
bash-2.04# mknod /dev/FC4700_VG/group c 64 0x020000
```

Import the volume group or groups back onto the system. To do this you need to identify the LUNs that are part of each of the volume groups to be imported. In our example LUNs 0, 1 and 2 that are presented from SAN Volume Controller make up the FC4700_VG volume group.

Run the **ioscan -kfnC disk** command.

This allows you to relate the LUN ID of the SVC presented devices to their disk device names and thus their vpath ID. Example 6-44 shows the output from this command on our system. From this you can determine that the SAN Volume Controller is presenting 3 LUNs with SCSI IDs 0, 1 and 2. There are four logical paths to each LUN. HP-UX has defined each path with a different disk identifier.

For LUN 0 the disk identifiers are:

c44t0d0, c45t0d0, c46t0d0 and c47t0d0

For LUN 1 they are:

c44t0d1, c45t0d1, c46t0d1 and c47t0d1

For LUN 2 they are:

c44t0d2, c45t0d2, c46t0d2 and c47t0d2

Looking back at Example 6-41 on page 155, you will see that the corresponding vpaths are vpath1 for LUN0, vpath 2 for LUN1, and vpath 3 for LUN2.

Example 6-44 Identification of SCSI LUN ID and disk identifier

```
bash-2.04# ioscan -kfnC disk
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	2	0/0/2/0.6.0	sdisk	CLAIMED	DEVICE	SEAGATE ST336704LC
			/dev/dsk/c2t6d0		/dev/rdisk/c2t6d0	
disk	3	0/0/2/1.6.0	sdisk	CLAIMED	DEVICE	SEAGATE ST318404LC
			/dev/dsk/c3t6d0		/dev/rdisk/c3t6d0	
disk	4	0/4/0/0.33.4.0.0.0	sdisk	CLAIMED	DEVICE	IBM 2145
			/dev/dsk/c44t0d0		/dev/rdisk/c44t0d0	

disk	5	0/4/0/0.33.4.0.0.0.1	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c44t0d1	/dev/rdisk/c44t0d1			
disk	6	0/4/0/0.33.4.0.0.0.2	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c44t0d2	/dev/rdisk/c44t0d2			
disk	10	0/4/0/0.33.6.0.0.0.0	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c45t0d0	/dev/rdisk/c45t0d0			
disk	11	0/4/0/0.33.6.0.0.0.1	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c45t0d1	/dev/rdisk/c45t0d1			
disk	12	0/4/0/0.33.6.0.0.0.2	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c45t0d2	/dev/rdisk/c45t0d2			
disk	7	1/10/0/0.97.9.19.0.0.0	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c46t0d0	/dev/rdisk/c46t0d0			
disk	8	1/10/0/0.97.9.19.0.0.1	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c46t0d1	/dev/rdisk/c46t0d1			
disk	9	1/10/0/0.97.9.19.0.0.2	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c46t0d2	/dev/rdisk/c46t0d2			
disk	13	1/10/0/0.97.11.19.0.0.0	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c47t0d0	/dev/rdisk/c47t0d0			
disk	14	1/10/0/0.97.11.19.0.0.1	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c47t0d1	/dev/rdisk/c47t0d1			
disk	15	1/10/0/0.97.11.19.0.0.2	sdisk	CLAIMED	DEVICE	IBM	2145
			/dev/dsk/c47t0d2	/dev/rdisk/c47t0d2			

Now that you understand the mapping, you can re-create the volume group using the vpath definitions. For our example we know that all three SCSI LUNs make up the FC4700_VG volume group. To re-create the volume group we used the command shown in Example 6-45.

Example 6-45 Recreating the volume

```
vgimport -p -v -m mapfile /dev/volume group name /dev/dsk/vpathA
/dev/dsk/vpathB...../dev/dsk/vpathD
```

The mapfile is created when you export a volume group. By default it is called *mapfile* and is placed in the directory where the command is run. You need to include the vpaths for all the disks that make up the volume group. The **-p** option allows you to test the command to ensure it will run successfully. If you get a successful completion, you need to re-run the command without the **-p** option.

Once you have successfully imported the volume group, you need to reactivate the volume group and remount the filesystems. Example 6-46 shows this.

Example 6-46 Reactivating the volume

```
vgchange -a y /dev/volume group name
mount filesystem(s)
```

It is also advisable to create a configuration backup of the volume group using this command:

```
vgcfgbackup /dev/volume group name
```

Example 6-47 shows the commands run on our system.

Example 6-47 Example commands

```
bash-2.04# vgimport -p -v -m mapfile /dev/FC4700_VG /dev/dsk/vpath1
/dev/dsk/vpath2 /dev/dsk/vpath3
Beginning the import process on Volume Group "/dev/FC4700_VG".
```

```

Logical volume "/dev/FC4700_VG/FC4700_LV" has been successfully created
with lv number 1.
Volume group "/dev/FC4700_VG" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating
the volume group.
bash-2.04# vgimport -v -m mapfile /dev/FC4700_VG /dev/dsk/vpath1 /dev/dsk/vpath2
/dev/dsk/vpath3
Beginning the import process on Volume Group "/dev/FC4700_VG".
Logical volume "/dev/FC4700_VG/FC4700_LV" has been successfully created
with lv number 1.
Volume group "/dev/FC4700_VG" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating
the volume group.

bash-2.04# vgchange -a y /dev/FC4700_VG
Activated volume group
Volume group "/dev/FC4700_VG" has been successfully changed.

bash-2.04# mount /EMC
bash-2.04# cd /EMC

bash-2.04# vgcfgbackup /dev/FC4700_VG

```

Repeat these steps for each of the volume groups and file systems you have integrated. On completion you should now have access to your original data.

AIX

After you have successfully installed the SDD, run the **cfgmgr** command to discover the SVC disks.

The **cfgmgr** command produces no feedback; therefore, use the **lsdev -Cc disk** to display the disk devices.

You should see a SAN Volume Controller Device for each path to the SVC VDisks and a Data Path Optimizer Pseudo Device Driver for each logical device.

Example 6-48 shows the output of these commands.

Example 6-48 Discovery of SDD devices on AIX

```

[root@wibbler]/
>cfgmgr

[root@wibbler]/
>lsdev -Cc disk
hdisk0 Available 10-60-00-0,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-60-00-1,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 10-68-01 SAN Volume Controller Device
hdisk3 Available 10-68-01 SAN Volume Controller Device
hdisk4 Available 10-68-01 SAN Volume Controller Device
hdisk5 Available 10-68-01 SAN Volume Controller Device
hdisk6 Available 10-68-01 SAN Volume Controller Device
hdisk7 Available 10-68-01 SAN Volume Controller Device
hdisk8 Available 10-68-01 SAN Volume Controller Device

```



```

hdisk9 Available 10-68-01 SAN Volume Controller Device
vpath0 Available Data Path Optimizer Pseudo Device Driver
vpath1 Available Data Path Optimizer Pseudo Device Driver
[root@wibbler]/
>

```

If you do not see the devices as expected, check that the SDD has been installed correctly. If necessary, check that the SAN zoning is correct and enabled. Also check that the SVC is presenting the LUNs to the host.

AIX writes a Physical Volume Identifier (PVID) to each LUN, so this means that the newly presented disks should be recognized as part of the *ClarVG* volume group.

Example 6-49 shows the output from the **lspv** command. You should see that the *vpaths* are assigned to their respective volume groups and you will also notice that they are not yet active.

Example 6-49 List physical volumes

```

[root@wibbler]/
>lspv
hdisk0          0051a28abb619770      rootvg      active
hdisk1          0051a28a1f16b48c      None
hdisk2          none                  None
hdisk3          none                  None
hdisk4          none                  None
hdisk5          none                  None
hdisk6          none                  None
hdisk7          none                  None
hdisk8          none                  None
hdisk9          none                  None
vpath0          0051a28a471f2ffb      ClarVG
vpath1          0051a28a472a5224      ClarVG
[root@wibbler]/
>

```

You now need to reactivate the volume group. To do this issue the **varyonvg ClarVG** command.

Example 6-50 shows the output from this command.

Example 6-50 List physical volumes output following varying on the volume groups

```

[root@wibbler]/
>varyonvg ClarVG
[root@wibbler]/
>lspv
hdisk0          0051a28abb619770      rootvg      active
hdisk1          0051a28a1f16b48c      None
hdisk2          none                  None
hdisk3          none                  None
hdisk4          none                  None
hdisk5          none                  None
hdisk6          none                  None
hdisk7          none                  None
hdisk8          none                  None

```

hdisk9	none	None	
vpath0	0051a28a471f2ffb	ClarVG	active
vpath1	0051a28a472a5224	ClarVG	active

```
[root@wibbler]/  
>
```

You can now remount the logical volumes.

In our example we have one mount point, /emc, so the command is **mount /emc/**.

The data stored on the EMC Clariion is now available to the host system again.



Integrating SVC with HP

In this chapter we show how to integrate the IBM System Storage SAN Volume Controller in an existing environment with a Hewlett-Packard (HP) storage controller.

7.1 Integrating the SVC in a SAN with a Hewlett-Packard EVA

In this section we demonstrate the integration of the IBM System Storage SAN Volume Controller (SVC) in an existing environment with a Hewlett-Packard EVA5000 storage controller. The existing storage volumes on the Hewlett-Packard EVA5000 storage controller (HP EVA) will be virtualized using the SAN Volume Controller. Note that the same steps can be used for any HP EVA storage controller because they share a common management interface.

7.2 Scenario description

In this scenario we start with an HP EVA5000 storage controller and a dual redundant SAN connected AIX and Microsoft® Windows host. We then introduce the SAN Volume Controller to virtualize the back-end storage.

The systems described in this section include:

- ▶ Hewlett-Packard EVA5000 storage controller (HP EVA)
- ▶ IBM b-type fibre channel switches
- ▶ Host system running AIX/Microsoft Windows

Figure 7-1 shows the relationships between the systems used in this scenario.




Figure 7-1 SAN Volume Controller integration example with an existing HP EVA controller

7.3 Summary of integration process

At a high level, the steps to integrate the SVC in the existing production environment are the following:

1. Preparation
 - a. General
 - i. Review SAN levels of software, firmware, and supported hardware using information from the configuration guide.
 - ii. Determine SVC worldwide port names and Host worldwide port names.
 - b. SAN
 - i. Create host zone with SVC (enable now or later).
 - ii. Create EVA controller zone with SVC (enable now or later).

Note: In this example the SVC Cluster is already set up and installed. For details on how to set up and install an SVC cluster consult *IBM System Storage SAN Volume Controller*, SG24-6423.

<http://www.redbooks.ibm.com/abstracts/sg246423.html?Open>

- c. EVA Controller
 - i. Identify the SCSI LUN IDs and Unique IDs (UID) that are to be integrated under SAN Volume Controller.
2. Perform the integration
 - a. Host (dependent on host type)
 - i. Stop applications.
 - ii. Remove existing multipath driver if not supported by SAN Volume Controller.
 - iii. Install SDD or equivalent driver.
 - iv. Shut down machine or unmount file systems and vary off volume groups.
 - b. SAN
 - i. Enable new zone for host-SAN Volume Controller and EVA controller-SAN Volume Controller.
 - ii. Disable EVA controller-host zone.
 - c. EVA Controller
 - i. Unmap LUNs from host.
 - ii. Remap LUNs to SVC.
 - d. SAN Volume Controller
 - i. Discover EVA controller LUNs.
 - ii. Create MDisk Group for MDisk in image mode.
 - iii. Create Image Mode VDisks.
 - iv. Create host with associated WWPNs.
 - v. Map VDisks to host.
 - e. Host (dependent on host)
 - i. Reboot or unmount.

- ii. Scan for new devices.
- iii. Vary on volume group.
- iv. Remount file systems.

Note: We recommend that you review 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization example.

7.4 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SAN Volume Controller is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment, it may require coordination of effort between various support teams and an understanding of several technical domains.

Confirm support for your environment

This section provides information about configuring the SAN Volume Controller in an existing environment with a Hewlett-Packard EVA storage controller.

Prior to integrating the SAN Volume Controller, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SAN Volume Controller supports all the current models of the HP EVA range.

The following models of HP EVA are supported by the SAN Volume Controller:

- ▶ EVA3000
- ▶ EVA5000
- ▶ EVA4000
- ▶ EVA6000
- ▶ EVA8000

This list is current at the time of writing. See the following Web site for the latest Hewlett-Packard EVA supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Concurrent host connection to the HP EVA and the SVC

There are restrictions for sharing an HP EVA between a host and the SVC.

Sharing ports

The HP EVA can be shared between a host and a SAN Volume Controller under the following conditions:

- ▶ The same host cannot be connected to both a SAN Volume Controller and an HP EVA at the same time because the multipath drivers (HP Secure Path and SDD) cannot coexist on the host.
- ▶ Logical units (LUNs) cannot be shared between a host and a SAN Volume Controller.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each of the two HSV controllers within the HP EVA in order to have no single point of failure.

Quorum disk support

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the HP EVA controllers as quorum disks.

Further information

It is strongly recommended that you refer to the SAN Volume Controller V4.1 Configuration Guide for further information regarding integrating the SAN Volume Controller with an existing storage controller:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

Preparing the SAN fabric

You need to prepare two additional zones, one to allow the SVC to be visible to the host whose data is to be virtualized (the BLUE zone in Figure 7-1 on page 162), and the second to enable the SAN Volume Controller to be visible to the HP EVA (the RED1 zone in Figure 7-1). Depending upon your environment, you may choose to simply create these zones at this time and only enable them later when the host applications have been stopped.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Note: It is assumed that the BLACK zone, which is the SVC node to SVC node zone, was created when creating the SVC cluster and therefore already exists in the SAN.

Record the SAN Volume Controller WWPNs

You will need the SAN Volume Controller worldwide port name (WWPN) for each of the SAN Volume Controller nodes. This will enable you to configure the HP EVA to make LUNs visible to SAN Volume Controller. There are several methods of collecting this information. You can use the SAN switch interface to identify the WWPNs provided you know which ports the SAN Volume Controller nodes are connected to. Alternatively, you can use the SAN Volume controller GUI or CLI. The following example shows how to use the GUI.

1. Open a browser window and connect to the SAN Volume Controller GUI.
2. Select **Work With Nodes** → **Nodes**. Figure 7-2 shows the resulting GUI panel on our system.




Figure 7-2 Selecting the node to determine the associated WWPNs

3. Select each node in turn by clicking the node name link shown in Figure 7-2; the panel shown in Figure 7-3 is then presented, showing general data about that node.




Figure 7-3 Selecting the Port Information for the node

4. Select the Port information link (Figure 7-3) to show the four WWPNs associated with that node. You will see a panel similar to Figure 7-4. Record the WWPNs to enable you to map the HP EVA LUNs to the SVC later on.




Figure 7-4 List of WWPNs for the node

Repeat these steps for each of the SAN Volume Controller nodes in your SAN Volume Controller cluster.

Record the host WWPN

There are several methods of identifying the WWPNs of the host involved in the integration. The simplest method is to capture the information through the SAN switch. Alternatively, you can access the information through the host using the host bus adapter BIOS. On some platforms you can capture the information in a configuration file.

Preparing the HP EVA

You need to identify the LUNs to be integrated and capture their SCSI LUN identifier and their Unique IDs. This is important when you present the LUNs through SAN Volume Controller because some host platforms will expect to see the LUNs presented with the same SCSI LUN identifier. The Unique ID helps to identify the LUNs, having presented them to SAN Volume Controller. If you have a large number of LUNs, this is the best method to confirm which they are.

Perform the following steps to prepare the HP EVA:

1. Access the administrative interface on the HP EVA. This can be done through a Web browser using the HP Storage Works Command View EVA software.
2. Log on and select the EVA system you want to integrate.
3. Expand the Virtual Disks folder in the left menu to display the LUNs currently presented. Identify those currently presented to the host to be integrated. For each of these, expand the name and open the **Vdisk Active Member Properties**.
4. Click the **General** tab. The LUN Unique ID is displayed in the identification section. Register this ID to later be able to identify this LUN from the SVC.
5. Select the **Presentation** tab. The LUN presentation properties are displayed. Register the SCSI LUN ID together with the UUID in the previous step (so you can later map it in the exact same way to the host from the SVC).

7.5 Begin integration

In this section we describe the steps to integrate to SVC managed storage.

Host

The SAN Volume Controller is provided with a freely available multi-path driver. This is known as the IBM Subsystem Device Driver (SDD). The driver manages the resiliency and connectivity between the host and the virtualized storage. Multi-path drivers from other vendors are also tested and supported. For the most up to date support matrix, refer to the SAN Volume Controller supported hardware list:

<http://www.ibm.com/storage/support/2145>

In the examples in this section we have to remove the existing multi-path drivers, which are specific to the HP storage controller, and replace them with the IBM Subsystem Device Driver. The procedure is similar for other device drivers as well, but we recommend that you read the relevant documentation available for a specific driver.

Microsoft Windows

In our environment it was necessary to uninstall the existing HP EVA MPIO device driver before installing the IBM SDD driver.

Note: If SecurePath for Windows is the driver used in your environment (for VCS 3.x and earlier versions) it is uninstalled in the same way.

Stop all applications using the host.

Select **Add or Remove Programs** from the Control Panel and click **Remove** to start the uninstall process. Figure 7-5 shows the removal of the HP MPIO driver from our Windows host.

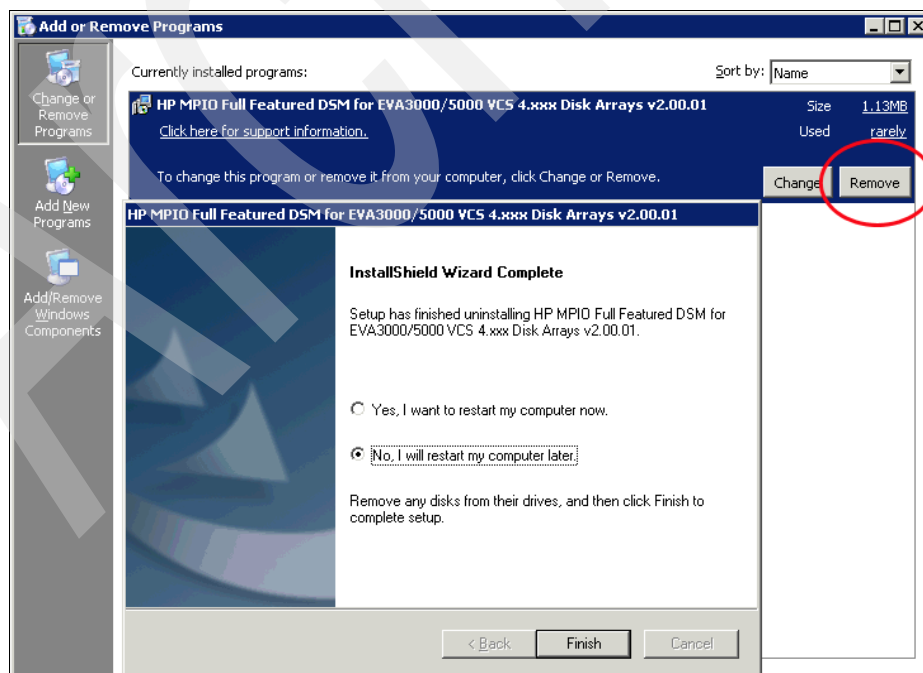


Figure 7-5 Uninstalling HP MPIO for Windows

Be sure to follow the existing storage controller vendor's recommendations when removing the existing multi-path device driver.

Install the new multipath driver, IBM Subsystem Device Driver. The latest version of the filesets can be downloaded from the Web. Refer to the SAN Volume Controller recommended software levels document, which will direct you to the relevant download site:

<http://www.ibm.com/storage/support/2145>

It is advisable to configure the server to prevent the automatic starting of any applications on the next boot. This is to prevent any problems with applications trying to access the SAN volumes visible to it before you complete the configuration.

Shut down the host

The host should remain briefly offline while you reconfigure the SAN. While the host is offline you remove direct access to the existing LUN and enable access to the virtualized LUN through the SAN Volume Controller.

Note: Although some parts of the storage virtualization process may be conducted with the host system online, we choose to shut down the host to minimize risk during the change.

AIX

On this platform, all the changes to the host system can be made without a full shutdown. In our example, the AIX system had the HP EVA MPIO driver installed.

Use the **lshsv** command to determine what disks are being presented to the host system.

You will use this and other specific commands to the HP MPIO device driver in the following steps. You might need to add the directory where this command is stored to the path list. You can do this by entering **PATH=\$PATH:/opt/hphsv/bin**.

Example 7-1 shows the output of this command on our system.

Example 7-1 List disks presented by HP EVA to AIX host

```
[wibbler]>lshsv
```

Disk	UUID	NWWN	Status	Location	Description
hdisk2	600508B400013A....	50001fe150027610	Available	10-68-01	HP HSV111...
hdisk3	600508B400013A....	50001fe150027610	Available	10-68-01	HP HSV111...
hdisk4	600508B400013A....	50001fe150027610	Available	10-68-01	HP HSV111...

The output of this command has been edited for legibility. The UUID corresponds to the one recorded earlier when you checked information about the EVA.

Use the **lspv** command to determine what volume group the disks are part of.

Example 7-2 shows the output of this command on our system.

Example 7-2 List physical volume to volume group mapping

```
[wibbler]>lspv
```

hdisk0	0051a28abb619770	rootvg	active
hdisk1	0051a28a1f16b48c	None	
hdisk2	0051a28a8b6742bf	HP_Volume	active
hdisk3	0051a28a8b67446c	HP_Volume	active
hdisk4	0051a28a8b674603	HP_Volume	active

On the system we are using you will see that the HP EVA is presenting three disks: hdisk2, hdisk3, and hdisk4, which are part of the HP_Volume volume group.

You now need to stop the applications that access this volume group. In your environment there may be more than one volume group involved, in which case you should execute the following set of commands with that in mind.

Use the `lsvg -o |lsvg -i -l` command to determine what mount points the volume group has.

Example 7-3 shows the logical volumes associated with the host system volume groups and the relevant mount points.

Example 7-3 Mapping between volume group and logical volumes

```
[wibbler]>lsvg -o |lsvg -i -l
```

HP_Volume:						
LV NAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
fslv00	jfs2	2	2	1	closed/syncd	N/A
loglv00	jfs2log	1	1	1	closed/syncd	N/A
fslv01	jfs2	10	10	1	closed/syncd	/HP
rootvg:						
LV NAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
hd5	boot	1	1	1	closed/syncd	N/A
hd6	paging	8	8	1	open/syncd	N/A
hd8	jfslog	1	1	1	open/syncd	N/A
hd4	jfs	1	1	1	open/syncd	/
hd2	jfs	18	18	1	open/syncd	/usr
hd9var	jfs	1	1	1	open/syncd	/var
hd3	jfs	33	33	1	open/syncd	/tmp
hd1	jfs	1	1	1	open/syncd	/home
hd10opt	jfs	4	4	1	open/syncd	/opt
lv00	jfs	4	4	1	open/syncd	/usr/vice/cache

You will see that the HP_Volume has three logical volumes associated. fslv01 is mounted to /HP.

Unmount the file system and vary off the volume group in readiness for integrating it under the SVC environment.

To do this run `/usr/sbin/umount Filesystem` followed by `varyoffvg Volume Group`.

Example 7-4 shows this set of commands being run on our system.

Example 7-4 Unmount of file system and vary off of volume group

```
[wibbler]>/usr/sbin/umount /HP
[wibbler]>
[wibbler]>/usr/sbin/varyoffvg HP_Volume
[wibbler]>
```

Remove the existing LUN definitions and HP MPIO device driver. The HP MPIO installation provides a set of commands that enable this.

The `rmhsv` command will remove the disk definitions. Example 7-5 shows the command output.

Example 7-5 Removing the HP EVA related hdisks

```
[wibbler]>rmhsv  
hdisk2 deleted  
hdisk3 deleted  
hdisk4 deleted
```

If for some reason the command does not delete the hdisks, run `rmdev -d1 hdiskn`, where *n* represents the corresponding number of each hdisk presented by the HP EVA.

You can now remove the HP MPIO driver by running the following command:

```
installp -u devices.fcp.disk.HP.hsv.mpio.rte
```

If this command is unsuccessful, it may be that you are running a different version of the MPIO driver for AIX. Refer to the HP-supplied documentation on the required set of commands for removing the driver.

Install the IBM Subsystem Device Driver and host attachment scripts onto the host system. The latest version of these filesets can be downloaded from the Web. Refer to the SAN Volume Controller recommended software levels document, which will direct you to the relevant download site:

<http://www.ibm.com/storage/support/2145>

If you are uncertain which level of AIX you have installed, the command `oslevel` will return this. This will determine which filesets you should download. The system in our example is running AIX v5.3, so the two filesets to be downloaded are:

```
devices.fcp.disk.ibm.rte.tar  
devices.sdd.53.rte.tar
```

Place the filesets in the same directory. We recommend creating a new directory called IBM.

Untar the files by running the commands shown in Example 7-6.

Example 7-6

```
tar -xvf devices.fcp.disk.ibm.rte.tar  
tar -xvf devices.sdd.53.rte.tar
```

Note that you will need to change the filenames to suit the filesets for your version of AIX.

This will generate two unpacked filesets. Example 7-7 shows the output from listing the directory in which they were untarred.

Example 7-7 List of files following untar command

```
[wibbler]/IBM>ls -l  
total 6168  
-rwxrwx---1 155350 449628 31744 09 May 23:59 devices.fcp.disk.ibm.rte  
-rw-r-----1 root system 40960 21 Jul 16:15 devices.fcp.disk.ibm.rte.tar  
-rw-r-----1 root system 1536000 09 May 20:02 devices.sdd.53.rte  
-rw-r-----1 root system 1546240 21 Jul 16:15 devices.sdd.53.rte.tar
```

You can now install these filesets by using `smit/smitty` or through the command line. Example 7-8 shows how we set the install source directory as IBM and then installed the filesets in that directory.

Example 7-8 Commands to install through the command line

```
[wibbler]
>cd /IBM
[wibbler]/IBM
>inutoc $PWD
[wibbler]/IBM
>installp -acd . -einstallp.log ALL
```

Check the output to ensure that the filesets were installed successfully. The output will show the cause of any failures.

Now you we have removed the HP drivers and installed the IBM Subsystem Device Driver, you can reconfigure the SAN fabric and map the existing HP EVA LUNs to the SAN Volume Controller.

Configure the SAN fabric

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SAN Volume Controller. (Figure 7-1 on page 162 shows the zones. You should remove the GREEN zone).

You should now enable the RED1 and BLUE zones (Figure 7-1) so that the SAN Volume Controller can access the HP EVA LUNs you are about to remap and so that the host under integration can access the SAN Volume Controller.

Refer to Appendix B, “SAN zoning changes” on page 257 for a description of the required SAN zones for integrating the SAN Volume Controller in an existing environment. This appendix also provides examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Configure the HP EVA

Unmap the LUNs that are currently presented directly to the host and remap them to the SAN Volume Controller nodes.

Un-presenting the LUNs from the host

Log on to the StorageWorks Command View EVA and select the EVA you are integrating. On the left side, select the Vdisk you are going to integrate and choose to un-present the Vdisk from the host.

When prompted, confirm the removal of the host to LUN access. Repeat this step for each LUN that is being unmapped.

Adding an SVC host group to the HP EVA

Create a new host group representing your SVC, to which you will use to present the LUNs you are about to integrate, using the following steps:

1. Select the **Hosts** folder in the left pane and click the **Create folder** button in the Host Folder Properties pane.
2. When prompted, enter a logical name for the host group, for example *SVC1*, and set the host mode. The SVC requires that the host mode is set to Windows.
3. Click **Finish**.

Create the host (SAN Volume Controller) that will be on the host group using these steps:

1. In the left pane select the host group you created.

2. Enter a logical name for the first SVC node, for example *SVC1_N1*, and the WWPN for the first port.
3. Click **Add host** to create the host.
4. Repeat steps 1 through 3 for each SAN Volume Controller node, as well as adding the WWPN for the other ports for each SVC node.

Presenting the LUNs to the SVC

You are now ready to map the LUNs to the SAN Volume Controller. When you present the LUNs to the SVC you must make sure that LUNs are exactly the same as the ones you un-presented in the previous procedure.

1. In the left pane select the LUN you want to present to the SVC and click the **Presentation** tab.
2. Select the SVC host group and click **Present Vdisk**. The LUN is now accessible from the SVC.
3. Repeat steps 1 and 2 for each of the LUNs that is to be integrated with the SVC.

Tip: If you are not familiar with configuring the HP EVA storage controller, refer to the installation and configuration documentation available on line at:

http://h20000.www2.hp.com/bizsupport/TechSupport/Product.jsp?locale=en_US&prodTypeId=12169&prodCatId=304617&prodSubCatId=304697

Configure the SAN Volume Controller

Confirm that the SAN Volume Controller can see the storage controllers and the LUNs that are to be virtualized.

Access the SAN Volume Controller GUI and select **Work With Managed Disks** → **Disk Controller Systems**.

Figure 7-6 shows the output of this command on our system. You will see that a new controller has been identified, *controller1*. We recommend that you rename the controller to something more memorable.

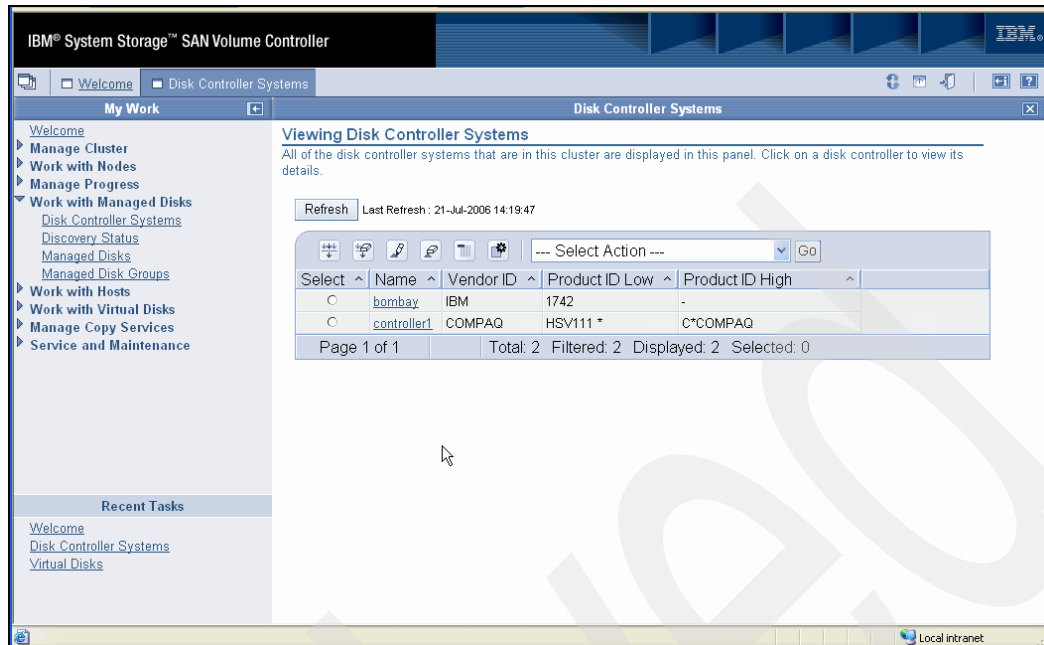


Figure 7-6 List of attached storage controllers

To do this check the **Select** button to the left of the controller name, click the drop-down menu button, and select **Rename a disk controller system**. Figure 7-7 shows this selection.

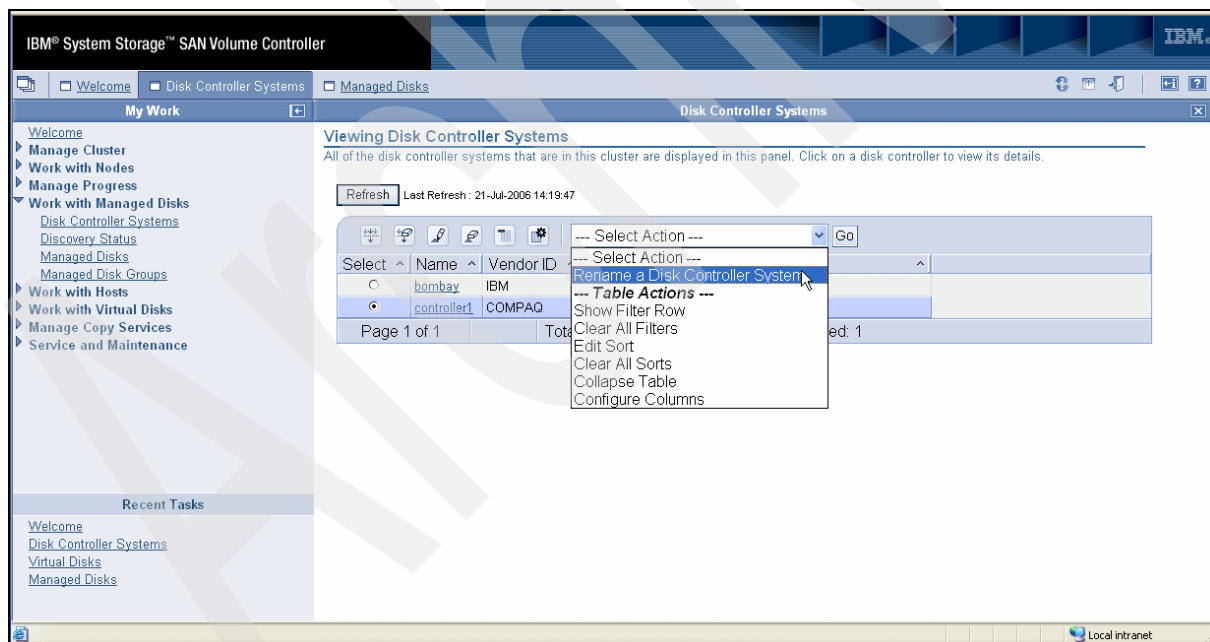


Figure 7-7 Renaming a storage controller

Check what LUNs the HP EVA storage controller is presenting. To do this select **Work With Managed Disks** → **Managed Disks**.

A filter option page is displayed. Click the **Bypass Filter box**.

Figure 7-8 shows the Managed Disk view.

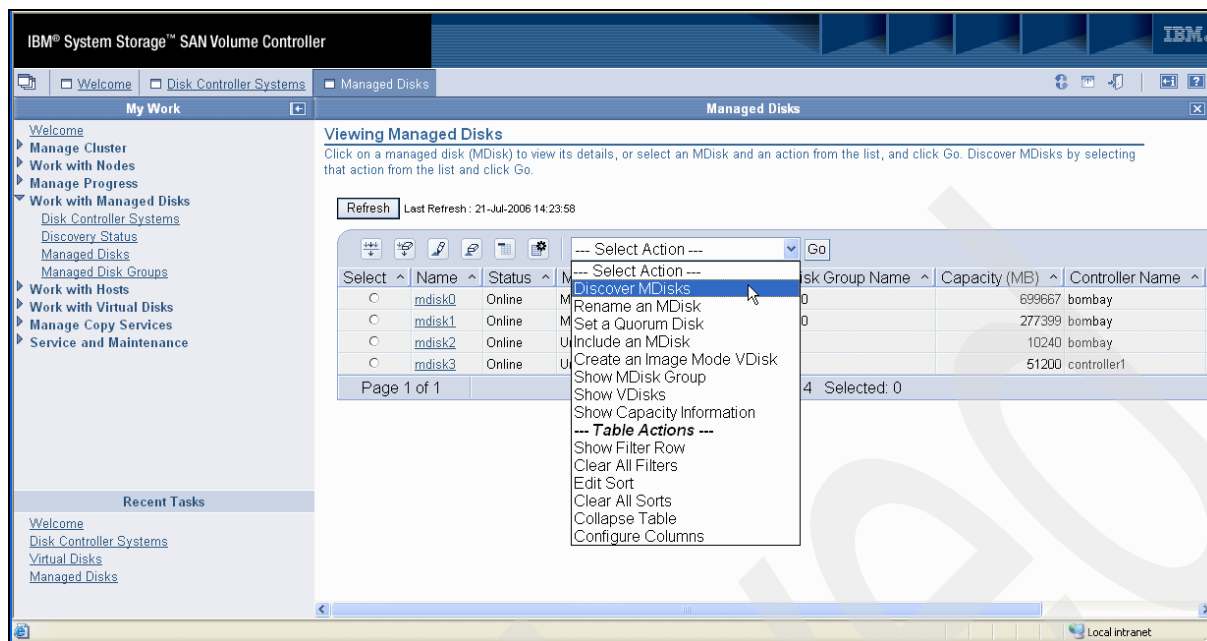


Figure 7-8 Managed Disk view

Since the HP EVA controller has been recently added, it may be that the LUNs it is presenting to the SAN Volume Controller have not yet been discovered. If this is the case, run the **Discover MDisks** command to scan for new devices. Figure 7-9 shows the result for this command.

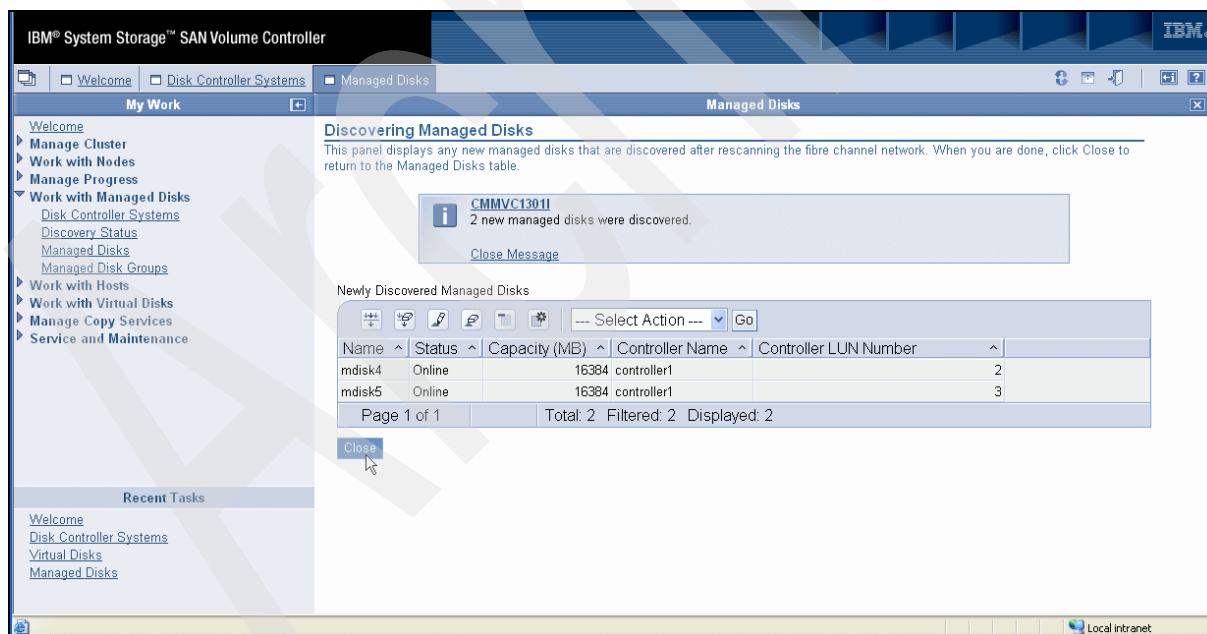


Figure 7-9 Result of discover MDisks command

Note that on our system two new MDisks were found. Click the **Close** button to return to the view panel.

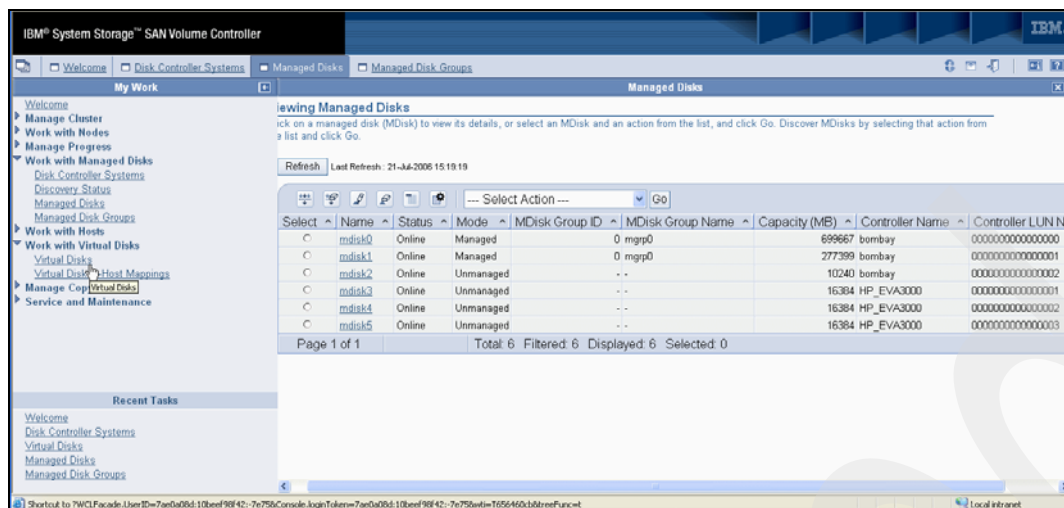


Figure 7-10 Managed Disk View after discovery

Figure 7-10 shows three disks (MDisk 3, 4, and 5) being presented from the EVA controller. At this time they are in the *Unmanaged* state. Create a managed disk group and define a VDisk in image mode for each MDisk to preserve the existing data on the LUNs. In this case, create three image mode VDisks. Unlike the SVC CLI, the GUI enables you to create a managed disk group as part of the process of creating an Image Mode VDisk.

Note: It is recommended that you create a Managed Disk Group (MDG) for the MDisks in image mode.

Important: Do not put an MDisk into the MDG at this time, while there is existing data on the MDisk. It must only become a member of an MDG when you create an image mode VDisk which points to this MDisk. Therefore, the MDisks remain as unmanaged until the VDisks in image mode are defined.

Earlier in the process you collected information about the SCSI LUN ID and the Unique ID (UID) to differentiate between the LUNs presented by the EVA. In Figure 7-10 you can determine the SCSI LUN ID that the EVA controller has assigned to each of the MDisks. However, you cannot determine the UID that each MDisk has. If you need to do this to identify each one, then you must use the Command Line Interface. The command is `svcinfo lsmdisk name`, where *name* is the MDisk in question. The UID is clearly identified. It is important to know the LUN numbers for each of the MDisks you are about to make into Image Mode Disks because you cannot access this information through the GUI once you begin the process.

Use the following steps to create the image mode VDisk:

1. Select **Work With Virtual Disks** → **Virtual Disks**.
2. Bypass the filter.
3. From the drop-down menu, select **Create Image Mode VDisk**. Figure 7-11 shows selecting this option.

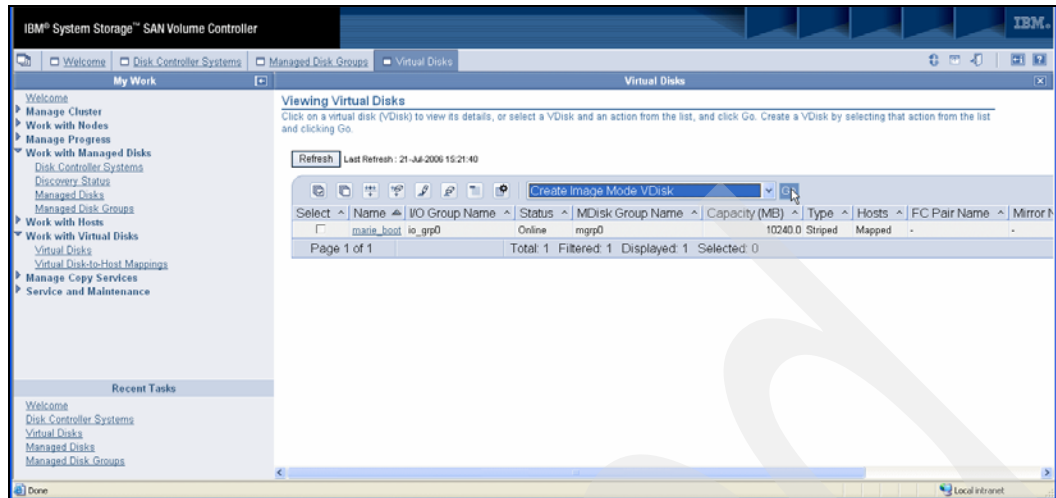


Figure 7-11 Create Image Mode VDisk

4. You are presented with a wizard that will take you through the whole process. Figure 7-12 shows the overview of the steps.

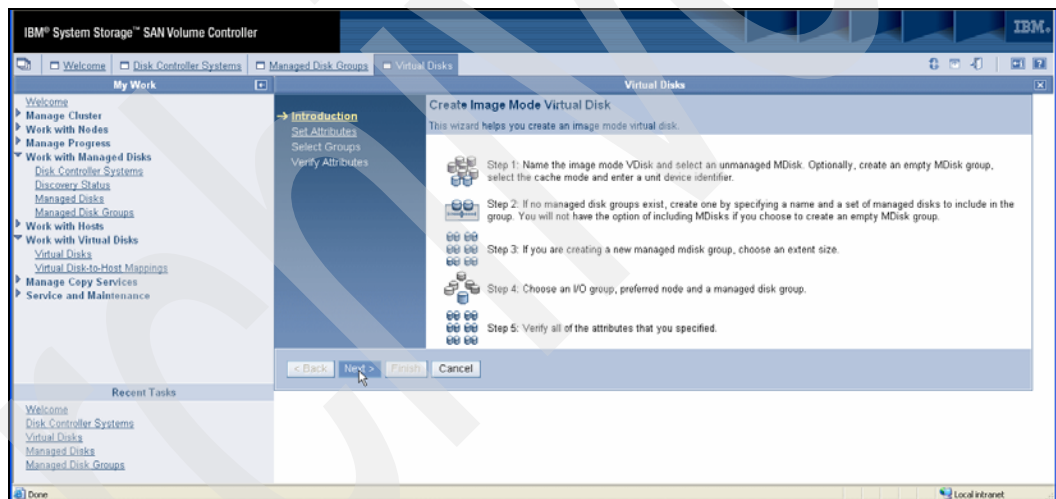


Figure 7-12 Create Image Mode Virtual Disk introduction

5. Figure 7-13 shows the second panel in this wizard. Specify the attributes for the VDisk you are creating in this panel:

- Give the Image Mode VDisk a name. We recommend that you give it a memorable name to identify it more easily in the future.
- Select the cache mode as **Read/Write**.
- The ID is the **SCSI LUN ID** that SVC will use to present this to the host. You should retain the IDs that were used by the HP EVA when presenting the LUNs directly to the host.
- If this is the first image mode VDisk you have created from the HP EVA, check the box **Create an empty MDisk group**. If you need to repeat these steps for further EVA LUNs you can choose not to check this box.
- Select the HP EVA LUN that you want to create an image mode VDisk from.

Click **Next**.

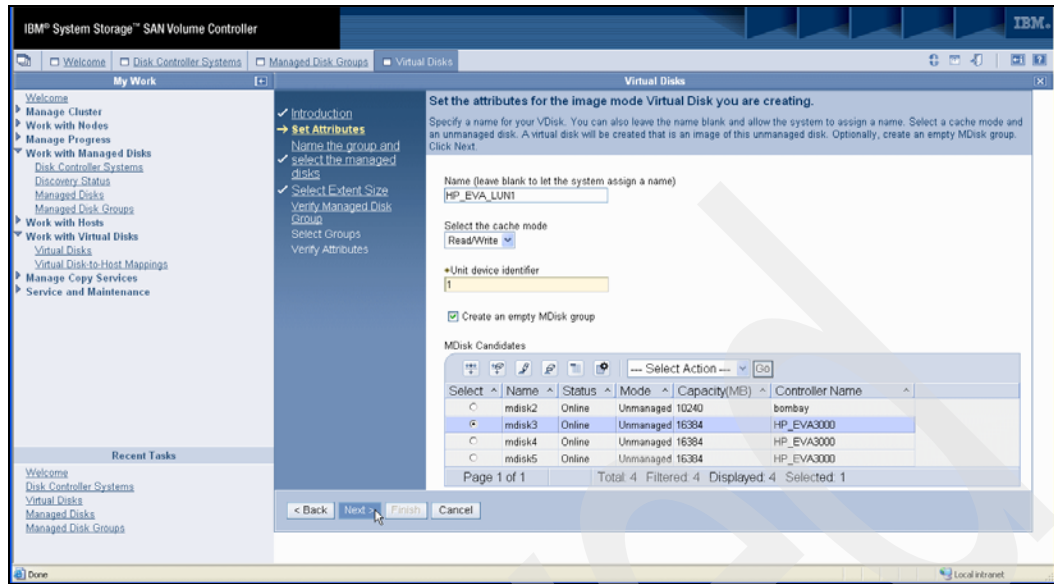


Figure 7-13 Create Image Mode VDisk set attributes

- Figure 7-14 shows the next panel, where you must enter an MDisk group name. This panel will only appear if you checked the **Create an empty MDisk group** box. We recommend that you use a descriptive name; in this example we entered *HP_EVA-ImgMod*. Click **Next**.

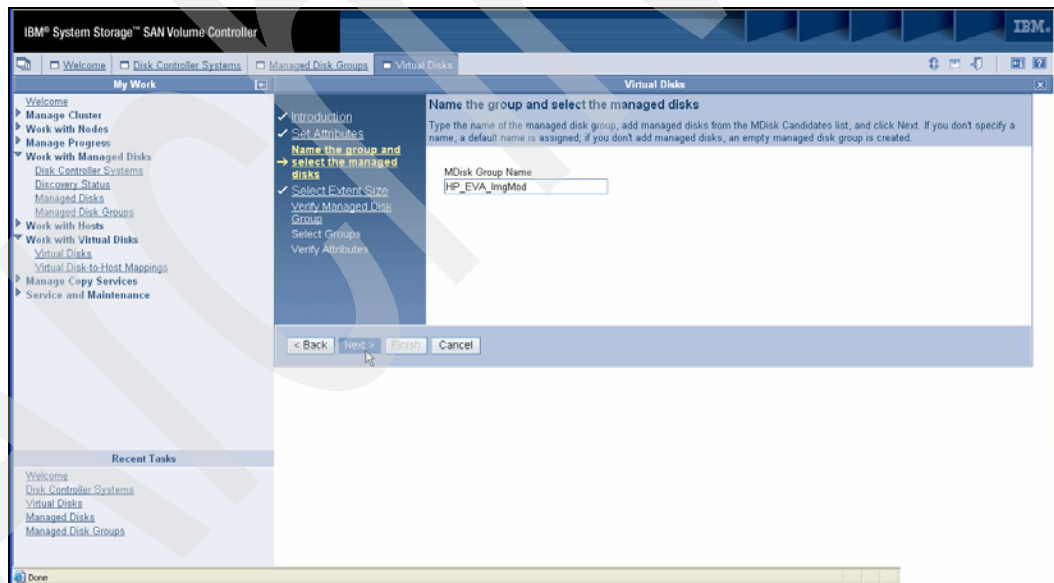


Figure 7-14 Create Image Mode VDisk name MDisk group

- On the next panel (Figure 7-15), define the extent size for the MDisk group. Once again this will only appear if you are creating a new MDisk group as part of the process. Click **Next**.

Note: The extent size is not specifically an attribute associated with Image Mode VDisks. The extent size is an attribute of the MDG and it is recommended that you use the same extent size for all MDGs across the SVC cluster because the extent size of both MDGs has to be the same if you want to migrate a VDisk from one MDG to another.



Figure 7-15 Create Image Mode VDisk select extent size

8. On the next panel (Figure 7-16) verify the MDisk group attributes before it is created. If everything is correct, click **Next**.

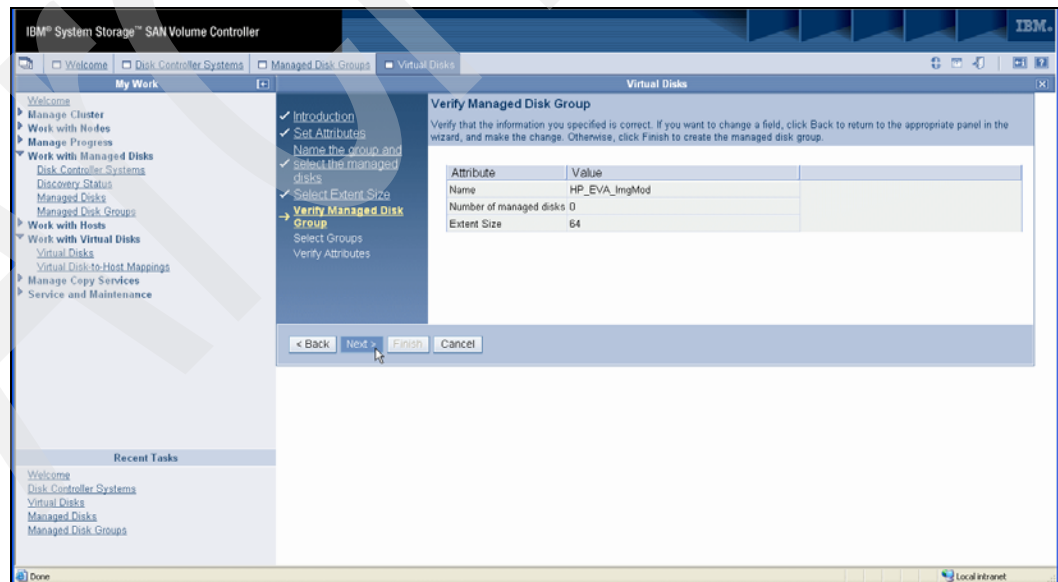


Figure 7-16 Create Image Mode VDisk verify managed disk group attributes

9. On the next panel (Figure 7-17) you can choose the following attributes:

- Which I/O group the Image Mode VDisk will be part of
- Which node is preferred
- To which MDisk group it will be assigned

We recommend that you allow the system to choose the preferred node and I/O group. Select the MDisk group that you have created for this image mode VDisk.

After making your entries, click **Next**.




Figure 7-17 Create Image Mode VDisk choose I/O group and managed disk group

10. On the final panel (Figure 7-18), carefully check all the parameters you have entered. This is your last opportunity to change the configuration before the Image Mode VDisk is created. If everything is in order, click **Finish**.

Once you have completed the creation process, repeat the steps for each HP EVA LUN that you want to define as an Image Mode VDisk.

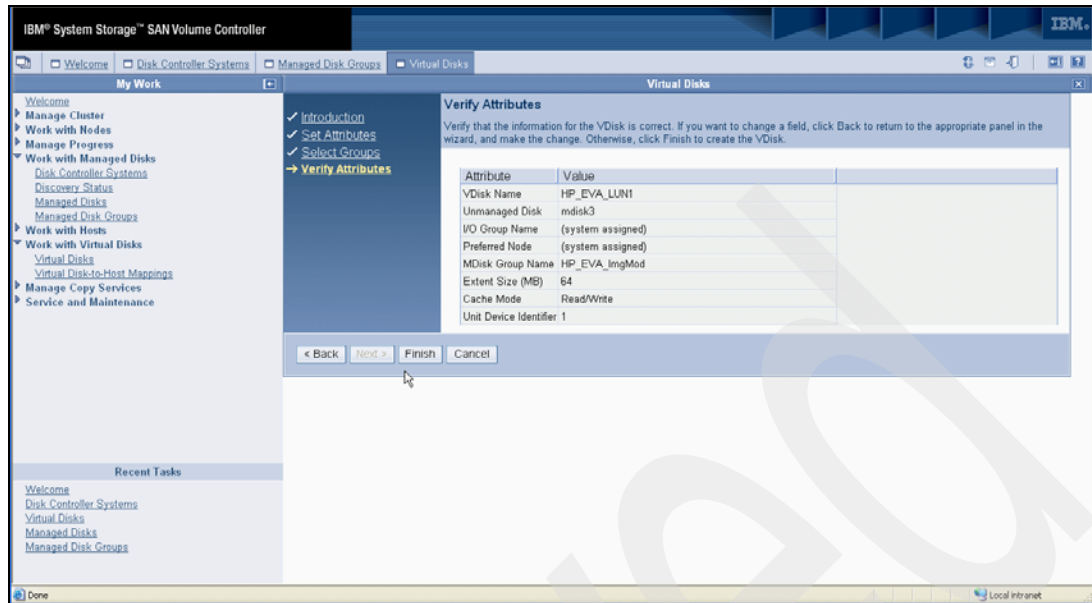


Figure 7-18 Create Image Mode VDisk verify attributes

Define a host to the SAN Volume Controller

Before the host can access any VDisks it must be defined to the SAN Volume Controller. Confirm that all the SAN rezoning work is complete and that the host WWPNs have been made visible to the SAN Volume Controller cluster before starting the process. Then use the following steps to define the host to the SVC:

1. From the SAN Volume Controller GUI, select **Work With Hosts** → **Hosts**.
2. Bypass the filter. The panel shown in Figure 7-19 is displayed.




Figure 7-19 Host View

3. From the drop-down menu select **Create a host**.
4. On the next panel (Figure 7-20), specify the following attributes of the host:
 - A host name.
 - The host type, either HP-UX or Generic.
 - Port mask is a new feature in V4.1 of the SAN Volume Controller software. It allows you to mask which ports present devices to the defined host. Unless you specifically need to use this function, we recommend that you leave the default setting of 1111.
 - The WWPNs of the host. Move the appropriate WWPNs from the Available Ports list to the Selected Ports list. If you are unsure what the WWPNs are for your host, refer to “Record the host WWPN” on page 167.
 - You can manually enter the WWPNs if they are not listed.

Note: If the WWPNs of the host you are attempting to define are not visible then this indicates that either the SAN zoning is not correct or that the host is not correctly logged into the switch.




Figure 7-20 Create a host

5. Click **OK** to complete the host creation. Figure 7-21 shows the host view panel after the new host has been added.

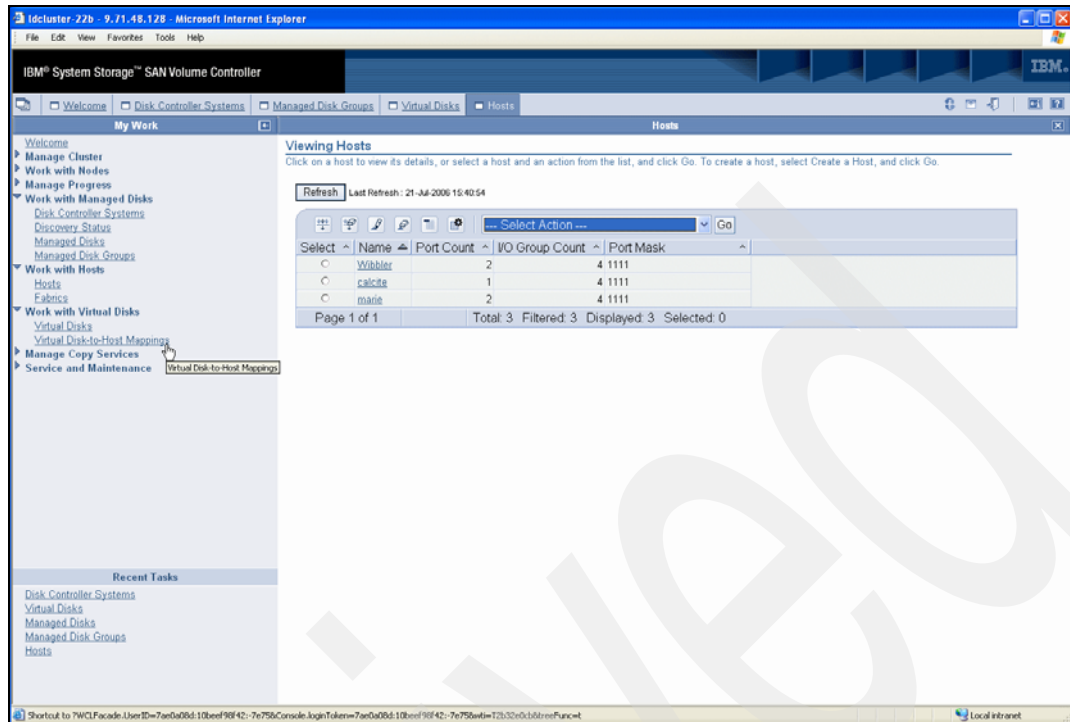


Figure 7-21 Host view after adding new host

Map the Image Mode VDisks

The final procedure is mapping the newly created Image Mode VDisks to the host. Use the following steps to do this:

1. Select **Work With Virtual Disks** → **Virtual Disks**
2. Bypass the filter and continue.
3. On the next panel (Figure 7-22) check the boxes in the **Select** column for all the VDisks you want to map to the host. In our example, this is the first three VDisks listed. Select **Map VDisks to a Host** from the drop-down menu as shown in Figure 7-22.

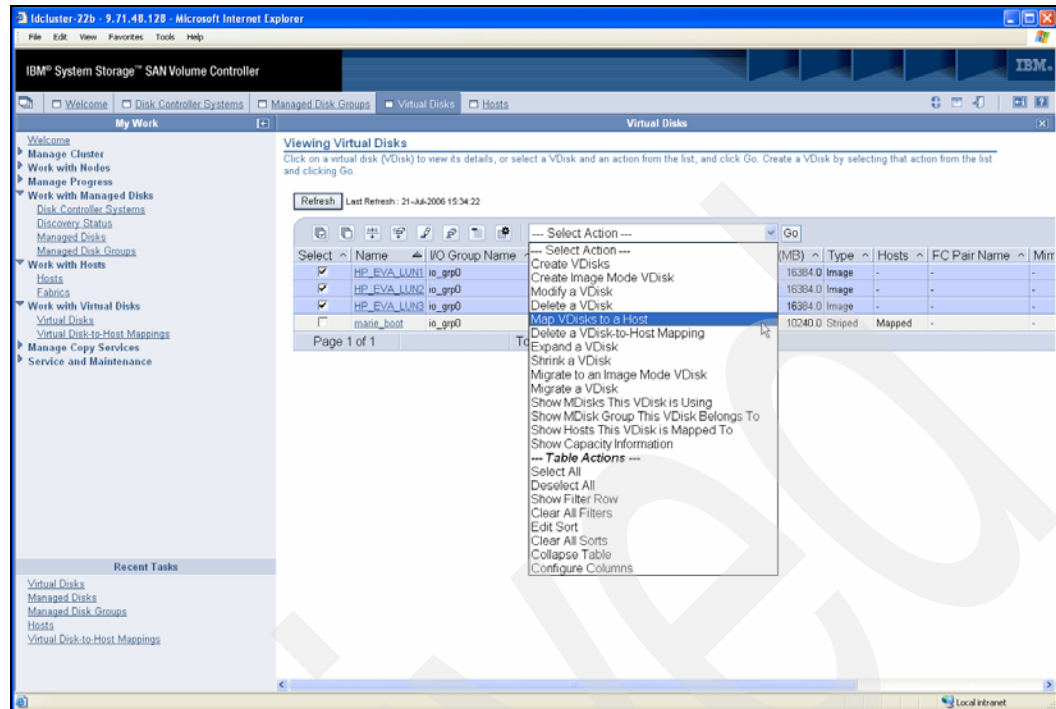


Figure 7-22 Map VDisks to host

4. On the next panel (Figure 7-23), select the host you want to map the VDisks to. Since you defined the host previously, you should see the host name listed. Select the host and click OK.

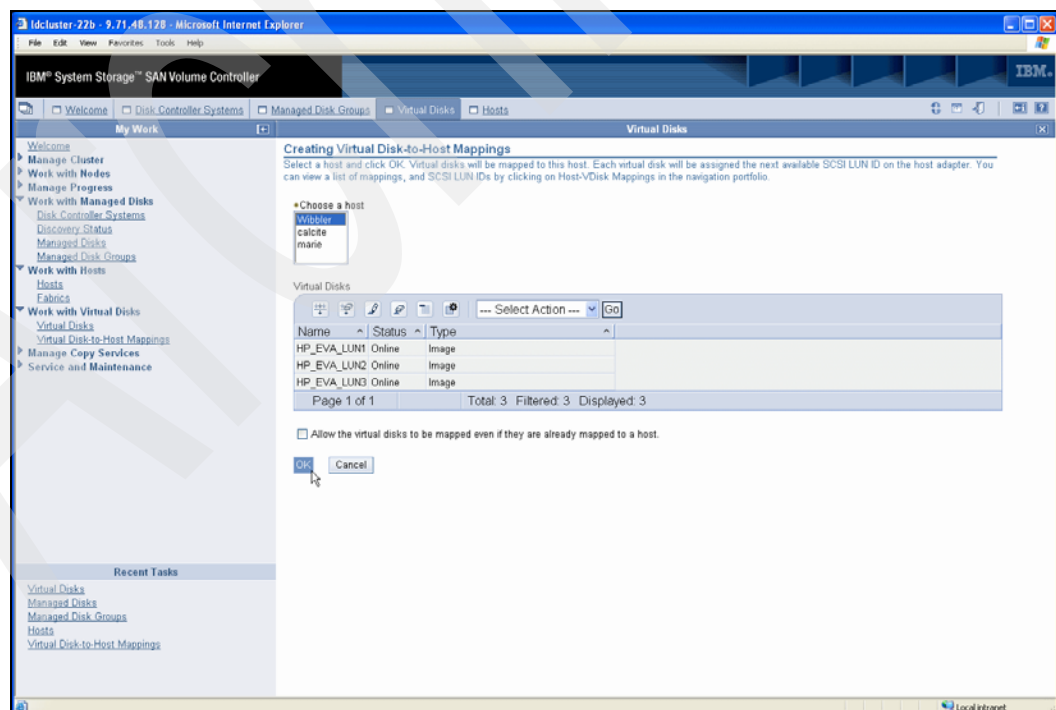


Figure 7-23 Selecting the host to map the VDisks to

5. The next panel (Figure 7-24) indicates that your VDisks have been mapped successfully.




Figure 7-24 Completion panel from map VDisk to Host operation

Since image mode volumes (from the existing storage controller) cannot be used as Quorum disks, the new managed mode MDiskS will be used as the SAN Volume Controller quorum volumes.

The preparation is now complete and you are ready to proceed with the storage virtualization.

7.6 Integration into the SVC environment

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SAN Volume Controller.

Configure the host system

Once the SAN fabric has been rezoned and the SAN Volume Controller has been configured to present the host LUNs, you can install the IBM Subsystem Device Driver and restore the LUNs.

Microsoft Windows

Restart the host system. When the host reboots it should now see the same LUNs as when directly attached, but now they are being presented through SAN Volume Controller.

Use the IBM SDD command line to check that the SAN volumes are available.

Example 7-9 shows a sample output of this command on a Windows host.

Example 7-9 Checking SAN devices with SDD

```
C:\Program Files\IBM\SDDDSM>datapath query device
```

Total Devices : 1

```
DEV#: 0  DEVICE NAME: Disk1 Part0  TYPE: 2145          POLICY: OPTIMIZED
SERIAL: 600507680191002AE0000000000000980
```

=====						
Path#	Adapter/Hard Disk	State	Mode	Select	Errors	
0	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	0	0	
1	Scsi Port2 Bus0/Disk1 Part0	OPEN	NORMAL	879	0	
2	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	0	0	
3	Scsi Port3 Bus0/Disk1 Part0	OPEN	NORMAL	877	0	

The data on the SAN disk is now available on the host for application and user access.

Note: The storage LUN has now been virtualized, but still resides on the original storage controller. Access to the LUNs is managed through the SAN Volume Controller

AIX

Once you have successfully installed the Subsystem Device Driver and the host attachment script, and have presented the LUNs through San Volume Controller, run the **cfgmgr** command.

When this has completed run the **lsdev -Cc disk** command.

Example 7-10 shows the output of this command. You should see vpath devices that are the HP EVA LUNs and hdisk devices that indicate all the available routes to those vpaths.

Example 7-10 Checking the Controller LUNs have been identified and configured

```
[root@wibbler]/IBM>lsdev -Cc disk
hdisk0 Available 10-60-00-0,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-60-00-1,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 10-68-01     SAN Volume Controller Device
hdisk3 Available 10-68-01     SAN Volume Controller Device
hdisk4 Available 10-68-01     SAN Volume Controller Device
hdisk5 Available 10-68-01     SAN Volume Controller Device
hdisk6 Available 10-68-01     SAN Volume Controller Device
hdisk7 Available 10-68-01     SAN Volume Controller Device
hdisk8 Available 20-58-01     SAN Volume Controller Device
hdisk9 Available 20-58-01     SAN Volume Controller Device
hdisk10 Available 20-58-01    SAN Volume Controller Device
hdisk11 Available 20-58-01    SAN Volume Controller Device
hdisk12 Available 20-58-01    SAN Volume Controller Device
hdisk13 Available 20-58-01    SAN Volume Controller Device
vpath0 Available             Data Path Optimizer Pseudo Device Driver
vpath1 Available             Data Path Optimizer Pseudo Device Driver
vpath2 Available             Data Path Optimizer Pseudo Device Driver
```

If you do not see the devices as expected, check that the Subsystem Device Driver and host attachment script filesets have been installed correctly. If necessary, check that the SAN zoning is correct and enabled. Also check that the SAN Volume Controller is presenting the LUNs to the host.

AIX writes a Physical Volume Identifier (PVID) to each LUN, which means that the newly presented disks should be recognized as part of the *HP_Volume* volume group.

Example 7-11 shows the output from the **lspv** command. The vpaths are assigned to their respective volume groups and you will also notice that they are not yet active.

Example 7-11 List physical volumes

[wibbler]>lspv			
hdisk0	0051a28abb619770	rootvg	active
hdisk1	0051a28a1f16b48c	None	
hdisk2	none	None	
hdisk3	none	None	
hdisk4	none	None	
hdisk5	none	None	
hdisk6	none	None	
hdisk7	none	None	
hdisk8	none	None	
hdisk9	none	None	
hdisk10	none	None	
hdisk11	none	None	
hdisk12	none	None	
hdisk13	none	None	
vpath0	0051a28a8b6742bf	HP_Volume	
vpath1	0051a28a8b67446c	HP_Volume	
vpath2	0051a28a8b674603	HP_Volume	

To reactivate the volume group, issue the **varyonvg HP_Volume** command.

Example 7-12 shows the output from **lspv** following the **varyon** command; it has been edited for legibility.

Example 7-12 List physical volumes output following varying on the volume groups

[wibbler]>lspv			
hdisk0	0051a28abb619770	rootvg	active
.....
vpath0	0051a28a8b6742bf	HP_Volume	active
vpath1	0051a28a8b67446c	HP_Volume	active
vpath2	0051a28a8b674603	HP_Volume	active

You can now remount the logical volumes.

In the example we have one mount point, /HP, so the command is **mount /HP**.

The data stored on the HP EVA is now available to the host system again, and we have completed the integration with the SVC.

Archived

Integrating SVC with NetApp

In this chapter we describe several implementations of an SVC in an existing environment with Network Appliance™ (NetApp) products:

- ▶ Integration with a NetApp V-Series NAS gateway. The existing storage volumes on the back-end SAN storage controller will be virtualized using the SAN Volume Controller and migrated to a new storage controller.
- ▶ Integration with a NetApp storage controller. In this scenario, the existing iSCSI storage volumes on the NetApp storage controller are converted to Fibre Channel (FCP) volumes and virtualized using the SAN Volume Controller. The virtualized storage volumes will then be migrated to an IBM DS4800 storage controller.

8.1 Integration with NetApp V-Series

In this scenario we start with a NetApp V-Series NAS Gateway system that is connected to an HDS AMS500 disk storage system. We then introduce a SAN Volume Controller to virtualize the back-end storage. The virtualized storage is then migrated to a new IBM DS4800 storage controller with no further disruption to the NAS Gateway.

8.1.1 Scenario description

The systems described in this section include:

- ▶ NetApp V3020c NAS Gateway cluster
- ▶ IBM b-type family switches
- ▶ HDS AMS500 disk storage system
- ▶ IBM DS4800 disk storage system

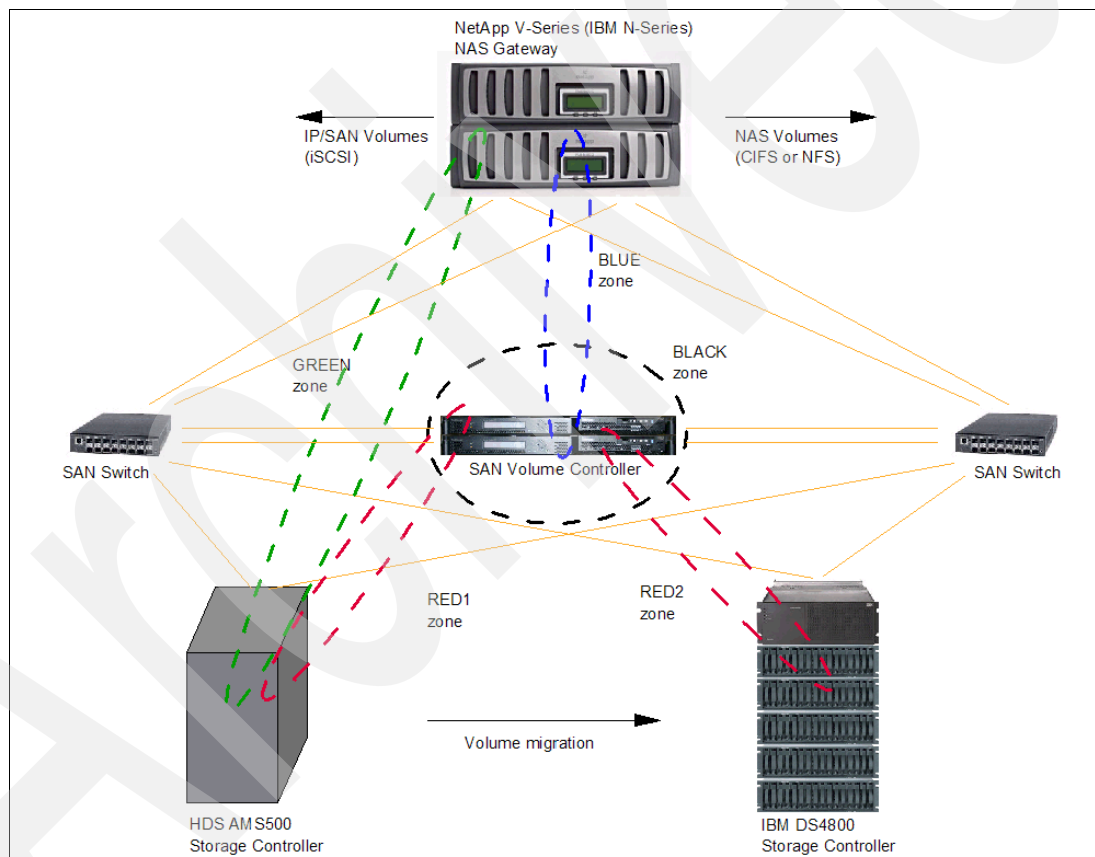


Figure 8-1 SVC implementation example with an existing NetApp V-Series storage controller

After the storage volumes have been migrated to the IBM DS4800 storage controller the original HDS AMS500 storage controller can be removed from the environment.

8.1.2 Summary of the SVC migration process

At a high level, the steps to migrate to SVC managed storage are the following:

- ▶ Preparation

- General
 - Review supported hardware and software from the configuration guide.
- NetApp V-series cluster
 - Review the existing storage configuration.
- SAN
 - Create a NetApp-SVC zone (enable now or later).
 - Create an HDS-SVC zone (enable now or later).
 - Create a DS4000-SVC zone (enable now).
- Existing storage controller
 - Identify the SCSI LUN IDs and unique IDs (UID) for the existing storage volumes.
- New storage controller
 - Configure some LUNs and present these to SVC
- SAN Volume Controller
 - Create a new managed disk group for the DS4000 LUNs.
 - Add the new DS4000 LUNs to this managed disk group.
- ▶ Migration to SVC managed storage
 - NetApp V-series cluster
 - Stop external access.
 - Shut down both clustered NetApp controllers.
 - SAN
 - Enable the NetApp-SVC and HDS-SVC zones.
 - Disable the NetApp-HDS controller zone.
 - Existing storage controller
 - Unmap the existing LUNs from the NetApp cluster.
 - Remap the existing LUNs to the SVC.
 - SAN Volume Controller
 - Create *image mode* VDisks from the HDS LUNs.
 - Define the NetApp host object.
 - Map the new VDisks to the NetApp host.
 - NetApp V-series cluster
 - Reboot.
 - Confirm correct access to the SVC managed storage.
- ▶ Migration of SVC managed volumes from the HDS AMS500 to the IBM DS4800
 - SAN Volume Controller
 - Migrate image mode VDisks to the DS4000 MDisk group.
 - Existing storage controller
 - On completion you can remove the existing HDS controller if desired.

Note: Be sure to review 1.4, “Planning and sizing considerations” on page 10 before continuing with the storage virtualization.

8.1.3 Confirm support for your environment

This section provides information about configuring the Network Appliance (NetApp) V-series NAS gateway system for attachment to the SVC.

Prior to implementing the SVC, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SAN Volume Controller supports models of the NetApp GF270, GF980, V3000, and V6000 series of subsystems.

Some models of the NetApp V-series range are equivalent to the IBM System Storage N-series; therefore, the information in this section also applies to the supported models of the IBM N-series.

Table 8-1 lists the supported models of the NetApp V-series and equivalent IBM N-series controllers.

Table 8-1 Supported models of the NetApp V-series and IBM N-series controllers

NetApp models	IBM models
GF270c	-
V3020c	N5200
V3050c	N5500
GF980c	-
V6030a	-
V6070a	-

See the following Web site for specific NetApp supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Concurrent host connection of the NetApp V-Series to the SVC

There are restrictions for sharing the SVC between a NetApp V-Series and other hosts.

Sharing ports

The SVC can be shared between a NetApp V-Series and other hosts under the following conditions:

- ▶ Both controllers in a NetApp NAS cluster must be able to see all paths (two maximum) to the LUs being provided by the back-end storage controller.
- ▶ No other hosts may share access to the LUs that are being used by the V-series.

Supported topologies

The NetApp V-series supports connection to the SAN Volume Controller under the following conditions:

- ▶ A maximum of two SAN paths are supported to any specific LU on the back-end storage controller. One is the primary path, and the other is the secondary (fail-over) path.
- ▶ The two paths must be to different controller ports (for example, two SVC nodes in an IO group) in order to avoid a single point of failure.

Further information

It is strongly recommended that you refer to the SVC version 4.1 Configuration Guide for further information regarding implementing the SVC with an existing storage controller:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

You must also refer to the following Network Appliance product documentation:

- ▶ *V-Series Planning Guide*, Part No. 210-00780_B0
- ▶ *V-Series Integration Guide for IBM Storage*, Part No. 210-00777_B0

Important: At the time of this writing, official support for NetApp V-series connection to the SVC was still undergoing support tests.

This section describes our experiences and is not an official commitment to provide support for any given configuration. You must confirm support for your particular configuration prior to implementation in your own environment.

8.1.4 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment, it may require coordination of effort between various support teams and an understanding of several technical domains.

Prerequisite conditions

In this scenario, we started with a NetApp V3020c that was already connected to an HDS AMS500 storage controller. The NetApp was configured to provide NAS and iSCSI storage over a LAN.

Since basic installation of the SVC, DS4800, and SAN switches is outside the scope of this book, we start our scenario at the point where all this preparatory work has already been completed.

We have assumed the following initial conditions:

- ▶ The NetApp V-series is configured as an active-active cluster.
- ▶ The NetApp V-series is already connected to the existing storage controller.
- ▶ The NetApp V-series is already configured to provide NAS and iSCSI LUN storage.
- ▶ The SVC nodes have been installed and minimally configured as a SVC cluster.
- ▶ The DS4800 storage controller has been installed, but no LUs have been created.
- ▶ The fibre channel switches have been installed and zoned for the existing environment.

Preparing the NetApp system

No significant changes are required on the NetApp V-series to enable the migration of back-end storage from the existing SAN storage controller to the SVC. However, before we proceed with the storage migration we should review the existing NetApp configuration.

Connect to the NetApp RLM console

A convenient way to administer the NetApp V-series is through the Remote LAN Module (RLM).

The RLM provides remote administration of the NetApp V-series regardless of the operational state of the system. This allows us to retain our CLI session even as we reboot the controller.

Example 8-1 shows connecting to the RLM on one of our V-series controllers.

Example 8-1 Connecting to the NetApp RLM console

```
c:\> putty netapp1-rlm

login as: naroot
Last login: Tue Aug  1 10:44:17 2006 from xxx.xxx.xxx.xxx
RLM netapp1> system console
Type Ctrl-D to exit.

Data ONTAP (netapp1.ssd.hursley.ibm.com)
login: root
Password: xxxxxxxxxxxx
Tue Aug  1 10:46 GMT [netapp1: rlm_console_login_m:info]: root logged in from RLM

netapp1>
```

Only one RLM session can be active to each NetApp controller at the same time.

List the storage configuration

Example 8-2 shows the storage configuration on one of our V-series controllers. Note that it also shows which disks are owned by the cluster partner.

Example 8-2 List the NetApp storage configuration

```
netapp1> sysconfig -r

Aggregate aggr0 (online, raid0) (block checksums)
  Plex /aggr0/plex0 (online, normal, active)
  RAID group /aggr0/plex0/rg0 (normal)

  RAID Disk Device      HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)  Phys (MB/blks)
  -----
  data    bittern:20.126L2 0a  -  -      -  LUN  N/A  8870/18165760  8960/18350080

Aggregate aggr1 (online, raid0) (block checksums)
  Plex /aggr1/plex0 (online, normal, active)
  RAID group /aggr1/plex0/rg0 (normal)

  RAID Disk Device      HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)  Phys (MB/blks)
  -----
  data    bittern:20.126L4 0a  -  -      -  LUN  N/A  4435/9082880   4480/9175040
  data    bittern:20.126L6 0a  -  -      -  LUN  N/A  4435/9082880   4480/9175040

Aggregate aggr2 (online, raid0) (block checksums)
  Plex /aggr2/plex0 (online, normal, active)
  RAID group /aggr2/plex0/rg0 (normal)

  RAID Disk Device      HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)  Phys (MB/blks)
  -----
  data    bittern:5.126L7  0a  -  -      -  LUN  N/A  6060/124125184 1220/125380584
  data    skua:14.126L8   0b  -  -      -  LUN  N/A  6060/124125184 6122/125380584
```

Spare disks (empty)

Partner disks

RAID Disk	Device	HA	SHELF	BAY	CHAN	Pool	Type	RPM	Used (MB/blks)	Phys (MB/blks)
partner	bittern:20.126L5	0a	-	-		-	LUN	N/A	0/0	4480/9175040
partner	bittern:20.126L1	0a	-	-		-	LUN	N/A	0/0	4480/9175040
partner	bittern:5.126L3	0a	-	-		-	LUN	N/A	0/0	8960/18350080

Example 8-3 shows another way to view the V-series disk configuration.

Example 8-3 List the NetApp storage configuration

```
netapp1> disk show
```

DISK	OWNER		POOL	SERIAL NUMBER
bittern:5.126L3	netapp2	(101166340)	Pool0	6005076801AB80065000000000002B67
bittern:5.126L7	netapp1	(101167306)	Pool0	6005076801AB80065000000000002BE1
skua:14.126L8	netapp1	(101167306)	Pool0	6005076801AB80065000000000002BE2
bittern:20.126L1	netapp2	(101166340)	Pool0	6005076801AB80065000000000002B5B
bittern:20.126L2	netapp1	(101167306)	Pool0	6005076801AB80065000000000002B6D
bittern:20.126L4	netapp1	(101167306)	Pool0	6005076801AB80065000000000002B59
bittern:20.126L5	netapp2	(101166340)	Pool0	6005076801AB80065000000000002B5A
bittern:20.126L6	netapp1	(101167306)	Pool0	6005076801AB80065000000000002B5C

The DISK column represents the LUNs that are being provided by the back-end storage controller. The name of each disk is prefixed with the name of the SAN switch through which the LUN is currently being accessed. This prefix can change over time as the V-Series balances the workload across the two paths to each LUN.

If you have a clustered NetApp V-series you will need to review the configuration of both V-Series controllers.

Preparing the SAN fabric

First we need to change the SAN zoning to allow the SVC to access the NetApp V-series, the existing storage controller, and the new IBM DS4800 storage controller.

Refer to Chapter 3, “Storage migration with the SVC” on page 29 for a description of the required SAN zones when implementing the SVC in an existing environment. In our example environment these are the RED1 and BLUE zones. You may choose to enable these zones now, or only create them at this point and enable them later.

If you are planning to migrate the data to a new storage controller (as in our example) then you should also zone the SAN to allow the SVC to access the new storage controller (RED2). You should enable this zone now because it will be needed during the preparation phase.

Refer to “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Note: It is assumed that the BLACK zone, which is the SVC node to SVC node zone, was created when creating the SVC cluster and therefore already exists in the SAN.

The NetApp V-series must have a maximum of two paths defined for access to the LUNs being provided by the back-end storage controller. On the SVC, all LUNs are visible on all ports. Therefore we must use SAN zoning to limit the number of paths to each LUN.

Figure 8-2 shows an example SAN zoning configuration to restrict the number of paths between the V-series and any LUNs that are to be provided by the SVC.




Figure 8-2 Zoning example for V-series to SVC connectivity

In a larger environment, with multiple SVC I/O groups, the SAN could be zoned so as to provide access to the VDisks on other I/O groups through the remaining (OC/OD) ports on the NetApp controllers.

Note: The SVC version 4.1 introduced a new option to the `mkhost` command, which can limit how many SVC ports any particular LUN is visible from. Although this method can be used, we recommend in an initial setup phase to use SAN zoning to accomplish this.

Example 8-4 shows how to confirm that each LUN has two paths from the V-series to the storage. This command shows all the disks across both NetApp cluster nodes.

Example 8-4 View the SAN paths to each LUN from the NetApp V-series controller

```
netapp1> storage show disk -p all
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	ADAPTER
bittern:5.126L3	-	skua:14.126L3	-	-	-	0a
bittern:5.126L7	-	skua:14.126L7	-	-	-	0a
bittern:20.126L2	-	skua:0.126L2	-	-	-	0a
bittern:20.126L4	-	skua:0.126L4	-	-	-	0a
skua:14.126L8	-	bittern:5.126L8	-	-	-	0b
skua:0.126L1	-	bittern:20.126L1	-	-	-	0b
skua:0.126L5	-	bittern:20.126L5	-	-	-	0b
skua:0.126L6	-	bittern:20.126L6	-	-	-	0b

In our environment the two SAN switches were named *bittern* and *skua* respectively.

If any disk does not have two SAN paths, then you must rectify this problem before continuing.

Preparing the existing storage controller

We need to review the existing LUN configuration on the existing storage controller.

The administrative interface on the HDS AMS500 can be accessed through the *Storage Navigator Modular GUI*, which must be installed on your workstation before it can be used.

In the Storage Navigator, double-click the existing storage controller to open the next management window.

From the Logical Status tab, list the existing LUN IDs (defined on each storage controller port). To do this, open the Array System View window, and perform the following steps:

1. Click **Logical Status**.
2. Click **Host Groups** and select the NetApp host group for each port.
3. Record the SCSI LUN ID numbers.

Open each of the ports to ensure that you record all of the LUNs for the NetApp cluster.

Record the SCSI LUN IDs that the existing storage controller presents to the host for the LUNs that are to be virtualized. You will need these ID numbers later to allow the SVC to present the same SCSI IDs to the host. This will allow the host to see the virtualized volumes as though they were the original, directly connected, volumes.

Note: It is not strictly necessary to assign the same SCSI ID number to the VDisks as the existing storage controller currently presents to the NetApp because the NetApp can identify the correct disks even if their SCSI ID changes.

However, we believe that it is good practice to maintain the same SCSI ID post migration. This ensures that the cluster's existing storage configuration remains valid. The SCSI IDs must be unique across the two clustered NetApp controllers.

Preparing the new storage controller

We need to create one or more LUNs on the new storage controller. These LUNs will be used as the target volumes for our volume migration from the existing storage controller.

The general recommendation for the SVC is that each of these LUNs be configured to the full capacity of the containing RAID array. Therefore, if possible, create one LUN per RAID array.

Change the LUN masking on the new storage controller to allow the SVC to access the LUNs that are to be virtualized.

The administrative interface on the IBM DS4800 can be accessed through the IBM System Storage DS4000 management interface.

Figure 8-3 shows the DS4800 management interface being used to configure volume LUN masking.

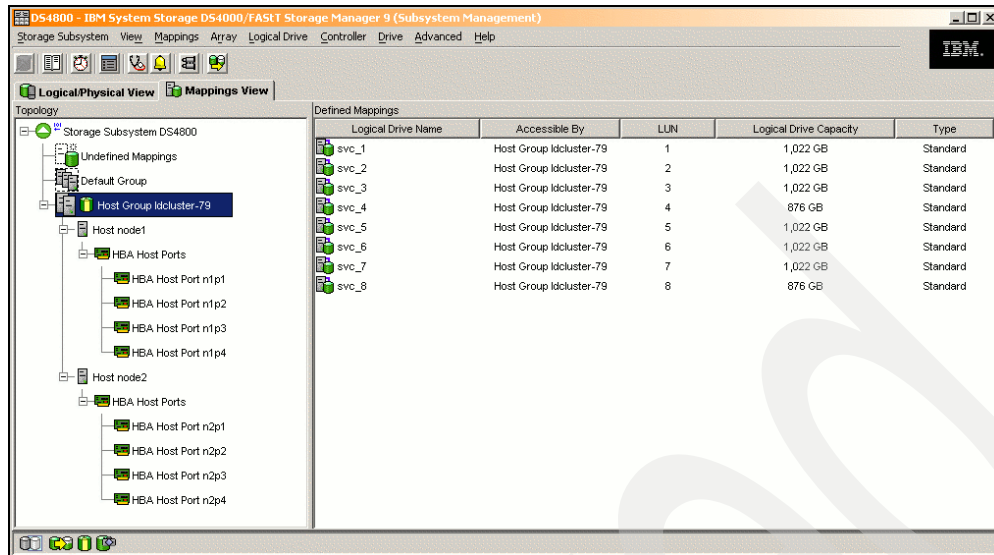


Figure 8-3 Enabling SVC access to new LUNs on the IBM DS4000 storage controller

Refer to the IBM DS4000 administration guide for detailed instructions on how to configure LUN masking on the DS4000 storage controller.

Since image mode volumes (from the existing storage controller) cannot be used as Quorum disks, the new managed mode MDiskS may also be used as the SVC quorum volumes.

Preparing the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized.

Use the `svcinfo lscontroller` command to list the storage controllers.

Example 8-5 shows the output of this command on our system.

Example 8-5 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfo lscontroller
```

id	controller_name	ctrl_s/n	vendor_id	product_id_low	product_id_high
0	bhindi		IBM	1742	
1	acadaca		HDS		

If both storage controllers are not visible to the SVC, check the SAN zone configuration.

Note: If you choose not to enable the SAN zone between the SVC and the existing storage controller (described earlier in this section), then the existing controller will not yet be visible with this command.

Record the SVC port numbers

First, you need to connect to the SVC and record the WWPN for each of the SVC nodes.

Connect to the SVC, then use the `svcinfo lsnode` command to list the SVC nodes.

Example 8-6 shows the output of this command on our system.

Example 8-6 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfolnode
```

id	name	UPS_serial_number	WWNN	status	IO_group_id
0	node1	10004BC047	500507680100017B	online	0
1	node2	100061P100	5005076801000124	online	0

To list further information about each of the SVC nodes issue the **svcinfolnode nodename** command.

Example 8-7 shows the output of this command for one node on our system.

Example 8-7 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfolnode node1
```

```
id 0
name node1
UPS_serial_number 10004BC047
WWNN 500507680100017B      (WWNN for node1)
status online
IO_group_id 0
IO_group_name iogroup0
partner_node_id 5
partner_node_name node2
config_node yes
UPS_unique_id 20400001124C0006
port_id 500507680140017B    (WWPN for port #4)
port_status active
port_speed 2Gb
port_id 500507680130017B    (WWPN for port #3)
port_status active
port_speed 2Gb
port_id 500507680110017B    (WWPN for port #1)
port_status active
port_speed 2Gb
port_id 500507680120017B    (WWPN for port #2)
port_status active
port_speed 2Gb
hardware 4F2
```

Repeat this process for each of the SVC nodes in your SVC cluster. You will need to know the WWPNs for all SVC cluster nodes when zoning the SAN.

Note that the WWPNs can be determined from each node's WWNN (shown underlined).

Create an MDG for the new storage controller's LUNs

Since we have already zoned the SAN and configured LUN masking on the new storage controller, the new LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

You can use the **svctask detectmdisk** command to rescan for newly available MDisk.

Note: This command produces no user feedback.

If the expected MDisks are still not visible to the SVC, re-check the SAN zoning and storage controller LUN masking.

Use the **svcinfolsmdisk** command to list the MDisks on the new DS4800 storage controller.

Example 8-8 shows the output of this command on our system.

Example 8-8 List MDisks on the DS4800 storage controller

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue  
controller_name=bhindi
```

```
id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
```

```
1,mdisk1,online,unmanaged,,,956.9GB,0000000000000007,bhindi,600...000  
2,mdisk2,online,unmanaged,,,956.9GB,0000000000000008,bhindi,600...000  
3,mdisk3,online,unmanaged,,,956.9GB,0000000000000008,bhindi,600...000  
4,mdisk4,online,unmanaged,,,820.2GB,0000000000000008,bhindi,600...000  
5,mdisk5,online,unmanaged,,,956.9GB,0000000000000008,bhindi,600...000  
6,mdisk6,online,unmanaged,,,956.9GB,0000000000000008,bhindi,600...000  
7,mdisk7,online,unmanaged,,,956.9GB,0000000000000008,bhindi,600...000  
8,mdisk8,online,unmanaged,,,820.2GB,0000000000000008,bhindi,600...000
```

This output has been edited for legibility. Note that the MDisks for the existing LUN are in *unmanaged* mode.

Use the **svctask mkmdiskgrp** command to add the unmanaged MDisks to a new MDG.

Example 8-9 shows the output of this command on our system.

Example 8-9 Add the new controller's LUNs to an MDG

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name bhindi_grp -ext 256 -mdisk  
mdisk1:mdisk2:mdisk3:mdisk4:mdisk5:mdisk6:mdisk7:mdisk8  
MDisk Group, id [0], successfully created
```

Use the **svctask lsmdisk** command to check that the MDG was created successfully.

Example 8-10 shows the output of this command on our system.

Example 8-10 List the MDisks that compose a given MDG

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue  
controller_name=bhindi
```

```
id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
```

```
1,mdisk1,online,managed,0,bhindi_grp,956.9GB,0000000000000007,bhindi,600...000  
2,mdisk2,online,managed,0,bhindi_grp,956.9GB,0000000000000008,bhindi,600...000  
3,mdisk3,online,managed,0,bhindi_grp,956.9GB,0000000000000008,bhindi,600...000  
4,mdisk4,online,managed,0,bhindi_grp,820.2GB,0000000000000008,bhindi,600...000  
5,mdisk5,online,managed,0,bhindi_grp,956.9GB,0000000000000008,bhindi,600...000  
6,mdisk6,online,managed,0,bhindi_grp,956.9GB,0000000000000008,bhindi,600...000  
7,mdisk7,online,managed,0,bhindi_grp,956.9GB,0000000000000008,bhindi,600...000  
8,mdisk8,online,managed,0,bhindi_grp,820.2GB,0000000000000008,bhindi,600...000
```

This output has been edited for legibility. Note that the MDisks for the new LUNs are now in *managed* mode.

The preparation is now complete and we are ready to proceed with the storage virtualization.

8.1.5 Migration to the SAN Volume controller

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SVC.

Configure the NetApp system

The NetApp V-series provides its own integrated multipath device support. It is not necessary to install any additional software to support connection to the SVC managed storage.

Shut down the NetApp cluster

Stop all user access to the NetApp V-series cluster. This includes NAS access (CIFS and NFS), and any iSCSI or FCP SAN access.

When external access to the NetApp cluster has been quiesced, shut down both of the NetApp controllers.

Example 8-11 shows how to shut down the first NetApp controller and prevent the other clustered controller from taking over its storage and services.

Example 8-11 Shutdown of the first NetApp controller

```
netapp1> halt -t 1 -f

Total number of connected CIFS users: 1
Total number of open CIFS files: 0
Warning: Terminating CIFS service while files are open may cause data loss!!
1 minute left until termination (^C to abort)...
Mon Jul 31 14:29:46 GMT [netapp1: cf.fsm.takeoverOfPartnerDisabled:notice]:
Cluster monitor: takeover of netapp2 disabled (local halt in progress)
Station xxx.xxx.xxx.xxx() could not be notified
Mon Jul 31 14:30:00 GMT [netapp1: monitor.globalStatus.critical:CRITICAL]: Cluster
failover of netapp2 is not possible: local halt in progress.

CIFS local server is shutting down...

waiting for CIFS shut down (^C aborts)...

CIFS local server has shut down...
Mon Jul 31 14:30:57 GMT [netapp1: kern.shutdown:notice]: System shut down because
: "halt".
Mon Jul 31 14:30:58 GMT [netapp1: pvif.allLinksDown:CRITICAL]: netapp1_vif: all
links down

CFE version 2.0.0 based on Broadcom CFE: 1.0.40
Copyright (C) 2000,2001,2002,2003 Broadcom Corporation.
Portions Copyright (c) 2002-2005 Network Appliance, Inc.

CPU type 0xF29: 2800MHz
Total memory: 0x80000000 bytes (2048MB)
CFE>
```

Repeat the shutdown for the other NetApp controller in the cluster.

Note: If you logged in through the RLM console, then you should be able to remain connected to the RLM even though the NetApp controller itself is shut down.

The NetApp cluster should remain briefly offline while we reconfigure the SAN. While the NetApp is offline we will remove direct access to the existing LUN and enable access to the virtualized LUN through the SVC.

Note: Although it should be possible to keep one half of the NetApp cluster online while the other half is offline and its storage is being migrated to the SVC, we decided to proceed with the entire cluster shutdown in order to minimize risk during the change.

Configure the SAN fabric

Change the SAN zoning to block the NetApp from directly accessing the existing LUN. From now on the NetApp will access the virtualized LUN through the SAN volume controller.

Refer to Chapter 3, “Storage migration with the SVC” on page 29 for a description of the required SAN zones when implementing the SVC in an existing environment. In our example environment we needed to disable the GREEN zone.

If you earlier chose not to enable the SAN zones from the NetApp to the SVC (BLUE zone), and the SVC to the storage (RED zone), then you should enable them now.

Refer to Appendix B., “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendor’s SAN equipment.

Configure the existing storage controller

Next, we need to change the LUN masking on the existing storage controller to block the host from directly accessing the existing storage, and to allow the SVC to access the LUNs that are to be virtualized.

Remove the NetApp LUN mappings

Return to the HDS AMS500 administration tool to perform the procedures in this section.

You will need to change the AMS500 administrative interface to *Management Mode* before you can make any changes to the controller’s configuration.

Open the Storage Navigator Modular window and do the following:

1. Click the **Management Mode** icon (pencil icon). Enter the management password.
2. Double-click the existing storage controller to open the next management window.

Open the Array System View window and do the following:

1. Click the **Logical Status** tab.
2. Click **Host Groups** and select the NetApp host group for each port.

You can now proceed to remove the LUN mappings from the NetApp host group.

Select the NetApp host group for each port, right-click **Logical Unit** and select **Modify Mapping**. This will open the Modify Mapping window.

After opening the Modify Mapping window, perform the following actions:

1. Click on each of the existing LUN mappings.
2. Click the down arrow to remove the mapping.
3. Click **OK** when all the LUN mappings have been removed.

Open each of the ports to ensure that you removed all of the LUNs for the NetApp cluster.

Create an SVC host group

We need to add a new SVC host group to each port for the SVC connection to the existing LUNs. Return to the Array System View window (Logical Status tab), right-click each port, and select **Add New Host Group**. You will then be prompted for the new host group information.

After entering the new SVC host group name, click **OK**. The main screen will update with the new host group. Open each of the ports to ensure that you have added an SVC host group to each.

Add the SVC LUN mappings

We can now add the LUN mappings for the SVC host group.

Select the SVC host group for each port, then right-click **Logical Unit** and select **Modify Mapping**. This will open the Modify Mapping window.

After starting the Modify Mapping window, perform the following actions:

1. Click each of the LUNs that were previously assigned to the NetApp.
2. Select the correct LUN ID for each LUN (as previously assigned).
3. Click the up arrow to add the mapping.
4. Click **OK** when all the LUN mappings have been added.

Open each of the ports to ensure that you have added all of the LUNs for the SVC host groups.

Add the SVC's WWPNS

We can now associate the SVC's WWPNS with the new SVC host groups.

Select the SVC host group for each port, right-click the **WWN**, and select **Modify WWN Information**. This will open the Modify WWN Information window.

After starting the Modify WWN Information window, perform the following actions:

1. Select the SVC node's WWPNS.
2. Click **Add** to add the ports.
3. Click **OK** when all the WWPNS have been added.

Open each of the ports to ensure that you have added all of the WWPNS for the SVC host groups.

Configure the AMS500 for SVC access

Finally, ensure that the HDS AMS500 is configured as per the SVC requirements.

Select the SVC host group for each port, right-click **Options**, select **Change**, and then both **Simple** and **Detail**.

Refer to the HDS administration guide for further instructions on how to configure the AMS500 storage controller.

Configure the SAN Volume Controller

Now that the other systems have been prepared we can proceed with the virtualization of the existing storage volume.

Define a host to the SVC

Before the host can access any VDisks it needs to be defined to the SVC.

Use the **svcinfolshbaportcandidate** command to list any unassigned host WWPNs that are visible to the SVC.

Example 8-12 shows the output of this command on our system.

Example 8-12 List unassigned host WWPNs that are visible to the SVC

```
IBM_2145:ldcluster:admin>svcinfolshbaportcandidate
```

```
id
500A098200004104    (This is the 0A HBA on the first NetApp host)
500A098300004104    (This is the 0B HBA on the first NetApp host)
500A098200004060    (This is the 0A HBA on the second NetApp host)
500A098300004060    (This is the 0B HBA on the second NetApp host)
```

Note: You will need to check on the NetApp (or the SAN switch) to determine which WWPN belongs to the NetApp (or host) with which you are working.

Use the **svctask mkhost** command to define the host to the SVC.

Example 8-13 shows the output of this command on our system.

Example 8-13 Define a host to the SVC

```
IBM_2145:ldcluster:admin>svctask mkhost -name netapp -hbawwpn
500A098200004104:500A098300004104:500A098200004060:500A098300004060
Host id [3] successfully created
```

Use the **svcinfolshost** command to check that the host was defined successfully.

Example 8-14 shows the output of this command on our system.

Example 8-14 Check a host definition on the SVC

```
IBM_2145:avcluster-06:admin>svcinfolshost netapp
```

```
id 1
name netapp
port_count 4
type generic
mask 1111
iogrp_count 4
WWPN 500A098300004104
node_logged_in_count 1
state active
WWPN 500A098300004060
```

```
node_logged_in_count 1
state active
WWPN 500A098200004060
node_logged_in_count 2
state active
WWPN 500A098200004104
node_logged_in_count 2
state active
```

Note: Although the NetApp V-series cluster is really two distinct NAS controllers, we can define it as a single host as far as the SVC is concerned. This is because the V-series requires that each LUN be visible through every HBA (for a maximum of two paths per LUN).

List the MDisks on the new storage controller

Since we have already correctly zoned the SAN and configured LUN masking on the existing storage controller, the existing LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

You can use the **svctask detectmdisk** command to rescan for newly available MDisks.

Note: This command produces no user feedback.

If the expected MDisks are still not visible to the SVC, double check the SAN zoning and storage controller LUN masking.

Use the **svcinfolsmdisk** command to list the MDisks on the new DS4800 storage controller.

Example 8-15 shows the output of this command on our system.

Example 8-15 List MDisks on the existing storage controller

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue
controller_name=acadaca
```

```
id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
```

```
11,mdisk11,online,unmanaged,,,68.3GB,0000000000000001,acadaca,600...000
12,mdisk12,online,unmanaged,,,68.3GB,0000000000000002,acadaca,600...000
13,mdisk13,online,unmanaged,,,68.3GB,0000000000000003,acadaca,600...000
14,mdisk14,online,unmanaged,,,68.3GB,0000000000000004,acadaca,600...000
15,mdisk15,online,unmanaged,,,68.3GB,0000000000000005,acadaca,600...000
16,mdisk16,online,unmanaged,,,68.3GB,0000000000000006,acadaca,600...000
17,mdisk17,online,unmanaged,,,68.3GB,0000000000000007,acadaca,600...000
18,mdisk18,online,unmanaged,,,68.3GB,0000000000000008,acadaca,600...000
```

This output has been edited for legibility. Note that the MDisk for the existing LUN is in *unmanaged* mode.

Create an MDG for the image mode volumes

Before we can configure the existing storage controller's LUNs as image mode VDisks we should define a new MDG to logically house those volumes.

Use the **svctask mkmdiskgrp** command to create an empty MDG for the image mode volumes.

Example 8-16 shows the output of this command on our system.

Example 8-16 Create an empty MDG for image mode volumes

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name imagegrp -ext 256
```

```
MDisk Group, id [1], successfully created
```

Although it is not strictly necessary to create a separate MDG for image mode volumes, we recommend doing so since it is good practice to logically separate image mode MDisk.

Important: To be able to migrate volumes between MDGs the extent size of the source and target MDGs must be the same.

Create image mode VDisks from the existing storage controller's LUNs

Now we can create the image mode VDisk using the existing unmanaged MDisk (which were previously connected directly to the NetApp controllers).

Use the **svctask mkvdisk** command to create the image mode VDisk.

Example 8-17 shows the output of this command on our system.

Example 8-17 Define a new image mode VDisk

```
IBM_2145:ldcluster:admin>svctask mkvdisk -mdiskgrp imagegrp -iogrp 0 -vtype image  
-name ntap1 -mdisk mdisk11
```

```
Virtual Disk, id [0], successfully created
```

Repeat the creation of image mode VDisks for all of the LUNs that were previously directly connected to the NetApp.

Attention: If you were to add the existing volumes to an MDG as normal striped MDisk then all data on those volumes would be lost. Since we intend to virtualize previously existing volumes with data intact, we must create *image mode virtual disks* instead.

Use the **svcinfn lsvdisk** command to verify that the VDisk was created successfully.

Example 8-18 shows the output of this command on our system.

Example 8-18 Display the properties of the new image mode VDisk

```
IBM_2145:ldcluster:admin>svcinfn lsvdisk ntap1
```

```
id 10  
name ntap1  
IO_group_id 0  
IO_group_name io_grp0  
status online  
mdisk_grp_id 1  
mdisk_grp_name imagegrp  
capacity 68.3GB  
type image
```



```
formatted no
mdisk_id
mdisk_name
FC_id
FC_name
RC_id
RC_name
vdisk_UID 6005076801AB80065000000000002BE1
throttling 0
preferred_node_id 2
fast_write_state empty
cache readwrite
udid
```

Map the image mode VDisk to the Host

Next, we need to define a relationship between the new image mode VDisk and our host system. This is the equivalent of LUN masking (or storage partitioning) on a normal storage controller.

Use the **svctask mkvdiskhostmap** command to map the new VDisk to the host.

Example 8-19 shows the output of this command on our system.

Example 8-19 Map a VDisk to a Host

```
IBM_2145:ldcluster:admin>svctask mkvdiskhostmap -host netapp -scsi 1 ntap1
```

```
Virtual Disk to Host map, id [0], successfully created
```

You should be aware that the NetApp V-series places restrictions on the size of the LUNs that it can utilize. However, since we are only migrating existing LUNs, which were previously connected to the NetApp, this will not affect us now. For future reference, at the time of writing, the maximum LUN size supported by the NetApp V-series is 500 GB, and the minimum size is 1 GB.

Note: It is not strictly necessary to assign the same SCSI ID number to the VDisk that the existing storage controller previously presented to the NetApp because the NetApp can identify the correct disks even if their SCSI ID changes.

However, we believe that it is good practice to maintain the same SCSI ID post migration. This ensures that the cluster's existing storage configuration remains valid. The SCSI IDs must be unique across the two clustered NetApp controllers.

Configure the host system

Now that the SAN has been re-configured you can restart the NetApp controllers.

Enter the **CFE> autoboot** command from the RLM console to restart the first NetApp controller.

The NetApp controller will then proceed to boot. You should restart the second NetApp controller at the same time.

When the NetApp controllers reboot they will see multiple paths to each of the SAN volumes.

The NetApp V-series provides its own integrated multipath device support. It is not necessary to install any additional software to support connection to the SVC managed storage.

Example 8-20 shows how to confirm that each LUN has two paths from the V-series to the SVC managed storage.

Example 8-20 View the SAN paths to each LUN from the NetApp V-series controller

```
netapp1> storage show disk -p all
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	ADAPTER
-----	---	-----	---	-----	----	-----
bittern:5.126L3	-	skua:14.126L3	-	-	-	0a
bittern:5.126L7	-	skua:14.126L7	-	-	-	0a
bittern:20.126L2	-	skua:0.126L2	-	-	-	0a
bittern:20.126L4	-	skua:0.126L4	-	-	-	0a
skua:14.126L8	-	bittern:5.126L8	-	-	-	0b
skua:0.126L1	-	bittern:20.126L1	-	-	-	0b
skua:0.126L5	-	bittern:20.126L5	-	-	-	0b
skua:0.126L6	-	bittern:20.126L6	-	-	-	0b

In our environment the two SAN switches were named *bittern* and *skua* respectively.

If any disks do not have two SAN paths then you must rectify this problem before continuing.

The data on the NetApp is now available for user access NAS or iSCSI access. If you disabled any external access to the NetApp hosted NAS services, then you should re-enable that access now.

Note: The storage volume has now been virtualized, but still resides on the original storage controller. Access to the volumes is managed through the SAN Volume Controller.

In the next section we migrate the storage volume to the new storage controller.

8.1.6 Move SVC managed volumes between storage controllers

The SAN volume controller enables the migration of a storage volume between storage controllers with no further disruption to the NetApp controllers.

Configure the SAN Volume Controller

In our example, we will migrate an image mode VDisk from the image mode MDG (on the original storage controller) to another MDG with MDisk from the new storage controller. During the migration the image mode VDisk is automatically converted to striped mode.

Migrate the VDisk to the new storage controller

Use the **svctask migratevdisk** command to start the volume migration.

Example 8-21 shows this command on our system.

Example 8-21 Migrating a VDisk between MDGs

```
IBM_2145:ldcluster:admin>svctask migratevdisk -vdisk ntap1 -mdiskgrp bhindi_grp
```

Note: This command produces no user feedback.

Remember that for volume migration to occur, the extent size of the source and target MDGs must be the same.

Monitor the VDisk migration

Use the **svcin** **lsmigrate** command to monitor the progress of the migration.

Example 8-22 shows the output of this command on our system.

Example 8-22 Monitoring the VDisk migration process

```
IBM_2145:ldcluster:admin>svcin lsmigrate
```

```
migrate_type MDisk_Group_Migration
progress 45
migrate_source_vdisk_index 141
migrate_target_mdisk_grp 0
max_thread_count 4
```

You can also use the **svcin** **lsdiskextent** command to view the location of the MDisk extents that comprise the VDisk.

Example 8-23 shows the output of this command on our system, while the Vdisk is being migrated (from MDisk 11).

Example 8-23 View the MDisk extents of a VDisks during migration

```
IBM_2145:ldcluster:admin>svcin lsdiskextent ntap1
```

id	number_extents
1	28
2	28
3	28
4	27
5	28
6	27
7	28
8	26
11	300 (This is the original image mode MDisk)

As the VDisk migration proceeds, the number of extents on the original MDisk will decrease. At the same time, the number of extents on the MDisk in the target MDG will increase.

When the volume migration is complete, no extents will remain on the original LUN (MDisk 11).

Note: The storage volume now resides on the new storage controller. Access to the virtualized volume is managed through the SAN Volume Controller.

Remove the original volume

After the volume migration is complete the original image mode MDisk remains defined in the original MDG even though it contains no extents. This MDisk can now be removed from the MDG, which changes its status to unmanaged mode.

Use the **svctask** **rmvdisk** command to remove the original MDisk from the MDG.

Example 8-24 shows the output of this command on our system.

Example 8-24 Remove a MDisk from a MDG

```
IBM_2145:ldcluster:admin>svctask rmmdisk -mdisk mdisk11 imagegrp
```

From that point the MDisk could be reassigned to another MDG or deleted from the storage controller.

Note: The `rmmdisk` command does *not* delete the MDisk or the LUN which it represents, it only removes it from the MDG and returns it to *unmanaged* mode. If you wish to remove the MDisk entirely, you will need to do so on the storage controller.

Use the `svcinfolsmdisk` command to view the status of the MDisk.

Example 8-25 shows the output of this command on our system.

Example 8-25 Display the properties of a MDisk

```
IBM_2145:ldcluster:admin>svcinfolsmdisk mdisk11

id 11
name mdisk11
status online
mode unmanaged
mdisk_grp_id
mdisk_grp_name
capacity 68.3GB
quorum_index
block_size 512
controller_name thali
ctrl_type 4
ctrl_WWNN 200200A0B80F06B1
controller_id 2
path_count 4
max_path_count 4
ctrl_LUN_# 00000000000000007
UID 600a0b80000f05cd0000000044cd910d0000000000000000000000000000000000
preferred_WWPN 200300A0B80F06B2
active_WWPN 200300A0B80F06B2
```

If you wish to, you can now delete the original LUN from the original storage controller.

Note: You should ensure that you do not destroy any controller LUNs until you have deleted them from the MDisk group and returned them to *unmanaged* mode.

8.2 Integration with NetApp FAS (back-end to SVC)

In this section we demonstrate the implementation of a SAN Volume Controller in an existing environment with a Network Appliance storage controller. The existing iSCSI storage volumes on the NetApp storage controller will be converted to Fibre Channel (FCP) volumes and virtualized using the SAN Volume Controller. The virtualized storage volumes will then be migrated to an IBM DS4800 storage controller.

Note: Be sure to review 1.4.2, “Planning and sizing” on page 11 before continuing with the storage virtualization.

8.2.1 Scenario description

In this scenario, we start with an existing environment where a NetApp FAS system has been configured to provide both file-level (NAS) and block-level (SAN) storage. The NAS services are a combination of CIFS and NFS, while the SAN services are accessed through the iSCSI protocol over a Gigabit ethernet network. There is no Fibre Channel infrastructure in the existing environment.

Our plan is to split the NAS and SAN storage onto different storage controllers, to transition from the Gigabit ethernet-based iSCSI SAN to a new fibre channel SAN. This will allow the FAS to be dedicated to NAS storage, while a new high performance DS4800 storage controller will be introduced to provide SAN storage through a new fibre channel infrastructure.

We demonstrate using the SAN Volume Controller to migrate the block-level volumes from the FAS to an IBM DS4800 storage controller. This involves changing the connection to the NetApp's block-level volumes from iSCSI protocol to Fibre Channel protocol, and then virtualizing those volumes with the SVC. The virtualized storage can then be migrated to another SAN storage controller with no disruption to the connected hosts.

The systems described in this section include:

- ▶ NetApp FAS-3050C NAS controller
- ▶ IBM SAN32M-2 fibre channel switches
- ▶ IBM DS4800 disk storage system

Figure 8-4 on page 212 shows the relationships between the various systems used in this scenario.




Figure 8-4 SVC implementation example with an existing NetApp FAS storage controller

After the block-level storage volumes have been migrated to the IBM DS4800 storage controller NetApp FAS controller could then be dedicated to providing NAS services.

8.2.2 Summary of the SVC migration process

At a high level, the steps to migrate to SVC managed storage are the following:

- ▶ Preparation
 - General
 - Review supported hardware and software from the configuration guide.
 - NetApp V-series cluster
 - Review the existing storage configuration.
 - SAN
 - Create a NetApp-SVC zone (enable now or later).
 - Create an HDS-SVC zone (enable now or later).
 - Create a DS4000-SVC zone (enable now).
 - Existing storage controller
 - Identify the SCSI LUN IDs and unique IDs (UID) for the existing storage volume.
 - New storage controller
 - Configure some LUNs and present these to SVC.

- SAN Volume Controller
 - Create a new managed disk group for the DS4000 LUNs.
 - Add the new DS4000 LUNs to this managed disk group.
- ▶ Migration to SVC managed storage
 - NetApp V-series cluster
 - Stop external access.
 - Shut down both clustered NetApp controllers.
 - SAN
 - Enable the NetApp-SVC and HDS-SVC zones.
 - Disable the NetApp-HDS controller zone.
 - Existing storage controller
 - Unmap the existing LUNs from the NetApp cluster.
 - Remap the existing LUNs to the SVC.
 - SAN Volume Controller
 - Create *image mode* VDisks from the HDS LUNs.
 - Define the NetApp host object.
 - Map the new VDisks to the NetApp host.
 - NetApp V-series cluster
 - Reboot.
 - Confirm correct access to the SVC managed storage.
- ▶ Migration of SVC managed volumes from the NetApp FAS to the IBM DS4800
 - SAN Volume Controller
 - Migrate image mode VDisks to the DS4000 MDisk group.
 - Existing storage controller
 - On completion you can remove the existing HDS controller if desired.

8.2.3 Confirm support for your environment

This section provides information about configuring the NetApp FAS systems for attachment to a SAN Volume Controller.

Prior to implementing the SVC it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Supported models

The SAN Volume Controller supports models of the NetApp FAS200, FAS900, FAS3000, and FAS6000 series of subsystems.

Some models of the NetApp FAS range are equivalent to the IBM System Storage N5000 series; therefore, the information in this section also applies to the supported models of the IBM N5000 series.

Table 8-2 lists the supported models of the NetApp FAS and IBM N5000 series controllers.

Table 8-2 Supported models of the NetApp FAS and IBM N-series controllers

NetApp models	IBM models
FAS250	-
FAS270	-
FAS920	-
FAS940	-
FAS960	-
FAS980	-
FAS3020	N5200
FAS3050	N5500
FAS6030	-
FAS6070	-

See the following Web site for the most up to date details about specific NetApp supported models and firmware levels:

<http://www.ibm.com/storage/support/2145>

Concurrent host connection to the NetApp FAS and the SVC

There are restrictions for sharing a NetApp FAS between a host and a SAN Volume Controller.

Sharing ports

The NetApp FAS can be shared between a host and a SAN Volume Controller under the following conditions:

- ▶ The SAN Volume Controller switch zone must include at least one target port from each filer to avoid a single point of failure.
- ▶ Target ports can be shared between the SAN Volume Controller and other hosts. However, you must define separate initiator groups (igroups) on the FAS for the SAN Volume Controller ports and the host ports.

Supported topologies

The SAN Volume Controller supports connection to the NetApp FAS under the following conditions:

- ▶ A single host cannot be connected to both the SAN Volume Controller and the NetApp FAS (block-level) to avoid the possibility of an interaction between multipathing drivers.
- ▶ You can connect other hosts directly to both the NetApp FAS and the SAN Volume Controller under the following conditions:
 - Target ports are dedicated to each host or are in a different igroup than the SAN Volume Controller.
 - LUNs that are in the SAN Volume Controller igroup are not included in any other igroup.

Quorum disk support

The SAN Volume Controller can use logical units (LUNs) that are exported by the NetApp FAS as quorum disks.

Further information

It is strongly recommended that you refer to the SVC version 4.1 Configuration Guide for further information regarding implementing the SVC with an existing storage controller:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

8.2.4 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Prerequisite conditions

In this scenario, we started with a NetApp FAS, which was already configured to provide NAS and iSCSI storage over LAN. There were no fibre channel devices in the existing environment.

Since the installation of the SVC, DS4800, and SAN switches is outside the scope of this book, we have started our example at the point where all this preparatory work has been completed.

We have assumed the following initial conditions:

- ▶ The NetApp FAS is already configured to provide NAS and iSCSI LUN storage.
 - The FCP feature has been purchased and is ready to install.
- ▶ The SVC nodes have been installed and minimally configured as an SVC cluster.
- ▶ The DS4800 storage controller has been installed, but no LUs have been created.
- ▶ The Fibre Channel switches have been installed, but no zones have been created.
 - There are some free FC ports to connect to the existing FAS.

Note: Our migration scenario uses the Fibre Channel target function on the NetApp FAS. This is a licensed function and may not already be configured in your environment.

Figure 8-5 shows the properties of the existing iSCSI volume as reported by Windows disk manager.




Figure 8-5 Properties of the existing iSCSI volume

Compare Figure 8-5 to the properties of the SVC managed FC volume as shown in Figure 8-20 on page 233.

Preparing the host system

The iSCSI connection to the existing NetApp LUN must be detached and prevented from attempting to reattach at system boot time.

Since the existing SAN was entirely iSCSI based, and there were no existing fibre channel devices, it was also necessary to install a fibre channel HBA into our host.

Disconnect from the iSCSI volume

Stop any applications that are accessing the existing volume before you detach it.

Start the Microsoft iSCSI initiator configuration tool by clicking the icon shown in Figure 8-6.



Figure 8-6 The Microsoft iSCSI initiator icon

The iSCSI initiator can also be accessed from the Windows Start menu by selecting: **Start** → **Programs** → **Microsoft iSCSI Initiator** → **Microsoft iSCSI Initiator**.

Figure 8-7 shows the iSCSI configuration tool.

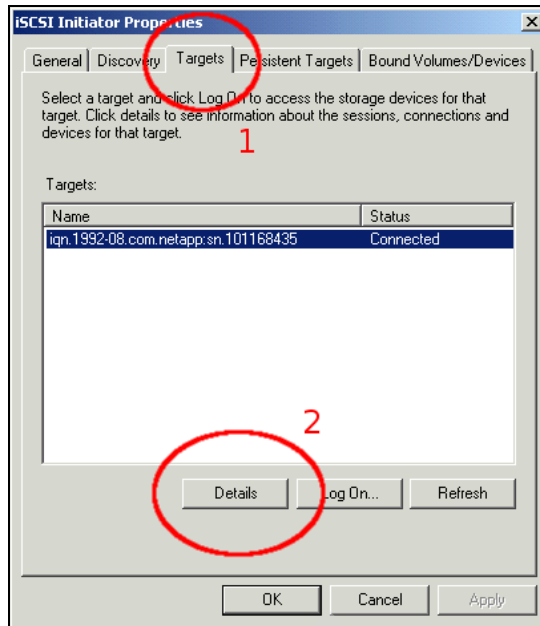


Figure 8-7 Show details for an iSCSI volume

After starting the iSCSI management tool, perform the following actions:

1. Select **Targets**, and highlight the iSCSI LU that is to be migrated.
2. Click **Details**.

Figure 8-8 shows disconnecting from the iSCSI volume.

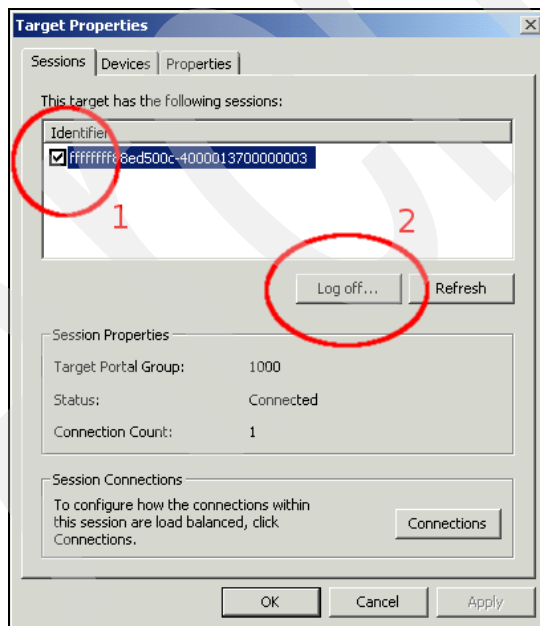


Figure 8-8 Disconnecting from an iSCSI volume

When the *details* dialog opens, perform the following actions:

1. Select the iSCSI volume from the list.
2. Click **Log off** to disconnect the iSCSI volume.

3. Click **OK** to complete the removal of the iSCSI volume.

The iSCSI volume should no longer appear as a drive in Windows Explorer or in the Device Manager.

Note: If you are definitely not going to use iSCSI volumes in the future, you can also uninstall the iSCSI initiator software from the host system.

Install the HBAs

Since our existing SAN was entirely iSCSI based, and there were no existing fibre channel devices, it was necessary to install a fibre channel HBA into our host.

The following procedure briefly describes the fibre channel HBA installation:

1. Shut down the host
 - a. Configure any applications to *not* start automatically at system boot.
2. Physically install a pair of supported fibre channel HBAs into the host.
 - a. Either two single port, or one dual port, Fibre Channel HBA.
 - b. If the host was using a dedicated iSCSI adapter, you may choose to remove it.
3. Boot the host.
4. Install the software device drivers for the HBA.
 - a. Install the IBM FC HBA device driver.
 - b. Install the IBM SDD multi-path driver.
5. Reboot the host.

When the host reboots, it will be ready to access the existing LUN (after its access is changed from the iSCSI protocol to FC protocol, and it is migrated to the SVC).

Note: No applications should be running at this point because the SAN volumes are not yet available.

It may be advisable to configure the server to prevent the automatic starting of any applications on the next boot. This is to prevent any problems with applications trying to access the SAN volumes before we can load the IBM SDD multi-path device driver.

Note: Although some parts of the storage virtualization process can be performed with the host system online, we chose to shut down the host so as to minimize risk during the change.

Preparing the SAN fabric

Change the SAN zoning to allow the SVC to access the existing storage controller and the host whose data is to be virtualized.

Refer to Chapter 3, “Storage migration with the SVC” on page 29 for a description of the required SAN zones for implementing the SVC in an existing environment. In our example environment these are the RED1 and BLUE zones.

If you are planning to migrate the data to a new storage controller (as in our example) then you should also zone the SAN to allow the SVC to access the new storage controller (RED2).

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Note: It is assumed that the BLACK zone, which is the SVC node to SVC node zone, was created when creating the SVC cluster and therefore already exists in the SAN.

Record the SVC port numbers

Connect to the SVC and record the WWPN for each of the SVC nodes.

Use the **svcinfo lsnode** command to list the SVC nodes. Example 8-26 shows the output of this command on our system.

Example 8-26 List the storage controllers

IBM_2145:ldcluster:admin>svcinfo lsnode					
id	name	UPS_serial_number	WWNN	status	IO_group_id
0	node1	UV201A0188	50050768010000C1	online	0
1	node2	UV201A0190	5005076801002838	online	0

To obtain the additional required information, use the **svcinfo lsnode nodename** command to list the properties of each of the SVC nodes.

Example 8-27 shows the output of this command for one node on our system.

Example 8-27 List the storage controllers

IBM_2145:ldcluster:admin>svcinfo lsnode node1	
id 0	
name node1	
UPS_serial_number UV201A0188	
WWNN 500507680100010E	(WWNN for node1)
status online	
IO_group_id 0	
IO_group_name iogroup0	
partner_node_id 5	
partner_node_name node2	
config_node yes	
UPS_unique_id 20400001124C0005	
port_id 50050768014000C1	(WWPN for port #4)
port_status active	
port_speed 2Gb	
port_id 50050768013000C1	(WWPN for port #3)
port_status active	
port_speed 2Gb	
port_id 50050768011000C1	(WWPN for port #1)
port_status active	
port_speed 2Gb	
port_id 50050768012000C1	(WWPN for port #2)
port_status active	
port_speed 2Gb	
hardware 4F2	

Repeat this process for each of the SVC nodes in your SVC cluster. You will need to know the WWPN for all SVC cluster nodes when zoning the SAN. Note that the WWPNs can be determined from each node’s WWNN.

Configure the existing storage controller

The administrative interface on the NetApp FAS can be accessed through a Web browser at the following URL:

http://netapp-controller-address/na_admin/

The NetApp FAS also has a command line interface that can be accessed through *telnet* (or SSH, if enabled) to the FAS controller.

Figure 8-9 shows starting FilerView® from the NetApp FAS administrative Web page.

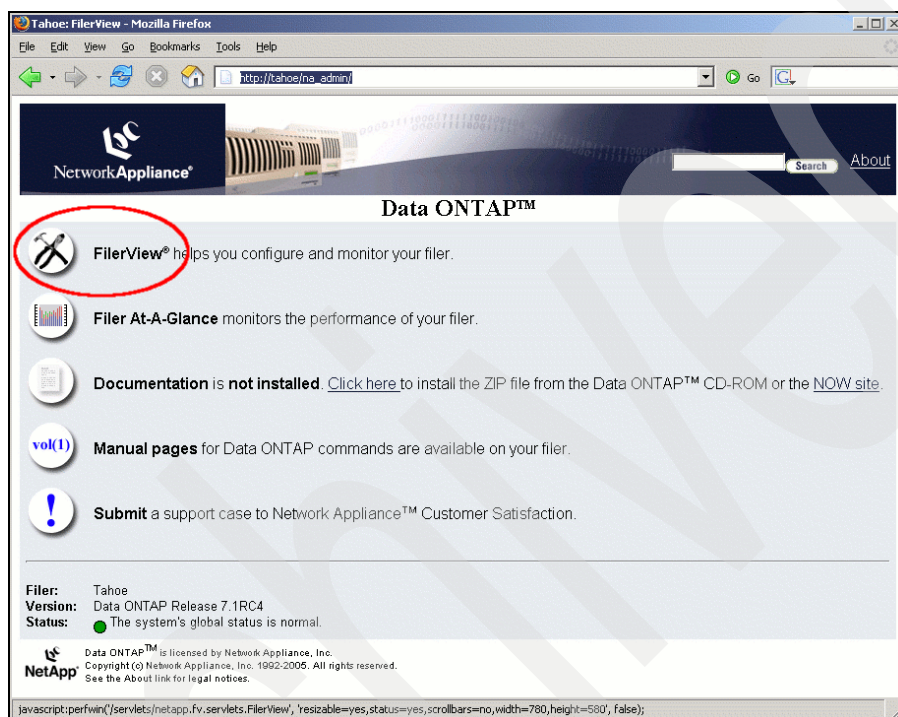


Figure 8-9 The NetApp FAS administrative interface

Click **FilerView** to start the administrative interface.

License the FCP feature on the NetApp FAS

In our environment, the FAS was already configured with NAS (CIFS and NFS) and iSCSI protocol licenses, but was not yet licensed for the FC protocol.

Before you can migrate the existing LU to the SVC, you must license the FC protocol.

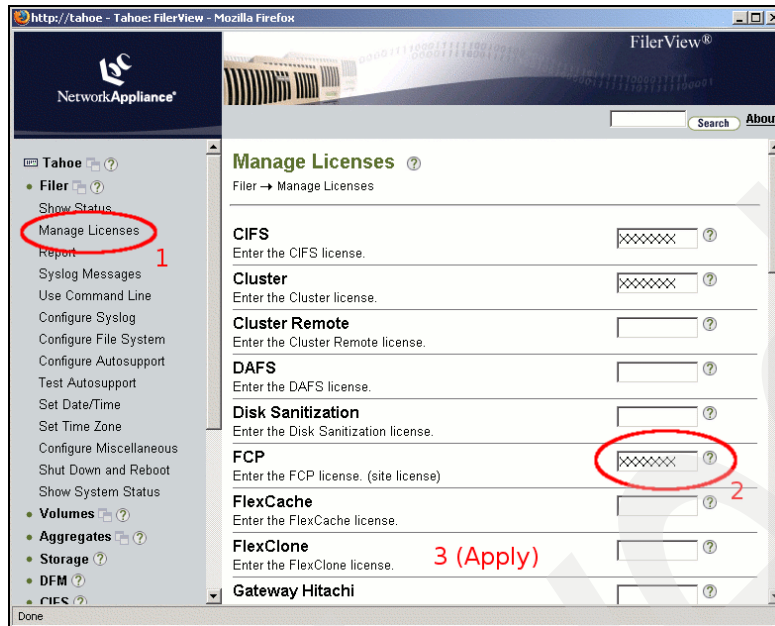


Figure 8-10 Licensing the FC protocol on the NetApp FAS

After starting the NetApp FilerView management tool, perform the following actions:

1. Click **Filer** and **Manage Licenses** from the left menu.
2. Enter the FC license key.
3. Click **Apply** to enable the license.

Note: The FC protocol is a licensed feature that must be purchased from your NetApp supplier.

Disable access to the existing LUN through the iSCSI protocol

The existing LU is already mapped to an iSCSI initiator group. You must remove this mapping before remapping it to a new FC initiator group.

Figure 8-11 shows the exiting LUN IDs on the NetApp FAS.

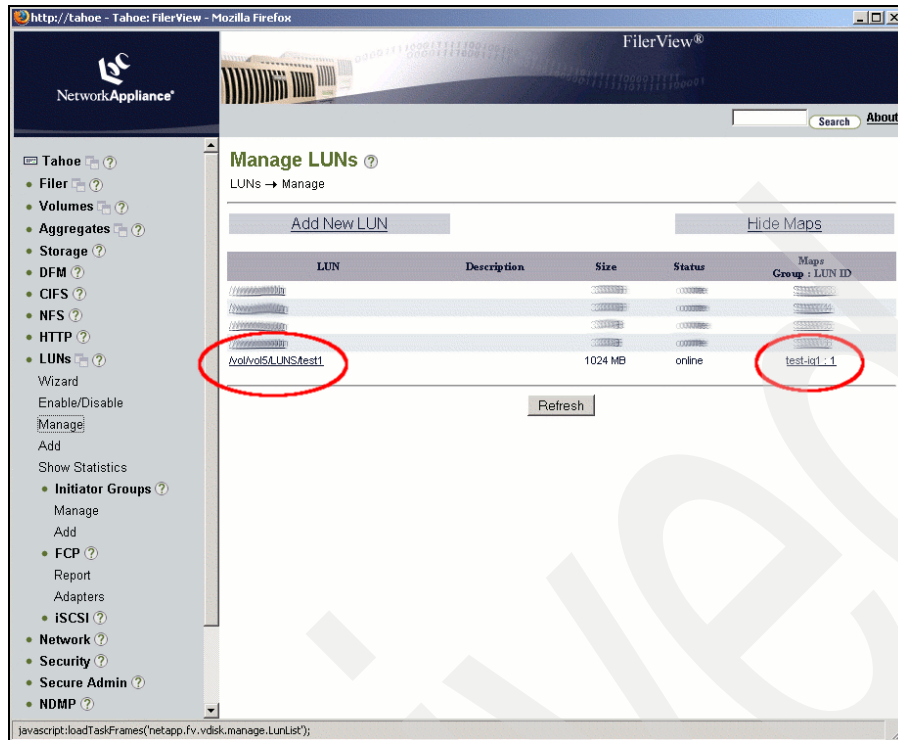


Figure 8-11 Managing LUNs on the NetApp FAS

Compare Figure 8-11 to Figure 8-18 on page 227, which shows the same view after the changes are made.

After starting the NetApp FilerView management tool, perform the following actions:

1. Click **LUNs** and **Manage** from the left menu.
2. Click the LUN that you are working with.
3. Click **Map LUN** in the Modify LUN panel (Figure 8-12).

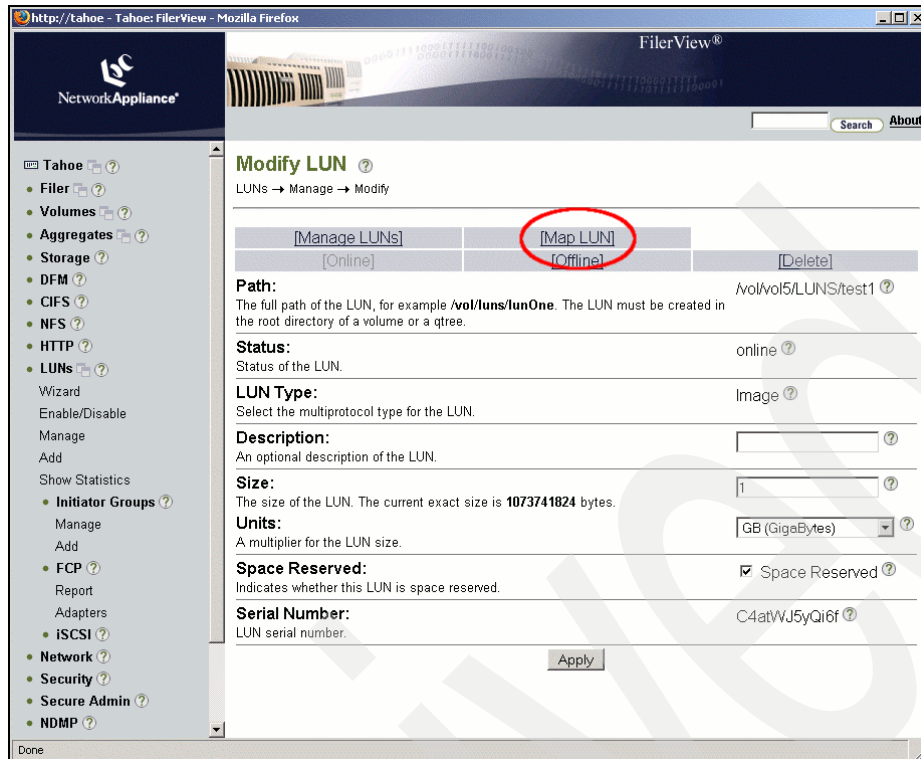


Figure 8-12 LUN details on the NetApp FAS

Select the LUN to unmap and click **Apply** (Figure 8-13).

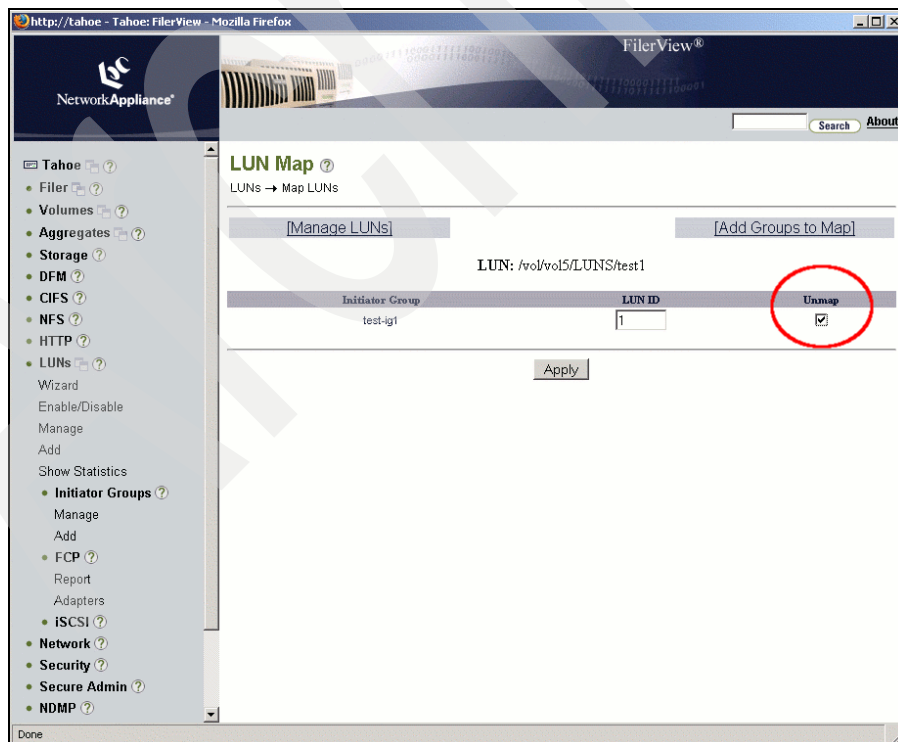


Figure 8-13 Unmapping a LUN on the NetApp FAS

The selected LUN is no longer accessible through the previous initiator group (iSCSI protocol).

Record the SCSI LUN IDs that the existing storage controller presents to the host for the LUNs that are to be virtualized. You will need these ID numbers later to allow the SVC to present the same SCSI IDs to the host. This will allow the host to see the virtualized volumes as though they were the original, directly connected, volumes.

Enable access to the existing LUN through the FC protocol

Now that the existing LU is no longer accessible through the iSCSI protocol, you can enable access through a new initiator group (and FC protocol) by performing the following steps:

Create a new initiator group on the FAS that uses the FC protocol. This initiator group will be associated with the WWPNs of every node in the SVC cluster.

Figure 8-14 shows the addition of a FC initiator group to the NetApp FAS.

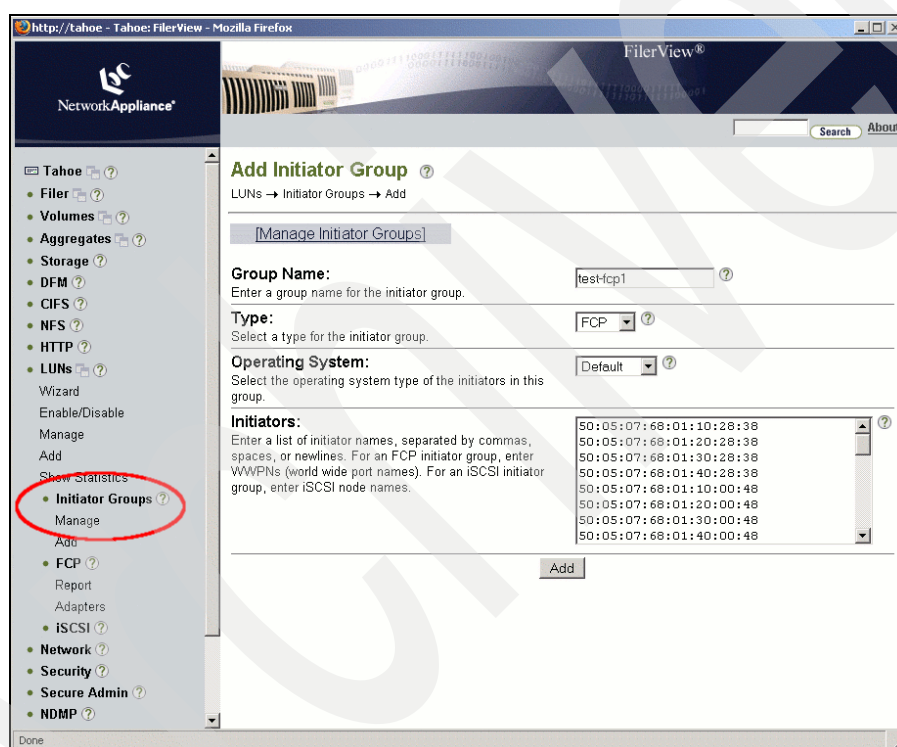


Figure 8-14 Adding an SVC host group to the HDS Lightning

After starting the NetApp FilerView management tool, perform the following actions:

1. Click **LUNs**, then **Initiator Groups**, and **Add** from the left menu.
2. Enter a logical name for the new group.
3. Select **FCP** as the protocol type.
4. Select **Default** as the Operating System (required for SVC access).
5. Enter the WWPN for every node in the SVC cluster.
6. Click **Add** to create the new group.

Now that the FCP initiator group has been defined, we can map our LU to the new group.

Figure 8-15 shows mapping an initiator group to a LUN on the NetApp FAS.

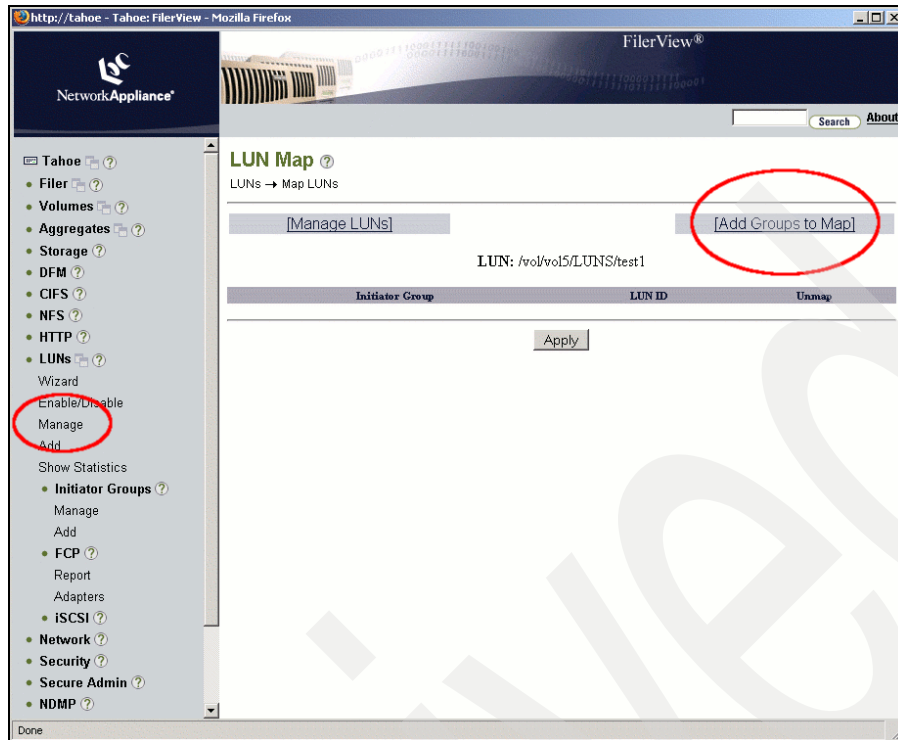


Figure 8-15 Mapping an initiator group to the LUN

After starting the NetApp FilerView management tool, perform the following actions:

1. Click **LUNs** and **Manage** from the left menu.
2. Select the LUN to Manage.
3. Click **Add Groups to Map** from the LUN Map window.

Select the new initiator group from the menu as shown in Figure 8-16, then click **Add**.

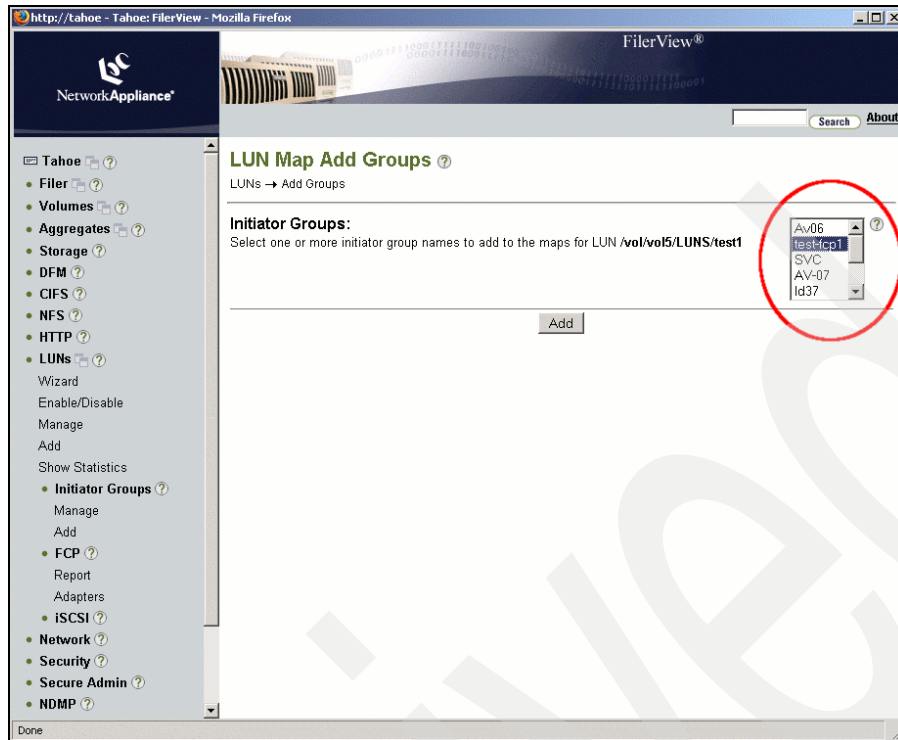


Figure 8-16 Add an initiator group to a LUN on the NetApp FAS

Figure 8-17 shows where to enter the SCSI LUN ID that the NetApp FAS will use to present this LU to the SVC.

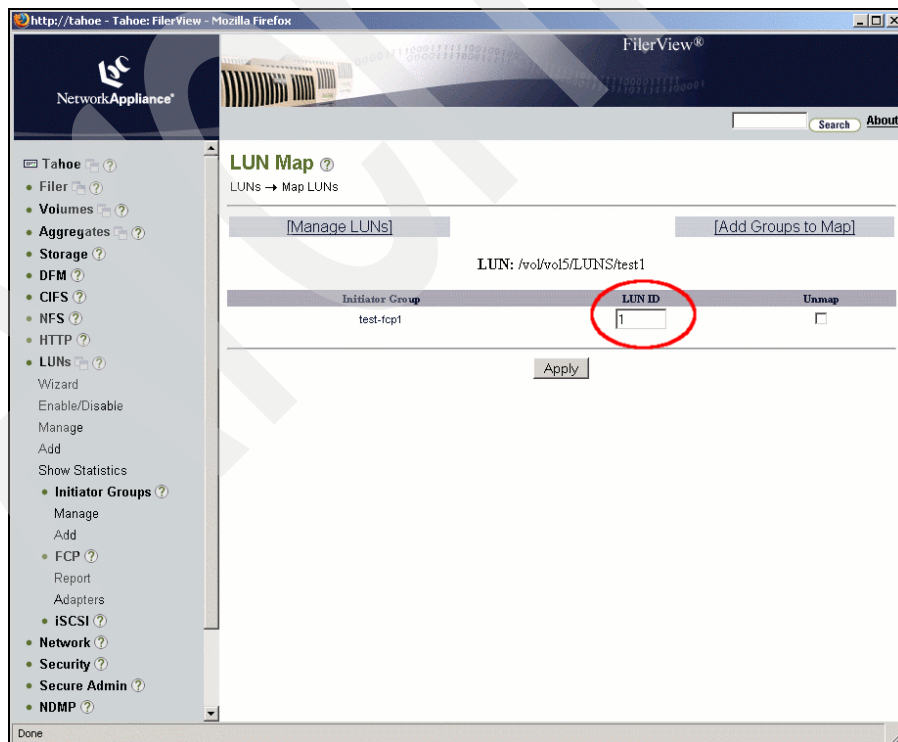


Figure 8-17 Setting the SCSI LUN ID on the NetApp FAS

Note: This is only the LUN ID that the FAS will present to the SVC. It does not need to be the same ID as the FAS previously presented to the host (through iSCSI protocol). However, it does need to be unique between the two FAS nodes (in a clustered FAS system).

Figure 8-18 shows the LUN and initiator group mappings after our changes.

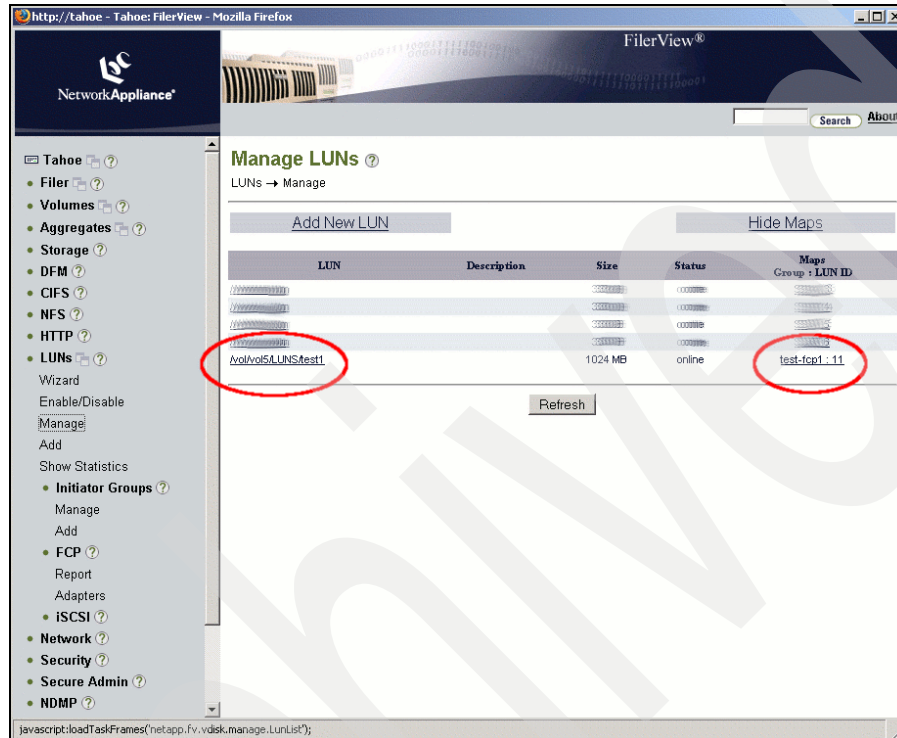


Figure 8-18 Review changes to the LUN and iGroup mappings

Compare Figure 8-18 to the same view before we made our changes (see Figure 8-11 on page 222). Note that the initiator group has changed.

Refer to the NetApp administration guide for detailed instructions on how to configure Fibre Channel LUNs on the FAS storage controller.

Configure the new storage controller

We need to create one or more LUNs on the new storage controller. These LUNs will be used as the target volumes for our volume migration from the existing storage controller.

The general recommendation for SVC use is that each of these LUNs be configured to the full capacity of the containing RAID array. Therefore, if possible, create one LUN per RAID array.

Change the LUN masking on the new storage controller to allow the SVC to access the LUNs that are to be virtualized.

The administrative interface on the IBM DS4800 can be accessed through the IBM System Storage DS4000 management interface.

Figure 8-19 shows the DS4000 administrative interface being used to configure volume LUN masking.

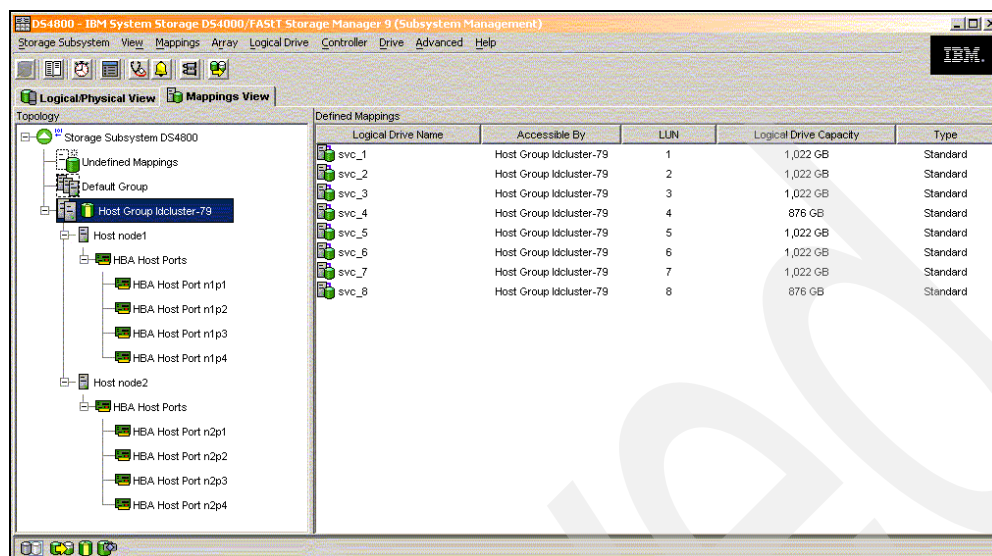


Figure 8-19 Enabling SVC access to new LUNs on the IBM DS4000 storage controller

Refer to the IBM DS4000 administration guide for detailed instructions on how to configure LUN masking on the DS4000 storage controller.

Since image mode volumes (from the existing storage controller) cannot be used as Quorum disks, the new managed mode MDisks will also be used as the SVC quorum volumes.

Preparing the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized.

Use the **svcinfolscntroller** command to list the storage controllers.

Example 8-28 shows the output of this command on our system.

Example 8-28 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfolscntroller
```

id	controller_name	ctrl_s/n	vendor_id	product_id_low	product_id_high
0	laverbread		IBM	1742	
1	fas-netapp	C4ar0J5BAU0e	NETAPP	LUN	

Use the **svcinfolsmdisk** command to list the MDisk on the existing storage controller.

Example 8-29 shows the output of this command on our system.

Example 8-29 List MDisk on the existing storage controller

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -filtervalue controller_name=fas-netapp
```

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
o,mdisk0,online,unmanaged,1.0GB,000000000000000B,fas-netapp,60a...000

This output has been edited for legibility. Note that the MDisk for the existing LUN is in *unmanaged* mode.

Create a MDG for the new storage controller's LUNs

Since we have already correctly zoned the SAN and configured LUN masking on the back-end storage controllers, the LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

Use the **svctask detectmdisk** command to rescan for newly available MDisks.

Note: This command produces no user feedback.

If the expected MDisks are still not visible to the SVC, then you should double check the SAN zoning and storage controller LUN masking.

Use the **svcinfolsmdisk** command to list the MDisks on the new storage controller.

Example 8-30 shows the output of this command on our system.

Example 8-30 List MDisks on the new storage controller

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue  
controller_name=laverbread
```

id	name	status	mode	mdisk_grp_id	mdisk_grp_name	capacity	ctrl_LUN_#	controller_name	UID
----	------	--------	------	--------------	----------------	----------	------------	-----------------	-----

1	mdisk1	online	unmanaged	0	laver1	135.4GB	0000000000000000	laverbread	600...000
2	mdisk2	online	unmanaged	0	laver1	135.4GB	0000000000000000	laverbread	600...000
3	mdisk3	online	unmanaged	0	laver1	135.4GB	0000000000000000	laverbread	600...000
4	mdisk4	online	unmanaged	0	laver1	135.4GB	0000000000000000	laverbread	600...000

This output has been edited for legibility. Note that the MDisks for the new LUNs are in *unmanaged* mode.

Use the **svctask mkmdiskgrp** command to add the unmanaged MDisks to a new MDG.

Example 8-31 shows the output of this command on our system.

Example 8-31 Add the new controller's LUNs to a MDG

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name sharkgrp -ext 32 -mdisk  
mdisk1:mdisk2:mdisk3:mdisk4
```

```
MDisk Group, id [1], successfully created
```

Use the **svctask lsmdisk** command to check that the MDG was created successfully.

Example 8-32 shows the output of this command on our system.

Example 8-32 List the MDisks that compose a given MDG

```
IBM_2145:ldcluster:admin>svcinfolsmdisk -delim , -filtervalue controller_name=laverbread
```

id	name	status	mode	mdisk_grp_id	mdisk_grp_name	capacity	ctrl_LUN_#	controller_name	UID
----	------	--------	------	--------------	----------------	----------	------------	-----------------	-----

1	mdisk1	online	managed	0	laver1	135.4GB	0000000000000000	laverbread	600...000
---	--------	--------	---------	---	--------	---------	------------------	------------	-----------

```
2,mdisk2,online,managed,0,laver1,135.4GB,0000000000000000,laverbread,600...000
3,mdisk3,online,managed,0,laver1,135.4GB,0000000000000000,laverbread,600...000
4,mdisk4,online,managed,0,laver1,135.4GB,0000000000000000,laverbread,600...000
```

This output has been edited for legibility. Note that the MDisks for the new LUNs are now in *managed* mode.

Define a host to the SVC

Before the host can access any VDisks, it needs to be defined to the SVC.

Use the **svcinio lshbaportcandidate** command to list any unassigned host WWPNs that are visible to the SVC.

Example 8-33 shows the output of this command on our system.

Example 8-33 List unassigned host WWPNs that are visible to the SVC

```
IBM_2145:ldcluster:admin>svcinio lshbaportcandidate
id
210000E08B10C091      (This is the first HBA on our host)
210100E08B30C091      (This is the second HBA on our host)
```

Note: You will need to check on the host (or the SAN switch) to determine which WWPN belongs to the host with which you are working.

Use the **svctask mkhost** command to define the host to the SVC.

Example 8-34 shows the output of this command on our system.

Example 8-34 Define a host to the SVC

```
IBM_2145:ldcluster:admin>svctask mkhost -name general -hbawwpn
210000E08B10C091:210100E08B30C091 -type generic
```

```
Host id [3] successfully created
```

Use the **svcinio lshost** command to check that the host was defined successfully.

Example 8-35 shows the output of this command on our system.

Example 8-35 Check a host definition on the SVC

```
IBM_2145:avcluster-06:admin>svcinio lshost pauline
id 0
name pauline
port_count 2
type generic
mask 1111
iogrp_count 4
WWPN 210000E08B10C091
node_logged_in_count 8
state active
WWPN 210100E08B30C091
node_logged_in_count 8
state active
```

The preparation is now complete and you are ready to proceed with the storage virtualization.

8.2.5 Migration to the SAN Volume controller

This section describes the process of virtualizing the volume on the existing storage controller. Access to this volume will then be managed by the SVC.

Configure the SAN fabric

If the SAN has not yet been configured with the correct zones for Host to SVC, and SVC to FAS access, you should configure it now.

Change the SAN zoning to block the host from directly accessing the existing LUN. From now on the host will access the virtualized LUN through the SAN volume controller.

Refer to Chapter 3, “Storage migration with the SVC” on page 29 for a description of the required SAN zones for implementing the SVC in an existing environment. In our example environment we needed to disable the GREEN zone.

At this point the SAN should already be configured to allow the host to access the SVC.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendor’s SAN equipment.

Configure the SAN Volume Controller

Now that the other systems have been prepared you can proceed with the virtualization of the existing storage volume.

Create an MDG for the image mode volumes

Before you can configure the existing storage controller’s LUNs as image mode VDisks, you should define a new MDG to logically house those volumes.

Use the **svctask mkmdiskgrp** command to create an empty MDG for your image mode volumes.

Example 8-36 shows the output of this command on our system.

Example 8-36 Create an empty MDG for image mode volumes

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name imagegrp -ext 32
```

```
MDisk Group, id [9], successfully created
```

Although it is not strictly necessary to create a separate MDG for image mode volume, we believe that it is good practice.

Important: To be able to migrate volumes between MDGs, the extent size of the source and target MDGs must be the same.

Create image mode VDisks from existing storage controller’s LUNs

Create the image mode VDisk using the existing unmanaged MDisk that was previously connected directly to the host.

Use the **svctask mkvdisk** command to create the image mode VDisk.

Example 8-37 shows the output of this command on our system.

Example 8-37 Define a new image mode VDisk

```
IBM_2145:ldcluster:admin>svctask mkvdisk -mdiskgrp imagegrp -iogrp 0 -vtype image  
-name pauline-01 -mdisk mdisk0
```

```
Virtual Disk, id [1029], successfully created
```

Attention: If you were to add the existing volumes to an MDG as normal *striped* MDisks then all data on those volumes would be lost. Since we intend to virtualize previously existing volumes with data intact, we must create image mode virtual disks instead.

Use the **svcinfolsvdisk** command to check that the VDisk was created successfully.

Example 8-38 shows the output of this command on our system.

Example 8-38 Display the properties of the new image mode VDisk

```
IBM_2145:ldcluster-65:admin>svcinfolsvdisk pauline-01
```

```
id 1435  
name pauline-01  
IO_group_id 1  
IO_group_name iogroup1  
status online  
mdisk_grp_id 7  
mdisk_grp_name fas-image  
capacity 1.0GB  
type image  
formatted no  
mdisk_id 28  
mdisk_name mdisk28  
FC_id  
FC_name  
RC_id  
RC_name  
vdisk_UID 6005076801BD000870000000000005A4  
throttling 0  
preferred_node_id 18  
fast_write_state empty  
cache readwrite  
udid
```

Map the image mode VDisk to the Host

Define a relationship between the new image mode VDisk and your host system. This is the equivalent of LUN masking (or storage partitioning) on a normal storage controller.

Use the **svctask mkvdiskhostmap** command to map the new VDisk to the host.

Example 8-39 shows the output of this command on our system.

Example 8-39 Map a VDisk to a Host

```
IBM_2145:ldcluster-65:admin>svctask mkvdiskhostmap -host pauline -scsi 1  
pauline-01
```

Virtual Disk to Host map, id [1], successfully created

Note: It is important to define the same SCSI ID number to the VDisk as the existing storage controller previously presented to the host. This will allow the virtualized LUN to appear as the original LUN when later exposed to the host.

Configure the host system

Now that the VDisk has been mapped, you should be able to see the LUN from the host.

You can use the IBM SDD command line to check that the SAN volumes are available. Example 8-40 shows the output of this command on our system.

Example 8-40 Checking SAN devices with SDD

```
C:\Program Files\IBM\Subsystem Device Driver>datapath query device  
Total Devices : 1
```

```
DEV#: 0 DEVICE NAME: Disk0 Part0 TYPE: 2145 POLICY: OPTIMIZED  
SERIAL: 6005076801BD000870000000000005A4
```

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port3 Bus0/Disk8 Part0	OPEN	NORMAL	0	0
1	Scsi Port3 Bus0/Disk8 Part0	OPEN	NORMAL	15628	0
2	Scsi Port4 Bus0/Disk8 Part0	OPEN	NORMAL	0	0
3	Scsi Port4 Bus0/Disk8 Part0	OPEN	NORMAL	15870	0

The data on the SAN disk is now available on the host for application and user access.

Figure 8-20 shows the properties of the existing iSCSI volume as reported by Windows disk manager.




Figure 8-20 Properties of the remapped FCP volume

Compare Figure 8-20 to the properties of the iSCSI volume as shown in Figure 8-5 on page 216.

Note: The storage volume has now been virtualized, but still resides on the original storage controller. Access to the volumes is managed through the SAN Volume Controller

In the next section we describe how to migrate the storage volume to the new storage controller.

8.2.6 Move SVC managed volumes between storage controllers

The SAN volume controller enables the migration of a storage volume between storage controllers with no disruption to the host systems.

Configure the SAN Volume Controller

In this example, we migrate an image mode VDisk between an MDG on one storage controller and another MDG on the new storage controller. During the migration the image mode VDisk is automatically converted to striped mode.

Migrate the VDisk to the new storage controller

Use the **svctask migratevdisk** command to start the volume migration.

Example 8-41 shows the output of this command on our system.

Example 8-41 Migrating a VDisk between MDGs

```
IBM_2145:ldcluster:admin>svctask migratevdisk -vdisk pauline-01 -mdiskgrp sharkc11
```

Note: This command produces no user feedback.

Remember that for volume migration to occur, the extent size of the source and target MDGs must be the same.

Monitor the VDisk migration

Use the **svcinfolsmigrate** command to monitor the progress of the migration.

Example 8-42 shows the output of this command on our system.

Example 8-42 Monitoring the VDisk migration process

```
IBM_2145:ldcluster:admin>svcinfolsmigrate
```

```
migrate_type MDisk_Group_Migration
progress 21
migrate_source_vdisk_index 1435
migrate_target_mdisk_grp 1
max_thread_count 4
```

You can also use the **svcinfolsvdiskextent** command to view the location of the MDisk extents that comprise the VDisk.

Example 8-43 shows the output of this command on our system.

Example 8-43 View the MDisk extents of a VDisks during migration

```
IBM_2145:ldcluster:admin>svcinfolsvdiskextentgeneral01
```

id	number_extents	
0	16	(This was the original image mode VDisk/MDisk)
1	4	
2	4	
3	4	
4	4	

As the VDisk migration proceeds, the number of extents on the original MDisk will decrease. At the same time, the number of extents on the MDisk in the target MDG will increase.

When the volume migration is complete, no extents will remain on the original LUN.

Note: The storage volume now resides on the new storage controller. Access to the virtualized volume is managed through the SAN Volume Controller.

Remove the original volume

After the volume migration is complete the original image mode MDisk remains defined in the original MDG even though it contains no extents. This MDisk can now be removed from the MDG, which changes its status to unmanaged mode.

Use the **svctask rmdisk** command to remove the original MDisk from the MDG.

Example 8-44 shows the output of this command on our system.

Example 8-44 Remove an MDisk from an MDG

```
IBM_2145:ldcluster:admin>svctask rmdisk -mdisk mdisk0 imagegrp
```

At this point the MDisk could be reassigned to another MDG or deleted from the storage controller.

Note: The **rmdisk** command does not delete the MDisk or the LUN that it represents, it only removes it from the MDG and returns it to *unmanaged* mode. If you wish to remove the MDisk entirely you must do so on the storage controller.

Use the **svcinfolsmdisk** command to view the status of the MDisk.

Example 8-45 shows the output of this command on our system.

Example 8-45 Display the properties of a MDisk

```
IBM_2145:ldcluster:admin>svcinfolsmdiskmdisk0
id 0
name mdisk0
status online
mode unmanaged
mdisk_grp_id
mdisk_grp_name
capacity 203.5GB
quorum_index
block_size 512
```

```
controller_name controller0
ctrl_type 4
ctrl_WWNN 50060E8003755400
controller_id 0
path_count 8
max_path_count 8
ctrl_LUN# 0000000000000001
UID 4849544143484920523435313735353430303046000000000000000000000000
preferred_WWPN 50060E8003755400
active_WWPN 50060E8003755400
```

If you wish, you can now delete the original LUN from the original storage controller.

Note: Do not destroy any controller LUNs until you have deleted them from the MDisk group and returned them to *unmanaged* mode.



SAN boot migration to the SVC

In this chapter, we demonstrate the implementation of a SAN Volume Controller in an existing environment with SAN booting Windows and Linux hosts. The existing SAN storage controller was an IBM DS4700.

9.1 Scenario description

In this scenario, we start with an IBM DS4700 storage controller, a SAN connected Windows host, and a SAN connected Linux host. The two hosts are already configured to boot directly from the SAN. We then introduce the SAN Volume Controller to manage the back-end storage. The existing SAN boot volumes on the DS4700 storage controller will be migrated to and managed by the SAN Volume Controller. The existing Windows and Linux hosts will be reconfigured so as to continue to boot from the SVC managed SAN boot volume.

The systems described in this section include:

- ▶ IBM DS4700 storage controller
- ▶ IBM b-type family switches
- ▶ Windows and Linux hosts

Figure 9-1 shows the relationships between the systems used in this scenario.

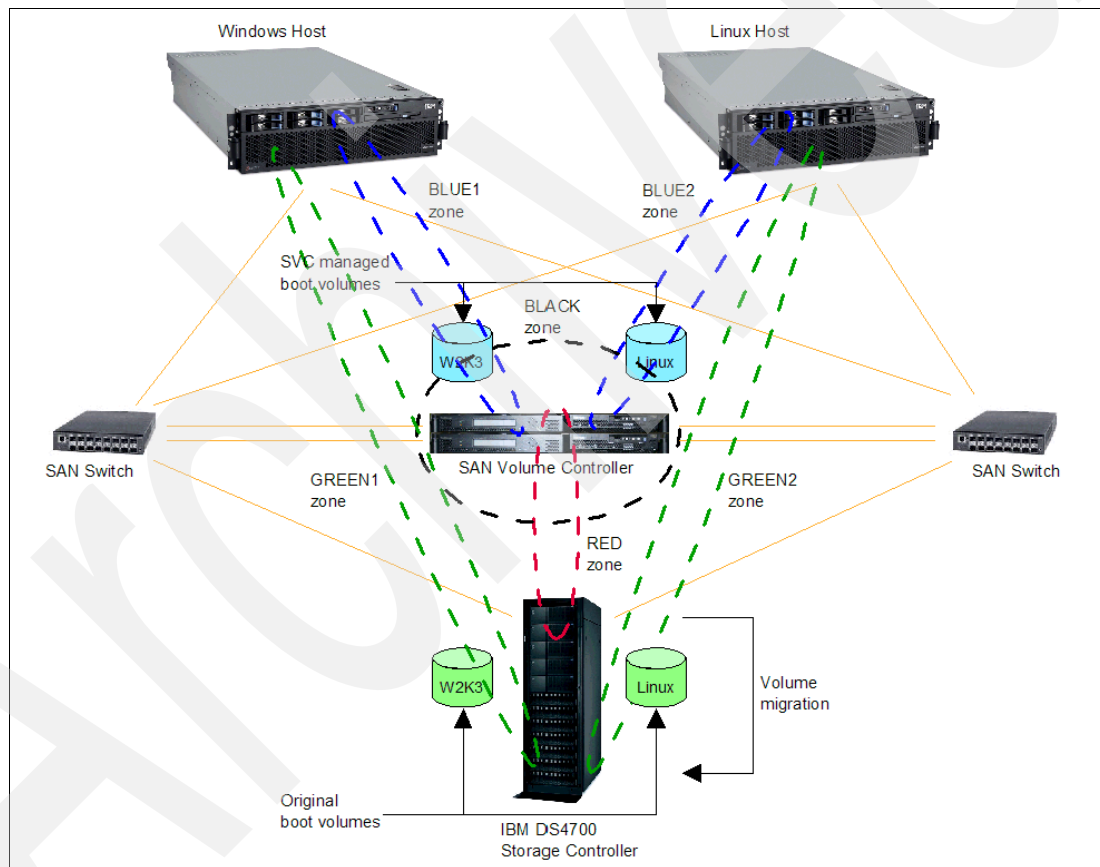


Figure 9-1 SVC implementation example with a SAN booting Windows host

After the boot volume has been migrated to the SVC managed disk, it can then be migrated to any other SVC managed volume or another storage controller with no disruption to the host.

9.1.1 Confirm support for your environment

This section provides information about configuring the SVC in an existing environment with a SAN booting Windows host.

Prior to implementing the SVC, it is important that you understand the existing environment, including the component systems, their interdependencies and business functions.

Windows requirements

The SAN Volume Controller supports boot from SAN for recent versions of Microsoft Windows. There are some restrictions on the type of HBA and multi-path device driver that can be used. In addition, Microsoft publishes some restrictions that are specific to Windows Clustered systems (MSCS) in a SAN boot environment.

Table 9-1 lists the versions of Windows that are supported for SAN booting with the SAN Volume Controller.

Table 9-1 Supported version of Windows using SAN boot with the SVC

Operating System	Multi-path driver	SAN boot support	Clustering support
Windows 2000	SDD 1.6.1.0-2	Yes (See the supported hardware list for specific HBA models)	Yes (See the SVC software restrictions document for information regarding Windows SAN Boot Clusters)
Windows 2003 (32-bit only)	SDD 1.6.1.0-2		
Windows 2003 (32- and 64-bit)	SDDDSM 2.1.1.0-2 (uses MPIO)		

When you configure your Windows system, keep in mind the maximum configurations specific to this platform, as shown in Table 9-2.

Table 9-2 Configuration maximums for Windows

Object	Maximum	Description
VDisk	512 (See Note 1)	The maximum number of VDisks that the SVC can expose to a Windows host
Paths per VDisk	8 max (4 recommended) (See Note 2)	The maximum number of SAN paths to each VDisk
Notes: 1. You can only assign a maximum of 26 drive letters to a Windows host. However, recent Windows versions support submounting drives as directories within other drives. 2. Both SDD and SDDDSM actually support 16 paths per VDisk, but the SVC only supports a maximum of 8 paths (to ensure a reasonable path failover time).		

See the following Web sites for Microsoft white papers on using SAN boot for Windows:

<http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/BootfromSANinWindows.msp>

<http://www.microsoft.com/windowsserver2003/techinfo/overview/san.msp>

Linux requirements

The SAN Volume Controller supports boot from SAN for recent versions of Linux (SLES). There are some restrictions on the type of HBA and multi-path device driver that can be used.

Table 9-3 lists the versions of Linux that are currently supported for SAN booting with the SAN Volume Controller. Consult the IBM support Web site for an updated list:

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002865#_Multi_Host

Table 9-3 Supported version of Linux (x86) using SAN boot with the SVC

Operating System	Multi-path driver	SAN boot support	Clustering support
Red Hat EL 3.0 + 4.0	SDD 1.6.1.0-3	No (One-time support is available on request)	No (One-time support is available on request)
SUSE SLES 8 (x86 or ppc)	SDD 1.6.1.0-3		
SUSE SLES 8 (zSeries®)	LVM version lvm-1.0.5-125	Yes	No
SUSE SLES 9	SDD 1.6.1.0-3	Yes	

When you configure your Linux system, keep in mind the maximum configurations specific to this platform, as shown in Table 9-4.

Table 9-4 Configuration maximums for Linux

Object	Maximum	Description
VDisk	512 (for 2.6 kernels) 256 (for 2.4 kernels)	The maximum number of VDisks that the SVC can expose to a single Linux host
Paths per VDisk	4	The maximum number of SAN paths to each VDisk

Further information

See the following Web sites for specific SAN boot supported Operating System levels:

<http://www.ibm.com/storage/support/2145>

http://www-1.ibm.com/support/docview.wss?rs=591&uid=ssglS1002865#_Multi_Host

It is strongly recommended that you also refer to the following documents for further information regarding performing SAN boot with the SVC:

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_Configuration_Guide.pdf

ftp://ftp.software.ibm.com/storage/san/sanvc/V4.1.0/pubs/English/English_HAG.pdf

http://www-1.ibm.com/support/docview.wss?rs=540&context=ST52G7&q=ssgl*&uid=ssglS7000303&loc=en_US&cs=utf-8&lang=en+en

9.1.2 Preparing the SAN for storage virtualization

Although the process of storage virtualization with the SVC is relatively straightforward, it should not be considered lightly or without adequate preparation. In an existing storage environment it may require coordination of effort between various support teams and an understanding of several technical domains.

Prerequisite conditions

In this scenario, we started with a Windows host that was already configured to boot from SAN on an existing IBM DS4700 storage controller.

Since the installation of the SVC itself is outside the scope of this book, we have started our example at the point where all this preparatory work has been completed.

We have assumed the following initial conditions:

- ▶ The Windows host is already configured to boot from SAN on the DS4700 controller.
- ▶ The DS4700 is already configured to provide SAN storage, and LUN security is enabled.
- ▶ The SVC nodes have been installed and minimally configured as an SVC cluster.
- ▶ The FC switches have been installed, and appropriate zones have been created for the existing environment.

Preparing the SAN fabric

Change the SAN zoning to allow the SVC to access the existing storage controller and the host whose data is to be virtualized.

Refer to Chapter 3, “Storage migration with the SVC” on page 29 for a description of the required SAN zones for implementing the SVC in an existing environment. In our example environment these are the RED and BLUE zones.

If you are planning to migrate the data to a new storage controller (not in our example) then you should also zone the SAN to allow the SVC to access the new storage controller.

Refer to Appendix B, “SAN zoning changes” on page 257 for examples of the commands used to define the required SAN zones on some common vendors’ SAN equipment.

Note: It is assumed that the BLACK zone, which is the SVC node to SVC node zone, was created when creating the SVC cluster and therefore already exists in the SAN.

Preparing the storage controllers

Configure access from the SVC to the LUN on the existing storage controller. You also need to create some LUNs on the storage controller to be your migration targets. The new LUNs can also be used as the SVC quorum disks.

Record the SVC port numbers

Connect to the SVC and record the WWPN for each of the SVC nodes.

Use the **svcinfo lsnode** command to list the SVC nodes.

Example 9-1 shows the output of this command on our system.

Example 9-1 List the SVC nodes

```
IBM_2145:ldcluster:admin>svcinfo lsnode
```

id	name	UPS_serial_number	WWNN	status	IO_group_id
0	node1	100062L102	50050768010008CD	online	0
1	node2	100062L097	5005076801001013	online	0

To obtain required additional information use the **svcinfo lsnode nodename** command to list the properties of each of the SVC nodes.

Example 9-2 on page 242 shows the output of this command for one node on our system.

Example 9-2 List the SVC node details

```
IBM_2145:ldcluster:admin>svcinfolsnode node1

id 0
name node1
UPS_serial_number 100062L102
WWNN 50050768010008CD          (WWNN for node1)
status online
IO_group_id 0
IO_group_name io_grp0
partner_node_id 1
partner_node_name node2
config_node no
UPS_unique_id 2040000182701002
port_id 50050768014008CD        (WWPN for port #4)
port_status active
port_speed 4Gb
port_id 50050768013008CD        (WWPN for port #3)
port_status active
port_speed 2Gb
port_id 50050768011008CD        (WWPN for port #1)
port_status active
port_speed 2Gb
port_id 50050768012008CD        (WWPN for port #2)
port_status active
port_speed 2Gb
hardware 8F4
```

Repeat this process for each of the SVC nodes in your SVC cluster. Note that the WWPN numbers can be determined from the SVC node's WWNN.

Configure the existing storage controller

You need to create one or more LUNs on the DS4700 storage controller. These LUNs will later be used as the target volumes for your migration from the SVC managed *image mode* boot volumes to SVC managed *striped mode* volumes.

The general recommendation for SVC use is that each of these LUNs be configured to the full capacity of the containing RAID array. Therefore, if possible, create one LUN per RAID array.

The creation of the new LUNs on the DS4700 is not shown here.

You also need to change the LUN masking on the storage controller to allow the SVC to access the new LUNs.

The administrative interface for the IBM DS4700 is a stand-alone application called *Storage Manager*.

Figure 9-2 shows an example of the IBM DS4700 Storage Manager administrative interface being used to define LUN mappings for access by the SVC to the target volumes.

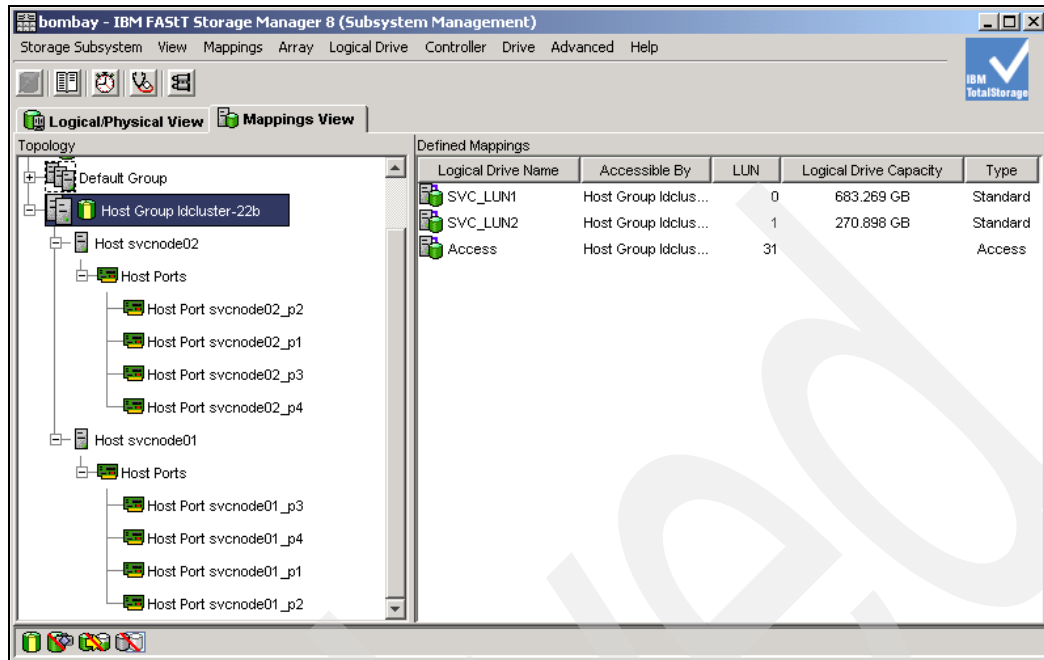


Figure 9-2 LUN mapping on the IBM DS4700

Refer to the IBM DS4000 administration guide for detailed instructions on how to configure LUN masking on the DS4700 storage controller.

Preparing the SAN Volume Controller

Confirm that the SVC can see the storage controllers and the LUNs that are to be virtualized.

Use the `svcinfolscntroller` command to list the storage controllers.

Example 9-3 shows the output of this command on our system.

Example 9-3 List the storage controllers

```
IBM_2145:ldcluster:admin>svcinfolscntroller
```

id	controller_name	ctrl_s/n	vendor_id	product_id_low	product_id_high
0	bombay		IBM	1742	

If the storage controller is not visible to the SVC, then you should double check the SAN zoning.

Create an MDG for the new storage controller's LUNs

Since you have already correctly zoned the SAN and configured LUN masking on the back-end storage controllers, the LUNs should already be visible to the SVC. However, if the SVC has not yet detected the LUNs, you may need to manually force their discovery.

Use the `svctask detectmdisk` command to rescan for newly available MDisk.

Note: The command `svctask detectmdisk` produces no user feedback.

If the expected MDisk are still not visible to the SVC, then you should double check the SAN zoning and storage controller LUN masking.

Use the **svcin** **lsmdisk** command to list the MDisk on the existing storage controller.

Example 9-4 shows the output of this command on our system.

Example 9-4 List MDisk on the existing storage controller

```
IBM_2145:ldcluster:admin>svcin lsmdisk -delim ,

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
0,mdisk0,online,unmanaged,,,683.3GB,0000000000000000,bombay,600...000
1,mdisk1,online,unmanaged,,,270.9GB,0000000000000001,bombay,600...000
```

This output has been edited for legibility. Note that the MDisk for the existing LUN is in *unmanaged* mode.

Use the **svctask mkmdiskgrp** command to create a new MDG and add the unmanaged MDisk to the MDG.

Example 9-5 shows the output of this command on our system.

Example 9-5 Add the new controller's LUNs to a MDG

```
IBM_2145:ldcluster:admin>svctask mkmdiskgrp -name mgrp0 -ext 256 -mdisk mdisk0:mdisk1

MDisk Group, id [0], successfully created
```

Use the **svctask lsmdisk** command to check that the MDG was created successfully.

Example 9-6 shows the output of this command on our system.

Example 9-6 List the MDisk that compose a given MDG

```
IBM_2145:ldcluster:admin>svcin lsmdisk -delim ,

id,name,status,mode,mdisk_grp_id,mdisk_grp_name,capacity,ctrl_LUN_#,controller_name,UID
0,mdisk0,online,managed,0,mgrp0,683.3GB,0000000000000000,bombay,600...000
1,mdisk1,online,managed,0,mgrp0,270.9GB,0000000000000001,bombay,600...000
```

This output has been edited for legibility. Note that the MDisk for the new LUNs are now in *managed* mode.

Define a host to the SVC

Before the host can access any VDisk, it needs to be defined to the SVC.

Use the **svcin lshbaportcandidate** command to list any unassigned host WWPNs that are visible to the SVC:

Example 9-7 shows the output of this command on our system.

Example 9-7 List unassigned host WWPNs that are visible to the SVC

```
IBM_2145:ldcluster:admin>svcin lshbaportcandidate

id
210000E08B0515D4      (This is the first HBA on our host)
210100E08B2515D4      (This is the second HBA on our host)
```

Note: You will need to check on the host (or the SAN switch) to determine which WWPN belongs to the host with which you are working.

Use the **svctask mkhost** command to define the host to the SVC.

As an interim step, while configuring the SAN boot to the SVC managed LUN, you need to define a single path to the boot LUN. Therefore you need to define the new host with only a single HBA. It is important to define the same HBA that was included in the SAN zone that you defined earlier.

Example 9-8 shows the output of this command on our system.

Example 9-8 Define a host to the SVC

```
IBM_2145:ldcluster:admin>svctask mkhost -name marie -hbawwn 210000E08B0515D4
```

```
Host, id [0], successfully created
```

Use the **svcinfolshost** command to list all hosts and verify that the host was defined successfully.

Example 9-9 shows the output of this command on our system.

Example 9-9 Check a host definition on the SVC

```
IBM_2145:ldcluster:admin>svcinfolshost marie
```

```
id 0
name marie
port_count 1
type generic
mask 1111
iogrp_count 4
WWPN 210000E08B0515D4
node_logged_in_count 0
state inactive
```

The preparation is now complete and you are ready to proceed with the storage virtualization.

9.1.3 Move SVC managed volumes between storage controllers

The SAN volume controller enables the migration of a storage volume between storage controllers with no disruption to the host systems.

Configure the SAN Volume Controller

In our example, we will migrate an image mode VDisk between an MDG on one storage controller and another MDG on the new storage controller. During the migration, the image mode VDisk is automatically converted to striped mode.

Migrate the VDisk to the new storage controller

Use the **svctask migratevdisk** command to start the volume migration.

Example 9-10 shows the output of this command on our system.

Example 9-10 Migrating a VDisk between MDGs

```
IBM_2145:ldcluster:admin>svctask migratevdisk -vdisk marie_boot -mdiskgrp mgrp0
```

Note: This command produces no user feedback.

Remember that for volume migration to occur, the extent size of the source and target MDGs must be the same.

Monitor the VDisk migration

Use the **svcinio lsmigrate** command to monitor the progress of the migration.

Example 9-11 shows the output of this command on our system.

Example 9-11 Monitoring the VDisk migration process

```
IBM_2145:ldcluster-22b:admin>svcinio lsmigrate
```

```
migrate_type MDisk_Group_Migration
progress 22
migrate_source_vdisk_index 0
migrate_target_mdisk_grp 0
max_thread_count 4
```

You can also use the **svcinio lsvdiskextent** command to view the location of the MDisk extents that comprise the VDisk.

Example 9-12 shows the output of this command on our system.

Example 9-12 View the MDisk extents of a VDisks during migration

```
IBM_2145:ldcluster:admin>svcinio lsvdiskextent marie_boot
```

id	number_extents
0	3
1	2
2	35 (This was the original image mode VDisk/MDisk)

As the VDisk migration proceeds, the number of extends on the original MDisk will decrease. At the same time, the number of extents on the MDisk in the target MDG will increase.

When the volume migration is complete, no extents will remain on the original LUN.

Note: The storage volume now resides on the new storage controller. Access to the virtualized volume is managed through the SAN Volume Controller.

Remove the original volume

After the volume migration is complete, the original image mode MDisk remains defined in the original MDG, even though it contains no extents. This MDisk can now be removed from the MDG, which changes its status to unmanaged mode.

Use the **svctask rmvdisk** command to remove the original MDisk from the MDG.

Example 9-13 shows the output of this command on our system.

Example 9-13 Remove an MDisk from an MDG

```
IBM_2145:ldcluster:admin>svctask rmmdisk -mdisk mdisk2 imagegrp
```

At this point, the MDisk could be reassigned to another MDG or deleted from the storage controller.

Note: The `rmmdisk` command does not delete the MDisk or the LUN that it represents; it only removes it from the MDG and returns it to *unmanaged* mode. If you wish to remove the MDisk entirely, you will need to do so on the storage controller.

Use the `svcinfo lsmdisk` command to view the status of the MDisk.

Example 9-14 shows the output of this command on our system.

Example 9-14 Display the properties of a MDisk

```
IBM_2145:ldcluster-22b:admin>svcinfo lsmdisk mdisk2
```

```
id 2
name mdisk2
status online
mode unmanaged
mdisk_grp_id
mdisk_grp_name
capacity 10.0GB
quorum_index
block_size 512
controller_name bombay
ctrl_type 4
ctrl_WWNN 200200A0B80F0621
controller_id 0
path_count 2
max_path_count 2
ctrl_LUN_# 0000000000000002
UID 600a0b80000f02c70000003244be1d010000000000000000000000000000000000
preferred_WWPN 200300A0B80F0622
active_WWPN 200300A0B80F0622
```

If you wish, you can now delete the original LUN from the original storage controller.

Note: Do not destroy any controller LUNs until you have deleted them from the MDisk group and returned them to *unmanaged* mode.

9.2 Unsupported OEM controllers

In some circumstances, you may find that one of your storage controller subsystems is not supported with SAN Volume Controller. This does not imply that you will be unable to implement a virtualized environment.

There are two different scenarios:

- ▶ You have a storage controller that you wish to connect to SAN Volume Controller in order to migrate the data from it before it is decommissioned.

- You have a storage controller that you want to continue to use under a virtualized environment with SAN Volume Controller.

For both of these situations, you should use the RPQ process to request support for your controller.

In the event that you are just looking for support during a data migration process, it is likely that the RPQ will be approved on the basis that you only connect a single port to the SAN. This is with the assumption that there are no other known limitations for connection of your storage controller to SAN Volume Controller.

For ongoing support of your controller, the RPQ will be reviewed and a decision will be made on whether to validate the system with SAN Volume Controller. The testing of new controller types is very rigorous, in order to ensure that the system performs with SAN Volume Controller under all situations, especially during error conditions (for example, faulty SFPs or fibres).

9.3 Host cluster migration

All the examples discussed in chapters 4 through 9 have focused specifically on single host systems. You may also have clustered host systems that you want to migrate under SAN Volume Controller.

If this is the case, you should carry out the migration process as defined in the relevant chapter. For all the host operations you should carry them out on all the hosts in the cluster at the same time. For example, if you are instructed to shut down the host, shut down the cluster.

If you modify each cluster host one at a time, there is a risk that the view they share of the underlying disk storage will be inconsistent. This may cause unexpected behavior and may result in the need for recreating the cluster.



Troubleshooting

This chapter provides examples of some problems that we encountered while writing this book. It also suggest possible ways to resolve these types of issues as well as commonly encountered problems when implementing the SVC for the first time.

10.1 Overview

SAN environments can be quite complex to debug when things go wrong. To aid in making the environment as simple as possible, try to stick to these rules:

- ▶ Use nicknames for the worldwide port names wherever possible.
- ▶ Make the SAN zone names meaningful.
- ▶ SAN zone configuration:
 - Keep host zones small, with one host per zone.
 - Have each storage controller individually zoned with the SAN Volume Controller nodes.
 - Always have a zone that contains only the SAN Volume Controller nodes.
 - Double check any manually entered worldwide port names.

10.2 Solving errors

This section describes some of the more commonly encountered errors and how to solve them.

10.2.1 SAN Volume Controller

The SAN Volume Controller clearly displays errors on both the front panel and through the error log. Errors are identified as either three or four digit numbers. Rather than explaining each of the possible errors, we suggest that you have a copy of the service guide at hand when configuring the system. This document provides descriptions for each of the error codes and recommended procedures to resolve the problem. To get access to the latest service guide, go to the following Web site and look in the documentation section:

<http://www.ibm.com/storage/support/2145>

10.2.2 AIX

Some common problems encountered on AIX and their solutions are as follows:

- ▶ The host WWPNs are not visible to SAN Volume Controller when creating a new host and assigning ports.
 - Run **cfgmgr** on the host. This will cause the host to log back into the switch.
- ▶ The host can see the LUNs being presented by SAN Volume Controller, but they are identified as *Other FC disk*.
 - It is likely that the host attachment script is not installed. Make sure that the `devices.fcp.disk.ibm.rte.tar` file is installed.
- ▶ The host can see the LUNs presented by SAN Volume Controller and they are identified as IBM 2145 devices. There are no vpath devices.
 - It is likely that the Subsystem Device Driver module is not installed. Make sure that the `devices.sdd.xx.rte.tar` file is installed.

10.2.3 HP-UX

Some common problems encountered on HP-UX and their solutions are as follows:

- ▶ The host WWPNS are not visible to SAN Volume Control when creating a new host and assigning ports.
 - HP-UX does not log its ports into a device unless it is accessing it, so in many instances, the WWPNS will not be listed under SAN Volume Controller. Enter them manually when configuring the host.
- ▶ The host can see the LUNs presented by SAN Volume Controller and they are identified as IBM 2145 devices. There are no vpath devices.
 - Have you run the **cfgvpath** command?
 - Is the Subsystem Device Driver module installed?

10.2.4 EMC CLARiiON

A common problem encountered on EMC CLARiiON and the solution is as follows:

- ▶ Although you have presented some LUNs to the SAN Volume Controller, not all of them are visible.
 - Check that you have run the **detectmdisk** command on the SAN Volume Controller.
 - Within the Navisphere GUI, de-allocate the LUNs from the SAN Volume Controller and then re-allocate them.

10.2.5 HDS Thunder

A common problem encountered on HDS Thunder and the solution is as follows:

- ▶ You are unable to modify the LUN parameters as specified for SAN Volume Controller. The values are greyed out on the management GUI.
 - You must de-allocate all LUNs from the host system before you can change their parameters.

Archived



SVC administration commands

This appendix provides a reference to the SVC administration commands.

SAN Volume Controller commands

Table A-1 shows the SVC administration commands grouped by function.

Table A-1 SVC commands

Task	svcinfo	svctask
Cluster configuration	lscluster lsclustercandidate lssshkeys lstimezones showtimezone lsnode lsnodecandidate lsnodevpd lsiogrp lsiogrpcandidate lsiogrpghost lslicense lsdiscoverystatus lscopystatus	chcluster addsshkeys rmsshkey rmallsshkeys setclustertime settimezone setpwdreset addnode chnode rmnode chiogrp chlicense stopcluster
Work with storage controllers	lscontroller lsfabric	chcontroller
Work with hosts	lshbaportcandidate lshost lshostiogrp lshostvdiskmap	mkhost chhost rmhost addhostport rmhostport addhostiogrp rmhostiogrp
Work with MDisk	lsmdisk lsmdiskmember lsmdiskextent lsfreeextents	detectmdisk chmdisk includemdisk setquorum
Work with MDisk groups	lsmdiskgrp lsmdiskcandidate	addmdisk rmdisk mkmdiskgrp chmdiskgrp rmdiskgrp

Task	svcinfo	svctask
Work with VDisks	lsvdisk lsvdiskhostmap lsvdiskextent lsvdiskmember lsvdiskprogress	mkvdisk chvdisk rmvdisk mkvdiskhostmap rmvdiskhostmap expandvdisksize shrinkvdisksize
Migrate VDisks (extents)	lsmigrate	migrateexts migratetoimage migratevdisk
Work with flashcopy	lsfcconsistgrp lsfcmap lsfcmapcandidate lsfcmapprogress	mkfcconsistgrp chfcconsistgrp rmfcconsistgrp prestartfcconsistgrp startfcconsistgrp stopfcconsistgrp mkfcmap chfcmap rmfcmap prestartfcmap startfcmap stopfcmap
Work with metro/global mirror	lsrconsistgrp lsrcrelationship lsrcrelationshipcandidate lsrcrelationshipprogress	mkrconsistgrp chrconsistgrp rmrconsistgrp mkrcrelationship chrcrelationship rmrcrelationship startrcrelationship stoprcrelationship switchrcrelationship startrcconsistgrp stoprcconsistgrp switchrcconsistgrp mkpartnership chpartnership rmpartnership

Task	svcinfo	svctask
Troubleshooting and reporting	catauditlog caterrlog caterrlogbyseqnum ls2145dumps lsauditlogdumps lsioatdumps lsioatdumps lssoftwaretdumps lsfeaturetdumps lserrlogbyfcconsistgrp lserrlogbyfcmap lserrlogbyhost lserrlogbyiogrp lserrlogbydisk lserrlogbydiskgrp lserrlogbynode lserrlogbyrcconsistgrp lserrlogbyrcrelationship lserrlogbyvdisk lserrlogtdumps	startstats stopstats starttrace stoptrace settrace setdisktrace setevent cleardumps clearerrlog cpdumps dumperrlog dumpinternallog dumppauditlog finderr cherrstate enablecli

Using the -? parameter will allow you to obtain online help for the various SVC commands, as shown in Example A-1:

Example: A-1 SVC command line help example

```
IBM_2145:ldcluster:admin>svctask mkhost -?
mkhost
```

The mkhost command creates a logical host object.

Syntax

```
>>- svctask -- -- mkhost -- --+-----+-- >>
                                '- -name -- new_name_arg -'

>-- -hbawpn -- wwpn_list -- --+-----+-- >
                                '- -iogrp--iogrp_list -'

>-- --+-----+-- --+-----+-- >
                                '- -mask--host_port_mask -'      '- -force -'

>--+-----+-- ><
    '- -type --+ hpux ----+-'
        '- generic -'
```

The example shown here has been truncated for readability.

SAN zoning changes

In this appendix we describe how to perform the necessary zoning changes when integrating the SVC in an existing environment. We show how to change the zoning to remove host-to-storage controller access, and instead enable the SVC to access the storage controller and the host to access the SVC. We provide examples of how to do this using the command line interface. We use different interfaces to demonstrate how it is done.

Note: We recommend that you consult *IBM System Storage: Implementing an IBM SAN*, SG24-6116 for a comprehensive description of how to set up and operate SAN switches and directors.

<http://www.redbooks.ibm.com/redpieces/abstracts/sg246116.html?Open>

B.1 Basic steps before zoning

We highly recommend that you ensure your environment complies with the following checklist prior to making any changes. All of the following should be true:

- ▶ Your environment is supported.
- ▶ There are no pending errors on any of the switches.
- ▶ There are no unfixed errors on your SVC cluster.
- ▶ The SVC cluster is not degraded.
- ▶ There are no errors on any of the storage controllers.
- ▶ You have taken into consideration the number of hosts used in your setup and planned how you want your zoning layout.

Before you change the zoning for a particular host when integrating the SVC into the environment, you must ensure that the host does not have any volumes mounted that reside on LUNs on the storage controller that you are about to integrate. Be sure you can shut down the particular host while you change the zoning.

Note: It is technically possible to allow the host to be zoned to the SVC and the storage at the same time. We do *not* recommend this. At any given time a LUN on a storage controller should not be accessible for both the host and the SVC because you risk corrupting the data on that LUN.

In general, most SAN environments consist of redundant fabrics, mostly dual fabric designs. In the following examples, we only show the steps to be performed for one of the fabrics in a dual fabric environment. The same steps must be performed on the other fabrics if not already configured.

Figure B-1 shows the setup.

The following applies for all our scenarios:

- ▶ All the zoning in Fabric 1 is complete.
- ▶ The zone containing all the ports of the SVC nodes in each fabric is already made.
- ▶ The aliases for the different nodes in the fabric are already created, since they are already in use.

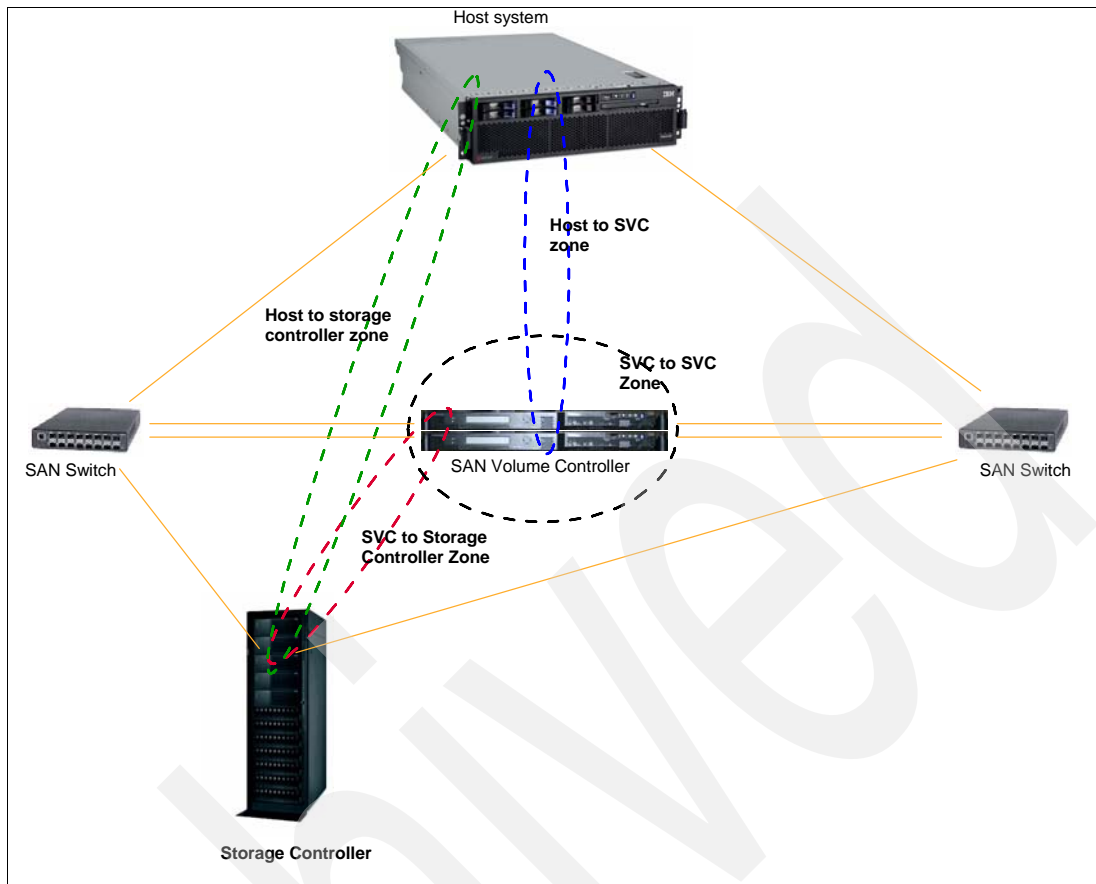


Figure B-1 The setup

B.2 Basic zoning steps

When changing the zoning, there are some basic steps that we always follow:

1. Create zones that contains the ports for:
 - a. SVC to storage
 - b. SVC to host
2. Remove the host to storage zone
3. Add the SVC to storage zone (svc_storage)
4. Add the SVC to host zone (svc_host)
5. Add the SVC zone

Important: It is good practice to create a zone that only contains all the SVC nodes ports in a fabric. This is to ensure that SVC node to node communication persists regardless of other zone changes, and to avoid nodes losing contact with one another as a result of misconfiguration.

B.2.1 Using the CLI for the IBM b-type family

In our scenario, we have a host connected to a storage controller. We bring the SVC into the environment and show how the transition will be handled from a zoning point of view.

In the first example, we telnet to an IBM 2109-B32 and use the CLI to configure the zoning.

Logging in to the switch

Telnet to the ip address (or DNS name) and enter the login credentials to log in to the switch, shown in Example B-1.

Example B-1: Login to the b-type SAN switch.

Fabos Version 5.1.0b

```
itsosan02 login: admin
Password:
itsosan02:admin>
```

Make the zones

Next we create the desired zones. We recommend that you use aliases and a simple naming standard that describes each alias, as well as descriptive names for the zones.

To create an alias, we use the command **alcreate**. Example B-2 shows the command we used to create the alias for hba2 on host1.

Example B-2: Creating an alias

```
itsosan02:admin> alcreate Host1_HBA2, 21:00:00:e0:8b:06:c5:0b
```

In our example, we created an alias for each SVC node port, using the command **alcreate**, as well as an alias for each storage controller port and host controller port. Example B-3 shows the defined aliases in fabric 2.

Example B-3: Defined aliases

```
alias: ITSOSVC01_N1_P2
      50:05:07:68:01:20:18:8e
alias: ITSOSVC01_N1_P4
      50:05:07:68:01:40:18:8e
alias: ITSOSVC01_N2_P2
      50:05:07:68:01:20:18:83
alias: ITSOSVC01_N2_P4
      50:05:07:68:01:40:18:83
alias: DS4500_CTLA_P2
      20:04:00:a0:b8:17:44:32
alias: DS4500_CTLB_P2
      20:05:00:a0:b8:17:44:32
alias: Host1_HBA2
      21:00:00:e0:8b:06:c5:0b
```

Attention: Always ensure that you have added the correct aliases to the zone. If wrong aliases are added to a zone, it could compromise your data and lead to the loss of data.

To allow the SVC to access the storage controller, we create an *SVC to Storage* zone with the SVC node ports and the storage controller ports as members, using the **zonecreate** command as shown in Example B-4.

Example B-4: Creating the SVC to Storage zone

```
itsosan02:admin> zonecreate SVC_to_DS4500, "ITS0SVC01_N1_P2; ITS0SVC01_N1_P2;  
ITS0SVC01_N2_P2; ITS0SVC01_N2_P4; DS4500_CTLA_P2; DS4500_CTLB_P2"
```

To allow the host to ccess the SVC node ports, we create a *Host to SVC* zone with the host port, and in this scenario port 2 on the SVC nodes, shown in Example B-5.

Example B-5: Creating the Host to SVC zone

```
itsosan02:admin> zonecreate Host1_to_SVC, "Host1_HBA2; ITS0SVC01_N1_P2;  
ITS0SVC01_N2_P2"
```

Note: For advice on selecting the number of paths for host to SVC access consult:

<http://www.redbooks.ibm.com/abstracts/sg246423.html?0open>

Saving and activating the zoning

We have now created the zones needed, but they are not saved or activated yet. We must add the new zones to the configuration and save the configuration before we can activate the new zones. Example B-6 shows how the defined configuration is saved.

Example B-6: Saving the defined configuration

```
itsosan02:admin> cfigsave  
You are about to save the Defined zoning configuration. This  
action will only save the changes on Defined configuration.  
Any changes made on the Effective configuration will not  
take effect until it is re-enabled.  
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y  
Updating flash ...  
itsosan02:admin>
```

Note that the only *active* zone for the host at the moment is the zone connecting the host to the storage. Next, we remove the old zone and activate the new host to the SVC zone.

Now that we have saved the configuration, we can remove the old host to storage zone.

Attention: Your host should be shut down at this point (or volumes defined on the LUNs should be unmounted).

We start by removing the zone for the host and storage controller as shown in Example B-7.

To remove a zone from a configuration we use the command **cfigRemove**.

Example B-7: Removing the host to storage controller zone

```
itsosan02:admin> cfigRemove "ITS0_Fabric2", "Host1_DS4500"
```

Now that we have removed the old host to storage zone, we save and enable the configuration for the changes to apply.

Example B-8: Saving and activating the zone configuration

```
itsosan02:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
Nothing changed: nothing to save, returning ...
itsosan02:admin> cfgenable
error: Usage: cfgenable cfgName
itsosan02:admin> cfgenable ITS0_Fabric2
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'ITS0_Fabric2' configuration (yes, y, no, n): [no] y
zone config "ITS0_Fabric2" is in effect
Updating flash ...
itsosan02:admin>
```

Note: Before committing the configuration to the switch, make sure that you have only removed or added the zones that you needed. Removing or adding the wrong zones could lead to the loss of connection and data, or allow different hosts to have access to the same data.

This completes the zoning configurations for the described scenario.

B.2.2 Using the CLI for the IBM m-type family

In our scenario, we have a host connected to a storage controller. We will bring the SVC into the environment and show how the transition will be handled from a zoning point of view.

We use an IBM SAN32M-2 to demonstrate the basic zoning steps needed when migrating.

Logging in to the switch

The first step is to log in to the switch using telnet as shown in Example B-9.

Example B-9: Logging in using Telnet

```
c:\telnet 9.71.50.36
```

```
Username: Administrator
password:
```

Be aware that the m-type switch has a hierarchical interface.

Making the zones

Since it is not possible to create aliases for the m-type switch, we have to have the worldwide port name (WWPN) for the following before we start zoning:

- ▶ SVC nodes port
- ▶ Storage Controllers
- ▶ Host system HBAs

We then create the zones needed using the command **addzone "zone name"**.

For that we need to go into *config zoning* mode as demonstrated in Example B-10.

Example B-10: Creating the zones

```
Root> config
Config> zoning
Config.Zoning>
Config.Zoning> addzone svc_storage
Config.Zoning> addzone svc_host
```

After we have created the zones, we need to add the WWPN for the different nodes. To display the different nodes that are logged in to the switch, we use the *nameserver* options as displayed in Example B-11.

Example B-11: Displaying the nodes logged in to the switch

```
Root>show
show>nameserver
```

We now have to add the WWPN to the zones. We do this using the **addWwnMem** command.

The switch will have to be in **Config.Zoning** mode.

Example B-12 shows adding the WWPN for the SVC, Storage controller and SVC, and host zones.

Important: Be careful when adding the WWPN. There are no warnings or limitations when using the CLI, so double check that you have used the correct WWPN.

Example B-12: Creating the zones

```
Config.Zoning> addWwnMem svc_storage 50:05:07:68:01:20:08:cd
Config.Zoning> addWwnMem svc_storage 50:05:07:68:01:10:08:cd
Config.Zoning> addWwnMem svc_storage 50:05:07:68:01:20:10:13
Config.Zoning> addWwnMem svc_storage 50:05:07:68:01:10:10:13
Config.Zoning> addWwnMem svc_storage 20:14:00:A0:B8:0B:DE:24
Config.Zoning> addWwnMem svc_host 21:01:00:E0:8B:2B:79:34
Config.Zoning> addWwnMem svc_host 50:05:07:68:01:20:08:cd
Config.Zoning> addWwnMem svc_host 50:05:07:68:01:10:08:cd
Config.Zoning> addWwnMem svc_host 50:05:07:68:01:20:10:13
Config.Zoning> addWwnMem svc_host 50:05:07:68:01:10:10:13
```

The next step is to remove the old zone and add the new zones.

Saving and activating the zoning

We remove the *host_storage* zone and add *svc_host* and *svc_storage*.

The m-type switch uses an active and pending area. All zones are displayed in the pending area. The pending area becomes the active zoneset when the **activatezoneset** command is issued. When a zone is removed from the pending area, it will not be removed from the active area before the pending is made active. The pending is also called the workarea and works as the saved configuration, so you have to be very careful when removing zones from it.

Example B-13: Activating the zones

```
Config.Zoning> showactive
Default Zone Enabled:   False
Zone Set:  NEW_ZONE_SET
  Zone: host_storage
    Zone Member: 21:01:00:E0:8B:2B:79:34
    Zone Member: 20:14:00:A0:B8:0B:DE:24
  Zone: svc
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13

Config.Zoning> showpending
Zone Set:  NEW_ZONE_SET
  Zone: svc
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13
  Zone: host_storage
    Zone Member: 21:01:00:E0:8B:2B:79:34
    Zone Member: 20:14:00:A0:B8:0B:DE:24
  Zone: svc_storage
    Zone Member: 20:14:00:A0:B8:0B:DE:24
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13
  Zone: svc_host
    Zone Member: 21:01:00:E0:8B:2B:79:34
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13
Config.Zoning> deletezone host_storage
Config.Zoning> activateZoneSet
Config.Zoning> showactive
Zone Set:  NEW_ZONE_SET
  Zone: svc
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13
  Zone: svc_storage
    Zone Member: 20:14:00:A0:B8:0B:DE:24
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
    Zone Member: 50:05:07:68:01:10:10:13
  Zone: svc_host
    Zone Member: 21:01:00:E0:8B:2B:79:34
    Zone Member: 50:05:07:68:01:20:08:cd
    Zone Member: 50:05:07:68:01:10:08:cd
    Zone Member: 50:05:07:68:01:20:10:13
```

This completes the zoning configurations.

B.2.3 Using the CLI for Cisco

In this scenario we have a host connected to a storage controller. We will bring the SVC into the environment and show how the transition will be handled from a zoning point of view.

We show the basic commands using an MDS 9120.

Logging in to the switch

The first step is to log in to the switch using telnet, so that we can administer the switch.

Note that Cisco refers to a worldwide port name (WWPN) using port worldwide name (pWWN).

Example B-14: Login

```
alpha login: admin
Password: xxxxxx
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2004, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
alpha#
```

Creating the zones

We start by looking at the active zoneset using the command **show zoneset name "name of zoneset"**. Example B-15 shows the output in our example.

Example B-15: Examining the current active zoneset

```
cisco10# show zoneset name firebird
zoneset name test vsan 1
  zone name host_storage vsan 1
    fcalias name storage vsan 1
      pwwn 20:14:00:a0:b8:0b:de:24

    fcalias name host vsan 1
      pwwn 21:01:00:e0:8b:2b:79:34

  zone name svc vsan 1

    fcalias name node1port1 vsan 1
      pwwn 50:05:07:68:01:10:08:cd

    fcalias name node1port2 vsan 1
      pwwn 50:05:07:68:01:20:08:cd

    fcalias name node2port1 vsan 1
```

```
pwwn 50:05:07:68:01:10:10:13
```

```
fcalias name node2port2 vsan 1  
pwwn 50:05:07:68:01:20:10:13
```

You can see that there already is a zone containing all the SVC nodes ports. We will use this zone again, since it is considered good practice.

We now create the new zoneset and zones. We operate in *config*, *config-zoneset* and *config-zone* mode using the commands **zoneset name "name of zoneset"**, **vsan "vsan number between 1 and 4096"**, **zone name "name of zone"**, **vsan "vsan number matching zoneset vsan number"** and **member fcalias "name of alias"**.

Next, we make the zones *svc_storage* and *svc_host* as shown in Example B-16.

Example B-16: creating the zones

```
cisco10(config)# zoneset name bigblue vsan 1  
cisco10(config-zoneset)# zone name svc_storage vsan 1  
cisco10(config-zone)# member fcalias storage  
cisco10(config-zone)# member fcalias node1port1  
cisco10(config-zone)# member fcalias node1port2  
cisco10(config-zone)# member fcalias node2port1  
cisco10(config-zone)# member fcalias node2port2  
cisco10(config-zone)# zone name svc_host vsan 1  
cisco10(config-zone)# member fcalias node1port1  
cisco10(config-zone)# member fcalias node1port2  
cisco10(config-zone)# member fcalias node2port1  
cisco10(config-zone)# member fcalias node2port2  
cisco10(config-zone)# member fcalias host
```

```
cisco10(config-zoneset)# member svc
```

Saving and activating the zoning

The zones are saved in the zoneset. So no further steps needs to be taken to save the configuration; we only need to activate the configuration. There can only be one *zoneset* active for each *vsan* at any time, so we do not need to remove the existing zoneset *firebird*.

We are in *config* mode using the **zoneset activate name "name of zoneset" vsan 1** command.

We show activating the configuration in Example B-17.

Example B-17: Activating the zoneset

```
cisco10(config)# zoneset activate name bigblue vsan 1  
  
cisco10# show zoneset active vsan 1  
zoneset name bigblue vsan 1  
  zone name svc_storage vsan 1  
    pwwn 50:05:07:68:01:10:08:cd  
    pwwn 50:05:07:68:01:20:08:cd  
    pwwn 50:05:07:68:01:10:10:13  
    pwwn 50:05:07:68:01:20:10:13  
    pwwn 20:14:00:a0:b8:0b:de:24
```

```
zone name svc_host vsan 1
  pwwn 50:05:07:68:01:10:08:cd
  pwwn 50:05:07:68:01:20:08:cd
  pwwn 50:05:07:68:01:10:10:13
  pwwn 50:05:07:68:01:20:10:13
  pwwn 21:01:00:e0:8b:2b:79:34
```

```
zone name svc vsan 1
  pwwn 50:05:07:68:01:10:08:cd
  pwwn 50:05:07:68:01:20:08:cd
  pwwn 50:05:07:68:01:10:10:13
  pwwn 50:05:07:68:01:20:10:13
```

The zoning is now complete.

It would now be safe to remove the zoneset *firebird*.

Archived

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 270. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM System Storage SAN Volume Controller*, SG24-6423
- ▶ *Get More Out of Your SAN with IBM Tivoli Storage Manager*, SG24-6687
- ▶ *IBM Tivoli Storage Area Network Manager: A Practical Introduction*, SG24-6848
- ▶ *Virtualization in a SAN*, REDP-3633
- ▶ *SAN: What is a VSAN?*, TIPS0199

Other publications

- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Planning Guide*, GA22-1052
- ▶ *IBM System Storage Master Console: Installation and User's Guide*, GC30-4090
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Installation Guide*, SC26-7541
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Service Guide*, SC26-7542
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Configuration Guide*, SC26-7543
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Command-Line Interface User's Guide*, SC26-7544
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: CIM Agent Developers Reference*, SC26-7545
- ▶ *IBM System Storage Multipath Subsystem Device Driver User's Guide*, SC30-4096
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Host Attachment Guide*, SC26-7563

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ IBM System Storage home page:
<http://www.storage.ibm.com>
- ▶ SAN Volume Controller supported platform:
<http://www-1.ibm.com/servers/storage/support/software/sanvc/index.html>

- ▶ Download site for Windows SSH freeware:
<http://www.chiark.greenend.org.uk/~sgtatham/putty>
- ▶ IBM site to download SSH for AIX:
<http://oss.software.ibm.com/developerworks/projects/openssh>
- ▶ Open source site for SSH for Windows and Mac:
<http://www.openssh.com/windows.html>
- ▶ Cygwin Linux-like environment for Windows:
<http://www.cygwin.com>
- ▶ IBM Tivoli Storage Area Network Manager site:
<http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageAreaNetworkManager.html>
- ▶ Microsoft Knowledge Base Article 131658:
<http://support.microsoft.com/support/kb/articles/Q131/6/58.asp>
- ▶ Microsoft Knowledge Base Article 149927:
<http://support.microsoft.com/support/kb/articles/Q149/9/27.asp>
- ▶ Sysinternals home page:
<http://www.sysinternals.com>
- ▶ Subsystem Device Driver download site:
<http://www-1.ibm.com/servers/storage/support/software/sdd/index.html>
- ▶ IBM System Storage Virtualization Home Page:
<http://www-1.ibm.com/servers/storage/software/virtualization/index.html>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

advanced copy functions 5, 31
AIX 250

B

back-end storage controller 114, 192
Brocade 8

C

Cisco 265–266
Cisco MDS 8
clustered host systems 248
CNT 8
controller 43

D

data migration 17
data migration methods 18
Dense Wavelength Division Multiplexing 12
Disk Magic 16
Disk Magic, planning SVC implementation with 16
distance-extender technology 12
DS4800 53–80
DS8000 53, 59

E

EMC 6, 101
EMC CLARiiON 51, 120, 123–136, 251
 controller 122–124
 CX200 123
 CX600 124
 FC4700 124
 switch zoning considerations 123
EMC FC4700 120–136
EMC PowerPath 105
 device driver 112
EMC Symmetrix 51, 102, 105
 administrative interface 107
 controller 105, 107
 existing storage volumes 102
 LUNs 104
 model 104
 port flag setting 114
 related hdisks 112
 storage controller 102
 subsystem 105
 Switch zoning considerations 105
 volume 108
environments, supported 5, 30–32
environments, unsupported 8
etApp FAS 210
extenders, FCIP 12

extent size 15, 26, 37, 116, 206

F

fabric 12
fabric, SAN 8
FC protocol 218
 existing LUN 224
FCIP extenders 12
Fibre Channel
 adapter 105
 director 105
 infrastructure 211
 LUNs 227
 protocol 211
 target function 215

G

groups, MDisk 15

H

HDS 81
HDS AMS500
 administration tool 202
 disk storage system 190
 storage controller 190
HDS Lightning ??–92
HDS Lightning 82
HDS TagmaStore USP 93
HDS Thunder 93–100, 251
HDS-SVC zone 191
Hewlett-Packard 6–7, 161–187
Hewlett-Packard EVA 162
Hitachi Data Systems 81
 see HDS 6
host cluster migration 248
host operating systems 6
host-based data migration 20–21
HP 161
HP-UX 251

I

I/O group 39, 147, 196
 SVC cluster 39
 SVC nodes 46
 VDisk migration 39
IBM DS4800 190
 administration guide 198
 administrative interface 197
 disk storage system 190
 storage controller 189–190
IBM N-series 192
 supported models 192

- IBM SDD 109
 - command line 233
 - multi-path device driver 218
- IBM System Storage
 - DS4000 management interface 197, 227
 - N5000 series 213
 - N-series 192
- image mode 18, 34, 37, 104, 115
- image mode VDisk 15, 37, 47, 116, 141, 206
- installation, SVC 16
- Inter-Switch Links 12
- IO groups 16
- iSCSI initiator
 - group 221
 - software 218
- iSCSI protocol 211
 - block-level volumes 211
 - existing LUN 221
- iSCSI volume 215, 217
- ISL 12–13

L

- level 6
- logical volume 111, 118
- Logical Volume Management (LVM) 27

M

- maintaining availability 4
- maintenance window 18
 - maximum length 22
- managed disk group
 - extent size 145
- Managed Mode migration 26
- Managing storage 4
- McDATA 8
- MDisk 15, 35, 37, 43, 116, 191
- MDisk Group
 - extent size 145
- MDisk group 15, 116, 145, 200, 254
- member fc aliases
 - host 266
 - node1port1 266
 - node1port2 266
 - node2port1 266
 - node2port2 266
 - storage 266
- migration
 - monitoring 47
 - planning and sizing 10
- migration types
 - MDisk 43–45
 - non-standard 45–47
 - VDisk between I/O groups 39–41
 - VDisk between MDGs 35–37
 - VDisk extents 41–43
 - VDisk to image mode 37–39
- migration, volume 35
- monitoring migration 47
- multipath device driver 9, 32, 103

- multipath device driver, supported 9

N

- NetApp 6, 189, 192
- NetApp FAS 52, 210–236
 - existing LUN IDs 221
 - FC protocol 221
 - FCP feature 220
 - Fibre Channel target function 215
 - SCSI LUN ID 226
 - supported models 213
- NetApp host
 - group 197, 202
 - object 191
- NetApp storage controller 189
- NetApp V-Series 189–190
 - Concurrent host connection 192
 - Integration 190
 - remote administration 194
 - supported models 192
- NetApp V-Series NAS gateway 189
- Network Appliance 189
- nodes, SVC 13
- Novell 7

O

- operating systems, host 6
- OS support 6

P

- physical installation, SVC 16
- planning and sizing migration 10
- planning SVC implementation with Disk Magic 16
- preferred node 147
- PV Name 128

Q

- quorum disk support 30

R

- RAID 15, 59
- Red Hat 7
- Redbooks Web site 270
- Remote LAN Module (RLM) 193
- RPQ process 9, 32

S

- SAN boot migration 237–247
- SAN fabric 8, 12, 31, 58, 74, 106, 113, 124, 195
- SAN fabric environment 8
- SAN switch 106, 193, 197, 257
- SAN zone 102, 198
 - EMC Symmetrix 114
- SAN zoning 257–258
- see NetApp 189
- sharing a NetApp 189

- solution assurance review 8
- storage controller-based data migration 19–20, 27
- storage virtualization 25, 31–32, 61, 104, 191
- StorageTek 6
- Subsystem Device Driver (SDD) 103
- SUN 7
- SUN StorEdge 6
- supported environments 5, 11, 30–32
- supported multipath drivers 9
- SuSE Linux 7
- SVC
 - OS support 6
 - supported environments 5
 - unsupported environments 8
- SVC administration commands 254–256
- SVC benefits 4–5
- SVC cluster 26, 30, 103, 193, 258
 - I/O group 39
 - Managed Disk Groups 35
- SVC Image Mode 23
- SVC Image Mode data migration 23–25
- SVC implementation with Disk Magic, planning 16
- SVC initiator port 123
- SVC Managed Mode 25
- SVC Managed Mode data migration 25, 27
- SVC migration plan 29
- SVC node 13, 39, 46, 57–58, 106, 124, 192, 258
- SVC physical installation 16
- SVC port
 - number 219
 - numbers First 195
- SVC port numbers 198, 241
- svcinfo lsmdisk 115, 200
 - command 115, 200
 - mdisk0 235
 - mdisk11 210
- svcinfo lsnode
 - command 106, 198
 - node1 106, 199
 - nodename command 106, 199
- svctask mkhost 116, 204, 256
- svctask mkvdisk 116, 206
- svctask mkvdiskhostmap 117, 207
- Switch zoning 31

T

- target MDG 48, 209
- target port 105, 214
- temporary storage-based data migration 22–23, 27

U

- unsupported environments 8
- unsupported OEM controllers 247

V

- VDisk 41
- Virtualization 5
- VMware 7

- volume group 103–104, 164
 - import process 157
 - vgcfbackup command 158
- volume migration 35, 197, 208

Z

- Zone Member 264
- zone name
 - host_storage vsan 1 265
 - svc vsan 1 265
 - svc_host vsan 1 266
 - svc_storage vsan 1 266
- zones, host 14
- zoning 14, 58, 60, 74, 257–262
- zoning changes 257
- zoning configuration 261–262, 265
- zoning layout 258
- zoning rules 14
- zoning steps 259–264, 266

Archived

IBM



Implementing the SVC in an OEM Environment

(0.5" spine)

0.475" <-> 0.873"

250 <-> 459 pages



Implementing the SVC in an OEM Environment



Redbooks

Planning an SVC implementation in an OEM environment

Practical integration scenarios and detailed procedures

Guidance for migrating data to the SVC

This IBM Redbooks document describes how to plan and perform integration of the IBM SAN Volume Controller (SVC) in existing OEM storage environments.

The first part of the book covers topics that apply to any OEM environment. It begins with a brief overview of the benefits of using the SVC to manage virtualized storage, and identifies the large number of currently supported environments. Planning and sizing considerations are discussed, along with data migration methodologies, generic SVC migration plans, and the types of SVC migration.

The second part of the book covers integration of the SVC in some of the most commonly used OEM storage environments from IBM, HDS, EMC, HP and NetApp. Practical scenarios are described for each environment, and step-by-step details and code samples guide you through the planning, preparation, and performance of the migration.

The final section presents an overview of SVC administration commands and a detailed discussion of SAN zoning, along with sample code for performing the actual zoning changes.

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7275-00

ISBN 0738486035