

Using the IBM System Storage N series with Mail Servers

Using N series with
Lotus Domino

Using N series with
Microsoft Exchange

Using N series software with
mail servers



Alex Osuna
Patrick Március Médice Bisi
Prashanta Kumar Goswami
Sven Schaffranneck
Haijiang Sha

Redbooks



International Technical Support Organization

Using the IBM System Storage N series with Mail Servers

December 2007

Archived

Note: Before using this information and the product it supports, read the information in Figure on page ix.

Archived

First Edition (December 2007)

This edition applies to Data ONTAP Version 7.1 and above.

© 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	ix
Trademarks	x
Preface	xi
The team that wrote this book	xi
Become a published author	xiii
Comments welcome	xiii
Part 1. Introduction	1
Chapter 1. Introduction to IBM System Storage N series	3
1.1 IBM N series hardware	4
1.2 Comparison of IBM N series Gateway to the IBM System Storage N series	5
1.2.1 IBM N series A models hardware quick reference	6
1.2.2 IBM N series A and G models hardware quick reference	7
1.3 The IBM N series Standard Software Features	8
1.4 Optional software	9
1.5 Software quick reference	11
1.6 IBM System Storage N series A models	13
1.6.1 IBM System Storage N3000 Introduction	15
1.6.2 IBM System Storage N5000 introduction	24
1.6.3 IBM System Storage N7000 introduction	29
1.7 IBM System Storage N series Gateways (G models)	38
1.7.1 IBM System Storage N series Gateway highlights	40
1.7.2 Gateway RAID	40
1.7.3 IBM N5200, N5300, N5500, and N5600 Gateway models	41
1.7.4 IBM Gateway models N7600 and N7800	43
1.7.5 LUN sizing	44
1.7.6 LUN mapping	44
1.8 Interoperability between G and A models	45
1.9 EXN2000	46
1.9.1 Switched Hub architecture	47
1.9.2 EXN1000	48
1.9.3 EXN4000	49
Chapter 2. Introduction to IBM Lotus Domino 7	51
2.1 IBM Lotus Domino 7 overview	52
2.1.1 Terms and definitions	53
2.1.2 Server options	53
2.1.3 Client options	54
2.2 Key features	54
2.2.1 Enterprise messaging system	54
2.2.2 Platform independent	54
2.2.3 High availability and load balancing	56
2.2.4 Security	57

2.3 Notes Storage Files (NSF) and transaction logging	57
Chapter 3. Introduction to Microsoft Exchange	61
3.1 Challenges of DAS-based Exchange infrastructures	62
3.2 Five steps to leverage your Microsoft Exchange infrastructure.	63
3.2.1 Deploy either an IP SAN or an FC SAN	63
3.2.2 Implement N series Snapshot for backup and recovery	64
3.2.3 Virtualize storage for faster provisioning and better capacity utilization	64
3.2.4 Implement low-overhead remote mirroring for DR	66
3.2.5 Implement tiered storage for archival and compliance	67
Part 2. Preparing the IBM System Storage N series for Lotus Domino and Microsoft Exchange.	69
Chapter 4. Preparing the IBM System Storage N series for IBM Lotus Domino 7 ...	71
4.1 Using Lotus Domino with an IBM System Storage N series storage system	72
4.1.1 Advantages of IBM System Storage N series storage systems for Lotus Domino	72
4.1.2 Lotus Domino storage requirements overview	73
4.1.3 Capacity calculation of operating system, swap area, and program files	75
4.1.4 Capacity calculation for Lotus Domino databases	76
4.1.5 Lotus Domino transactional log calculation.	77
4.1.6 Fractional space reservation.	80
4.1.7 Hard disk performance and I/O tuning	80
4.2 Storage configuration for Lotus Domino server.	83
4.2.1 Aggregates	83
4.2.2 Creating the Domino database aggregate with Data ONTAP FilerView.	85
4.2.3 Creating the Domino log aggregate with Data ONTAP CLI	89
4.2.4 Volumes	91
4.2.5 Creating the Domino database volume with Data ONTAP FilerView	92
4.2.6 Creating the Domino transaction log volume with Data ONTAP CLI	96
4.2.7 Creating the LUNs	97
4.3 Zoning	99
4.3.1 Hard zoning.	99
4.3.2 Soft zoning	100
4.3.3 Zoning architecture	100
4.3.4 Zoning recommendations	102
4.3.5 Paths.	103
4.4 Snapshots	104
4.5 Protocols	105
4.5.1 Lotus Domino protocol recommendations.	105
4.6 Requirements for Lotus Domino with N series	107
4.7 SnapDrive	107
4.7.1 Benefits of SnapDrive	108
4.7.2 SnapDrive requirements	109
4.8 SnapMirror	110
4.8.1 Asynchronous mode	110
4.8.2 Synchronous mode	110
4.8.3 Semi-synchronous mode	111
4.8.4 Benefits of SnapMirror	111
4.8.5 Effects on sizing	111
4.8.6 SnapMirror requirements	111
4.9 SnapVault	112
4.9.1 Benefits of SnapVault	113
4.9.2 SnapVault requirements	113
4.9.3 Effects on sizing	114

4.10 FlexClone	114
4.10.1 FlexClone use cases and benefits	115
4.10.2 FlexClone requirements	116
4.10.3 Effects on sizing	117
Chapter 5. Preparing the IBM System Storage N series for Microsoft Exchange...	119
5.1 Storage requirements for a Microsoft Exchange server	120
5.1.1 Capacity	120
5.1.2 I/O performance	126
5.2 Storage configuration for Microsoft Exchange server	126
5.2.1 Aggregates	127
5.2.2 Volumes	136
5.2.3 LUNs	144
5.2.4 Protocols	144
5.2.5 Role-based access control	146
5.2.6 Creating a new role	147
5.2.7 Creating a new group	148
5.2.8 Creating a new user	148
5.3 Zoning	149
5.3.1 Hard zoning	149
5.3.2 Soft zoning	150
5.3.3 Zoning architecture	150
5.3.4 Zoning recommendations	152
5.3.5 Paths	153
5.4 Snapshots	154
5.5 Requirements for Microsoft Exchange with N series	156
5.6 SnapDrive	156
5.6.1 Benefits of SnapDrive	158
5.6.2 SnapDrive requirements	159
5.7 SnapMirror	159
5.7.1 Benefits of SnapMirror	161
5.7.2 Effects on sizing	161
5.7.3 SnapMirror requirements	161
5.8 SnapVault	162
5.9 FlexClone	162
5.9.1 FlexClone use cases and benefits	163
5.9.2 FlexClone requirements	165
5.9.3 Effects on sizing	166
5.10 SnapManager for Microsoft Exchange	166
Part 3. Installing Lotus Domino Server and Microsoft Exchange on the IBM System Storage N series storage system.	167
Chapter 6. Installing Microsoft Exchange on IBM System Storage N series.	169
6.1 Accessing the IBM System Storage N series storage system	170
6.1.1 Fibre Channel Protocol (FCP)	170
6.1.2 Internet SCSI Protocol (iSCSI)	187
6.2 SnapManager for Microsoft Exchange	217
Chapter 7. Installing Lotus Domino on IBM System Storage N series.	239
7.1 IBM System Storage N series protocol setup	240
7.1.1 Fibre Channel Protocol (FCP)	240
7.2 Role-based access control of the IBM System Storage N series	242
7.2.1 Creating a new role	243

7.2.2	Creating a new group	243
7.2.3	Creating a new user	244
7.3	Host connection	245
7.3.1	Configuring a Linux Host for FCP	246
7.3.2	Configuring an AIX Host for FCP	250
7.4	SnapDrive installation and usage	257
7.4.1	Install SnapDrive on Linux	259
7.4.2	Install SnapDrive on AIX	261
7.4.3	Create and map LUNs with SnapDrive	264
7.4.4	Create Snapshots with SnapDrive	267
7.4.5	Resize the volume	269
7.4.6	Best practice: Create storage with LVM enabled	271
7.5	Lotus Domino storage partitioning	271
7.5.1	Database placement	272
7.5.2	Transactional log placement	273
7.5.3	Program file placement	274
Part 4.	IBM System Storage N series system operations with Lotus Domino Server and Microsoft Exchange	277
Chapter 8.	IBM System Storage N series administration with Lotus Domino Server and Microsoft Exchange	279
8.1	FilerView	280
8.1.1	Resize aggregate	281
8.1.2	Resize volume	288
8.1.3	Resizing the LUN Using SnapDrive	292
8.2	Working with Snapshots	294
8.2.1	Restoring Snapshot copy with SnapDrive	298
8.2.2	Scheduling Snapshot Using SnapDrive	300
8.2.3	Scheduling Snapshot Using SnapManager for Microsoft Exchange	301
8.3	Scheduling Snapshot using FilerView	302
8.4	Performance commands and tools	303
8.4.1	Data ONTAP commands	304
8.5	Operations Manager	306
8.5.1	Operations Manager	307
8.6	FlexClone	310
8.6.1	Creating FlexClone	310
8.6.2	Status of FlexClone	311
8.6.3	Splitting the FlexClone	311
8.6.4	Space utilization with FlexClone	311
8.7	Reallocation	312
8.8	User management	312
8.9	Autosupport	313
8.9.1	Configuring Autosupport from FilerView	314
Chapter 9.	Backup and restore of Microsoft Exchange servers using IBM System Storage N series	315
9.1	Backup and restore	316
9.1.1	Backup	316
9.1.2	Restore	335
9.1.3	Single Mailbox Recovery	377

Chapter 10. Backup and restore of Lotus Domino server	383
10.1 Backup and restore	384
10.1.1 Backup	385
10.1.2 Restoring from Snapshot	400
 Appendix A. Integrating Lotus Domino for Windows into the IBM System Storage N series	 407
Introduction	408
Configuration	408
Microsoft iSCSI software initiator	408
SnapDrive software	409
Benefits of SnapDrive	410
SnapDrive requirements	411
Network	411
IBM System Storage N series storage system	412
Using Lotus Domino with an IBM System Storage N series storage system	413
Advantages of IBM System Storage N series storage systems	413
Lotus Domino transaction logging	415
Configure the Domino environment	417
Configuring the storage system	417
Create an aggregate	418
Creating a volume	424
Specify the volume security settings to use	429
Disable the automatic Snapshot feature	430
Setting the snap reserve option for the volume	430
Obtain a node name for the storage system	430
Create a user with administrator privileges	431
Configuring the Windows server	431
Installing the Microsoft iSCSI software	431
Configuring the Microsoft iSCSI software	435
Installing SnapDrive for Windows	444
Creating Virtual disks	452
Installing and configuring a Lotus Domino server	458
Enabling Lotus Domino Transaction Logging	462
Migrating an existing Lotus Domino server from local disk to a IBM System Storage N series storage system	464
 Related publications	 467
IBM Redbooks publications	467
Other publications	467
Online resources	467
How to get IBM Redbooks publications	468
Help from IBM	468
 Index	 469

Archived

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:

AIX 5L™
AIX®
Domino®
DB2®
IBM®

Lotus Notes®
Lotus®
LotusScript®
Notes®
Power PC®

Redbooks®
Redbooks (logo) ®
System Storage™
Tivoli®
TotalStorage®

The following terms are trademarks of other companies:

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

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Snapshot, SecureAdmin, RAID-DP, Network Appliance, LockVault, FlexShare, WAFL, SyncMirror, SnapVault, SnapValidator, SnapRestore, SnapMover, SnapMirror, SnapManager, SnapLock, SnapDrive, NearStore, FlexVol, FlexClone, FilerView, Data ONTAP, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

AMD, AMD Opteron, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

Flex, and Portable Document Format (PDF) are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

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

Active Directory, Microsoft, Outlook, SQL Server, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

"Microsoft product screen shot(s) reprinted with permission from Microsoft Corporation."

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

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

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

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

Preface

This IBM® Redbooks® publication gives mail database administrators a introduction to the IBM System Storage™ N series. It also gives both storage administrators and mail database administrators the information to prepare, install, and administer the IBM System Storage N series when used in conjunction with Microsoft® Exchange and IBM Lotus® Domino®. Often the lines of responsibility with regards to these tasks are blurred or crossed over, so some sections of this book may be redundant for experienced N series administrators but of importance to mail database administrators. Conversely, some sections that may be redundant for mail database administrators may help storage administrators.

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.

Alex Osuna is a Project Leader with the San Jose International Technical Support Organization. He has over 27 years in the I/T industry and 22 years of experience in the hardware/software storage area dealing with maintenance, development, early ship programs, education, publishing, performance analysis, and technical sales support. He holds 10 certifications from IBM, Microsoft, and Red Hat.

Patrick Március Médice Bisi is an IT Specialist at IBM Global Services in IBM Brazil. He has over 10 years of experience in the IT industry. He holds a degree in Business Administration from SEDES/UVV and holds several certifications from Microsoft, including MCSE: Messaging and MCSA: Messaging. His areas of expertise include Microsoft Windows® operating systems, Microsoft Exchange servers, and Active Directory® directory services.

Prashanta Kumar Goswami is working as a Technical Lead of the Functional Verification and Testing team for the Virtual Input/Output System. He has been with IBM India for two years and has over five years of experience in the IT industry. He holds a Bachelor of Computer Science and Technology degree from Vidayasagar University India.

Sven Schaffranneck is an IT project leader at Netzlink Informationstechnik GmbH, an IBM Business Partner from Germany. He has over 10 years of experience in several areas of the IT industry. He is responsible for a high availability hosting environment that includes customer IT outsourcing, Microsoft Exchange, and RIM BES infrastructure. In addition, his current focus is presales and postsales work on Microsoft Windows, Active Directory, Microsoft Exchange, and IBM system storage installations. Sven holds a Dipl.-Ing. degree in Computer Sciences from the University of Applied Science in Braunschweig/Wolfenbüttel, Germany.



Figure 1 Haijiang Sha

Haijiang Sha is a Advisory IT Specialist in IBM China. He has 10 years of experience in the network support field. He holds a Bachelor's degree of Engineering from the Institute of Command and Technology of COSTIND. His areas of expertise include NAS, SAN fabric, and networks.



Figure 2 From left to right: Patrick Bisi, Sven Schaffranneck, Alex Osuna, and Prashanta Goswami

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

Jawahar Lal
Network Appliance™ Inc.

Robert Campbell
Network Appliance Inc.

Shannon Flynn
Network Appliance Inc.

Lee Dorrier
Network Appliance Inc.

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- Use the online **Contact us** review book 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

Introduction

In this part, we introduce IBM System Storage N series, Lotus Domino Server, and Microsoft Exchange.

Archived

Introduction to IBM System Storage N series

In this chapter, we introduce the IBM System Storage N series and describe the hardware and software. The reader will also be introduced to storage architectures and file systems.

The IBM Storage System N series provides a range of reliable, scalable storage solutions for a variety of storage requirements. These capabilities are achieved by using network access protocols such as NFS, CIFS, HTTP, and iSCSI, as well as Storage Area technologies, such as Fibre Channel. Utilizing built-in RAID technologies, all data is well protected with options to add additional protection through mirroring, replication, snapshots, and backup. These storage systems are also characterized by simple management interfaces that make installation, administration, and troubleshooting uncomplicated and straightforward.

The IBM System Storage N series is designed from the ground up as a stand-alone storage system. The advantages of using this type of flexible storage solution include:

- ▶ The ability to tune the storage environment to a specific application while maintaining the flexibility to increase, decrease, or change access methods with a minimum of disruption.
- ▶ The capability to react easily and quickly to changing storage requirements. If additional storage is required, being able to expand it quickly and non-disruptively is needed. When existing storage exists but is deployed incorrectly, the capability to reallocate available storage from one application to another quickly and simply cannot be done.
- ▶ The ability to maintain availability and productivity during upgrades. If outages are required, keeping them to the shortest time possible.
- ▶ The ability to create effortless backup/recovery solutions that operate commonly across all data access methods.
- ▶ File and block level services in a single system, helping to simplify your infrastructure.
- ▶ The ability to tune the storage environment to a specific application while maintaining its availability and flexibility.
- ▶ The deployment of storage resources can be changed non-disruptively, easily and quickly. Online storage resource redeployment is possible.
- ▶ An easy and quick upgrade process. Non-disruptive upgrade is possible.

- Strong data protection solutions and support online backup/recovery.

1.1 IBM N series hardware

In the following sections, we will discuss the N series models available today (see Figure 1-1 and Figure 1-2 on page 5).

- N3000
- N5000 series
- N7000 series



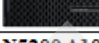
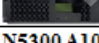
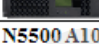
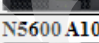

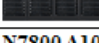

Use [↗]	Model [↗]	Capacity [↗]
Entry level [↗]	N3700 A10&A20 [↗] 	16TB [↗]
Entry level [↗]	N3300 A10&A20 [↗] 	24TB [↗]
Entry level [↗]	N3600 A10&A20 [↗] 	69TB [↗]
Midrange [↗]	N5200 A10&A20 [↗] 	84TB [↗]
Midrange [↗]	N5300 A10&A20 [↗] 	126TB [↗]
Midrange [↗]	N5500 A10&A20 [↗] 	168TB [↗]
Midrange [↗]	N5600 A10&A20 [↗] 	210TB A10 [↗] 252TB A20 [↗]
Enterprise Class [↗]	N7600 A10&A20 [↗] 	336TB A10 [↗] 420TB A20 [↗]
Enterprise Class [↗]	N7800 A10&A20 [↗] 	336TB A10 [↗] 504TB A20 [↗]

Figure 1-1 IBM System Storage N series systems







Use ^o	Model ^o	Capacity ^o
Midrange ^o	N5200 G10&G20 ^o 	50TB ^o
Midrange ^o	N5300 G10&G20 ^o 	126TB ^o
Midrange ^o	N5500 G10&G20 ^o 	84TB ^o
Midrange ^o	N5600 G10&G20 ^o 	252TB ^o
Enterprise Class ^o	N7600 G10&G20 ^o 	420TB ^o
Enterprise Class ^o	N7800 G10&G20 ^o 	504TB ^o

Figure 1-2 IBM N series Gateway models

1.2 Comparison of IBM N series Gateway to the IBM System Storage N series

- ▶ Identical core NAS feature/functionality.
- ▶ Identical iSCSI feature/functionality.
- ▶ Identical FCP feature/functionality.
- ▶ Filer SAN host support matrix applies.
- ▶ Identical behavior for WAFL® file system.
- ▶ Identical data availability characteristics.
- ▶ Identical data integrity characteristics.
- ▶ Identical data management characteristics.
- ▶ Identical serviceability characteristics.
- ▶ Supports the same version of Data ONTAP®.
- ▶ Differences exist in system initialization and storage expansion.
- ▶ The A N5000 and N 7000 series Gateway physical attributes are the same as the N5000 and N7000 models A10 and A20 storage systems.
- ▶ The IBM N series Gateway does not use the following features of Data ONTAP®:
 - Disk sanitization and disk scrubbing.
 - SnapLock® Compliance.
 - LockVault™ Compliance.
 - Nearstore option.
 - RAID-DP™.
 - RAID 4.
- ▶ N5000 storage systems use disk storage provided by IBM only; the Gateway models support heterogeneous storage.
- ▶ Data ONTAP is enhanced to enable the Gateway Series solution.

- A RAID array provides LUNs to the Gateway:
 - Each LUN is equivalent to an IBM disk.
 - LUNs are assembled into aggregates/volumes, then formatted with the WAFL file system, just like the IBM System Storage N series systems.

1.2.1 IBM N series A models hardware quick reference

Table 1-1 shows the IBM N series A models hardware quick reference.

Table 1-1 IBM N series A models hardware quick reference

Function	N3700	N3300	N3600	N5200	N5500	N5300	N5600	N7600	N7800
Maximum Raw Capacity in TB A10 models	16	24	69	84	168	126	210	336	336
Maximum Raw Capacity in TB A20 models	16	24	69	84	168	126	252	420	504
Fibre Channel Disk drives	144 GB 10K RPM, 144 GB 15K RPM, 300 GB 10K RPM - EXN2000 144GB 15K, 300GN 10K - EXN4000								
SATA Disk Drives	250 GB 7.2K RPM., 320 GB 7.2K RPM., 500 GB 7.2K RPM, 750GB 7.2KRPM								
Maximum number of disks	56	40	104	168	336	252	420(A10) 504(A20)	672(A10) 840(A20)	672 (A10) 1008 (A20)
Maximum Raw Capacity in TB based on SATA Drive type	16	24 ¹	69 ²	42 @ 250GB 54 @ 320GB 84 @ 500GB	168	63 @2 50GB 126 @ 500GB	105 @250 GB(A10) 126 @250 GB(A20) 134 @320 GB(A10) 161 @320 GB(A20) 210 @500 GB(A10) 252 @500 GB(A20)	168 @250 GB(A10) 210 @250 GB(A20) 215 @320 GB(A10) 268 @320 GB(A20) 336 @500 GB(A10) 420 @500 GB(A20)	168 @250 GB(A10) 252 @250 GB(A20) 215 @320 GB(A10) 322 @320 GB(A20) 336 @500 GB(A10) 504 @500 GB(A20)
Expansion units supported	EXN1000 (SATA), EXN2000 (FC), EXN4000 4Gbps (FC)								

1. N3300 and N3600 also support SAS disks.

2. N3300 and N3600 also support SAS disks.

Table 1-2 shows the IBM N series G models hardware quick reference.

Table 1-2 N series G models hardware quick reference

Function	N5200	N5300	N5500	N5600	N7600	N7800
Maximum Raw Capacity in TB G10 models	50	126	84	252	336	336

Function	N5200	N5300	N5500	N5600	N7600	N7800
Maximum Raw Capacity in TB G20 models	50	126	84	252	420	504
Max. number of — Logical Units (LUNs) on back-end disk storage array	168	252	336	420 for A10 and 504 for A20	840	1008
Max LUN size in GB	500	500	500	500	500	500
Maximum Volume size in TB	16	16	16	16	16	16

Note: A stand-alone gateway must own at least one LUN. A cluster configuration must own at least two LUNs.

1.2.2 IBM N series A and G models hardware quick reference

Table 1-3 shows the IBM N series A and G models hardware quick reference.

Table 1-3 IBM N series A and G models hardware quick reference

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7800
Network Protocol support	NFS V2/V3/V4 over UDP or TCP, PCNFSD V1/V2 for (PC) NFS client authentication, Microsoft CIFS, iSCSI, FCP, VLD, HTTP 1.0, HTTP1.1 Virtual Host								
Other Protocol Support	SNMP, NDMP, LDAP, NIS, DNS								
Onboard I/O ports per node	2 X GbE 2 X Optical FC	4 x GbE 4 x Optical FC	4 x GbE 4 x Optical FC	4 X GbE 4 X FC 1 X LVD SCSI	4 X GbE 4 X FC 1 X LVD SCSI	4 X GbE 4 X FC 1 X LVD SCSI	4 X GbE 4 X FC (4Gbps)	6 X GbE 8 X FC	6 X GbE 8 X FC
PCI expansion slots per node	N/A	N/A	1 X PCI-E	3 X PCI-X	3 x PCI-E	3 X PCI-X	3 X PCI-E	5 X PCI-E, 3 X PCI-X	5 X PCI-E, 3 X PCI-X
NVRAM in MB per node	128	128	256	512	512	512	512	512	2048
Memory in GB per node	1	1	2	2	4	4	8	16	32
Redundancy/ High Availability	CompactFlash, dual-redundant hot-plug integrated cooling fans, hot-swappable autoranging power supplies, clustered storage controllers, hot-swappable disk bays**								

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7800
Required rack space	3U	2U	4U	3U per node	3U per node	3U per node	3U per node	6U per node	6U per node
Processors (A10)	Two Broadcom MIPS-based	2.2 GHz 64-bit processors	2 2.2 GHz 64-bit processors	One 2.8 GHz Xeon	Two 1.8 GHz AMD™	Two 2.8 GHz Xeon	Two AMD 1.8GHz dual-core	Two 2.6 GHz AMD Opteron™	Four 2.6 GHz AMD Opteron
Processors (A20)	Four Broadcom MIPS-based	Two 2.2 GHz 64-bit processors	Four 2.2 GHz 64-bit processors	Two 2.8 GHz Xeon	Four 1.8 GHz AMD	Four 2.8 GHz Xeon	Four AMD 1.8GHz dual-core	Four 2.6 GHz AMD Opteron	Eight 2.6 GHz AMD Opteron

1.3 The IBM N series Standard Software Features

These are licensed, no charge features that are available with the IBM N series.

Table 1-4 Standard Software Features

Feature	Description
Data ONTAP	Operating system software that optimizes data serving and allows multiple protocol data access.
FTP	File Transfer Protocol (FTP), which is a standard internet protocol that is a simple way to exchange files between computers on the internet.
Telnet	The TELNET Protocol provides a general, bi-directional, eight-bit byte oriented communications facility. It provides user oriented command-line login sessions between hosts.
Snapshot™	Enables online backups, providing near instantaneous access to previous versions of data without requiring complete, separate copies.
FlexVol®	FlexVol creates multiple flexible volume on a large pool of disks. This provides dynamic, nondisruptive (thin) storage provisioning, which is very space and time efficient. These flexible volumes can span multiple physical volumes without regard to size.
FlexShare™	FlexShare gives administrators the ability to leverage the existing infrastructure and increase processing utilization without sacrificing the performance of critical business needs. With the use of FlexShare, administrators can confidently consolidate different applications and data sets on a single storage system. FlexShare gives administrators the ability to prioritize applications based on how critical they are to the business.
FlexCache	FlexCache has the ability to distribute files to remote locations without the need for continuous hands-on management. Storage systems deployed in remote offices automatically replicate, store, and serve the files or file portions that are requested by remote users without the need for any replication software or scripts.

Feature	Description
Disk sanitization	Disk sanitization is the process of physically obliterating data by overwriting disks with specified byte patterns or random data in a manner that prevents recovery of current data by any known recovery methods. This feature enables you to carry out disk sanitization by using three successive byte overwrite patterns per cycle. By default, six cycles are performed.
FilerView®	A Web-based administration tool that allows IT administrators to fully manage N series storage systems from remote locations. This is simple and intuitive Web-based single-appliance administration.
SnapMover®	Migrates data among N series clusters with no impact on data availability and no disruption to users.
AutoSupport	AutoSupport is a sophisticated, event-driven logging agent featured in the Data ONTAP operating software and inside each N series system that continuously monitors the health of your system and issues alerts if a problem is detected. These alerts can also be in the form of e-mails.
SecureAdmin™	SecureAdmin is a Data ONTAP module that enables authenticated, command-based administrative sessions between an administrative user and Data ONTAP over an intranet or the internet.
DNS	The N series supports using a host naming file or a specified DNS server and domain.
Cluster	<ul style="list-style-type: none"> ► Ensures high data availability for business-critical requirements by eliminating a single point of failure. ► Must be ordered for A20 clustered configurations or upgrades from A10 to A20 ► Active-active pairing delivers even more “nines to right of the decimal point”.
NIS	The N series does provide NIS client support and can participate in NIS domain authentication.
Integrated automatic RAID manager	The IBM N series and Data ONTAP provide integrated RAID management with RAID-Double Parity (default) and RAID 4.
iSCSI Host Attach Kit for AIX®, Windows, and Linux®	A Host Support Kit includes support software and documentation for connecting a supported host to an iSCSI network. The support software includes programs that display information about storage, and programs to collect information needed by customer support to diagnose problems.

1.4 Optional software

Table 1-5 shows some of the optional software for IBM N series.

Table 1-5 Optional software

Feature	Description
CIFS	Provides file system access for Microsoft Windows environments.
NFS	Provides file system access for UNIX® and Linux environments.
HTTP	Hypertext Transfer Protocol allows a user to transfer displayable Web pages and related files.

Feature	Description
FlexClone®	Designed to provide instant replication of data volumes/sets without requiring additional storage space at the time of creation
Multistore	<ul style="list-style-type: none"> ▶ Permits an enterprise to consolidate a large number of Windows, Linux, or UNIX file servers onto a single storage system. ▶ Many “virtual filers” on one physical appliance ease migration and multi-domain failover scenarios.
SnapLock	Provides non-erasable and non-rewriteable data protection that helps enable compliance with government and industry records retention regulations.
LockVault	Designed to provide non-erasable and non-rewriteable copies of Snapshot data to help meet regulatory compliance needs for maintaining backup copies of unstructured data.
SnapMirror®	<ul style="list-style-type: none"> ▶ Remote mirroring software that provides automatic block-level incremental file system replication between sites. ▶ Available in synchronous, asynchronous, and semi synchronous modes of operation.
SnapRestore®	Allows rapid restoration of the file system to an earlier point in time, typically in only a few seconds.
SnapVault®	Provides disk based backup for N series systems by periodically backing up a Snapshot copy to another system.
SnapDrive®	SnapDrive enables Windows and UNIX applications to access storage resources on N series storage systems, which are presented to the Windows 2000 or later operating system as locally attached disks. For UNIX, it allows you to create storage on a storage system in the form of LUNs, file systems, logical volumes, or disk groups.
SnapManager®	Host software for managing Exchange and SQL Server™ and SAP® backup and restore. SnapManager software simplifies Exchange data protection by automating processes to provide hands-off, worry-free data management.
SnapValidator®	For Oracle® deployments, SnapValidator can be used to provide an additional layer of integrity checking between the application and N series storage. SnapValidator allows Oracle to create checksums on data transmitted to N series storage for writes to disk and include the checksum as part of the transmission.
SyncMirror®	SyncMirror is a synchronous mirror of a volume. It maintains a strict physical separation between the two copies of your mirrored data. In case there is an error in one copy, the data is still accessible without any manual intervention.
Single Mailbox Recovery for Exchange	SMBR is a software option from SnapManager that is designed to make near-instantaneous online backups of Exchange databases, verify that the backups are consistent, and rapidly recover Exchange within levels: storage group, database, folder, single mailbox, or single message. The potential results are improved service to internal clients, reduced infrastructure expenses, and significant time savings for Exchange administrators.
Operations Manager	Operations Manager provides remote, centralized management of IBM N series data storage infrastructure, including global enterprise, storage network, and so on.
MetroCluster	MetroCluster software provides an enterprise solution for high availability over wide area networks.

Feature	Description
NearStore® option	A disk-based, secondary storage device for enterprise applications.
Advanced Single Instance Storage	Designed to significantly improve physical storage efficiency and network efficiency by enabling the sharing of duplicate data blocks.

1.5 Software quick reference

Table 1-6 Software quick reference

Product/Feature /Function	Included/ optional	N3000 A10 & A20	N5X00 A10 & A20	N5x00 G10 & G20	N7x00 A10 & A20	N7x00 G10 & G20
Data ONTAP	Included	X	X	X	X	X
iSCSI protocol	Included	X	X	X	X	X
FTP protocol	Included	X	X	X	X	X
NDMP protocol	Included	X	X	X	X	X
FlexVol	Included	X	X	X	X	X
Snapshot	Included	X	X	X	X	X
SecureAdmin	Included	X	X	X	X	X
iSCSI Host Attach Kit for AIX, Windows, Linux	Included	X	X	X	X	X
FlexShare	Included				X	X
SnapMover	Included	X	X	X	X	X
CIFS protocol	Optional	X	X	X	X	X
NFS protocol	Optional	X	X	X	X	X
HTTP protocol	Optional	X	X	X	X	X
FCP protocol	Optional	X	X	X	X	X
FlexClone	Optional	X	X	X	X	X
Clustered Failover	Optional	X(A20)	X	X	X	X
Multistore	Optional	X	X	X	X	X
SnapMirror	Optional	X			X (Requires special HBA card.)	X (Requires special HBA card.)
SnapRestore	Optional	X	X	X	X	X
Open Systems SnapVault (OSSV)	Optional	X	X	X	X	X
SnapVault	Optional	X	X	X	X	X

Product/Feature /Function	Included/ optional	N3000 A10 & A20	N5X00 A10 & A20	N5x00 G10 & G20	N7x00 A10 & A20	N7x00 G10 & G20
SnapDrive for Windows and UNIX: AIX, Solaris™, HP-UX, and Linux	Optional	X	X	X	X	X
SnapValidator	Optional	X	X	X	X	X
SyncMirror	Optional		X	X	X	X
SnapManager for SQL Server	Optional	X	X	X	X	X
SnapManager for SAP	Optional					
SnapManager for Exchange	Optional	X	X	X	X	X
Single Mailbox Recovery for Exchange (SMBR)	Optional	X	X	X	X	X
Operations Manager Core, BC & SRM License	Optional	X	X	X	X	X
SnapLock Enterprise	Optional	X	X	X	X	X
LockVault Enterprise	Optional	X	X	X	X	X
MetroCluster A20 Models only	Optional		X	X	X	X
Disk Sanitization	Included	X	X		X	
SnapLock Compliance	Optional	X	X		X	
LockVault Compliance	Optional	X	X		X	
NearStore Option Bundle	Optional	X	X		X	
RAID 4, RAID-DP	Included	X	X		X	
Advanced Single Instance Storage	Optional		X	X	X	X

1.6 IBM System Storage N series A models

The A models of the IBM System Storage N series systems offer multiprotocol connectivity using internal storage or storage provided by expansion units (see Figure 1-3). The IBM System Storage N series systems are designed to provide integrated block- and file-level data access, allowing concurrent operation in IP SAN (iSCSI), FC SAN, NFS, and CIFS environments. Other storage vendors may require the operation of multiple systems to provide this functionality. IBM N series systems are designed to avoid costly downtime, both planned and unplanned, and improve your access to mission-critical data, thereby helping you gain a competitive advantage.

The IBM N series A models are a specialized, “thin server” storage system with a customized operating system, similar to a stripped down UNIX kernel, hereafter referred to as Data ONTAP. With a reduced operating systems, many of the server operating system functions that you are familiar with are not supported. The objective is to improve performance and reduce costs by eliminating unnecessary functions normally found in the standard operating systems.

The IBM N series come with pre-configured software and hardware, and with no monitor or keyboard for user access. This is commonly termed a “headless” system. A storage administrator accesses the systems and manages the disk resources from a remote console using a Web browser or command line.

One of the typical characteristics of a IBM System Storage N series product is its ability to be installed rapidly using minimal time and effort to configure the system. It is integrated seamlessly into the network. This approach makes IBM N series products especially attractive when lack of time and skills are elements in the decision process.

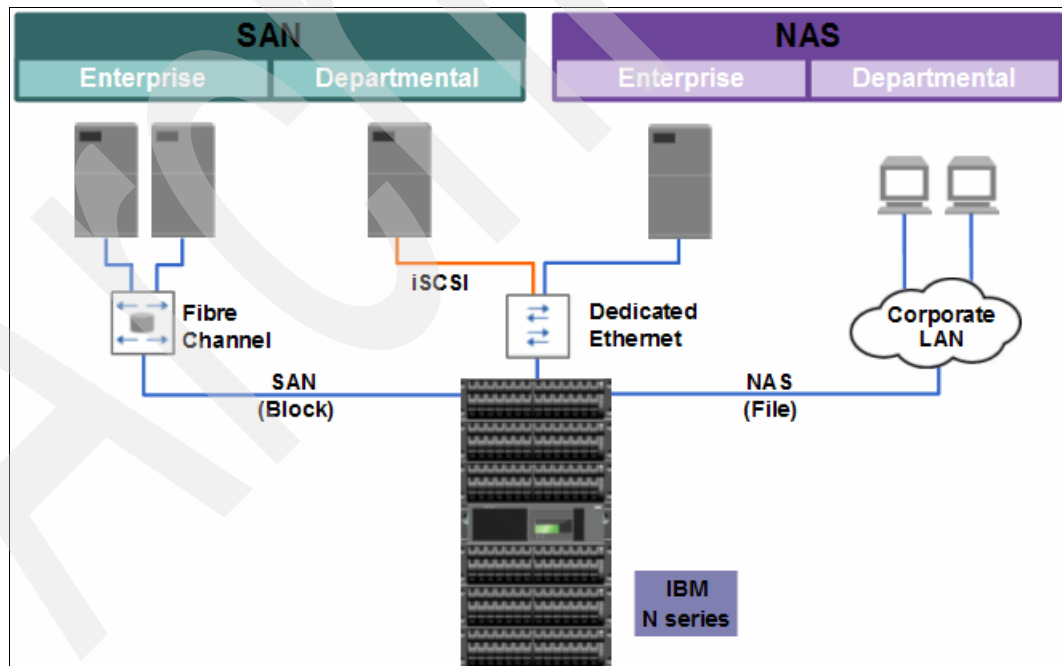


Figure 1-3 IBM N series A models

Drive flexibility

The IBM System Storage N series products are designed to provide network attached storage for environments where customers have a need to utilize their storage investment in a multifaceted environment. The IBM System Storage N series storage systems provide customers a tremendous amount of versatility by allowing this solution to be populated with both Fibre Channel disk drives and SATA disk drives. An IBM System Storage N series system populated with Fibre Channel disk drives may be suitable for mission-critical high-performance data transaction environments, while an IBM System Storage N series system populated with SATA disk drives may be attractive to customers who wish the use the platform for disk to disk backup scenarios, disaster recovery scenarios, archive data, or data like home directories that do not require high-performance transactional environments.

Table 1-7 gives examples of drive positioning.

Table 1-7 Drive positioning

Requirement	Fibre Channel Drives	SATA drives
On-line, high-performance, and mission critical production data repository	X	
Near-line storage used for tiered storage or infrequently accessed data	X	X
Data retention to help meet the needs of customers required to store data in non-erasable and non-rewriteable (WORM) formats		X

Near-line storage

The IBM System Storage N series with SATA drives offers near-line storage. In Figure 1-4 on page 15, you see an example of “Traditional disk based backup and recovery”. On the left, you see primary storage, which is characterized by a higher cost and very fast performance. On the far right, you have archive targets that have traditionally been tape or optical jukeboxes with reduced access times to read and write data. Two years ago, the concept of near-line storage in the middle for disk staging was introduced. This enables organizations to do daily backups to disk and back up to tape weekly or bi-weekly, which reduces the amount of data that needs to be written to tape. Also, data is online for faster recovery. The other advantage this provides is that you can leverage your existing investment in primary storage, your backup application, and tape libraries.

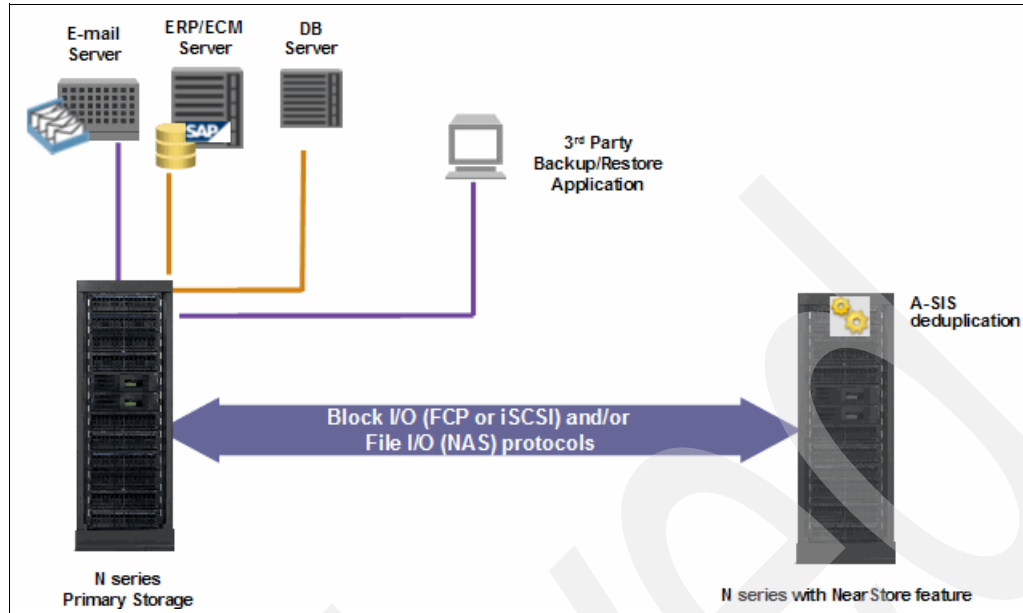


Figure 1-4 Near-line storage

1.6.1 IBM System Storage N3000 Introduction

The IBM N3000 systems are designed to provide primary and secondary storage for midsize enterprises. IT administrators can consolidate their fragmented application-based storage and unstructured data into one unified, easily managed, and expandable platform. N3000 systems offer integrated block- and file-level data access, and intelligent management software and data protection capabilities, such as higher-end N series systems, in a cost-effective package. IBM N series innovations include Serial-Attached SCSI (SAS) drive support, expandable I/O connectivity, and onboard remote management.

The N3000 systems are designed as the entry point to the entire N series family. The systems can provide the following key advantages:

- ▶ High availability leverages proven features, including a high performing and scalable operating system, data management software, and redundancy features.
- ▶ The Backup and Recovery Features are designed to support disk-based backup, with file or application-level recovery with Snapshot and SnapRestore software features.
- ▶ Simple Replication and Disaster Recovery is designed to provide an easy-to-deploy mirroring solution that is highly tolerant of WAN interruptions.
- ▶ Management Simplicity allows self-diagnosing systems designed to enable on-the-fly provisioning.
- ▶ A versatile, single, and integrated architecture designed to support concurrent block I/O and file serving over Ethernet and Fibre Channel SAN infrastructures.

The IBM N3000 is compatible with the entire family of N series unified storage systems, which feature a comprehensive line-up from top-to-bottom of hardware and software designed to address a variety of possible deployment environments.

- ▶ N3700
 - 2863-A10 Single Filer
 - 2863-A20 Clustered

- ▶ N3300
 - 2859-A10 Single Filer
 - 2859-A20 Clustered
- ▶ N3600
 - 2862-A10 Single Filer
 - 2862-A20 Clustered

The N3000 supports Ethernet and Fibre Channel environments, enabling economical NAS, FC, and iSCSI deployments. The N3000 system functions as a “unification engine,” which is designed to allow you to simultaneously serve both file and block-level data across a single or multiple networks, demanding procedures that for some solutions require multiple separately managed systems.

N3000 storage systems can offer significant advantages for distributed enterprises with remote and branch office sites. These organizations and others can leverage the SnapVault and SnapMirror software functions to implement a cost-effective data protection strategy by mirroring data back to a corporate data center.

N3700

The N3700 storage system (see Figure 1-5) is a 3U solution designed to provide NAS and iSCSI functionality for entry to mid-range environments. The basic N3700 offering is a single-node model A10, which is upgradeable to the dual-node model A20 and requires no additional rack space. The dual-node clustered A20 is designed to support failover and failback functions to maximize reliability. The N3700 filer can support 14 internal hot-plug disk drives with scalability being provided through attachment to up to three expansion units, each with a maximum of 14 drives. The N3700 also has the capability to connect to a Fibre Channel tape for backup.

A list of supported Tape drives can be found at:

<http://www.ibm.com/totalstorage/nas>

Refer to the IBM System Storage and TotalStorage® N series interoperability matrix found at:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>



Figure 1-5 N3700

The type of controller defines the model. Figure 1-6 shows a single control unit. The single node A10 uses a single control unit with a dual node clustered A20 using two control units (see Figure 1-7).



Figure 1-6 N3700 A10

Note: The only upgrade for an N3700 model A10 is to a model A20. This upgrade requires no additional rack space.



Figure 1-7 N3700 A20

The N3700 comes with redundant power supplies for higher reliability (see Figure 1-8).




Figure 1-8 Redundant Power supplies

From the rear of the N3700 on the power supply you can see the Diagnostic and Operational LEDs (Figure 1-9). Table 1-8 explains what the LEDs and possible configurations mean.




Figure 1-9 Diagnostic and Operational LEDs

Table 1-8 LED status table

LEDs visible from the rear of the system unit		
PSU Status Normal	On	Normal
AC missing for this PSU	Off	
Fan Fault	Off	
Output voltage, current, or temperature fault	Off	
PSU Status Normal	Off	Power Supply failure
AC missing for this PSU	Off	
Fan Fault	Off	
Output voltage, current, or temperature fault	On	
PSU Status Normal	Off	Fan Failure
AC missing for this PSU	Off	
Fan Fault	On	
Output voltage, current, or temperature fault	Off	
PSU Status Normal	Off	No power to this PSU
AC missing for this PSU	On	
Fan Fault	Off	
Output voltage, current, or temperature fault	On	

The CPU module shown in Figure 1-10 controls connectivity to the storage and connectivity to the clients. If a power supply fails or is turned off while the other power supply is still providing DC power, then both cooling fans will continue to operate.




Figure 1-10 CPU Tray module front view

The N3700 is based around a MIPS dual core processor. It has 1GB of system memory, of which 128 MB is defined as non-volatile due to having a battery backup. The battery is a three cell Li-Ion Figure 1-11.

Note: The NVRAM is a battery backed up portion of the main system memory




Figure 1-11 CPU tray module showing battery backup for memory

There is a 256 MB compact flash card located on the bottom of the CPU tray module (Figure 1-12). It contains a copy of the Data ONTAP operating system along with firmware. The operating system is also stored on each disk drive.




Figure 1-12 Bottom of CPU tray module showing Compact Flash card

Rear ports on the N3700 can be seen in Figure 1-13. Each CPU tray module has two integrated 2 Gbps Fibre Channel ports. Both of these ports are initially configured in “initiator mode” and do not use Small Form factor Pluggable SFPs.

The first port, channel C, is optical and intended for direct or SAN attachment to a tape library. A standard LC-LC short wave optic cable should be used for this connection.

The second port, channel B, is copper and is exclusively used for connection of expansion unit. A special copper cable (option X6531-C) is used for this connection.




Figure 1-13 External ports on CPU tray module

The CPU tray module also contains two on board 10/100/1000 Mb copper Ethernet ports. Each port has two LED lights, for activity and speed.

The final connection allows connection of an ASCII terminal through an RJ45 to DB-9 cable.

N3700 hardware features

- ▶ 3U integrated storage system
- ▶ 3U optional storage expansion shelf - up to three
- ▶ Redundant hot plug power supplies
- ▶ Redundant Cooling
- ▶ Integrated 10/100/1000 full duplex Ethernet
- ▶ Two Integrated Fibre Channel adapters
- ▶ Compact flash

► Diagnostic LEDs/OPS

Though ESH2 modules in the EXN2000 would support up to 84 drives per loop or one N3700 and five expansion shelves, only 56 drives and three expansion shelves and one N3700 are supported. This is a firmware limitation for backward compatibility by the manufacturer for the previous shelf module the Loop Redundant Circuit (LRC). This module is not available in the IBM N series.

Optional hardware

A second CPU tray supports Cluster Failover.

Table 1-9 Additional N3700 specifications

Storage System Specifications	N3700 A10	N3700 A20
Clustered Failover-Capable	No (Requires upgrade to A20)	Yes
Max Number of Expansion Units EXN1000/SATA disk drives or EXN2000/FC disk drives	3	3

N3300 and N3600

The N3300/3600 (see Figure 1-14 and Figure 1-15) systems provide multiple I/O connectivity options, a small footprint to hold high density SAS drives, and external expansion using low-cost SATA drives and Fibre Channel disks for production applications, and utilize Data ONTAP Snapshot technology. Serial-Attached SCSI (SAS) is the Next Generation of SCSI and it combines the advantages of parallel SCSI and serial FC. For further systems administration time and cost advantages, the systems come standard with Remote Onboard Management capabilities to help simplify remote system monitoring, cycle power, execute firmware upgrades, enter console commands, and run diagnostics to help maintain the reliability of the system and your business-critical data.




Figure 1-14 N3300



Figure 1-15 N3600

Figure 1-16 and Figure 1-17 shows the rear panel of N3300/N3600. The single node A10 uses a single control unit with a dual node clustered A20 using two control units (see Figure 1-7 on page 17).



Figure 1-16 N3300 rear view




Figure 1-17 N3600 rear view

N3300 is a 2U high device. It has 12 Internal SAS Drive Bays. It can support up to Two External Disk Expansion Units. Each Controller has dual Gigabit Ethernet ports and dual 4Gbps Fibre Channel ports (see Figure 1-18). It also has one console port and one Remote Management port.

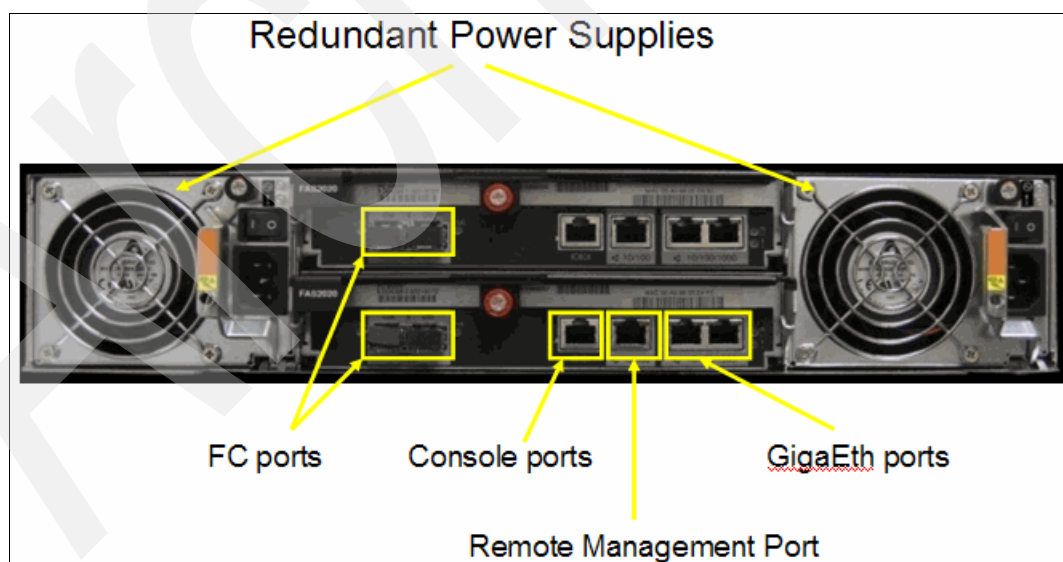


Figure 1-18 External ports on N3300

Tips: The N3300 series supports SAS, FC and SATA disk technologies. 12 SAS disk drives are supported in the controller chassis. The N3300 can be configured with 0 disk drives in the controller and use the storage from disk expansion units like the EXN1000 for SATA or EXN4000 for Fibre Channel disks.

The N3600 also has redundant power supplies (see Figure 1-19).




Figure 1-19 N3600 power supplies and expansion slot




Figure 1-20 external ports on N3600

N3600 is a 4U high device. It has 20 Internal SAS Drive Bays. N3600 can support up to six External Disk Expansion Units. Each Controller has dual Gigabit Ethernet ports and dual 4 Gbps Fibre Channel ports (see Figure 1-20). It also has one console port and one Remote Management port. N3600 has a PCIe Slot on each Controller.

Tips: The N3600 series supports SAS, FC, and SATA disk technologies. 20 SAS disk drives are supported in the controller chassis. The N3600 requires a minimum of six SAS drives in the controller chassis.

N3300/N3600 key specifications:

- ▶ 2U High (N3300) / 4U high (N3600)
- ▶ Up to two External Disk Expansion Units for N3300 and Up to six External Disk Expansion Units for N3600
- ▶ High Performance SAS infrastructure
- ▶ Single Controller or Dual Controller (for HA)
- ▶ Unified Storage: iSCSI, NAS, and Fibre Channel
- ▶ Each Controller: Dual Gigabit Ethernet Ports and Dual 4 Gbps Fibre Channel Ports
- ▶ Onboard Remote Platform Management
- ▶ Internal SAS Drive Bays

N3000 is a small form-factor appliance that conserves scarce and valuable space in data centers or remote office locations. It is engineered for Small to Medium Enterprises.

1.6.2 IBM System Storage N5000 introduction

The N5200, N5300, N5500, and N5600 are suitable for environments that demand data in high availability, high capacity, and highly secure data storage solutions. The IBM System Storage N5000 series offers an additional choice to organizations for enterprise data management. The IBM System Storage N5000 series is designed to deliver high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

- ▶ The IBM N5000 A series comes in four models:
 - N5200
 - 2864-A10 Single Filer
 - 2864-A20 Clustered
 - N5300
 - 2869-A10 Single Filer
 - 2869-A20 Clustered
 - N5500
 - 2865-A10 Single Filer
 - 2865-A20 Clustered
 - N5600
 - 2868 -A10 Single Filer
 - 2868 -A20 Clustered
- ▶ FC or SATA (both may be used behind a single controller but not in the same drawer)

The N5000 A10 models come in a compact 3U rack mountable unit that can coexist in the same rack as a EXN1000, EXN2000, and EXN4000 storage expansion unit (see Figure 1-23 on page 26 and Figure 1-22 on page 25). The A20 models require 6U of space.

There are no visible external differences between the N5200 and N5500. The differences are in maximum storage capacity and CPU processing power, as illustrated in Table 1-2 on page 6. From the front, the N5600 and N5300 also looks very similar to the N5200 and N5500; some of the differences are seen from the rear (Figure 1-25 on page 27), especially

the absence of a LVD SCSI connector. The N5600 and N5300 also uses a BIOS prompt upon boot rather than a Common Firmware Environment (CFE) prompt.



Figure 1-21 N5200




Figure 1-22 N5500 with EXN2000 shelves

Depending on the model you have, you will see one (N5200) or two (N5500) modules internally (see Figure 1-23).




Figure 1-23 N5200 One CPU module

The easily accessible rear of the N5000 series provides I/O connectivity and power supply access and status indications (see Figure 1-24 and Figure 1-24).

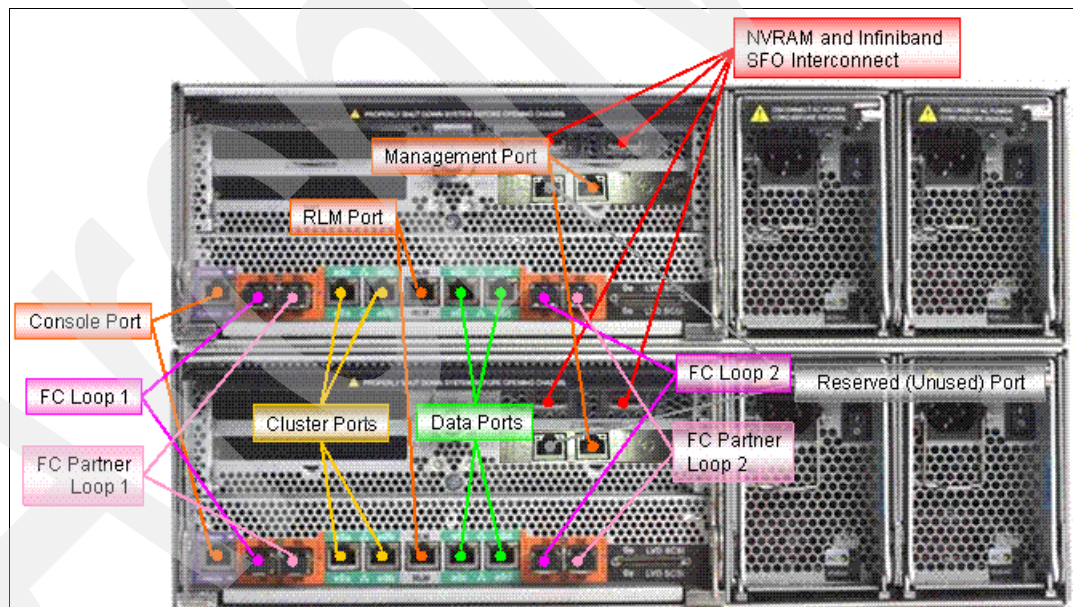


Figure 1-24 Rear View of N5500




Figure 1-25 Rear view of N5600 and N5300

This top view of the N5200 or N5500 (Figure 1-26) shows the modular design and Field Replaceable Unit capabilities.




Figure 1-26 Top view of N5000

The motherboard of the N5200 is a self-contained unit holding components such as memory, the CPU, and interfaces (see Figure 1-27).

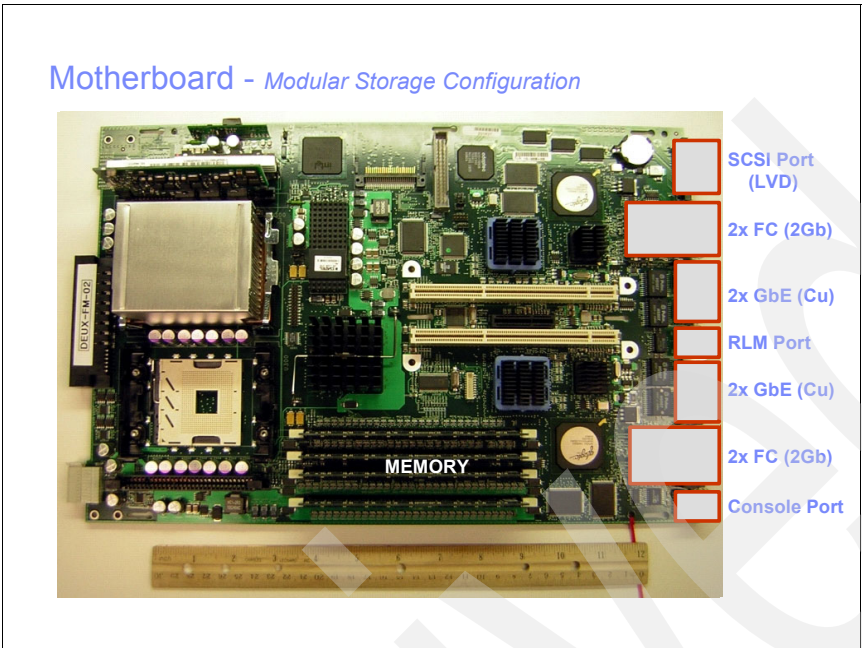


Figure 1-27 N5200 and N5500 motherboard

RAID group sizes

Table 1-10 RAID group size in drive type

Model	FC-AL drives default	FC-AL drives maximum	ATA drives default	ATA drives maximum
RAID 4	8	14	7	7
RAID DP	16	28	14	16

The N5200, N5300, and N5500 multi-disk drive options offer Mission Critical, Near-line, and Compliance Storage Solutions, as shown in Figure 1-28 on page 29.

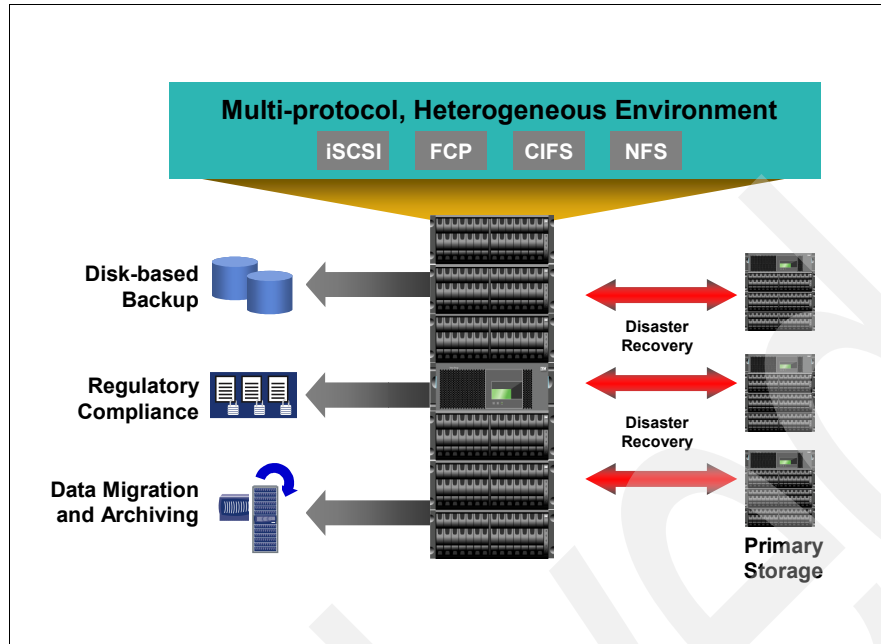


Figure 1-28 Multi storage options

1.6.3 IBM System Storage N7000 introduction

The IBM System Storage N7000 series offers an additional choice to organizations facing the challenges of enterprise data management. The IBM System Storage N7000 series is designed to deliver high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

- ▶ The IBM N7000 A series comes in two Models:
 - N7600
 - 2866-A10 Single Node
 - 2866-A20 Clustered
 - N7800
 - 2867-A10 Single Node
 - 2867-A20 Clustered
- ▶ FC or SATA (both may be used behind a single controller but not in the same drawer)

Like its N5000 predecessor, the N7000 series' unit front has the LCD display and the standard three LEDs indicating system activity, status, and power (see Figure 1-29). Externally, the N7600 and N7800 appear in the SAN; the differences lie internally with the increased CPU, memory, and NVRAM capability of the N7800, as compared to the N7600.



Figure 1-29 Front view of N7000

From the rear of the N7000 you can see the redundant power supplies, the NVRAM card, and the Gigabit Ethernet interfaces, as well as the Fibre Channel interfaces. The console port and RLM port are also located on the rear (see Figure 1-30 on page 31).




Figure 1-30 N 7000 rear view

Each N7000 node requires 6U of rack space, and each expansion unit requires 3U of rack space. Each N7000 node requires at least one expansion unit (see Figure 1-30).



Figure 1-31 N7000 racked

A dual-node N7600 supports a maximum of 60 storage expansion units (EXN1000 or EXN2000). A dual-node N7800 supports a maximum of 72 storage expansion units (EXN1000 or EXN2000) (see Figure 1-33). Each rack holds a maximum of 12 expansion units. The N7000 products are installed by IBM service, not Customer Setup.



Figure 1-32 Clustered N7000 with multiple expansion units

When you remove the bezel, you see the CompactFlash card reader, and directly below it, the Remote LAN Module (RLM). The RLM is required in all N 7000 series systems. The systems will not boot unless the card is present. Also, you will see five fan units. The fans are hot swappable and are numbered here for your reference (Figure 1-33).




Figure 1-33 Front of N7000 with bezel removed

On the side of the system you will notice that there are two handles on each side to help you lift the system (Figure 1-34 on page 33). The system is very heavy. Fully loaded, it weighs 120 pounds. IBM recommends that before lifting the system, you remove the fan units and the two power supplies. This reduces the weight to slightly over 90 pounds. It is recommended that you use three people to lift the system.

Caution

- Remove fan units and power supplies before lifting
- Three people required to lift system





Figure 1-34 Lifting N7000

The two hot swappable power supplies can be seen and removed from the rear of the N7000 (see Figure 1-35).




Figure 1-35 N7000 power supplies

In the middle there are nine-PCI slots. They are numbered 1-9, from left to right, and all the slots are slot-specific (see Figure 1-36). When installing any adapter cards, always use the System Configuration Guide on the IBM site for reference.




Figure 1-36 PCI slots

Below the PCI slots, there is a console port and RLM port (see Figure 1-37 on page 35).




Figure 1-37 RLM and Console ports

Below the Ethernet ports is the Fibre Channel Tray, referred to as the “FC Tray”, with eight on-board Fibre Channel ports (Figure 1-38). This tray is actually a Field Replaceable Unit.




Figure 1-38 Fibre Channel ports

Inside the N7000, looking from the top (see Figure 1-39), you can see the PCI slots and system memory, but you cannot see the processors. They are on the other side of the motherboard tray. Recall that the N7800 has four CPUs and the N7600 has two processors.

From this perspective, you can see the nine PCI slots. You will notice that the slots are color-coded. Slots 3, 4, and 9 are black and represent PCI-X.

Slots 1, 2, 5, 6, 7, and 8 are white and represent PCI-Express.

Notice the NVRAM6 adapter resides in slot 2 on this stand-alone system. If this were an active/ active configuration, the NVRAM6 adapter would reside in slot 1 and would also be used as the cluster interconnect card.




Figure 1-39 Top of N7000

Keep in mind that the N7800 uses an NVRAM6 adapter with 2 GB of memory, and the N7600 uses an NVRAM6 adapter with 512 MB of memory (see Figure 1-40 on page 37).

- **NVRAM6 adapter for N7800 contains 2 GB memory**
- **NVRAM6 adapter for N7600 contains 512 MB memory**



Figure 1-40 NVRAM

There are some new LEDs of which you should be aware: the fan units, PCI slots, and memory DIMMs. The LEDs indicate a failed component (see Figure 1-41).

- **LEDs for Fan units, PCI adapters, and DIMMs**




Figure 1-41 New LEDs

RAID group size

Table 1-11 RAID group size in drive type

Model	FC-AL drives default	FC-AL drives maximum	ATA drives default	ATA drives maximum
N7600 RAID 4	8	14	7	7
N7600 RAID DP	16	28	14	16
N7800 RAID 4	8	14	7	7
N7800 RAID DP	16	28	14	16

1.7 IBM System Storage N series Gateways (G models)

The IBM System Storage N series Gateways, an evolution of the N 5000 series product line, are a network-based virtualization solution that virtualizes tiered, heterogeneous storage arrays, allowing customers to leverage the dynamic virtualization capabilities available in Data ONTAP across multiple tiers of IBM and vendor acquired storage (see Figure 1-42 on page 39). Like all IBM System Storage N series storage systems, the IBM System Storage N series Gateway family is based on the industry-hardened Data ONTAP microkernel operating system, which unifies block and file storage networking paradigms under a common architecture and brings a complete suite of IBM System Storage N series advanced data management capabilities for consolidating, protecting, and recovering mission-critical data for enterprise applications and users.

The industry's most comprehensive virtualization solution, the N series Gateways provides proven and innovative data management capabilities for sharing, consolidating, protecting, and recovering mission-critical data for enterprise applications and users and seamlessly integrates into mission-critical enterprise-class SAN infrastructures. These innovative data management capabilities, when deployed with disparate storage systems, simplify heterogeneous storage management.

The IBM System Storage N series Gateway will present shares, exports, or LUNs that are built on flexible volumes that reside on aggregates. The N series Gateway is also a host on the storage array SAN. Disks are not shipped with the N series Gateway. N series Gateways take storage array LUNs (which are treated as disks) and virtualize them through Data ONTAP, presenting a unified management interface.

The IBM System Storage N series Gateway offers customers new levels of performance, scalability, and a robust portfolio of proven data management software for sharing, consolidating, protecting, and recovering mission critical data. IBM System Storage N series storage systems seamlessly integrate into mission-critical SAN environments and provide a simple, elegant data management solution decreasing management complexity, improving asset utilization, and streamlining operations to increase business agility and reduce total cost of ownership. Organizations that are looking for ways to leverage SAN-attached storage create a consolidated storage environment for the various classes of applications and storage needs throughout their enterprise. These prospects are looking for ways to increase utilization, simplify management, improve consolidation, enhance data protection, enable rapid recovery, increase business agility, deploy heterogeneous storage services, and broaden centralized storage usage by provisioning SAN capacity for business solutions requiring NAS, SAN, or IP SAN data access (see Figure 1-43 on page 39).

These prospects have:

- ▶ Significant investments or a desire to invest in a SAN architecture
- ▶ Excess capacity or an attractive storage cost for SAN capacity expansion
- ▶ Increasing requirements for both block (FCP and iSCSI) and file (NFS, CIFS, and so on) access
- ▶ Increasing local or remote shared file services and file access workloads

They are seeking solutions to cost effectively increase utilization, consolidate distributed storage, Direct Access Storage, and file services to SAN storage, simplify storage management, and improve storage management business practices.

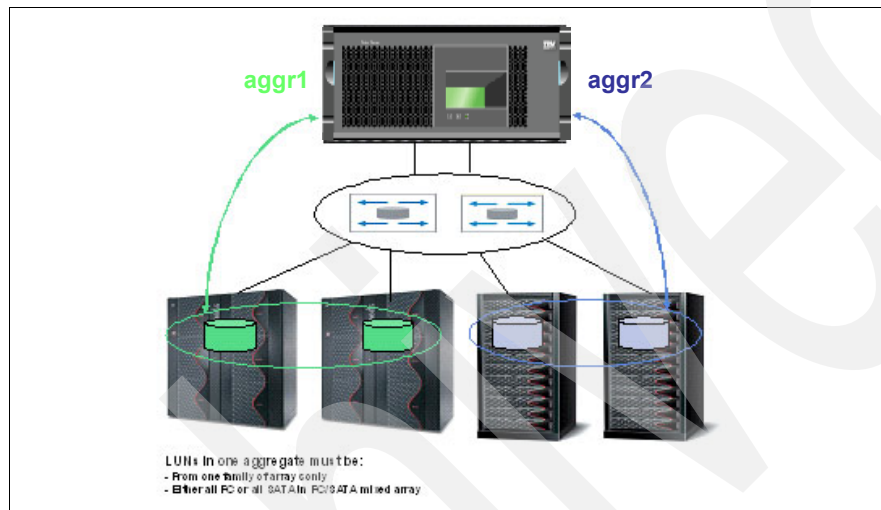


Figure 1-42 Heterogeneous storage

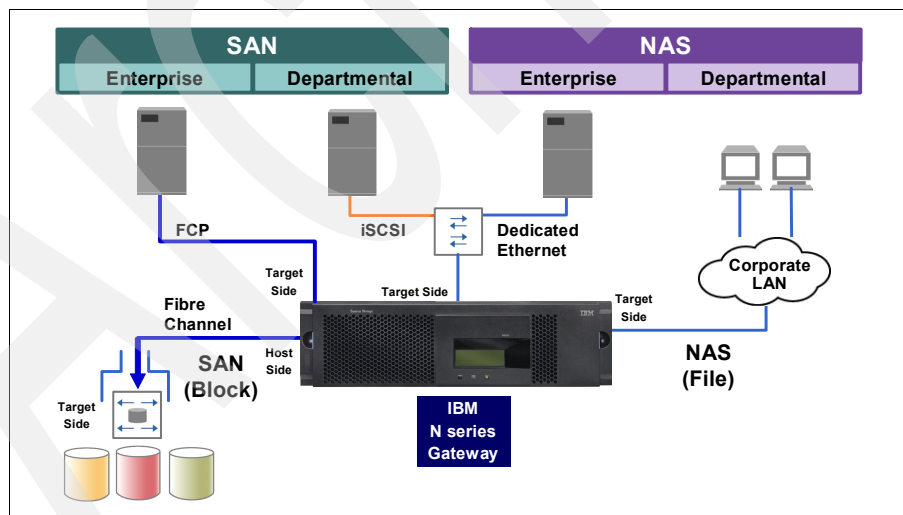


Figure 1-43 Gateway Topology

Two halves to set up

A IBM System Storage N series Gateway implementation can be thought of as a front-end implementation and a back-end implementation. The front-end setup includes configuring the N series Gateway for all protocols (NAS or FCP), implementing any snap features (Snapshot, SnapMirror, SnapVault, and so on), and setting up backup, including NDMP dumps to tapes.

The back-end implementation includes all tasks required to set up the IBM System Storage N series Gateway system up to the point where it is ready for Data ONTAP installation. These tasks include array LUN formatting, port assignment, cabling, switch zoning, assigning LUNs to the V-Series system, creating aggregates, and loading Data ONTAP.

1.7.1 IBM System Storage N series Gateway highlights

IBM System Storage N series Gateway systems provide a number of key features that enhance the value and reduce the management costs of utilizing a Storage Area Network (SAN). A IBM System Storage N series Gateway system will do the following:

- ▶ Simplifies storage provisioning and management.
- ▶ Lowers storage management and operating costs.
- ▶ Increases storage utilization.
- ▶ Provides comprehensive simple-to-use data protection solutions.
- ▶ Improves business practices and operational efficiency.
- ▶ Transforms conventional storage systems into a better managed storage pool (see Figure 1-44).




Figure 1-44 Tiered Heterogeneous Storage

1.7.2 Gateway RAID

Gateways use RAID 0 on top of RAID 1, RAID 5, or RAID 10 on RAID storage subsystems (see Figure 1-45 on page 41). Physical disk operations, such as scrubbing, are disabled.




Figure 1-45 RAID configuration

RAID 0 is used to write data (see Example 1-1 for an example of a volume's status with the Gateway; it looks very similar to what you would see on a IBM System Storage N series model A except for the RAID status).

Example 1-1 Vol status with gateway volumes

```

itsotuc2*> vol status -v vol3
      Volume State      Status      Options
      vol3 online      raid0, flex  nosnap=on, nosnapdir=off,
                                     minra=off,
                                     no_atime_update=off,
                                     nvfail=off,
                                     snapmirrored=off,
                                     create_ucose=on,
                                     convert_ucose=on,
                                     maxdirsize=31457,
                                     fs_size_fixed=off,
                                     guarantee=volume,
                                     svo_enable=off,
                                     svo_checksum=off,
                                     svo_allow_rman=off,
                                     svo_reject_errors=off,
                                     fractional_reserve=100,

      Containing aggregate: 'aggr0'

      Plex /aggr0/plex0: online, normal, active
      RAID group /aggr0/plex0/rg0: normal
  
```

1.7.3 IBM N5200, N5300, N5500, and N5600 Gateway models

The N5000 Gateway models are a good value for those wishing to extend the reach of their SANs. The N5000 Gateway incorporates a variety of reliability and availability features designed to support high demand operations. It houses hot swappable, redundant power supplies and fans, and supports multipath failover protection and host dual pathing between the unit and its SAN-attached storage device. In addition, the clustering feature between two storage systems is designed to help reduce system downtime. From a hardware perspective,

the G10 and G20 models are identical to the A10 and A20 models of the N5200, N5300, N5500, and N5600. The differences lie in the spectrum of Data ONTAP features supported and enabled. The models are:

- ▶ N5200
 - 2864-G10
 - 2864-G20 Clustered model
- ▶ N5300
 - 2869-G10
 - 2869-G20 Clustered model
- ▶ N5500
 - 2865-G10
 - 2865-G20 Clustered model
- ▶ N5600
 - 2868-G10
 - 2868-G20 Clustered model

Table 1-12 shows the capacities for the various Gateway models.

Table 1-12 Gateway capacity

Model	Maximum capacity
2864-G10	50 TB
2864-G20	50 TB per node
2869-G10	126 TB
2869-G20	126 TB
2865-G10	80 TB
2865-G20	80 TB per node
2868	252 TB
2868	252 TB

Important: If you are going to enable the `cf.takeover.on_panic` option, ensure that a spare LUN is available for core dumps. If the `cf.takeover.on_panic` option is enabled and no spare LUN is available, no core dump file is produced upon failure. (The `cf.takeover.on_panic` option controls whether a cluster partner immediately takes over for a panicked partner.)

Table 1-13 on page 43 shows LUN information for the various Gateway models.

Table 1-13 LUN information for Gateway models

Model	Maximum number of LUNs	Minimum number of LUNs
N5200 2864-G10 (non cluster model)	168.	A stand-alone gateway must own at least one LUN. A cluster configuration must own at least two LUNs.
N5200 2864-G20 (cluster model)	For each node, single node N5200 2864-G10 values apply.	
N5300 2869 - G10	252.	
N5300 2869 - G20	For each node, single node N5300 2869-G10 values apply.	
N5500 2865-G10 (non cluster model)	336.	
N5500 2865-G20 (clustermodel)	For each node, single node N5500 2865-G10 values apply.	
N5600 2868-G10 (non cluster model)	504.	A stand-alone gateway must own at least one LUN.
N5600 2868-G20 (clustermodel)	For each node, single node N5600 2868-G10 values apply.	A cluster configuration must own at least two LUNs.

1.7.4 IBM Gateway models N7600 and N7800

The IBM System Storage N7000 series Gateway models offers additional choice to organizations facing the challenges of enterprise data management. The IBM System Storage N7000 series is designed to deliver high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy. The IBM System Storage N series Gateway models N7600 and N7800 deliver all the feature function that the N5000 series does, but with increased processing, memory, NVRAM, and total storage capacity. The N7600 and N7800 are designed with the high end of the enterprise environments in mind. The N7000 series Gateway hardware is identical to the A10 and A20 models, with the difference being in the enabled features and disk attachment by Data ONTAP.

► The IBM N7000 A series comes in two models:

- N7600
 - 2866-G10 Single Node
 - 2866-G20 Clustered
- N7800
 - 2867-G10 Single Node
 - 2867-A20 Clustered
- 2865-G20 Clustered model

1.7.5 LUN sizing

Gateway support for LUN sizes is as follows:

- ▶ Maximum LUN size: 500 GB
- ▶ Minimum LUN size: 1 GB

Note: The Data ONTAP definition of a GB is as follows: one GB is equal to 1000 x 1024 x 1024 bytes. Therefore, the maximum LUN size that Data ONTAP supports means $500 * 1000 * 1024 * 1024 = 524,288,000,000$ bytes.

1.7.6 LUN mapping

Storage Subsystem LUNs are converted to disks for the IBM System Storage N series Gateway systems (see Figure 1-46). The Gateway disk count is equal to the LUN count.

Figure 1-46 is a example of a Array LUN mapped to Gateway disk.

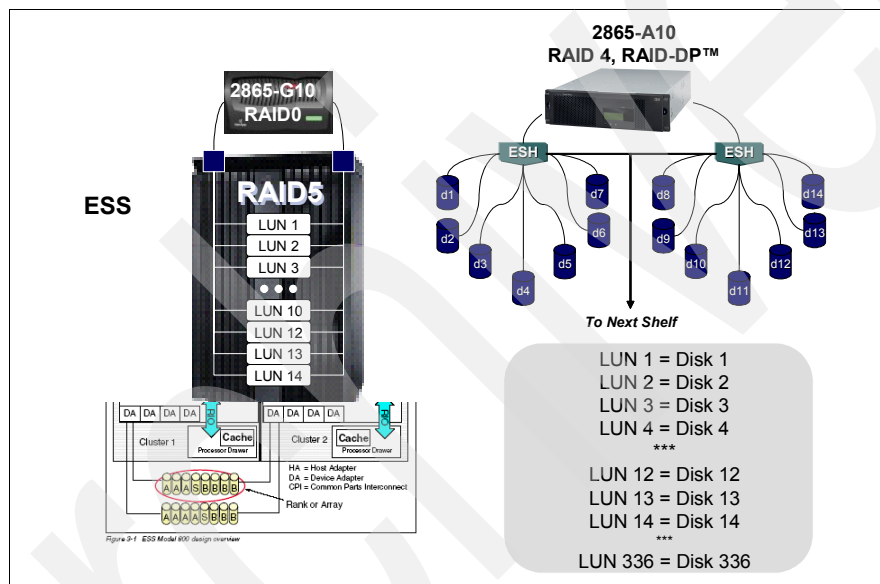


Figure 1-46 LUN to N series Gateway disk relationship

LUNs are added to the Gateway through the same volume wizard we use on the IBM System Storage N series A models (Figure 1-47 on page 45).

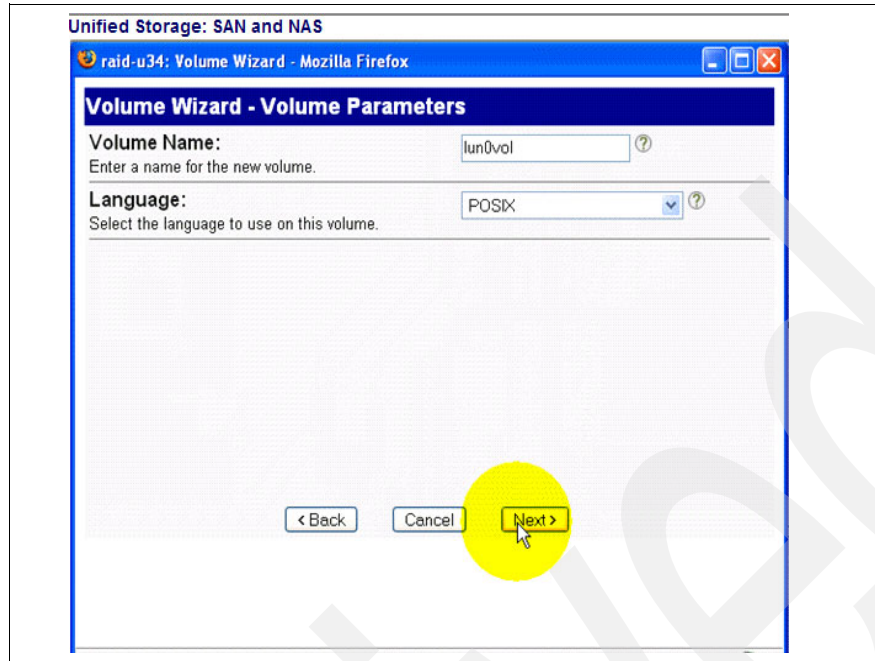


Figure 1-47 Volume wizard

1.8 Interoperability between G and A models

- Replication between SnapMirror on a G model and SnapMirror on an A model includes async, semisync, and synchronous (Figure 1-48).

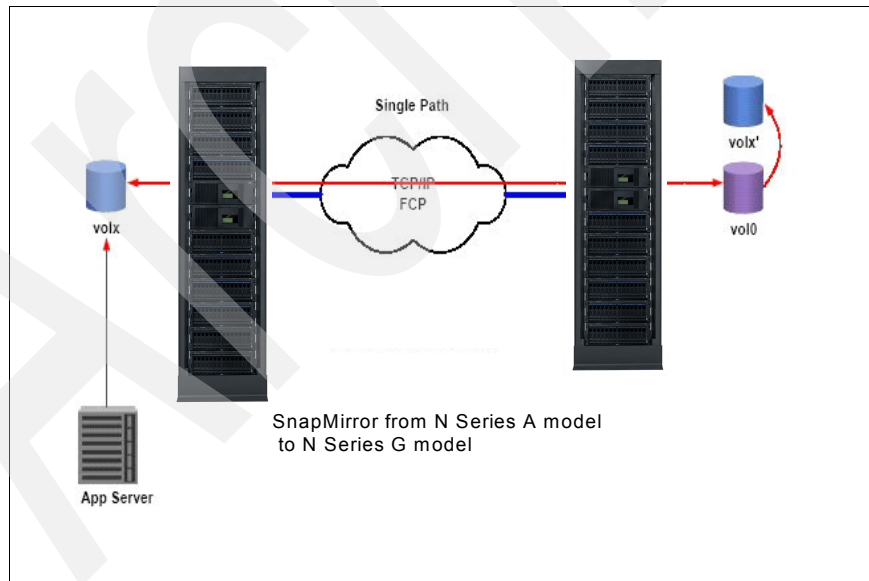


Figure 1-48 SnapMirror Interoperability

- There is disk-to-disk backup from the SnapVault primary on a G model to the SnapVault secondary on an A model (Figure 1-49).




Figure 1-49 SnapVault interoperability

- There is disk-to-disk backup from the SnapVault primary on an A model to the SnapVault secondary on a G model.

1.9 EXN2000

The EXN2000 is a fibre expansion unit for the A models of the IBM System Storage N series. The EXN2000 looks very similar to the N3700, but unlike the N3700, which has the CPU modules, the EXN2000 supports only the disk modules and the connectivity to them (see Figure 1-50).



Figure 1-50 EXN2000

The EXN2000 is identical to the N3700 chassis except that the slot holding the CPU tray is replaced with an Electronically Switched Hub (ESH2). ESH2 provides a point to point connection to the drives rather than the traditional arbitrated loop. This is illustrated in Figure 1-51 on page 47. The maximum number of drives per shelf is unaffected by the capacity of the individual drive modules. Mixing of drives of different capacity in the same shelf is not recommended because of the effects it has on sparing, RAID groups, and flex

volumes. The maximum number of drives on a loop are 84 or six shelves using the ESH2 module.




Figure 1-51 Rear of the EXN2000 schematic: Arbitrated loop versus Switch hub

1.9.1 Switched Hub architecture

Switched Hub architecture has the benefit of additional availability, boosted performance in high I/O environments, and more powerful diagnostic abilities (see Figure 1-52). Figure 1-53 on page 48 shows the ESH2 module. From a purely technical viewpoint, Fibre Channel loops support 126 devices. From a practical position, traditional FC-AL daisy-chain topologies (for example, loop resiliency circuits) require limits on the number of devices for performance reasons. The performance impact is directly attributable to loop overhead traffic. In the past, recommendations for LRC topologies is 56 devices per loop. Advanced FC-AL topologies allow the “cost” of loop overheads to be minimized, thereby increasing the number of supported disk drives. This is true for system configurations that include switched Hub architecture. Switched Hub architecture is a hub and spoke arrangement with local neighborhoods surrounding each device. Loop overhead is minimized by the fact that traffic no longer flows through each disk drive. The hubs are capable of local communication to the disk drives and therefore more efficiently convey this information to the storage system. There are two Fibre Channel ports on each module. The PS/2 port is for IBM service only and provides no functionality. The units have LED status lights that indicate speed and fault status are hot swappable, thereby allowing maximum availability.

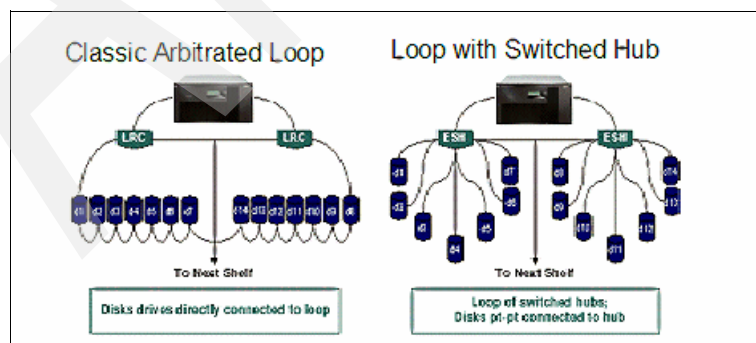


Figure 1-52 Switched Hub architecture




Figure 1-53 External ports on ESH2 module

EXN2000 was withdrawn from marketing May 22, 2007

1.9.2 EXN1000

The EXN1000 uses the same shelf and hardware as the EXN2000 and EXN4000, so it has the same dimensions. It also supports the same number of disks per shelf (14) (see Figure 1-54). The main differences are:

- ▶ It uses drive type SATA instead of Fibre Channel.
- ▶ It uses the AT-FCX interface module instead of the ESH2.



Figure 1-54 EXN1000 expansion unit

AT-FCX refers to the controller module (Figure 1-55) of the serial advanced technology attachment (SATA) storage expansion unit.



Figure 1-55 AT-FCX module

Data ONTAP supports up to 400 RAID groups per storage system or cluster. When configuring your aggregates, keep in mind that each aggregate requires at least one RAID group and that the total of all RAID groups in a storage system cannot exceed 400.

1.9.3 EXN4000

The EXN4000 uses the same shelf and hardware as the EXN2000, so it has the same dimensions. EXN4000 also supports the same number of disks per shelf (14). EXN4000 uses ESH4 as its controller module. ESH4 refers to the third-generation, multiloop speed ESH module. ESH4 can function at 1 Gb, 2 Gb, or 4 Gb loop speed when it works with EXN4000. The ESH4 has LEDs that indicate whether the module is functioning normally (refer to Figure 1-58 on page 50), whether there are any problems with the hardware, and the loop speed operation of the EXN4000. The main differences are:

- ▶ A 4 Gbps capable Fibre Channel (FC) disk enclosure, that is, twice the maximum loop bandwidth of EXN2000
- ▶ Higher bandwidth for heavy sequential workload
- ▶ Fewer HBAs or slots used to achieve higher bandwidth needs

The EXN4000 FC Storage Expansion Unit will run at 2 Gbps FC when attached to systems that do not have 4 Gbps capability. It can be added to EXN2000 FC loops.


Figure 1-56 shows the front of the EXN4000 expansion unit, while Figure 1-56 shows the rear of the unit.



Figure 1-56 EXN4000 expansion unit



Figure 1-57 2xESH4, 2xPSU/Fans



14

Figure 1-58 Location of the LEDs for an ESH4

EXN4000 is the replacement for EXN2000 FC Storage Expansion Unit.



Introduction to IBM Lotus Domino 7

This chapter is meant for the storage administrator or anyone unfamiliar with IBM Lotus Domino. It gives a brief introduction to IBM Lotus Domino 7 and its architecture, mailbox, and database design.

2.1 IBM Lotus Domino 7 overview

IBM Lotus Domino is a collaboration solution that goes far beyond messaging. It contains directory services, integrated e-mail, calendaring, scheduling, discussion databases, and multiple address book capabilities (see Figure 2-1). Additionally, you can build and run collaboration solutions, such as Customer Relationship Management (CRM), project management, or document management, that help you to be more efficient.

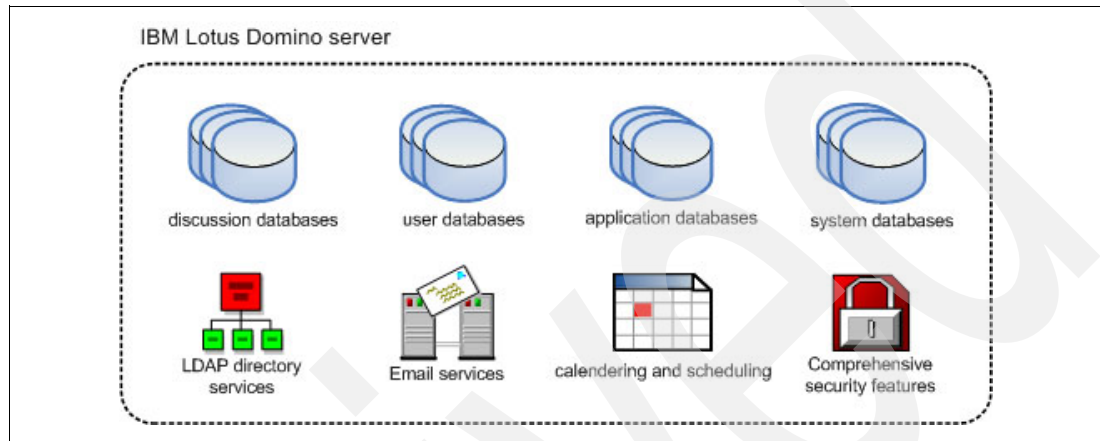


Figure 2-1 Inside Lotus Domino server

For enterprise needs, Lotus Domino includes comprehensive real-time replication, clustering, and load balancing features. Figure 2-2 gives you an example of a common Lotus Domino Infrastructure.

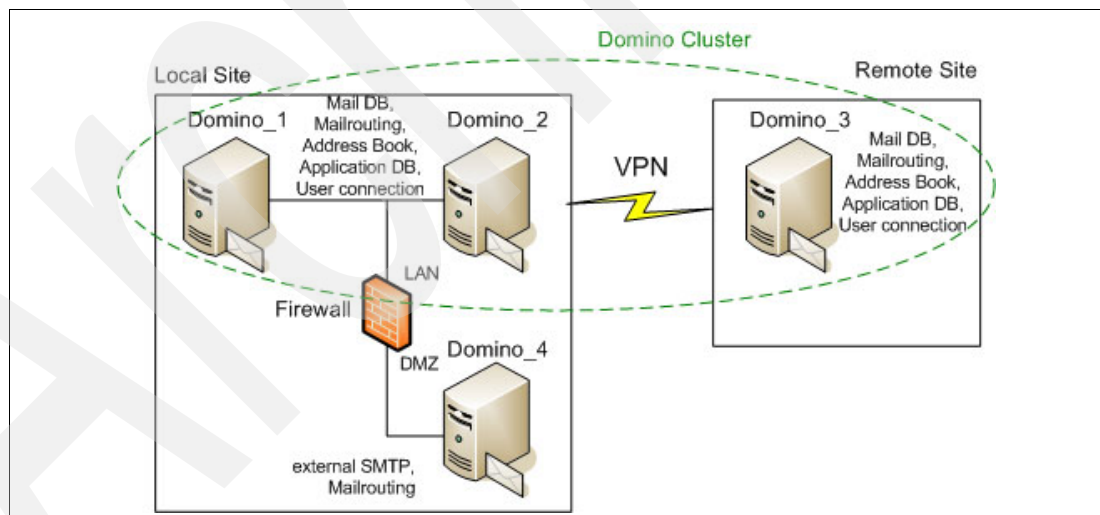


Figure 2-2 IBM Lotus Domino infrastructure with two sites

IBM Lotus Domino software is known for backward compatibility and smooth upgrades, even in complex environments. The enhanced support for open standards and Web services in Version 7 let you extend the use of your IBM Lotus Domino application (and investment) for a longer period of time. New technologies for publishing information, such as Really Simple Syndication (RSS), online journals (blogs), or Web services that are based on Simple Object Access Protocol (SOAP), are available.

The design of IBM Lotus Domino contains various security features to prevent unauthorized access. For details, see 2.2.4, “Security” on page 57.

Note: For more information about IBM Lotus Domino, see:

<http://www.lotus.com/products/product4.nsf/wdocs/dominohomepage>

2.1.1 Terms and definitions

This chapter contains important definitions to help you better understand IBM Lotus Domino:

- ▶ *Notes database*

A Notes database is a single, physical file that contains a set of documents and a copy of the application design elements.

- ▶ *Application design elements*

The application design elements control the creation and modification of documents in a Notes database.

2.1.2 Server options

IBM provides several Lotus Domino Server license options, depending on your company needs:

- ▶ Lotus Domino Messaging Server

This license option allows you to deploy a company-wide e-mail and scheduling infrastructure, including basic collaboration tools, such as a discussion database and more.

- ▶ Lotus Domino Enterprise Server

The Enterprise Server contains all of the features of the Messaging Server, plus clustering, partitioning, load balancing, and integrated administration and system management tools.

- ▶ Lotus Domino Utility Server

If you do not need e-mail and calendaring features, but you need non-mail applications for a large number or unknown number of users, this is a cost-effective license model. It is typically used for Web application deployment inside and outside of your organization. Web browser access to non-mail applications does not need client access licenses.

- ▶ Lotus Domino Express

The Express option is designed for companies with 1,000 employees or fewer and provides a feature-rich enterprise class collaboration tool at a small and medium business price.

Note: For more information about IBM Lotus Domino Server options, see:

<http://www.lotus.com/products/product4.nsf/wdocs/dominooverview>

2.1.3 Client options

For individual needs, there are various ways to get connected to a IBM Lotus Domino server. The primary client options are:

- ▶ IBM Lotus Notes®

This is the recommended and full-featured rich client for IBM Lotus Domino. It allows users to access their e-mails, calendar, tasks, and collaboration tools. Through replication, it provides off-line access to desired databases.

- ▶ IBM Lotus Domino Web Access

It allows you to access the IBM Lotus Domino databases with various support for Windows, Linux, and Mac OS Web browsers.

- ▶ IBM Lotus Domino Access for Microsoft Outlook®

Users, who cannot or will not use IBM Lotus Notes for accessing IBM Lotus Domino, can use the familiar Microsoft Outlook client. This is even a good choice for migration from Microsoft Exchange to IBM Lotus Domino.

- ▶ IBM Lotus Domino WebMail

You can get basic e-mail features through standard internet mail protocols, such as Post Office Protocol (POP3) or Internet Message Access Protocol (IMAP).

2.2 Key features

This section discusses the important features of IBM Lotus Domino.

2.2.1 Enterprise messaging system

IBM Lotus Domino provides several features for messaging:

- ▶ This product provides a strong client/server messaging backbone, which is scalable and reliable.
- ▶ There is support for different internet mail standards and protocols, such as IMAP, POP3, and SMTP.
- ▶ Lightweight Directory Access Protocol (LDAP) support is available for addressing external users in third-party Web-based directories, such as Four11 and Bigfoot. Also, other LDAP clients can query the Domino directory.
- ▶ You have access to mail from a browser or over the internet.
- ▶ Lotus Domino can coexist with other mail systems.
- ▶ This product includes the ability to mail-enable applications that are developed in Lotus Domino.

2.2.2 Platform independent

One of the key IBM Lotus Domino competitive strengths has always been its support for multiple operating system platforms. Domino 7 maintains and extends this tradition.

Table 2-1 on page 55 and Table 2-2 on page 56 summarize the various operating system platforms that support IBM Lotus Notes and IBM Lotus Domino 7.

Table 2-1 Lotus Notes clients

Platform	Microsoft Windows 2000	Microsoft Windows XP	Macintosh
Supported operating system versions	Microsoft Windows 2000 Professional	Microsoft Windows XP Professional; Microsoft Windows XP Tablet PC Edition 2005 (Digital Ink Input will not be supported in the Notes client.)	Macintosh OS Power PC® 10.4.2 Intel® 10.4.4
RAM	128 MB minimum 256 MB or more recommended	128 MB minimum 256 MB or more recommended	128 MB minimum 256 MB or more recommended
Disk space (The minimum amount is the disk space that is required for installing default files. More disk space is required if databases are replicated locally or copied locally.)	512 MB or more recommended	512 MB or more recommended	512 MB or more recommended
NetBIOS over IP	Yes	Yes; No (64-bit)	No
TCP/IP	Yes	Yes	Yes

Important: The Domino Administrator Client and Domino Designer are not yet supported on Macintosh OS.

Table 2-2 IBM Lotus Domino server

Platform	Microsoft Windows	Linux	Unix
Supported operating system versions	Microsoft Windows 2000 Server; Microsoft Windows 2000 Advanced Server Microsoft Windows 2003 Server Standard Edition; Microsoft Windows 2003 Server Enterprise Edition; Microsoft Windows 2003 Server x64 Edition	Novell SUSE Linux Enterprise Server (SLES) 8 Novell SUSE Linux Enterprise Server (SLES) 9; Novell SUSE Linux Enterprise Server (SLES) 10 (32- and 64-bit); Red Hat Enterprise Linux (RHEL) 4 (32- and 64-bit)	IBM AIX 5L™ V5.2 (64-bit kernel) IBM AIX 5L V5.3 (64-bit kernel) Sun™ Solaris 9 (64-bit kernel) Sun Solaris 10 (64-bit kernel)
RAM	256 MB minimum 512 MB or more recommended per CPU	512 MB minimum; 512 MB or more recommended per CPU	512 MB minimum; 512 MB or more recommended per CPU
Disk space (The minimum amounts are the disk space required for installing default files. More disk space is required if databases are replicated locally or copied locally.)	1.5 GB minimum per partition	1.5 GB minimum 1.5 GB or more recommended	1.5 GB minimum 1.5 GB or more recommended
NetBIOS over IP	Yes	No	No
TCP/IP	Yes	Yes	Yes

Note: For a complete and up-to-date list of platform and system requirements for IBM Lotus Domino 7.0.2 and IBM Lotus Notes 7.0.2, see:

http://www-12.lotus.com/ldd/doc/domino_notes/7.0.2/rn702.nsf

2.2.3 High availability and load balancing

With its clustering technology, IBM Lotus Domino provides fail-over protection to IBM Lotus Notes clients in the event that any server goes down. Real-time replication keeps the data on all servers synchronized.

Beside the cluster configuration, IBM Lotus Domino dynamically balances the user load between several servers. For that task, the cluster member server constantly monitors its own workload. It also polls all the other servers in the cluster to determine their workload. When the workload on a server exceeds a certain level designated by the administrator, the server becomes "busy", and the Domino server rejects subsequent database open requests until the workload falls back below the specified level. If the cluster-aware IBM Lotus Notes client gets this "busy" answer, it asks the Cluster Manager for next least heavily used server that also holds the requested database.

In addition, IBM Lotus Domino supports various other load balancing features such as server availability threshold, quantity of user per server, or user's home server.

Note: For more information about Workload balancing with Domino clusters, see:
http://www-128.ibm.com/developerworks/lotus/library/ls-Workload_balancing

2.2.4 Security

IBM Lotus Domino has a comprehensive security design from the beginning of its development in 1984. Some of the key security features are:

- ▶ User ID files that contain encryption keys and certificates that IBM Lotus Domino servers use to verify the authorization of the person who use it. Additionally, these ID files can be secured with an individual password.
- ▶ Communication between IBM Lotus Domino and IBM Lotus Notes is secured by encryption.
- ▶ The server security is based on the information stored in Domino Directory. It specifies which user, server, or group has rights to create new databases, can use pass through connections, or create new replicas.
- ▶ The Lotus Domino server implements role concepts for administration and access.
- ▶ Every database has its own Access Control List (ACL) that specifies which role, user, and server has the right to access the database and perform a task on it.
- ▶ All design elements could be signed and protected with an ID file, down to a single viewing-field.

2.3 Notes Storage Files (NSF) and transaction logging

As we discuss IBM System Storage N series with IBM Lotus Domino 7, it is important to understand how Domino works with its data.

The data of every user or application is stored in an independent Notes Storage File (NSF) (see Figure 2-3). This makes IBM Lotus Domino very robust against database corruptions on unexpected hardware or software failure; not every database file is affected.

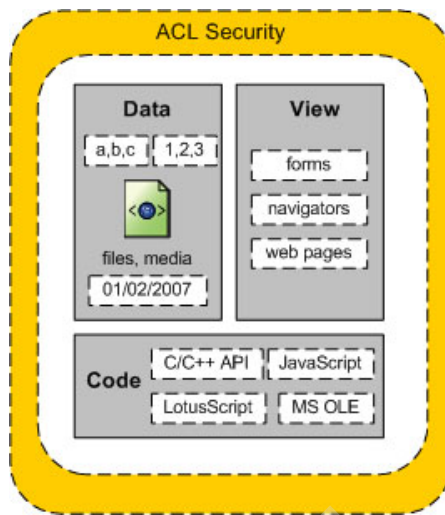


Figure 2-3 Inside the NSF

Note: Since IBM Lotus Domino 7, you have the option to use DB2® instead of NSF on a per-server and per-database basis. For more information about DB2 support in IBM Lotus Domino 7, visit:

<http://www.lotus.com/products/product4.nsf/wdocs/nsfdb2>

The major features of the NSF are:

- ▶ **Shared**
Several users are able to access one database at the same time.
- ▶ **Distributed**
The databases may exist on several servers at the same time. Their content can be replicated on schedule or in real time.
- ▶ **Mail enabled**
Every document in a database can be sent to another database.
- ▶ **Object oriented**
In contrast to relational database systems, IBM Lotus Domino does not use tables for saving information. Information is saved as objects with masks and views on it. Thus, the databases are very flexible and are able to save different kinds of information, such as text, pictures, time, or video.

Every database change is recorded to a transactional log file. This enables Domino to maintain the integrity of its databases in the event of a power loss or other uncontrolled shutdown and to restart quickly after such a failure. The transactional logs can also participate in your backup strategy. While this feature is optional, we recommend it as a best practice.

A Lotus Domino server requires disk space for several purposes. It could be divided in three parts:

- ▶ Disk space for the operating system, IBM Lotus Domino application files, and swap area
- ▶ Disk space for the IBM Lotus Domino data directories and users mail databases
- ▶ Disk space for the IBM Lotus Domino transactional log files

The directories for user data and transaction log files are independent of your operating system choice, while the space for operating system and swap area may vary. Figure 2-4 gives you an high-level overview of an common storage segmentation.




Figure 2-4 Common storage segmentation

The three storage volumes *a*, *b*, and *c* represent the required disk space mentioned above and are described as follows:

1. Volume *a* holds the operating system and IBM Lotus Domino program files. It is recommended that these disks are protected with a local RAID 1.
2. Volume *b* holds the IBM Lotus Domino user database files. Accessing these files will create heavy random I/O operations. The size of the database volume depends on the mailbox size of your user.
3. Volume *c* holds the optional but recommended transactional log files. Every change to the Domino database files (NSF) are written to a transactional log file before writing them to the actual NSF. Because these log files are written in a serial sequence, the I/O profile of the transactional log volume differs from the database volume. Separated and fast disks dedicated for transactional logging are the keys for increasing overall performance of your IBM Lotus Domino server.

Note: If you decided to use transactional logging, it is absolutely essential to place the transactional log directory on a separate physical device.

Archived

Introduction to Microsoft Exchange

Today more than ever, people and businesses are dependent on the constant availability of e-mail systems such as Microsoft Exchange. The explosion of e-mail content, its business-critical nature, and the enforcement of new compliance regulations have created a renewed interest in examining storage infrastructures that service messaging systems. The increasingly stringent requirements of Exchange are a substantial challenge to IT organizations, as they struggle with maintaining and improving Exchange's quality of service while operating within tight IT budgets, an already lean IT staff, and limited time available to perform extensive data management functions, such as backup and restore operations.

Direct-attached storage (DAS) systems are a common storage scenario for Microsoft Exchange. However, DAS has fundamental limitations that make it very difficult to meet client requirements for demanding Exchange environments. Five key steps can be taken to transform a DAS-based infrastructure to a new network storage-based architecture that meets the availability, scalability, and manageability requirements of highly demanding Exchange environments, with the cost structure and ease of use appropriate for Windows.

3.1 Challenges of DAS-based Exchange infrastructures

Earlier DAS-based storage infrastructures for Microsoft Exchange typically consist of multiple servers, each with direct-attached storage. Typical DAS-based Exchange environments utilize tape for backup and recovery, with tape drives connected either directly or through a network (see Figure 3-1).




Figure 3-1 Earlier DAS-based Exchange servers

Direct-attached storage infrastructures present inherent issues for meeting cost and service level requirements for Microsoft Exchange:

- ▶ **Poor scalability and capacity utilization:** Administrators face a frequent need to increase the capacity of their Microsoft Exchange-related storage to accommodate both rapid growth in e-mail stores and the demanding I/O requirements of Microsoft Exchange. Storage utilization with DAS is often as low as 30%, implying substantial investment in unused capacity.
- ▶ **Disruptive backups and long recovery times:** Administrators struggle with compressed or nonexistent Microsoft Exchange backup windows. Perhaps more importantly, it can take several hours to several days to recover either the contents of individual user mailboxes or entire mailbox stores from tape, a factor that makes tape-based recovery times inconsistent with user requirements. Corrupted backup tapes is also a common situation when recovering Exchange systems.
- ▶ **Disaster Recovery infrastructure is problematic:** For many clients utilizing DAS, advanced disaster recovery (DR) capabilities, such as remote mirroring are impractical to implement due to the inherent complexity of replicating and managing the DAS and tape infrastructure.
- ▶ **Complexity surrounding e-mail retention and regulatory compliance:** With the continued growth of e-mail stores, the importance of retaining e-mails as business records for reference and legal protection, and new regulatory requirements that put added focus on preserving critical e-mail trails, administrators are compelled to deploy more costly storage. At the same time, administrators must manage retention policies that require more planning and manual intervention to assess when e-mail messages should be deleted, moved, or archived to other platforms.

For many clients, including those that are migrating from previous releases to Microsoft Exchange 2003, it is appropriate to consider whether to upgrade their Exchange storage infrastructure to better support business goals.

3.2 Five steps to leverage your Microsoft Exchange infrastructure

Companies that are willing to overcome the limitations of DAS and tape can quickly move their Exchange operations to a new quality of service level by implementing the five steps outlined in this chapter. This five-step process can transform a DAS-based Exchange infrastructure into an environment that effectively addresses key business requirements.

The five steps are:

1. Deploy either an IP Storage Area Network (SAN) or a Fibre Channel (FC) SAN.
2. Implement N series Snapshot for backup and recovery.
3. Virtualize storage for faster provisioning and better capacity utilization.
4. Implement low-overhead remote mirroring for DR.
5. Implement tiered storage for archival and compliance.

Depending on the needs of their environment, some companies will take only the first few steps, while those companies that have demanding business requirements, including compliance and DR, will complete all of the steps.

3.2.1 Deploy either an IP SAN or an FC SAN

Many organizations that need to simplify and improve their Microsoft Exchange environment will realize substantial benefits by consolidating their existent DAS infrastructures to a SAN, based on either the iSCSI (IP SAN) or FC (FC SAN) data transfer protocol. These storage networks allow storage to be moved from each Exchange server for more centralized and efficient management and control of disk resources, as shown in Figure 3-2.




Figure 3-2 Exchange servers accessing IBM N series SAN

For Microsoft Windows in general, and Exchange environments in particular, many companies consider IP SANs to be a particularly good fit. IP SANs offer a more affordable option for upgrading to a networked storage architecture by using low-cost, standards-based commodity components and enabling IT organizations to make use of both existing investments in IP infrastructure and their current IP expertise.

A SAN that is based on FC also provides the benefits of storage networking and will be more suitable for those who prefer FC technology or have already deployed FC.

3.2.2 Implement N series Snapshot for backup and recovery

Tape media is often used to back up Microsoft Exchange environments. Although tape media is a low-cost means of data protection, tape backup has fundamental limitations that make it suboptimal for environments that require nondisruptive backup and rapid restore. A minor recovery from tape can take hours, while recovery from a more catastrophic failure can easily take days. The use of Snapshot technology addresses these issues by allowing administrators to take point-in-time copies of data and save them to disk.

Snapshot technology is uniquely efficient in its use of disk space compared to other snapshot implementations in the industry, which tend to consume large quantities of storage. This unique efficiency is because Snapshot copies consist of only data updates and additions, not full copies of the data. More importantly, Snapshot incurs no I/O or performance penalties on Exchange.

SnapManager for Exchange (SME), in conjunction with SnapRestore, further automates and streamlines the coordination of backups and restores from earlier snapshots. Fully integrated with Microsoft Volume Shadow Copy Service (VSS), SnapManager for Exchange ensures that snapshots taken of the Exchange database will be consistent and able to be fully restored.

Single Mailbox Recovery for Microsoft Exchange (SMBR) fully automates the process to restore a single mailbox or even a single e-mail message or contact from a mailbox without the need to use Exchange Recovery Storage Groups (RSG).

3.2.3 Virtualize storage for faster provisioning and better capacity utilization

To deal with exponential growth of Microsoft Exchange message store needs and Exchange's high I/O requirements, administrators in a DAS-based environment often find themselves bringing Exchange systems offline in order to add a new server with additional disk storage. The actual issue with adding more storage for Exchange is not only the amount of additional disk space required, but the number of physical disk spindles that Exchange requires to support its high I/O demands. To meet I/O needs, DAS-based Exchange clients often acquire sufficient disk spindles but then underutilize much of the acquired disk capacity. This sizing for performance and cumbersome provisioning result in DAS environments that often have capacity utilization as low as 30%.

While promising better utilization than DAS, many traditional SAN solutions often suffer from the same inability to maximize the use of all available disk spindles and all available storage. Usage of disks on a traditional SAN solution is shown on Figure 3-3 on page 65.

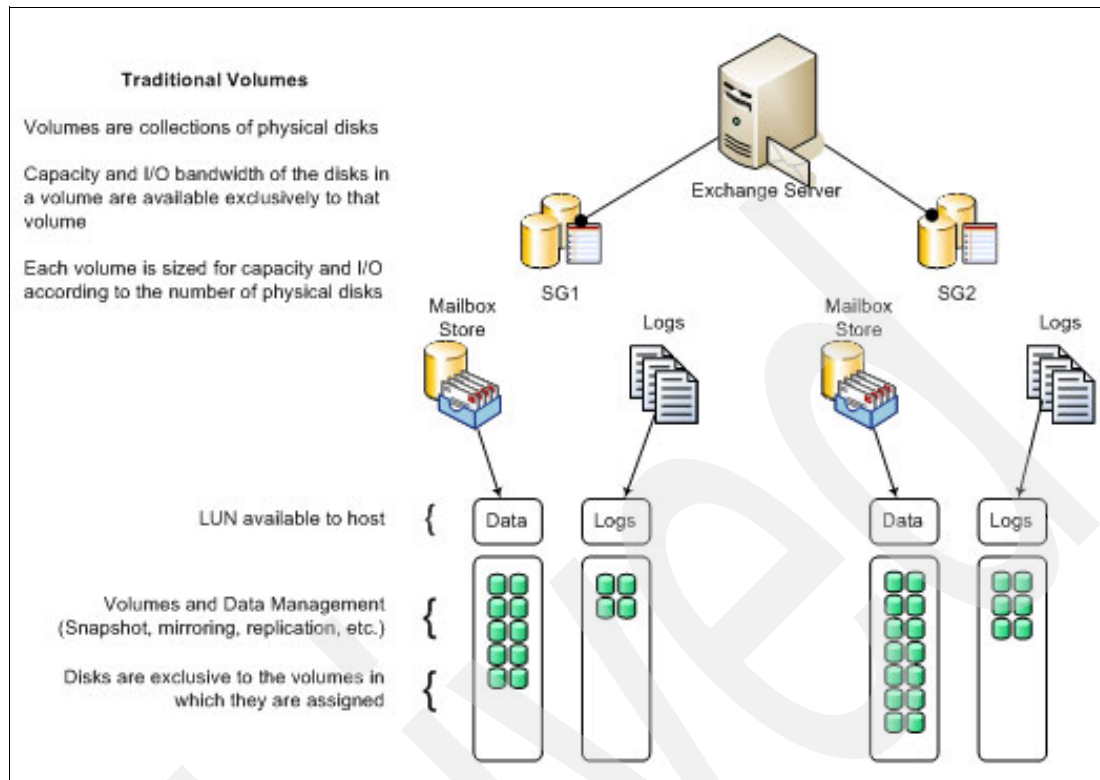


Figure 3-3 Disk utilization on traditional volumes

N series virtualization technology can address this issue by allowing the administrator to manage a set of physical storage devices as one large logical pool of storage. This type of technology can accelerate provisioning and increase capacity utilization.

N series offers both FlexVol and SnapDrive software functionality that lets administrators add new storage rapidly and nondisruptively to their Exchange environments. FlexVol technology allows the creation of one simplified, logical pool of storage, known as an *Aggregate*. After the logical aggregate is in place, you can perform provisioning storage to an Exchange server in minutes. The administrator just creates a “flexible volume” from the aggregate that is sized as required. You can also use a single command to increase the size of an existing flexible volume in real time, without incurring any Exchange application downtime or IT staff overtime. Because FlexVol allows storage to be provisioned dynamically, when and where it is needed, disk capacity utilization can increase to 70% or greater. On the server, a simple application called SnapDrive that is integrated with Windows MMC nondisruptively maps the new flexible volume to a drive letter on the Exchange server. No rebooting is required.

Figure 3-4 shows how disks and volumes are managed by N series.

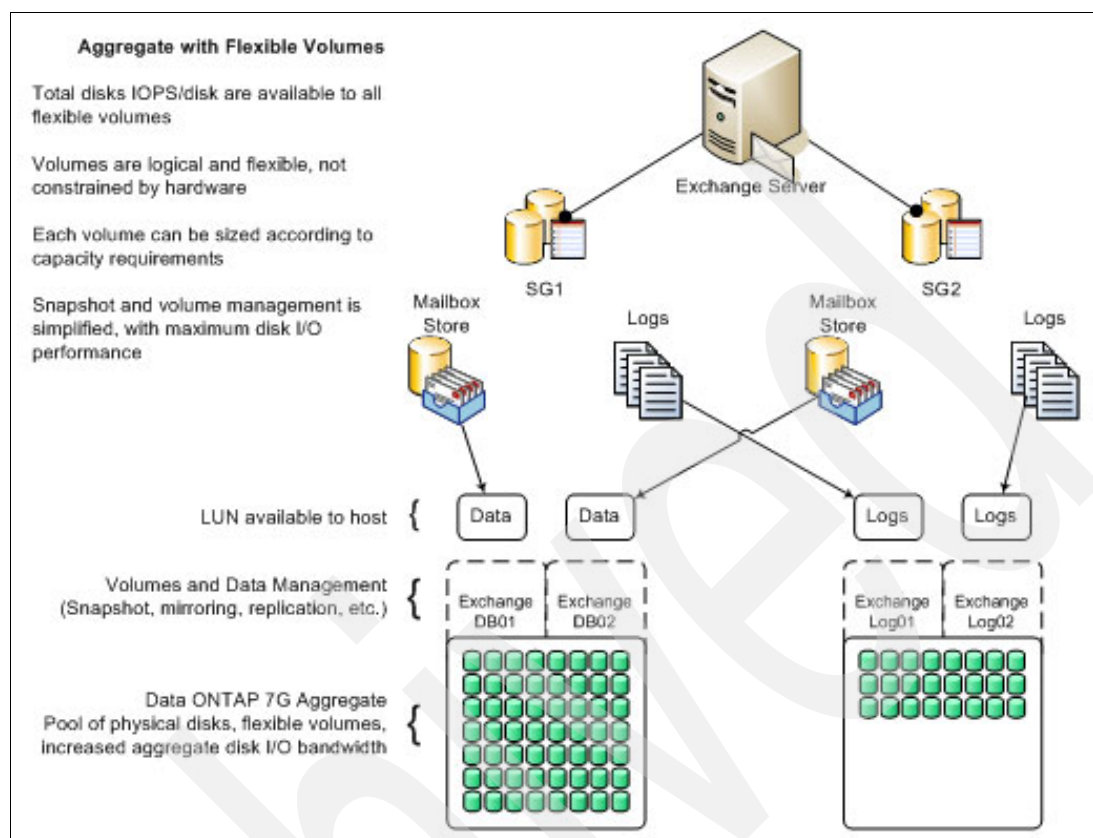


Figure 3-4 Disk utilization on N series with flexible volumes

3.2.4 Implement low-overhead remote mirroring for DR

Ensuring business continuity is one of the more challenging tasks faced by today's Exchange administrators. DAS architectures with a reliance on tape backup usually require additional functionality to ensure application availability in the event of some type of catastrophic failure at the main data center. Remote mirroring to an off-site DR location is the main method Exchange environments adopt to ensure application availability after a site disaster or some type of catastrophic failure. While these types of solutions are available, many are impractical to deploy for the average Exchange environment, due to their complexity and extensive bandwidth requirements. Many even require expensive and proprietary software, not to mention ongoing professional services to aid in the design, installation, and ongoing management of these types of systems.

Using SnapMirror technology, N series storage systems provide a flexible, affordable way to perform remote mirroring of Exchange data to another N series system over an existing IP-based LAN, MAN, or WAN architecture. Compared to other alternatives, this solution is quick to deploy and very cost effective, making it an excellent fit for Windows environments. SnapMirror is designed to utilize as little bandwidth as possible, since only changed blocks are transmitted over the network. With SnapMirror, Exchange data can be mirrored in one of three modes: synchronously, semi synchronously, or asynchronously, based on an automated schedule set by the administrator. Management and integration of SnapMirror simply requires selecting a single check box within SnapManager for Exchange.

3.2.5 Implement tiered storage for archival and compliance

Many administrators have begun to consider e-mail archival to lower-cost tiers of storage as a practical way to extend the use of e-mail and gain better utilization from their underlying storage. With the continued decrease in cost per megabyte of disk drives, disk-based e-mail archiving gives IT organizations and users rapid access to archived e-mail content.

Companies seeking to archive e-mails in compliance with new regulations can also benefit from disk-based archival solutions that preserve e-mail records from further edits or revisions, through write-once, read-many (WORM) technology. This approach combines the retention capabilities of WORM, with the ease of use and rapid access of disk storage.

For compliance-driven environments that need the retention capabilities of WORM, SnapLock can be added to existing N series systems through a simple software license key. With SnapLock, administrators just need to configure one additional attribute setting during the provisioning process that turns a flexible volume into a WORM volume. A SnapLock volume can hold snapshots, files sent to it, or, in the case of Exchange, can provide a natural repository for archival applications from third-party providers like VERITAS, Open Text-IXOS, and others. Based on open protocols, no extra API integration is required. This means upgrades to Data ONTAP software or third-party archival software will have no negative effect on standing archival procedures.

Figure 3-5 illustrates a simple, yet powerful, Microsoft Exchange architecture integration with a IBM System Storage N series storage system.

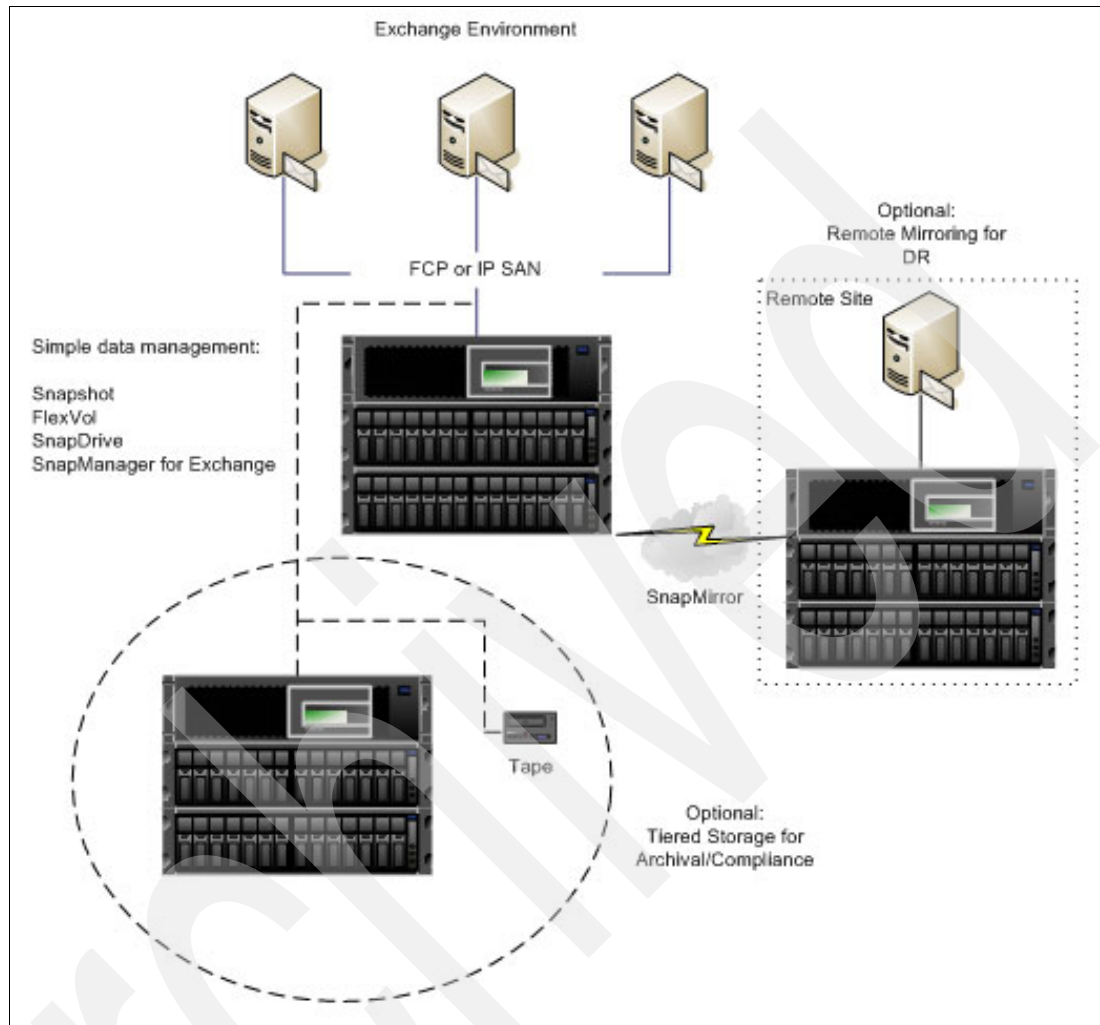


Figure 3-5 Exchange and N series solution integration



Part 2

Preparing the IBM System Storage N series for Lotus Domino and Microsoft Exchange

In this part, we discuss those hardware and software functions of the IBM System Storage N series that need to be installed, set up, or modified for integration of Lotus Domino Server and Microsoft Exchange server.

Archived

Preparing the IBM System Storage N series for IBM Lotus Domino 7

This chapter discusses the steps to prepare the IBM System Storage N series for IBM Lotus Domino 7. It also contains sizing and performance information. As Lotus Domino is multi-platform capable, we consider different operation systems and protocols, followed by our recommendations for performance and reliability when integrating the N series storage system.

In this chapter, we give preparation information for the Red Hat Enterprise Linux and IBM AIX operating systems with Lotus Domino and IBM System Storage N series.

4.1 Using Lotus Domino with an IBM System Storage N series storage system

There are several advantages to storing Lotus Domino databases and transaction logs on an IBM System Storage N series storage system. This section examines some of those advantages and explores steps that must be taken to prepare the installation and ensure overall Domino performance.

4.1.1 Advantages of IBM System Storage N series storage systems for Lotus Domino

Running Domino with databases and transaction log files stored on an IBM System Storage N series storage system has several advantages:

- **Extremely Fast Backup**

Snapshot copies (see Figure 4-1) can be created in a matter of seconds, regardless of the size of the Domino database or the level of activity on the IBM System Storage N series storage system used. This reduces the Domino backup window from hours to seconds and allows Domino administrators to take frequent full backups without having to take Domino server offline. Keep in mind that the effects of compaction and its high change rate will require a higher Snapshot reserve for preservation of snapshots.

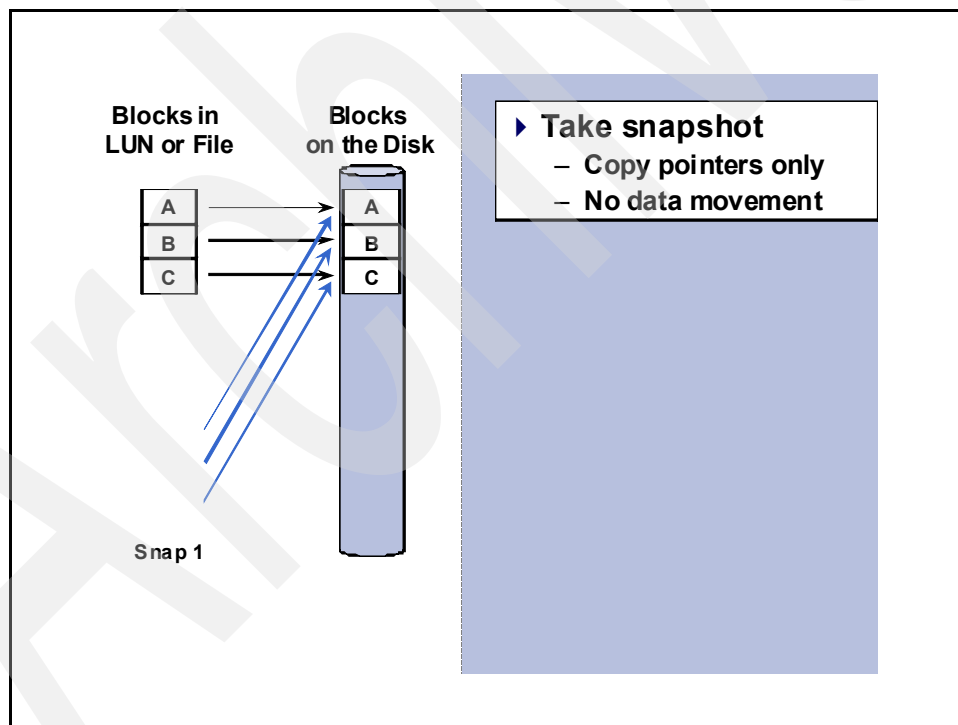


Figure 4-1 Snapshot

- **Quick recovery**

Using the Data ONTAP SnapRestore feature, an entire Domino environment can be restored in a matter of seconds. Since there is no data copying involved, an incredible amount of time is saved as the file system is put back to the original state it was in at the time the Snapshot was created. Data ONTAP supports 255 Snapshot copies per WAFL

volume. The ability to store a large number of low-impact, frequently created Snapshot copies brings the time needed to perform a roll-forward recovery operation down to minutes or seconds. In many circumstances, it allows the Domino administrator to restore the Domino server immediately without the necessity to restore from tape.

- ▶ **High availability**

The need for 24x7 availability is fast becoming a reality for organizations of all sizes. Organizations cannot tolerate scheduled downtime or afford extended periods of slow system response caused by traditional Domino server backup methods. Snapshot technology that can create Domino server backup in a matter of seconds without bringing the server down can be used as complementary technology to ensure higher system uptime.

- ▶ **High reliability**

The RAID architecture used for N series storage systems is unique and provides greater reliability than many traditional RAID implementations. If a disk in a N series RAID group fails, it is reconstructed automatically without any user intervention. Additionally, N series supports RAID-DP architecture. RAID-DP is considered approximately 4,000 times more reliable than traditional RAID.

- ▶ **No impact on system response time during backup**

A Snapshot copy is simply a picture of the file system at a specific point-in-time. Therefore, creating a Domino server backup using Snapshot does not involve actual data movement (data I/O), so the backup process has virtually no performance impact on system response time.

- ▶ **Minimum storage requirement**

Two Snapshot copies created in sequence differ from each other by the blocks added or changed in the time interval between their creation. This block-incremental behavior limits associated storage capacity consumption.

- ▶ **Load balance**

Many of the tasks associated with load balancing between multiple Domino directories can be eliminated. Because of the high performance of the N series storage system, only one volume needs to be defined for each directory used.

4.1.2 Lotus Domino storage requirements overview

To prepare IBM System Storage N series storage system, it is necessary to calculate and size the capacity and I/O requirements of Lotus Domino.

A Domino server requires storage space for several purposes. It could be divided in three parts:

- ▶ Storage space for the operating system, IBM Lotus Domino application files, and swap area
- ▶ Storage space for the IBM Lotus Domino data directories and users mail databases
- ▶ Storage space for the IBM Lotus Domino transactional log files

The directories for user data and transaction log files are independent of your operating system choice, while the space for operating system and swap area may vary. Figure 4-2 gives you an overview of our recommendations of the directory separation.

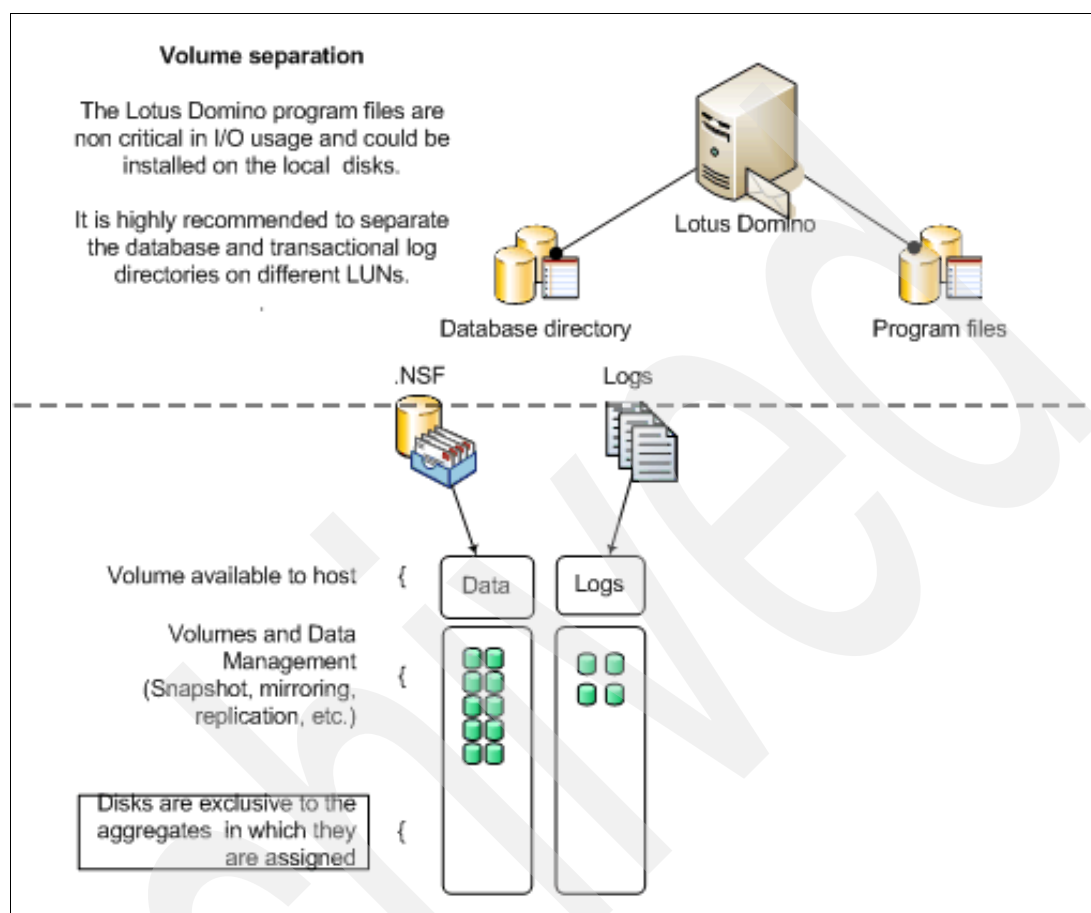


Figure 4-2 Recommended volume separation with Lotus Domino

The three storage volumes represent the required disk space mentioned in Figure 4-2 and are described as follows:

1. The program file volume holds the operating system and IBM Lotus Domino program files. It is recommended that these disks be protected with RAID1.
2. The database (NSF) volume holds the IBM Lotus Domino user database files. Accessing these files will create heavy random I/O operations. The size of the database volume depends on the mailbox size of your user and application database usage.

Note: IBM Lotus Domino forces the flushing of the file system buffer that maintains the operating system for strategic points. This flushing ensures that all write operations to the database are committed to the physical disk or non-volatile memory.

3. The transactional log volume holds the optional but recommended transactional log files. Every change to the Domino database files (NSF) is written to a transactional log file before writing them to the actual NSF. Separate and fast disks dedicated for transactional logging are the keys for increasing overall performance of your IBM Lotus Domino server.

Figure 4-3 shows our lab Linux and AIX file systems with separated program files, user databases, and transactional logging volumes. Remember that it is essential to place the transactional logs on a different volume and disks than the user database files.

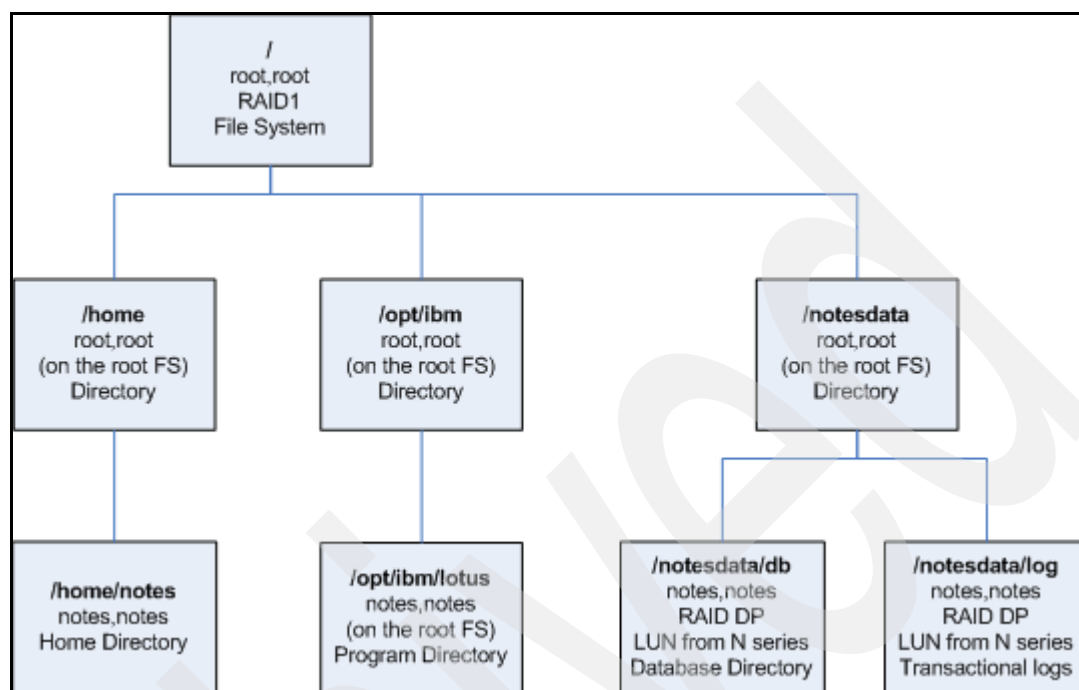


Figure 4-3 Linux file system structure used for this book

The capacity and I/O calculation in 4.1.3, “Capacity calculation of operating system, swap area, and program files” on page 75 and 4.1.7, “Hard disk performance and I/O tuning” on page 80 gives you an example of sizing a IBM System Storage N series storage system for IBM Lotus Domino.

4.1.3 Capacity calculation of operating system, swap area, and program files

The required disk space for various OSes is mentioned in Table 4-1.

Table 4-1 Disk space required for operating system and Lotus Domino program files

Operating system	Windows 2000 / 2003	Linux	UNIX
Required disk space for OS	Windows 2000: 2 GB Windows 2003: 2 GB	SLES 8, 9, 10: 4 GB RHEL 4: 5 GB	AIX 5.2, 5.3: 2.2 GB Solaris 9: 3 GB Solaris 10: 6.8 GB
Required disk space for Lotus Domino program files	1.5 GB	1.5 GB	1.5 GB
Required disk space for swap area	One-third the size of physical RAM installed, but at least 2 GB	One-third the size of physical RAM installed, but at least 2 GB	One-third the size of physical RAM installed, but at least 2 GB

There are several reasons why you might want to configure more than the minimum disk space. Considering the decreasing costs for disk capacity and the need for room for updates, patches, or release changes, it should be no challenge to allocate at least 15 or more gigabytes for this volume.

Tip: As a best practice, allocate at least 15 GB or more for your volume for the operating system, Lotus Domino program files, and swap area.

At the time this book was written, the Domino 7.0.2¹ release notes specified that the swap areas should be two or three times the physical memory. Based on our experience in different customer situations, we did not require that large of a swap area.

4.1.4 Capacity calculation for Lotus Domino databases

A Lotus Domino partition requires about one gigabyte of disk space for its database directory. It contains essential server databases such as names.nsf and mail.box. For organizations larger than a few thousand registered users, it is necessary to scale up the required disk space, for example, an additional 25 MB per thousand users.

Usually, most of the consumed space belongs to the user databases. Add up the total size of the users' mail databases to estimate the total space required. If you do not have access to this information, estimate the level of user activity.

- ▶ For each light user, estimate 100 MB.
- ▶ For each medium user, estimate 300 MB.
- ▶ For each heavy user, estimate 500 MB.

Note: This is just an approximate value and may differ in your environment.

Additionally, space for metadata should be regarded. Depending on the type of file system, no more than 10% should be required. For good performance, it is a best practice to prevent overloading the file system, so never let the database volume become more than 80% full. If you have more specific information for your file system, you might be able to refine this estimate, but unless you are sizing a very large system, the difference should not be to much.

Lotus Domino database LUN calculation

These steps summarize the steps to size your Domino database LUN:

1. Take one gigabyte per Domino partition plus the adjustments for a user count larger than a few thousand.
2. Add the estimated space for users databases.
3. Add the estimated space for application databases, if necessary.
4. Multiply by 1.1 to add the file system metadata.
5. Divide by 0.8 to leave 20% additional space for performance reasons

Example LUN calculation

- ▶ Number of light users: 20
- ▶ Number of medium users: 100

¹ For the Lotus Domino 7.0.2 release notes, see:
http://www-12.lotus.com/ldd/doc/domino_notes/7.0.2/rn702.nsf

- ▶ Number of heavy users: 30
 - ▶ Assumed size of additionally applications such as CRM or DMS: 25 GB
- $$\text{DB LUN} = ((1 \text{ GB} + (200 * 100 \text{ MB} + 150 * 300 \text{ MB} + 30 * 500 \text{ MB}) + 25 \text{ GB}) * 1.1) / 0.8$$
- $$\text{DB LUN} = 146 \text{ GB (We assumed 1 GB = 1000 MB for simplify this calculation.)}$$

Note: This calculation does not include database growth. Depending on your daily business, the database can double within a year. This should be considered when you size the database LUN and volume.

Lotus Domino database volume calculation

The volume calculation for Lotus Domino databases on N series depends on some additional factors:

- ▶ Quotas
- ▶ Mail sizes / rates
- ▶ Full-text indexing
- ▶ Clustering and replication
- ▶ Message tracking
- ▶ Third-party product databases

In our case, we did not consider these factors. For additional information about sizing, refer to *Sizing your IBM Lotus Domino mail servers* at:

<http://www.ibm.com/developerworks/lotus/library/domino-mail-sizing/>

We only added a extra 15% of space for growth.

Formula

Minimum database volume size = ((2 * (LUN size + 15% for growth)) + (number of Snapshots * data change percentage * LUN size))

Example calculation

- ▶ Using the database LUN size from above (+ 15% for growth): 146 GB + 15% = 168 GB
- ▶ 10% data change between backups
- ▶ Seven backups kept online

Minimum database volume size = (2 * 168 GB) + (7 * 10% * 146 GB) = 438 GB

Note: We doubled the LUN size to have an additional 100% space reservation on the volume.

4.1.5 Lotus Domino transactional log calculation

Domino supports transaction logging and recovery. With this feature enabled, the system captures database changes and writes them to the transaction log. Then, if a system or media failure occurs, using the transaction log and a third-party backup utility (such as IBM Tivoli® Storage Manager) are one method of recovering your database. A very fast backup and recovery solution provided by the N series Snapshot feature can protect transactional logging as well.

Tip: Enabling the transactional logging feature can improve server performance.

The transactional logging feature saves processing time because it allows Domino to defer database updates to disk during periods of high server activity. Because the transactions are already recorded, Domino can safely defer database updates until a period of low server activity. Additionally, the N series storage system is very efficient at writing sequential data to its volumes.

Note: For more information about transactional logging on Domino servers, see the *Lotus Domino Administrator 7 Help*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=swg27003543>

Three logging styles are available: *circular*, *linear*, and *archived*. The logging style you choose is also dependent on your disk size and backup strategy:

- ▶ Circular logging is the default when you activate the transactional logging feature. It continuously re-uses the log files (up to 4 GB) and overwrites old transactions. You are limited to restoring only the transactions stored in the transaction log.
- ▶ Linear logging is like circular, but it allows more than 4 GB of transactional logging. You need this style of logging if the space for the transactional logs between your backup intervals is greater than 4 GB and you do not want to archive the transactional logs with a third-party backup tool.
- ▶ The archive logging style is recommended. Domino does not re-use the log files until they are archived. You have to use a third-party backup tool for this style (IBM Tivoli Storage Manager, for example).

Figure 4-4 on page 79 shows the default circular logging configuration in the Lotus Domino Administrator. You will find it at Configuration Server Current Server Document Transactional Logging.

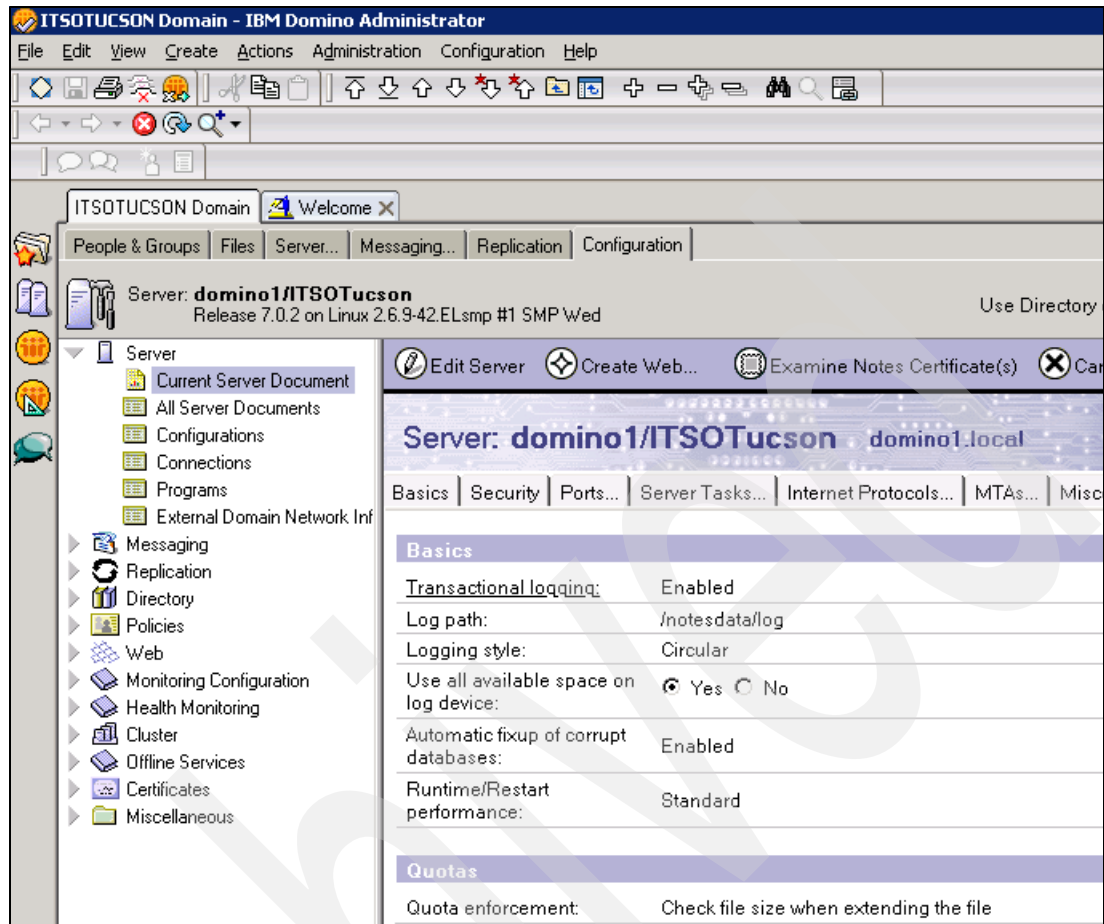


Figure 4-4 Transactional logging configuration in the Lotus Domino Server Document

Because the amount of log file space used with linear and archive logging is determined by the volume of database activity and your backup concept, the amount of storage needed for this type of logging cannot be predetermined.

Lotus Domino transactional log LUN calculation

The Lotus Domino transactional log LUN calculation relies on this formula:

Transaction log LUN size = (number of transaction logs generated * logsize) / 0.8 for file system performance

Example calculation

- Logging style: Circular
- Maximal log size: 4 GB

$$\text{Transaction log LUN size} = (1 * 4 \text{ GB}) / 0.8 = 5 \text{ GB}$$

Lotus Domino transactional log volume calculation

The Domino transaction log volume contains additional space for the Snapshots as shown in Example 4-4 on page 92. The formula and an example of using the formula are covered in the following paragraphs.

Formula

Transaction log volume size = (transaction log LUN size * 2) + (number of transaction logs generated * logsize * number of backups keep online)

Example calculation

- ▶ Using the transaction log LUN size from above: 5 GB (Because of circular style only one transaction log is generated.)
- ▶ Logsize: 4 GB
- ▶ Number of backups kept online: Seven

Transaction log volume size = (5 GB * 2) + (1 * 4 GB * 7) = 38 GB

4.1.6 Fractional space reservation

Fractional space reservation enables you to reserve less than 100% on volumes hosting LUNs, allowing for growing space utilization for those volumes. However, you can choose less than 100%, but you run a higher risk of unexpectedly running out of disk space on a volume, which in turn will halt any writes to a LUN.

If you want to use fractional space reservation, it affects two volume sizing formulas:

- ▶ Database volume sizing

Minimum database volume size = (((1 + fractional space reservation percentage) * (LUN size + 15% for growth)) + (number of Snapshots * data change percentage * LUN size))

- ▶ Transactional log volume sizing

Transaction log volume size = ((1 + fractional space reservation percentage) * transaction log LUN size) + (number of transaction logs generated * logsize * number of backups keep online)

We decided on a fractional space reservation of 65% for the database volume sizing.

Example volume size calculation with modified fractional space reservation

- ▶ Fractional space reservation: 65%
- ▶ Using the database LUN size from above (+ 15% for growth): 146 GB + 15% = 168 GB
- ▶ 10% data change between backups
- ▶ Seven backups kept online

Min. database volume size = ((1 + 0.65) * 168 GB) + (7 * 10% * 146 GB) = 379 GB

Note: Keeping the volumes at 100% space reservation is a best practice. It minimizes the risk of running out of space on the volume and shutting down your Lotus Domino Server. It is not recommended to use fractional reserve. Using fractional reserve is the same thing as reducing the size of a safety net.

4.1.7 Hard disk performance and I/O tuning

The performance of CPU, memory, and other core components is important. But in many ways the performance of a system is only equal to that of its poorest-performing component. Compared to the solid state components in a server, hard disks have, by far, the worst performance. Hard disks are improving more in size than speed, but so are CPUs, memory, and motherboards, and these other components tend to get faster more often, therefore

widening the gap. Therefore, hard disks are still the overall main constraint of the performance of servers.

The speed of the hard disks becomes very important for applications where hard disk operations are intensive. The Lotus Domino server is disk intensive in terms of reading and writing to the hard disk. For the database and transactional logging volumes, you should take the fastest disks available. Another benefit available with N series flexible volumes is that a small FlexVol will still share the I/O operations (IOPs) available through the large number of disks in an aggregate. The result of IOPs sharing with a large disk pool is that small flexible volumes can be created without incurring the performance penalty that occurs with smaller but busy traditional volumes. And since the parity disks are a part of the aggregate underlying RAID group and not the flexible volume itself, there is no capacity penalty lost to additional parity disks when creating a smaller Flexible volume.

The caching capabilities of N series is an outstanding performance attribute. Lotus Domino database transactions are written to the non-volatile random access memory (NVRAM) and cache of the N series. A write is not acknowledged until data is written in both places. Data is written from cache unless there is a controller failure that prevents it from doing so, at which point it is written from NVRAM. NVRAM and memory have a significant impact on Lotus Domino performance. You should contact your IBM representative to do proper sizing for your installation.

To separate transactional log files from database files, choose different aggregates on the N series, as shown in Figure 4-5.

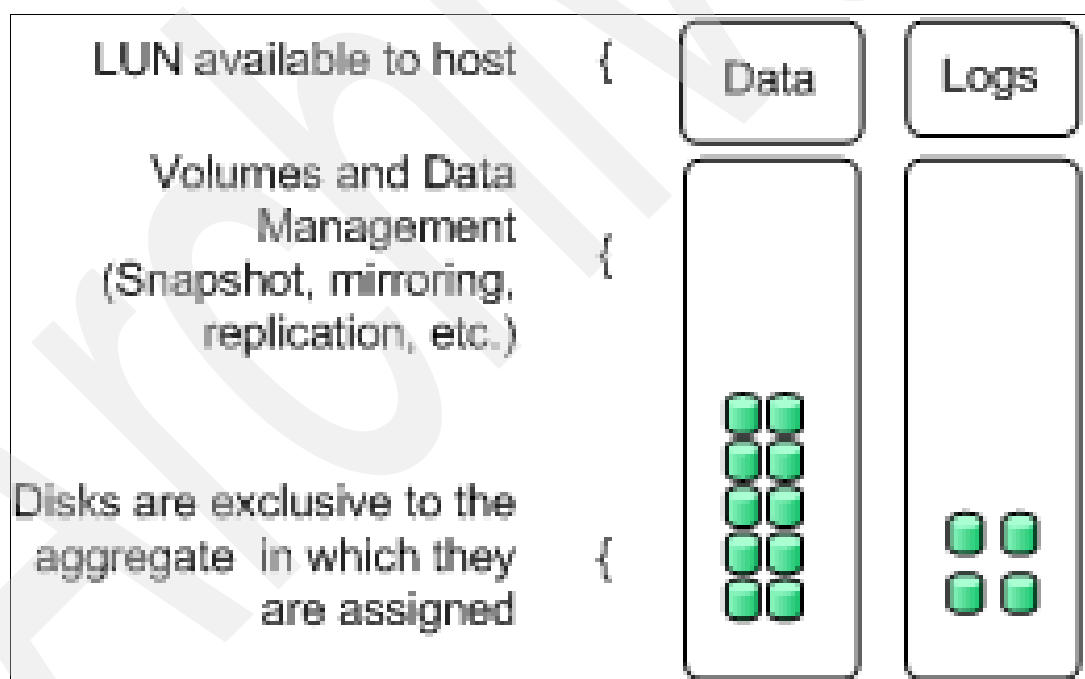


Figure 4-5 Separating database volume and log volume for best I/O performance

In many companies, it is common to have more than one Lotus Domino server at each site. For that scenario, increase the size of your aggregate with additional disks. As a positive side effect, this will also increase the overall performance of the existing LUNs located on the aggregate. Figure 4-6 shows an example with two Lotus Domino servers, which provides load balancing and high availability in a cluster configuration.

Instead of traditional volumes with dedicated disks, Domino server 1 and 2 gain the overall performance of the higher spindles count. This will improve the access time for Domino databases, particularly in peak times. Because the I/O profile of both database volumes are the same, they can be accumulated on the same aggregate. This applies for the transactional log volumes as well.

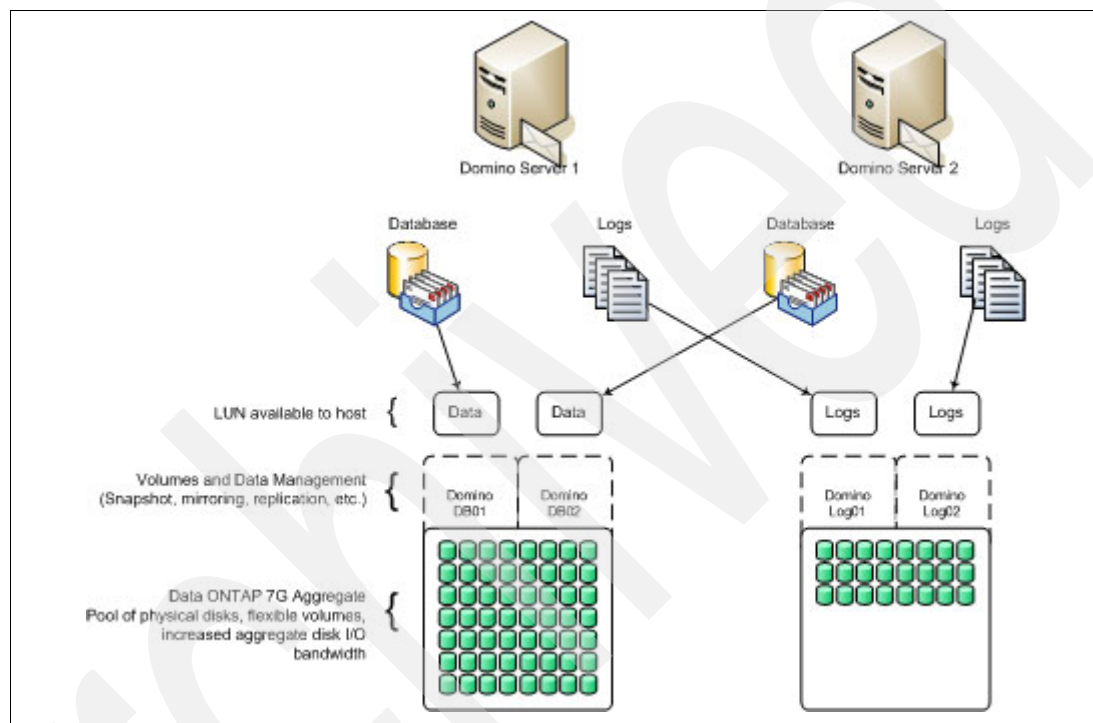


Figure 4-6 Best practice of database and transactional log volume segmentation with several Domino server

As shown in Figure 4-7 on page 83, we chose and recommend a Fibre Channel (FC) connection between the Lotus Domino Server and N series storage for performance reasons. At the time of publication, SnapDrive 2.2.1 does not support Host Multipathing on SUSE or Red Hat Linux, so we used a single path configuration in our lab to connect the Linux Domino server to N series. For performance reasons, we recommend a second HBA to separate the connection of database and transactional logging LUN.

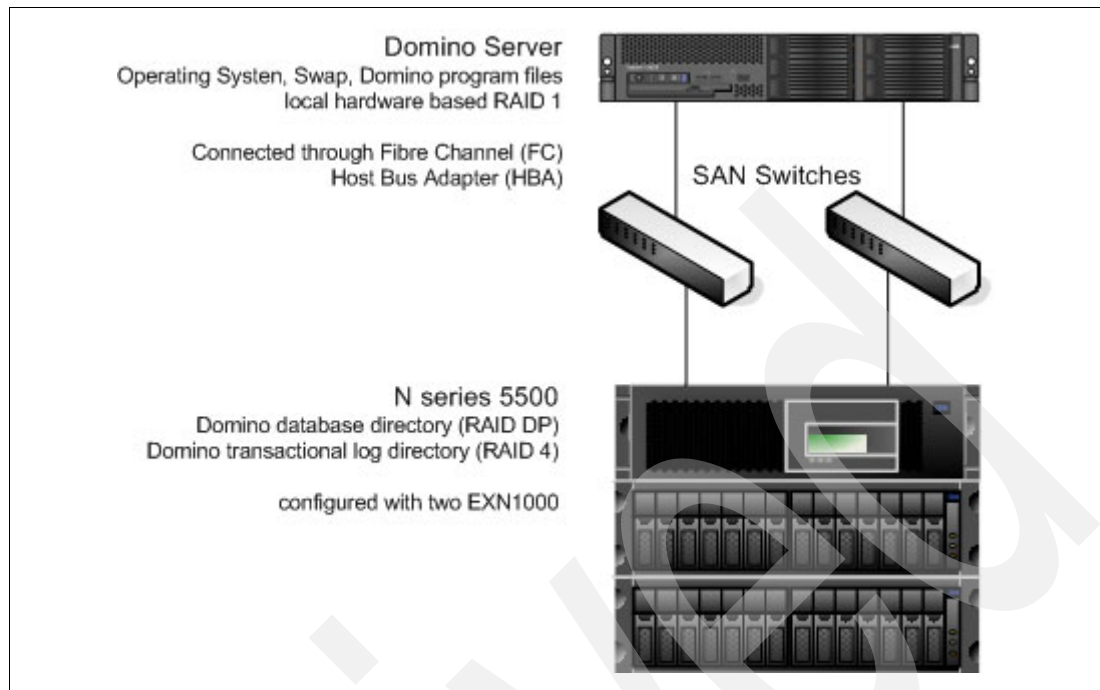


Figure 4-7 Lotus Domino server and N series lab infrastructure used in this book

Note: In a production environment, we recommend a redundant SAN path from the server to the clustered N series.

4.2 Storage configuration for Lotus Domino server

The IBM System Storage N series storage system must be configured prior to have the Lotus Domino server running on it. For this, the aggregates, volumes and LUNs must be created and configured to support the Lotus Domino server environment.

4.2.1 Aggregates

An aggregate is a collection of physical disks from which the space is allocated to the volumes. When creating the aggregates on the IBM System Storage N series storage system, there are some considerations:

- ▶ On each aggregate, one or more Flexible volumes can be created.
- ▶ Each aggregate has its own RAID configuration and set of assigned physical disks.
- ▶ The available space on the aggregate can be increased by simply adding disks to the existing RAID group or by adding new RAID groups.
- ▶ Performance is proportional to the number of disk spindles on the aggregate.

The aggregate consists of one or more plexes. A plex is a collection of one or more RAID groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system volumes. By default, the aggregates are created on plex0.

For detailed information about aggregates, the WAFL file system, and Data ONTAP V7.2, refer to *IBM System Storage N series Data ONTAP 7.2 Storage Management Guide*, found at: <http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001595&aid=1>

There are two types of aggregates: unmirrored aggregates and mirrored aggregates:

1. Unmirrored aggregates consist of a single plex. In Figure 4-8, an unmirrored aggregate named *aggr A* is being shown. This aggregate is made up of three RAID-DP groups named *rg0*, *rg1*, and *rg2* on Plex *plex0*.

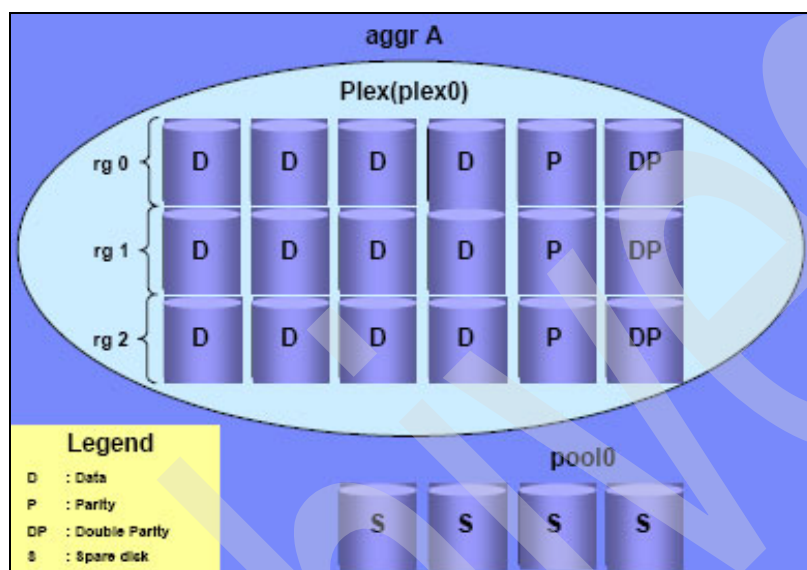


Figure 4-8 Unmirrored aggregate

2. Mirrored aggregates consist of two plexes that are mirrored to each other. A mirrored aggregate named *aggr A* is shown in Figure 4-9. This aggregate is made up of three RAID-DP groups named *rg0*, *rg1*, and *rg2* on both plexes *plex0* and *plex1*.




Figure 4-9 Mirrored aggregate

In the purposed scenario, two aggregates are required for the Lotus Domino server infrastructure: one aggregate for the databases and one aggregate for the transactional log files. This configuration can vary depending on many factors. For large deployments on IBM System Storage N series storage systems, the databases and the transactional log files should always be placed on different aggregates for performance purposes. For smaller deployments, performance may be better if Flexible volumes for databases and transactional logs are on a large single aggregate.

4.2.2 Creating the Domino database aggregate with Data ONTAP FilerView

When creating the aggregate, a name should be defined. There are some naming conventions for the aggregate's name. It should:

- ▶ Begin with either a letter or an underscore
- ▶ Contain only letters, digits, and underscores
- ▶ Contain no more than 255 characters

Disks to include in the aggregate must follow these rules:

- ▶ Disks must be of the same type, such as FC, SATA, or SCSI
- ▶ Disks must have the same RPM

To create a new aggregate through the Data ONTAP FilerView, follow these steps:

1. Type the IP address plus administration-path into your browser (an example is shown in Example 4-1):

`http://ip_address/na_admin`

Example 4-1 Data ONTAP Web front-end URL

`http://9.11.218.237/na_admin`

2. Click the FilerView Icon: A new browser window with the FilerView appears.
3. Click **Aggregates Add** to open the Aggregate Wizard (see Figure 4-10). A new browser window with the Aggregate Wizard appears.

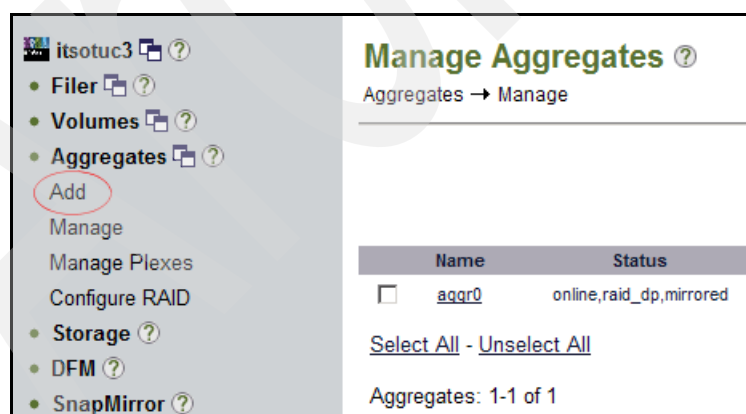
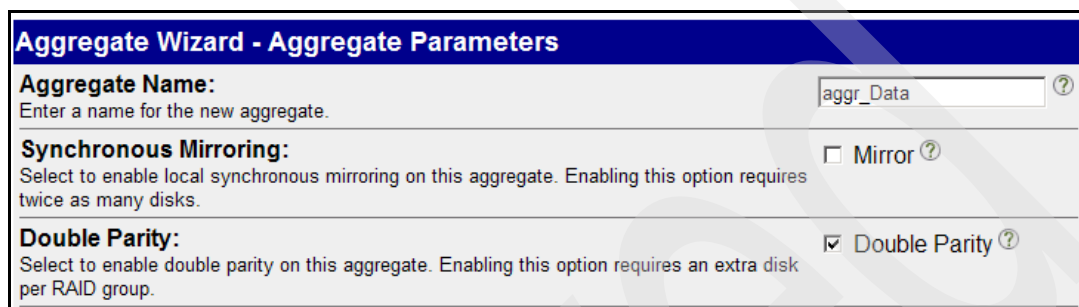


Figure 4-10 Adding a new aggregate

4. As seen in Figure 4-11, first the Aggregate Wizard asks for a name for the new aggregate. We suggest a prefix that identifies it as an aggregate for further usage, such as *aggr_*. In our case, we named the IBM Lotus Domino database aggregate *aggr_Data*.

Synchronous Mirroring enables the disk mirroring SyncMirror feature, which provides additional reliability. We do not use this in our case, but we selected **Double Parity** to gain the benefits of the N series RAID DP redundancy technology.

Click **Next**.



Aggregate Wizard - Aggregate Parameters

Aggregate Name:
Enter a name for the new aggregate. ?

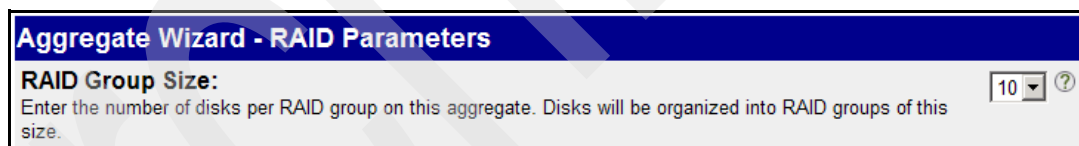
Synchronous Mirroring:
Select to enable local synchronous mirroring on this aggregate. Enabling this option requires twice as many disks. ☐ Mirror ?

Double Parity:
Select to enable double parity on this aggregate. Enabling this option requires an extra disk per RAID group. ☒ Double Parity ?

Figure 4-11 Aggregate Wizard - Aggregate Parameters

5. On the RAID Parameters window (Figure 4-12), select the number of disks that will be used on each RAID Group created for the aggregate. The recommended number of disks per RAID group with RAID-DP is 16 disks. If you are using less than 16 disks per RAID Group, protection against disk failure is increased, but performance will decrease because there will be less disk spindles for accessing the data. If you are using more than 16 disks per RAID Group, performance will be increased (more disk spindles to access the data), but protection against disk failure will decrease.

Click **Next**.



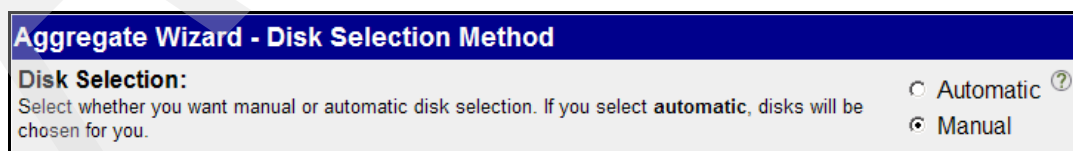
Aggregate Wizard - RAID Parameters

RAID Group Size:
Enter the number of disks per RAID group on this aggregate. Disks will be organized into RAID groups of this size. ?

Figure 4-12 Aggregate Wizard - RAID Parameters

6. The next step (Figure 4-13) gives us the option for either the automatic or manual disk selection method. Because we want to specify which disks are used, our choice is **Manual**.

Click **Next**.



Aggregate Wizard - Disk Selection Method

Disk Selection:
Select whether you want manual or automatic disk selection. If you select **automatic**, disks will be chosen for you. ☐ Automatic ? ☒ Manual

Figure 4-13 Aggregate Wizard - Disk Selection Method

7. Because we choose the manual selection method, we now select the disks added to the aggregate (see Figure 4-14 on page 87). Depending on your RAID group size selection (Figure 4-12), Data ONTAP will create one or more RAID groups based on your selection. Click **Next**.

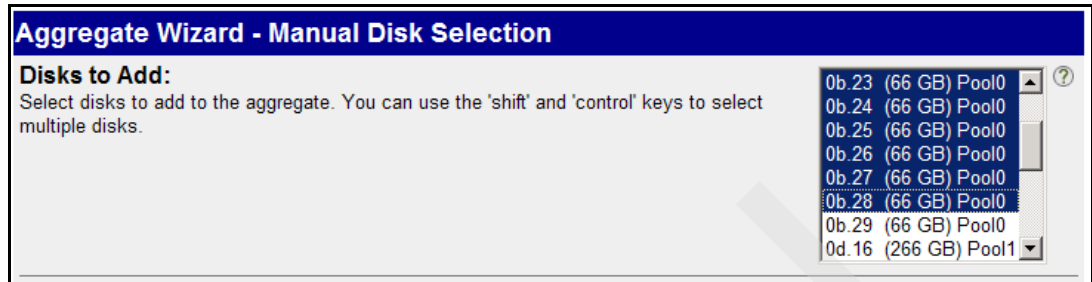


Figure 4-14 Aggregate Wizard - Manual Disk Selection

8. The configuration is done. Check the summary (see Figure 4-15) and click **Commit**.



Figure 4-15 Aggregate Wizard - Commit

9. Figure 4-16 shows the successfully end of aggregate creation. Click **Close**.

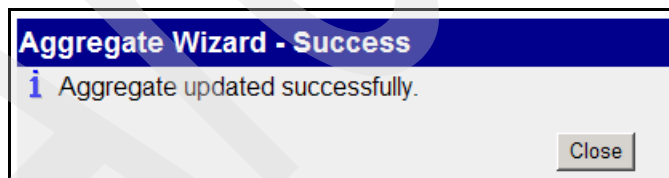


Figure 4-16 Aggregate Wizard - Success

Back in the Data ONTAP FilerView window, click **Aggregates** → **Manage**. The status of our aggregate (see Figure 4-17) is now:

creating raid_dp, initializing

Manage Aggregates ?

Aggregates → Manage

Filter by:

All Aggregates

View

Name	Status	Root	Avail	Used	Total	Disks	Files	Max Files	Checksums
<input type="checkbox"/> aggr0	online,raid_dp,mirrored	✓	14.2 GB	75%	56.8 GB	6	110	31.1 k	block
<input type="checkbox"/> aggr_Data	creating,raid_dp,initializing	-	-	-	-	0	-	-	block

Select All - Unselect All

Online

Restrict

Offline

Destroy

Figure 4-17 Initializing the aggregate

The initializing means that the disks are now initialized by *zeroing*. To track the status of the creation process, select **Storage** → **Disks** → **Manage**. Figure 4-18 shows an example of the zeroing process on disks 0b.19, 0b.20, 0b.21, and 0b.22.

Manage Disks ?

Storage → Disks → Manage

View Type: All Disks

View

Disk	Type	Checksum	Shelf	Bay	Chan	Size
<input type="checkbox"/> 0b.16	data	zoned/block	1	0	FC:A	66 GB
<input type="checkbox"/> 0b.17	parity	zoned/block	1	1	FC:A	66 GB
<input type="checkbox"/> 0b.18	dparity	zoned/block	1	2	FC:A	66 GB
<input type="checkbox"/> 0b.19	pending zeroing (4% done)	zoned/block	1	3	FC:A	66 GB
<input type="checkbox"/> 0b.20	pending zeroing (4% done)	zoned/block	1	4	FC:A	66 GB
<input type="checkbox"/> 0b.21	pending zeroing (4% done)	zoned/block	1	5	FC:A	66 GB
<input type="checkbox"/> 0b.22	pending zeroing (4% done)	zoned/block	1	6	FC:A	66 GB

Figure 4-18 Zeroing disks

Depending on the disk size, it will take about an hour or more to initialize the aggregate. After zeroing the disks, they belong to the newly created aggregate, in our case *aggr_Data* (see Figure 4-19 on page 89). Click **Refresh** to refresh the status.

Manage Disks ?

Storage → Disks → Manage

View Type: All Disks View

	Disk	Type	Checksum	Shelf	Bay	Chan	Size	Physical	Pool	Aggregate
<input type="checkbox"/>	0b.16	data	zoned/block	1	0	FC:A	66 GB	68 GB	Pool0	aggr0
<input type="checkbox"/>	0b.17	parity	zoned/block	1	1	FC:A	66 GB	68 GB	Pool0	aggr0
<input type="checkbox"/>	0b.18	dparity	zoned/block	1	2	FC:A	66 GB	68 GB	Pool0	aggr0
<input type="checkbox"/>	0b.19	parity	zoned/block	1	3	FC:A	66 GB	68 GB	Pool0	aggr_Data
<input type="checkbox"/>	0b.20	dparity	zoned/block	1	4	FC:A	66 GB	68 GB	Pool0	aggr_Data
<input type="checkbox"/>	0b.21	data	zoned/block	1	5	FC:A	66 GB	68 GB	Pool0	aggr_Data
<input type="checkbox"/>	0b.22	data	zoned/block	1	6	FC:A	66 GB	68 GB	Pool0	aggr_Data
<input type="checkbox"/>	0b.23	data	zoned/block	1	7	FC:A	66 GB	68 GB	Pool0	aggr_Data
<input type="checkbox"/>	0b.24	data	zoned/block	1	8	FC:A	66 GB	68 GB	Pool0	aggr_Data

Figure 4-19 Zeroing disks finished

Select **Aggregates** → **Manage**. The aggregate *aggr_Data* is now ready to use and has 454 GB of free space available (see Figure 4-20).

Manage Aggregates ?

Aggregates → Manage

Filter by: All Aggregates View

	Name	Status	Root	Avail	Used	Total	Disks	Files	Max Files	Checksums
<input type="checkbox"/>	ExchangeDB	creating,raid_dp,initializing	-	-	-	-	0	-	-	block
<input type="checkbox"/>	ExchangeLGs	creating,raid_dp,initializing	-	-	-	-	0	-	-	block
<input type="checkbox"/>	aggr0	online,raid_dp,mirrored	✓	14.2 GB	75%	56.8 GB	6	110	31.1 k	block
<input type="checkbox"/>	aggr_Data	online,raid_dp		454 GB	0%	454 GB	10	98	31.1 k	block

[Select All](#) - [Unselect All](#) Online Restrict Offline Destroy

Aggregates: 1-4 of 4

Figure 4-20 Aggregate created

Even the use of Data ONTAP FilerView is an easy way to create a new aggregate; the command line interface provides similar features.

4.2.3 Creating the Domino log aggregate with Data ONTAP CLI

The following steps describe the creation of a new aggregate with the Data ONTAP command-line interface (CLI):

1. Log in to your Filer with SSH (you have to start it first) or Telnet.

2. View a list of the spare disks that are left on the storage system. These could be used to create a new aggregate. The column *Device* provides the required device information for the creation process. Use the following command to view the status:

```
aggr status -s
```

3. Aggregate creation example: The command described in Example 4-2 creates an aggregate with no more than two disks in a RAID 4 group consisting of the disks with device IDs 0d.24 and 0d.25.

Example 4-2 Aggregate creation with command line interface

```
itsotuc3> aggr create aggr_Log -r 2 -t raid4 -d 0d.24 0d.25
```

Creation of an aggregate with 2 disks has been initiated. The disks need to be zeroed before addition to the aggregate. The process has been initiated and you will be notified via the system log as disks are added.

```
itsotuc3>
```

- Option *-r*: *raidsize* is the number of disks that you want in the RAID group created for this aggregate.
- Option *-t*: *RAID type*, *raid4* or *raid_dp* is allowed.
- Option *-d*: *device ID* of one or more available disks. Use a space to separate multiple disks.

4. Verify the aggregate exists as you specified:

```
aggr status aggr_Log -r
```

Example 4-3 shows the status of our created aggregate.

Example 4-3 Aggregate status

```
itsotuc3*> aggr status aggr_Log -r
Aggregate aggr_Log1 (online, raid4) (block checksums)
Plex /aggr_Log/plex0 (online, normal, active, pool0)
RAID group /aggr_Log/plex0/rg0 (normal)
```

RAID	Disk	Device	HA	SHELF	BAY	CHAN	Pool	Type	RPM
parity	0d.24	0d	2	4	FC:A	0	FCAL	10000	
data	0d.25	0d	2	5	FC:A	0	FCAL	10000	

```
itsotuc3*>
```

Note: For more information about the aggregate creation command, run the following command:

```
itsotuc3> aggr help create
```

You should get the following output:

```
aggr create <aggr-name>
      [-f] [-l <language-code>] [-L [compliance | enterprise]]
      [-m] [-n] [-r <raid-group-size>]
      [-R <rpm>] [-T {ATA | EATA | FCAL | LUN | SCSI}]
      [-t {raid4 | raid_dp}] [-v] <disk-list>
- create a new aggregate using the disks in <disk-list>;
  <disk-list> is either
      <ndisks>[@<disk-size>]
or
      -d <disk-name1> <disk-name2> ... <disk-nameN>
      [-d <disk-name1> <disk-name2> ... <disk-nameN>].
If a mirrored aggregate is desired, make sure to specify
even number for <ndisks>, or to use two '-d' lists.
```

4.2.4 Volumes

Volumes on the IBM System Storage N series storage system can be designated as Traditional volumes or flexible volumes.

Traditional volumes are tied to the physical disks on the aggregate on which they are created. This means that the disks used on a Traditional volume cannot be used on a different volume, whether it is a Traditional volume or a flexible volume.

Traditional volumes do not allow much flexibility and the only way to increase the size of a Traditional volume is to add disk spindles to the volume array. This type of volume does not allow downsizing.

On the other hand, flexible volumes created on aggregates can use disks from and share disks with different flexible volumes. This is due to the fact that flexible volumes are not tied to the physical disks on which they are created but to the aggregates collection of disks.

Flexible volumes provide more management flexibility and allow for dynamic volume size expansion and shrinkage without impact on the host client.

There are some advantages to using flexible volumes with the Lotus Domino server:

- ▶ All volumes can be managed independently, while taking advantage of the maximum I/O performance benefit of a much larger pool of physical disks.
- ▶ Volume size can be dynamically increased and decreased without adding extra physical disks.
- ▶ A larger number of volumes can be created with independent SnapShot management, schedules, mirroring policies, and so on.
- ▶ There is less waste of disk space because the flexible volumes relies on the aggregates physical disks, instead of having a separate array of disks.

The recommendation for a Lotus Domino server environment is to always use the Flexible volumes instead of the Traditional volumes because the flexible volumes provide a better performance and manageability at the same time that they provide a better physical disk

resources utilization. This recommendation is independent of the operating system or used protocol of the Lotus Domino server implementation.

4.2.5 Creating the Domino database volume with Data ONTAP FilerView

Every volume on the IBM System Storage N series storage system must be created on an aggregate. Because of the recommendation for Lotus Domino servers, the two volumes should be created on different aggregates.

The volume name must follow these naming conventions:

- ▶ Begin with either a letter or an underscore
- ▶ Contain only letters, digits, and underscores
- ▶ Contain no more than 255 characters

To create a new flexible volume through the Data ONTAP FilerView, follow these steps:


1. Type the IP address plus administration-path into your browser (an example is given in Example 4-4):

`http://ip_address/na_admin`

Example 4-4 Data ONTAP Web front-end URL

`http://9.11.218.237/na_admin`

2. Click the FilerView Icon. A new browser window with the FilerView appears.
3. On the left menu bar (see Figure 4-21), click **Add** to start the volume Wizard.



Click **Add** for creating a new volume

Figure 4-21 Adding a new volume

4. A new browser window with the volume Wizard appears (see Figure 4-22). Click **Next** to continue.

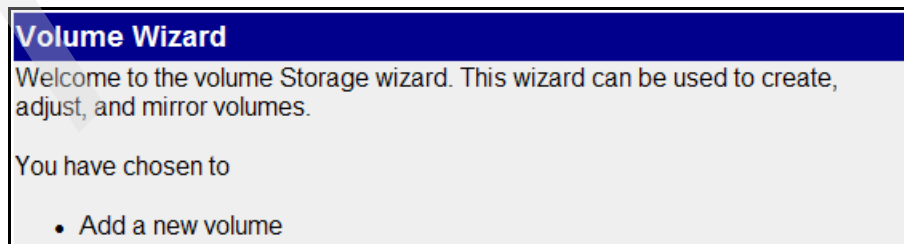


Figure 4-22 Volume Wizard welcome page

5. The next step lets you choose the volume type. To take advantage of the benefits of the FlexVol feature, we selected **Flexible** (see Figure 4-23). For Lotus Domino database and transactional logging files, this is recommended. Select your volume type and click **Next**.

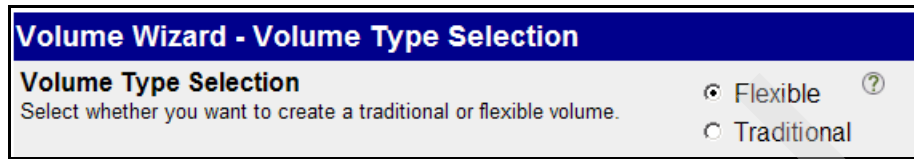


Figure 4-23 Volume Wizard - Volume Type Selection

6. The volume name matches our lab naming convention and starts with *vol_*. We suggest using a prefix that identifies it as an volume for further usage. In our case (see Figure 4-24), we named the IBM Lotus Domino database volume *vol_DominoDB*.
- The language option should be the same as set for the root volume. This is selected as the default.

Note: We strongly recommend that all volumes have the same language as the root volume, and that you set the volume language at volume creation time. Changing the language of an existing volume can cause some files to become inaccessible.

Set the volume name and choose your language option. Then click **Next** to continue.

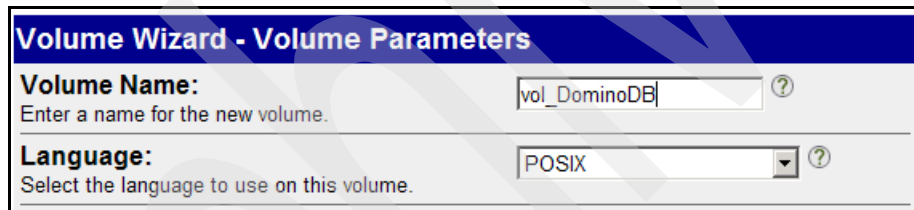


Figure 4-24 Volume Wizard - Volume Parameters

7. The volume parameters are set in the next step (see Figure 4-25 on page 94). With FlexVols, the total volume size can be increased and decreased after creation. We chose the calculated Lotus domino database volume size of 438 GB.

Each flexible volume has a space guarantee attribute that controls how its storage is managed in relation to its containing aggregate. There are three possible settings for this attribute:

- volume

Data ONTAP pre-allocates space in the aggregate for the volume. The pre-allocated space cannot be allocated to any other volume in that aggregate.

Space management for a flexible volume with a space guarantee of volume is equivalent to a traditional volume.

- file

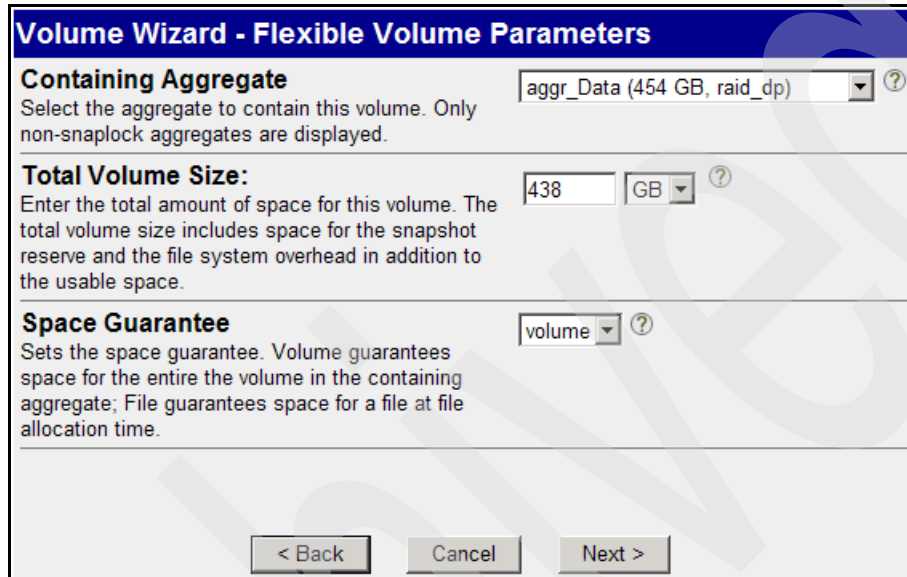
Data ONTAP pre-allocates space in the volume so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are pinned for a snapshot.

- none

Data ONTAP reserves no extra space for the volume. Writes to LUNs or files contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

Note: Use the space guarantee option *none* only in testing environments.

Chose your aggregate, volume size, and space guarantee, and then click **Next**.



Volume Wizard - Flexible Volume Parameters

Containing Aggregate
Select the aggregate to contain this volume. Only non-snaplock aggregates are displayed.
aggr_Data (454 GB, raid_dp) ?

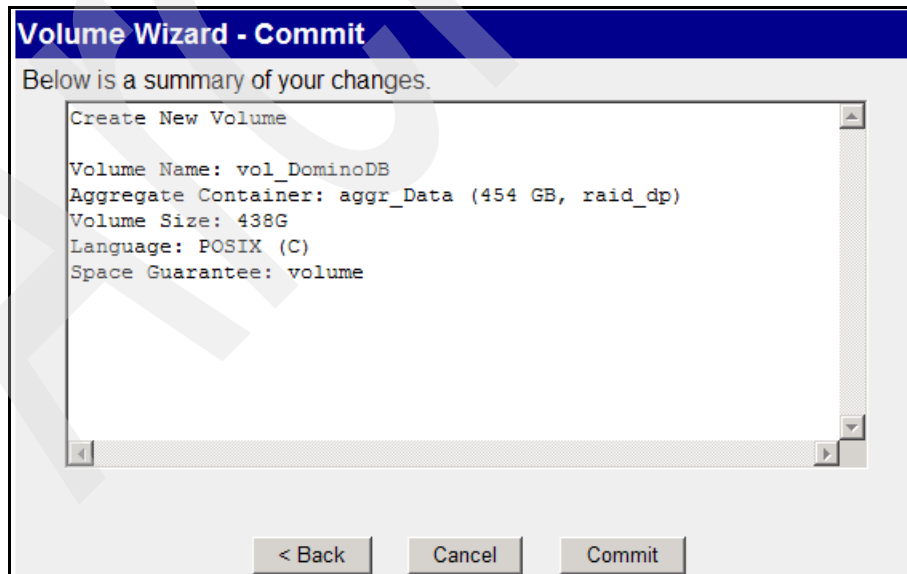
Total Volume Size:
Enter the total amount of space for this volume. The total volume size includes space for the snapshot reserve and the file system overhead in addition to the usable space.
438 GB ?

Space Guarantee
Sets the space guarantee. Volume guarantees space for the entire the volume in the containing aggregate; File guarantees space for a file at file allocation time.
volume ?

< Back Cancel Next >

Figure 4-25 Volume Wizard - Flexible Volume Parameters

8. The configuration is done. Check the summary (see Figure 4-26) and click **Commit**.



Volume Wizard - Commit

Below is a summary of your changes.

```

Create New Volume
Volume Name: vol_DominoDB
Aggregate Container: aggr_Data (454 GB, raid_dp)
Volume Size: 438G
Language: POSIX (C)
Space Guarantee: volume
  
```

< Back Cancel Commit

Figure 4-26 Volume Wizard - Commit

Figure 4-27 on page 95 shows the successful creation of the volume. Click **Close**.




Figure 4-27 Volume Wizard - Success

Back in the Data ONTAP FilerView window, select **Volumes** → **Manage**. Figure 4-28 shows the list of all available volumes on the N series. It gives a short overview of the important information, such as online-status, RAID configuration, size, and containing aggregate. Using the buttons Online, Restricted, Offline, and Destroy, you are able to modify the volume status in special cases, such as creating a SnapMirror target or deleting (destroying) a volume. The last row represent our instantly created *vol_DominoDB* volume with its status:

online, raid_dp

Manage Volumes ?
Volumes → Manage

Filter by:

	Name	Status	Root	Containing Aggregate	Clone	Avail	Used	Total	Files	Max Files
<input type="checkbox"/>	Vol Exch DB	online,raid_dp		Aggr Exch DB	-	155 GB	46%	288 GB	109	12.5 m
<input type="checkbox"/>	Vol Exch Log	online,raid_dp		Aggr Exch Log	-	2.71 GB	94%	44.8 GB	109	2.21 m
<input type="checkbox"/>	redhat	online,raid_dp,mirrored		aggr0	-	880 MB	95%	17.6 GB	39.6 k	761 k
<input type="checkbox"/>	vol0	online,raid_dp,mirrored	✓	aggr0	-	12 GB	25%	16 GB	5.61 k	692 k
<input type="checkbox"/>	vol1	online,raid_dp,mirrored		aggr0	-	294 MB	27%	400 MB	104	31.1 k
<input type="checkbox"/>	vol_DominoDB	online,raid_dp		aggr_Data	-	350 GB	0%	350 GB	101	15.1 m

[Select All](#) - [Unselect All](#)

Volumes: 1-6 of 6

Figure 4-28 Volume creation finished

Even the use of Data ONTAP FilerView is an easy way to create a new flexible volume; the command-line interface provides similar features.

4.2.6 Creating the Domino transaction log volume with Data ONTAP CLI

The **vol create** command will be used to create a new volume on the aggregate. See Example 4-5.

Example 4-5 Help output of the vol create command

```
itsotuc3*> vol create
vol create: No volume name supplied.
usage:
vol create <vol-name>
      { [-l <language-code>] [-s {none | file | volume}]
        <hosting-aggr-name> <size>[k|m|g|t]
        [-S remotehost:remotevolume] }
|
      { [-f] [-l <language-code>] [-m] [-n]
        [-L [compliance | enterprise]]
        [-r <raid-group-size>] [-t {raid4 | raid_dp}]
        <disk-list> }
- create a new volume, either a flexible volume from an existing
  aggregate, or a traditional volume from a disk list. A disk
  list is either
      <ndisks>[@<disk-size>]
or
      -d <disk-name1> <disk-name2> ... <disk-nameN>
      [-d <disk-name1> <disk-name2> ... <disk-nameN>].
itsotuc3*>
```

The following example describes the creation of a new volume with the Data ONTAP command-line interface (CLI):

1. Log in to your Filer with SSH (you have to start it first) or Telnet.
2. Verify the new volume does not already exist on the N series (see Example 4-6).

Example 4-6 List all volumes

```
itsotuc3*> vol status
Volume State      Status      Options
vol0 online      raid_dp, flex  root, maxdirsize=20971
vol1 online      raid_dp, flex  mirrored
redhat online     raid_dp, flex  mirrored
Vol_Exch_DB online raid_dp, flex  create_ucose=on,
convert_ucose=on
Vol_Exch_Log online raid_dp, flex  create_ucose=on,
convert_ucose=on
vol_DominoDB online raid_dp, flex
itsotuc3*>
```

3. Verify that the aggregate, which should hold the new volume, is online and if it has enough free space (see Example 4-7 on page 97). In our case, we create the volume on *aggr_Log*.

Example 4-7 Status of all aggregates

```
itsotuc3*> aggr status
      Aggr State      Status      Options
aggr_Log online    raid4, aggr    raidsize=2
      aggr_Data online  raid_dp, aggr    raidsize=10
      Aggr_Exch_DB online  raid_dp, aggr    raidsize=8
      Aggr_Exch_Log online  raid_dp, aggr    raidsize=8
      aggr0 online      raid_dp, aggr    root
                        mirrored
```

```
itsotuc3*> aggr status aggr_Log -r
Aggregate aggr_Log (online, raid4) (block checksums)
Plex /aggr_Log/plex0 (online, normal, active, pool1)
RAID group /aggr_Log/plex0/rg0 (normal)
```

RAID	Disk	Device	HA	SHELF	BAY	CHAN	Pool	Type	RPM	Used (MB/blks)
parity	0d.24	0d	1	8		FC:A	1	FCAL	10000	272000/557056000
data	0d.25	0d	1	9		FC:A	1	FCAL	10000	272000/557056000

```
tsotuc3*>
```

4. Create the volume *vol_DominoLog* on aggregate *aggr_Log* (see Example 4-8). The language of the volume will be the language set on the root volume, in our case *POSIX* declared as *C* on the command-line output.

Example 4-8 Creating the volume

```
itsotuc3*> vol create vol_DominoLog -s volume aggr_Log 38g
Tue Jun 5 14:54:39 PDT: Language on volume vol_DominoLog changed to C

The new language mappings will be available after reboot
Tue Jun 5 14:54:39 PDT: XL-Language of volume vol_DominoLog has been changed to C.
Creation of volume 'vol_DominoLog' with size 38g on containing aggregate
'aggr_Log' has completed.
itsotuc3*>
```

5. After the volume was created successfully, it will appear in the list generated by the **vol status** command.

4.2.7 Creating the LUNs

To use the FCP or iSCSI connection between the Lotus Domino server and IBM N series, it is necessary to create a *logical unit number (LUN)* on the N series. Usually this is done by the N series SnapDrive for Windows, Linux, or UNIX feature. In this case, the SnapDrive LUN creation process creates the LUN, *Initiator Group (iGroup)*, and file system, and associates the LUN with a mounting point (Linux and UNIX) or disk (Windows).

Note: We recommend the use of the N series SnapDrive feature for LUN creation.

For manual LUN creation through the FilerView, select **LUNs** → **Add** (Figure 4-29). The LUN creation page appears in the left frame. The Path to your LUN depends on the volume name and LUN name. As we chose *lun_DominoDB* for the LUN name, the complete path in our case is:

`/vol/vol_DominoDB/lun_DominoDB`

Set the LUN Protocol Type (also known as *Host Type*), size of the volume, and if the space should be reserved. After that, click **Add**.

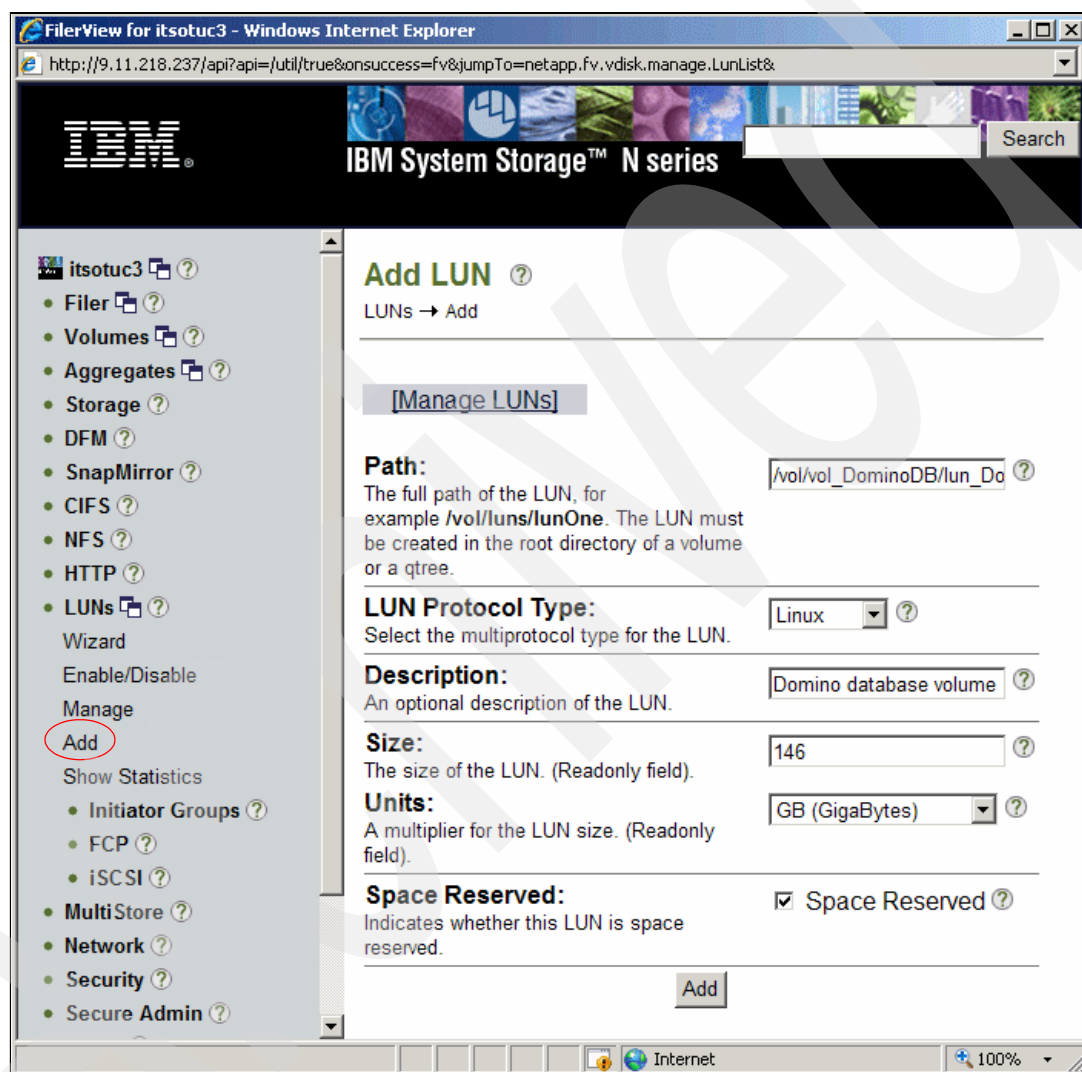


Figure 4-29 Creating the LUN

Note: The (Readonly field) advice at the Size and Units option applies only if you are changing the LUN. Although you can resize a LUN through the command-line interface, we highly recommend managing the LUN with SnapDrive

After the LUN is created, it appears in the LUN list. Select **LUNs** → **Manage**. The LUN overview, shown in Figure 4-30 on page 99, gives you important information about your LUN status, size, and mapping. The mapping follows the rule Groups: LUN ID. Before the target server can mount the N series LUN, it is necessary to create a Initiator Group and map the

LUN. As we use the N series SnapDrive for Linux and UNIX feature, we highly recommend letting SnapDrive automatically create and map the LUNs.

Manage LUNs ?
LUNs → Manage

Add New LUN
Hide Maps

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/Vol Exch DB/lun_exch_db		133.015 GB	online	viaRPC.21:00:00:e0:8b:12:0e:54.NAMGMT3 : 1 viaRPC.21:00:00:e0:8b:12:b9:56.NAMGMT3 : 1
/vol/Vol Exch Log/lun_exch_logs		42.002 GB	online	viaRPC.21:00:00:e0:8b:12:0e:54.NAMGMT3 : 0 viaRPC.21:00:00:e0:8b:12:b9:56.NAMGMT3 : 0
/vol/Vol DominoDB/lun_DominoDB	Domino database volume	146.019 GB	online	No Maps

Refresh

Figure 4-30 LUN creation finished

4.3 Zoning

A Fibre Channel zone consists of a group of Fibre Channel ports or nodes that can communicate with each other. It can be thought of as a logical fabric subset. Two Fibre Channel nodes can communicate with one another only when they are contained in the same zone. A node can be contained in multiple zones. For example, it is typical for a storage node (that is, N series target FC port) to be contained in multiple zones. There are generally considered to be two methods of zoning: hard and soft zoning.

4.3.1 Hard zoning

Port based, often referred to as *port zoning*, is where the zone is defined by specifying the fabric unique N_port IDs of the ports to be included. In other words, the switch and switch port (that is, switch3/Port 4) are used to define the zone members.

Security

Hard zoning typically is considered to offer improved security, since it is not possible to breach the zoning using WWN spoofing. However, if someone has physical access to the switch, simply replacing a cable can allow access when hard zoning is used.

Manageability

In many environments, hard zoning is easier to create and manage since only the switch or switch domain and port number need to be worked with instead of the 16-digit WWNs.

4.3.2 Soft zoning

In World Wide Name (WWN) based zoning, the zone is defined by specifying the WWN of the members to be included within the zone. Depending on the switch vendor, either World Wide Node Names or World Wide Port Names can be used, although we recommend World Wide Port Name zoning.

Flexibility

Since access is not determined by where the device is physically plugged into the fabric, it is possible with soft zoning to simply move a cable from one port to another without needing to reconfigure the zoning. This can be useful in troubleshooting situations.

4.3.3 Zoning architecture

Because of the nature of zoning, it is often easiest to understand by illustration. Here we show several diagrams that show several of the benefits of zoning.

Figure 4-31 shows an example where each host is shown in a separate zone. This should be the standard zoning configuration for a simple environment. The zones are overlapping since the storage ports are included in each zone to allow each server to access the storage. Each host can see all of the Fibre Channel target ports, but cannot see or interact with the other host ports.




Figure 4-31 Hosts in individual zones

Using hard or port zoning, this zoning configuration can be done in advance even if all the servers are not present. Assuming that the storage is connected to ports 1 through 4, each zone can be defined to contain a single port for the server and ports 1 through 4. For example, one zone would consist of port 1, 2, 3, 4, and 5 while the next zone would consist of ports 1, 2, 3, 4, and 6 and so forth.

This diagram shows only a single fabric, but a supported configuration would have two fabrics. The second fabric would have the exact same zone structure.

In Figure 4-32, Host1 and Host2 do not have multipathing software and therefore have to be zoned so that only one path to each LUN is available to them. Therefore, the zone containing these hosts contains only one of the two storage ports. Even though the host has only one HBA, both storage ports were included in the zone, so the LUNs would be visible through two different paths, one going from the host FC port to storage port 0 and the other going from host FC port to storage port 1.




Figure 4-32 Single fabric, no multipathing zoning

Since this figure contains only a single fabric (may consist of one or more switches), it is not considered fully redundant, because a switch or fabric failure would lead to an outage. IBM only considers dual, physically independent fabrics to be fully redundant. However, as shown, Host3 and Host4 have multipathing software and, to protect against a possible storage controller failure, are zoned so that a path to the LUN(s) is available through each of the storage controllers.

Figure 4-33 shows a configuration where Host1 will be accessing LUNs from appliance #1 and Host2 will be accessing LUNs from appliance #2. Both storage appliances are highly available with two storage controllers. This is a fully redundant configuration and both fabrics are shown in this example. Multiple IBM System Storage N series storage server are shown in this diagram, but this is not necessary for high availability. Even a single N3300 A20 system offers high availability.




Figure 4-33 Multiple storage system zoning

This zoning separates the hosts to eliminate initiator (host HBA) cross talk and prevents Host1 from accessing appliance #2, which increases security while improving reliability, manageability, and problem solving.

4.3.4 Zoning recommendations

For Fibre Channel SAN zoning, we recommend the following setups.

Zoning

Anytime four or more hosts are connected to a SAN, zoning should be implemented.

Hard or soft zoning

World Wide Node Name zoning is possible with some switch vendors, but for N series FC Target connections, WWPN zoning is recommended. Because of the various trade-offs, there is no specific recommendation between hard versus soft zoning.

Zone size

For N series, it is generally recommended to keeping the zone size as small as possible while still maintaining manageability. It is not a problem to have many multiple overlapping zones to help keep the individual zones smaller. Ideally, a zone will be defined for each host or host cluster.

4.3.5 Paths

There is one or more paths between the Lotus Domino server and the N series storage. We recommend at least two paths between the server and storage, that is, Multi Path IO (MPIO) setup.

Note: Before implementing the paths, check the compatibility matrix for a supported environment at the following Web site:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

The basic premise behind MPIO is to provide a redundant path to a given LUN. If it is implemented correctly, it can mitigate the issues that are associated with a Single Point of Failure. Typically, these types of failures will be switch or HBA failures. A cable failure is also possible, but those are pretty rare.

In Figure 4-34, the network and associated hardware is properly configured for MPIO, but that alone is not enough. The Initiators and SnapDrive must be configured as well. If we just tell SnapDrive to create a session and a LUN on the filer, there will be only one logical path to the LUNs despite the network configuration.




Figure 4-34 MPIO setup

4.4 Snapshots

A Snapshot copy is a space-efficient, point-in-time image of the data in a volume or an aggregate. Snapshot copies are used for such purposes as backup and error recovery.

Data ONTAP automatically creates and deletes Snapshot copies of data in volumes to support commands related to Snapshot technology. These Snapshots are based on the IBM System Storage N series storage server unique Write Anywhere File Layout ((WAFL); see Figure 4-35).




Figure 4-35 WAFL file system

Data ONTAP also automatically creates Snapshot copies of aggregates to support commands related to the SyncMirror software, which provides RAID-level mirroring. For example, Data ONTAP uses aggregate Snapshot copies when data in two plexes of a mirrored aggregate need to be re-synchronized. In this book, we are usually talking about volume related Snapshots.

A Snapshot copy is a frozen, read-only image of a traditional volume, a flexible volume, or an aggregate that reflects the state of the file system at the time the Snapshot copy was created. Snapshot copies are your first line of defense for backing up and restoring data.

Some facts about Snapshot copies:

- ▶ Data ONTAP maintains a configurable Snapshot schedule that creates and deletes Snapshot copies automatically for each volume.
- ▶ For taking Snapshot copies of LUNs, use SnapDrive. It takes care of flushing all host operating system buffers.
- ▶ You can store up to 255 Snapshot copies at one time on each volume.
- ▶ You can specify the percentage of disk space that Snapshot copies can occupy. The default setting is 20% of the total (both used and unused) space on the disk.

Note: For more information Snapshots, see the *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide* at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

4.5 Protocols

The user data on IBM System Storage N series storage server is accessible through one or more of the access protocols supported by Data ONTAP, including:

- ▶ Network File System (NFS)
- ▶ Common Internet File System (CIFS)
- ▶ HyperText Transfer Protocol (HTTP)
- ▶ File Transfer Protocol (FTP)
- ▶ Fibre Channel Protocol (FCP)
- ▶ Internet SCSI (iSCSI)

The SCSI and NAS protocols differ in their type of I/O operation, that is, *File I/O* and *Block I/O*. As the IBM Lotus Domino server uses File I/O for data access, you can configure a storage system as a target device in an NAS network using Network File System (NFS) or CIFS. You can also configure a storage system as a storage device in an iSCSI network using the SCSI protocol over TCP/IP (using the iSCSI service) and in a SAN network using the SCSI protocol over FC (using the FCP service) to communicate with one or more hosts.

Note: Block storage is normally abstracted by a file system or database management system for use by applications and users. The physical or logical volumes accessed through Block I/O may be devices internal to a server, direct attached, or through iSCSI and FCP. Some database management systems often use their own Block I/O for improved performance and recoverability as compared to layering the DBMS on top of a file system.

IBM Lotus Domino uses File I/O operations. The database files rely on the file system managed by the operating system. This is why we can use the NAS (File I/O; the Data ONTAP operating system will handle the File I/O requests) and SCSI (Block I/O; the Lotus Domino server operating system will handle the File I/O requests and translate them to SCSI Block I/O requests) features of the N series for connecting the Lotus Domino with the IBM System Storage N series storage server.

4.5.1 Lotus Domino protocol recommendations

The IBM Lotus Domino server is able to use the iSCSI, FCP, NFS and CIFS protocols together with the IBM System Storage N series storage server. The recommended protocol depends on your business needs.

Figure 4-36 gives an overview of the NFS, iSCSI, and FCP stacks on the client side. CIFS is similar to NFS, so we only talk about NFS. As seen in the figure, iSCSI needs an additional protocol layer for transporting the SCSI commands, compared to FCP. This makes a slightly performance difference between these protocols.

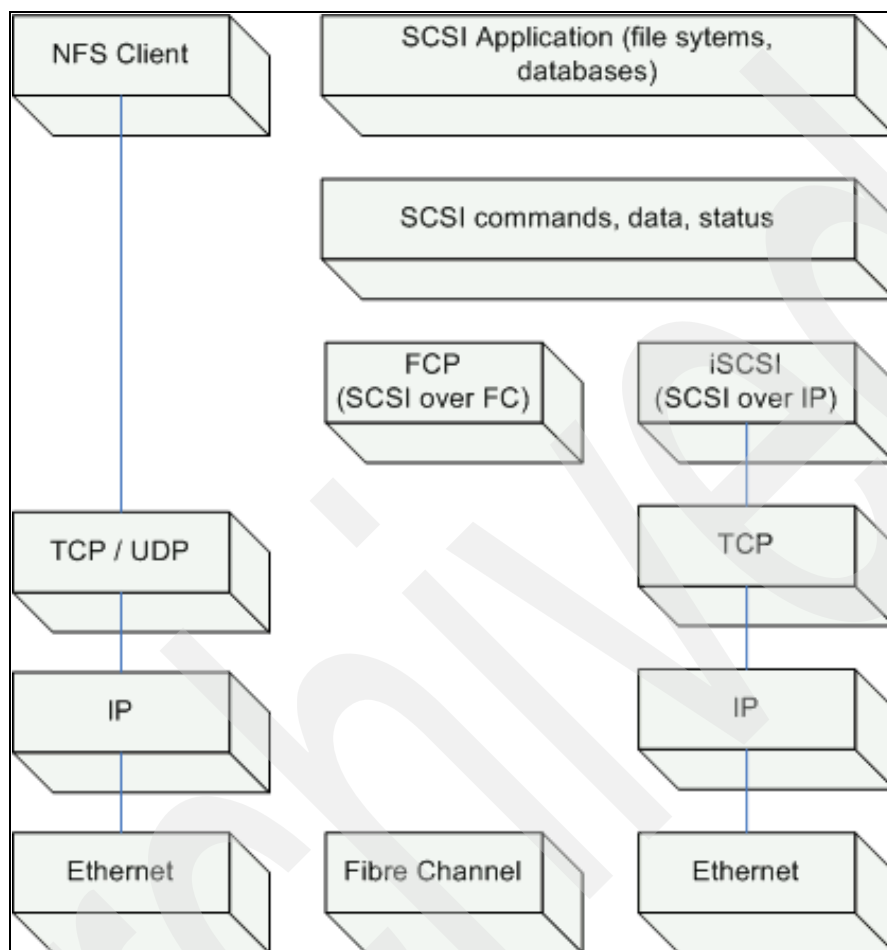


Figure 4-36 NFS, iSCSI, and FCP protocol comparison, client side

The Block I/O based protocols iSCSI and FCP differ in their protocol-overhead. While FCP is designed and optimized for storage operations, iSCSI uses TCP/IP as transport layer. Only if you use iSCSI-HBA, which off loads the TCP/IP and iSCSI protocol handling, can you reduce the additional CPU time needed for protocol overhead.

In our experience, NAS-based approaches, such as CIFS and NFS, might have serious performance bottlenecks due to the layering of file system and application-level structures on top of the storage transport layer, depending on the application.

To have a productive Lotus Domino mail server infrastructure, FCP is the recommended protocol. If you are not able or do not want to spend money on a SAN infrastructure, you might use the iSCSI protocol. In this case, use iSCSI-HBAs or at least network interfaces with the TCP/IP offload engine (TOE). Also, you should use a physical dedicated network for your iSCSI environment.

If you need to share your Lotus Domino database files, the N series provides NFS (or CIFS) capability.

Note: For best performance and reliability, we recommend the Fibre Channel protocol.

4.6 Requirements for Lotus Domino with N series

IBM Lotus Domino on Red Hat Linux or IBM AIX requires the following N series features, licences, and software packages; make sure you have them available:

- ▶ FCP connection

In this book, we use FCP to map the database and transactional log LUN to the IBM Lotus Domino server. Linux and AIX requires the Host Attach Kits to use FCP. Refer to the following links for more information:

- IBM System Storage N series FCP Linux Host Attach Kit

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/allproducts?brandind=5000029>

- IBM System Storage N series FCP AIX Host Attach Kit

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/allproducts?brandind=5000029>

Note: The Host Attach Kits are available from the IBM support for licensed customers only.

- ▶ License

For the following features, you need a license:

- FCP protocol and Host Attach Kit
- SnapDrive for your operating System (Windows, Linux, or UNIX)

Optionally, you might want to use and license the SnapMirror or FlexClone features.

- ▶ Check the compatibility matrix for your environment to make sure you use a supported configuration. For more information, see:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

4.7 SnapDrive

The IBM System Storage N series SnapDrive feature provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file, down through the volume manager, to the storage-system-side logical unit numbers (LUNs) providing the actual repository. In addition, it simplifies the backup of data and helps you decrease the recovery time.

SnapDrive provides a layer of abstraction between an application running on the host operating system and the underlying IBM System Storage N series storage systems (see Figure 4-37). Applications that are running on a server with SnapDrive use virtual disks (or LUNs) on IBM System Storage N series storage systems as though they were locally connected drives or mount points. This allows applications that require locally attached storage, such as IBM Lotus Domino and Microsoft Exchange, to benefit from the N series technologies, including Snapshot, flexible volumes, cloning, and space management technologies.




Figure 4-37 Example of a typical SnapDrive deployment

SnapDrive includes all the necessary drivers and software to manage interfaces, protocols, storage, and Snapshot copies. Snapshot copies are nondisruptive to applications and functions on execution. Snapshot backups can also be mirrored across LAN or WAN links for centralized archiving and disaster recovery.

4.7.1 Benefits of SnapDrive

Most of today's enterprises use business-critical applications such as IBM Lotus Domino, and their storage management team face a number of challenges. They must:

- ▶ Support new business initiatives with minimal increases in their operating budget
- ▶ Protect data from corruption, disaster, and attacks
- ▶ Back up data without any performance degradation, quickly and consistently without any errors

SnapDrive addresses these problems by providing simplified and intuitive storage management and data protection from a host/server perspective. The following list highlights some of the important benefits of SnapDrive:

- ▶ Allows host's and application's administrators to quickly create virtual disks with a dynamic pool of storage that can be reallocated, scaled, and enlarged in real time, even while system are accessing data.

- ▶ Dynamic on-the-fly file system expansion new disks are usable within seconds.
- ▶ Snapshot copies provide rapid backup and recovery capability with minimal resource and capacity requirements.
- ▶ Supports multipath technology for high performance. (Check the compatibility matrix first. At the time the book was written, SnapDrive for Linux did not support multipath technology.)
- ▶ Enables connecting to existing Snapshot copies from the original host or different host.
- ▶ Independent of underlying storage access media and protocol; SnapDrive supports FCP, iSCSI, and NFS as the transport protocols (NFS supports only Snapshot management).
- ▶ Robust and easy-to-use data and storage management feature and software.

4.7.2 SnapDrive requirements

IBM System Storage N series SnapDrive is a licensed feature and is available by contacting IBM Support.

These are some general requirements for SnapDrive:

- ▶ Host operating system and appropriate patches
- ▶ Host file systems
- ▶ IP access between the host and storage system
- ▶ Storage system licenses
- ▶ FCP Host Utilities or iSCSI Host Utilities required software
- ▶ For security reasons, we recommend a separate user account on the IBM System Storage N series storage server.
- ▶ Depending on the operating system, the system requirements can be found on the following links:
 - *IBM System Storage N series SnapDrive for Windows 4.2.1 Release Notes*, found at: <http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001650&aid=1>
 - *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at: <http://www.ibm.com/servers/storage/support/nas/snapdrive/>

4.8 SnapMirror

The Data ONTAP SnapMirror feature allows IBM System Storage N series Snapshot images to mirror SnapShot images either asynchronously or synchronously over the network for backup or disaster recovery purposes.

Figure 4-38 shows two steps for creating and running the SnapMirror feature. At first, a baseline SnapShot image replication is done. Several transport media could be used, such as LAN, FCP, or Tape.

Step two shows the asynchronous, synchronous, or semi-synchronous mirroring process from source to target.

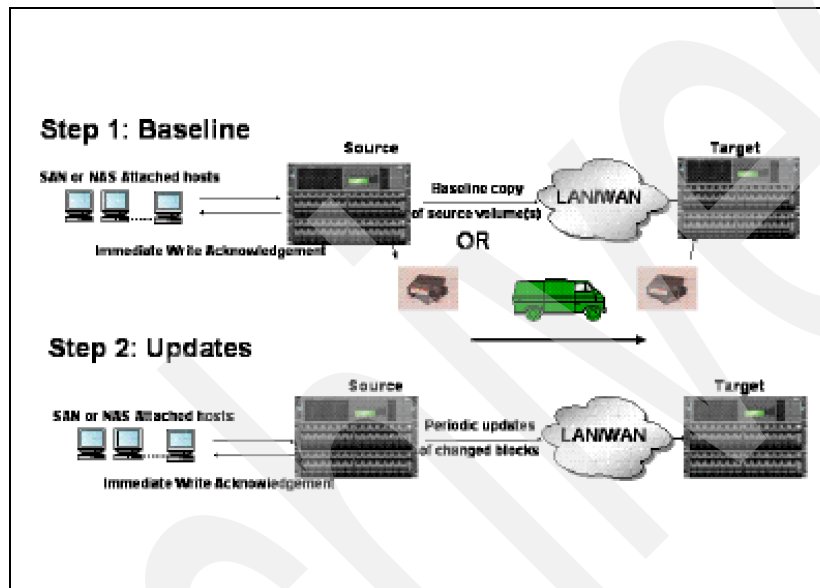


Figure 4-38 SnapMirror

The three SnapMirror replication modes are described in the following sections.

4.8.1 Asynchronous mode

In asynchronous mode, SnapMirror performs incremental, block-based replication as frequently as once per minute. The performance impact on the source IBM System Storage N series storage system is minimal as long as the system is configured with sufficient CPU and disk I/O resources. The administrator can decide on the replication interval depending on the business needs.

4.8.2 Synchronous mode

The Synchronous SnapMirror mode is a SnapMirror feature that replicates data from a source volume to a partner destination volume at or near the same time that it is written to the source volume, rather than according to a predetermined schedule. The chance of data loss is minimized in this mode.

Note: Use the same type and amount of physical hard disks and aggregate configuration on the SnapMirror source and destination for reducing the performance impact.

4.8.3 Semi-synchronous mode

To improve the performance of the synchronous SnapMirror feature, it can be configured to lag behind the source volume by a user-defined number of write operations or milliseconds, which is known as semi-synchronous SnapMirror mode.

This mode is like asynchronous mode in that the application does not need to wait for the secondary storage to acknowledge the write before continuing with the transaction. This mode is like synchronous mode in that updates from the primary storage to the secondary storage occur right away, rather than waiting for scheduled transfers.

4.8.4 Benefits of SnapMirror

The IBM Data ONTAP SnapMirror feature replicates data not only for backup and disaster recovery purposes. The following list shows some use cases:

- ▶ Using the SnapMirror target for local read access at remote sites
Fast access to corporate data on the remote site and offloading the source side
- ▶ Off load tape backup CPU cycles to a mirror
- ▶ Isolate testing from the production volume
ERP testing and Offline Reporting
- ▶ Cascading Mirrors
Replicated mirrors on a larger scale
- ▶ Disaster recovery
Replication to “hot site” for mirror failover and eventual recovery
- ▶ You can use the Data ONTAP SnapMirror feature in combination with a FlexClone volume to perform migration faster and more efficiently:
 - For corporations with a warm backup site or that need to off load backups from production servers
 - For generating queries and reports on near-production data

4.8.5 Effects on sizing

This Data ONTAP feature mirrors your selected data, so you need the same size used on the source site for the target location of your mirror. As a best practice, use the same aggregate and volume configuration on both sites when synchronous mirroring is used.

4.8.6 SnapMirror requirements

The following prerequisites must be met before you can run SnapMirror:

- ▶ You must purchase and enable the SnapMirror license.
If the SnapMirror source and destination are on different systems, you must purchase a license and enter the SnapMirror license code on each system.

- ▶ For SnapMirror volume replication, you must create a restricted volume to be used as the destination volume.
- ▶ For SnapMirror volume replication, the destination volume must run under a version of Data ONTAP that is equal to or later than that of the SnapMirror source volume.

If you configure volume SnapMirror to support replication for the purpose of disaster recovery, both the source and destination systems should run the same version of ONTAP software.

Note: If you upgrade your systems to a later version of Data ONTAP, upgrade the systems of SnapMirror destination volumes before you upgrade the systems of the SnapMirror source volumes.

- ▶ The name and IP address of the source system must be in the `/etc/hosts` file of the destination system or must be resolvable by way of DNS or yp.

4.9 SnapVault

SnapVault is a low impact, disk-based online backup of heterogeneous storage systems for fast and simple restores. SnapVault is a separately licensed feature in Data ONTAP that provides disk-based data protection for storage systems.

The SnapVault server runs on the IBM System Storage N series platform. The SnapVault feature replicates selected Snapshots from multiple storage systems or backups data from an open system platform to a common Snapshot on the SnapVault server. Figure 4-39 shows a common SnapVault implementation with an open system platform and N series storage server on the primary site. The changed data is transferred to the N series secondary storage system, where regularly made Snapshots are saved on tape.

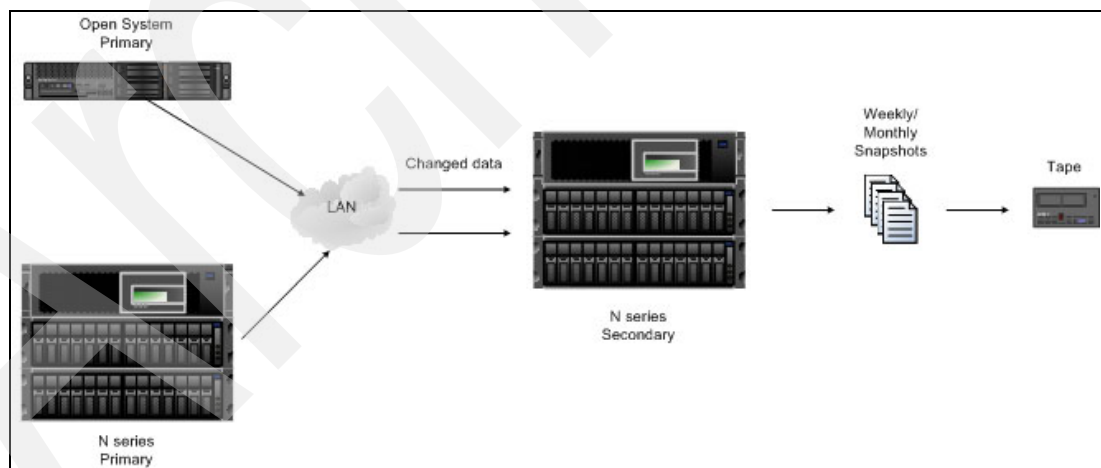


Figure 4-39 Common SnapVault implementation

Note: Because Data ONTAP SnapDrive uses *qtrees* as backup storage, it is not supported for volumes containing Data ONTAP LUNs, as we use it in this book for Lotus Domino and Microsoft Exchange. For more information about SnapVault, see the *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

4.9.1 Benefits of SnapVault

SnapVault is the IBM System Storage N series storage server solution for a fast disk-based backup and recovery solution.

IBM Lotus Domino benefits from these advantages when used with NFS or CIFS. Databases reside on an external disk based backup system for fast recovery. The space consumption is less because of the SnapShot technology. Also, this makes SnapVault available as a WAN backup solution for Lotus Domino.

The Data ONTAP feature contains the following benefits:

- ▶ It avoids the bandwidth limitations of tape drives so the restore can be faster.
- ▶ It does not require full dumps from the primary storage, so there is no need for a backup window.
- ▶ It is a data protection solution for heterogeneous storage environments.
- ▶ It performs disk-to-disk backup and recovery.
- ▶ It is designed to address pain points associated with tape.
- ▶ Intelligent data movement reduces:
 - Network traffic.
 - Impact on production systems.
- ▶ It has frequent backups to ensure superior data protection.
- ▶ It is based on Data ONTAP Snapshot technology:
 - Significantly reduces the amount of backup media.
 - Reduced backup overhead: Incrementals and changed blocks only.
 - Instant single file restore: Snapshot directory displays SnapVault Snapshots.
 - Can protect remote sites over a WAN.

4.9.2 SnapVault requirements

The Data ONTAP SnapVault feature needs to be licensed. Additionally, you need to complete the following tasks to activate SnapVault:

- ▶ Enter the license code on both the secondary storage system and the primary storage system.
- ▶ Turn on the `snapvault.enable` option for both the secondary storage system and the primary storage system. Setting this option enables SnapVault data transfers and Snapshot copy creation.
- ▶ Turn on the NDMP service on both the secondary storage system and the primary storage system.
- ▶ Set the `snapvault.access` option on both the secondary storage system and the primary storage system to allow primary storage system and secondary storage system access.
 - Setting this option on the SnapVault secondary storage system determines which SnapVault primary storage systems can access the secondary storage system.
 - Setting this option on the SnapVault primary storage system determines which secondary storage system can access data from that primary storage system.

4.9.3 Effects on sizing

The storage space required to implement the SnapVault feature will depend upon the number of backup copies created. On the SnapVault back-end server, you need at least the size of your source storage. Every further SnapShot based backup increases the storage space requirement on the SnapVault target server for the count of changed blocks.

4.10 FlexClone

This feature allows you to instantly create clones of a FlexVol flexible volume in few seconds without interrupting the parent flexible volume. A FlexClone volume is a writable point-in-time image of a flexible volume of another FlexClone volume (see Figure 4-40). They use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between parent and clone. In addition to these benefits, clone volumes have the same high performance as other kinds of volumes.




Figure 4-40 FlexClone

A FlexClone LUN clone shares space with the LUN in the backup SnapShot copy. The clone does not require additional disk space until changes are made to it. You cannot delete the backing SnapShot copy until you split the clone from it. When you split the clone from the backing SnapShot copy, you copy the data from the SnapShot copy to the clone. After the splitting operation, both the backup SnapShot copy and the clone occupy their own space.

Note: For more information about using volume cloning with LUNs, see the *IBM System Storage N series Block Access Management Guide for FCP and iSCSI* at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

4.10.1 FlexClone use cases and benefits

The main difference between a Snapshot and cloning is that a cloned volume or LUN is writable. This is an important and powerful feature for Lotus Domino administrators. In the case of a database or software upgrade, release, or configuration change, a clone of the production system can be made. This clone can be used for testing with your real data, without worrying about damaging the production system.

Volume cloning with the FlexClone feature provides similar results to volume copying, but cloning offers some important advantages over volume copying:

- ▶ Volume cloning is instantaneous, whereas volume copying can be time consuming.
- ▶ If the original and cloned volumes share a large amount of identical data, considerable space is saved because the shared data is not duplicated between the volume and the clone.

Cloning is not recommended as a primary backup method, because all data still resides on the same volume and aggregate. You might use LUN cloning for following reasons:

- ▶ **Application Testing:** Make the necessary changes to the infrastructure without worrying about crashing the production system. Avoid making untested changes on the system under tight maintenance window deadlines.
- ▶ You need to make a copy of your data available to additional users without giving them access to the production data.
- ▶ You want to create a clone of the Lotus Domino databases for manipulation and projection operations, while preserving the original data in unaltered form.
- ▶ You want to access a specific subset of a LUN's data (a specific logical volume or file system in a volume group, or a specific file or set of files in a file system) and copy it to the original LUN, without restoring the rest of the data in the original LUN. This works on operating systems that support mounting a LUN and a clone of the LUN at the same time. SnapDrive for Linux and UNIX allows this with the **snap connect** command (this could be done with a Snapshots, too).
- ▶ **Data Mining:** Data mining operations and software can be implemented more flexibly because both reads and writes are allowed.
- ▶ **Parallel Processing:** Multiple FlexClone volumes of a single milestone/production data set can be used by parallel processing applications across multiple servers to get results more quickly.
- ▶ **System Deployment:** Maintain a template environment and use FlexClone volumes to build and deploy either identical or variation environments.
- ▶ **IT Operations:** Maintain multiple copies of production systems: live, development, test, reporting, and so on. Refresh working FlexClone volumes regularly to work on data as close to live production systems as is practical.

Note: A FlexClone volume can be created from a Snapshot copy in a SnapMirror destination, but a FlexClone volume cannot be the destination of a SnapMirror relationship.

Figure 4-41 shows a common situation for Data ONTAP FlexClone deployment. The IT staff needs to make substantive changes to a production environment. The cost and risk of a mistake are too high to do it on the production volume. Ideally, there would be an instant writable copy of the production system available at minimal cost in terms of storage and service interruptions.

By using FlexClone volumes, the IT staff gets just that: an instant point-in-time copy of the production data that is created transparently by SnapShot and uses only enough space to hold the desired changes. They can then try out their upgrades using the FlexClone volumes.

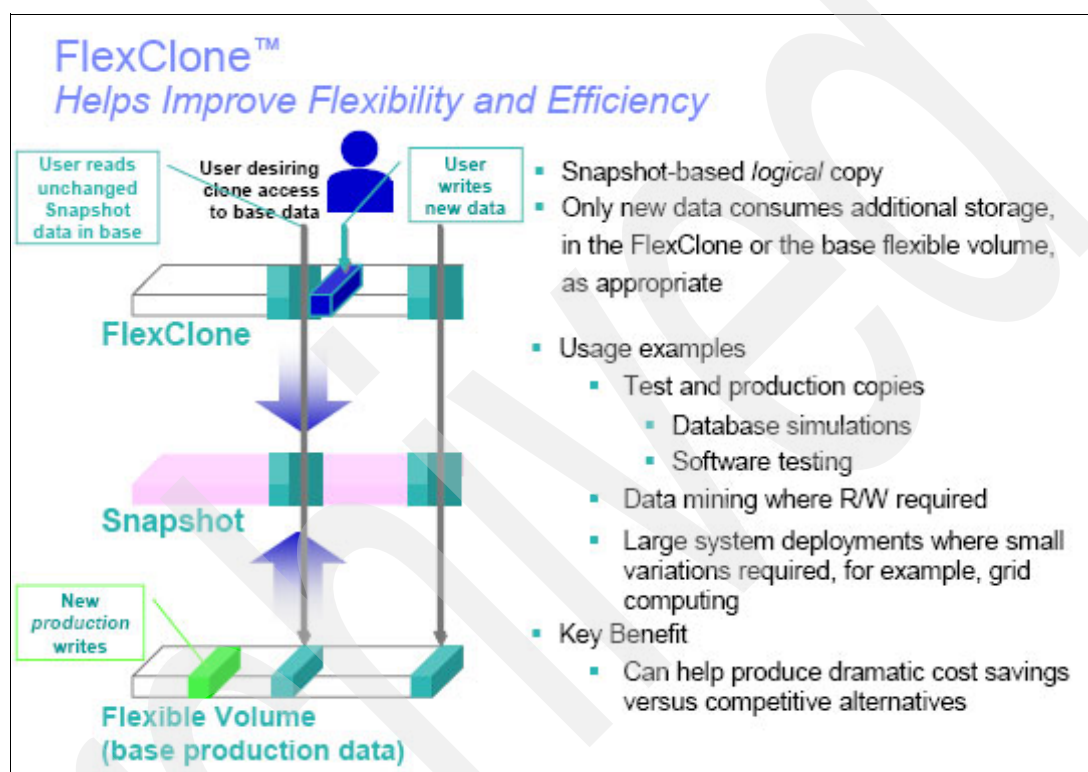


Figure 4-41 FlexClone basics

At every point where they make solid progress, they clone their working FlexClone volume to lock in the successes. At any point where they get stuck, they just destroy the working clone and go back to the point of their last success. When everything is finally working just the way they want it to work, they can either split off the clone to replace their current production volumes or modify their successful upgrade process to use on the production system during the next maintenance window.

The FlexClone feature allows them to make the necessary changes to their infrastructure without worrying about crashing their production systems or making untested changes on the system under tight maintenance window deadlines. The results are less risk, less stress, and higher levels of service for the IT customers.

4.10.2 FlexClone requirements

The Data ONTAP FlexClone feature needs to be licensed. Additionally, you need to follow these points to use FlexClone:

- You must install the license for the FlexClone feature before you can create FlexClone volumes.

- ▶ FlexClone volumes and their parent volumes share the same disk space (aggregate) for any data common to the clone and parent. This means that creating a FlexClone volume is instantaneous and requires no additional disk space (until changes are made to the clone or parent).
- ▶ While a FlexClone volume exists, some operations on its parent are not allowed. For more information, refer to *IBM System Storage N series Data ONTAP 7.2 Storage Management Guide* at:
<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001595&aid=1>
- ▶ Only flexible volumes can be cloned. To create a copy of a traditional volume, you must use the **vol copy** command, which creates a distinct copy with its own storage.

4.10.3 Effects on sizing

When a FlexClone volume is created, it shares all of its data with its parent volume. So even though its logical size is the same as its parent's size, depending on how much data it contains, it can use very little free space. As data is written to either the parent or the FlexClone volume, that data is no longer shared between the parent and FlexClone volumes, and the FlexClone volume starts to require more space from its containing aggregate, depending on the size of your changed data.

If you want to split the clone from its parent, it removes any space optimizations that are currently employed by the FlexClone volume. After the split, both the FlexClone volume and the parent volume require the full space allocation determined by their space guarantees.

Note: Because the clone-splitting operation is a copy operation that might take considerable time to carry out, Data ONTAP also provides commands to stop or check the status of a clone-splitting operation.

Archived



Preparing the IBM System Storage N series for Microsoft Exchange

This chapter will help you configure the IBM System Storage N series to host a Microsoft Exchange environment. Some features will be presented in this chapter and their use will be explained.

5.1 Storage requirements for a Microsoft Exchange server

Microsoft Exchange servers require responsiveness from the underlying storage infrastructure to support user activities, backup and restore operations, database maintenance, and so on. The requirements for all of these actions can be divided into two basic categories:

- ▶ Capacity
- ▶ I/O performance

There are a number of situations that impact the capacity of the storage infrastructure, and most of them are defined by the Microsoft Exchange administrators. Here are a few examples:

- ▶ Number of users on the server
- ▶ Mailbox quota limits
- ▶ Deleted items in cache
- ▶ Mailbox retention policy
- ▶ Estimated growth in number of users
- ▶ Public folders
- ▶ Third-party Microsoft Exchange add-on applications that would require additional storage

Also, there are some situations that impact the I/O performance of the server. Some examples are:

- ▶ User activity, such as sending/receiving e-mails, calendar sharing, attachments, and accessing public folders
- ▶ Smart mobile devices accessing user data
- ▶ Server specific events, such as online database maintenance, indexing, and backup and restore

5.1.1 Capacity

Proper storage sizing is the key to a successful Microsoft Exchange deployment. If you size too small, you will obviously run out of space and the Microsoft Exchange Storage Groups will be dismounted and unavailable for users. If you size too large, you are not efficiently using your storage capacity. Despite the fact that the IBM System Storage N series storage system allows for dynamic resizing of aggregates, volumes and LUNs, an optimized size plan helps reduce wasted storage, increases backup and restore operations, and improves performance.

Microsoft Exchange servers access the LUN created on the volumes on an IBM System Storage N series storage system. For efficient capacity sizing, you must first determine the LUN size for the Microsoft Exchange server. After the LUN size is determined, we can determine the volume size needed to support that LUN for Microsoft Exchange servers and all the other needed features.

A best practice is to store the Microsoft Exchange databases on different physical disks (or aggregates) than the log files for that database (for both performance and reliability purposes). Figure 5-1 on page 121 shows the Microsoft Exchange components' distribution on an IBM System Storage N series storage system.

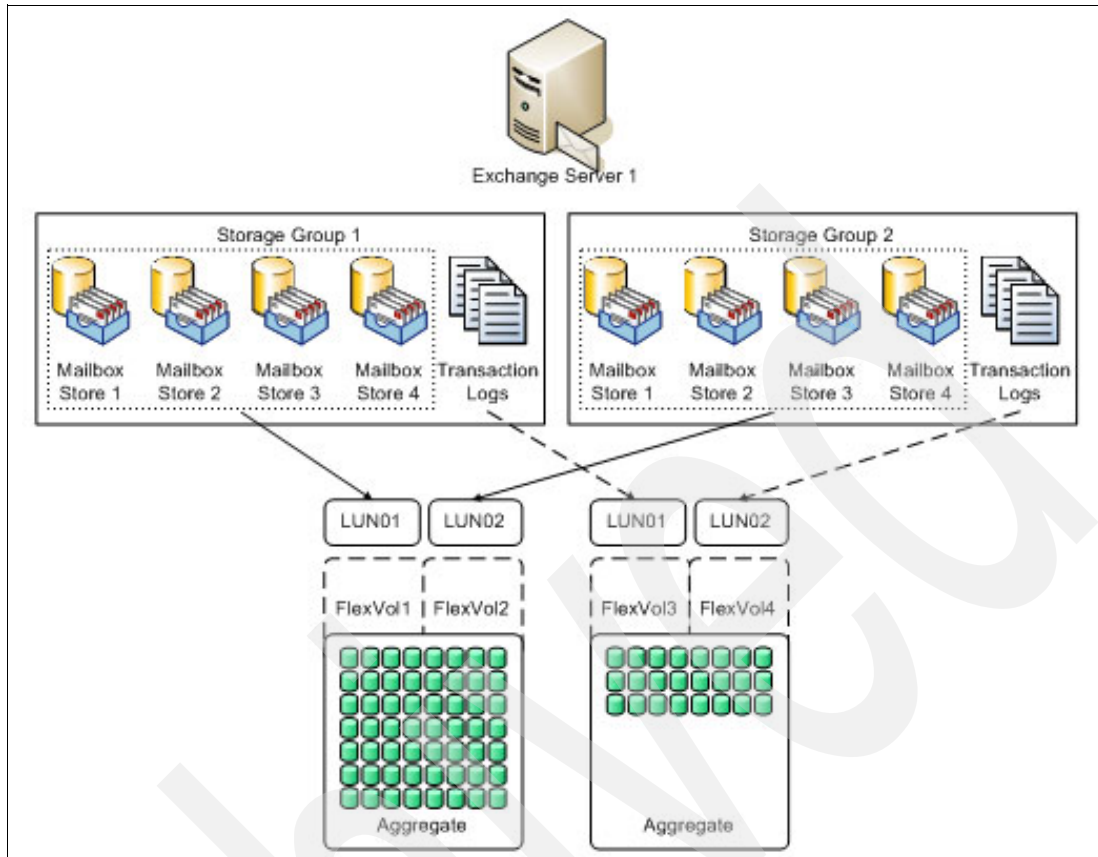


Figure 5-1 Microsoft Exchange databases and transaction logs on an IBM System Storage N series storage system

There are some rules enforced by SnapManager for Microsoft Exchange when running the Configuration Wizard:

- ▶ Microsoft Exchange databases cannot share a LUN with databases from different storage groups.
- ▶ Microsoft Exchange databases cannot share the same LUN containing transaction log files.
- ▶ Microsoft Exchange transaction logs from different storage groups can share the same LUN.

As a best practice, place your SMTP and MTA queues on any LUN or disk that is not hosting any Microsoft Exchange database files. Depending on the business need, you may have to create a LUN on the N series storage or you may use your local disks for this. Figure 5-2 shows the MTA and SMTP queues on the server but optionally on the same LUN as the transaction logs.

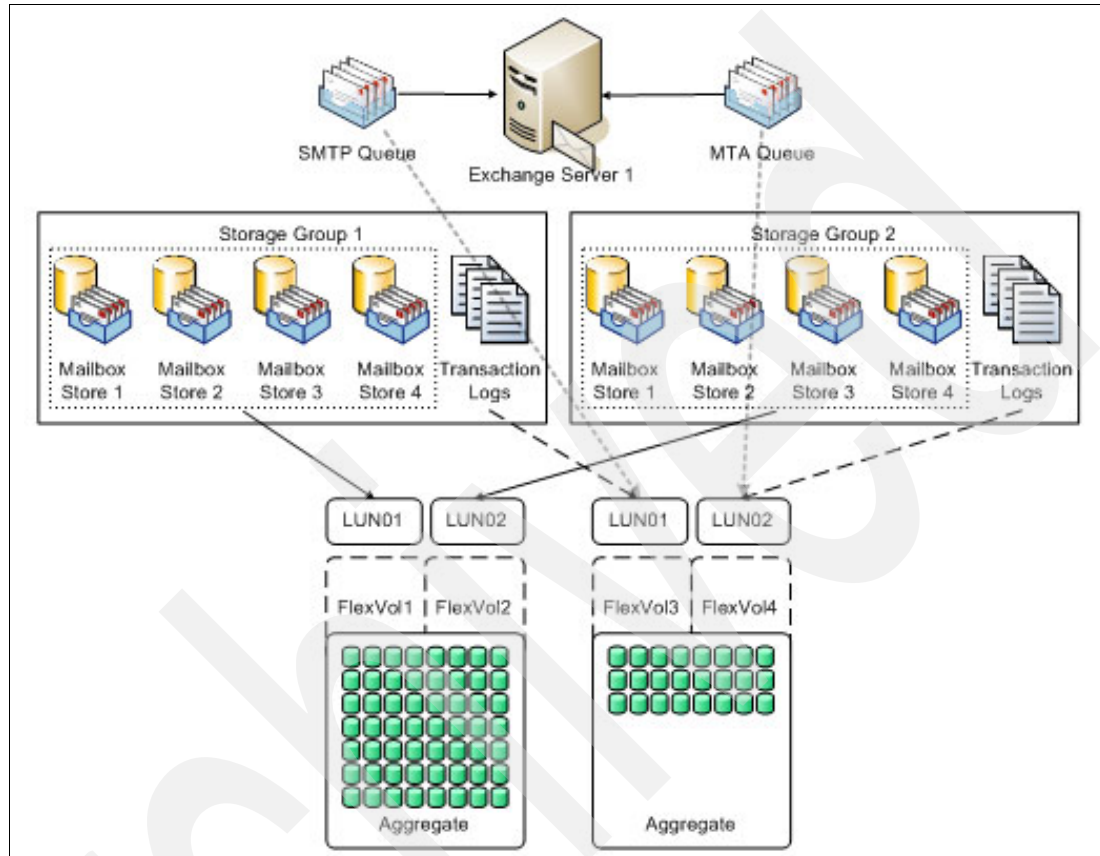


Figure 5-2 SMTP and MTA queues on IBM System Storage N series storage system

Note: In the above scenario, a LUN is being created for each Storage Group database on the Microsoft Exchange Server and another LUN is being created for each Storage Group transaction logs. A deeper level of granularity for restore (such as single Mailbox Store restore) may be needed. In this case, you can set up one LUN for each Mailbox Store on the same FlexVol. The architecture may vary, depending on the business needs.

Database LUN sizing

To calculate the LUN size needed to host your Microsoft Exchange database, you need the following information:

- ▶ Number of users on the Storage Group
- ▶ Mailbox size limit policy
- ▶ Deleted items in cache percentage
- ▶ Single instance ratio (average number is 20%)

The formula to calculate the Database LUN size is:

Database LUN size = ((number of mailboxes * mailbox quota) / single instance ratio) + deleted items cache

For example:

- ▶ 1200 users on the Storage Group.
- ▶ The mailbox size limit is 100 Mb.
- ▶ The deleted items cache is 15%.

Database LUN size = $((1200 * 100) / 1.20) + 15\% = 115000 \text{ Mb}$ or 115 Gb

Transaction logs LUN sizing

To calculate the LUN size needed for the transaction logs, you need to determine an average number of log files that will be generated between backups. Microsoft Exchange Server 2003 log files are 5 MB in size, and Microsoft Exchange Server 2007 log files are 1 MB in size.

The formula to calculate the Transaction Logs LUN size is:

Transaction Logs LUN size = number of transaction log files generated between backups * log size

For example:

- ▶ 1200 transaction logs generated between backups
- ▶ Microsoft Exchange Server 2003 (which means 5 MB transaction logs size)

Transaction Logs LUN size = $1200 * 5 = 6000 \text{ MB}$ or 6 GB

SnapInfo Directory LUN sizing

SnapInfo Directory is a special storage directory where the transaction log files are copied to when a backup is started from SnapManager for Microsoft Exchange (SME).

Note: The recommended configuration is to store both the Transaction Logs and the SnapInfo Directory on the same LUN.

Remember that Microsoft Exchange Server 2003 log files are 5 MB in size, and Microsoft Exchange Server 2007 log files are 1 MB in size.

The formula to calculate the SnapInfo Directory size differs if the directory will be placed on its own LUN or in the Transaction Logs LUN.

SnapInfo Directory LUN size = (number of transaction log files generated between backups * log size) * number of backups to keep online

For example:

- ▶ 1200 transaction logs generated between backups
- ▶ Microsoft Exchange Server 2003 (which means 5 MB transaction logs size)
- ▶ Seven backups kept online

SnapInfo Directory LUN size = $(1200 * 5) * 7 = 42000 \text{ MB}$ or 42 GB

The formula to calculate the SnapInfo Directory size, when stored on the same LUN as the Transaction Logs, is just like adding the two results:

SnapInfo Directory and Transaction Logs LUN size = $((\text{number of transaction log files generated between backups} * \text{log size}) * (\text{number of backups to keep online} + 1))$

For example:

- ▶ 1200 transaction logs generated between backups
- ▶ Microsoft Exchange Server 2003 (which means 5 MB transaction logs size)
- ▶ Seven backups kept online

SnapInfo Directory and Transaction Logs LUN size = $(1200 * 5) * (7 + 1) = 48000$ MB or 48 GB

All of the LUN sizing above shows the exact LUN size for a specific situation. It is always a best practice to add enough free space to support growth on the database or log files. Make sure you have at least 10% to 15% more space on the LUN.

Database volume sizing

After calculating the LUN size for the database, you need to determine the volume sizing for the LUN or LUNs that will be stored on it. In our scenario, one LUN is being stored per volume. If you have different LUNs per volume, the number should be added to calculate the volume size.

The formula to calculate the Database volume size is:

Database volume size = $(2 * \text{database LUN size}) + (\text{number of online backups} * \text{data change percentage} * \text{max database size})$

For example:

- ▶ Database LUN size is 115 GB + 15% = 133 GB
- ▶ 10% data change between backups
- ▶ Seven backups kept online

Database volume size = $(2 * 133) + (7 * 10\% * 133) = 360$ GB

The database LUN size is being multiplied by 2 to guarantee 100% space reservation on the volumes hosting the LUN. In this manner, if you need to increase the LUN size, you can do this by expanding the LUN from SnapDrive. You will not even need to restart the Microsoft Exchange server.

If for any reason you do not want or do not have available storage space to guarantee 100% growth, you can use fractional space reservation on your sizing. Simply replace 2 (which means 100%) by $1 + \text{fractional space reservation needed}$:

Database volume size = $((1 + \text{fractional space needed}) * \text{LUN size}) + (\text{number of online backups} * \text{data change percentage} * \text{max database size})$

Transaction Logs and SnapInfo Directory volume sizing

Remember that Microsoft Exchange Server 2003 log files are 5 MB in size, and Microsoft Exchange Server 2007 log files are 1 MB in size.

To calculate the Transaction Logs and SnapInfo Directory volume size, we need to consider the LUN size for both the transaction logs and the SnapInfo Directory.

The formula to calculate the Transaction Logs and SnapInfo Directory volume size is:

Transaction Logs and SnapInfo Directory volume size = $(2 * \text{transaction logs LUN size}) + (2 * \text{number of transactional logs generated between backups} * \text{log size} * \text{number of backups to keep online})$

For example:

- ▶ The Transaction Logs LUN size is 48 GB + 15% = 55.2 Gb or 55200 MB.
- ▶ 1200 transaction logs generated between backups.
- ▶ Microsoft Exchange Server 2003 (which means a 5 MB transaction logs size).
- ▶ Seven backups kept online.

Transaction Logs and SnapInfo Directory volume size = $(55200 * 2) + (2 * 1200 * 5 * 7) = 194400$ MB or 195 GB

The Transaction Log LUN size is being multiplied by 2 to guarantee 100% space reservation on the volumes hosting the LUN. In this manner, if you need to increase the LUN size, you can do this by provisioning more storage and the LUN from SnapDrive. You will not even need to restart the Microsoft Exchange server.

If for any reason you do not want or do not have available storage space to guarantee 100% growth, you can use fractional space reservation on your sizing. Simply replace 2 (which means 100%) by $1 + \text{fractional space reservation needed}$.

Transaction Logs and SnapInfo Directory volume size = $(\text{transaction logs LUN size} * (1 + \text{fractional space reservation})) + (2 * \text{number of transactional logs generated between backups} * \text{log size} * \text{number of backups to keep online})$

When you have the Transaction Logs and SnapInfo Directory on the same NTFS volume and LUN, which is the recommended configuration, SME automatically configures log archiving to use NTFS Hard Links. During a normal backup, SME creates copies of the transaction log files on the appropriate SnapInfo Directory for archiving. With NTFS Hard Links, and in supported configurations, when a backup is started, links are created to the transaction log files instead of copying the files.

The formula to calculate the Transaction Logs and SnapInfo Directory volume size is slightly different when using NTFS Hard Links:

Transaction Logs and SnapInfo Directory with NTFS Hard Links volume size = $(\text{transaction logs LUN size} * 2) + (\text{number of transactional logs generated between backups} * \text{log size} * \text{number of backups to keep online})$

For example:

- ▶ The Transaction Logs LUN size is 48 GB + 15% = 55.2 GB or 55200 MB.
- ▶ 1200 transaction logs generated between backups.
- ▶ Microsoft Exchange Server 2003 (which means 5 MB transaction logs size).
- ▶ Seven backups kept online.

Transaction Logs and SnapInfo Directory with NTFS Hard Links volume size = $(55200 * 2) + (1200 * 5 * 7) = 152400$ MB or 153 GB

Using the sizing information stated, this is the scenario for LUN and volume creation on the N series storage system:

- ▶ Aggregate 1
 - Database volume = 360 GB
 - Database LUN = 133 GB
- ▶ Aggregate 2
 - Transaction Logs and SnapInfo Directory volume = 153 GB
 - Transaction Logs and SnapInfo Directory LUN = 42 GB

SMTP and MTA queues will reside on the local disk of the server, without being moved to any LUN on the IBM System Storage N series storage system.

5.1.2 I/O performance

Most of the overall performance of a Microsoft Exchange environment is determined by the server disk subsystem. When moving your existing Direct Access Storage based Microsoft Exchange system to an IBM System Storage N series SAN, you will notice an improvement in Microsoft Exchange performance, because the disk subsystem will perform better and so will the Microsoft Exchange system. This better performance is due to the use of the aggregates.

An aggregate is a collection of disks that will be available for the volumes. This means that the disk subsystem's performance will become better with every disk you add to the aggregate, as I/Os are spread across more disks. In general, FlexVol on large aggregate performance is higher than the volumes on the separate aggregate.

After sizing the LUNs and volumes that should be available for the Microsoft Exchange systems, the aggregate should be created to accommodate the volumes. This aggregate sizing is a very important step because it will impact the performance of the systems and servers that will be using these disks. I/O is spread across the disks in an aggregate, and more disks help improve performance by creating a greater spread.

The aggregate should accommodate all the volumes that are planned to be on it, and have free space in case a volume expansion is needed. This free space will vary depending on the business need and disk availability, but 30% to 40% would be an average space reservation for the aggregate. Keep this in mind when planning the sizing for your Microsoft Exchange infrastructure.

Exchange storage has specific latency requirements for database and transaction log LUNs. The transaction log LUN should be placed on the fastest storage with a goal of less than 10 millisecond (<10ms) writes. The database LUN requires read and write response times of less than 20 milliseconds (<20ms).

Effect of host side memory on storage

As a result of the larger amounts of memory available in the 64-bit platform, the Exchange Server 2007 database cache can be larger. This can reduce the number of reads and writes to disk. In Exchange Server 2003, the maximum was 900 MB of database cache. With additional memory available in Exchange Server 2007, this amount can increase from 900 MB to multiple GBs, depending on the amount of server system memory installed. As additional memory is added to the host server, the Exchange server database cache can increase, further decreasing the number of disk reads. It is important to understand how the increased size of the database cache affects the host side memory when planning storage and IOPS.

Increasing the amount of the database cache also affects IOPS. When testing 4,000 users with 8 GB of server memory, the database cache is limited to 1.48 MB/user [(8 GB – 2048) / 4000]. Increasing the amount of server memory to 16 GB increases the amount of the database cache to 3.48 MB/user [(16GB – 2048) / 4000].

5.2 Storage configuration for Microsoft Exchange server

The IBM System Storage N series storage system must be configured prior to running the Microsoft Exchange server on it. The aggregates, volumes, LUNs and Snapshots must be created and configured to support the Microsoft Exchange server environment.

5.2.1 Aggregates

An aggregate is a collection of physical disks from which the space is allocated to the volumes. When creating the aggregates on the IBM System Storage N series storage system, there are some considerations:

- ▶ On each aggregate, one or more flexible volumes can be created.
- ▶ Each aggregate has its own RAID configuration and set of assigned physical disks.
- ▶ The available space on the aggregate can be increased by simply adding disks to the existing RAID group or by adding new RAID groups.
- ▶ Performance is proportional to the number of disk spindles on the aggregate.

The aggregate consists of one or more plexes. A plex is a collection of one or more RAID groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system volumes. By default, the aggregates are created on plex0.

For detailed information about aggregates, WAFL file system, and Data ONTAP V7.2, refer to the document *IBM System Storage N series Data ONTAP 7.2 Storage Management Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssglS7001595&aid=1>

There are two types of aggregates: unmirrored aggregates and mirrored aggregates.

Unmirrored aggregates consist of a single plex. In Figure 5-3, an unmirrored aggregate named aggr A is being shown. This aggregate is made up of three RAID-DP groups named rg0, rg1, and rg2 on plex plex0.

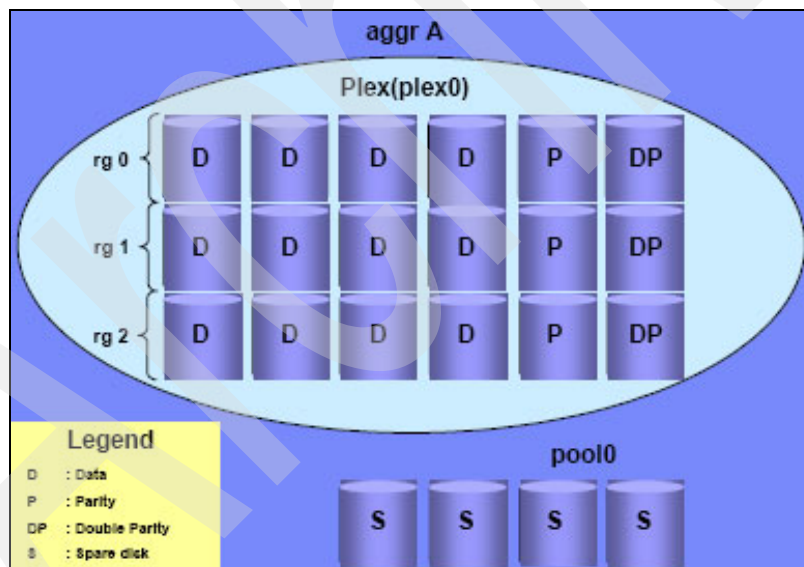


Figure 5-3 Unmirrored aggregate

Mirrored aggregates consist of two plexes that are mirrored to each other. A mirrored aggregate named aggr A is shown in Figure 5-4. This aggregate is made up of three RAID-DP groups named rg0, rg1, and rg2 on plexes plex0 and plex1.




Figure 5-4 Mirrored aggregate

In this scenario, two aggregates are required for the Microsoft Exchange server infrastructure: one aggregate for the databases and one aggregate for the transaction log files. This configuration can vary depending on many factors. Independent of the IBM System Storage N series storage system's configuration, the databases and the transaction log files should always be placed on different aggregates for performance and predict purposes.

Creating aggregates

When creating the aggregate, a name should be defined. There are some naming conventions for the aggregate's name. It should:

- ▶ Begin with either a letter or an underscore
- ▶ Contain only letters, digits, and underscores
- ▶ Contain no more than 255 characters

After you have the name, size, and disk configurations planned, these are the steps to create an aggregate:

1. Open the FilerView for the IBM System Storage N series storage system where you want to create the aggregate.

2. On the FilerView, select **Aggregates** → **Add**. This will bring up the Add New Aggregate window, as shown in Figure 5-5. Click **Next**.




Figure 5-5 Add New Aggregate window

3. The Aggregate Name window will be shown (Figure 5-6). Type in the name for the aggregate you are creating. Select whether this aggregate will be a mirrored aggregate (check box Mirror checked) or an unmirrored aggregate (check box Mirror unchecked). The parity should also be defined in this window. If you are creating a RAID-DP based aggregate, select the **Double Parity** check box. If the Double Parity check box is unchecked, the aggregate will be created using RAID 4. Click **Next**.

http://9.11.218.237/ - itsotuc3: Aggregate Wizard - Windows Internet Explorer

Aggregate Wizard - Aggregate Parameters

Aggregate Name:
Enter a name for the new aggregate.

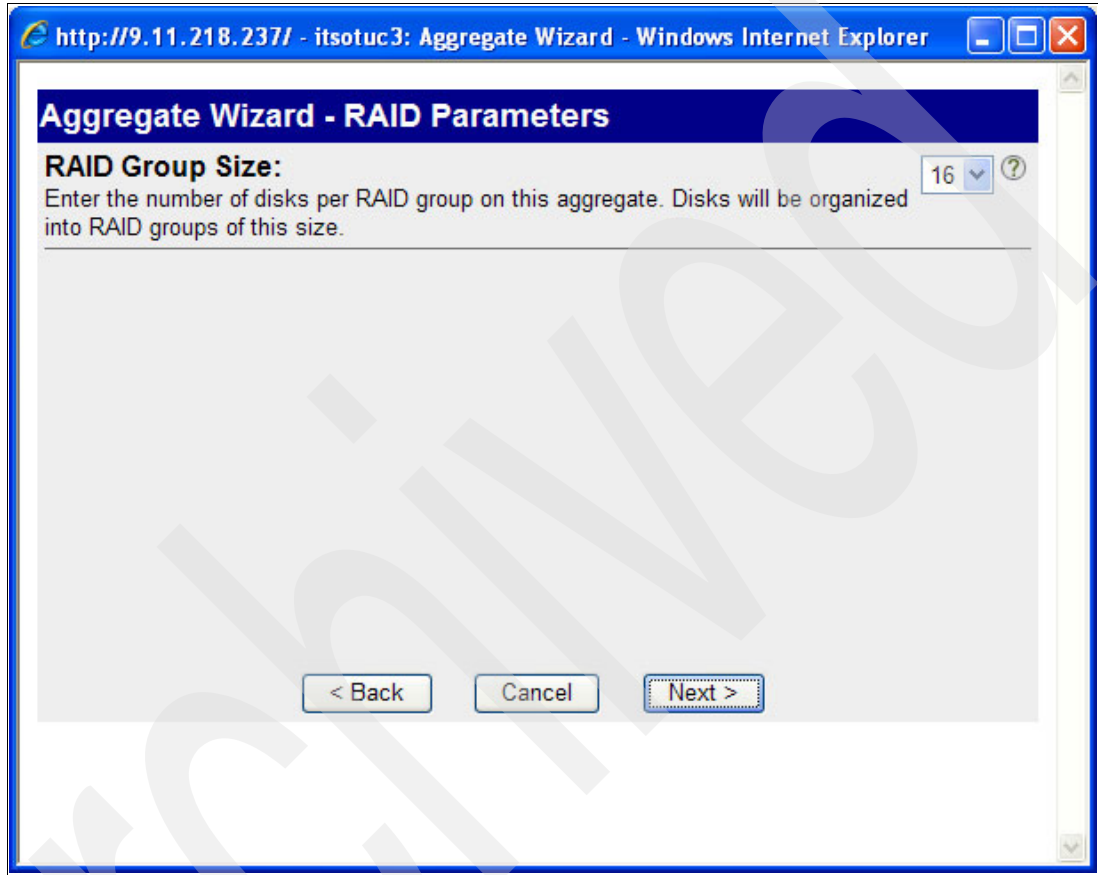
Synchronous Mirroring:
Select to enable local synchronous mirroring on this aggregate. Enabling this option requires twice as many disks. ☐ Mirror

Double Parity:
Select to enable double parity on this aggregate. Enabling this option requires an extra disk per RAID group. ☒ Double Parity

< Back Cancel Next >

Figure 5-6 Aggregate Name window

4. In the RAID Parameters window (Figure 5-7) select the number of disks that will be used on each RAID Group created for the aggregate. The recommended number of disks per RAID group is 16 disks. If you are using less than 16 disks per RAID Group, protection against disk failure is increased, but performance will decrease because there will be less disk spindles for accessing the data. If you are using more than 16 disks per RAID Group, performance will be increased (more disk spindles to access the data), but protection against disk failure will decrease. Click **Next**.



The screenshot shows a web browser window titled "http://9.11.218.237/ - itsotuc3: Aggregate Wizard - Windows Internet Explorer". The main content area is titled "Aggregate Wizard - RAID Parameters". It features a section labeled "RAID Group Size:" with a text input field containing the number "16" and a help icon. Below this, a descriptive text reads: "Enter the number of disks per RAID group on this aggregate. Disks will be organized into RAID groups of this size." At the bottom of the form, there are three buttons: "< Back", "Cancel", and "Next >".

Figure 5-7 RAID Parameters window

5. In the Disk Selection window (Figure 5-8) select the method that should be used to identify the disks used on the aggregate. The default method is Automatic so that the IBM System Storage N series storage system will automatically select the disks based on your choices for the size and number of disks from the next windows. If for any reason you need to select specific disks to compose the RAID Groups, click **Manual** and select the number and size of disks to be included on the aggregate. Click **Next**.




Figure 5-8 Disk Selection window

6. If you select the Automatic method selection for the disks, the Disk Size window will be shown, as shown in Figure 5-9. Select the disk size from the available options or select **Any Size** and click **Next**.




Figure 5-9 Disk Size window

7. Select the number of disks of the selected size to be used on the aggregate in the Number of Disks window (Figure 5-10). Click **Next**.

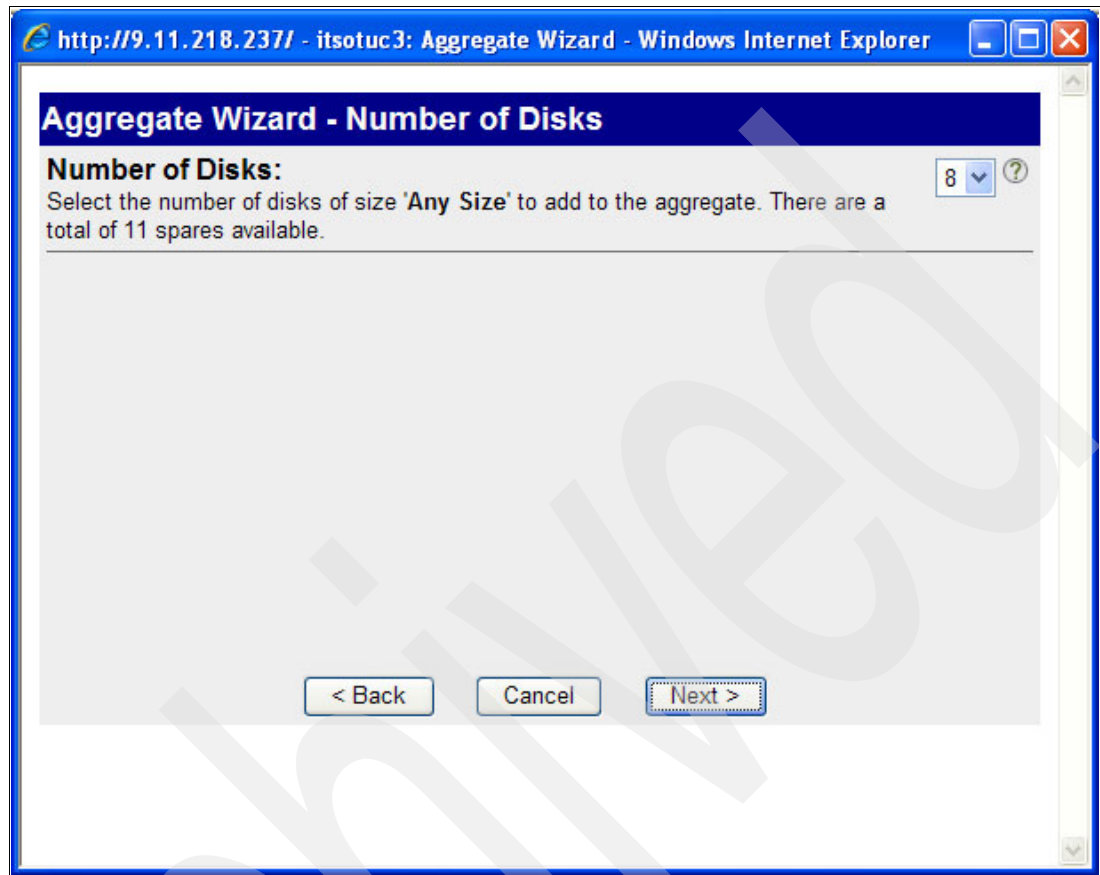


Figure 5-10 Number of Disks window

8. Review your selection in the Commit changes window (Figure 5-11) and click **Commit**.




Figure 5-11 Commit changes window

- The aggregate will be created. In the FilerView, select **Aggregates** → **Manage** and a list of the existent aggregates will be shown, along with their status, RAID level, size, available size, and other information (see Figure 5-12).

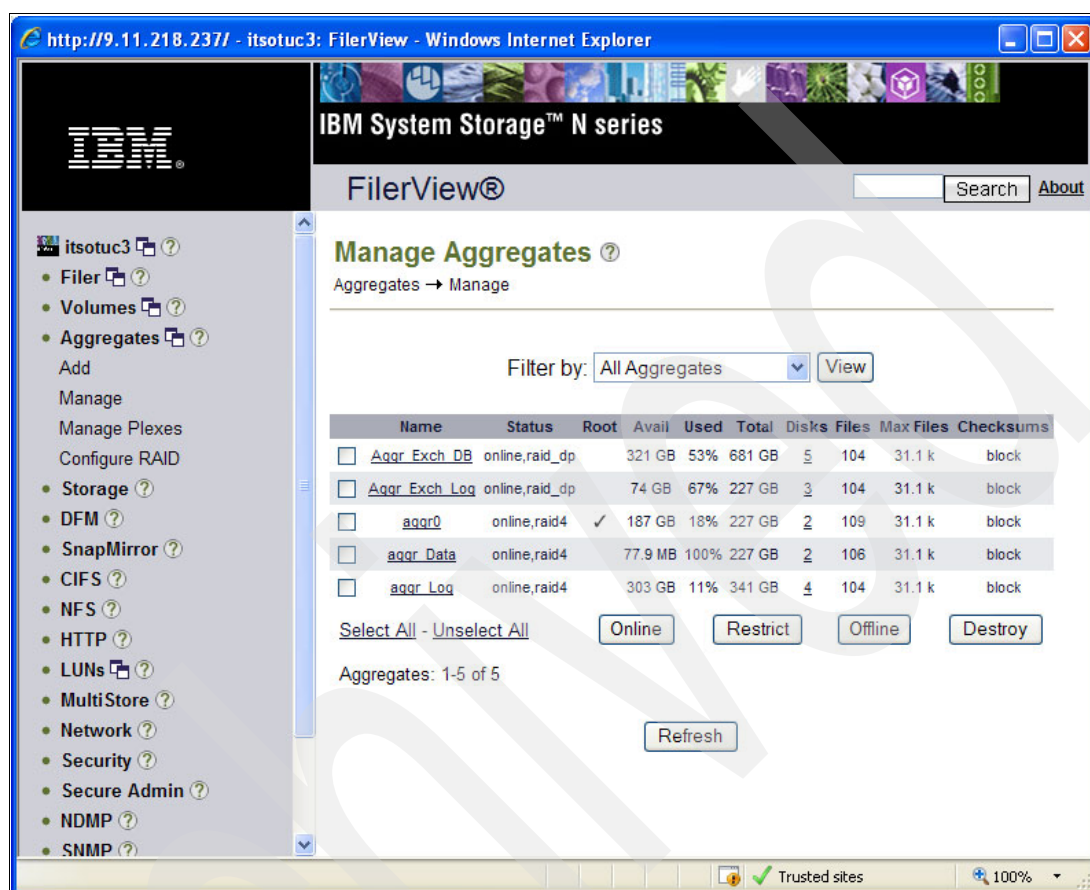


Figure 5-12 Manage Aggregates window

5.2.2 Volumes

Volumes on the IBM System Storage N series storage system can be designated as Traditional volumes or flexible volumes.

Traditional volumes are tied to the physical disks on the aggregate on which they are created. This means that the disks used on a Traditional volume cannot be used on a different volume, whether it is a Traditional volume or a flexible volume.

Traditional volumes do not allow much flexibility and the only way to increase the size of a Traditional volume is to add disk spindles to the volume array. This type of volumes does not allow downsizing.

On the other hand, flexible volumes created on aggregates can use disks from and share disks with different flexible volumes. This is due to the fact that flexible volumes are not tied to the physical disks on which they are created but to the aggregate's collection of disks.

Flexible volumes provide more management flexibility and allow for dynamic volume size expansion and shrink without impact on the host client.

There are some advantages of using flexible volumes on the Microsoft Exchange server:

- ▶ There is no waste of disk space because the flexible volumes rely on the aggregate's physical disks, instead of having a separate array of disks.
- ▶ Volume size can be dynamically increased and decreased without adding extra physical disks.
- ▶ A larger number of volumes can be created with independent Snapshot management, schedules, mirroring policies, and so on.
- ▶ All volumes can be managed independently, while taking advantage of the maximum I/O performance benefit of a much larger pool of physical disks.

The recommendation for the Microsoft Exchange server environment is to always use the flexible volumes instead of the Traditional volumes because the flexible volumes provide a better performance and manageability at the same time that they provide a better physical disk resources utilization.

In this scenario, two volumes (one on each aggregate) are necessary so that the Microsoft Exchange server's databases and transaction log files can be moved to different paths on the IBM System Storage N series storage system.

Creating volumes

Every volume on the IBM System Storage N series storage system must be created on an aggregate. Because of the recommendation for Microsoft Exchange servers, the two volumes should be created on different aggregates.

The volume name must follow these naming conventions:

- ▶ Begin with either a letter or an underscore
- ▶ Contain only letters, digits, and underscores
- ▶ Contain no more than 255 characters

These are the steps to create the volume on the aggregate:

1. Open the FilerView for the IBM System Storage N series storage system where you want to create the aggregate.

1. In the FilerView, select **Volumes** → **Add**. This will bring up the Add New Volume window, as shown in Figure 5-13. Click **Next**.

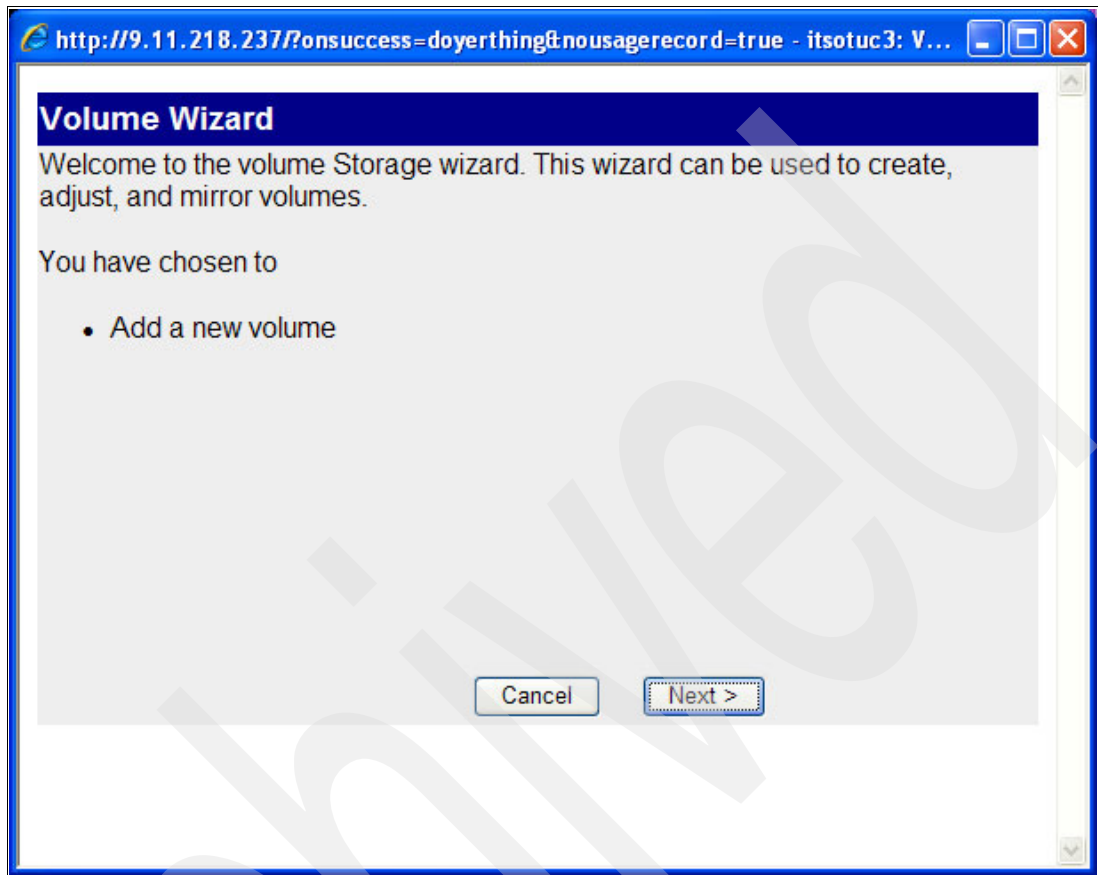


Figure 5-13 Add New Volume window

2. The Volume Type Selection window will be shown (Figure 5-14). Select **Flexible** for flexible volumes or **Traditional** for Traditional volumes. The recommended type for Microsoft Exchange Server is FlexVol. Click **Next**.


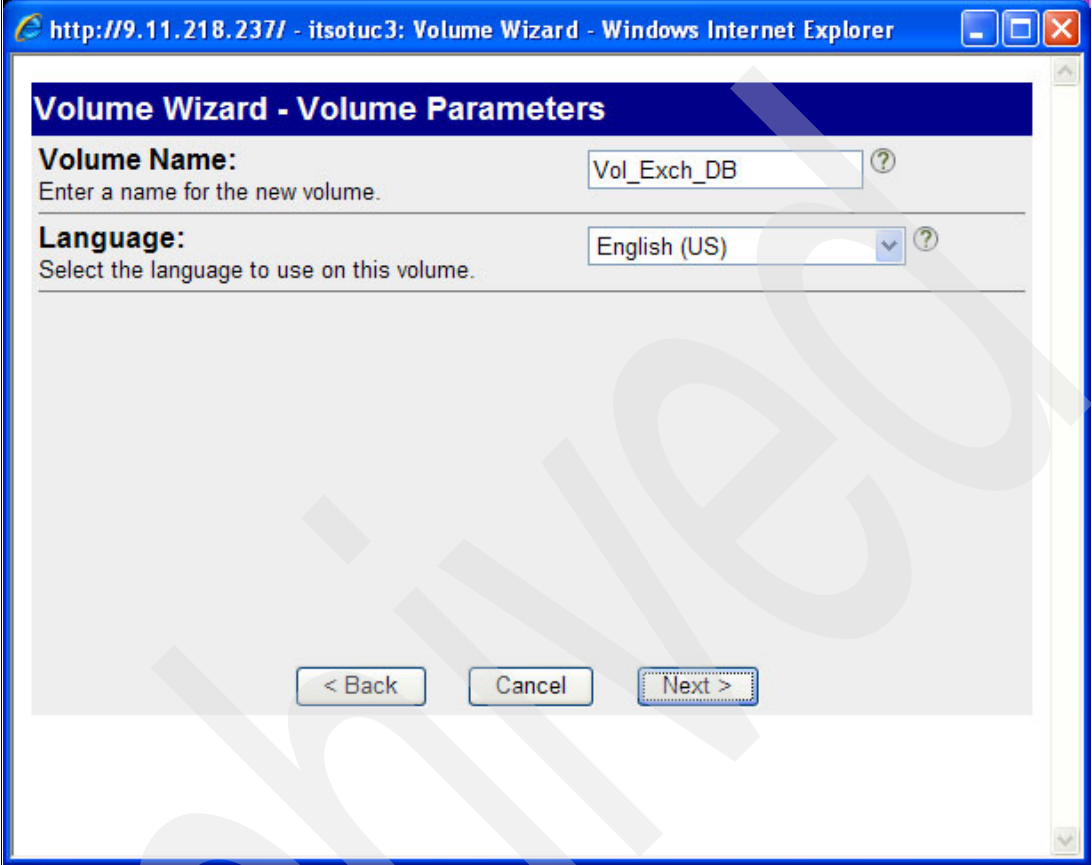


Figure 5-14 Volume Type Selection window

3. In the Volume Parameters window (Figure 5-15), type in the volume name and select the language used on the volume. By default, the root volume language will be selected. Click **Next**.



The screenshot shows a web browser window with the address bar displaying "http://9.11.218.237/ - itsotuc3: Volume Wizard - Windows Internet Explorer". The main content area is titled "Volume Wizard - Volume Parameters". It contains two input fields: "Volume Name:" with the text "Vol_Exch_DB" and a help icon, and "Language:" with a dropdown menu showing "English (US)" and a help icon. Below these fields are three buttons: "< Back", "Cancel", and "Next >".

Figure 5-15 Volume Parameters window

4. In the FlexVol Parameters window (Figure 5-16), select the aggregate on which you want to create the volume. Type in the size for the volume in KB, MB, GB, or TB. Select the type of Space Guarantee to be used. The default, recommended one is Volume. This option will pre-allocate the entire volume size on the aggregate. Other options are file space guarantee and none. Click **Next**.

http://9.11.218.237/ - itsotuc3: Volume Wizard - Windows Internet Explorer

Volume Wizard - Flexible Volume Parameters

Containing Aggregate
Select the aggregate to contain this volume. Only non-snaplock aggregates are displayed.

Aggr_Exch_DB (381 GB, raid_dp) ?

Total Volume Size:
Enter the total amount of space for this volume. The total volume size includes space for the snapshot reserve and the file system overhead in addition to the usable space.

360 GB ?

Space Guarantee
Sets the space guarantee. Volume guarantees space for the entire the volume in the containing aggregate; File guarantees space for a file at file allocation time.

volume ?

< Back Cancel Next >

Figure 5-16 FlexVol Parameters window

5. Review the selections in the Commit Changes window (Figure 5-17) and click **Commit**.

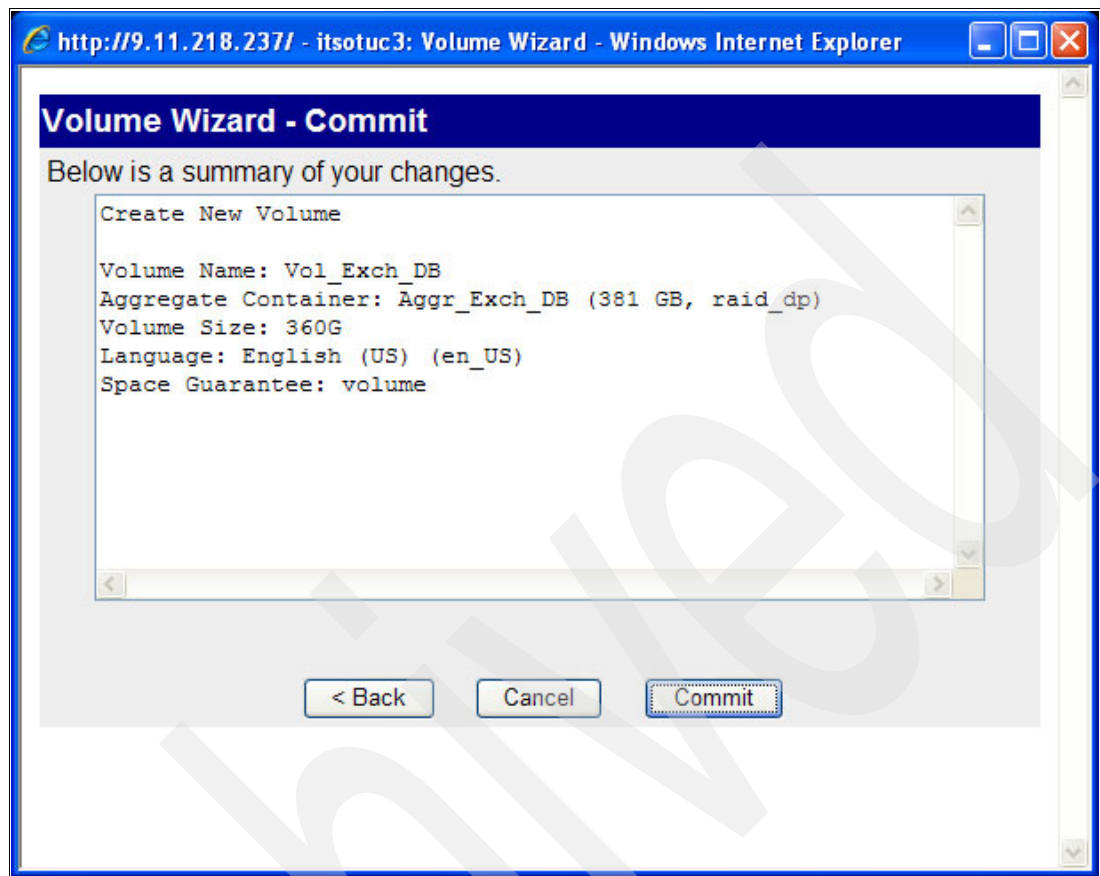


Figure 5-17 Commit Changes window

6. The volume will be created. In the FilerView, select **Volumes** → **Manage** and a list of the existing volumes will be shown, along with their status, RAID level, size, available size, and other information (see Figure 5-18 on page 143).

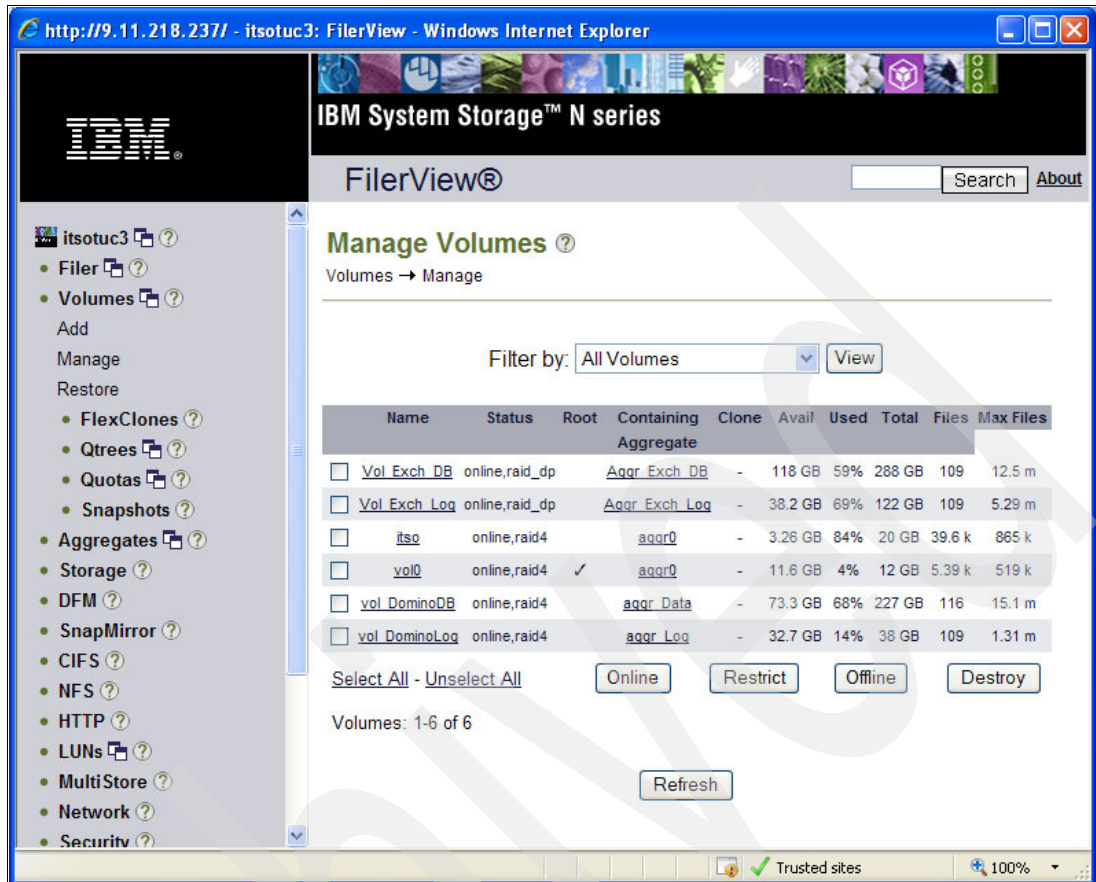


Figure 5-18 Manage Volumes window

After the volumes are created for the Microsoft Exchange server, they must be shared. This is done by using the CIFS option on the FilerView.

1. In the FilerView, select **CIFS** → **Shares** → **Add**. The Add a CIFS Share window will be shown (Figure 5-19 on page 144). Type in the following information:
 - a. Share Name: This is the name that will be used to access the volume for the LUN creation on the Microsoft Exchange server.
 - b. Mount Point: The path to connect to this volume on the N series storage system, such as /vol/Vol_Exch_DB or /vol/Vol_Exch_Log.
 - c. Share Description: General description for the Share.
 - d. Max. Users: Maximum number of concurrent users at a time on the Share.
 - e. Force Group: Not used for volumes accessed by Windows hosts.
2. Click **Add**.

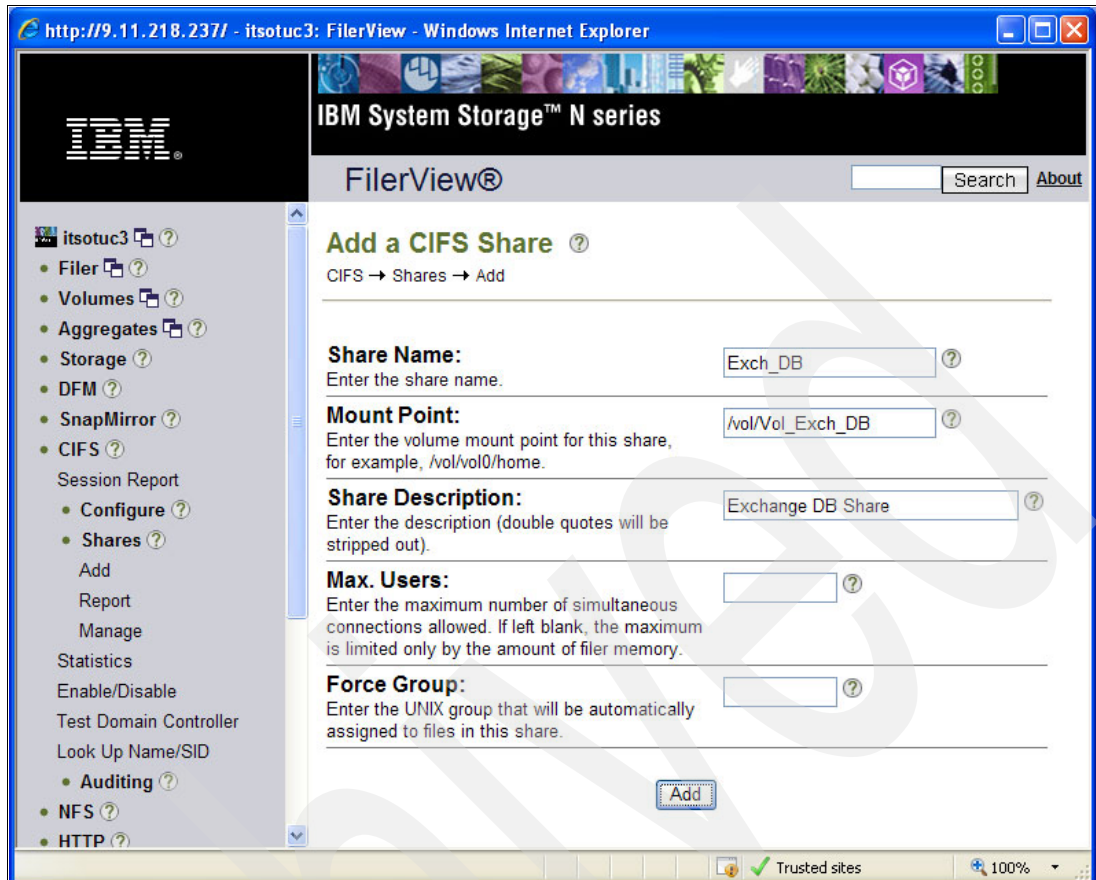


Figure 5-19 Add a CIFS Share window

5.2.3 LUNs

Logical Unit Numbers (LUNs) are the logical units of storage. They are created on the volumes and appear to host systems (in this case, the Microsoft Exchange Server) as SAN disks. The LUNs are actually the disks that will be accessed by the hosts.

The recommended way to create LUNs is using the SnapDrive utility on the Microsoft Exchange server. The steps to accomplish the LUN creation will be listed later in this book.

5.2.4 Protocols

The user data on an IBM System Storage N series storage system is accessible through one or more of the access protocols supported by Data ONTAP, including:

- ▶ Network File System (NFS)
- ▶ Common Internet File System (CIFS)
- ▶ Hyper Text Transfer Protocol (HTTP)
- ▶ Web-based Distributed Authoring and Versioning (WebDAV)
- ▶ File Transfer Protocol (FTP)
- ▶ Fibre Channel Protocol (FCP)
- ▶ Internet SCSI Protocol (iSCSI)

For a Microsoft Exchange environment, you can configure a storage system as a storage device in an iSCSI network using the SCSI protocol over TCP/IP (using the iSCSI service) and in a SAN network using the SCSI protocol over FC (using the FCP service).

Microsoft Exchange protocol recommendations

Microsoft Exchange server can use iSCSI or FCP protocols to communicate with the IBM System Storage N series storage system. The recommended protocol depends on your business needs, actual environment, and planned upgrades.

Figure 5-20 gives an overview of the iSCSI and FCP protocol stack on the client side. As seen in the figure, iSCSI needs an additional protocol layer for transporting the SCSI commands, compared to FCP. This makes for a slight performance difference between these protocols.

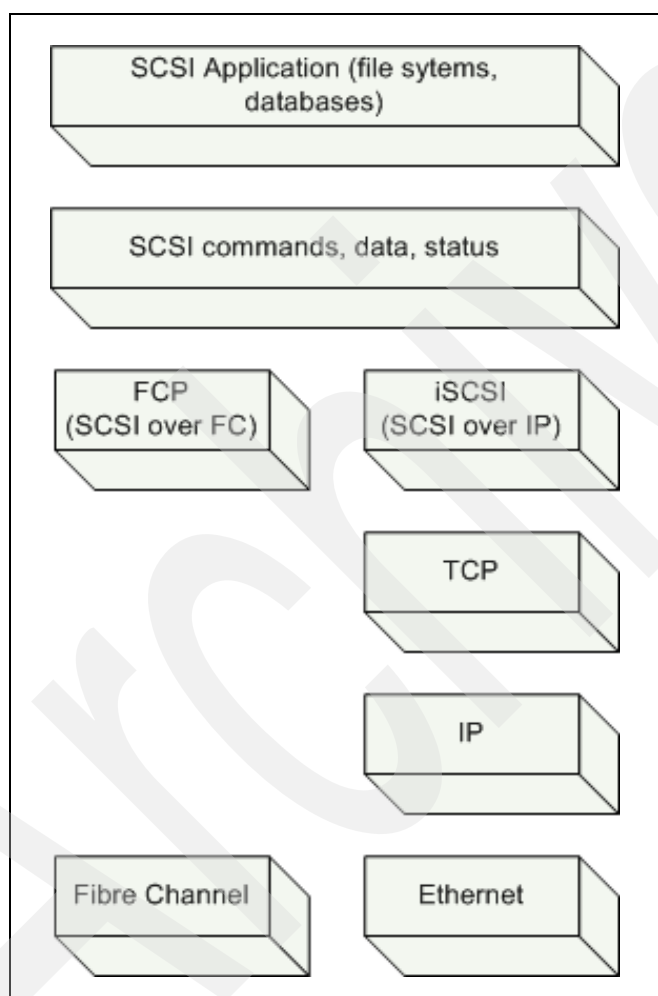


Figure 5-20 FCP and iSCSI protocol comparison, client side

The Block I/O based protocols iSCSI and FCP differs in their protocol-overhead. While FCP is designed and optimized for storage operations, iSCSI uses TCP/IP as the transport layer. If you only use iSCSI-HBA, which off loads the TCP/IP and iSCSI protocol handling, you reduce the additional CPU time needed for protocol processing.

For a production Microsoft Exchange server infrastructure, we recommend FCP for the protocol. If the environment is not able to support FCP and you do not plan to upgrade it, you might use the iSCSI protocol. In this case, use iSCSI-HBAs or at least network interfaces with

the TCP/IP offload engine. Also, you should use a physical dedicated network for your iSCSI environment.

Note: For best performance and reliability, we recommend the Fibre Channel Protocol (FCP).

5.2.5 Role-based access control

Role-based access control (RBAC) is a method for managing the set of actions that a user or administrator may perform in a computing environment.

While reserving certain functions for administrator-only access is a good start, additional problems need to be solved. Most organizations have multiple system administrators, some of which require more privileges than others. By selectively granting or revoking privileges for each user, you can customize the degree of access that an administrator has to the system.

We use the role-based access control feature for creating a separate user account that belongs to the SnapDrive feature on the Microsoft Exchange server. This account has specific rights to create, modify, and delete LUNs and Snapshots.

Users are members of groups and groups have one or more roles. Each role grants a set of capabilities. In this way Data ONTAP allows you to create flexible security policies that match your organizational needs. All configuration for role-based access controls occurs through the **useradmin** command provided by Data ONTAP (see Example 5-1).

Example 5-1 Data ONTAP useradmin command

```
itsotuc3> useradmin
Usage:
    useradmin user  add <login_name> [options]
                   modify <login_name> [options]
                   delete <login_name>
                   list [options]
    group add <group_name> [options]
          modify <group_name> [options]
          delete <group_name> [options]
          list [options]
    role  add <role_name> [options]
          modify <role_name> [options]
          delete <role_name> [options]
          list [options]
    domainuser add <user_name> [options]
              delete [options]
              list [options]
              load <filename>

For more detailed information about each subcommand, use:

    useradmin help { user | group | role | domainuser }
itsotuc3>
```

We use the **useradmin user add** or **useradmin user modify** commands for adding or modifying user accounts.

Note: For more information about accessing the storage system and role-based access control of Data ONTAP, refer to the *IBM System Storage N series Data ONTAP System Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001690&aid=1>

5.2.6 Creating a new role

A role is defined as a defined set of capabilities. Data ONTAP comes with several roles predefined, and users may create additional roles or modify the provided roles.

When creating new groups, Data ONTAP requires specification of the roles the new groups should have.

Therefore, it is better to create appropriate roles before defining groups. Follow these steps for creating a new role:

1. Log in to the N series with SSH or Telnet. SSH must be enabled prior to using it.
2. Create a new role with the **useradmin** command, as shown in Example 5-2. We use the role name **rl_snapdriveadmin** in our case.

Example 5-2 Creating a new role

```
itsotuc3*> useradmin role add rl_snapdriveadmin -c "SnapDrive Admin Role" -a  
api-*,login-http-admin  
Role <rl_snapdriveadmin> added.
```

```
itsotuc3*>
```

- The option **-c** specifies a comment.
- The option **-a** follows granted privileges. For SnapDrive, the **api-*** and **login-http-admin** privileges are needed.

Note: If you use HTTPS, grant the privilege **login-https-admin** instead of **login-http-admin**.

3. Verify the created role (see Example 5-3).

Example 5-3 Verify the created role

```
itsotuc3*> useradmin role list rl_snapdriveadmin  
Name:    rl_snapdriveadmin  
Info:  
Allowed Capabilities: api-*,login-http-admin
```

```
itsotuc3*>
```

Now it is time to create a group and map the role to it.

5.2.7 Creating a new group

A group is defined as a collection of users or domain users. Groups may be assigned one or more roles. It is important to remember that the groups defined within Data ONTAP are separate from the groups defined in other contexts, such as a Microsoft Active Directory server. This is true even if the groups within Data ONTAP have the same names as groups elsewhere within the environment.

When creating new users or domain users, Data ONTAP requires specification of group membership. Therefore, it is best to create appropriate groups before defining users or domain users. Follow these steps for creating a new group:

1. Log in to your N series with SSH or Telnet.
2. Create a new role with the **useradmin** command, as shown in Example 5-4. We use the group name `grp_snapadmins` in our case.

Example 5-4 Creating a new group

```
itsotuc3*> useradmin group add grp_snapadmins -c "This group is for SnapDrive  
usage only" -r rl_snapdriveadmin  
Group <grp_snapadmins> added.
```

```
itsotuc3*>
```

- The option `-c` specifies a comment.
- The option `-r` maps the role `snapdriveadmin` after the creation process.

3. Verify the created group (see Example 5-5).

Example 5-5 Verify the created group

```
itsotuc3*> useradmin group list grp_snapadmins  
Name: grp_snapadmins  
Info: This group is for SnapDrive usage only  
Rid: 131072  
Roles: rl_snapdriveadmin  
Allowed Capabilities: api-*,login-http-admin
```

```
itsotuc3*>
```

Tip: If you need to modify the group and role mapping later, use the **useradmin group modify -r** command.

5.2.8 Creating a new user

A user is defined as an account that is authenticated on the IBM System Storage N series storage system.

Follow these steps to create a new user:

1. Log in to your N series with SSH or Telnet.
2. Create a new user with the **useradmin** command, as shown in Example 5-6 on page 149. We use the user name `windows` in our case.

Note: The password must have at least eight characters.

Example 5-6 Creating the user account

```
itsotuc3*> useradmin user add windows -c "MS Windows SnapDrive user" -n "Microsoft Windows" -g grp_snapadmins
New password:
Retype new password:
User <windows> added.
```

```
itsotuc3*>
```

- The option -c specifies a comment.
- The option -n specifies the full name of the user.
- The option -g specifies the group name to which the user belongs.

3. Verify the created user (see Example 5-7).

Example 5-7 Verify the created user

```
itsotuc3*> useradmin user list notes
Name: windows
Info: MS Windows SnapDrive user
Rid: 131077
Groups: grp_snapadmins
Full Name: Microsoft Windows
Allowed Capabilities: api-*,login-http-admin
Password min/max age in days: 0/4294967295
Status: enabled
```

```
itsotuc3*>
```

Note: The command **useradmin user** will give you a list of possible options for user management on the N series Data ONTAP.

5.3 Zoning

A Fibre Channel zone consists of a group of Fibre Channel ports or nodes that can communicate with each other. It can be thought of as a logical fabric subset. Two Fibre Channel nodes can communicate one another only when they are configured on the same zone. A node can be configured on multiple zones. For example, it is a typical scenario for a storage node (such as the IBM System Storage N series target FC port) to be configured on multiple zones. There are generally considered to be two methods of zoning: hard and soft zoning.

5.3.1 Hard zoning

In port based zoning, often referred to as “port zoning”, the zone is defined by specifying the fabric unique N_port IDs of the ports to be included. In other words, the switch and switch port (such as switch3/Port 4) are used to define the zone members.

Security

Hard zoning typically is considered to offer improved security since it is not possible to breach the zoning using WorldWide Port Name (WWPN) spoofing. However, if someone has physical access to the switch, simply replacing a cable can allow access when hard zoning is used.

Manageability

In many environments, hard zoning is easier to create and manage because only the switch or switch domain and port number need to be used on the configuration instead of the 16-digit WWPNs.

5.3.2 Soft zoning

In WorldWide Port Name (WWPN) based zoning, the zone is defined by configuring the WWPN of the members to be included within the zone. Depending on the switch vendor, either the World Wide Node Name or World Wide Port Names can be used, although we recommend World Wide Port Name zoning.

Flexibility

Since access is not determined by which port on the switch the device is physically plugged into, it is possible with soft zoning to simply move a cable from one port to another without needing to reconfigure the zoning. This can be useful in troubleshooting situations.

5.3.3 Zoning architecture

Because of the nature of zoning, it is often easiest to understand by illustration. Below are several diagrams showing several of the benefits of zoning.

Figure 5-21 on page 151 shows an example where each host is shown in a separate zone. This should be the standard zoning configuration for a simple environment. The zones are overlapping because the storage ports are included on each zone configuration to allow each server to access both storages. Each host can access all of the Fibre Channel target ports but cannot interact with the other host ports.




Figure 5-21 Hosts in individual zones

By using hard or port zoning, the zoning configuration can be done in advance, even if all the servers are not present. Assuming that the storage is connected to ports 1 through 4, each zone can be defined to contain a single port for the server and ports 1 through 4. For example, one zone would consist of ports 1, 2, 3, 4, and 5 while the next zone would consist of ports 1, 2, 3, 4, and 6 and so forth.

Figure 5-22 shows only a single fabric, but a supported configuration would have two fabrics. The second fabric would have the exact same zone structure.




Figure 5-22 Single fabric, no multipathing zoning

In Figure 5-22 on page 151, Host1 and Host2 do not have multipathing software and therefore have to be zoned in a way that only one path to each LUN is available to them. Therefore, the zone containing these hosts contains only one of the two storage ports. Even though the host has only one HBA, both storage ports were included in the zone, and the LUNs would be visible through two different paths: one going from the host FC port to storage port 0, and the other going from the host FC port to storage port 1.

Because there is only one single fabric (which may consist of one or more switches), this configuration is not considered fully redundant because a switch or fabric failure would lead to an outage. IBM only considers dual, physically independent fabrics to be fully redundant. However, as shown, Host3 and Host4 have multipathing software and, to protect against a possible storage controller failure, are zoned so that a path to the LUN(s) is available through each of the storage controllers.

Figure 5-23 shows a configuration where Host1 will be accessing LUNs from N series 1 and Host2 will be accessing LUNs from appliance 2. Both storage appliances are configured for high availability with two storage controllers. This is a fully redundant configuration and both fabrics are shown in this example.




Figure 5-23 Multiple storage system zoning

This zoning separates the hosts to eliminate initiator (host HBA) cross talk and prevents Host1 from accessing appliance 2, which increases security while improving reliability, manageability, and problem solving.

5.3.4 Zoning recommendations

For Fibre Channel SAN zoning, we recommend the following:

Zoning

Anytime four or more hosts are connected to a SAN, zoning should be implemented.

Hard or WWPN zoning

World Wide Node Name zoning is possible with some switch vendors, but for IBM System Storage N series FC Target connections, WWPN zoning is recommended.

Because of the various trade-offs, there is no specific recommendation between hard versus soft zoning.

Zone size

For IBM System Storage N series storage systems, we generally recommend keeping the zone size as small as possible while still maintaining manageability. It is not a problem to have many multiple overlapping zones to help keep the individual zones smaller. Ideally, a zone will be defined for each host or host cluster.

5.3.5 Paths

There is one or more paths between the Microsoft Exchange server and the IBM System Storage N series storage system. We recommend at least two paths between the server and storage so that a Multi Path IO (MPIO) setup configuration can be done.

The basic premise behind MPIO is to provide a redundant path to a given LUN. If it is implemented correctly, it can mitigate the issues that are associated with a single point of failure. Typically, these types of failures will be switch or HBA failures. A cable failure is also possible, but those are very rare.

In Figure 5-24, the network and associated hardware is properly configured for MPIO, but that alone is not enough. The Initiators and SnapDrive must be configured as well. If we just tell SnapDrive to create a session and a LUN on the filer, there will be only one logical path to the LUNs despite the network configuration. The MPIO driver should also be installed on the Microsoft Exchange server.




Figure 5-24 MPIO setup

5.4 Snapshots

A Snapshot copy is a space-efficient, point-in-time image of the data in a volume or an aggregate. Snapshot copies are used for such purposes as backup and error recovery.

Data ONTAP automatically creates and deletes Snapshot copies of data in volumes to support commands related to Snapshot technology. These Snapshots are based on the IBM System Storage N series storage system unique Write Anywhere File Layout (WAFL) (Figure 5-25 on page 155).

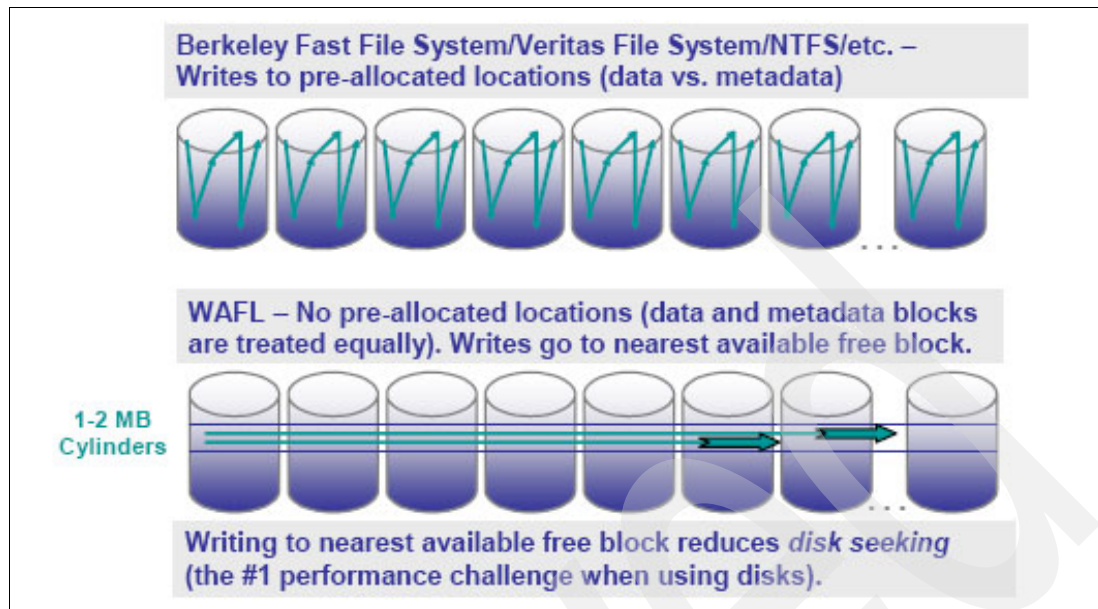


Figure 5-25 WAFL file system

Data ONTAP also automatically creates Snapshot copies of aggregates to support commands related to the SyncMirror software, which provides RAID-level mirroring. For example, Data ONTAP uses aggregate Snapshot copies when the data in two plexes of a mirrored aggregate need to be resynchronized. In this book, we are usually talking about volume related Snapshots.

A SnapShot copy is a frozen, read-only image of a traditional volume, a flexible volume, or an aggregate that reflects the state of the file system at the time the SnapShot copy was created. SnapShot copies are your first line of defense for backing up and restoring data.

Some facts about Snapshot copies:

- ▶ Data ONTAP maintains a configurable Snapshot schedule that creates and deletes Snapshot copies automatically for each volume.
- ▶ For taking Snapshot copies of LUNs, use SnapDrive. It will handle the flushing of the host operating system buffers.
- ▶ You can store up to 255 Snapshot copies at one time on each volume.
- ▶ You can specify the percentage of disk space that Snapshot copies can occupy. The default setting is 20% of the total (both used and unused) space on the disk.

Note: For more information about Snapshots, see the *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

5.5 Requirements for Microsoft Exchange with N series

A Microsoft Exchange server running on Microsoft Windows 2003 Server requires the following N series features, licences, and software packages. Make sure you have them available:

- ▶ FCP connection

In this IBM Redbooks publication, we used FCP to map the database and transactional log files LUNs to the Microsoft Exchange server. SnapDrive for Microsoft Windows was used to manage the connection to the LUNs on the IBM System Storage N series storage system. SnapDrive is a licensed feature and can be obtained by contacting IBM support.

- ▶ License

For the following features, you need a license:

- FCP protocol
- iSCSI protocol
- SnapDrive for Microsoft Windows
- Data ONTAP DSM for Windows MPIO (for Multipathing support) SnapManager for Microsoft Exchange servers
- Optionally, you might want to use and license the SnapMirror or FlexClone features (not covered in this book)

- ▶ Check the compatibility matrix for your environment to make sure you are using a supported configuration. For more information, see:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

5.6 SnapDrive

The N series SnapDrive feature provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file, down through the volume manager, to the storage-system-side logical unit numbers (LUNs) providing the actual repository. In addition, it simplifies the backup of data and helps you decrease the recovery time.

SnapDrive provides a layer of abstraction between an application running on the host operating system and the underlying IBM System Storage N series storage systems (see Figure 5-26 on page 157). Applications that are running on a server with SnapDrive use virtual disks (or LUNs) on N series storage systems, as though they were locally connected drives or mount points. This allows applications that require locally attached storage, such as IBM Lotus Domino and Microsoft Exchange to benefit from the N series technologies, including Snapshot, flexible volumes, FlexClone, and space management technologies.

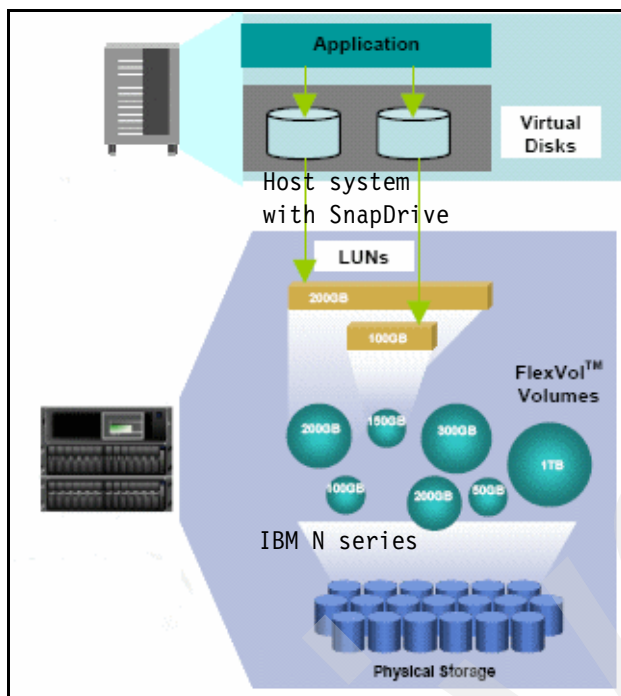


Figure 5-26 Example of a typical SnapDrive deployment

SnapDrive includes all the necessary drivers and software to manage interfaces, protocols, storage, and Snapshot copies. Snapshot copies are nondisruptive to applications and functions on execution. Snapshot backups can also be mirrored across LAN or WAN links for centralized archiving and disaster recovery.

SnapDrive MMC is totally integrated on the Computer Management MMC, as shown in Figure 5-27.

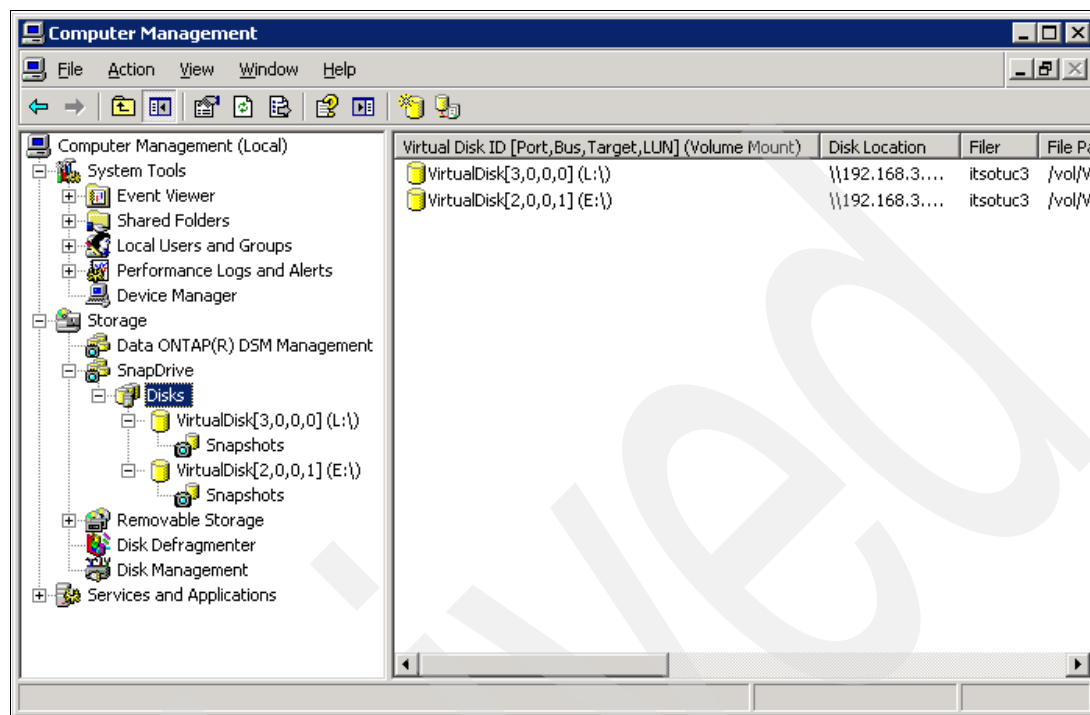


Figure 5-27 SnapDrive MMC integration with Computer Management MMC

5.6.1 Benefits of SnapDrive

Most of today's enterprises use business-critical applications, such as IBM Lotus Domino and Microsoft Exchange, and their storage management team faces a number of challenges. They must:

- ▶ Support new business initiatives with a minimal increase in the operating budget.
- ▶ Protect data from corruption, disaster, and attacks.
- ▶ Back up data without any performance degradation, quickly and consistently without any errors.

SnapDrive addresses these problems by providing simplified and intuitive storage management and data protection from a host/server perspective. The following list highlights some of the important benefits of SnapDrive for Windows:

- ▶ Allows hosts and applications administrators to quickly create virtual disks with a dynamic pool of storage that can be reallocated, scaled, and enlarged in real time, even while the systems are accessing data.
- ▶ Dynamic on-the-fly file system expansion, making new disks usable within seconds.
- ▶ Snapshot copies provide rapid backup and recovery capabilities with minimal resource and capacity requirements.
- ▶ Supports multipath technology for high performance.
- ▶ Provides connectivity to existing Snapshot copies from original host or different host.
- ▶ SnapDrive is independent of underlying storage access media and protocol. It supports FCP and iSCSI as the transport protocols.

- ▶ Robust and easy-to-use data and storage management feature and software.

5.6.2 SnapDrive requirements

IBM System Storage N series SnapDrive is a licensed feature and is available by contacting the IBM support.

These are some general requirements for SnapDrive:

- ▶ Host operating system and appropriate patches
- ▶ Name resolution infrastructure
- ▶ IP connectivity between the host and the storage system
- ▶ Storage system licenses
- ▶ FCP connectivity between the host and the storage system
- ▶ For security reasons, we recommend a separate user account on the IBM System Storage N series storage system. See 5.2.5, “Role-based access control” on page 146 for more details)

For more information about SnapDrive for Windows V4.2.1, refer to:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001650&aid=1>

5.7 SnapMirror

The Data ONTAP SnapMirror feature allows IBM System Storage N series Snapshot images to mirror SnapShot images either asynchronously or synchronously over the network for backup or disaster recovery purposes.

Figure 5-28 shows two steps for creating a SnapMirror and using the SnapMirror feature. At first, a baseline SnapShot image replication is done. Several transport media could be used, such as LAN, FCP, or Tape.

Step two shows the asynchronous, synchronous, or semi-synchronous mirroring process from source to target.

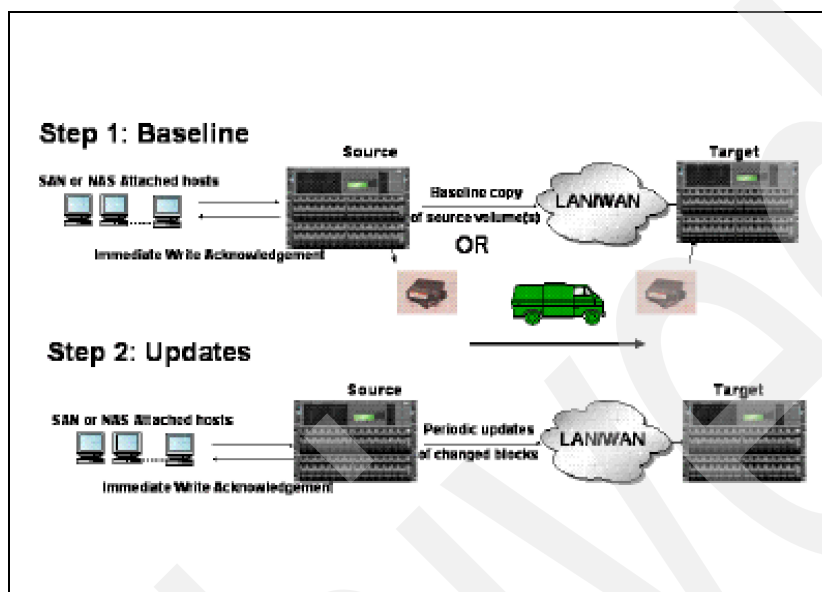


Figure 5-28 SnapMirror

► Asynchronous mode

In asynchronous mode, SnapMirror performs incremental, block-based replication as frequently as once per minute. The performance impact on the source IBM System Storage N series storage system is minimal as long as the system is configured with sufficient CPU and disk I/O resources. The administrator can decide on the replication interval depending on the business needs.

► Synchronous mode

The Synchronous SnapMirror mode is a SnapMirror feature that replicates data from a source volume to a partner destination volume at or near the same time that it is written to the source volume, rather than according to a predetermined schedule. The chance of data loss is minimized in this mode.

Note: Use the same type and amount of physical hard disks and aggregate configuration on the SnapMirror source and destination to reduce the performance impact.

► Semi-Synchronous mode

To improve the performance of the synchronous SnapMirror feature, it can be configured to lag behind the source volume by a user-defined number of write operations or milliseconds, which is mentioned as semi-synchronous SnapMirror mode.

This mode is like asynchronous mode in that the application does not need to wait for the secondary storage to acknowledge the write before continuing with the transaction. This mode is like synchronous mode in that updates from the primary storage to the secondary storage occur right away, rather than waiting for scheduled transfers.

5.7.1 Benefits of SnapMirror

The N series SnapMirror feature replicates data for more than backup and disaster recovery purposes. The following list shows some use cases:

- ▶ Using the SnapMirror target for local read access at remote sites.
Fast access to corporate data on remote sites and off loading the source side.
- ▶ Off load tape backup CPU cycles to mirror.
- ▶ Isolate testing from production volume.
ERP testing, Offline Reporting, and Disaster Recovery exercises.
- ▶ Cascading Mirrors.
Replicated mirrors on a larger scale.
- ▶ Disaster recovery.
Replication to “hot site” for mirror failover and eventual recovery.
- ▶ They can use the N series SnapMirror feature in combination with FlexClone volume to perform migration faster and more efficiently.
 - For corporations with a warm backup site, or a need to off load backups from production servers.
 - For generating queries and reports on near-production data.

5.7.2 Effects on sizing

This N series SnapMirror feature mirrors your selected data, so you need the same size on both the source site and the target location of your mirror. A best practice is to use the same aggregate and volume configuration on both sites when synchronous mirroring is used.

5.7.3 SnapMirror requirements

The following prerequisites must be met before you can run SnapMirror:

- ▶ You must purchase and enable the SnapMirror license.
If the SnapMirror source and destination are on different systems, you must purchase a license and enter the SnapMirror license code on each system.
- ▶ For SnapMirror volume replication, you must create a restricted volume to be used as the destination volume.
- ▶ For SnapMirror volume replication, the destination volume must run under a version of Data ONTAP that is the same or later version than that of the SnapMirror source volume.
If you configure volume SnapMirror to support replication for the purpose of disaster recovery, both the source and destination systems should run the same version of ONTAP software.

Note: If you upgrade your systems to a later version of Data ONTAP, upgrade the systems of SnapMirror destination volumes before you upgrade the systems of the SnapMirror source volumes.

- ▶ The name and IP address of the source system must be in the /etc/hosts file of the destination system or must be resolvable by way of Domain Naming System (DNS) or Network Information Service (NIS) (also called Yellow Pages (YP)).

5.8 SnapVault

SnapVault is a low overhead, disk-based online backup of a heterogeneous storage system for fast and simple restores. SnapVault is a separately licensed feature in Data ONTAP that provides disk-based data protection for storage systems.

The SnapVault feature runs on the IBM System Storage N series platform. The SnapVault feature replicates selected Snapshots from multiple storage systems or back ups data from an open system platform to a common Snapshot on the SnapVault server. Figure 5-29 shows a common SnapVault implementation with an open system platform and IBM System Storage N series storage server on the primary site. The changed data is transferred to the IBM System Storage N series secondary storage system, from where regular Snapshots are saved on tape.

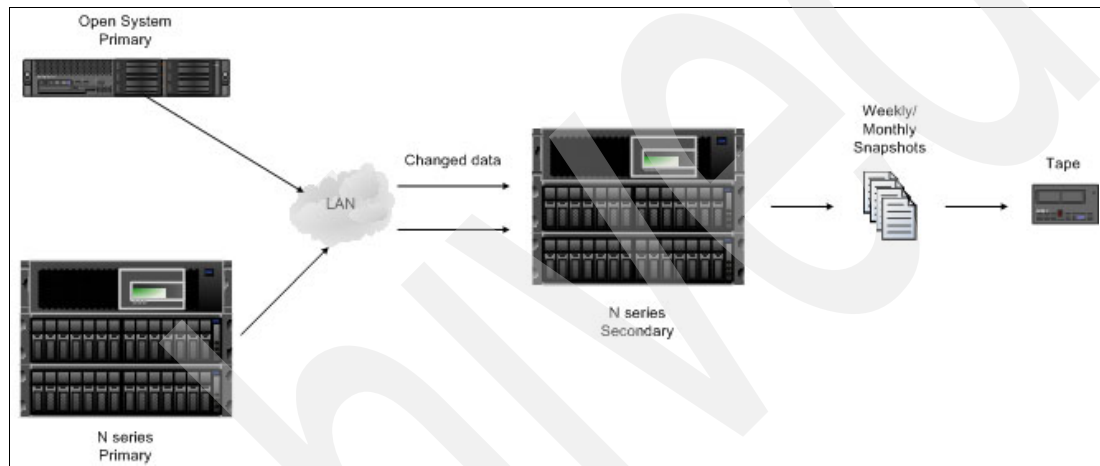


Figure 5-29 Common SnapVault implementation

Note: Because Data ONTAP SnapDrive uses qtrees as backup storage, it is not supported for volumes containing Data ONTAP LUNs, as we use it in this book for Lotus Domino and Microsoft Exchange. For more information about SnapVault, see the *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

5.9 FlexClone

This feature allows you to instantly create clones of a FlexVol flexible volume in a few seconds without interrupting the parent flexible volume. A FlexClone volume is a writable point-in-time image of a flexible volume of another FlexClone volume (see Figure 5-30 on page 163). They use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between parent and clone. In addition to these benefits, clone volumes have the same high performance as other kinds of volumes.




Figure 5-30 FlexClone

A FlexClone LUN clone shares space with the LUN in the backing SnapShot copy. The clone does not require additional disk space until changes are made to it. You cannot delete the backing SnapShot copy until you split the clone from it. When you split the clone from the backing SnapShot copy, you copy the data from the SnapShot copy to the clone. After the splitting operation, both the backing SnapShot copy and the clone occupy their own space.

Note: For more information about using volume cloning with LUNs, see the *IBM System Storage N series Block Access Management Guide for FCP and iSCSI*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

5.9.1 FlexClone use cases and benefits

The main difference between a Snapshot and cloning is that a cloned volume or LUN is writable. This is an important and powerful feature for Microsoft Exchange administrators. In the case of a database or software upgrade, release, or configuration change, a clone of the production system could be made. This clone could be used for testing with your real data, without worrying about damaging the production system.

Volume cloning with the FlexClone feature provides similar results to volume copying, but cloning offers some important advantages over volume copying:

- ▶ Volume cloning is instantaneous, whereas volume copying can be time consuming.
- ▶ If the original and cloned volumes share a large amount of identical data, considerable space is saved because the shared data is not duplicated between the volume and the clone.

Cloning is not recommended as a primary backup method, because all data still resides on the same volume and aggregate. You might use LUN cloning for the following reasons:

- ▶ Application Testing: Make the necessary changes to the infrastructure without worrying about crashing the production system. Avoid making untested changes on the system under tight maintenance window deadlines.

- ▶ You need to make a copy of your data available to additional users without giving them access to the production data.
- ▶ You want to create a clone of the Microsoft Exchange databases for manipulation and projection operations, while preserving the original data in unaltered form.
- ▶ Data Mining: Data mining operations and software can be implemented more flexibly because both reads and writes are allowed.
- ▶ Parallel Processing: Multiple FlexClone volumes of a single milestone/production data set can be used by parallel processing applications across multiple servers to get results more quickly.
- ▶ System Deployment: Maintain a template environment and use FlexClone volumes to build and deploy either identical or variant environments.
- ▶ Microsoft Exchange Operations: Maintain multiple copies of production systems: live, development, test, reporting, and so on. Refresh working FlexClone volumes regularly to work on data as close to live production systems as practical.

Note: A FlexClone volume can be created from a Snapshot copy in a SnapMirror destination, but a FlexClone volume cannot be the destination of a SnapMirror relationship.

Figure 5-31 on page 165 shows a common situation for FlexClone deployment. The IT staff needs to make substantive changes to a production environment. The cost and risk of a mistake are too high to do it on the production volume. Ideally, there would be an instant writable copy of the production system available at minimal cost in terms of storage and service interruptions.

By using FlexClone volumes, the Microsoft Exchange administrator gets just that: an instant point-in-time copy of the production data that is created transparently by SnapShot and uses only enough space to hold the desired changes. The administrator can then try out their upgrades using the FlexClone volumes.

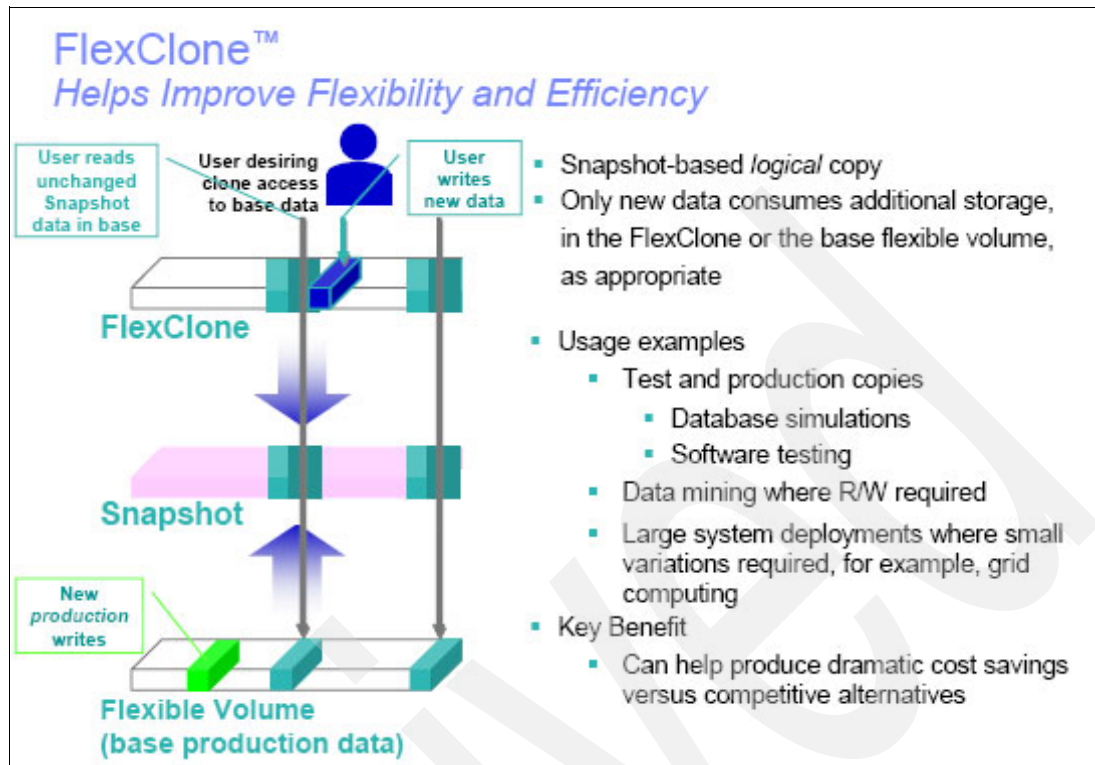


Figure 5-31 FlexClone basics

At every point that makes solid progress, the administrator clones the working FlexClone volume to lock in the successes. At any point where the administrator get stuck, they just destroy the working clone and go back to the point of their last success. When everything is finally working just the way it should be, the administrator can either split off the clone to replace the current production volumes or modify the successful upgrade process to use on the production system during the next maintenance window.

The FlexClone feature allows the administrator to make the necessary changes to the infrastructure without worrying about crashing the production systems or making untested changes on the system under tight maintenance window deadlines. The results are less risk, less stress, and higher levels of service for the IT customers.

5.9.2 FlexClone requirements

The FlexClone feature needs to be licensed. Additionally, you need to follow these points to use FlexClone:

- ▶ You must install the license for the FlexClone feature before you can create FlexClone volumes.
- ▶ FlexClone volumes and their parent volumes share the same disk space (aggregate) for any data common to the clone and parent. This means that creating a FlexClone volume is instantaneous and requires no additional disk space (until changes are made to the clone or parent).
- ▶ While a FlexClone volume exists, some operations on its parent are not allowed. For more information, refer to *IBM System Storage N series Data ONTAP 7.2 Storage Management Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001595&aid=1>

- Only flexible volumes can be cloned. To create a copy of a traditional volume, you must use the **vol copy** command, which creates a distinct copy with its own storage.

5.9.3 Effects on sizing

When a FlexClone volume is created, it shares all of its data with its parent volume. So even though its logical size is the same as its parent's size, depending on how much data it contains, it can use very little free space. As data is written to either the parent or the FlexClone volume, that data is no longer shared between the parent and FlexClone volumes, and the FlexClone volume starts to require more space from its containing aggregate, depending on the size of the changed data.

If you want to split the clone from its parent, it removes any space optimizations that are currently employed by the FlexClone volume. After the split, both the FlexClone volume and the parent volume require the full space allocation determined by their space guarantees.

Note: Because the clone-splitting operation is a copy operation that might take considerable time to carry out, Data ONTAP also provides commands to stop or check the status of a clone-splitting operation.

5.10 SnapManager for Microsoft Exchange

N series SnapManager for Microsoft Exchange is a software integrated with SnapDrive that offers a comprehensive data management solution for Microsoft Exchange.

SnapManager for Microsoft Exchange dramatically reduces the time it takes to back up and restore Exchange data by using SnapShot technology. Nearly instantaneous Snapshot backups are verified for data integrity using Microsoft tools after each backup.

Restoring entire Microsoft Exchange databases with SnapManager can be done in minutes. Restorations that used to take days can be accomplished in a matter of minutes and with complete confidence.

Using SnapManager for Exchange, server scalability is no longer limited by the time it takes to back up and restore data. SnapManager includes an intuitive graphical user interface with task wizards that help simplify administration.

SnapManager for Microsoft Exchange is a licensed feature and can be obtained by contacting IBM support.

If you are installing the SnapManager Version 4.0, Microsoft Powershell is required to support it in Microsoft Exchange 2003 and Microsoft Exchange 2007.

Note: SnapManager is the recommended software for online backups in a Microsoft Exchange environment installed on an IBM System Storage N series storage system. For external backups, use IBM Tivoli Storage Manager (ITSM) in the environment.



Part 3

Installing Lotus Domino Server and Microsoft Exchange on the IBM System Storage N series storage system

In this part, we discuss topics related to installing Lotus Domino Server and Microsoft Exchange on the IBM System Storage N series. It is not meant as a comprehensive guide to installing Lotus Domino Server and Microsoft Exchange but rather as a tool to pointing out differences or considerations when the IBM System Storage N series is involved.

Archived



Installing Microsoft Exchange on IBM System Storage N series

This chapter provides the steps to access the IBM System Storage N series storage system from the Microsoft Exchange server as well as how to install and configure the necessary features and softwares on the server.

6.1 Accessing the IBM System Storage N series storage system

After planning and implementing the necessary features on the IBM System Storage N series storage system as well as creating the aggregates and volumes, it is time to access the IBM System Storage N series storage system and creating the LUNs from the host client, which is the Microsoft Exchange Server.

The first step is to install and configure the protocol that will be used for this access. At this point you have already planned for the protocol you are going to use and have all the infrastructure in place for the use of either iSCSI or FCP. There is some configuration that need to be done on the IBM System Storage N series storage system and on the Microsoft Exchange Server as well.

Note: Because of its improved performance and reliability, the recommended protocol is FCP.

6.1.1 Fibre Channel Protocol (FCP)

Many companies already have a Fibre Channel infrastructure already in place. This eases the installation and configuration of FCP on the Microsoft Exchange Server. Also, the companies already have the knowledge for troubleshooting issues related to FCP.

The recommended configuration when using FCP is to have multiple paths configured. In this manner, should any of the Host Bus Adapters (HBAs) or any of the Fibre Channel cords or any of the FC Switches fail, you still have connectivity between the host and the IBM System Storage N series storage system, as shown on Figure 6-1 on page 171.




Figure 6-1 Multipathing configuration for Microsoft Exchange Server using FCP

Note: To enable the FCP protocol and the FCP adapter on the IBM System Storage N series storage system, you will need an FCP license to be installed on the IBM System Storage N series storage system. In the FilerView, select **Filer** → **Manage Licenses** and scroll down the right pane until you see the FCP license. Add the FCP License and click **Apply**.

Assuming that all the infrastructure is already in place and working, the first feature to install and enable on the server is the Data ONTAP DSM for Windows MPIO software. This software is part of the IBM System Storage N series solution and requires a license during the installation. DSM for Windows MPIO will provide the driver for SnapDrive multipathing capabilities either for high availability or load balancing on FCP infrastructures.

After installing Data ONTAP DSM for Windows MPIO, SnapDrive should be installed so that the LUNs can be created.

Note: Despite the fact that the LUNs can be created from the IBM System Storage N series storage system, the recommended procedure is to create the LUNs from the Microsoft Exchange client using SnapDrive.

Installing SnapDrive for Windows

Important: Before installing SnapDrive, ensure that you have the latest patches installed on your Microsoft Windows Server®.

These are the steps to install SnapDrive on the Microsoft Exchange Server:

1. In the SnapDrive installation welcome window (Figure 6-2), click **Next**.

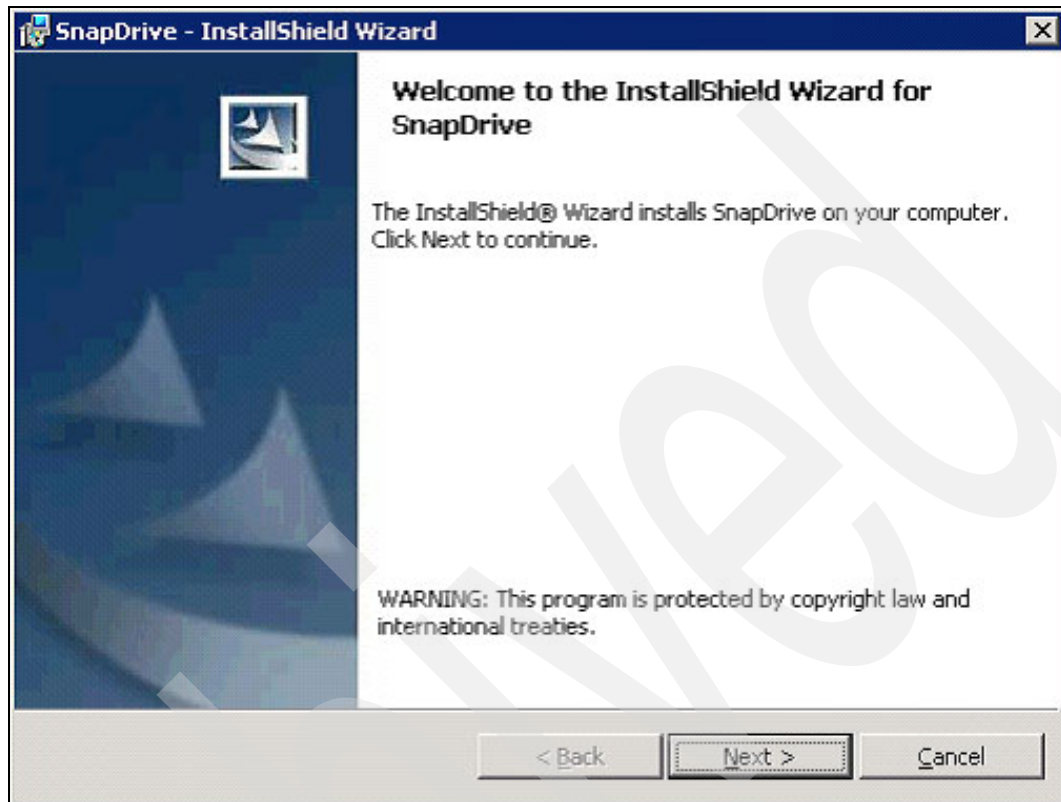


Figure 6-2 SnapDrive installation welcome window

2. In the License Agreement window (Figure 6-3), accept the terms on the License Agreement and click **Next**.

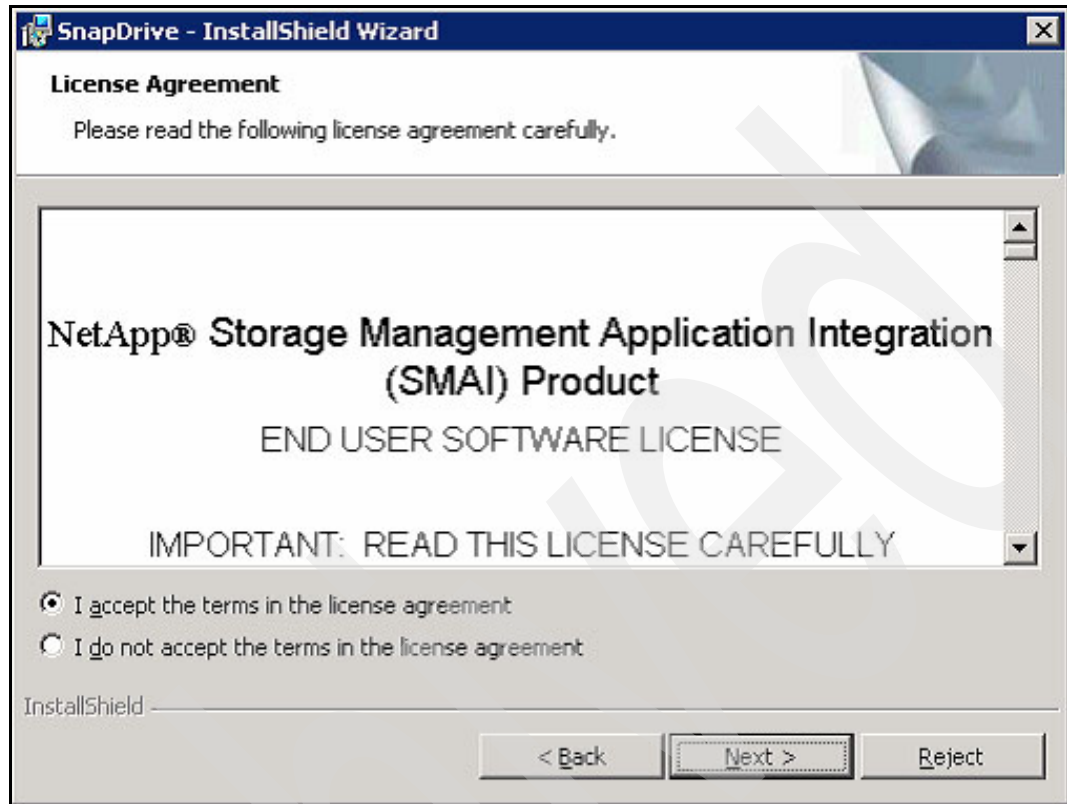


Figure 6-3 License Agreement

3. In the License key window (Figure 6-4), type in the License key for SnapDrive and click **Next**.

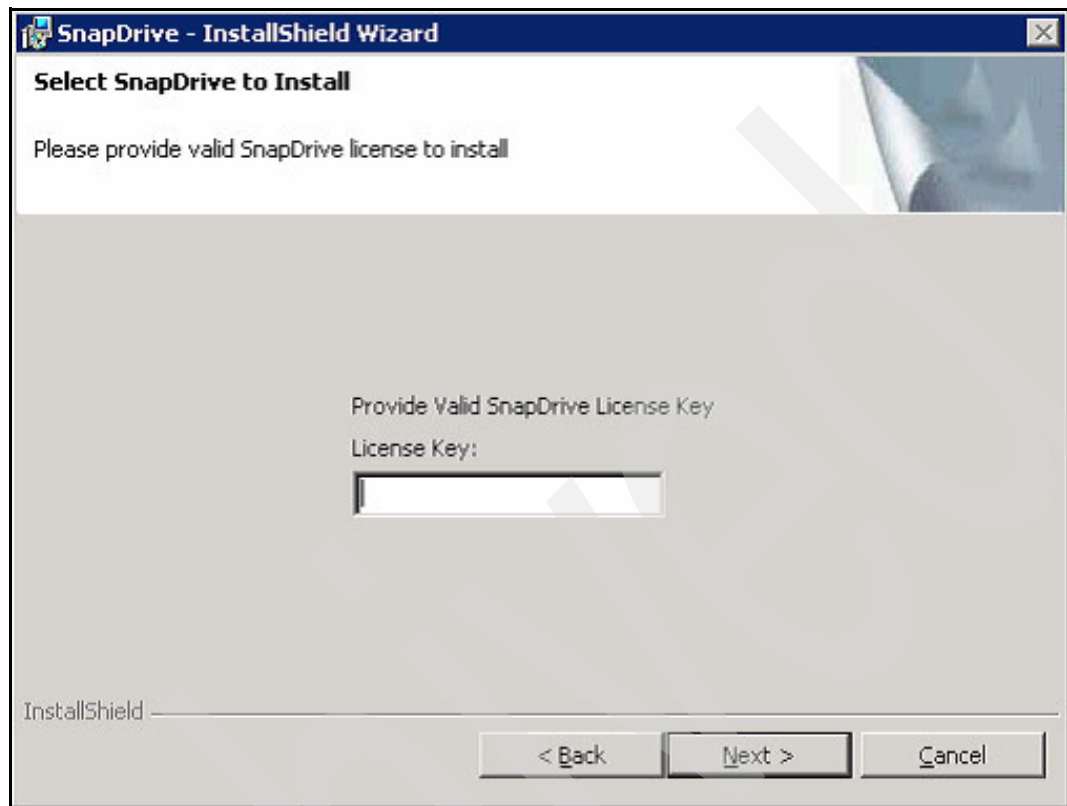
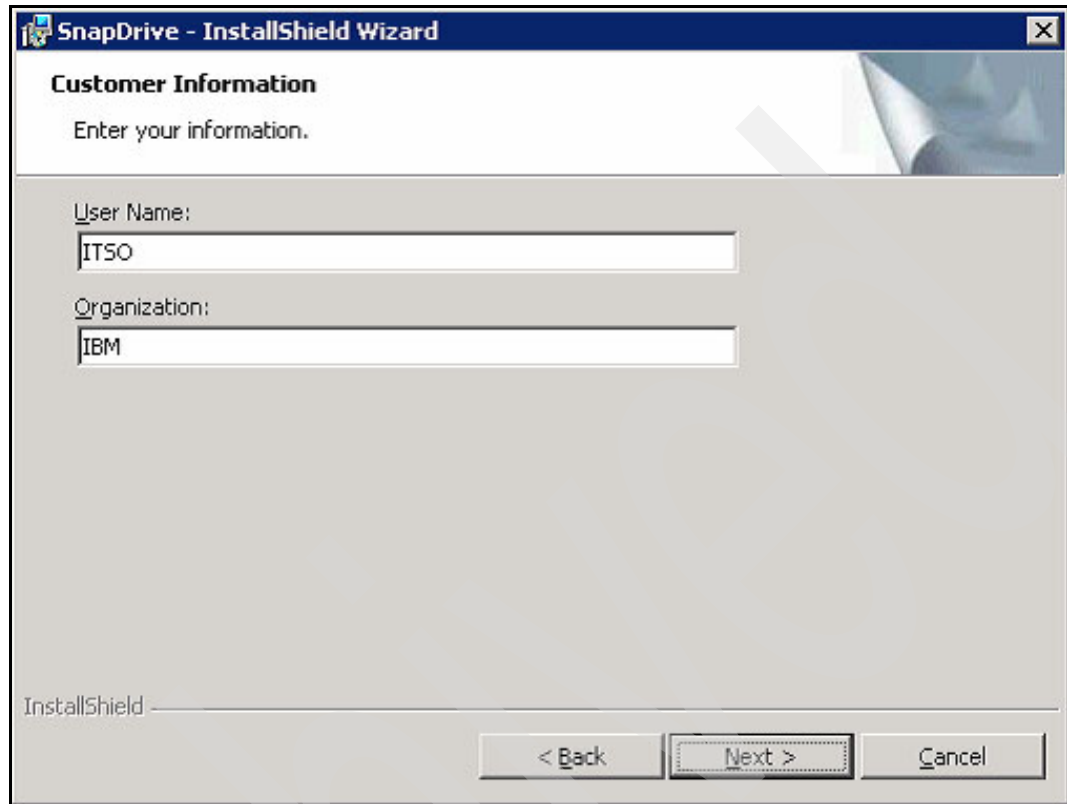


Figure 6-4 License Key

4. In the Customer information window (Figure 6-5), type in the User Name and Organization information and click **Next**.



The image shows a Windows-style dialog box titled "SnapDrive - InstallShield Wizard". The main heading is "Customer Information" with the instruction "Enter your information." below it. There are two text input fields: "User Name:" containing "ITSO" and "Organization:" containing "IBM". At the bottom, there is a progress bar labeled "InstallShield" and three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

Figure 6-5 Customer information

5. In the Destination Folder window (Figure 6-6), confirm or change the destination folder for the installation files and click **Next**.




Figure 6-6 Destination Folder

6. In the SnapDrive Service Credentials window (Figure 6-7 on page 177), type in the Account and Password for the user account to be used to start SnapDrive service and click **Next**.

Note: The SnapDrive service user account should be a member of the Active Directory's Domain Admins group and be a member of the filer's local administrators group.

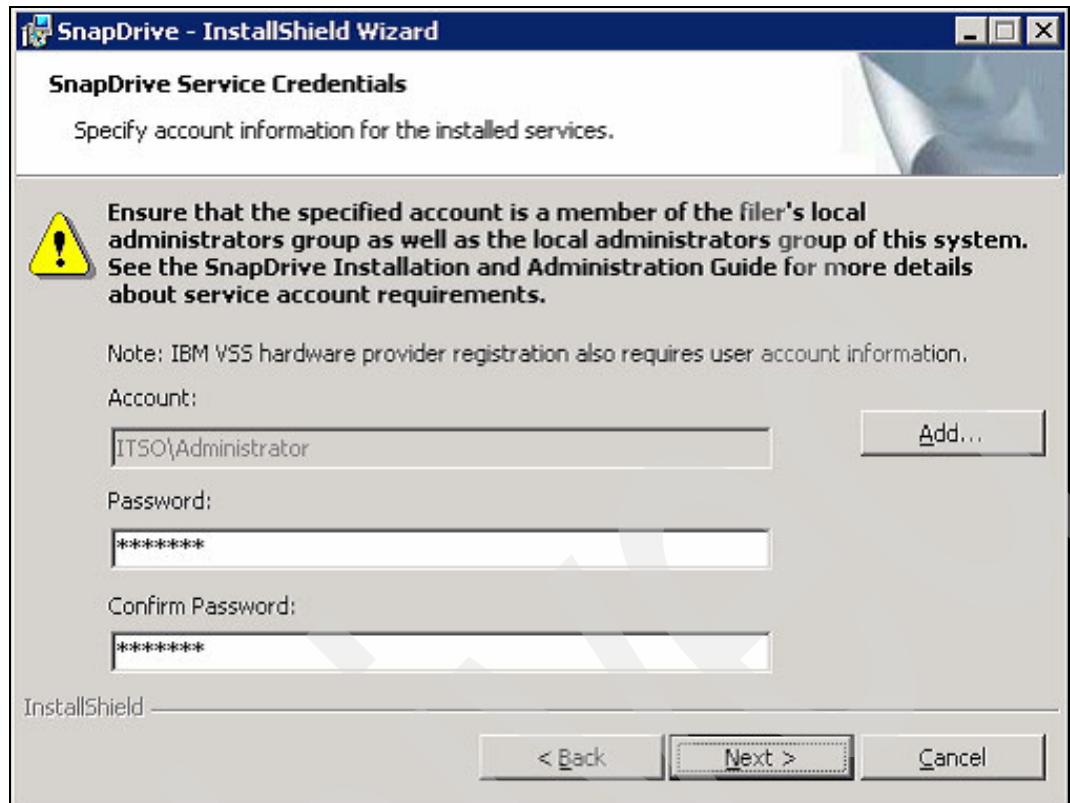


Figure 6-7 SnapDrive Service Credentials

7. Click **Finish**.

After the installation is done, no restart is needed in order to get SnapDrive for Windows working.

Creating disks from SnapDrive

Now that the Data ONTAP DSM for Windows MPIO and SnapDrive are installed, the disk drives can be added using the SnapDrive software.

In order to create the disks from SnapDrive, we assume that:

- ▶ An aggregate has been created on the filer.
- ▶ A volume has been created on the aggregate.
- ▶ A CIFS share has been created mapping the path to the volume.
- ▶ The Fibre Channel infrastructure is in place and working.

These are the steps to create the LUN from the SnapDrive:

1. Access the Computer Management MMC.
2. Expand **Storage** and then expand **SnapDrive**.

3. Right-click **Disks** and click **Create Disk**, as shown in Figure 6-8.

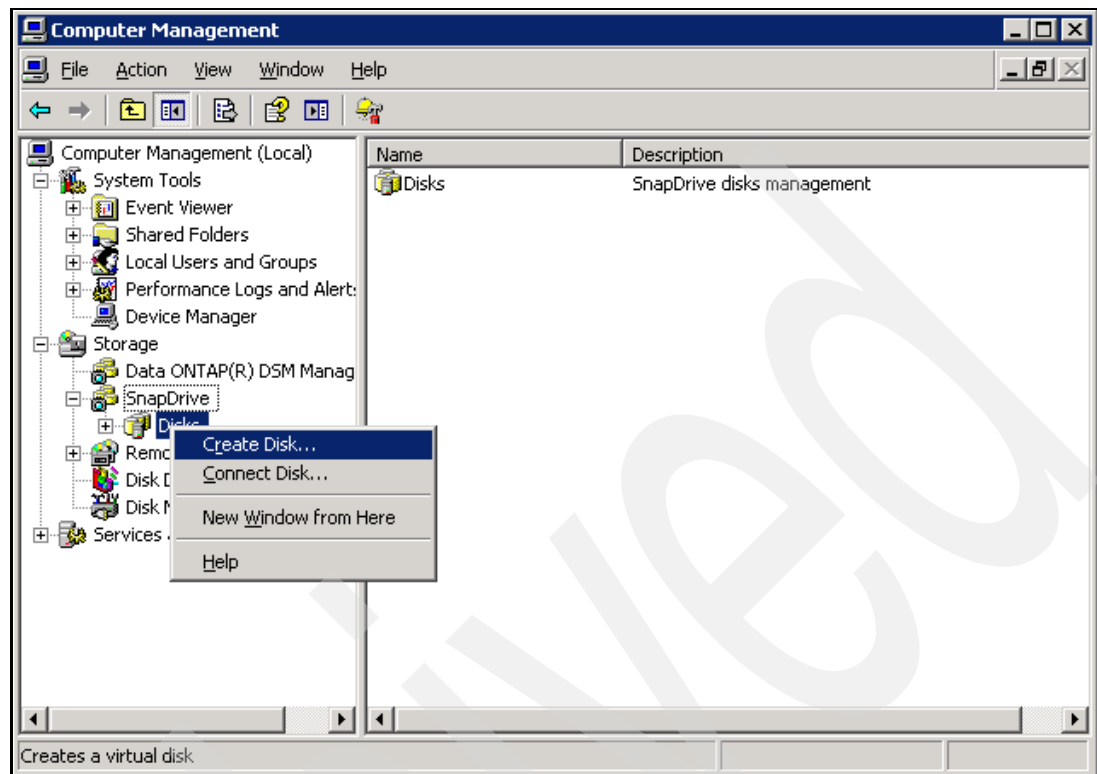


Figure 6-8 Disk creation using SnapDrive

4. In the Create Disk welcome window (Figure 6-9), click **Next**.

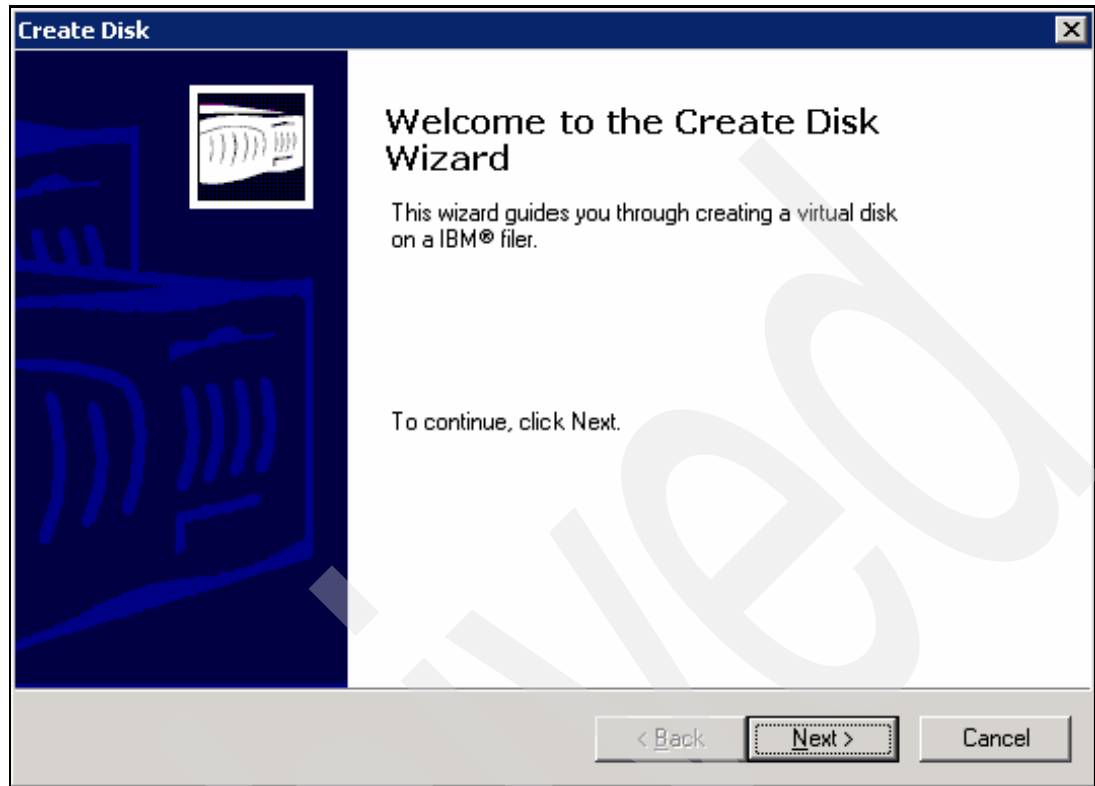
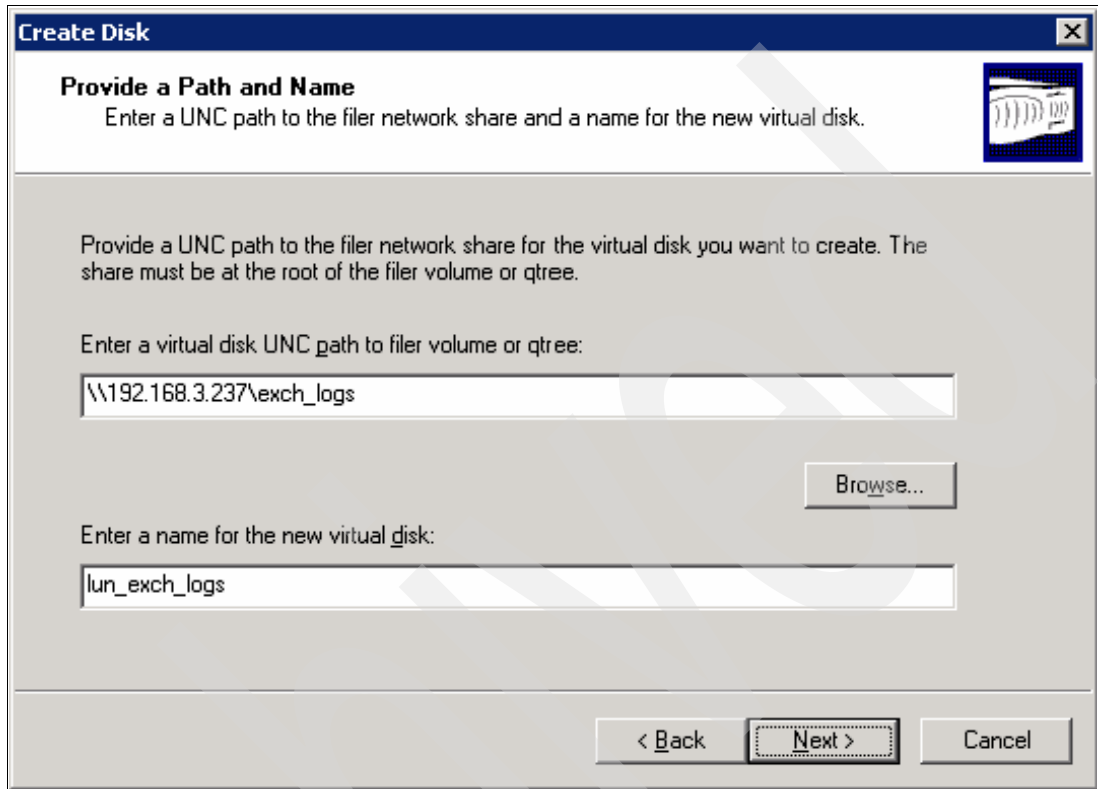


Figure 6-9 Create disk welcome window

5. On the Provide a Path and Name window (Figure 6-10), enter the information needed. For the UNC path, use \\Filer IP\CIFS Share on the filer. Remember that this CIFS share points to the volume you created. For the name for the new virtual disk, type in the name you want to assign to the LUN that will be created. Click **Next**.



The screenshot shows a Windows-style dialog box titled "Create Disk". The main heading is "Provide a Path and Name", followed by the instruction "Enter a UNC path to the filer network share and a name for the new virtual disk." Below this, a paragraph explains: "Provide a UNC path to the filer network share for the virtual disk you want to create. The share must be at the root of the filer volume or qtree." There are two input fields. The first is labeled "Enter a virtual disk UNC path to filer volume or qtree:" and contains the text "\\192.168.3.237\exch_logs". To the right of this field is a "Browse..." button. The second input field is labeled "Enter a name for the new virtual disk:" and contains the text "lun_exch_logs". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a dashed border.

Figure 6-10 Provide path and name window

6. In the Select a virtual disk window (Figure 6-11), select **Dedicated** if this disk will be accessed by only one server. Select **Shared** if this disk will be accessed by a Cluster Service. Click **Next**.




Figure 6-11 Select a virtual disk type window

7. In the Select Virtual Disk Properties window (Figure 6-12), select if you want to assign a drive letter (and which drive letter) for the disk being created or if you want to assign a Volume Mount Point. The next option will impact the size that will be available for the LUN creation. You need to select if you will reserve space for at least one snapshot of this LUN on the volume. Then, enter the size you want for this LUN. Click **Next**.

Create Disk

Select Virtual Disk Properties
Provide the drive letter and the size of the virtual disk to create.

Provide a Drive Letter or Volume Mount Point

☒ Assign a Drive Letter: L

☐ Use Volume Mount Point:

Do you want to limit the maximum disk size to accommodate at least one snapshot on the volume?

☐ Yes ☒ No

Enter a virtual disk size that is equal to or less than the maximum size, but greater than or equal to the minimum size.

Maximum Virtual Disk Size:	44 GB
Minimum Virtual Disk Size:	32 MB

Enter or Select Virtual Disk Size: 42 GB

< Back Next > Cancel

Figure 6-12 Select Virtual Disk Properties window

8. In the Select Initiators window (Figure 6-13), select the initiators from the Available Initiators column on the left and click the arrow to move them to the Selected Initiators on the right. Because we are using FCP, the initiators will be listed as the World Wide Port Number (WWPN) of the HBAs on the system. Click **Next**.

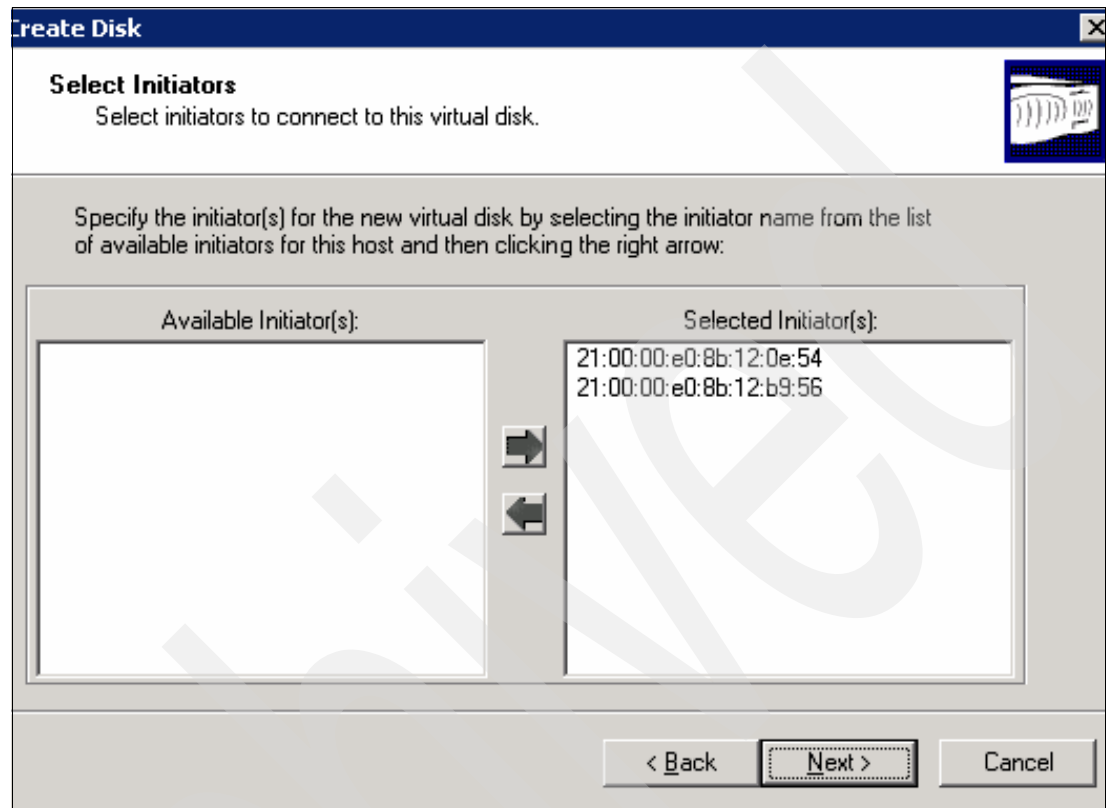


Figure 6-13 Select Initiators window

9. In the Summary window (Figure 6-14), verify if all the information is correct and click **Finish**. This will start the LUN creation process on the filer.




Figure 6-14 Summary window

10. SnapDrive will format the drive and a Disks window will appear (Figure 6-15). Click **OK**.




Figure 6-15 Disks window

11. Now the disk is available for the Microsoft Exchange Server using the drive letter you selected or using the Volume Mount Point you created.

After the LUN creation, when logging in to the Filer, you will notice that an Initiator Group has been created for each HBA interface on the server. In the FilerView, select **LUNs** → **Initiator Groups** → **Manage** for a list of the Initiator Groups on the Filer. A group named `viaRPC.ServerWWPN.ServerName` will be created for each of the HBAs interfaces. You can keep the configuration as is or you can create a new Initiator Group, add the server's WWPNs as members, and map it to the LUN. These are the steps to create the Initiator Group:

1. In the FilerView, select **LUNs** → **Initiator Groups** → **Add**.

2. As shown in Figure 6-16, enter the name for the group on the Group Name field. Select the type of the initiator as FCP. Select the operating system as Windows. In the list of initiators, add the WWPNs of the HBAs on the Exchange Server. Click **Add**.

The screenshot displays the 'Add Initiator Group' window in the IBM System Storage N series FilerView. The left sidebar shows a navigation tree with options like File, Volumes, Aggregates, Storage, DFM, SnapMirror, CIFS, NFS, HTTP, LUNs, Wizard, Enable/Disable, Manage, Add, Show Statistics, and Initiator Groups. The main area is titled 'Add Initiator Group' and includes a breadcrumb 'LUNs → Initiator Groups → Add'. Below this is a '[Manage Initiator Groups]' button. The 'Group Name' field is set to 'Exchange Server'. The 'Type' dropdown is set to 'FCP'. The 'Operating System' dropdown is set to 'Windows'. The 'Initiators' list contains two WWPNs: '21:00:00:e0:8b:12:0e:54' and '21:00:00:e0:8b:12:b9:56'. An 'Add' button is at the bottom right.

Figure 6-16 Add Initiator Group window

- After the Initiator Group has been created, select **LUNs** → **Manage** and click the LUN you have created. This will show the LUN management window (Figure 6-17). Click **Map LUN**.

IBM System Storage™ N series

FilerView® Search About

Modify LUN ?

LUNs → Manage → Modify

[Manage LUNs]	[Map LUN]
[Online]	[Offline]
[Delete]	

Path: /vol/Vol_Exch_Log/lun_exch_logs ?
The full path of the LUN, for example /vol/luns/lunOne. The LUN must be created in the root directory of a volume or a qtree.

Status: online ?
Status of the LUN.

LUN Protocol Type: Windows ?
Select the multiprotocol type for the LUN.

Description: ?
An optional description of the LUN.

Size: 45099210240 ?
The size of the LUN. (Readonly field). The current exact size is 45099210240 bytes.

Units: Bytes ?
A multiplier for the LUN size. (Readonly field).

Space Reserved: ☒ Space Reserved ?

Figure 6-17 LUN Management window

- Click the LUN Map window (Figure 6-18 on page 187) and click **Add Groups to Map**.

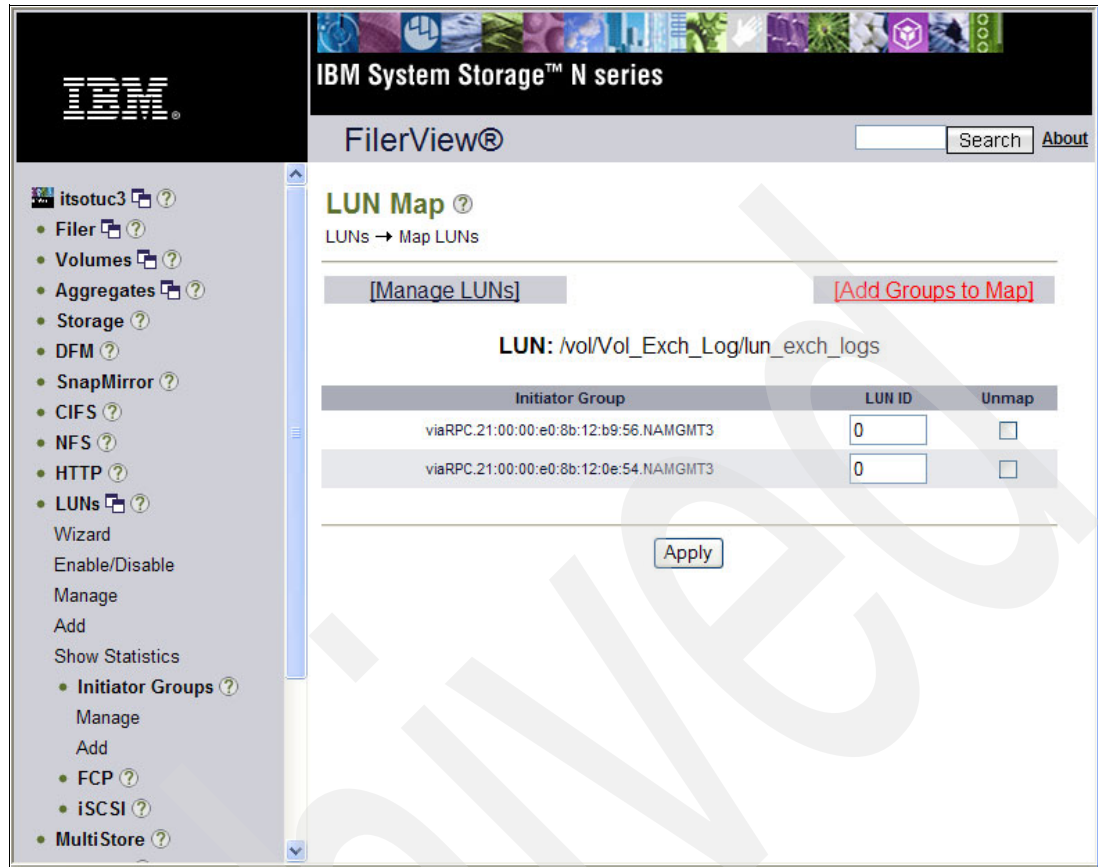


Figure 6-18 LUN Map window

5. Select the Initiator Group you have just created and click **Add**.
6. In the LUN Map window, remove the old mapping by selecting the **Unmap** check box and assign the same LUN ID to the new group. Click **Apply**.
7. Verify if the LUNs are accessible by the Microsoft Exchange Server.

6.1.2 Internet SCSI Protocol (iSCSI)

For companies that do not have an FCP infrastructure in place or for those that want to access the storage using their Ethernet infrastructure and knowledge, iSCSI can be used as the access protocol for communication between the Microsoft Exchange server and IBM System Storage N series storage system.

iSCSI Initiator adapters off load the largest amount of processing from the server's CPU.

In case there are no iSCSI adapters on your planned environment, the iSCSI Initiator software can be used to provide the same connectivity to the IBM System Storage N series storage system. The use of multipaths is also recommended when using a hardware or software based iSCSI solution, as shown in Figure 6-19.




Figure 6-19 Multipathing configuration for Microsoft Exchange Server using iSCSI

In the example, there are two interfaces on the server (either iSCSI hardware based or Gigabit Ethernet cards) that connect to two different LAN switches. It is important, for performance and reliability reasons, that these LAN segments and switches are other than the public ones. The IBM System Storage N series storage system will have two of their adapters also connected to both switches.

Note: To enable the iSCSI protocol and the iSCSI adapter on the IBM System Storage N series storage system, you will need an iSCSI license to be installed on the IBM System Storage N series storage system. In the FilerView, select **Filer** → **Manage Licenses** and scroll down the right pane until you see the iSCSI license. Add the iSCSI License and click **Apply**.

Assuming that the infrastructure is already in place and working, Microsoft iSCSI Software Initiator has to be installed on the server. And after installing and configuring it, SnapDrive should be installed and configured as well, so that the LUNs can be created.

Note: Despite the fact that the LUNs can be created from the IBM System Storage N series storage system, the recommended procedure is to create the LUNs from the Microsoft Exchange client using SnapDrive.

Installing Microsoft iSCSI Software Initiator

Microsoft iSCSI Software Initiator is the software installed on the server that allows SCSI communication over TCP/IP. This software is required if you are using the SCSI protocol to communicate with the IBM System Storage N series storage system, but do not have the hardware based iSCSI adapters.

Microsoft iSCSI Software Initiator will create additional layers in the network so that a layer for the iSCSI protocol and for SCSI drivers are present. In this manner, the regular Network Interface Card (NIC) could be used to communicate with the IBM System Storage N series storage system.

As a best practice, always use the latest versions of the software and drives on your environment, unless there are compatibility issues. For information about Microsoft iSCSI Software Initiator and the latest versions, visit the following Web site:

<http://www.microsoft.com/windowsserver2003/technologies/storage/iscsi/msfiSCSI.msp>
x

The following steps outline the Microsoft iSCSI Software Initiator software installation:

1. In the Welcome window (Figure 6-20), click **Next**.



Figure 6-20 Welcome window

2. In the Installation Options window (Figure 6-21), select the following options and click **Next**:
 - a. Initiator Service
 - b. Software Initiator
 - c. Microsoft MPIO Multipathing Support for iSCSI




Figure 6-21 Installation options window

3. In the License Agreement window (Figure 6-22), agree to the terms and click **Next**.

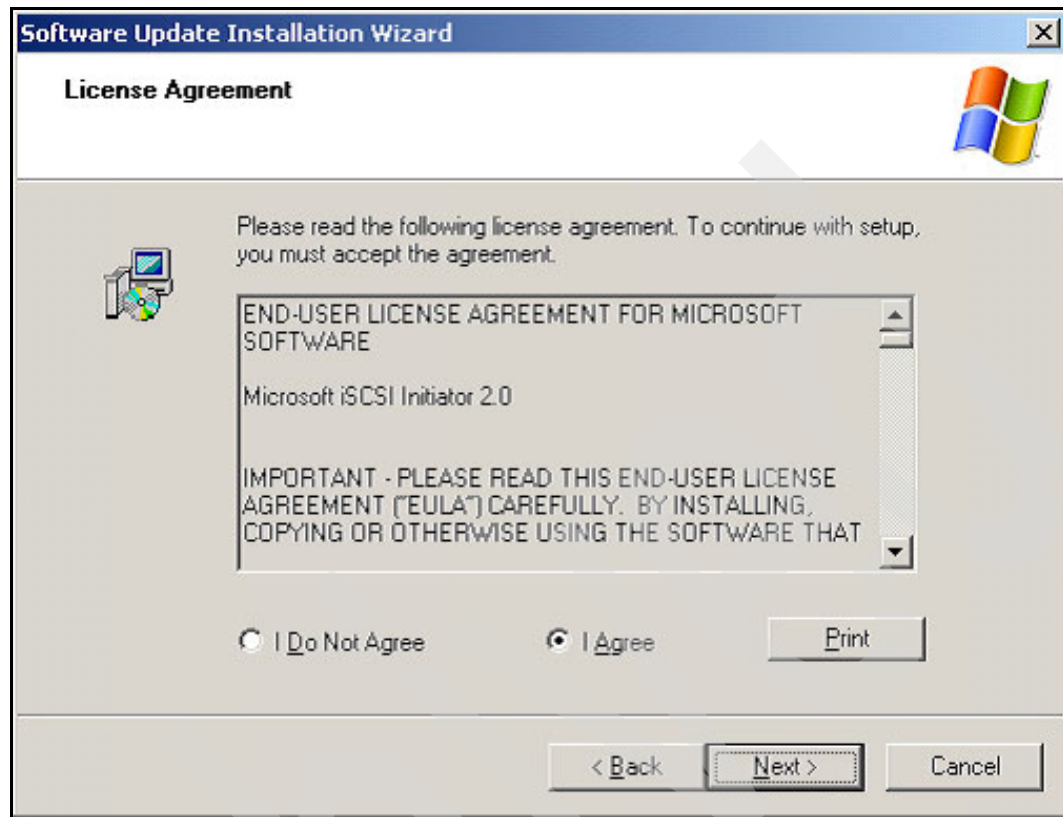


Figure 6-22 License Agreement window

4. The installation starts. At the end of the installation, in the Finish window (Figure 6-23), click **Finish**. If you do not want your server to reboot now, select the check box **Do not restart now**. Otherwise, your server will reboot immediately.

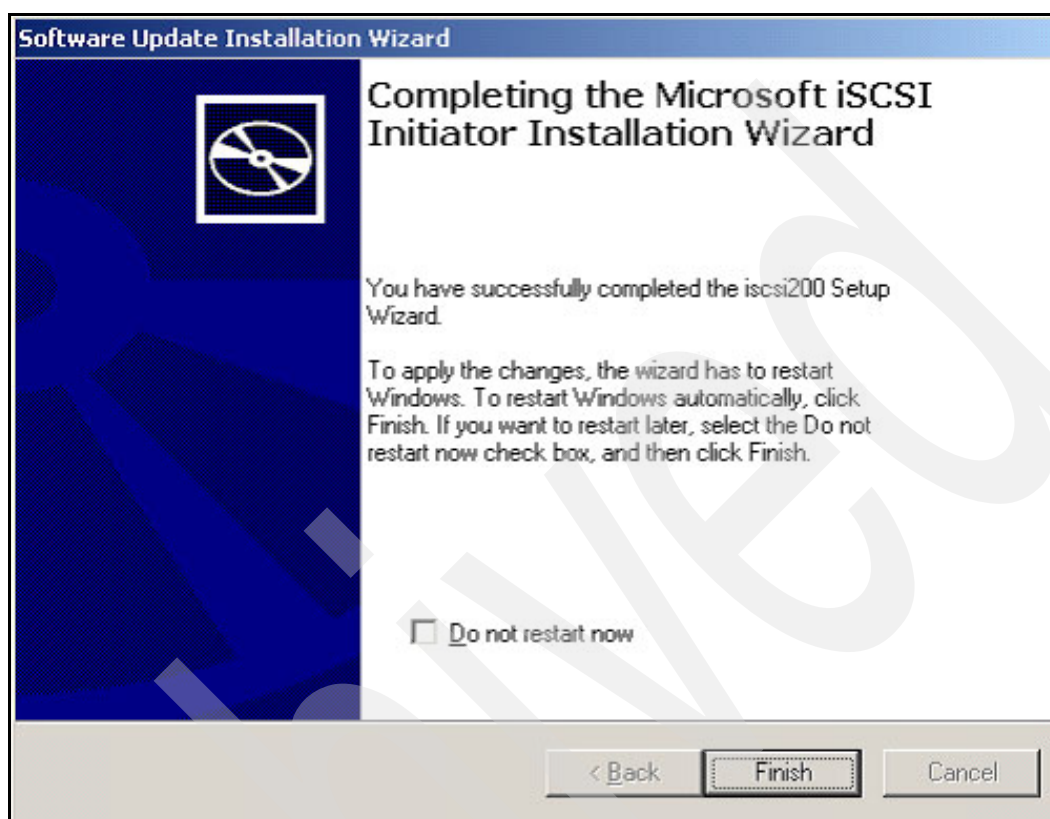


Figure 6-23 Finish window

Configuring iSCSI connectivity

1. Start Microsoft iSCSI Initiator by using the desktop icon or by selecting **Start → All Programs → Microsoft iSCSI Initiator → Microsoft iSCSI Initiator**. This will bring up the iSCSI Initiator Properties window (Figure 6-24 on page 193).

2. Copy the Initiator Node Name from the iSCSI Initiator Properties window (Figure 6-24).

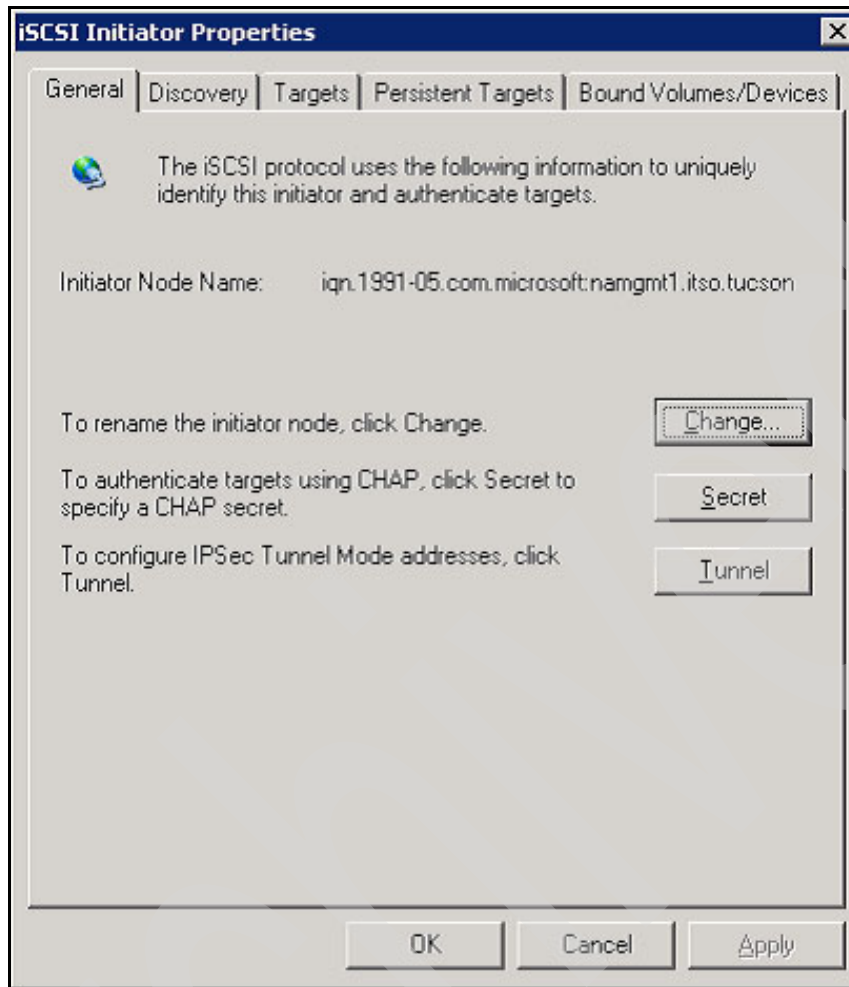


Figure 6-24 iSCSI Initiator Properties window

3. In the FilerView, select **LUNs** → **Initiator Groups** → **Add**. In the Add Initiator Group window (Figure 6-25), type in a name for the group you are creating, select **iSCSI** as the protocol for the group, select **Windows** as the operating system, and paste in the iSCSI initiator node name you copied on the previous step.

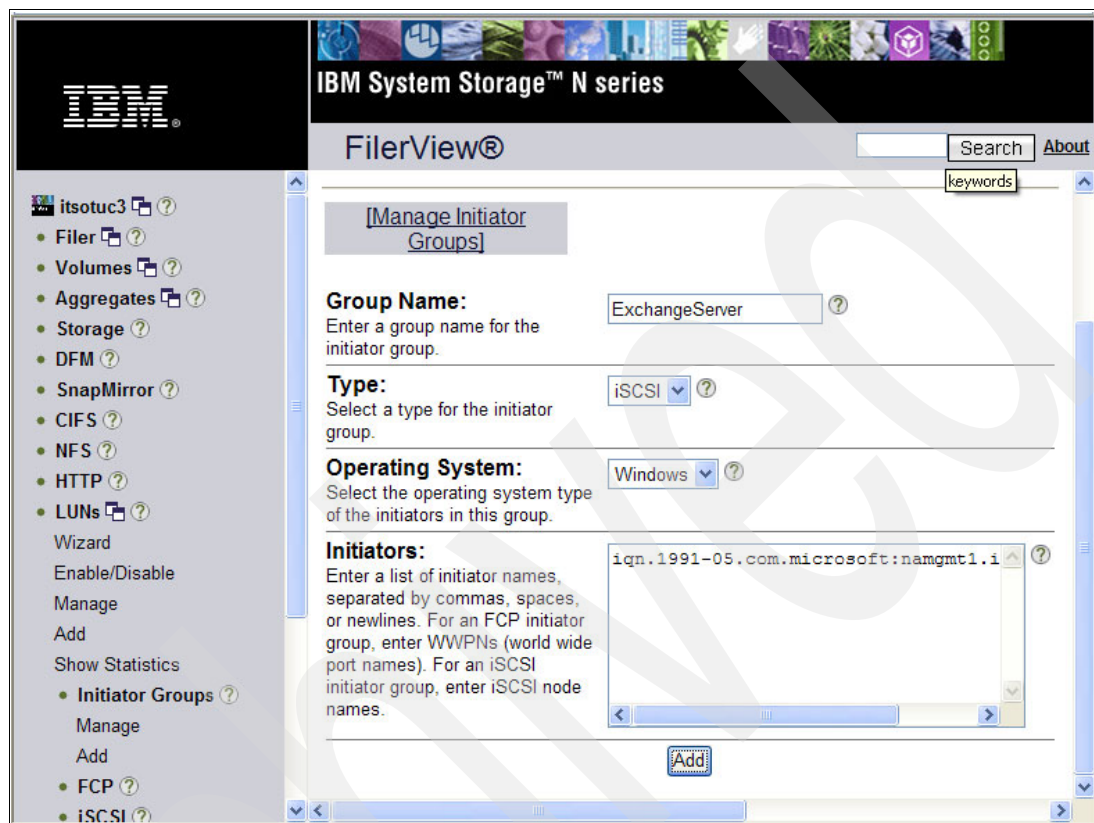


Figure 6-25 Add Initiator Group window

4. In the Microsoft Exchange Server, click the **Discovery** tab and the iSCSI Initiator Discovery window will be shown (Figure 6-26 on page 195). Click **Add** in the Target Portals session.

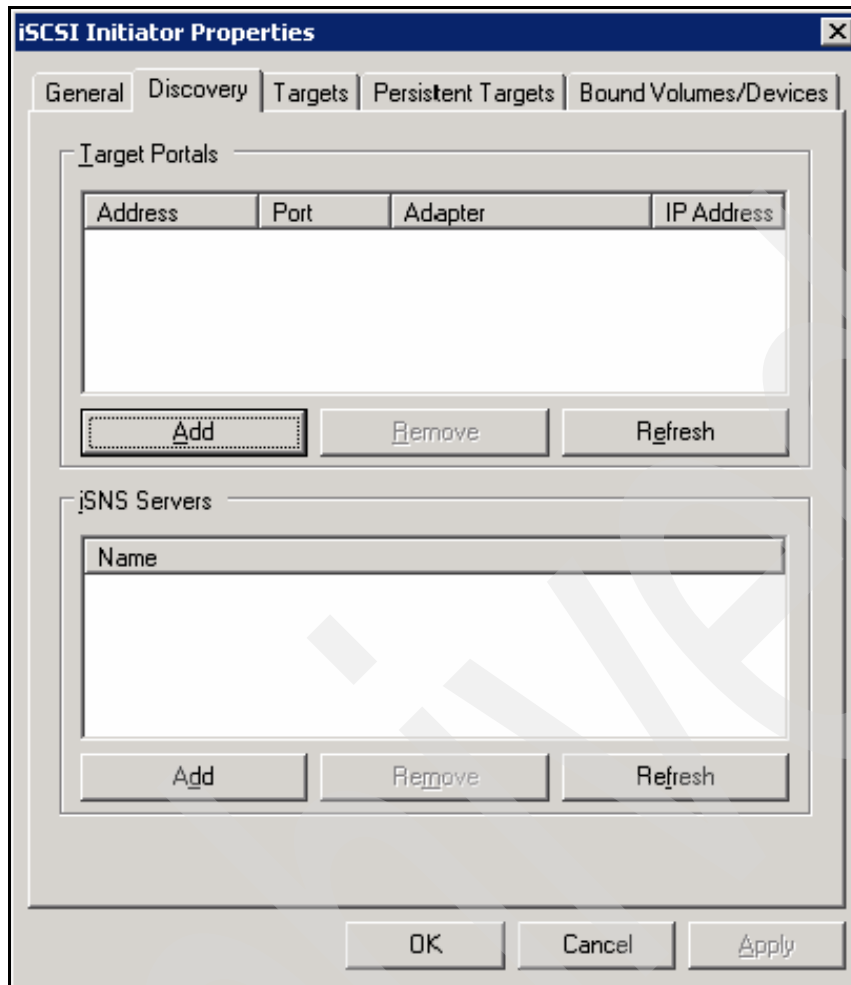


Figure 6-26 iSCSI Initiator Discovery window

5. In the Add Target Portal window (Figure 6-27), type in the IP address or DNS name of the filer and click **Advanced**.

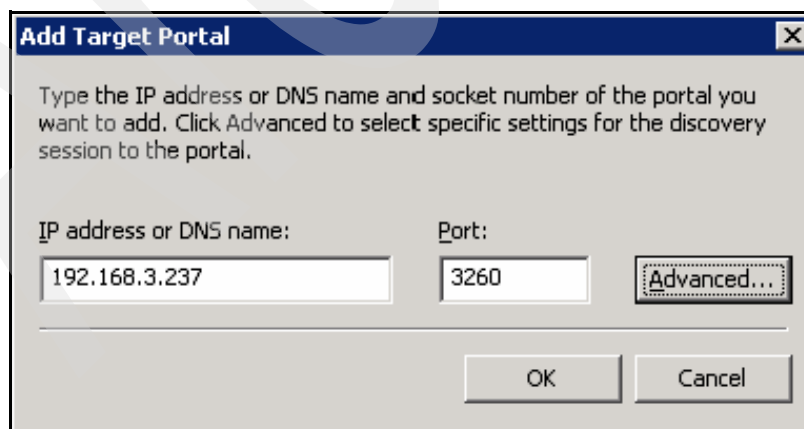


Figure 6-27 Add Target Portal window

6. In the Advanced Settings window (Figure 6-28), select **Microsoft iSCSI Initiator** as the Local Adapter and then select one of the IP addresses as the Source IP. At this time, the CHAP Authentication protocol may be configured, if it is defined on the IBM System Storage N series storage system configuration. For this scenario, we are not going to use CHAP. Click **OK**.

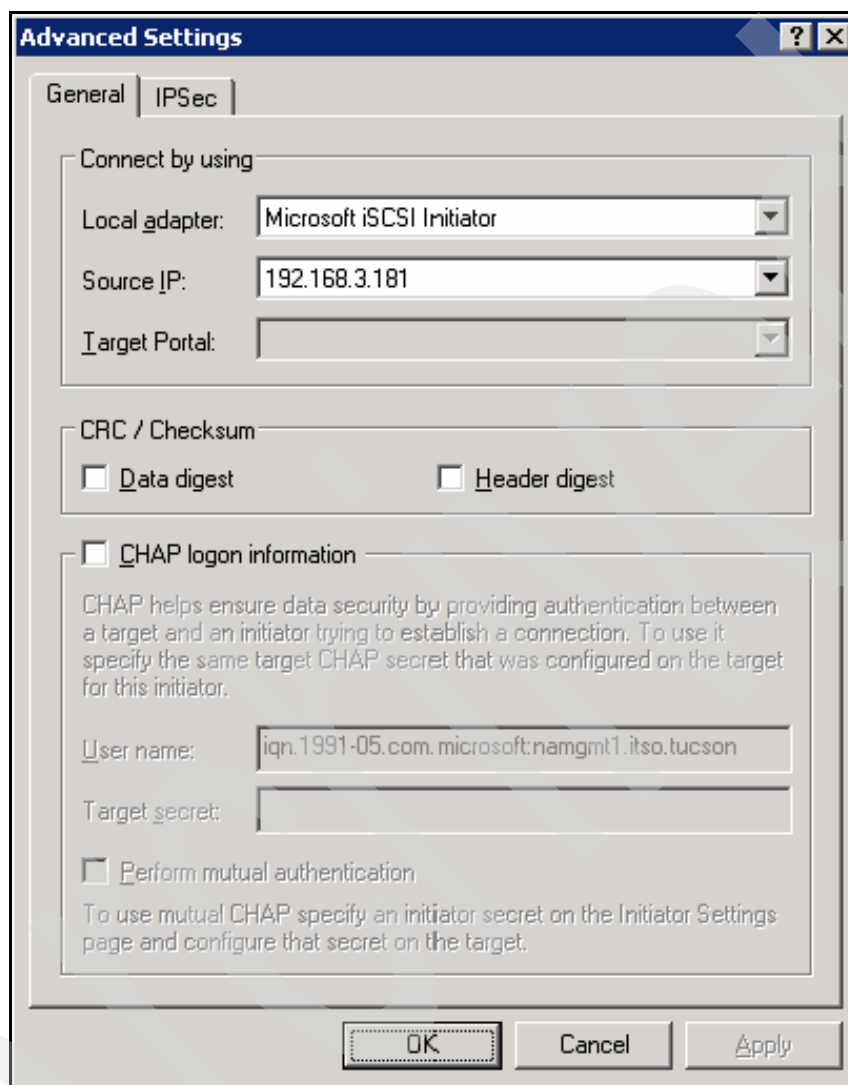
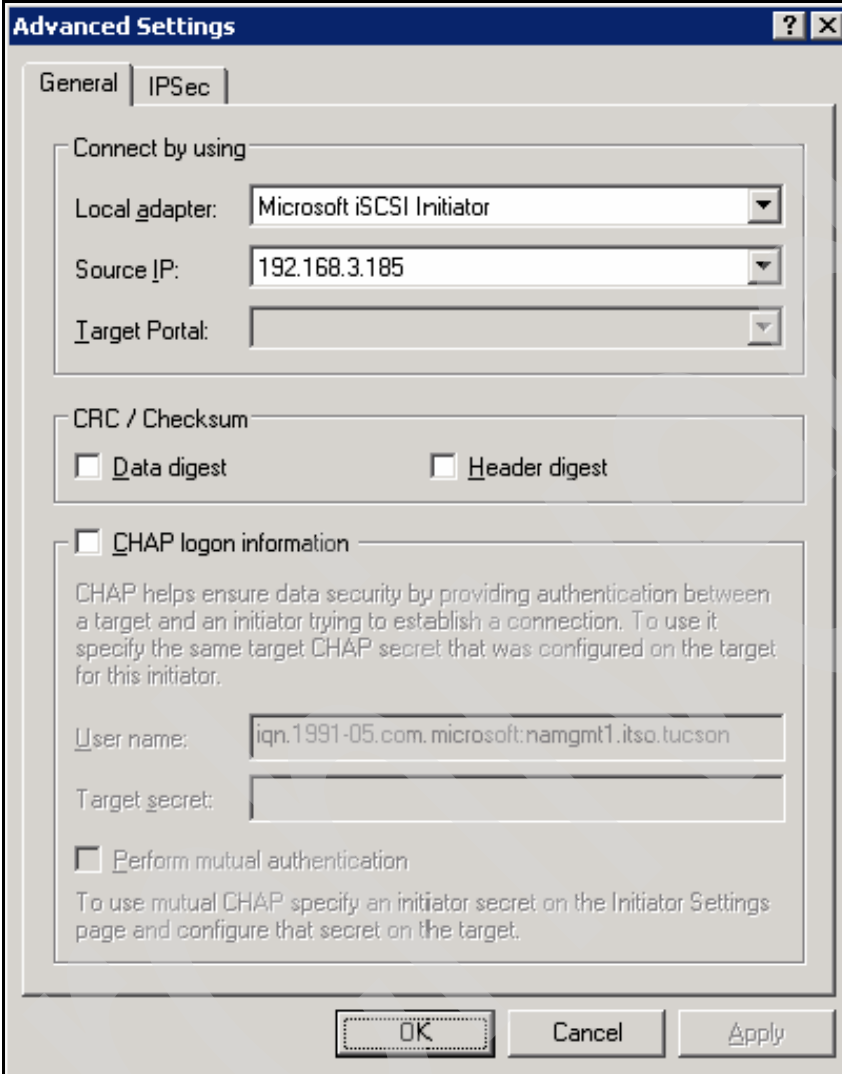


Figure 6-28 Advanced Settings window

7. Repeat steps 4 on page 194 , 5 on page 195, and 6 on page 196 to obtain the additional IP address, as shown in Figure 6-29. Click **OK**.



The image shows a Windows-style dialog box titled "Advanced Settings" with a question mark icon and a close button. It has two tabs: "General" and "IPSec", with "General" currently selected. The "Connect by using" section contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.185", and "Target Portal:" which is empty. Below this is a "CRC / Checksum" section with two unchecked checkboxes: "Data digest" and "Header digest". The "CHAP logon information" section is also unchecked and contains a text box for "User name:" with the value "iqn.1991-05.com.microsoft:namgmt1.itso.tucson" and an empty "Target secret:" text box. At the bottom of this section is an unchecked checkbox for "Perform mutual authentication" with explanatory text. At the very bottom of the dialog are three buttons: "OK", "Cancel", and "Apply".

Figure 6-29 Advanced Settings window

8. In the iSCSI Initiator Discovery window (Figure 6-30), you will notice that two paths were created for the same Target Portal, one for each interface. Depending on your infrastructure configuration, this may change due to more Target Portals configured on the IBM System Storage N series storage system.

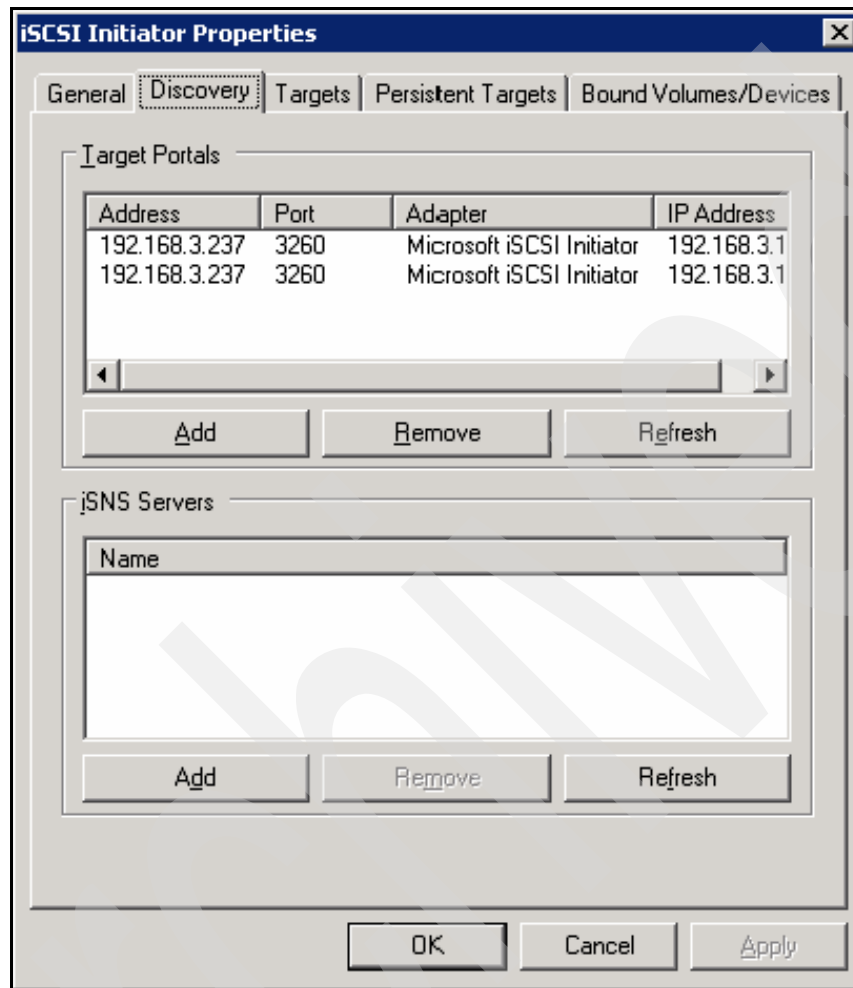


Figure 6-30 iSCSI Initiator Discovery window

9. Click the **Targets** tab and a list of the targets iSCSI storage devices will be shown, as shown in Figure 6-31 on page 199. Click **Log On** to configure the paths to the IBM System Storage N series storage system.

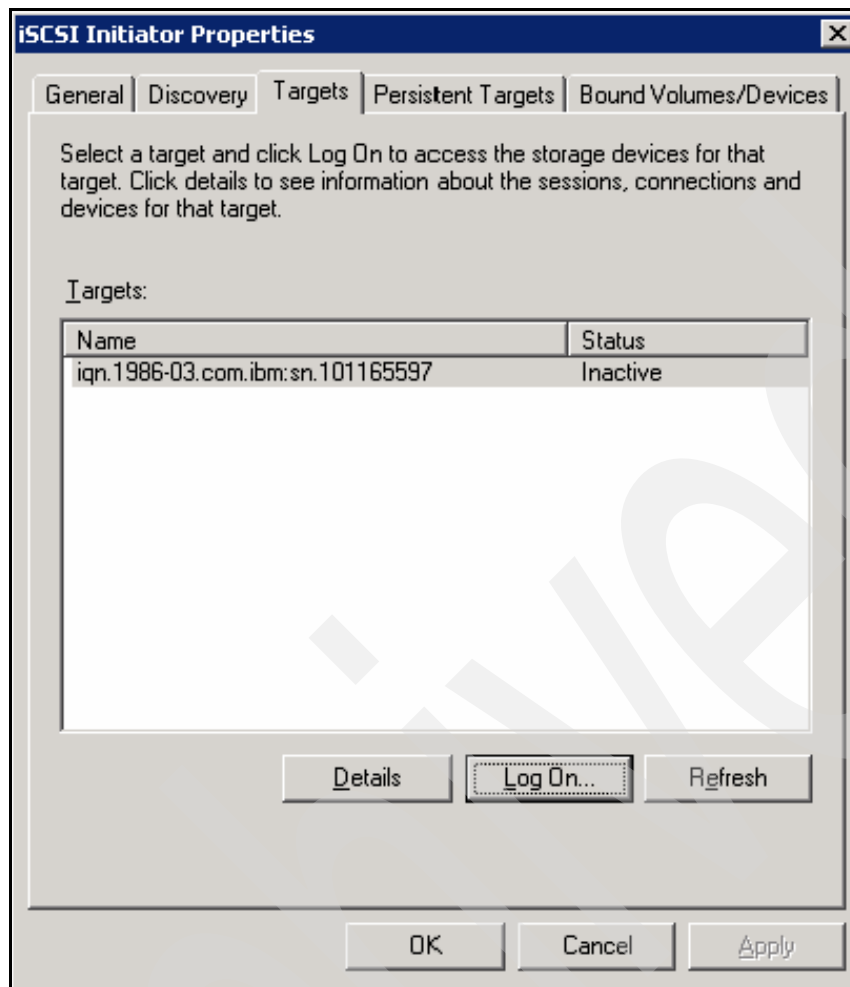


Figure 6-31 iSCSI Initiator Targets window

10. In the Log On to Target window, select both check boxes, as shown in Figure 6-32. This will configure the path be persistent and will enable the multipathing. Click **Advanced**.


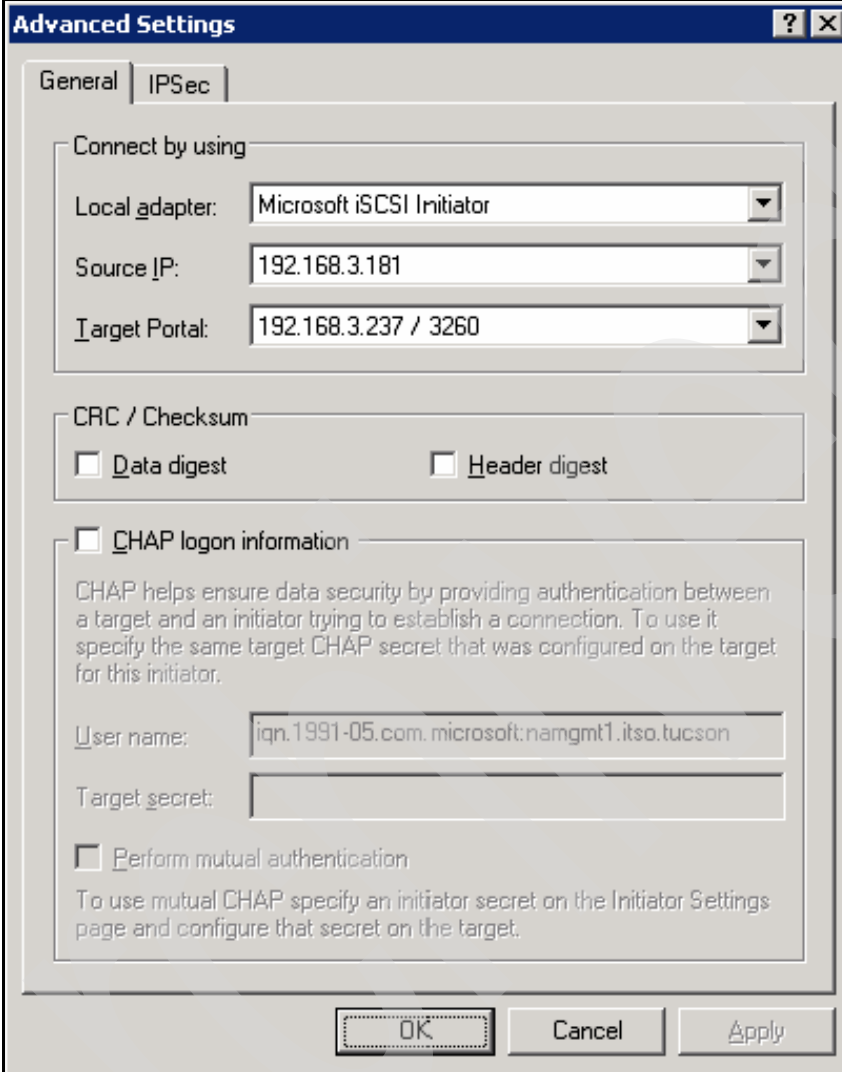


Figure 6-32 Log On to Target window

11. In the Advanced Settings window (Figure 6-33), select **Microsoft iSCSI Initiator** as the Local Adapter, select the first IP address as the Source IP, and select the proper Target Portal's combination of IP address and port number. Click **OK**.



The image shows a Windows-style dialog box titled "Advanced Settings" with a blue header bar containing a question mark and a close button. The dialog has two tabs: "General" and "IPSec", with "General" currently selected. The "Connect by using" section contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.181", and "Target Portal:" set to "192.168.3.237 / 3260". Below this is a "CRC / Checksum" section with two unchecked checkboxes: "Data digest" and "Header digest". The "CHAP logon information" section is also unchecked and contains a text box for "User name:" with the value "iqn.1991-05.com.microsoft:namgmt1.itso.tucson", an empty "Target secret:" text box, and an unchecked "Perform mutual authentication" checkbox. A descriptive paragraph for CHAP is present. At the bottom are "OK", "Cancel", and "Apply" buttons.

Figure 6-33 Advanced settings window

12. Repeat steps 9 on page 198, 10 on page 199, and 11 for the additional IP address on the Microsoft Exchange Server. Select this additional IP address as the Source IP, as shown in Figure 6-34 on page 201. Click **OK**.

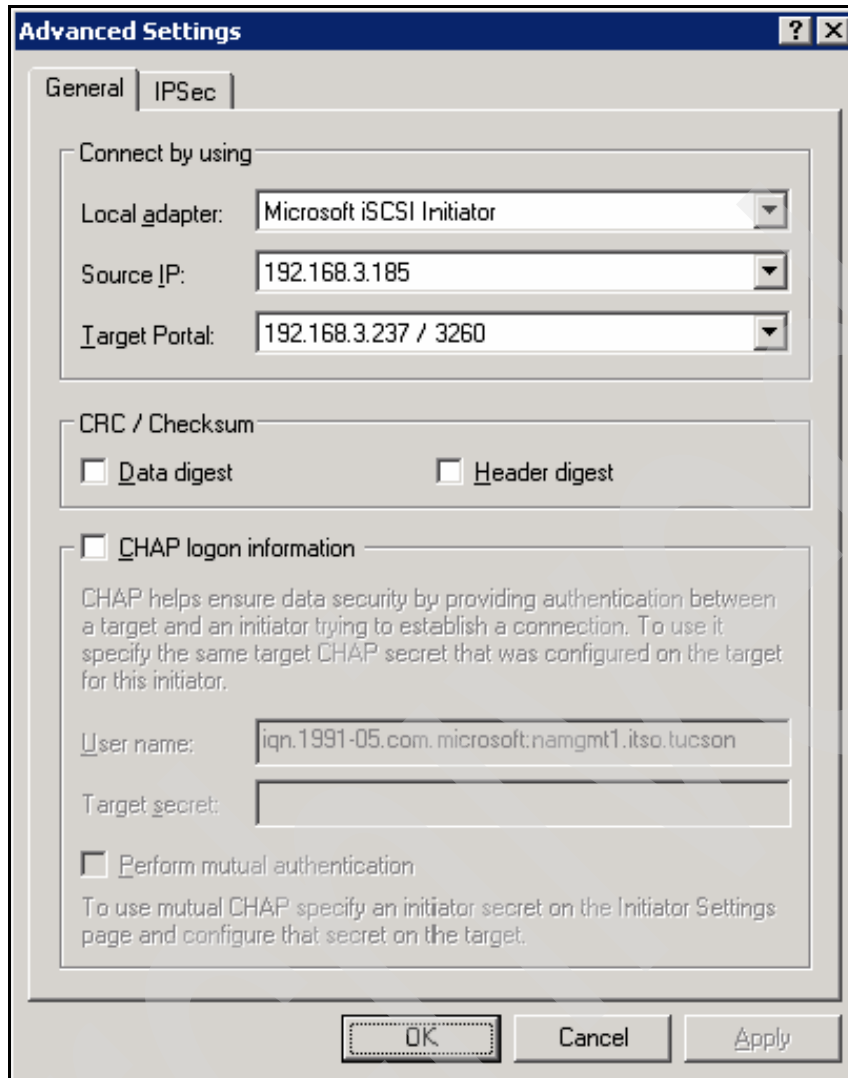


Figure 6-34 Advanced Settings window

13. Click **OK** to close the iSCSI Initiator Properties window.

At this point, the communication between the Microsoft Exchange Server and the IBM System Storage N series storage system is established. SnapDrive installation can take place.

Installing SnapDrive for Windows

These are the steps to install SnapDrive on the Microsoft Exchange Server:

1. In the SnapDrive installation welcome window (Figure 6-35), click **Next**.

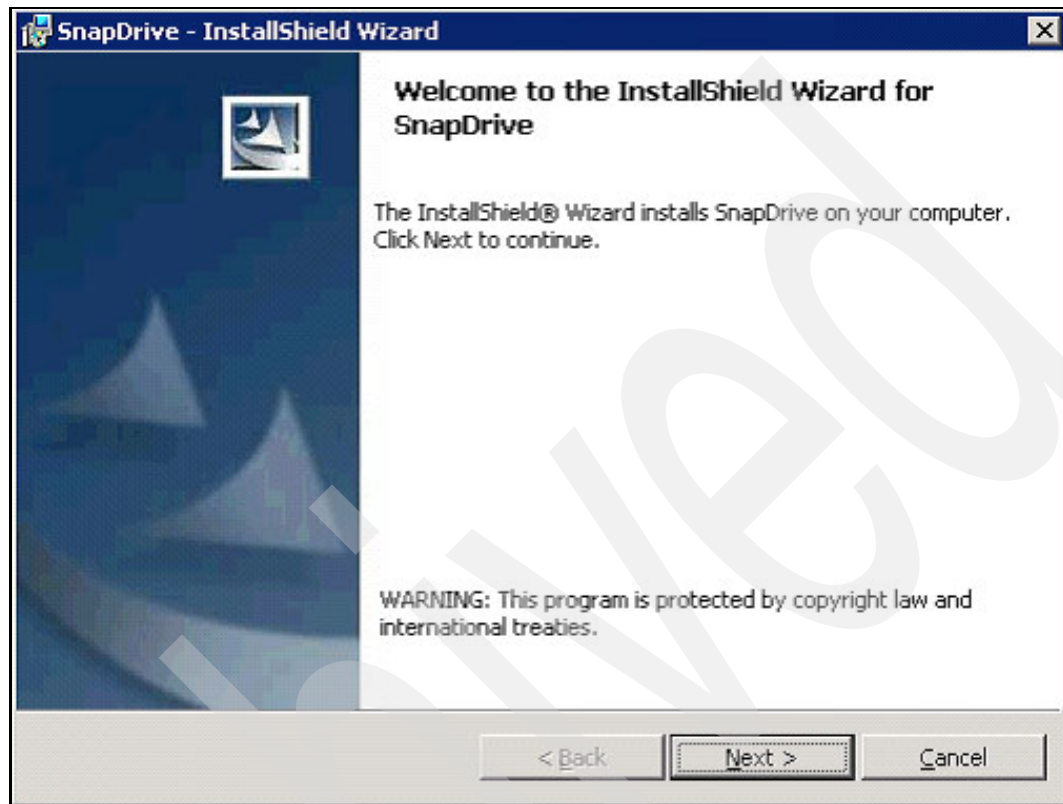


Figure 6-35 SnapDrive installation welcome window

2. In the License Agreement window (Figure 6-36), accept the terms in the License Agreement and click **Next**.

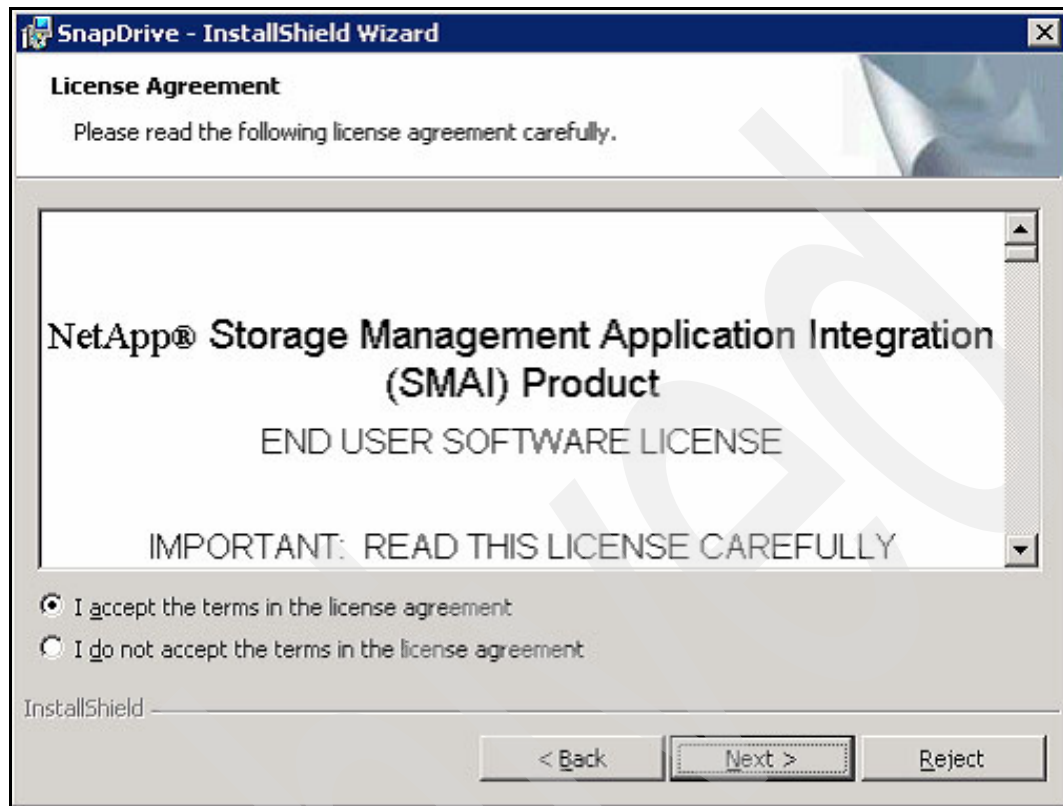


Figure 6-36 License Agreement

3. In the License key window (Figure 6-37), type in the License key for SnapDrive and click **Next**.

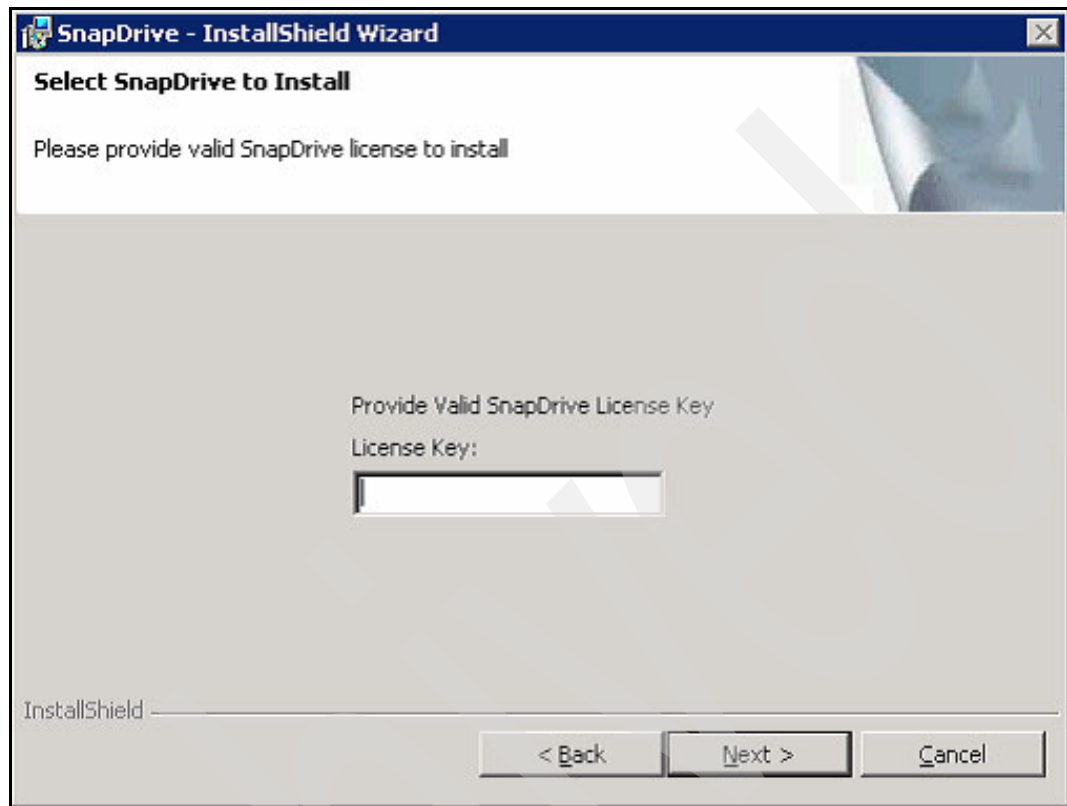


Figure 6-37 License Key

4. SnapDrive will check for the minimum requirements for the iSCSI driver version, as shown in Figure 6-38. Click **Next**.

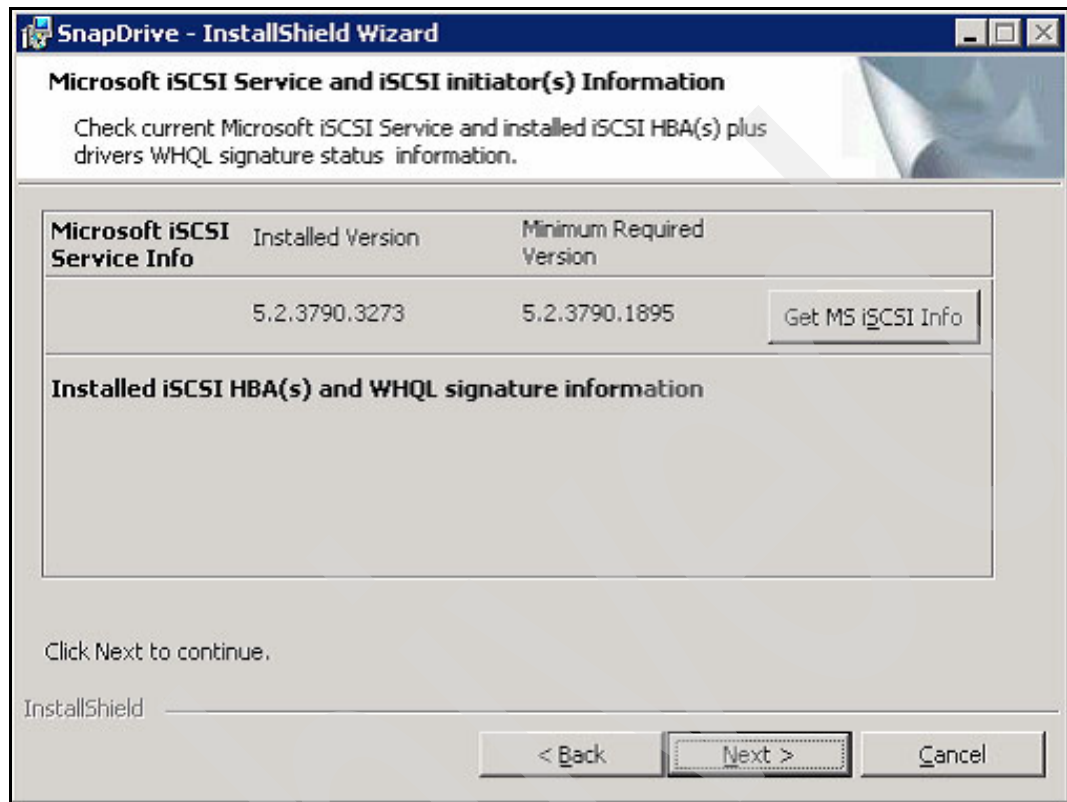
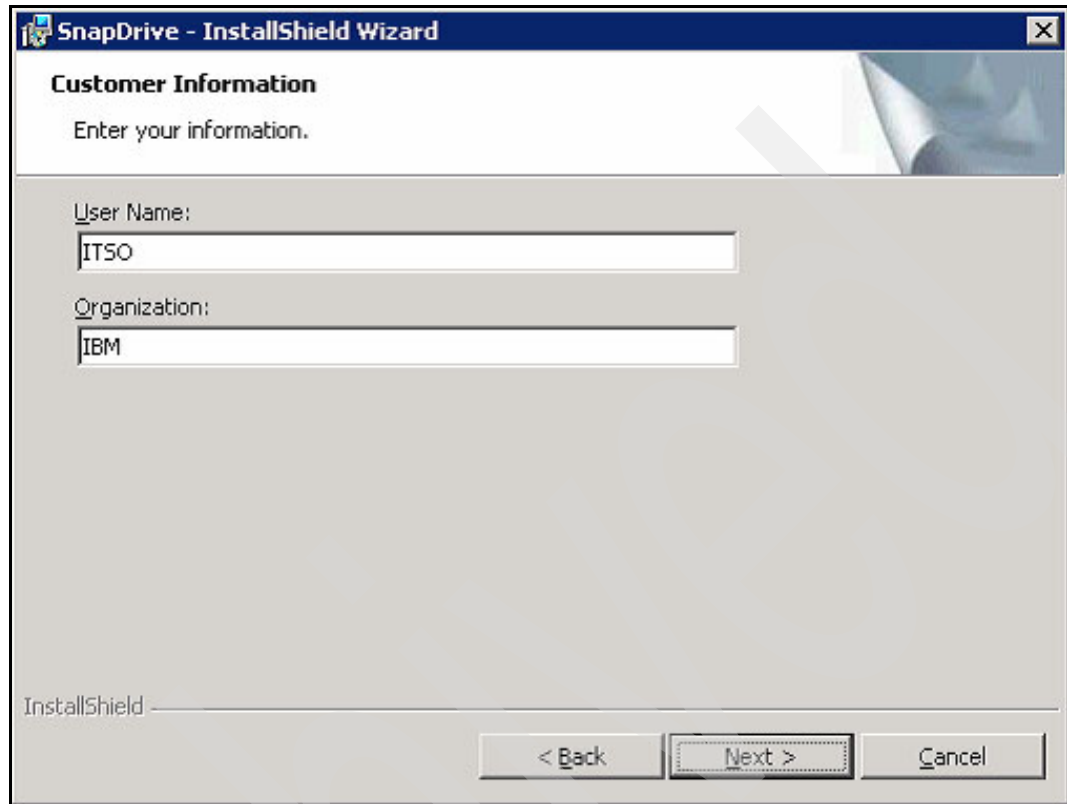


Figure 6-38 iSCSI Service and iSCSI Initiator Information window

5. In the Customer information window (Figure 6-39), type in the User Name and Organization information and click **Next**.



The image shows a Windows-style dialog box titled "SnapDrive - InstallShield Wizard". The window has a blue title bar with a close button in the top right corner. Below the title bar, the text "Customer Information" is displayed in bold, followed by the instruction "Enter your information." in a smaller font. The main area of the window contains two text input fields. The first field is labeled "User Name:" and contains the text "ITSO". The second field is labeled "Organization:" and contains the text "IBM". At the bottom of the window, there is a status bar that says "InstallShield". To the right of the status bar are three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

Figure 6-39 Customer information

6. In the Destination Folder window (Figure 6-40), confirm or change the destination folder for the installation files and click **Next**.

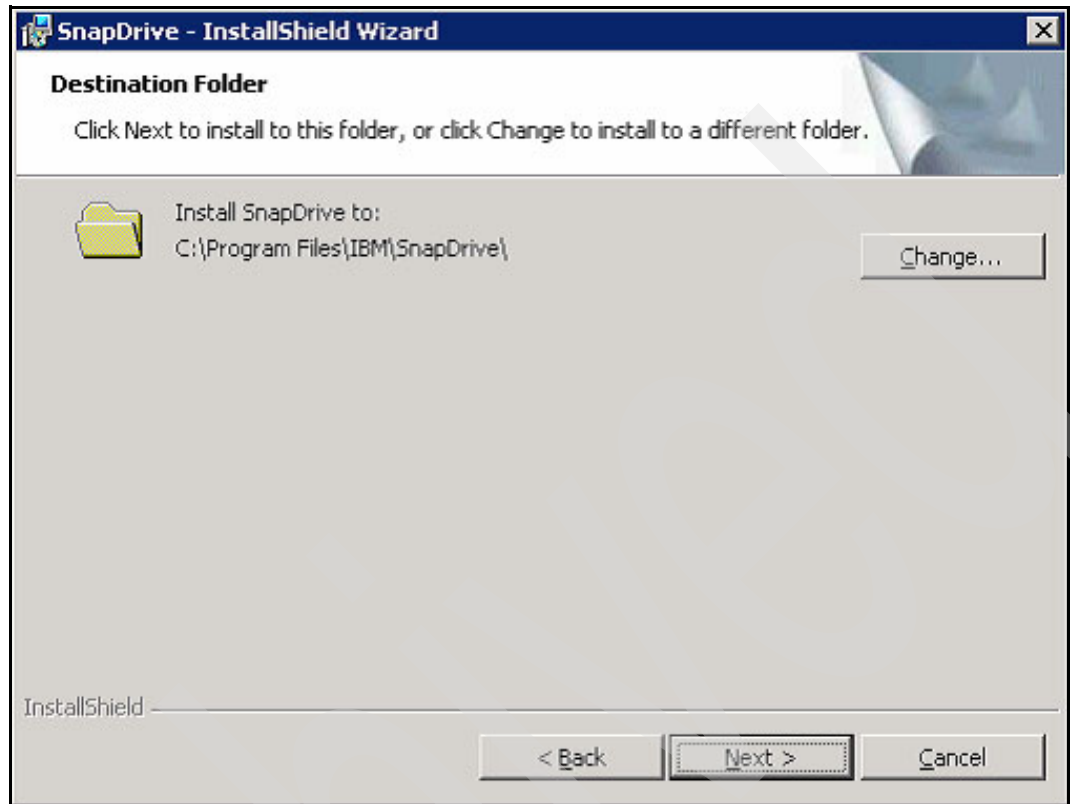
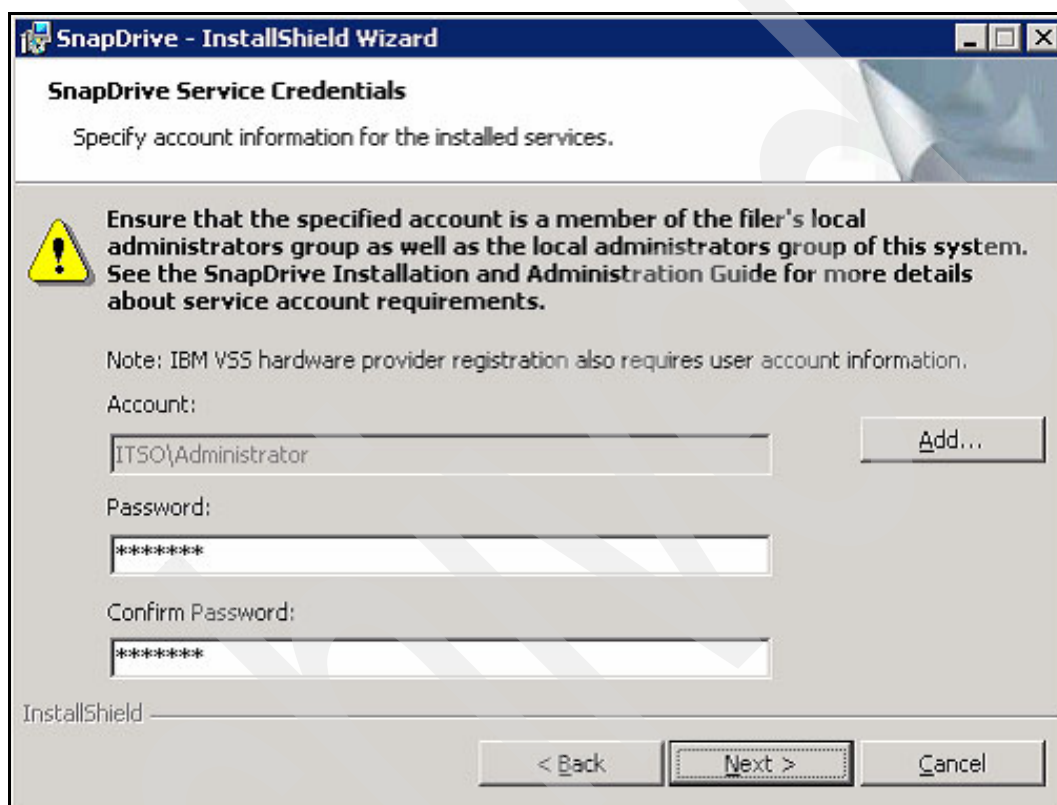


Figure 6-40 Destination Folder

7. In the SnapDrive Service Credentials window (Figure 6-41), type in the account and password for the user account to be used to start the SnapDrive service and click **Next**.

Note: The SnapDrive service user account should be a member of the Active Directory's Domain Admins group and be a member of the filer's local administrators group.



The image shows a Windows-style dialog box titled "SnapDrive - InstallShield Wizard". The main heading is "SnapDrive Service Credentials" with the instruction "Specify account information for the installed services." Below this is a yellow warning icon and a message: "Ensure that the specified account is a member of the filer's local administrators group as well as the local administrators group of this system. See the SnapDrive Installation and Administration Guide for more details about service account requirements." A note follows: "Note: IBM VSS hardware provider registration also requires user account information." There are three input fields: "Account:" with the text "ITSO\Administrator" and an "Add..." button; "Password:" with masked characters "*****"; and "Confirm Password:" with masked characters "*****". At the bottom left is the "InstallShield" logo, and at the bottom right are three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

Figure 6-41 SnapDrive Service Credentials

8. Click **Finish**.

After the installation is done, no restart is needed in order to get the SnapDrive for Windows working. When accessing the Computer Management MMC, you will notice the SnapDrive snap-in and the iSCSI Management snap-in showing the already configured connections for the IBM System Storage N series storage system (see Figure 6-42 on page 209).

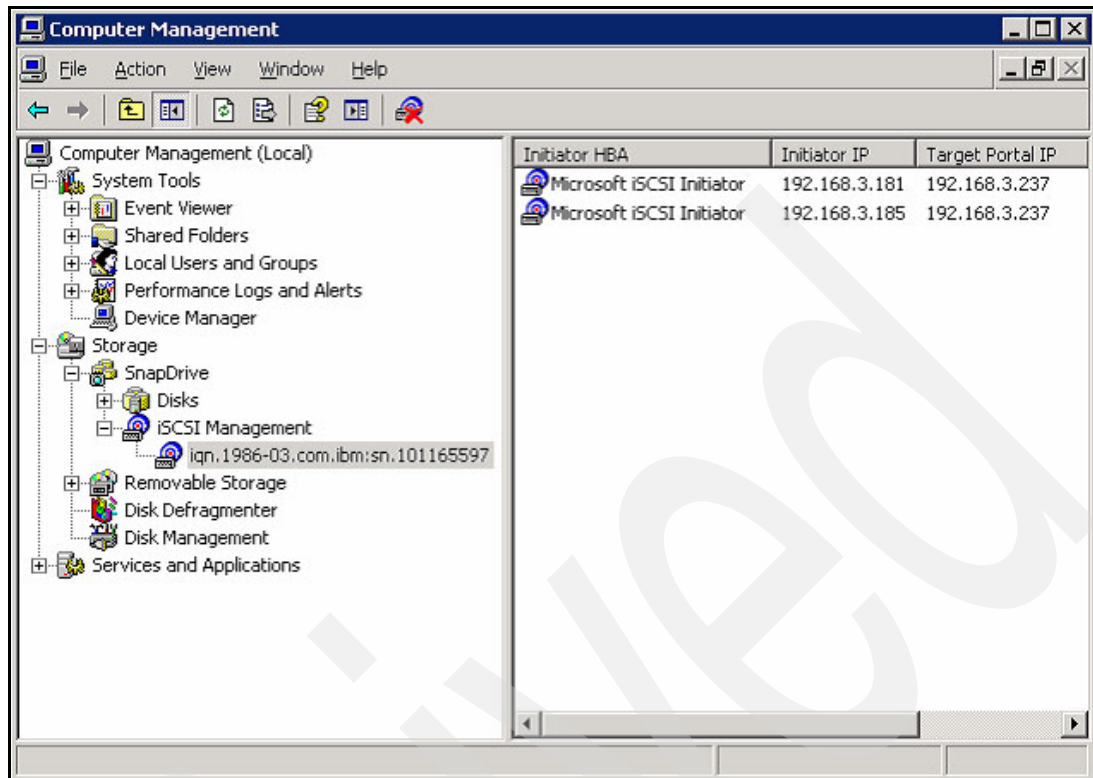


Figure 6-42 Computer Management MMC with SnapDrive

Creating disks from SnapDrive

Now that the Data ONTAP DSM for Windows MPIO, Microsoft iSCSI Initiator Software, and SnapDrive are installed, the disk drives can be added using the SnapDrive software.

In order to create the disks from SnapDrive, we assume that:

- ▶ An aggregate has been created on the N series.
- ▶ A volume has been created on the aggregate.
- ▶ A CIFS share has been created mapping the path to the volume.
- ▶ iSCSI and networking infrastructures are in place and working.

These are the steps to create the LUN from the SnapDrive:

1. Access the Computer Management MMC.
2. Expand **Storage** and then expand **SnapDrive**.

3. Right-click **Disks** and click **Create Disk**, as shown in Figure 6-43.

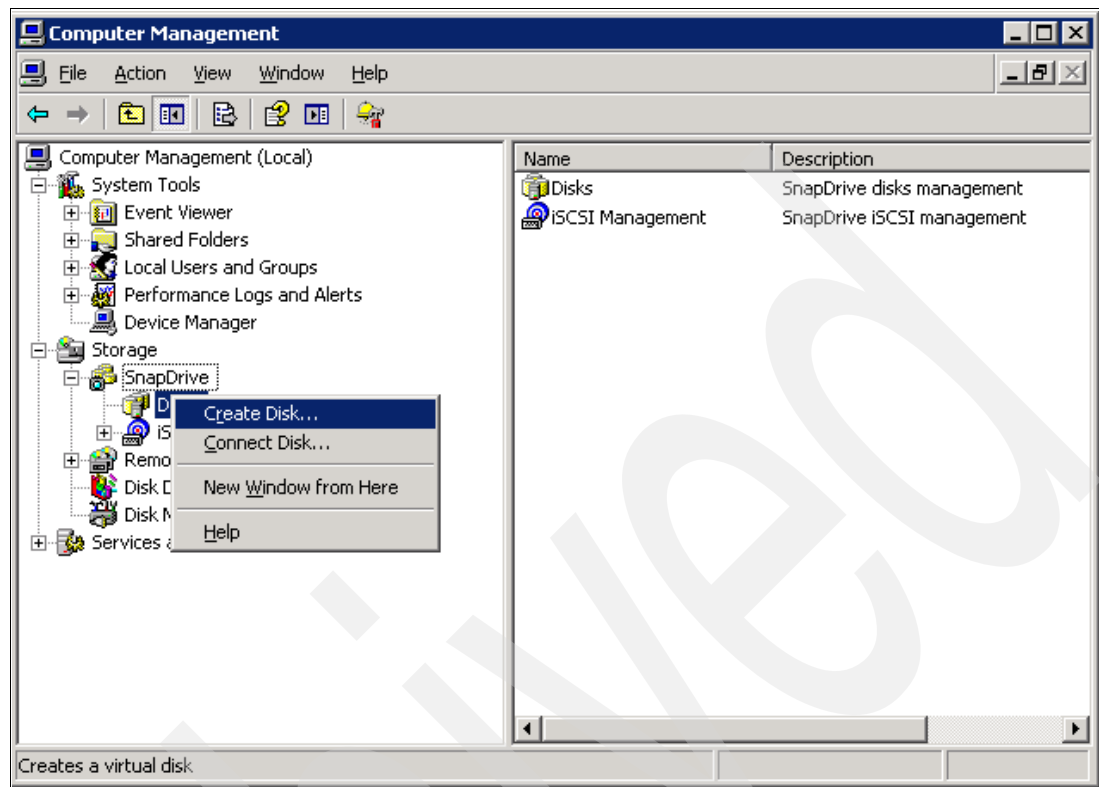


Figure 6-43 Disk creation using SnapDrive

4. In the Create Disk welcome window (Figure 6-44), click **Next**.


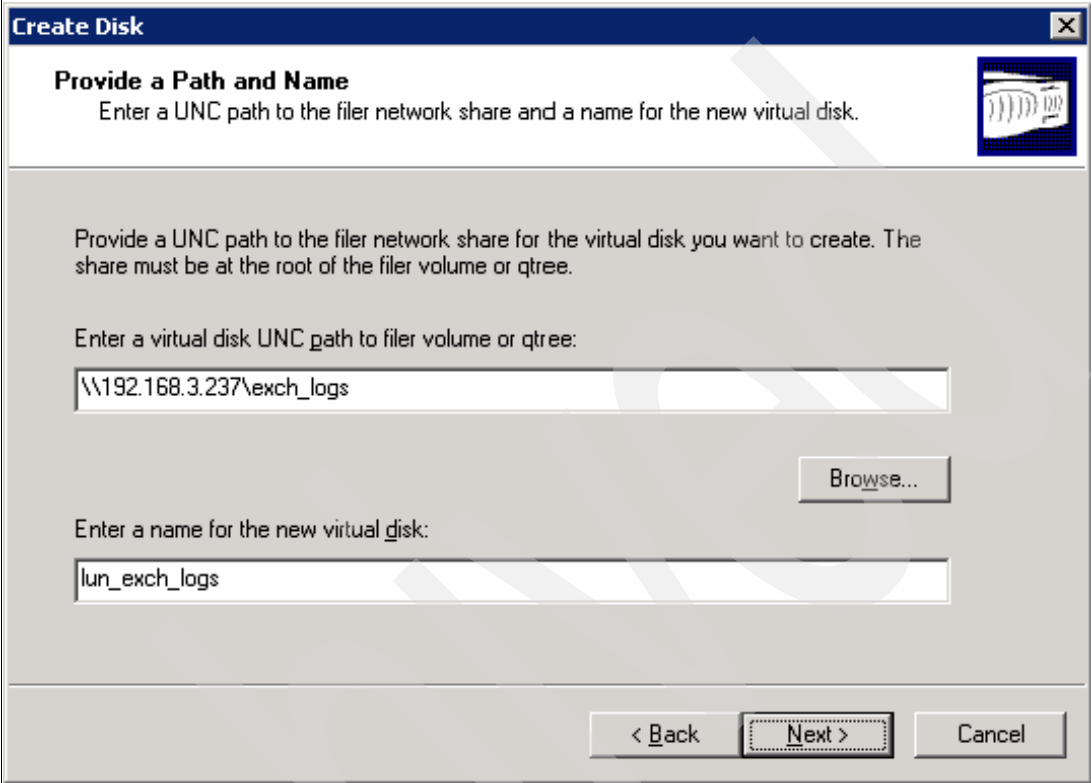


Figure 6-44 Create Disk welcome window

5. In the Provide a Path and Name window (Figure 6-45), enter the information needed. For the UNC path, use \\Filer IP\CIFS Share on the filer. Remember that this CIFS share points to the volume you created. For the name for the new virtual disk, type in the name you want to assign to the LUN that will be created. Click **Next**.



The screenshot shows a Windows-style dialog box titled "Create Disk". The main heading is "Provide a Path and Name", followed by the instruction "Enter a UNC path to the filer network share and a name for the new virtual disk." Below this, a paragraph explains: "Provide a UNC path to the filer network share for the virtual disk you want to create. The share must be at the root of the filer volume or qtree." There are two input fields. The first is labeled "Enter a virtual disk UNC path to filer volume or qtree:" and contains the text "\\192.168.3.237\exch_logs". To the right of this field is a "Browse..." button. The second input field is labeled "Enter a name for the new virtual disk:" and contains the text "lun_exch_logs". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a dashed border.

Figure 6-45 Provide path and name window

6. In the Select a virtual disk type window (Figure 6-46), select **Dedicated** if this disk will be accessed by only one server. Select **Shared** if this disk will be accessed by a Cluster Service. Click **Next**.




Figure 6-46 Select a Virtual Disk Type window

7. In the Select Virtual Disk Properties window (Figure 6-47), select if you want to assign a drive letter (and which drive letter) for the disk being created or if you want to assign a Volume Mount Point. The next option will impact the size that will be available for the LUN creation. You need to select if you will reserve space for at least one snapshot of this LUN on the volume. Then, enter the size you want for this LUN. Click **Next**.

Create Disk

Select Virtual Disk Properties
Provide the drive letter and the size of the virtual disk to create.

Provide a Drive Letter or Volume Mount Point

☒ Assign a Drive Letter: L

☐ Use Volume Mount Point:

Do you want to limit the maximum disk size to accommodate at least one snapshot on the volume?

☐ Yes ☒ No

Enter a virtual disk size that is equal to or less than the maximum size, but greater than or equal to the minimum size.

Maximum Virtual Disk Size: 44 GB

Minimum Virtual Disk Size: 32 MB

Enter or Select Virtual Disk Size: 42 GB

< Back Next > Cancel

Figure 6-47 Select Virtual Disk Properties window

8. In the Select Initiators window (Figure 6-48), select the initiators from the Available Initiators column on the left and click the arrow to move them to the Selected Initiators on the right. Because we are using iSCSI, the initiators will be listed as the Initiator node name on the Microsoft iSCSI Initiator software configuration. Click **Next**.

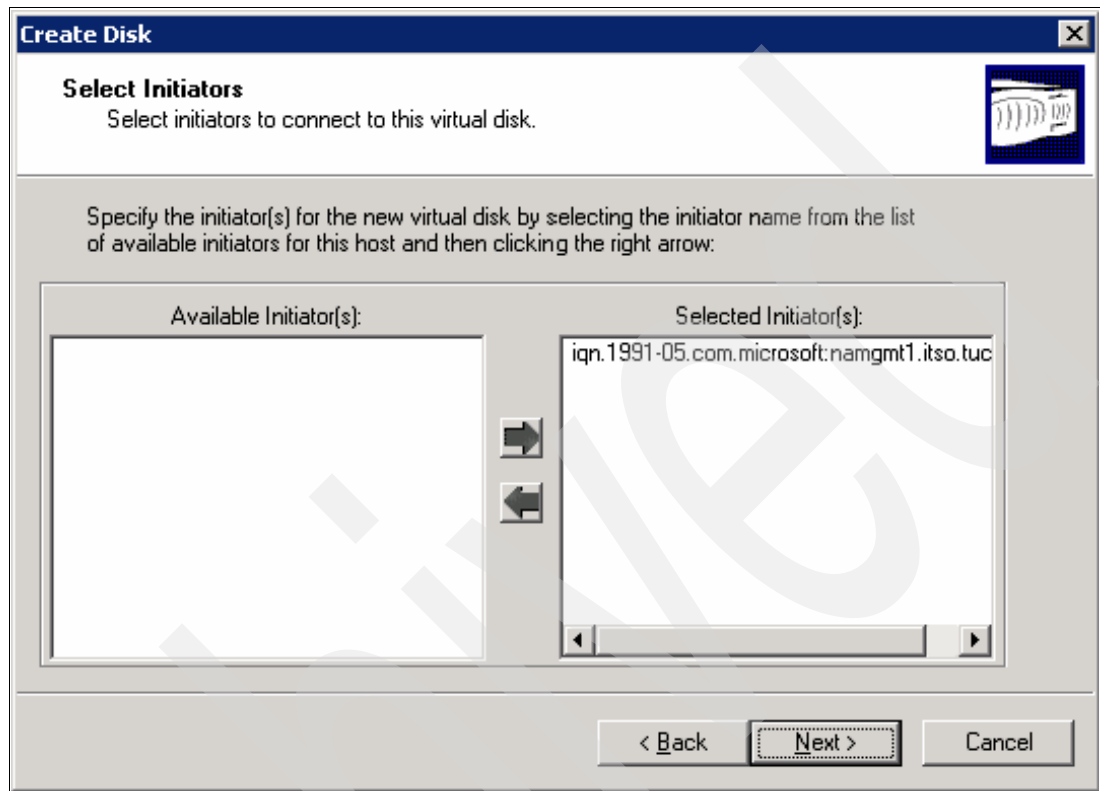


Figure 6-48 Select Initiators window

9. In the Summary window (Figure 6-49), verify if all the information is correct and click **Finish**. This will start the LUN creation process on the filer.



Figure 6-49 Summary window

10. SnapDrive will format the drive and a Disks window will appear (Figure 6-50). Click **OK**.




Figure 6-50 Disks window

11. Now the disk is available for the Microsoft Exchange Server using the drive letter you selected or using the Volume Mount Point you created.

Important: When using iSCSI to map the LUNs from the IBM System Storage N series storage system to the Microsoft Exchange Server, you should guarantee that the Server service does not start *before* the Microsoft iSCSI Initiator service. If that happens, the shares on the LUNs will be unavailable until you restart the Server service.

To resolve this problem, create a dependency between the Server service and the Microsoft iSCSI Initiator service by following these steps:

1. Select **Start** → **Run** and type **regedt32.exe** to start the Registry Editor.
2. Navigate to the key **HKLM\SYSTEM\CurrentControlSet\Services\LanManServer**.
3. Select **LanManServer** and select **Edit** → **New** → **Multi-String Value**.
4. Type **DependOnService** as the value name. Notice the lack of space between the words and capitalization of letters.
5. Double-click the **DependOnService** value and type in **MSiSCSI**.
6. Click **OK**.
7. In the iSCSI Initiator Properties window, shown in Figure 6-24 on page 193, click the tab **Bound Volumes/Devices**.
8. Click **Add** and type in the drive letters for all the LUN drives available to the Microsoft Exchange Server. This configuration will guarantee that the Microsoft iSCSI Initiator service will only be started after all resources listed on this tab are online as well. Click **OK**.
9. Restart the server.

6.2 SnapManager for Microsoft Exchange

SnapManager for Exchange is an licensed product that helps Exchange Administrators manage day by day activities. Backups, restores, management, archival, and many other options are facilitated by using the SME integrated with the SnapDrive on the Microsoft Exchange environment.

To take advantage of the SME features, you need to install it on the Microsoft Exchange Server. Optionally, you can install it on a second server only for management purposes.

Microsoft Exchange Server software installation will not be covered by this IBM Redbooks publication. Because the database files, transaction log files, SMTP queues, and MTA queues will be moved during SME configuration, the actual path for these files makes no difference to the configuration. It will make no difference if they are in the original installation path or if they had been moved for administrative purposes.

Installing SnapManager for Microsoft Exchange

These are the steps to install SnapManager for Exchange on the Microsoft Exchange server:

Important: SnapManager for Microsoft Exchange Version 4.0 requires Windows Powershell to be installed first because of its support for Microsoft Exchange 2007. Many automated functions on SME are executed by Powershell scripts.

1. Start the installation, and in the SME Welcome window (Figure 6-51), click **Next**.



Figure 6-51 SME Welcome window

2. In the License Agreement window (Figure 6-52), accept the terms and click **Next**.

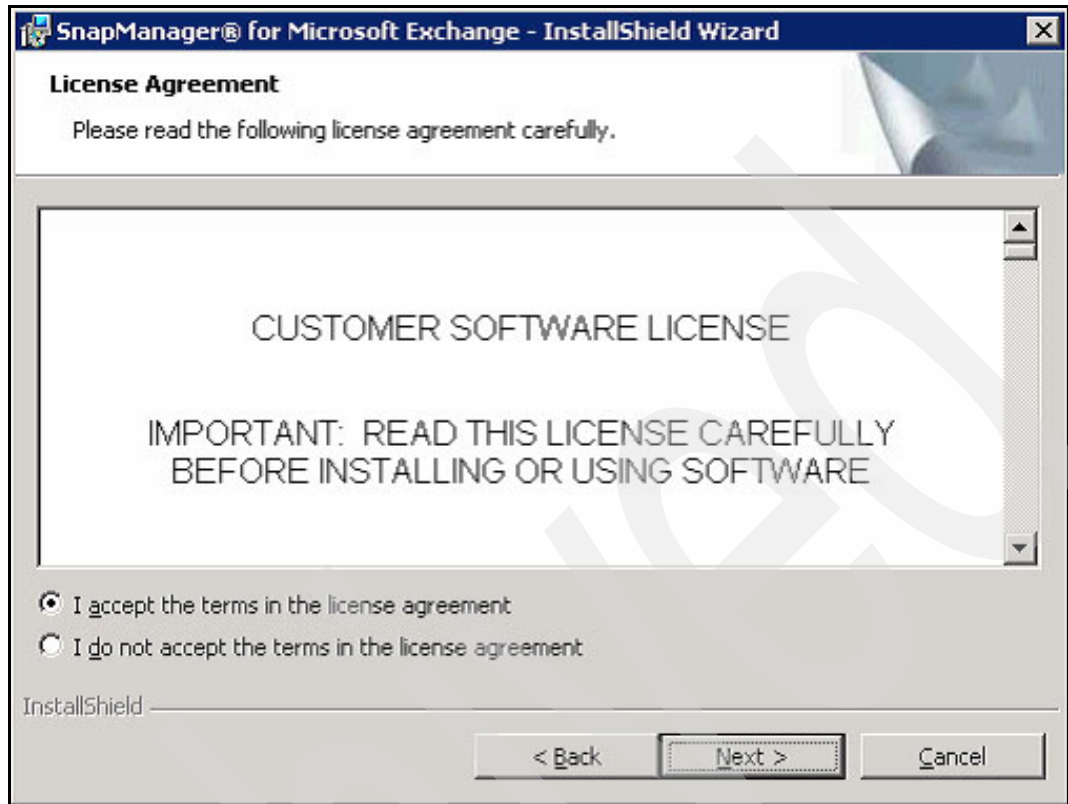
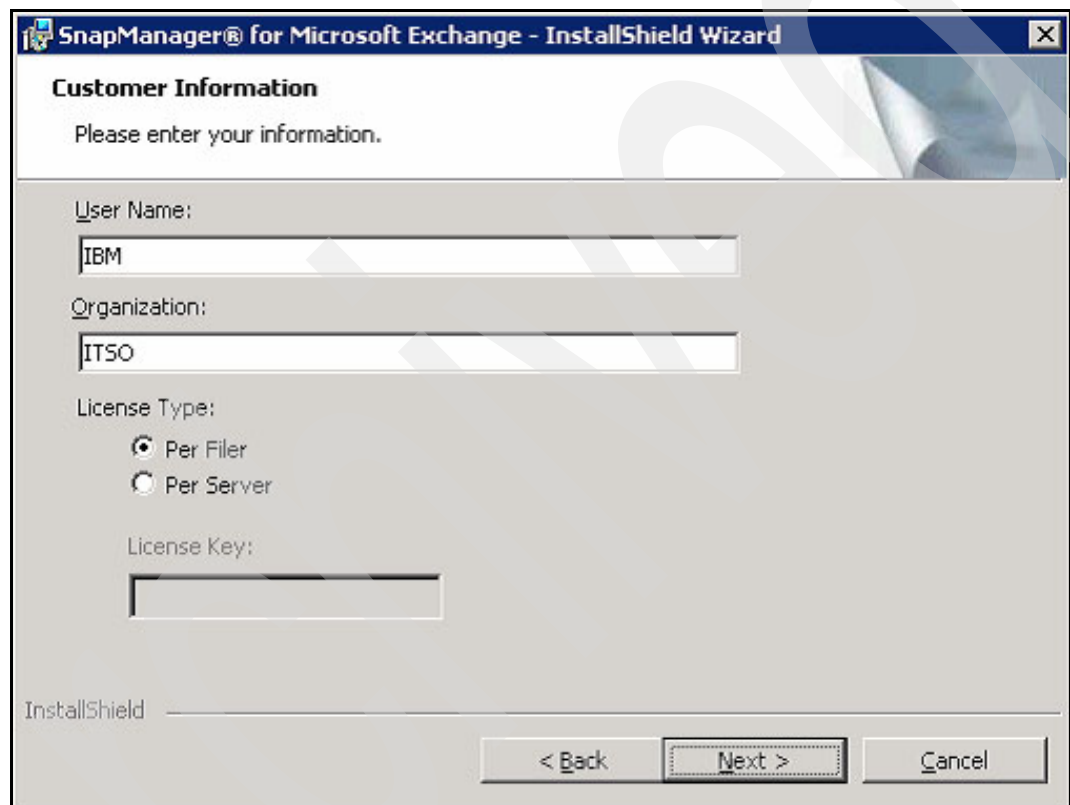


Figure 6-52 License Agreement window

3. In the Customer Information window (Figure 6-53), type in the information for the User Name and Organization. Select the type of License you have for the SnapManager for Microsoft Exchange. If it is a Per Server License, you need to type in the license before you proceed. If it is a Filer License, then the license is installed on the IBM System Storage N series storage system. Click **Next**.

Note: If you have a Filer License, this license has to be installed prior to the installation of SME. In the FilerView, select **Filer** → **Manage Licenses** and scroll down the right pane until you see the SnapManager Exchange license. Add the SnapManager Exchange license and click **Apply**.



The image shows a screenshot of the 'SnapManager® for Microsoft Exchange - InstallShield Wizard' window. The title bar includes the product name and a close button. The main heading is 'Customer Information' with a sub-instruction 'Please enter your information.' Below this, there are three input fields: 'User Name:' with 'IBM' entered, 'Organization:' with 'ITSO' entered, and 'License Key:' which is empty. Under the 'License Type:' section, there are two radio buttons: 'Per Filer' (which is selected) and 'Per Server'. At the bottom of the window, there is a progress bar labeled 'InstallShield' and three buttons: '< Back', 'Next >' (which is highlighted), and 'Cancel'.

Figure 6-53 Customer Information window

4. In the Destination Folder window (Figure 6-54), verify the path where the source files for SME will be installed. If you need to change this path, just click **Change** and select a new path. Click **Next**.

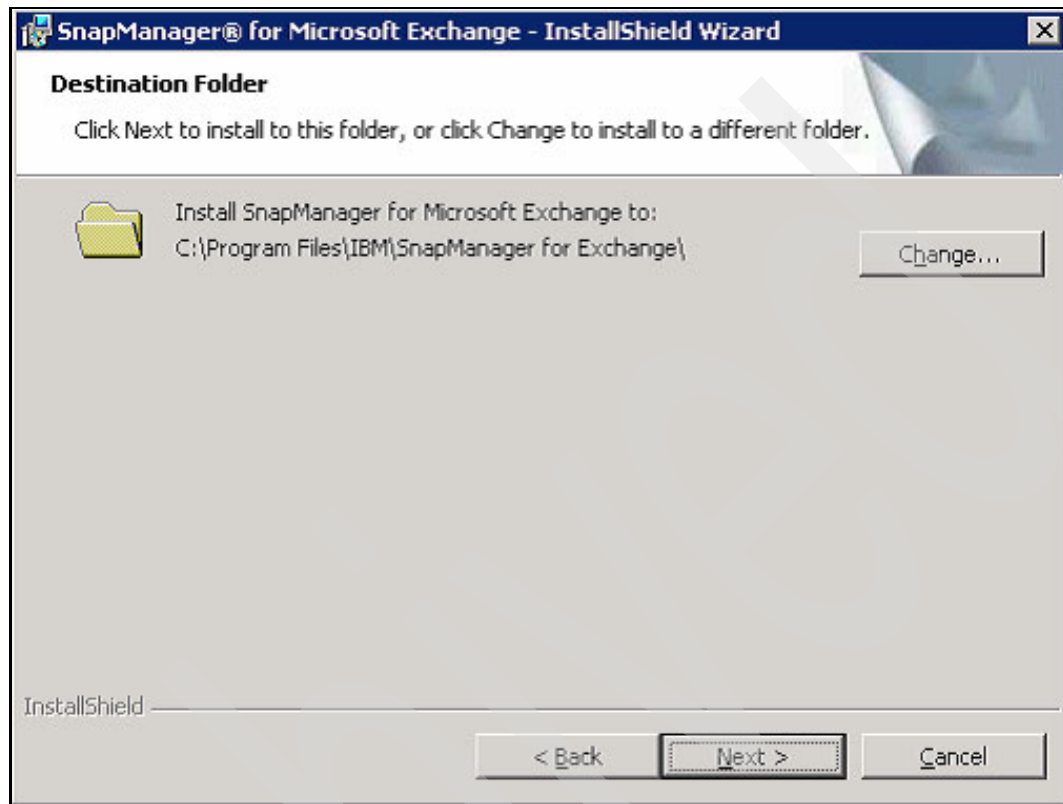


Figure 6-54 Destination Folder window

5. In the SnapManager Server Identity window (Figure 6-55), type in the User Account and Password that will be used to start the SME service. The recommended configuration is to use the same SnapDrive service account on the SME configuration, because this user account already has the permissions set on the IBM System Storage N series storage system and on the Microsoft Exchange Server. Click **Next**.

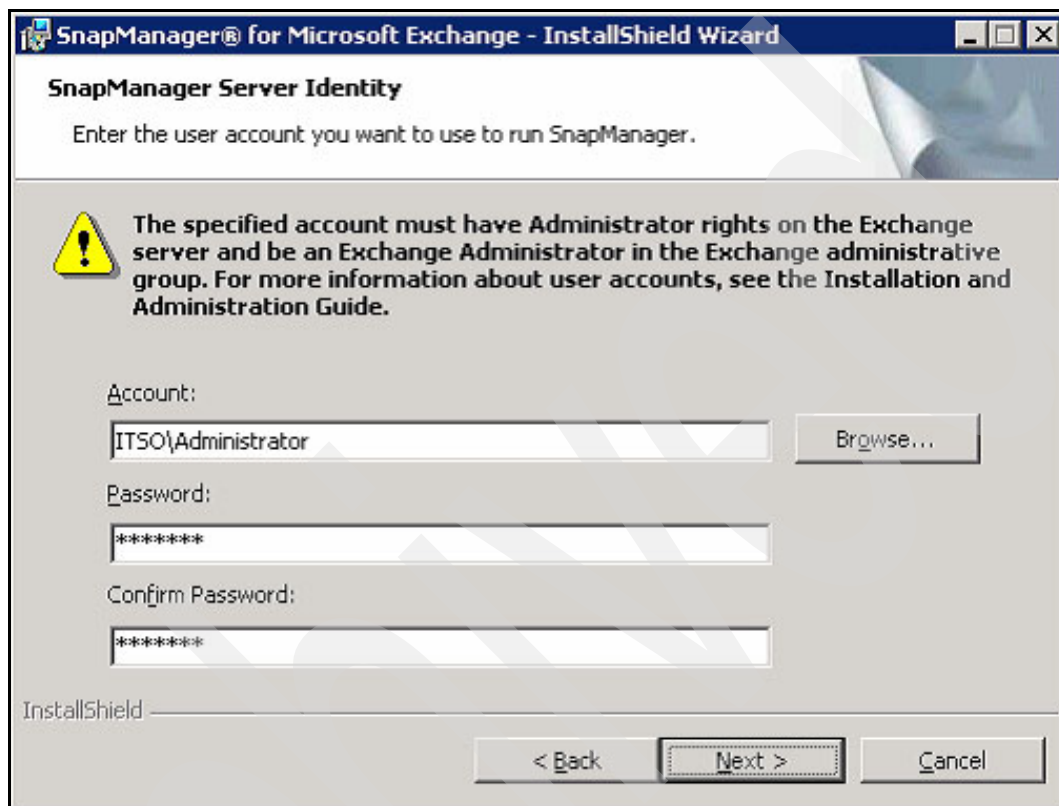


Figure 6-55 SnapManager Server Identity window

6. In the Ready to Install window (Figure 6-56 on page 223), click **Install** to start the installation process. After the installation, click **Finish**. No reboot is needed.

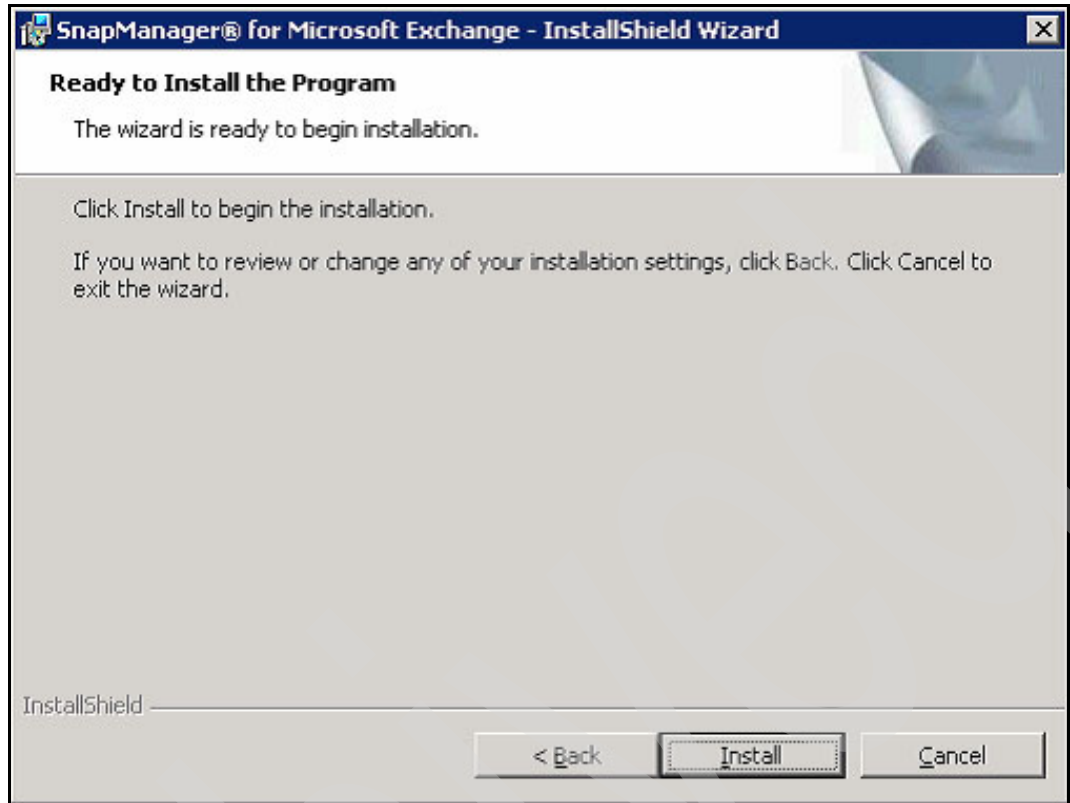


Figure 6-56 Ready to Install window

Configuring SnapManager for Microsoft Exchange

After installing the SnapManager for Microsoft Exchange, start the application by selecting **Start → All Programs → IBM → SnapManager for Exchange Management Console**.

Because this is the first time the MMC is being opened, it has to be customized and the SnapManager for Microsoft Exchange has to be configured on the Microsoft Exchange server. The following steps should be used to configure the SME upon first use:

1. The information window shown in Figure 6-57 appears. Click **OK**.




Figure 6-57 Information window

2. A Microsoft Exchange Server should be added to the Management Console. At this time, SME will search for the Microsoft Exchange Servers on the Active Directory Infrastructure and will bring you a list of them, as shown in Figure 6-58. Double-click the server's name in the list and click **Add**. Alternatively, you can click **Browse** and search for your server on the Network Places tree. After selecting the server, click **Add**.

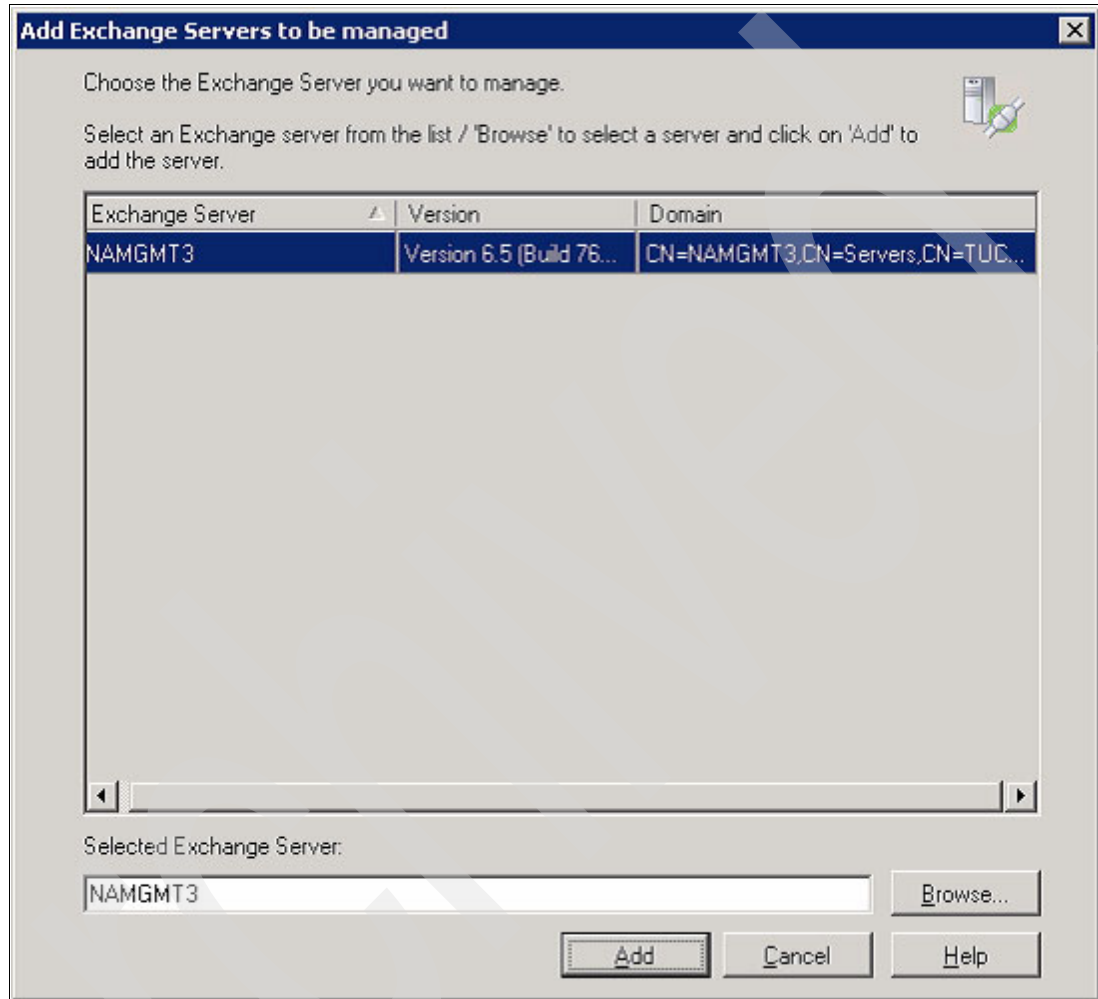


Figure 6-58 Add Exchange Servers window

3. If this is the first time the server is being added to the SME, a warning window, shown in Figure 6-59, appears indicating that the Configuration Wizard should be run prior to accessing the server. Click **OK**.




Figure 6-59 Warning window

4. The Configuration Wizard automatically starts and the Welcome window (Figure 6-60) is shown. Click **Next**.



Figure 6-60 Configuration Wizard welcome window

5. In the Verification Server Settings window (Figure 6-61), type in the name of the server that will be used as the verification server for the backed-up databases. The Microsoft Exchange server name will be filled in automatically. If you want to change the verification server type in the server's name or click **Browse** to select it from the Network Places tree. Note that if the server being designated as the verification server is not a Microsoft Exchange server, you need to copy the Eseutil.exe file to the server and then specify the full path to the file in the Verification Setting window. Because the server selected in Figure 6-61 is the Microsoft Exchange server, the option is grayed out. Click **Next**.



Figure 6-61 Verification Server Settings window

6. The full path to the Eseutil.exe file will be automatically discovered by SME, as shown in Figure 6-62, because the server is a Microsoft Exchange server. Click **OK**.




Figure 6-62 Eseutil path window

7. In the Microsoft Exchange Server to configure window (Figure 6-63), confirm the name of the Microsoft Exchange server that you are configuring and click **Next**.

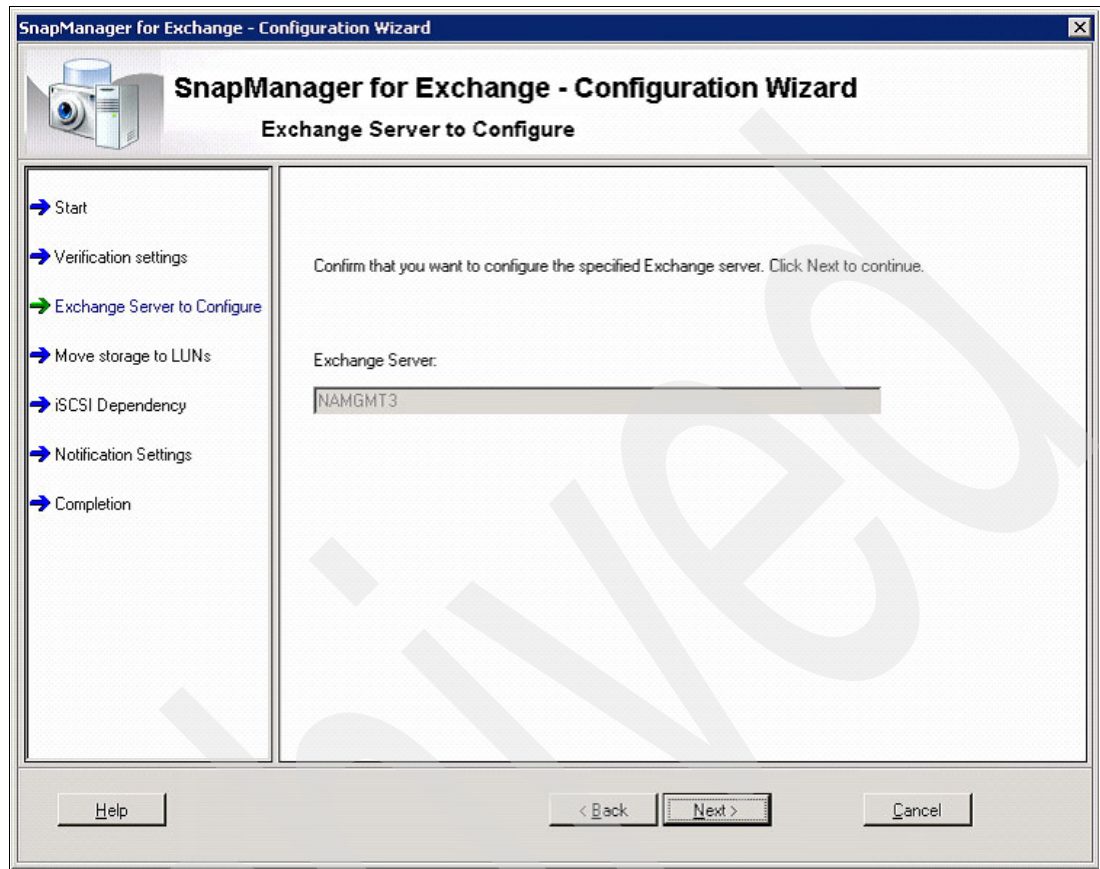


Figure 6-63 Exchange Server to configure window

8. The next step in the SME configuration is to move the Storage Groups to the IBM System Storage N series storage system's LUNs. The first window will show the databases on the Storage Groups: Mailbox Store and Public Folder Store (see Figure 6-64). If the Microsoft Exchange Server has more than one Storage Group, all of them will be shown. Expand the Storage Group container and select the database you want to move to the LUN. The whole Storage Group can also be selected. In the Available Disks window, select the LUN where you want the databases to be moved to and click the arrow button.

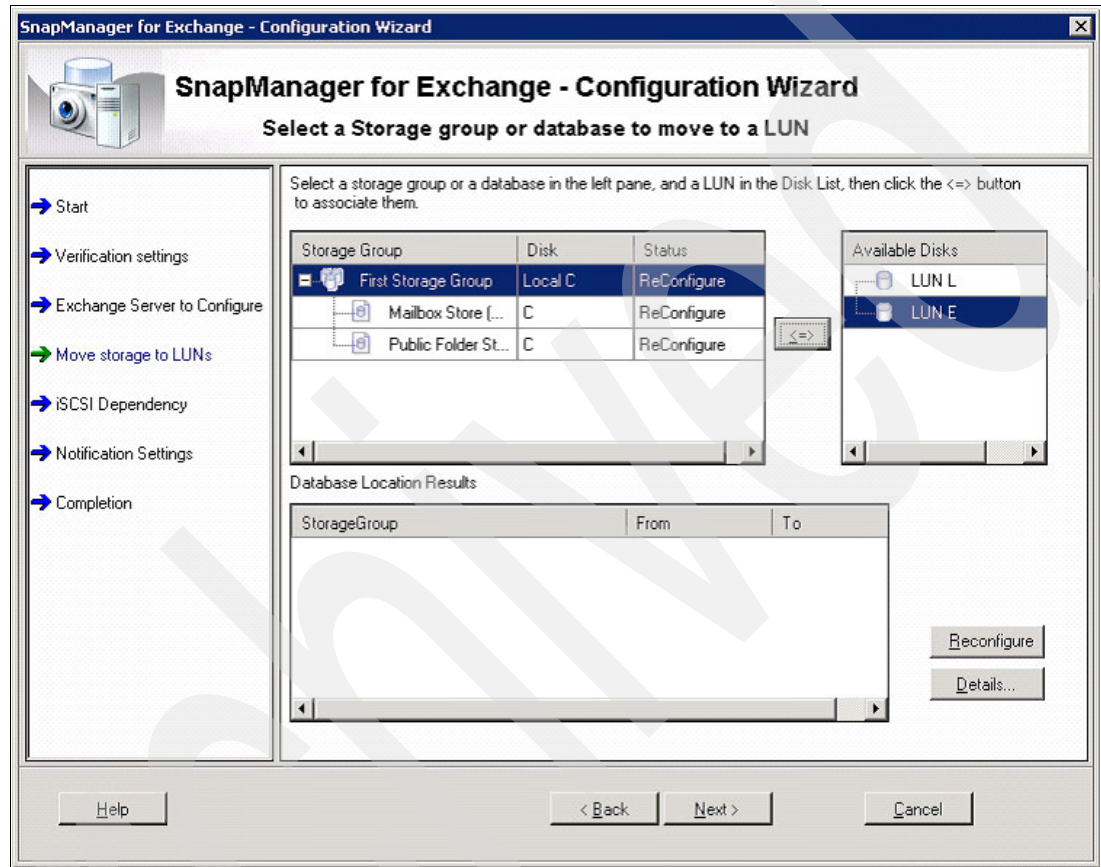


Figure 6-64 Database move window

9. The database will be marked to be moved and the result will be shown in the Database Location Results window (see Figure 6-65). Click **Next**.

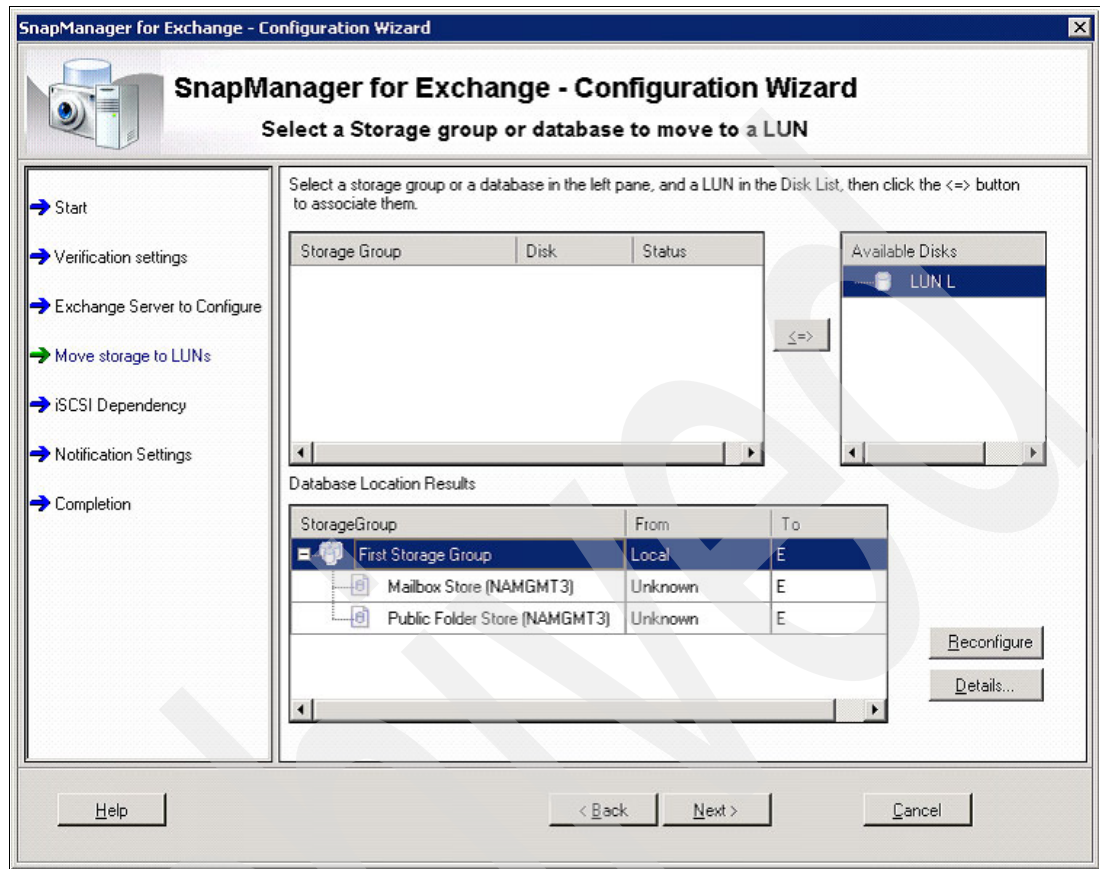


Figure 6-65 Database move window

10. The next step is to move the transaction log files to the IBM System Storage N series storage system's LUNs. Considering that you have only one Storage Group, only one set of transaction log files will be shown, as shown in Figure 6-66. Select the set of transaction log files to be moved to the left pane. In the Available Disks window, select the LUN where you want the transaction log files to be moved to and click the arrow button.

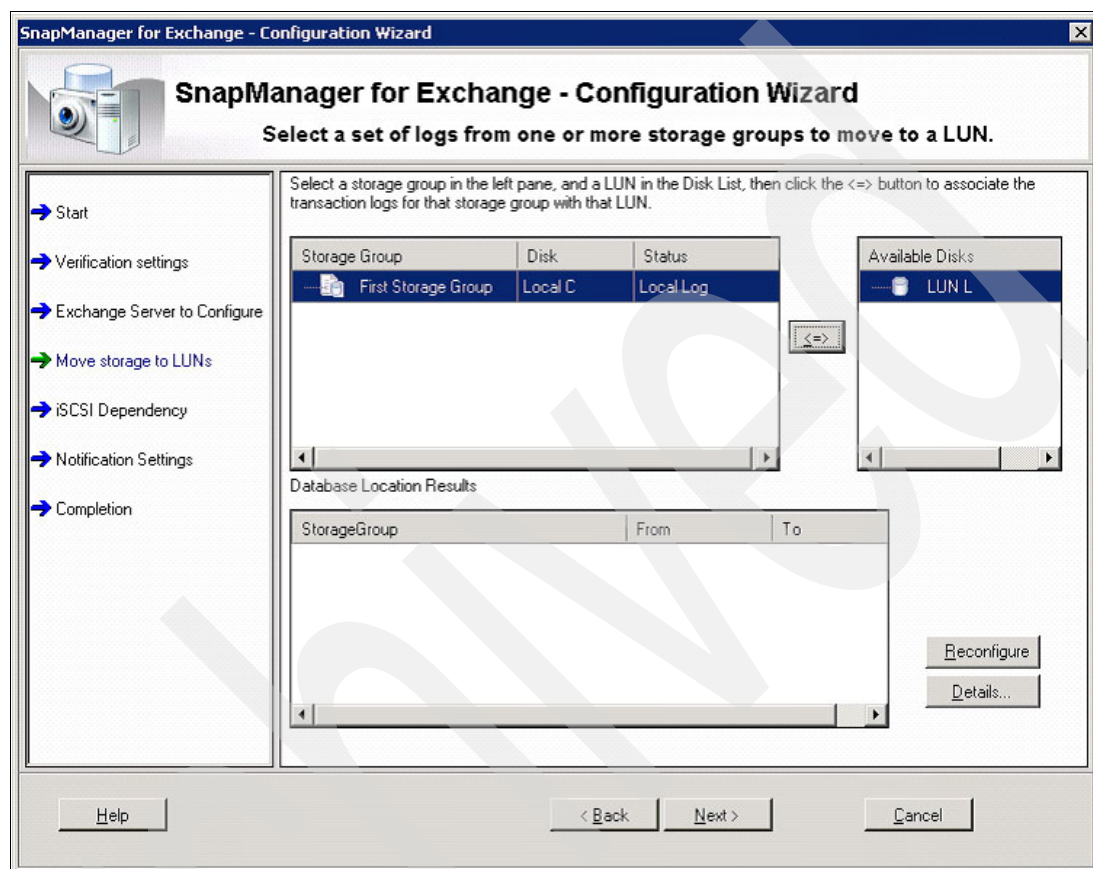


Figure 6-66 Transaction log files move window

11. The set of transaction log files will be marked to be moved and the results will be shown in the Database Location Results (see Figure 6-67). Click **Next**.

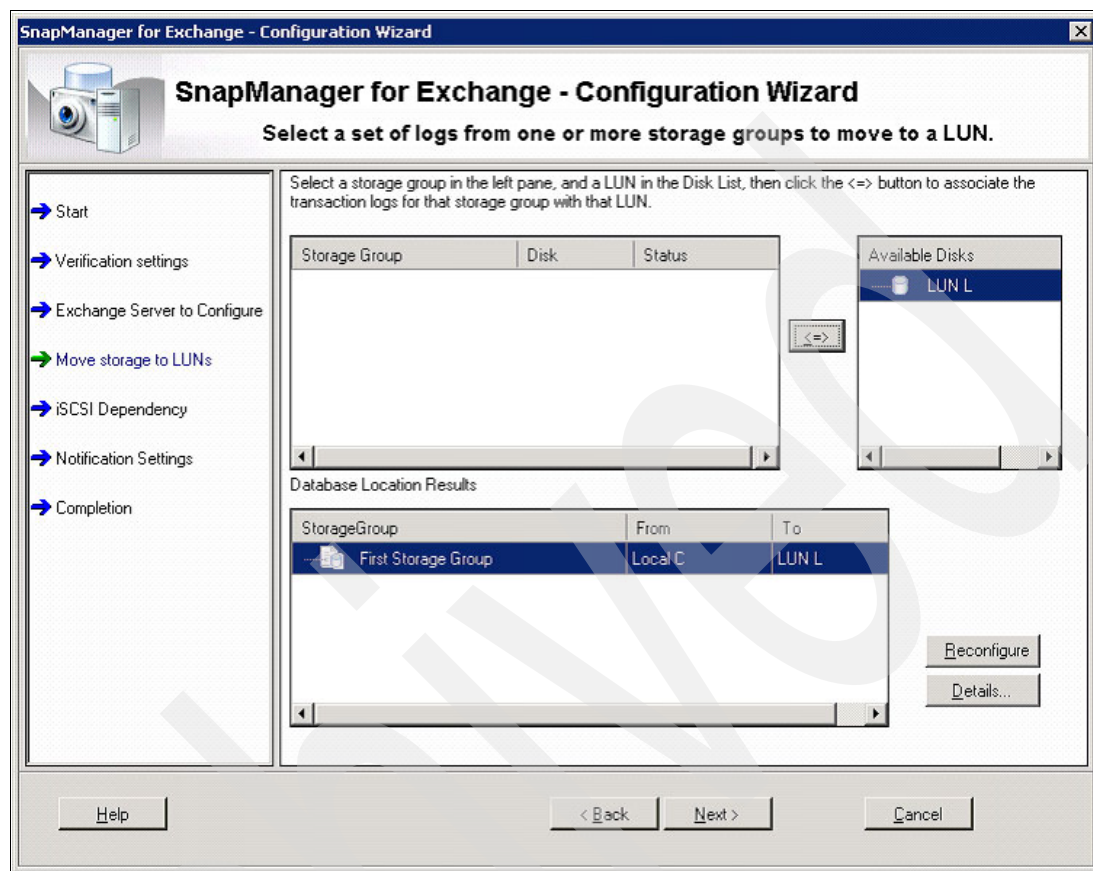


Figure 6-67 Transaction log files move results window

12. The next step is to move the SMTP and MTA queues. Depending on your planned environment these queues will or will not be moved to a LUN on the IBM N series storage system. If you planned to move them to the external LUN, notice that these queues should not be placed on the same LUN as the database files. By default, the SMTP and MTA queues will be preserved on the local disks, as shown in Figure 6-68. If you want to change the location of the external LUN on the IBM System Storage N series storage system, click **Reconfigure** and set the location accordingly. Notice that on our scenario, only the Transaction Log files' LUN is available to select. Click **Next** when you are done with the SMTP and MTA queues location configuration.

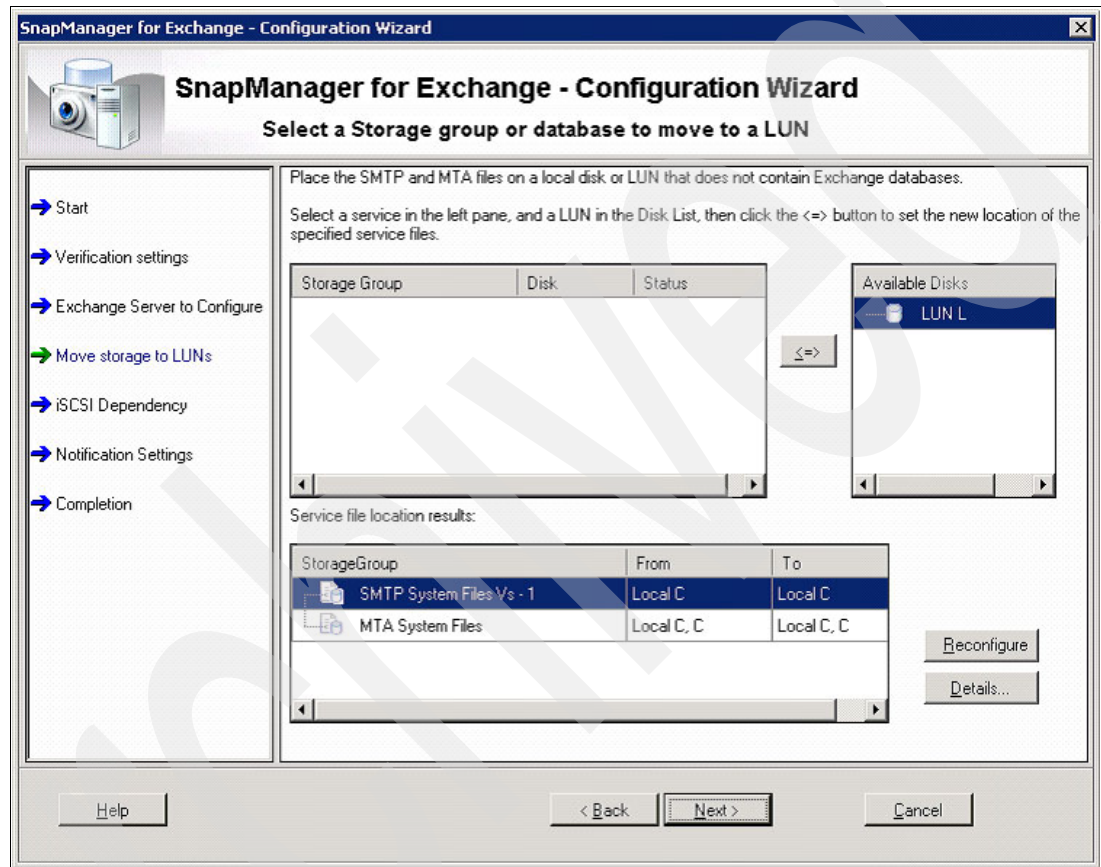


Figure 6-68 SMTP and MTA queues move results window

13. The next step is to configure the SnapInfo Directory location. By default and as recommended, SME configures the SnapInfo Directory to be on the same LUN as the Transaction Log files (see Figure 6-69). If for any reason you want to change the SnapInfo Directory location, click **Reconfigure** and select the new destination LUN. Click **Next**.

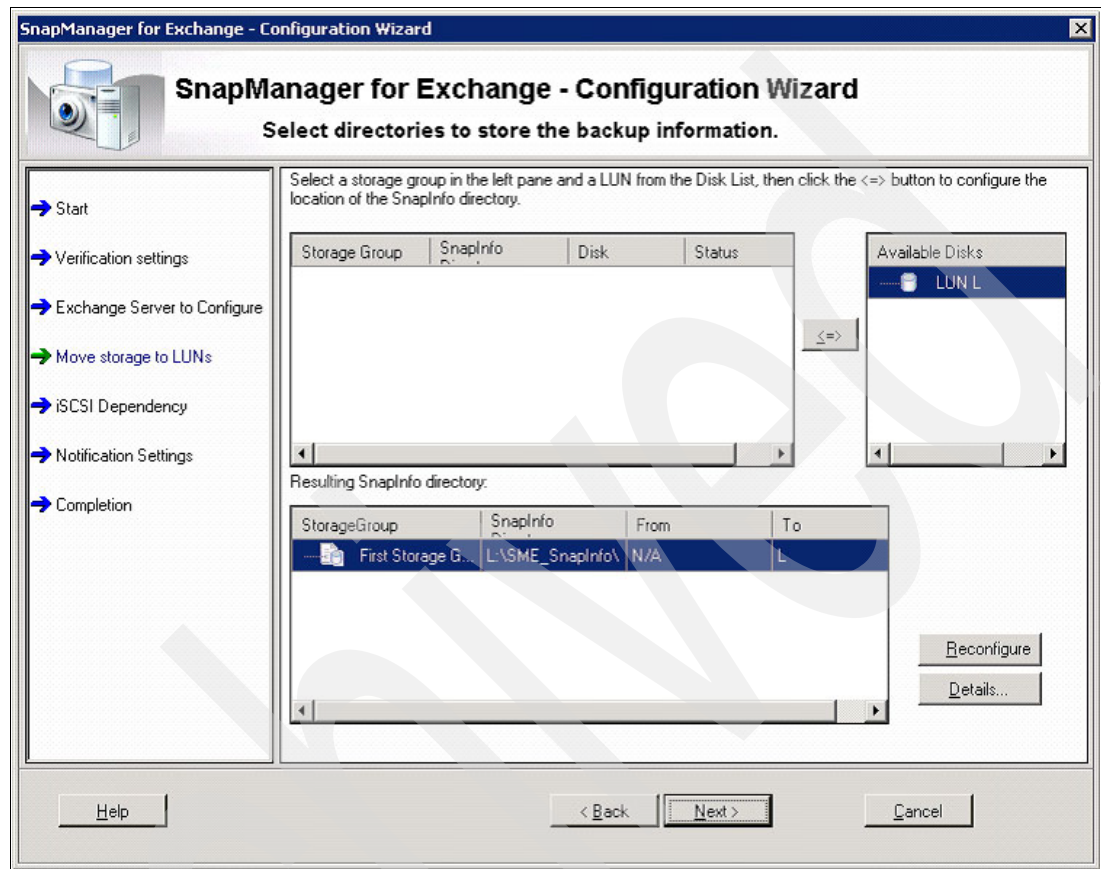


Figure 6-69 SnapInfo Directory window

14. The window seen in Figure 6-70 will appear only if SME detects that the Microsoft iSCSI Software Initiator is installed on the Microsoft Exchange server. This will not be shown for an FCP connected server. Accept the default to make the Microsoft Exchange System Attendant service dependent on the Microsoft iSCSI Initiator service. This configuration prevents the Microsoft Exchange services from trying to start up before the LUNs (through the connection made by the Microsoft iSCSI Software Initiator) are available. Click **Next**.

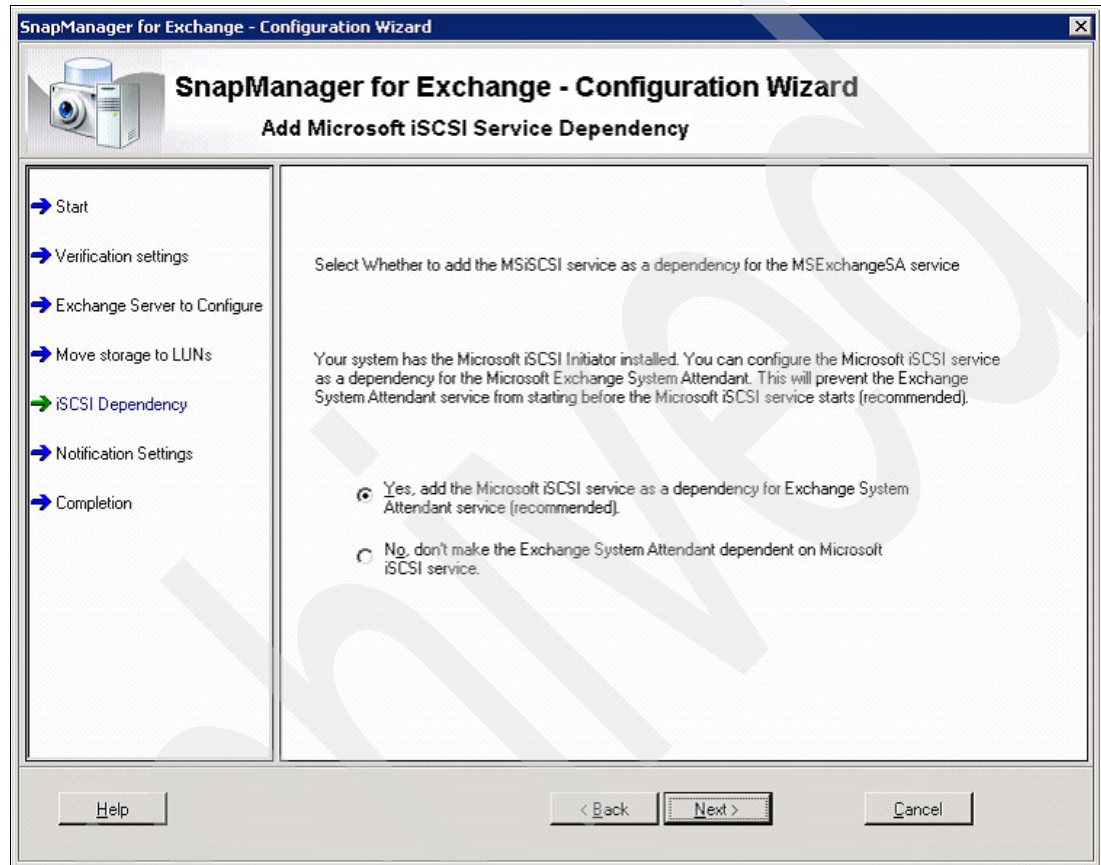


Figure 6-70 iSCSI Dependency window

15. The next window will show you the notification configuration for SME, as shown in Figure 6-71. SME can be configured to send mail notifications to administrators. After you have configured it as needed, click **Next**.

The image shows a Windows-style configuration wizard window titled "SnapManager for Exchange - Configuration Wizard". The window has a sidebar on the left with a list of steps: Start, Verification settings, Exchange Server to Configure, Move storage to LUNs, iSCSI Dependency, Notification Settings (highlighted with a green arrow), and Completion. The main area is titled "Configure Automatic Event Notification (Recommended)" and contains the following options:

- ☐ Send e-mail notification
 - SMTP Server: [text box] [Advanced...]
 - From: [text box with value "SMEAutoSender"] [Send a Test E-mail]
 - To: [text box]
 - Subject: [text box with value "SnapManager for Exchange"]
 - Note: Text will be added to default subject string.
- ☒ Log SnapManager events to storage system syslog
- ☒ Send AutoSupport Notification
 - ☒ On failure only
 - If AutoSupport is enabled on the storage system, SnapManager can alert of SnapManager errors and events.

At the bottom of the window are buttons for "Help", "< Back", "Next >", and "Cancel".

Figure 6-71 Event notification window

16. The Summary window will appear (Figure 6-72). Review the configuration and click **Finish**.

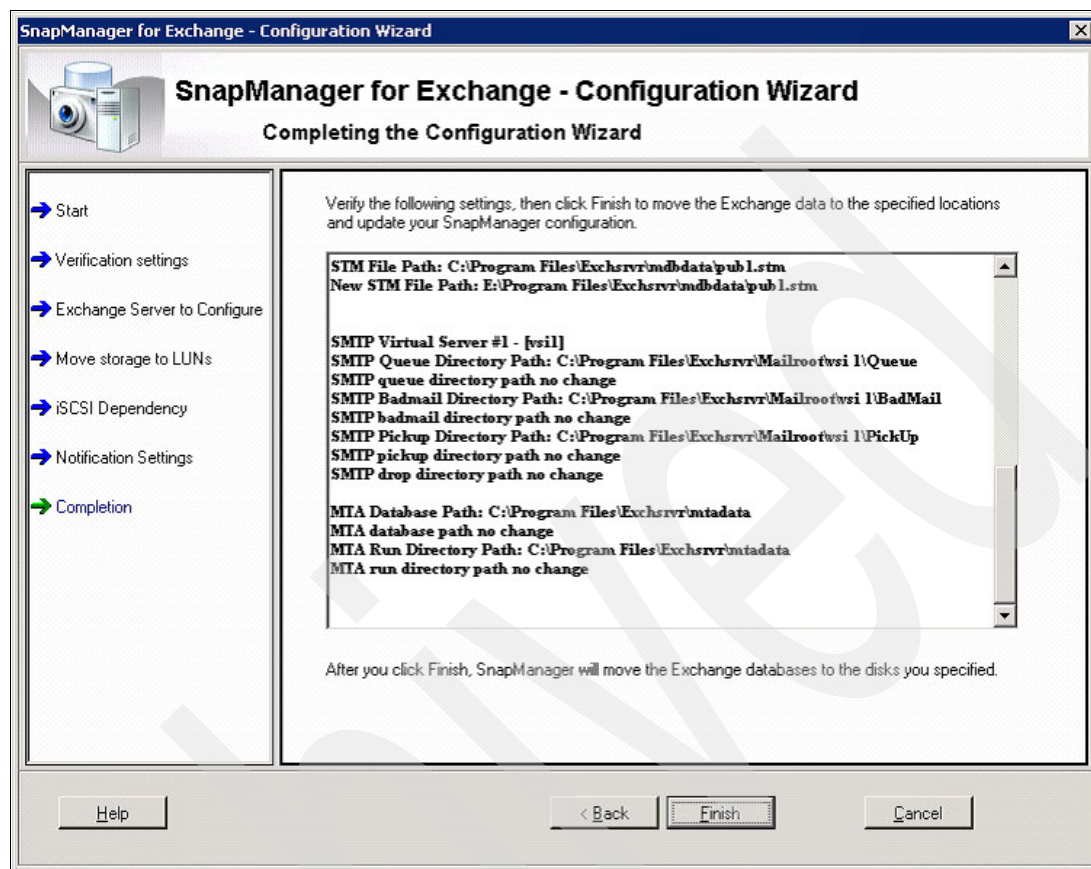


Figure 6-72 Summary window

17. The Configuration Status window will be shown (Figure 6-73) and the tasks will be listed on the Configuration Task List tab. Click **Start Now**.

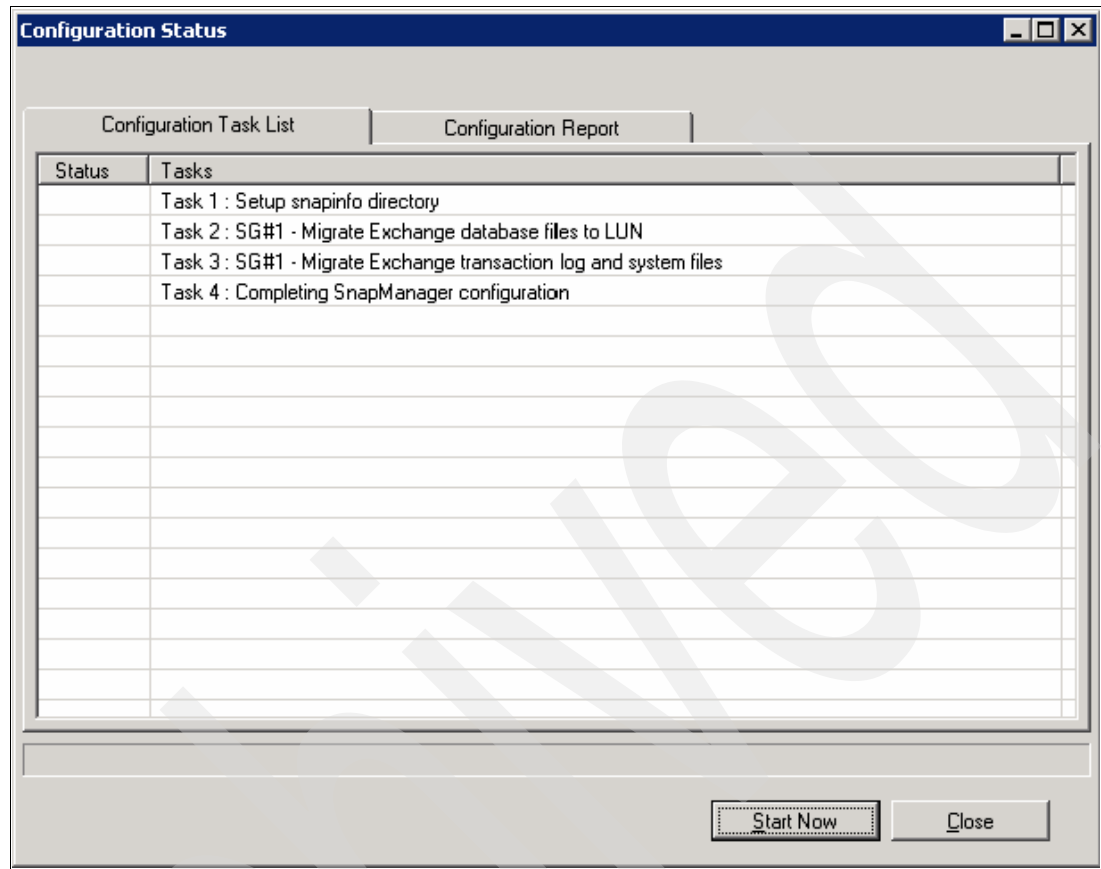


Figure 6-73 Configuration Status window

18. Because we configured the databases to be moved to the LUN on the IBM System Storage N series storage system, a Warning window (Figure 6-74) will be shown that advises you that the Exchange databases will be taken offline to be moved. Click **Yes** to get these databases taken offline and moved to the LUNs.




Figure 6-74 Warning window

19. After the tasks are completed, an Information window will appear stating the tasks completed successfully (Figure 6-75). Click **OK**.




Figure 6-75 Information window

20. The Configuration Report window (Figure 6-76) will be shown. You can review the log entries and click **Close**.

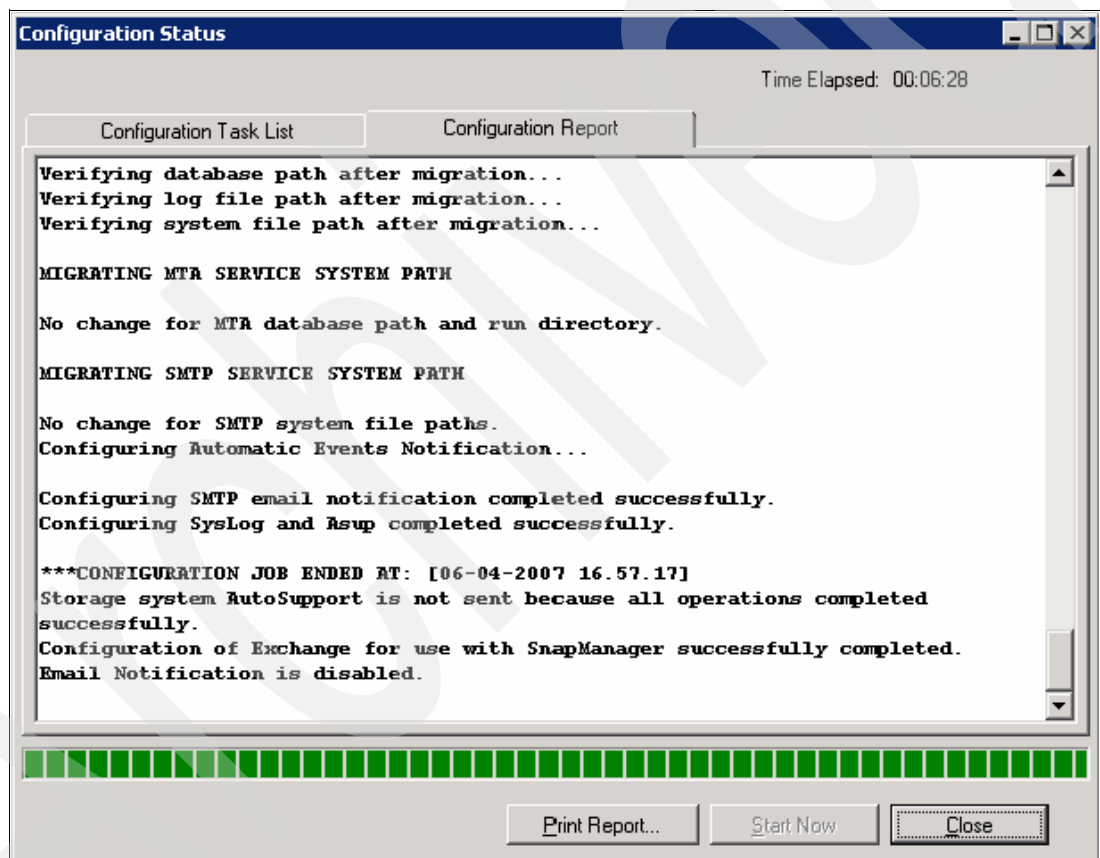


Figure 6-76 Configuration Report window

Installing Lotus Domino on IBM System Storage N series

This chapter provides the steps to access the IBM System Storage N series storage system from the IBM Lotus Domino server as well as how to install and configure the necessary features and software on the server.

7.1 IBM System Storage N series protocol setup

After planning and implementing the necessary features on the IBM System Storage N series storage system as well as creating the aggregates and volumes, it is time to access the IBM System Storage N series storage system and create the LUNs from the host client, which is the IBM Lotus Domino server.

The first step is to install and configure the protocol that will be used for this access. At this point, you have already planned for the protocol you are going to use and have the infrastructure in place for using either iSCSI, FCP, or CIFS. There are some configurations that need to be done on the IBM System Storage N series storage system and on the IBM Lotus Domino server as well.

In this book, we discuss the FCP connection between the Linux and UNIX IBM Lotus Domino servers and IBM System Storage N series.

Note: Because of its improved performance and reliability, the recommended protocol is FCP.

7.1.1 Fibre Channel Protocol (FCP)

Most companies now have a Fibre Channel infrastructure already in place. This eases the installation and configuration of FCP on the IBM Lotus Domino server. Also, these companies already have the knowledge for troubleshooting issues related to FCP.

The recommended configuration when using FCP is to have multiple paths configured. In this manner, should any of the Host Bus Adapters (HBAs) or any of the Fibre Channel cords or any of the FC switches fail, you still have connectivity between the host and the IBM System Storage N series storage system, as shown in Figure 7-1.

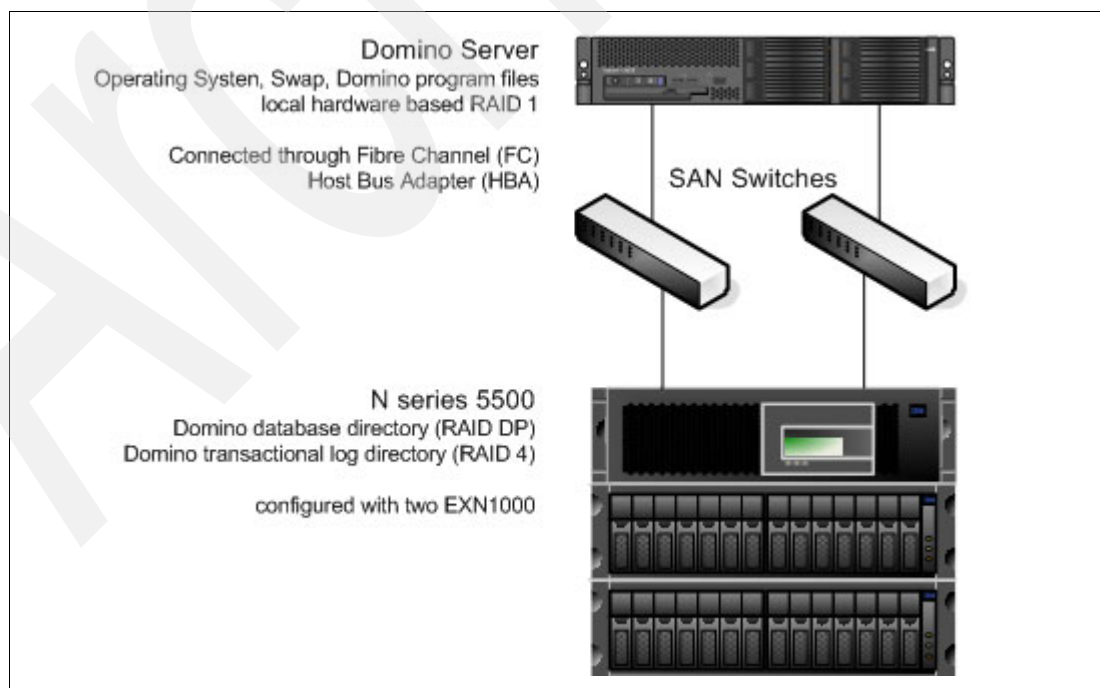


Figure 7-1 Multi-path configuration for IBM Lotus Domino server using FCP with IBM System Storage N series

The steps for connecting the IBM Lotus Domino server with the IBM System Storage N series are:

1. License and start the FCP service on IBM System Storage N series (see Example 7-1).

Note: To enable the FCP protocol and the FCP adapter on the IBM System Storage N series storage system, you will need an FCP license to be installed on the IBM System Storage N series storage system.

Example 7-1 Add FCP license and start the service

```
itsotuc3*> license add XXXXXXXX
A fcp site license has been installed.
Run 'fcp start' to start the FCP service.
Also run 'lun setup' if necessary to configure LUNs.
      FCP enabled.
itsotuc3*> fcp start
itsotuc3*> fcp status
FCP service is running.
itsotuc3*>
```

2. Configure Fibre Channel switches, zoning, and cabling.

For details, refer to 4.3, “Zoning” on page 99.

3. Check the SAN inter operability matrix for:

- Operating system version
- Host bus adapter (HBA) model
- HBA firmware and version
- Fibre Channel switch model and firmware version
- IBM System Storage N series Storage system and Data ONTAP version

You can find the interoperability matrix at:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

4. Install and configure the HBAs and software on the host.

For details, refer to 7.3, “Host connection” on page 245.

5. Configure the IBM System Storage N series storage system.

6. Discovering and accessing LUNs from the host.

We recommend using SnapDrive for creating, discovering, and accessing LUNs. See 7.4, “SnapDrive installation and usage” on page 257 for further information.

Note: Despite the fact that the LUNs can be created from the IBM System Storage N series storage system, the recommended procedure is to create the LUNs from the Lotus Domino server using SnapDrive.

7.2 Role-based access control of the IBM System Storage N series

Role-based access controls, or RBAC, are a method for managing the set of actions that a user or administrator may perform in a computing environment.

While reserving certain functions for administrator-only access is a good start, additional problems need to be solved. Most organizations have multiple system administrators, some of whom require more privileges than others. By selectively granting or revoking privileges for each user, you can customize the degree of access that an administrator has to the system.

We use the role-based access control feature for creating a separate user account that belongs to the SnapDrive feature on the Linux and Unix Lotus Domino server. This account has specific rights to create, modify, and delete LUNs and Snapshots.

Users are members of groups that have one or more roles, and each role grants a set of capabilities. In this way, Data ONTAP allows you to create flexible security policies that match your organizational needs. All configuration for role-based access controls occurs through the **useradmin** command provided by Data ONTAP (see Example 7-2).

Example 7-2 Data ONTAP useradmin command

```
itsotuc3> useradmin
```

Usage:

```
useradmin user  add <login_name> [options]
                modify <login_name> [options]
                delete <login_name>
                list [options]
group          add <group_name> [options]
                modify <group_name> [options]
                delete <group_name> [options]
                list [options]
role           add <role_name> [options]
                modify <role_name> [options]
                delete <role_name> [options]
                list [options]
domainuser    add <user_name> [options]
                delete [options]
                list [options]
                load <filename>
```

For more detailed information about each subcommand, use:

```
useradmin help { user | group | role | domainuser }
itsotuc3>
```

We use the **useradmin user add** or **useradmin user modify** commands for modifying or adding users.

Note: For more information about accessing the IBM System Storage N series storage system and role-based access control of Data ONTAP, refer to the *IBM System Storage N series Data ONTAP System Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001690&aid=1>

7.2.1 Creating a new role

A role is defined as a named set of capabilities. Data ONTAP comes with several roles predefined, and users may create additional roles or modify the provided roles.

When creating new groups, Data ONTAP requires specification of the roles the new groups should have.

Therefore, it is best to create the appropriate roles before defining the groups. Follow these steps for creating a new role:

1. Log in to your IBM System Storage N series with SSH (you have to start it first) or Telnet.
2. Create a new role with the **useradmin** command, as shown in Example 7-3. We use the role name **rl_snapdriveadmin** in our case.

Example 7-3 Creating a new role

```
itsotuc3*> useradmin role add rl_snapdriveadmin -c "SnapDrive Admin Role" -a
api-*,login-http-admin
Role <rl_snapdriveadmin> added.
```

```
itsotuc3*>
```

- The option -c follows a comment.
- The option -a follows granted privileges. For SnapDrive, the api-* and login-http-admin privileges are needed.

Note: If you use HTTPS, grant the privilege login-https-admin instead of login-http-admin.

3. Verify the created role, as shown in Example 7-4.

Example 7-4 Verify the created role.

```
itsotuc3*> useradmin role list rl_snapdriveadmin
Name:      rl_snapdriveadmin
Info:
Allowed Capabilities: api-*,login-http-admin
```

```
itsotuc3*>
```

Now it is time to create a group and map the role to it.

7.2.2 Creating a new group

A group is defined as a collection of users or domain users. Groups may be assigned one or more roles. It is important to remember that the groups defined within Data ONTAP are separate from the groups defined in other contexts, such as a Microsoft Active Directory server. This is true even if the groups within Data ONTAP have the same names as groups elsewhere within your environment.

When creating new users or domain users, Data ONTAP requires specification of group membership. Therefore, it is best to create appropriate groups before defining users or domain users. Follow these steps to create a new group:

1. Log in to your IBM System Storage N series with SSH (you have to start it first) or Telnet.

2. Create a new role with the **useradmin** command, as shown in Example 7-5. We use the group name **grp_snapadmins** in our case.

Example 7-5 Creating a new group

```
itsotuc3*> useradmin group add grp_snapadmins -c "This group is for SnapDrive  
usage only" -r rl_snapdriveadmin  
Group <grp_snapadmins> added.
```

```
itsotuc3*>
```

- The option **-c** follows a comment.
- The option **-r** maps the role **snapdriveadmin** after the creation process.

3. Verify the created group (Example 7-6).

Example 7-6 Verify the created group

```
itsotuc3*> useradmin group list grp_snapadmins  
Name: grp_snapadmins  
Info: This group is for SnapDrive usage only  
Rid: 131072  
Roles: rl_snapdriveadmin  
Allowed Capabilities: api-*,login-http-admin
```

```
itsotuc3*>
```

Tip: If you need to modify the group and role mapping later, use the **useradmin group modify -r** command.

7.2.3 Creating a new user

A user is defined as an account that is authenticated on the IBM System Storage N series storage system. We use a separate user account for SnapDrive usage on Linux and UNIX.

Follow these steps to create a new user:

1. Log in to your IBM System Storage N series with SSH (you have to start it first) or Telnet.
2. Create a new user with the **useradmin** command, as shown in Example 7-7. We use the user name **notes** in our case for the Lotus Domino user on Linux and UNIX.

Note: The password must have at least eight characters.

Example 7-7 Creating the user account

```
itsotuc3*> useradmin user add notes -c "Lotus Domino SnapDrive user" -n "Lotus  
Domino" -g grp_snapadmins  
New password:  
Retype new password:  
User <notes> added.
```

```
itsotuc3*>
```

- The option **-c** follows a comment.
- The option **-n** is the full name of the user.

- The option -g follows the group name to which the user belongs.

3. Verify the created user (Example 7-8).

Example 7-8 Verify the created user

```
itsotuc3*> useradmin user list notes
Name: notes
Info: Lotus Domino SnapDrive user
Rid: 131074
Groups: grp_snapadmins
Full Name: Lotus Domino
Allowed Capabilities: api-*,login-http-admin
Password min/max age in days: 0/4294967295
Status: enabled

itsotuc3*>
```

Note: The command **useradmin user** will give you a list of possible options for user management on the IBM System Storage N series Data ONTAP.

7.3 Host connection

This section discusses the LUN mapping through FCP for Red Hat Linux and AIX. Check the following requirements first:

- ▶ License and start the FCP service on IBM System Storage N series.
- ▶ Configure the zoning.
- ▶ Check the SAN interoperability matrix for:
 - Operating system version
 - Host bus adapter (HBA) model
 - HBA firmware and version
 - Fibre Channel switch model and firmware version
 - IBM System Storage N series storage system and Data ONTAP version

You can find the interoperability matrix at:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

The procedure for configuring Fibre Channel protocol on a host and mapping the LUN follows the same basic sequence for all hosts with Linux and AIX:

1. Install and configure HBAs and software on host.
2. Configure the IBM System Storage N series storage system.
3. Configure Fibre Channel switches and cabling.
4. Discovering and accessing LUNs from the host.

Note: For companies who would like to use Microsoft Windows with FCP or iSCSI for Lotus Domino, the operating system configuration is analogous to the Exchange installation section. Refer to 6.1, “Accessing the IBM System Storage N series storage system” on page 170 for detailed instructions.

7.3.1 Configuring a Linux Host for FCP

We recommend that you download and install the SAN (FCP) host attach kit for Linux provided by IBM. This product simplifies the configuration of a Linux machine as the host component of an IBM System Storage N series SAN environment. At this time, it supports the following Linux OS versions:

- ▶ Red Hat Enterprise Linux (RHEL) Version 4 update 2 and RHEL 4 update 3
- ▶ SUSE Linux Enterprise Server (SLES) Version 9 SP2 and SLES 9 SP 3

Note: New configuration components are also qualified between support kit releases. For the latest support information, see the appropriate interoperability matrix for your IBM System Storage N series product, available on the IBM support web site at:

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

If you are using the dm-multipath support for RHEL, you need the following userspace RPM versions:

1. device-mapper RPM: 1.02.02-3.0.RHEL4
2. device-mapper-multipath RPM: 0.4.5-12.0.RHEL4

These are the steps to install and configure the FCP Linux Host Attach Kit:

Note: For more information about the FCP Linux Host Attach Kit, see:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskind=3&brandind=5000029&familyind=5329809>

1. Install the IBM host Attach kit. In order to install this product, you need to complete the following basic steps:
 - a. IBM System Storage N series FCP Linux Host Attach Kit 3.0 binaries are provided to licensed customers only and can be obtained by contacting IBM support.
 - b. After obtaining and copying the downloaded attach kit, tar the file to a directory (/tmp/nseries) on the Lotus Domino server and extract the software by executing the following command:

```
tar -xvzf /tmp/nseries/ibm_linux_host_utils_3_0.tar.gz
```

The extracted software is placed in the /tmp/nseries/ibm_linux_host_utils_3_0 directory.
 - c. Change the directory to ibm_linux_host_utils_3_0, where you have extracted the attach kit software, and execute the `./install` command, as shown in Example 7-9.

Example 7-9 Installing the FCP host attach kit for Linux

```
[root@domino1 ibm_linux_host_utils_3_0]# ./install
opt/ontap/
opt/ontap/man/
opt/ontap/man/man1/
opt/ontap/man/man1/brocade_info.1
opt/ontap/man/man1/cisco_info.1
opt/ontap/man/man1/filer_info.1
opt/ontap/man/man1/linux_info.1
opt/ontap/man/man1/mcdata_info.1
```



```
opt/ontap/man/man1/qla2xxx_lun_rescan.1
opt/ontap/man/man1/sanlun.1
opt/ontap/santools/
opt/ontap/santools/sanlun
opt/ontap/santools/brocade_info
opt/ontap/santools/cisco_info
opt/ontap/santools/filer_info
opt/ontap/santools/linux_info
opt/ontap/santools/mcdata_info
opt/ontap/santools/uninstall
opt/ontap/santools/SHsupport.pm
opt/ontap/santools/san_version
opt/ontap/santools/Telnet.pm
opt/ontap/santools/qla2xxx_lun_rescan
opt/ontap/santools/mpath_prio_ontap
opt/ontap/santools/libHBAAPI.so
[root@domino1 ibm_linux_host_utils_3_0]#
```

- d. Add the following lines to your `/.bash_profile` file:

```
export PATH=$PATH:/opt/ontap/santools
```

For detailed installation steps, please refer to *Installation and Setup Guide for Fibre Channel Protocol on Linux* on the IBM support Web site at:

<http://www-03.ibm.com/servers/storage/support/nas/index.html>

2. Install the HBA and driver.

Note: For Linux, you might need development tools such as `gcc` to compile the newest driver for your system.

Before you install the HBA on the Lotus Domino server, you need to check for product compatibility on the IBM Web site. The compatibility matrix is available on the IBM support site. After confirming compatibility, install the HBA and its driver on the database server. For the HBA and its driver installation steps, please refer to product's installation and configuration guide on IBM support site, found at:

<http://www-03.ibm.com/servers/storage/support/nas/index.html>

3. Obtain the WWPN for the initiator.

Each initiator installed on the database server is uniquely identified by WWPN or WWNN. The WWPN is required to create an igroup for FC on the IBM System Storage N series storage system. The WWPN for the HBA installed on the database server can be obtained by completing the following steps:

- a. After installing the IBM System Storage N series attach kit, the HBA, and the required driver, you would execute the following command on the database server to obtain the WWPN:

```
sanlun fcp show adapter -c
```

Upon execution with `-c` option, the `sanlun` command generates the `igroup create` command. The output of the above command looks somewhat similar to the one below:

```
igroup create -f -t linux "domino1" 210000e08b07f0be
```

- b. The WWPN for adapter in the above example is 210000e08b07f0be. Execute the above generated command on the IBM System Storage N series Data ONTAP CLI to create the igroup on the IBM System Storage N series storage system (see Example 7-10).

Example 7-10 Creating the igroup on N series

```
itsotuc3> igroup create -f -t linux "domino1" 210000e08b07f0be
itsotuc3*> igroup show
    domino1 (FCP) (ostype: linux):
        21:00:00:e0:8b:07:f0:be (logged in on: 0a)
itsotuc3*>
```

- c. Map the LUN to the created igroup, as shown in Example 7-11.

Example 7-11 Map the LUN to the igroup

```
itsotuc3*> lun map /vol/vol_DominoDB/lun_DominoDB domino1
lun map: auto-assigned domino1=0
itsotuc3*>
```

- d. Refresh the FC adapter by executing the following commands on your database server:

```
modprobe -r qla2300
modprobe -v qla2300
```

4. Accessing LUNs as a regular file system table space container.

Once the FCP driver is refreshed, the database server should be able to discover target LUNs on the IBM System Storage N series storage system as new devices. In order to use these devices as file system table space containers, complete the following steps:

- a. Get the assigned names for the newly added devices by executing the following command on the database server:

```
sanlun lun show
```

Example 7-12 shows our lab configuration.

Example 7-12 Mapped LUNs

```
[root@domino1 ~]# sanlun lun show
filer: lun-pathname device filename adapter protocol lun size lun state
itsotuc3: /vol/vol_DominoDB/lun_DominoDB /dev/sdc host5 FCP 146.0g (156786229248)
GOOD
[root@domino1 ~]#
```

If your output contains no LUNs (see Example 7-13), there is no LUN mapped with the igroup.

Example 7-13 No LUN available

```
[root@domino1 ~]# sanlun lun show
no filer LUNs available
[root@domino1 ~]#
```

- b. Before you use the newly added devices as the database or transaction log containers, you need to create partition tables and format them by executing the following command on the Lotus Domino server:

```
fdisk devicename
```

For example, to create partition tables and format a device named `/dev/sdc`, you would execute the following command on the Lotus Domino server:

```
fdisk /dev/sdc
```

The **fdisk** command invokes the format wizard. You need to answer a series of questions in order to complete the formatting.

- c. After formatting, you need to create file system on each device by executing the following command on the Lotus Domino server:

```
mkfs devicename
```

For example, to create a file system on the device named `/dev/sdc1`, you would execute the **mkfs.ext3** command on the Lotus Domino server, as shown in Example 7-14.

Example 7-14 File system creation

```
[root@domino1 ~]# mkfs.ext3 /dev/sdc1
mke2fs 1.35 (28-Feb-2004)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
19152896 inodes, 38276862 blocks
1913843 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=41943040
1169 block groups
32768 blocks per group, 32768 fragments per group
16384 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 34 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
[root@domino1 ~]#
```

- d. After adding the file system, the LUN can be mounted as a shared file system by executing the following commands on the database server:

```
mkdir -p dir_name
mount devicename mount_point
```

In order to mount a device named `/dev/sdc1` on a mount point named `/notesdata/db`, you would execute the **mount** command, as shown in Example 7-15.

Example 7-15 Mounting the file system

```
[root@domino1 ~]# mount /dev/sdc1 /notesdata/db
[root@domino1 ~]#
```

- e. In order to put the Lotus Domino databases on the IBM System Storage N series storage system volumes, the database instance owner must have the appropriate permissions on all the mounted volumes to be used for the database. You do so by changing the ownership of each mounted IBM System Storage N series storage system volume to the database instance owner by executing the following command on the Lotus Domino server:

```
chown -R notes:notes mount_point
```

For example, to change the ownership of the N series storage system volume mounted on the mount point named /mnt/db2data, you would execute the following command on the Lotus Domino server:

```
chown -R notes:notes /notesdata/db
```

Note: The user name and group may differ on your installation.

- f. In order to make the file system mount persistent after a system reboot, you need to add an entry for each file system to the /etc/fstab file, which is found on the Lotus Domino server. The entry in /etc/fstab file should appear similar to the one indicated below:

```
<device name> <mountpoint> <FS type> <mount option> <bkup freq> <fsck pass>
```

We use the **echo** command to add a new line (Example 7-16).

Example 7-16 Adding additional line to /etc/fstab

```
[root@domino1 ~]# echo "/dev/sdc1 /notesdata/db ext3 defaults 0 0" >> /etc/fstab
```

The IBM System Storage N series LUN is mapped to the Lotus Domino server and ready for further use. With the IBM System Storage N series SnapDrive feature, most of the steps are done through the SnapDrive service.

Note: Although you can create and map LUNs manually, we highly recommend the usage of the IBM System Storage N series SnapDrive feature.

7.3.2 Configuring an AIX Host for FCP

Here we describe the configuration of FCP to connect an AIX host to an IBM System Storage N series storage server. We recommend that you download and install the N series provided SAN (FCP) Host Attach Kit for AIX on the Lotus Domino server. This product simplifies the configuration and management of the AIX host in an IBM System Storage N series SAN environment.

The FCP AIX Host Utility includes:

- ▶ The FCP AIX host settings software package. It contains FCP device definitions that are required for FCP AIX configurations.
- ▶ The Data ONTAP SAN Toolkit software package. It contains the **san1un** utility and diagnostic scripts. The **san1un** utility displays configuration information about LUNs and HBAs. The diagnostic scripts collect information about your system if a problem occurs. Customer support may ask you to run the **san1un** utility or a diagnostic script if a problem occurs.

You must download both the host settings software package and the SAN Toolkit software package.

1. Install the FCP AIX Host Utilities.

Note: IBM System Storage N series FCP AIX Host Attach Kit binaries are provided to licensed customers only and can be obtained by contacting IBM support.

- ### Example 7-17 Extracting the AIX host setting software package

Note: This installation example installs the host settings software first. You then repeat the steps to install the SAN Toolkit software package.


```
# installp -aXd Ontap.SAN_toolkit Ontap.SAN_toolkit
+-----+
+-----+
Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets
-----
Ontap.SAN_toolkit.sanlun 3.0.0.0      # Data Ontap SAN Toolkit sanlun
Ontap.SAN_toolkit.scripts 3.0.0.0    # Data Ontap SAN Toolkit Scripts

<< End of Success Section >>

FILESET STATISTICS
-----
2 Selected to be installed, of which:
2 Passed pre-installation verification
----
2 Total to be installed

+-----+
+-----+
Installing Software...
+-----+
installp: APPLYING software for:
        Ontap.SAN_toolkit.scripts 3.0.0.0
        Ontap.SAN_toolkit.sanlun 3.0.0.0

. . . . . << Copyright notice for Ontap.SAN_toolkit >> . . . . .
(C) Copyright Network Appliance, Inc. 2003.
All rights reserved.

. . . . . << End of copyright notice for Ontap.SAN_toolkit >>. . . .

Finished processing all filesets. (Total time: 2 secs).

+-----+
+-----+
Summaries:
+-----+

Installation Summary
-----
```

Name	Level	Part	Event	Result
Ontap.SAN_toolkit.scripts	3.0.0.0	USR	APPLY	SUCCESS
Ontap.SAN_toolkit.sanlun	3.0.0.0	USR	APPLY	SUCCESS

```
# sanlun version
sanlun version 3.0 (574942)
#
```

2. Verify the HBA initiator queue depth.

It is a good practice to verify the value of the HBA initiator queue depth. Check the HBA queue depth by using the **lsattr -El** command. We recommend using a value of 48 for the HBA queue depth.

Example 7-20 Checking the queue depth

```
# lsattr -El fcs0
bus_intr_lvl 105      Bus interrupt level      False
bus_io_addr  0x2ec00  Bus I/O address         False
bus_mem_addr 0xec040000 Bus memory address       False
init_link    al       INIT Link flags          True
intr_priority 3        Interrupt priority        False
lg_term_dma  0x800000  Long term DMA             True
max_xfer_size 0x100000 Maximum Transfer Size     True
num_cmd_elems 48      Maximum number of COMMANDS to queue to the adapter True
pref_alpa    0x1       Preferred AL_PA           True
sw_fc_class   2        FC Class for Fabric       True
#
```

3. Obtain the WWPN for the HBA.

Each FC HBA attached to the database server is uniquely identified by a WWPN. In order to create an igroup on the IBM System Storage N series storage system, the WWPN for the HBA is required. In order to obtain the WWPN, complete the following steps:

- a. After installing the Host Attach Kit for AIX, HBA, and the drivers on the database server, you can find the WWPN by executing the following command:

```
sanlun fcp show adapter -c
```

Upon execution with the -c option, the **sanlun** command generates the **igroup create** command, which can be used to create an igroup on the IBM System Storage N series storage system. The output from the above command looks similar to the one shown in Example 7-21.

Example 7-21 Show the FCP Adapter WWPN

```
# sanlun fcp show adapter -c
```

Enter this filer command to create an initiator group for this system:

```
igroup create -f -t aix "aixld2" 10000000c93316f8
```

```
#
```

- b. The WWPN for adapter in the above example is 10000000c93316f8. Use the generated command to create the igroup.
- c. Load the FC driver by executing the following command on the database server:

```
cfgmgr -l fcs0
```

Once the driver is loaded, the database server should be able to see the target LUNs on the IBM System Storage N series storage system as new devices.

When the FCP driver is refreshed, the database server should be able to discover target LUNs on the IBM System Storage N series storage system as new devices. In order to use these devices as file system table space containers, complete the following steps:

1. Check the newly added devices by executing the following command on database server:

```
lsdev -Cc disk
```

The output would look similar to Example 7-22.

Example 7-22 lsdev command

```
# lsdev -Cc disk
hdisk0 Available 1S-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1S-08-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 1D-08-01 NetApp LUN
#
```

2. Get the device names by executing the following command:

```
sanlun lun show
```

Upon execution, the **sanlun lun show** command shows the devices along with the assigned names (see Example 7-23).

Example 7-23 Getting the device names with the sanlun command

```
# sanlun lun show
filer:          lun-pathname      device filename  adapter  protocol
lun size        lun state
itsotuc3: /vol/vol_DominoDB/lun_DominoDB2 hdisk2      fcs1      FCP
100g (107374182400) GOOD
#
```

3. Verify the attributes of the new disks by entering the following command:

```
lsattr -El device
```

For example, to verify the attribute of newly added device named hdisk2, you would execute that command on the Lotus Domino server (see Example 7-24).

Example 7-24 Verify the disk attributes

```
# lsattr -El hdisk2
clr_q          no          Device CLEARS its Queue on error      False
location       Location Label                                True
lun_id         0x0          Logical Unit Number ID                 False
max_transfer   0x40000      Maximum TRANSFER Size                 True
node_name      0x500a09808647e7ba FC World Wide Node Name              False
pvid           none        Physical volume identifier             False
q_err          yes         Use QERR bit                          False
q_type         simple      Queuing TYPE                          False
queue_depth    12          Queue DEPTH                           True
reassign_to    120         REASSIGN time out value               True
reserve_lock   no          Use SCSI RESERVE when opening the device True
rw_timeout     30          READ/WRITE time out value             True
scsi_id        0x11700      SCSI ID                               False
start_timeout  60          START unit time out value             True
ww_name        0x500a09818647e7ba FC World Wide Port Name              False
```

#

Tip: Changing the queue depth for each disk results in better performance for FCP and iSCSI I/O.

4. Now create a volume group for each LUN device by issuing the following command on the database server:

```
mkvg -f -y 'VgName' device
```

For example, to create a volume group named domdbvg1 for the device named hdisk2, you would execute the command shown in Example 7-25 on the Lotus Domino server.

Example 7-25 Create the volume group

```
# mkvg -f -y 'dbvg1' hdisk2
dbvg1
#
```

Repeat this process until all the necessary volume groups are created.

5. You can verify the volume groups by executing the command shown in Example 7-26 on the database server.

Example 7-26 Verify the volume group

```
# lspv
hdisk0      00000000f1b495eb      rootvg      active
hdisk1      000000006d598dee      None
hdisk2      0007041a63e4b51c      dbvg1       active
#
```

If you assigned the wrong name to a volume group, it can be changed by executing the following commands:

```
varyoffvg VolumeGroupName
exportvg VolumeGroupName
mkvg -q -f -y 'NewVG_Name' Device
```

For example, to change a volume group's name from dbvg1 to domdbvg1, you would execute the following commands on the database server:

```
varyoffvg dbvg1
exportvg dbvg1
mkvg -q -f -y 'domdbvg1' hdisk2
```

6. Create mount points on the Lotus Domino server and assign the appropriate permissions to the database instance by executing the following commands:

```
mkdir -p mountpoint
chown -R user:group mountpoint
```

For example, to create a mount point named /notesdata/db and change the ownership to notes user, you would execute the following command on the database server:

```
mkdir -p /notesdata/db
chown -R notes:notes /notesdata/db
```

7. Create a new file system on the volume group by executing the following command on the database server:

```
crfs -v jfs -a bf=true -g 'VgName' -a size=size < M|G > -m 'mountPoint' -p 'rw'
-a nbpi=4096 -a ag=64
```

For example, to create a file system on the volume group named dbvg1 that is mounted on a mount point named /notesdata/db, you would execute the following command shown in Example 7-27 on the Lotus Domino server.

Example 7-27 File system creation

```
# crfs -v jfs -a bf=true -g dbvg1 -a size=100M -m '/notesdata/db' -p 'rw' -a  
nbpi=4096 -a ag=64
```

Based on the parameters chosen, the new /notesdata/db JFS file system is limited to a maximum size of 134217728 (512 byte blocks)

New File System size is 262144

```
#
```

Note: The value for nbpi must be set to 8192, 16384, 32768, 65546, or 131072 if you need to create a file system that is greater than 64 GB in size.

8. To make the file system mount persistent on system reboots, you need to update the file /etc/filesystems. Open the file /etc/filesystems and locate the entries that correspond to the file systems you want to make persistent and set the option mount = true. For example, to make a file system named /dev/lv04 persistent on database server reboot, you would modify the corresponding entry in the /etc/filesystems file; after the update, the entry should look somewhat similar to the entry shown in Example 7-28.

Example 7-28 /etc/filesystems file (abstract)

```
# cat /etc/filesystems  
/notesdata/db:  
    dev          = /dev/lv04  
    vfs          = jfs  
    log          = /dev/loglv04  
    mount        = true  
    options      = rw  
    account      = false  
  
#
```

Repeat this step for each mounted file system to be used for the Lotus Domino Server, such as the transactional log file system.

7.4 SnapDrive installation and usage

SnapDrive for Linux and UNIX provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file, down through the volume manager, to the storage-system-side logical unit numbers (LUNs) providing the actual repository. Additionally, it simplifies the backup of data and helps you decrease the recovery time.

The IBM System Storage N series SnapDrive feature provides a layer of abstraction between an application running on the host operating system and the underlying IBM System Storage N series storage systems (see Figure 7-2). Applications that are running on a server with SnapDrive use virtual disks (or LUNs) on IBM System Storage N series storage systems, as though they were locally connected drives or mount points. This allows applications that require locally attached storage, such as IBM Lotus Domino and Microsoft Exchange, to benefit from the IBM System Storage N series technologies, including Snapshot, flexible volumes, cloning, and space management technologies.




Figure 7-2 Example of a typical SnapDrive deployment

SnapDrive includes all the necessary drivers and software to manage interfaces, protocols, storage, and Snapshot copies. Snapshot copies are nondisruptive to applications and functions on execution. Snapshot backups can also be mirrored across LAN or WAN links for centralized archiving and disaster recovery.

This section describes the installation process of the Data ONTAP SnapDrive feature. Also, the LUN creation of Lotus Domino database and transactional logging disks are described.

Note: For more information about SnapDrive for Linux and UNIX requirements, see the IBM NAS support web page at:

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

SnapDrive for Linux and UNIX has the following requirements:

- ▶ Host operating system and appropriate patches
- ▶ Host file systems
 - AIX: JFS2
 - Linux: Ext3
- ▶ IP access between the host and IBM System Storage N series storage system

- ▶ IBM System Storage N series storage system licenses
 - FCP license
 - SnapRestore license on the IBM System Storage N series storage system
- ▶ FCP Host Utilities or iSCSI Host Utilities required software
- ▶ For security reason, we recommend a separate user account on the N series storage server.

Note: If your network is insecure, we recommend the usage of HTTPS instead of HTTP. For more information about enabling SSL encryption, see:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>

7.4.1 Install SnapDrive on Linux

The IBM System Storage N series SnapDrive for Linux and UNIX binaries are provided to licensed customers only and can be obtained by contacting IBM support. When you get the binaries, save them on your Lotus Domino Server.

Note: Review the compatibility matrix. For example, at this time, SnapDrive for Linux and UNIX actually supports Red Hat Enterprise Linux Version 4 Update 3, but not Update 4 or higher. If you install SnapDrive in a unsupported environment like Red Hat Enterprise Linux Version 4 Update 4, several features will not work.

Follow these steps to install SnapDrive on Red Hat Linux Enterprise Server 4:

1. Copy the downloaded file to your Linux host. You can place it in any directory on the host. In our case, it is /tmp/snapdrive.
2. Use the `rpm` command to install the software (Example 7-29).

Example 7-29 Installing the SnapDrive on Linux

```
[root@domino1 snapdrive]# rpm -U -v ontap.snapdrive.linux_2_2_1.rpm
Preparing packages for installation...
ontap.snapdrive-2.2.1-1
[root@domino1 snapdrive]#
```

3. Verify your installation, as shown in Example 7-30.

Example 7-30 Verify the SnapDrive installation on Linux

```
[root@domino1 snapdrive]# rpm -qa ontap.snapdrive
ontap.snapdrive-2.2.1-1
[root@domino1 snapdrive]# ls -R /opt/Ontap/
/opt/Ontap/:
snapdrive

/opt/Ontap/snapdrive:
bin diag docs snapdrive.conf

/opt/Ontap/snapdrive/bin:
snapdrive

/opt/Ontap/snapdrive/diag:
```

```
filer_info linux_info SHsupport.pm snapdrive.dc Telnet.pm
```

```
/opt/Ontap/snapdrive/docs:  
man1 snapdrive.1.html
```

```
/opt/Ontap/snapdrive/docs/man1:  
filer_info.1 linux_info.1 snapdrive.1 snapdrive.dc.1  
[root@domino1 snapdrive]# ls -l /usr/sbin/snapdrive  
lrwxrwxrwx 1 root root 39 Jun 6 12:38 ../../opt/Ontap/snapdrive/bin/snapdrive  
[root@domino1 snapdrive]#
```

4. Configure the SnapDrive user login information (Example 7-31). For security reasons, we use the manually created user notes instead of root. Our IBM System Storage N series server name is itsotuc3.

The command **snapdrive config set** expects the following options:

```
snapdrive config set username servername
```

Example 7-31 Configuration of the SnapDrive user login information

```
[root@domino1 snapdrive]# snapdrive config set notes itsotuc3  
Password for notes:  
Retype password:  
[root@domino1 snapdrive]#
```

Note: If you get an error, check if you can ping the IBM System Storage N series storage server name, and check the password and the user permission settings on the Data ONTAP operating system.

Here is an example failure output:

```
0001-136 Admin error: Unable to log on to filer: itsotuc3  
Please change user name and/or password for itsotuc3, i.e.  
snapdrive config set root itsotuc3
```

Possible reasons: The user does not exist, you are using the wrong password, or you are using the wrong permissions. The Data ONTAP syslog messages might help you figure out the problem.

Here is an example syslog message on the Data ONTAP storage server:

```
Wed Jun 6 14:19:37 PDT [itsotuc3: useradmin.unauthorized.user:warning]: User  
'notes' denied access - missing required capability: 'login-http-admin'
```

In this case, the permission login-http-admin was not granted.

The N series SnapDrive feature installation on the Lotus Domino server is completed.

Note: For more information about the installation process of SnapDrive on Linux, see:

<http://www-1.ibm.com/support/search.wss?tc=STCGUJQ&rs=1151&atr=SWPlatform&atr=unix&atrwcs=on&dc=DA400+DB300+DA100+DA110+DA120+DB100+DA700+DA450+DA300&dtm>

7.4.2 Install SnapDrive on AIX

The IBM System Storage N series SnapDrive UNIX binaries are provided to licensed customers only and can be obtained by contacting IBM support. When you get the binaries, save them on your Lotus Domino Server.

Follow these steps to install SnapDrive on AIX server:

1. Log in as root.
2. Start SMIT by entering the following command:

```
# smit
```

 - Select the option titled Software Installation and Maintenance. When you start SMIT, it displays the screen shown in Example 7-32. In this screen, the Software Installation and Maintenance option is the first menu option.

Example 7-32 SMIT command output

System Management

Move cursor to desired item and press Enter.

Software Installation and Maintenance

Software License Management
Devices
System Storage Management (Physical & Logical Storage)
Security & Users
Communications Applications and Services
Print Spooling
Advanced Accounting
Problem Determination
Performance & Resource Scheduling
System Environments
Processes & Subsystems
Applications
Installation Assistant
Cluster Systems Management
Using SMIT (information only)

- At the screen that appears, select the Install and Update Software menu option (see Example 7-33).

Example 7-33 Software Installation and Maintenance

Software Installation and Maintenance

Move cursor to desired item and press Enter.

Install and Update Software

List Software and Related Information
 Software Maintenance and Utilities
 Software Service Management
 Network Installation Management
 EZ NIM (Easy NIM Tool)
 System Backup Manager
 EFIX Management
 Thin Server Maintenance

- At the next screen, select the Install Software menu option (see Example 7-34)

Example 7-34 Install and Update Software screen

Install and Update Software

Move cursor to desired item and press Enter.

Install Software

Update Installed Software to Latest Level (Update All)
 Install Software Bundle
 Update Software by Fix (APAR)
 Install and Update from ALL Available Software

- At the Install Software screen, specify the location of the software in one of the following ways:
 - Manually enter the location by providing the path to the software package (for example, /tmp/Ontap.snapdrive_aix_2_2).

The installation process checks the version of AIX and displays a warning message before it completes if you are not running a version of AIX that SnapDrive for UNIX supports.

Example 7-35 is an example of entering the path to the software package when you are at the Install Software screen.

Example 7-35 Screen to enter the path of software package

Install Software

Type or select a value for the entry field.
 Press Enter AFTER making all desired changes.

	[Entry Fields]
* INPUT device / directory for software	</tmp/Ontap.snapdrive_aix_2_2] +

After you enter the path to the software package, SMIT displays the screen shown in Example 7-36 on page 263. This is the screen on which you are able to modify the software package installation (see Example 7-35). In our case, the default is fine.

Example 7-36 Screen to enter the name of software package

Install Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* INPUT device / directory for software	/tmp/Ontap.snapdrive_>	
* SOFTWARE to install	[_all_latest]	+
PREVIEW only? (install operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install requisite software?	yes	+
EXTEND file systems if space needed?	yes	+
OVERWRITE same or newer versions?	no	+
VERIFY install and check file sizes?	no	+
Include corresponding LANGUAGE filesets?	yes	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	no	+
Preview new LICENSE agreements?	no	+

Example 7-37 shows a successful installation. It is an example of the output you might see when an installation successfully completes.

Example 7-37 Output when installation is successful

```
File:
I:Ontap.snapdrive 2.2
+-----+
Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...
SUCSESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.
Selected Filesets
-----
Ontap.snapdrive 2.2 # Ontap Snapdrive
<< End of Success Section >>
FILESET STATISTICS
-----
1 Selected to be installed, of which:
1 Passed pre-installation verification
----
1 Total to be installed

+-----+
Installing Software...
+-----+
installp: APPLYING software for:
Ontap.snapdrive 2.2
. . . . . << Copyright notice for NetApp.snapdrive >> . . . . .
```

```

Copyright (c) 2006 Network Appliance, Inc.
All Rights Reserved.
. . . . . << End of copyright notice for NetApp.snapdrive >>. . . .
Finished processing all filesets. (Total time: 2 secs).
+-----+
Summaries:
+-----+
Installation Summary
-----
Name Level Part Event Result
-----
Ontap.snapdrive 2.2 USR APPLY SUCCESS

```

- ▶ Verify your installation of the package by entering the `lspp -l Ontap.snapdrive` command. It should produce the output shown in Example 7-38.

Example 7-38 Installation success

```

# lspp -l Ontap.snapdrive
Fileset                                Level  State    Description
-----
Path: /usr/lib/objrepos
Ontap.snapdrive                        2.2.0.0 COMMITTED SnapDrive
#

```

You can also check the SMIT log file (`smit.log` and `smit.script`) if your installation fails. These files are in the SMIT log directory (`$HOME`) and contains installation errors.

- ▶ Complete your setup by configuring SnapDrive for UNIX for your system. Most of this information is set by default; however, you need to specify the following information:
 - The login information for the IBM System Storage N series storage system.
 - The AutoSupport settings.
 - The correct configuration value for the following options based on whether you are using the FCP protocol or the iSCSI protocol:
 - `default-transport`: Set this equal to your protocol. The acceptable values are “`fc`” and “`iscsi`”.
 - `multipathing-type`: For FCP, set this to “`SANPath`”. For iSCSI, set this to “`none`”. Multipathing is not available with iSCSI for this release.

7.4.3 Create and map LUNs with SnapDrive

One of the greatest improvements SnapDrive provides is the storage provisioning feature. It allows you to create, modify, or delete a complete mount point with just one command.

If your operating system requires that you prepare it before you create new LUNs, you can use the `snapdrive config` command. This command lets you check information about how many LUNs can be created on a IBM System Storage N series storage system that is mapped to your host.

Note: The host preparation is currently only required on Linux and Solaris hosts.

Follow these steps to create new LUNs with SnapDrive:

1. Determine how many LUNs could be created.

SnapDrive for Linux and UNIX lets you determine how many LUNs could be created on the host without exceeding a host-local limit (Example 7-39).

Example 7-39 Determine available LUNs on the host

```
[root@domino1 ~]# snapdrive config check luns
LUN ID summary:
    available (prepared, unreserved and not in use): 7
    reserved via available-lun-reserve: 8
```

```
[root@domino1 ~]#
```

Note: If there are no LUN IDs available, SnapDrive cannot create new LUNs. This is an example output of the **snapdrive config check luns** command, where no LUN IDs are prepared:

```
[root@domino1 ~]# snapdrive config check luns
LUN ID summary:
    available (prepared, unreserved and not in use): 0
    reserved via available-lun-reserve: 0 (out of 8 requested)
```

snapdrive cannot create or connect more LUNs until more LUN IDs are available and unreserved.

Use 'snapdrive config prepare luns' or update host manually.

```
[root@domino1 ~]#
```

2. Prepare the host for the creation of a specific number of new LUNs. Use the following command:

```
snapdrive config prepare luns -count count
```

Where *count* is the number of new LUNs for which you want the host to be prepared. See Example 7-40.

Example 7-40 Prepare the host for creating new LUNs

```
[root@domino1 ~]# snapdrive config prepare luns -count 8
LUN ID update:
    LUN IDs prepared adjusted to be: 16
    now available (prepared, unreserved and not in use): 8
    now reserved via available-lun-reserve: 8
```

```
[root@domino1 ~]#
```

Note: Even if you give a count less than 8, SnapDrive will prepare at least eight LUNs for usage as the default. This value can be modified in the SnapDrive configuration file /opt/Ontap/snapdrive/snapdrive.conf:

```
[root@domino1 snapdrive]# grep lun-reserve /opt/Ontap/snapdrive/snapdrive.conf
#available-lun-reserve=8 # Number of LUNs for which to reserve host resources
[root@domino1 snapdrive]#
```

3. Create the LUN and file system with SnapDrive.

Note: SnapDrive is able to use the operating system Logical Volume Manager (LVM) for creating additional volume groups and logical volumes. If you prefer to create new file systems without LVM, you can use the **snapdrive storage create** command with the **-nolvm** option. Because only LVM volumes could be resized, we recommend the usage of LVM for the creation of new LUNs and file systems.

Execute the **snapdrive storage create** command, as in Example 7-41.

Example 7-41 LUN and file system creation

```
[root@domino1 ~]# snapdrive storage create -fs /notesdata/log -filervol  
itsotuc3:/vol/vol_DominoLog -vgsize 5g  
  
LUN log-1_SdLun ... created  
  
mapping new lun(s) ... done  
discovering new lun(s) ... done  
  
LUN to device file mappings:  
- itsotuc3:/vol/vol_DominoLog/log-1_SdLun => /dev/sdc  
  
disk group log_SdDg created  
host volume log_SdHv created  
file system /notesdata/log created  
  
[root@domino1 ~]#
```

Keep the following in mind when you execute this command:

- The **-fs** argument specifies the mount point of the file system you want to create. We use **/notesdata/log** in our case.
- Use the **-filervol** option to specify the IBM System Storage N series storage system and volume where you want the LUNs created.

Do not specify the LUN. SnapDrive for Linux and UNIX creates the LUN automatically when you use this form of the **snapdrive storage create** command. It uses system defaults to determine the LUN IDs, and the size of each LUN. It bases the names of the associated disk/volume groups on the name of the host volume or file system.

- To control the size of the host volume group, use the **-vgsize** size option to specify the size in bytes of the underlying LUN.

Note: For more details of the **snapdrive storage create** command, see the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide* at:

<http://www-1.ibm.com/support/docview.wss?uid=ssglS7001600&aid=1>

4. Verify the created LUN and file system (Example 7-42).

Example 7-42 Verify the created mount point

```
[root@domino1 etc]# mount | grep notesdata  
/dev/mapper/log_SdDg-log_SdHv on /notesdata/log type ext3 (rw)  
[root@domino1 etc]#
```

The Lotus Domino database mount point is created in the same way (Example 7-43).

Example 7-43 Creating the Lotus Domino database LUN and file system

```
[root@domino1 ~]# snapdrive storage create -fs /notesdata/db -filervol
itsotuc3:/vol/vol_DominoDB -vgsize 146g

LUN db_SdLun ... created

mapping new lun(s) ... done
discovering new lun(s) ... done

LUN to device file mappings:
- itsotuc3:/vol/vol_DominoDB/db_SdLun => /dev/sdd

disk group db_SdDg created
host volume db_SdHv created
file system /notesdata/db created

[root@domino1 ~]# mount|grep notesdata
/dev/mapper/log_SdDg-log_SdHv on /notesdata/log type ext3 (rw)
/dev/mapper/db_SdDg-db_SdHv on /notesdata/db type ext3 (rw)
[root@domino1 ~]#
```

7.4.4 Create Snapshots with SnapDrive

You use the **snapdrive snap create** command to create Snapshot copies, which are point-in-time, read-only images of data on IBM System Storage N series storage system volumes. The **snap create** operation ensures that you have backed up your LUNs files and directory trees. You can use the Snapshot copy you create to restore your data if you encounter corruption or other problems.

To ensure that a Snapshot copy is “application-consistent,” you usually need to stop or do whatever steps are required to quiesce the application before taking the Snapshot copy. IBM Lotus Domino has powerful database consistence routines. In this case, taking a Snapshot requires no application shutdown when transactional logging is used.

Note: To ensure database consistency, use Lotus Domino transactional logging in conjunction with the IBM System Storage N series Snapshot feature.

Follow these guidelines when you enter commands that create Snapshot copies:

- ▶ You can keep a maximum of 255 Snapshot copies per IBM System Storage N series storage system volume. This limit is set by the IBM System Storage N series storage system. The total number can vary depending on whether other tools use these Snapshot copies.

When the number of Snapshot copies has reached the maximum limit, the **snapshot create** operation fails. You must delete some of the old Snapshot copies before you can use SnapDrive for Linux and UNIX to take any more.

- ▶ SnapDrive for UNIX does not support Snapshot copies that it does not create. For example, it does not support Snapshot copies that are created from the IBM System Storage N series storage system console, because such a practice can lead to inconsistencies within the file system.

- You cannot use SnapDrive for UNIX to create Snapshot copies of the following:
 - Root disk groups. The **snap create** operation fails when you try to take a Snapshot copy of a root disk group for a logical volume manager.
 - Boot device or swap device. SnapDrive for UNIX does not take a Snapshot copy of a system boot device or a system swap device.

Note: For more information about requirements, hints, and commands, we highly recommend the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssglS7001600&aid=1>

To create a Snapshot copy from the Lotus Domino database and transactional log LUN, using SnapDrive for Linux and UNIX, complete the following steps:

1. Create the Snapshot with the **snapdrive snap create** command (Example 7-44).

Example 7-44 Creating the Snapshot of Lotus Domino database and transactional log

```
[root@domino1 /]# snapdrive snap create -vg db_SdDg log_SdDg -snapname snap1
Successfully created snapshot snap1 on 2 filer volumes:
    itsotuc3:/vol/vol_DominoDB
    itsotuc3:/vol/vol_DominoLog

    snapshot snap1 contains:
    disk group db_SdDg containing host volumes
        db_SdHv (filesystem: /notesdata/db)
    disk group log_SdDg containing host volumes
        log_SdHv (filesystem: /notesdata/log)
[root@domino1 /]#
```

- **-vg:** This switch follows the volume group from which a Snapshot should be created. It is allowed to stack more than one volume group for crash-consistency.
- **-snapname:** The short name of your Snapshot. In our case, it is *snap1*.

Note: Unless you specify otherwise, SnapDrive assumes that all entities that you specify on a given **snap create** command line are related, in other words, that the validity of updates to one entity may depend on updates to the other entities specified. When storage entities have *dependent writes* in this way, SnapDrive takes steps to create a Snapshot copy that is crash-consistent for all storage entities as a group.

2. Validate the created Snapshot. Use the **snapdrive snap list** command for a list of all Snapshots on the specified file system or volume (see Example 7-45).

Example 7-45 List all Snapshots

```
[root@domino1 /]# snapdrive snap list -vg db_SdDg log_SdDg

snap name host date snapped
-----
itsotuc3:/vol/vol_DominoDB:snap1 domino1 Jun  8 16:24 db_SdDg log_SdDg
itsotuc3:/vol/vol_DominoLog:snap1 domino1 Jun  8 16:24 db_SdDg log_SdDg
[root@domino1 /]#
```

- -vg: List all Snapshots of specified volume groups.

Note: In the end of every row, the Snapshot group is quoted. In the time that the Snapshot was created, SnapDrive stopped all I/O operations on db_SdDg and log_SdDg.

7.4.5 Resize the volume

SnapDrive for UNIX lets you increase the size of the IBM System Storage N series storage system volume group or disk group. You use the **snapdrive storage resize** command to do this task.

Note: This command does not let you resize host volumes or file systems. For example, you cannot use the **resize** command to change the size of a file system on a LUN. This makes SnapDrive for Linux and UNIX different than SnapDrive for Windows.

SnapDrive for UNIX adds a system-generated LUN. If you specify an amount by which you want to increase the storage, such as 20 GB, it makes the LUN 20 GB. If you specify a target size for the storage, it calculates the difference between the current size and the target size. The difference becomes the size of the LUN it then creates. These are the two options to increase the size of your storage.

Follow these hints when you use SnapDrive for Linux and UNIX for storage resizing:

- ▶ The storage resize operation can only increase the size of storage. You cannot use it to decrease the size of an entity.
- ▶ All LUNs must reside in the same IBM System Storage N series storage system volume.
- ▶ The resize operation is not supported directly on logical host volumes, or on file systems that reside on logical hosts volumes or on LUNs. In those cases, you must use the LVM commands to resize the storage.

Note: Additional guidelines are available for cluster environment configurations. For details, go to the following URL:

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

The **snapdrive storage resize** command applies only to IBM System Storage N series storage system disk groups and volume groups. If you want to increase the size of your host volume or file system, you must use LVM commands. Table 7-1 summarizes the LVM commands you can use on the different platforms. For more information about these commands, see their man pages.

Table 7-1 Host volume and file system resize commands

Host	Host volume	File systems
AIX	extendlv	chfs
Linux	lvextend	resize2fs

1. As shown in Example 7-46, we increase the size of our IBM Lotus Domino database volume, mounted in /notesdata/db, to about 20 GB for further growth. The name of the volume group in our case is db_SdDg. On Linux, you can review your available volume groups with the **lvm vgdisplay** command. In AIX, the **lsvg** command displays the volume groups.

Example 7-46 Resize the volume group by adding an additional LUN

```
[root@domino1 ~]# snapdrive storage resize -vg db_SdDg -growby 20g -addlun
discovering filer LUNs in disk group db_SdDg...done
LUN itsotuc3:/vol/vol_DominoDB/db-1_SdLun ... created

mapping new lun(s) ... done
discovering new lun(s) ... done.
initializing LUN(s) and adding to disk group db_SdDg...done
Disk group db_SdDg has been resized
Desired resize of host volumes or file systems
contained in disk group must be done manually
[root@domino1 ~]#
```

- -vg db_SdDg: The volume group db_SdDg, which should be affected by our resize command.
 - -growby 20g: Grow the volume group by 20 GB.
 - -addlun: This option is necessary, because SnapDrive for Linux and UNIX does not support resizing a LUN at the this book was written. A new LUN will be added to the volume group.
2. Extend the logical volume with the **lvextend** command. Use the **extendlv** command for AIX (Example 7-47).

Example 7-47 Extending the logical volume

```
[root@domino1 ~]# lvextend -L +20G /dev/db_SdDg/db_SdHv
Extending logical volume db_SdHv to 166.02 GB
Logical volume db_SdHv successfully resized
[root@domino1 ~]#
```

- -L +20G: Extend the logical volume about 20 GB.
- /dev/db_SdDg/db_SdHv: The device path of your logical volume. If you are not sure about it, use the **lvm lvdisplay** command for details about your volumes. Use **lslv** for AIX. Use **smit lslv** for the user interface.

Note: The man page of the **lvextend** command on RHEL4 Update 3 refers to the **--resizefs** option, which should resize the logical volume file system within one step. Unfortunately, this option was available in LVM1 and does not work with LVM2, which is used in newer Linux versions. You have to use the **resize2fs** command after resizing the logical volume. For more information, refer to the following Web site:

https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=117316

3. Resize the file system. For this task, you have to halt the Lotus Domino server and umount the file system (see Example 7-48 on page 271). As we use the ext3 journaling file system, an additional check with **e2fsck** is not needed. Use the **smit chfs** command to increase the size of the file system in AIX.

Note: To umount the file system, it is necessary to halt the Lotus Domino server first.

Example 7-48 Resize the file system

```
[root@domino1 ~]# umount /notesdata/db
[root@domino1 ~]#
[root@domino1 ~]# resize2fs -f /dev/db_SdDg/db_SdHv
resize2fs 1.35 (28-Feb-2004)
Resizing the filesystem on /dev/db_SdDg/db_SdHv to 43520000 (4k) blocks.
The filesystem on /dev/db_SdDg/db_SdHv is now 43520000 blocks long.
[root@domino1 ~]#
[root@domino1 ~]# mount /notesdata/db
[root@domino1 ~]#
```

- **-f**: It forces **resize2fs** to resize the file system on the logical volume. This option is necessary.
- **/dev/db_SdDg/db_SdHv**: Device path of the logical volume.

As we do not use any further option, **resize2fs** will expand the file system to the logical volume size.

4. Verify if the resizing operation affected your file system. For this, we use the **mount** command, (Example 7-49).

Example 7-49 Verify the resized file system

```
[root@domino1 ~]# df -h /notesdata/db
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/db_SdDg-db_SdHv
                164G   92M  159G   1% /notesdata/db
[root@domino1 ~]#
```

The file system is increased by 20 GB and the Lotus Domino server can be started again.

7.4.6 Best practice: Create storage with LVM enabled

Because the **snapdrive storage resize** command applies only to IBM System Storage N series storage system disk groups and volume groups, you should not use the **-nolvm** option within storage creation, such as a file system or LUN. SnapDrive is able to use the operating system Logical Volume Manager (LVM) for creating additional volume groups and logical volumes.

7.5 Lotus Domino storage partitioning

The right storage partitioning for Lotus Domino is important for data reliability and performance. While it is common to place the program files on the same drive as the operating system, we highly recommend creating additional volumes on different disks for the database and transactional logs. This has some benefits:

- In the unlikely case of an aggregate failure and loss of the database, you will be able to restore all transactions, from the last database backup up to the time of failure because the transactional log files reside on another aggregate.

- On a server failure, all configuration and user data resides on the external volumes. You can easily switch the LUNs to another server, install the Lotus Domino binaries, and your business will be online again.
- Improved I/O performance because of the Lotus Domino database and the transactional logs' different I/O profile.

Figure 7-3 shows the example file system and mount structure on the Red Hat Linux Enterprise Server in our lab.

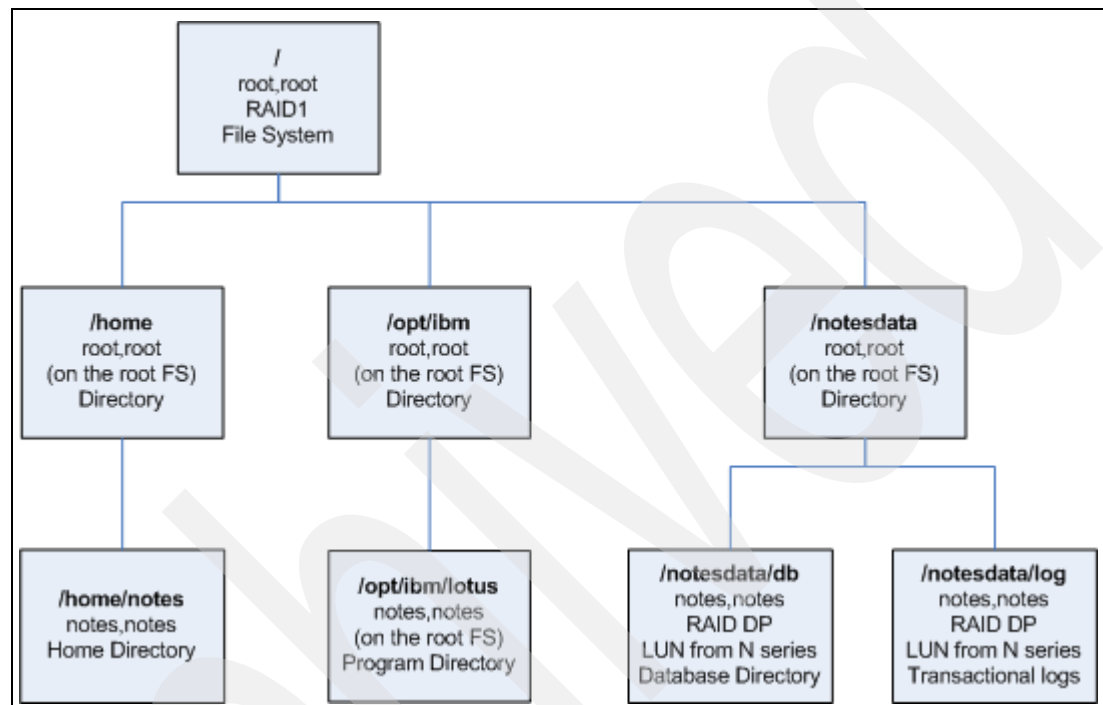


Figure 7-3 An example Linux file system structure we used for this book

Where:

- Database placement: `/notesdata/db`
- Transactional log placement: `/notesdata/log`
- Program file placement: `/opt/ibm/lotus`

Note: On AIX, the default installation path for Lotus Domino program files is `/opt/lotus`.

7.5.1 Database placement

The IBM Lotus Domino database is the fastest growing and most important storage item.

We create an additional path to the Linux root file system named `/notesdata`. Below this path, the `/notesdata/db` mount point is reserved for database placement (see the IBM Lotus Domino database directory installation step in Example 7-50 on page 273). It contains an external LUN from the IBM System Storage N series storage, connected through FCP. This LUN is created through the SnapDrive for Linux and UNIX application (see 7.4.3, “Create and map LUNs with SnapDrive” on page 264).

The Lotus Domino database is stored on different physical disks than the program and transactional logging files. We recommend the IBM System Storage N series RAID Double Parity (DP) protection for the aggregate.

Example 7-50 IBM Lotus Domino database directory setting

```
=====
Domino Server Installation
=====

The data directory is the path where the Install program
installs the Domino data files.

-----
Type h for help.
Type e to exit the Install program.
Press ESC to return to the previous screen.
Press ENTER to edit a setting.
Press TAB to accept a setting and continue to the next screen.
-----

Current data directory setting : /notesdata/db
=====
```

7.5.2 Transactional log placement

As with the Lotus Domino database, the transactional logging files should be separated by physical disks. If you lose or crash the database, the transactional logs will recover all transactions since your last database backup.

Since every transaction will first be saved in the transactional log files, performance is very important for the transactional log partition. This is why we recommend an exclusive Flexible volume for the transactional log files.

In our case, the /notesdata/db mount point is reserved for Lotus Domino transactional log placement. It contains an external LUN from the IBM System Storage N series storage, connected through FCP. This LUN is created through the SnapDrive for Linux and UNIX application (see 7.4.3, “Create and map LUNs with SnapDrive” on page 264).

To change the Lotus Domino transactional logging path, select **Configuration** → **Server** → **Current Server Document** → **Transactional Logging** (see Figure 7-4).

Set the Log path option to your transactional logging directory, in our case, /notesdata/log.

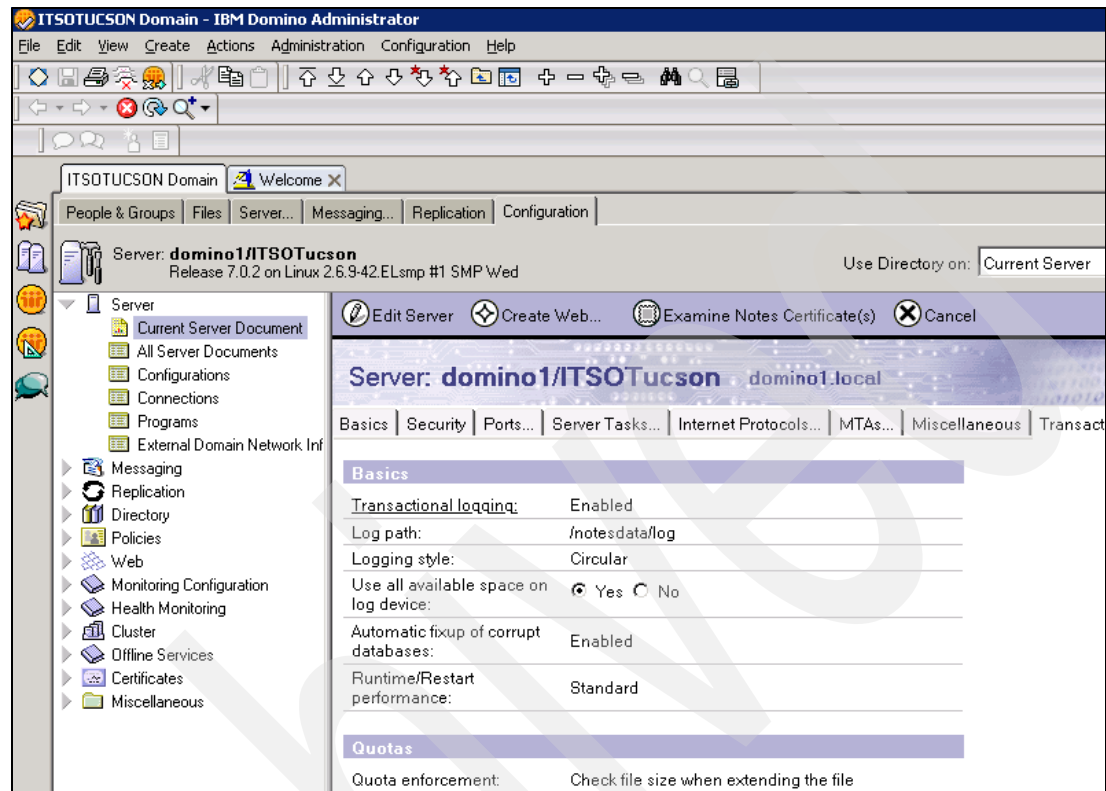


Figure 7-4 Transactional logging configuration in the Lotus Domino Server Document

The Lotus Domino transactional logs are stored on different physical disks than the program and database files. Because it is not a disaster if you lose your transactional log volume, we recommend the IBM System Storage N series RAID 4 protection for your Lotus Domino transactional logging aggregate, which saves you one more disk and gives slightly better performance.

7.5.3 Program file placement

The Lotus Domino program files are not changed much at run time. They should be protected by RAID, such as RAID 1. We recommend you put the program files on the local system volume.

In our case, we used the default installation path /opt/ibm/lotus for Lotus Domino program files. It is located on the operating system root volume.

Example 7-51 on page 275 shows the IBM Lotus Domino installation steps for setting the program directory.

Example 7-51 IBM Lotus Domino program file directory setting

=====
Domino Server Installation
=====

The program directory is the path where the Install program installs the Domino program files. The Install program automatically adds "lotus" to the path.

Type h for help.
Type e to exit the Install program.
Press ESC to return to the previous screen.
Press ENTER to edit a setting.
Press TAB to accept a setting and continue to the next screen.

Current program directory setting : /opt/ibm/lotus

Archived



Part 4

IBM System Storage N series system operations with Lotus Domino Server and Microsoft Exchange

This part will discuss those common administration tasks on the IBM System Storage N series when working with Lotus Domino Server and Microsoft Exchange.

Archived



IBM System Storage N series administration with Lotus Domino Server and Microsoft Exchange

This chapter discuss administering IBM System Storage N series storage systems day to day operations with Lotus Domino Server and Microsoft Exchange. Storage administrators need to monitor the system performance, do provisioning for expanding the storage space, and review reports.

8.1 FilerView

FilerView is the GUI system administrator user interface that runs on the IBM System Storage N series. Administrators can manage aggregates, volumes, LUNs, protocols, security, and many other storage features. Refer to Figure 8-1.

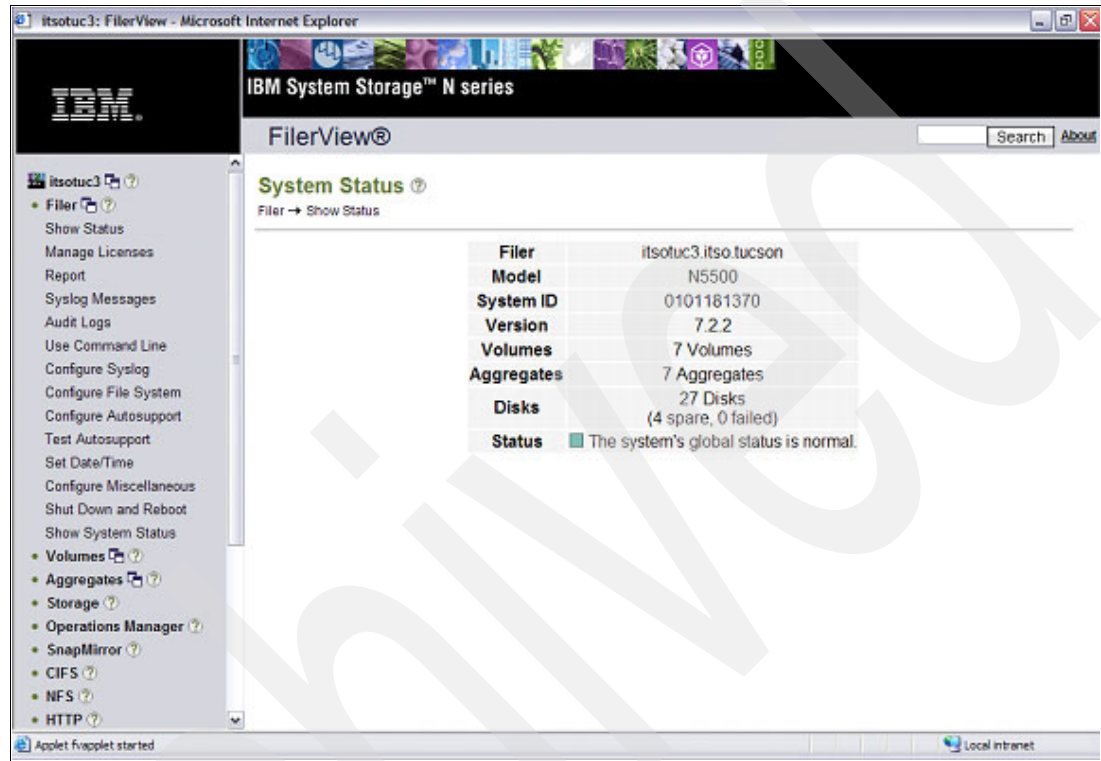


Figure 8-1 FilerView

8.1.1 Resize aggregate

Storage requirements are increasing every year with applications such as Lotus Domino and Microsoft Exchange serving more users and storing more data. As the volumes in the aggregate are completely used up, the network administrator needs to increase the aggregate storage space. The procedure to increase the aggregate space is defined as follows:

In the Manage Aggregate window of FilerView, select **Aggregate** → **Manage** (Figure 8-2).

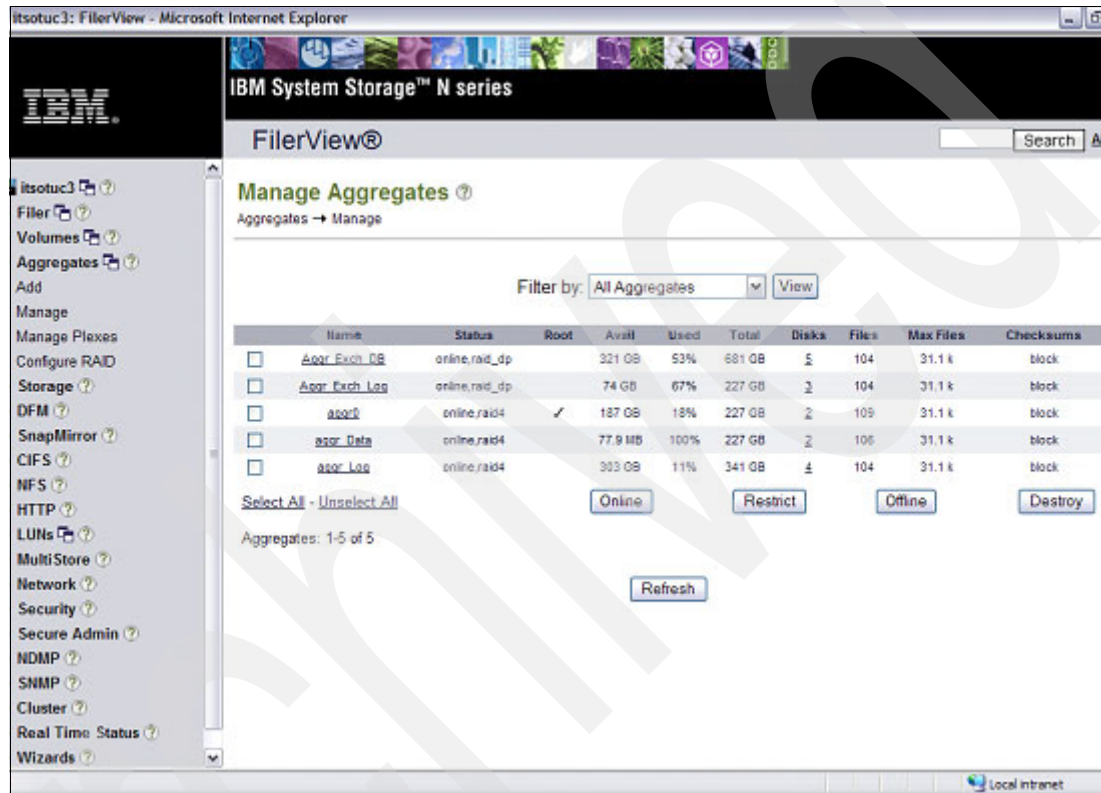


Figure 8-2 Manage Aggregate

Click the aggregate you want to resize (Figure 8-3).

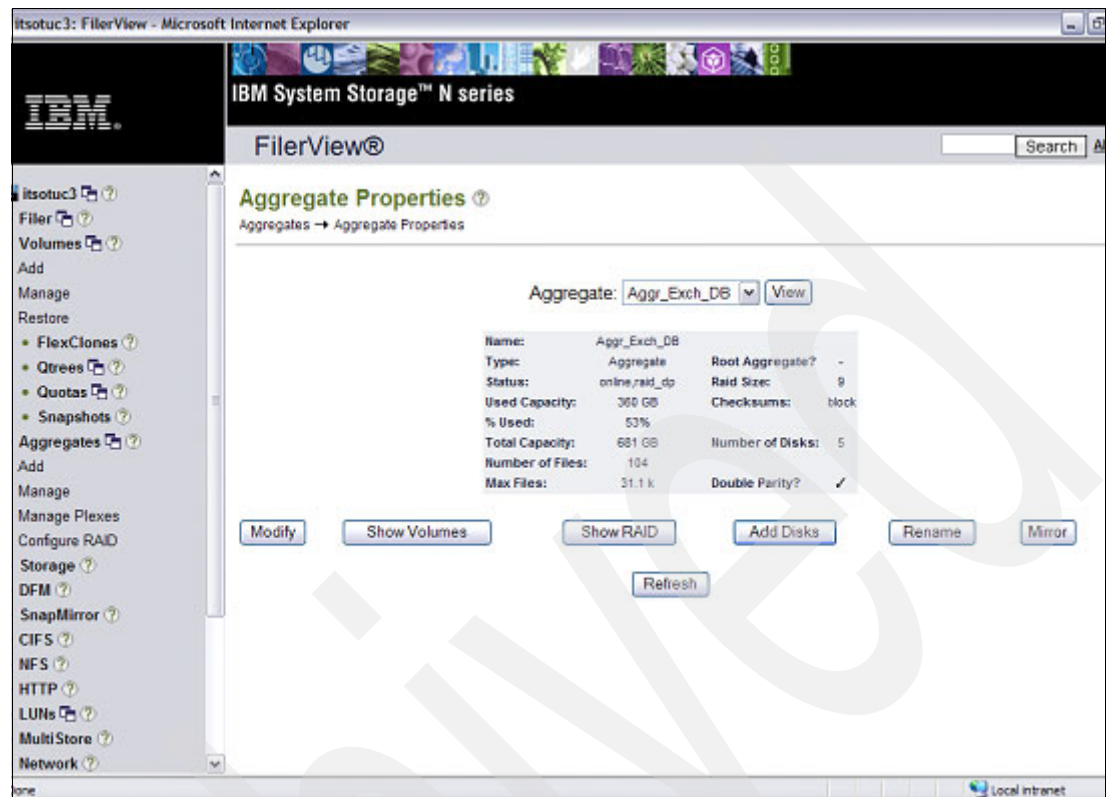


Figure 8-3 Aggregate Properties

Click **Next** (Figure 8-4).




Figure 8-4 Aggregate Storage Wizard

Select **Disk Selection Method** (Figure 8-5).

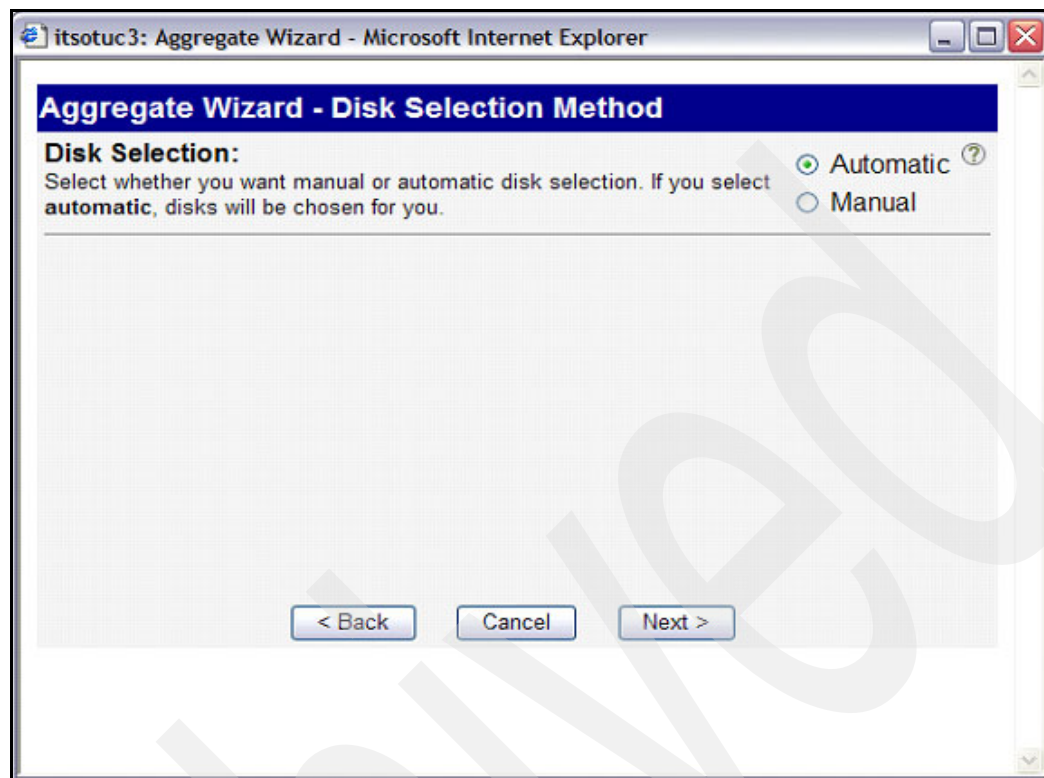


Figure 8-5 Disk Selection Method

Select **Disk Size** (Figure 8-6).

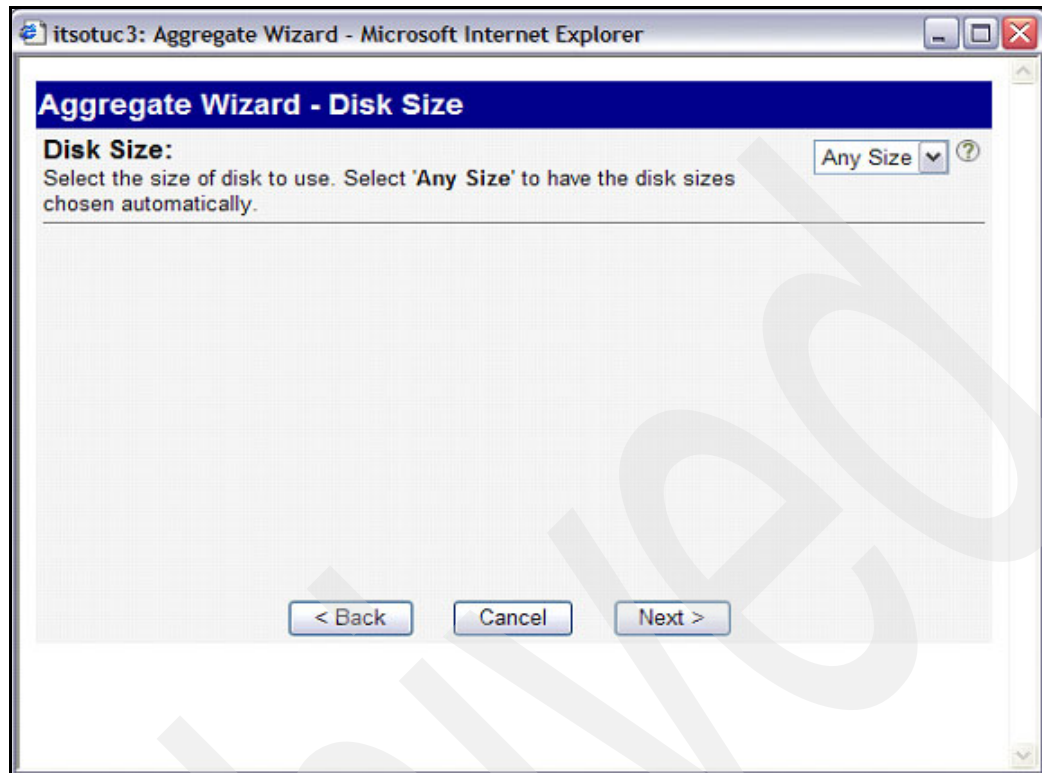


Figure 8-6 Select Disk Size

Select **Number of disks** (Figure 8-7).

The screenshot shows a web browser window titled "itsotuc3: Aggregate Wizard - Microsoft Internet Explorer". The main content area has a blue header bar with the text "Aggregate Wizard - Number of Disks". Below the header, the text "Number of Disks:" is followed by a description: "Select the number of disks of size 'Any Size' to add to the aggregate. There are a total of 11 spares available." To the right of this text is a dropdown menu showing the number "1" and a help icon (a question mark in a circle). Below the description is a large, empty rectangular area with a light gray grid pattern. At the bottom of the grid area are three buttons: "< Back", "Cancel", and "Next >".

Figure 8-7 Select Number of Disks

Click **Commit** to resize the aggregate (Figure 8-8).




Figure 8-8 Commit to resize the aggregate

8.1.2 Resize volume

As the capacity requirements increase for Lotus Domino and Microsoft Exchange, the administrator will need to increase the size of LUNs if there is no space left in the volume. The procedure to increase the volume size is defined in the following steps:

In the Manage Volumes window of FilerView, select **Volume** → **Manage** (Figure 8-9).

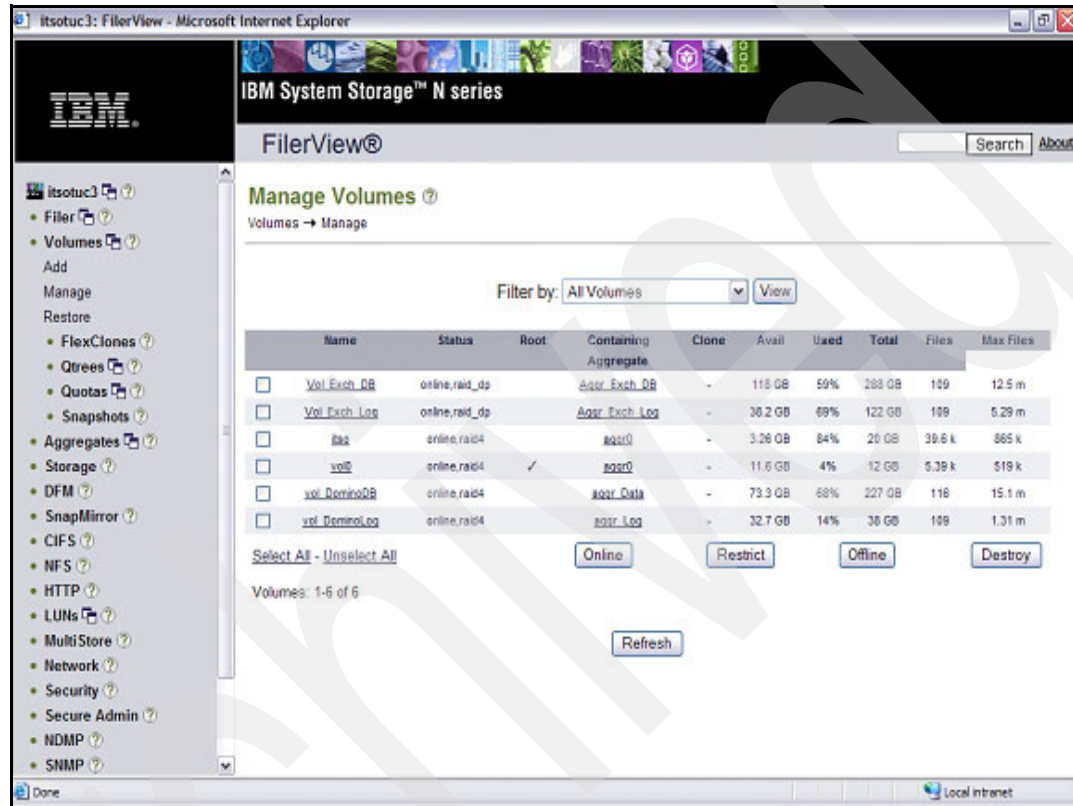


Figure 8-9 N series Manage Volume

Click the volume that you want to resize (Figure 8-10).

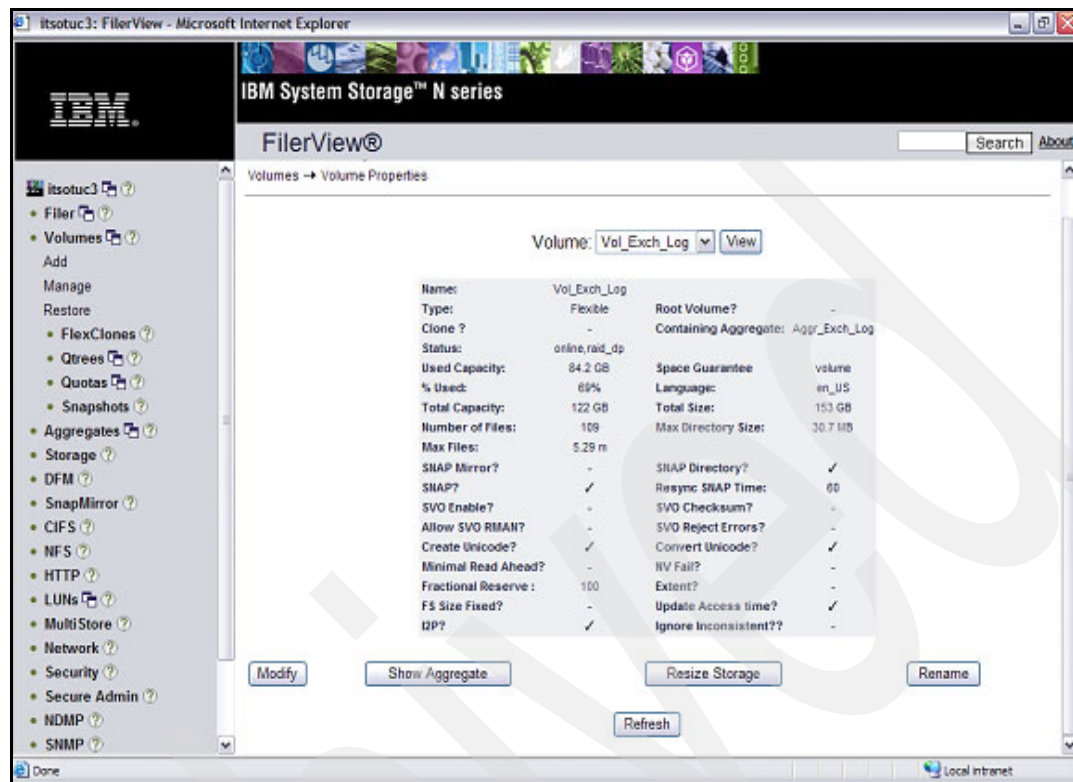


Figure 8-10 Volume Properties

Click **Resize Storage** (Figure 8-11).

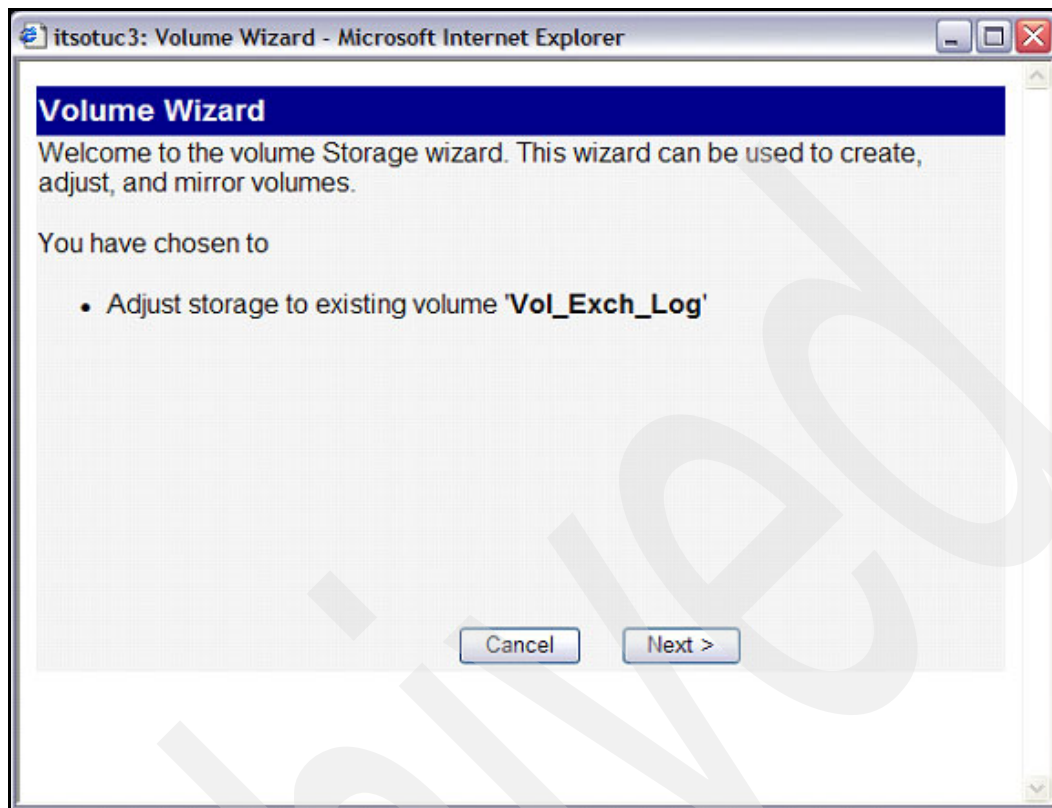


Figure 8-11 Resize volume

Click **Next** (Figure 8-12 on page 291).

Volume Wizard - Flexible Volume Parameters

Containing Aggregate: Aggr_Exch_Log ?
The containing aggregate of volume 'Vol_Exch_Log'.

Total Volume Size: 170 GB ?
Enter the total amount of space for this volume. The total volume size includes space for the snapshot reserve and the file system overhead in addition to the usable space. The volume is using a total of 84.2 GB out of its current 153 GB total volume size. The new total volume size cannot exceed 227 GB with the Space Guarantee set to volume.
227 GB (Max)

Space Guarantee: volume ?
Sets the space guarantee. Volume guarantees space for the entire the volume in the containing aggregate; File guarantees space for a file at file allocation time.

< Back Cancel Next >

Figure 8-12 Volume Parameters

Select the Parameter you need to change. Click **Next** (Figure 8-13).

Volume Wizard - Commit

Below is a summary of your changes.

```
Adjust Existing Volume
Volume Name: Vol_Exch_Log
New Volume Size: 170G
Space Guarantee: volume
```

< Back Cancel Commit

Figure 8-13 Summary of volume properties

Click **Commit** to resize the volume.

8.1.3 Resizing the LUN Using SnapDrive

Lotus Domino and Microsoft Exchange may require more storage space to meet the database growth due to an increasing number of users. If the volume or volumes do not have enough storage space, then the volume size needs to be increased and then the LUN size. The procedure to increase the LUN size is defined below.

SnapDrive for UNIX (including AIX) lets you increase the size of the storage system volume group or disk group. You use the **snapdrive storage resize** command to do this.

Note: This command does not let you resize host volumes or file systems. For example, you cannot use the **resize** command to change the size of a file system on a LUN. This is the difference with SnapDrive for Linux and UNIX in contrast to SnapDrive for Windows.

SnapDrive for UNIX adds a system-generated LUN. If you specify an amount by which you want to increase the storage, such as 20 GB, it makes the LUN 20 GB. If you specify a target size for the storage, it calculates the difference between the current size and the target size. The difference becomes the size of the LUN it then creates. These are the two options to increase the size of your storage.

Follow these hints when you use SnapDrive for Linux and UNIX for storage resizing:

- ▶ The storage resize operation can only increase the size of storage. You cannot use it to decrease the size of an entity.
- ▶ All LUNs must reside in the same storage system volume.
- ▶ The resize operation is not supported directly on logical host volumes, or on file systems that reside on logical host volumes or on LUNs. In those cases, you must use the LVM commands to resize the storage.
- ▶ You must use the **-addlun** option to add a new LUN. Refer to Example 8-1 on page 293.

Note: Additional guidelines are available for cluster environment configurations. For details, see:

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

The **snapdrive storage resize** command applies only to storage system disk groups and volume groups. If you want to increase the size of your host volume or file system, you must use LVM commands. Table 8-1 summarizes the LVM commands you can use on the different platforms. For more information about these commands, see their man pages.

Table 8-1 Host volume and file system resize commands

Host	Host volume	File systems
AIX	extendlv	chfs
Linux	lvextend	resize2fs

1. As shown in Example 8-1 on page 293, we increase the size of our IBM Lotus Domino database volume, mounted in /notesdata/db, to 20 GB for future growth. The name of the volume group in our case is db_SdDg. On Linux, you can review your available volume

groups with the **lvm vgdisplay** command. In AIX, the **lsvg** command displays the volume groups.

Example 8-1 Resize the volume group by adding an additional LUN

```
[root@domino1 ~]# snapdrive storage resize -vg db_SdDg -growby 20g -addlun
discovering filer LUNs in disk group db_SdDg...done
LUN itsotuc3:/vol/vol_DominoDB/db-1_SdLun ... created

mapping new lun(s) ... done
discovering new lun(s) ... done.
initializing LUN(s) and adding to disk group db_SdDg...done
Disk group db_SdDg has been resized
Desired resize of host volumes or file systems
contained in disk group must be done manually
[root@domino1 ~]#
```

- **-vg db_SdDg**: The volume group db_SdDg, which should be affected by our **resize** command.
- **-growby 20g**: Grow the volume group by 20 GB.
- **-addlun**: This option is necessary, because SnapDrive for Linux and UNIX does not support resizing a LUN at the time this book was written. A new LUN will be added to the volume group.

2. Extend the logical volume with the **lvextend** command. Use the **extendlv** command for AIX (Example 8-2).

Example 8-2 Extending the logical volume

```
[root@domino1 ~]# lvextend -L +20G /dev/db_SdDg/db_SdHv
Extending logical volume db_SdHv to 166.02 GB
Logical volume db_SdHv successfully resized
[root@domino1 ~]#
```

- **-L +20G**: Extend the logical volume to about 20 GB.
- **/dev/db_SdDg/db_SdHv**: The device path of your logical volume. If you are not sure about it, use the **lvm lvdisplay** command for details about your volumes. Use **ls1v** for AIX. Use **smit ls1v** for the user interface.

Note: The man page of the **lvextend** command on RHEL4 Update 3 refers to the **--resizefs** option, which should resize the logical volume file system within one step. Unfortunately, this option was available in LVM1 and does not work with LVM2, which is used in the newer Linux version. You have to use the **resize2fs** command after resizing the logical volume. For more information, refer to:

https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=117316

3. Resize the file system (Example 8-3). For this task, you have to halt the Lotus Domino server and umount the file system. As we use the ext3 journaling file system, an additional check with **e2fsck** is not needed. Use the **smitty chfs** command to increase the size of the file system in AIX.

Note: To umount the file system, it is necessary to halt the Lotus Domino server first.

Example 8-3 Resize the file system

```
[root@domino1 ~]# umount /notesdata/db
[root@domino1 ~]#
[root@domino1 ~]# resize2fs -f /dev/db_SdDg/db_SdHv
resize2fs 1.35 (28-Feb-2004)
Resizing the filesystem on /dev/db_SdDg/db_SdHv to 43520000 (4k) blocks.
The filesystem on /dev/db_SdDg/db_SdHv is now 43520000 blocks long.
[root@domino1 ~]#
[root@domino1 ~]# mount /notesdata/db
[root@domino1 ~]#
```

- **-f**: It forces **resize2fs** to resize the file system on the logical volume. This option is necessary.
- **/dev/db_SdDg/db_SdHv**: The device path of the logical volume.

If we do not use any additional options, **resize2fs** will expand the file system to the logical volume size.

4. Verify if the resizing operation affected your file system. For this task, we use the **mount** command (see Example 8-4).

Example 8-4 Verify the resized file system

```
[root@domino1 ~]# df -h /notesdata/db
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/db_SdDg-db_SdHv
                164G   92M  159G   1% /notesdata/db
[root@domino1 ~]#
```

The file system is increased by 20 GB and the Lotus Domino server can be started again.

8.2 Working with Snapshots

Often with Lotus Domino and Microsoft Exchange, you will have to manipulate IBM System Storage N series Snapshots in some fashion. This section describes how to work with Snapshots created by SnapDrive for Linux and UNIX and AIX. The following topics are discussed:

- ▶ Connecting the Snapshot
- ▶ Disconnecting the Snapshot
- ▶ Renaming the Snapshot
- ▶ Deleting the Snapshot

For more information about working with Snapshots, see the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>

Connecting the Snapshot

SnapDrive for UNIX lets you connect to a Snapshot copy from a different location on a host. This new location can be on the host where you took the Snapshot copy (the originating host) or on a different host (the non-originating host).

Lotus Domino and Microsoft Exchange need to connect to Snapshot to work with the backup files.

Being able to access the Snapshot copies from an new location allows you to back up a Snapshot copy to another medium, or simply access the Snapshot copy data without disrupting the original copy of the data.

Example 8-5 Connect to the Snapshot with SnapDrive

```
[root@domino1 /]# snapdrive snap connect -vg db_SdDg -autorename -snapname  
itsotuc3:/vol/vol_DominoDB:snap2
```

```
connecting db_SdDg:  
  LUN copy db_SdLun_0 ... created  
    (original: itsotuc3:/vol/vol_DominoDB/db_SdLun)  
  LUN copy db-1_SdLun_0 ... created  
    (original: itsotuc3:/vol/vol_DominoDB/db-1_SdLun)  
  
  mapping new lun(s) ... done  
  discovering new lun(s) ... done  
  Importing db-1_SdDg_0  
Successfully connected to snapshot itsotuc3:/vol/vol_DominoDB:snap2  
  disk group db-1_SdDg_0 containing host volumes  
    db_SdHv_0 (filesystem: /notesdata/db_0)  
[root@domino1 /]#
```

- -vg: The volume group name to which you wish to connect.
- -autorename: This option tells SnapDrive for Linux and UNIX to generate a new, unused name for the destination entity if the default name is in use.
- -snapname: Long name of the Snapshot.

Note: There are several command-line options for the **snapdrive snap** command:

```
[root@domino1 ~]# snapdrive snap help
snapdrive: For detailed syntax of individual commands, type
snapdrive: 'snapdrive command operation help'
snapdrive: Supported commands and operations are:
      snapdrive snap      show
      snapdrive snap      list
      snapdrive snap      create
      snapdrive snap      delete
      snapdrive snap      rename
      snapdrive snap      connect
      snapdrive snap      disconnect
      snapdrive snap      restore
      snapdrive storage  show
      snapdrive storage  list
      snapdrive storage  create
      snapdrive storage  delete
      snapdrive storage  resize
      snapdrive storage  connect
      snapdrive storage  disconnect
      snapdrive host     connect
      snapdrive host     disconnect
      snapdrive version
      snapdrive config   access
      snapdrive config   prepare
      snapdrive config   check
      snapdrive config   show
      snapdrive config   set
      snapdrive config   delete
      snapdrive config   list

[root@domino1 ~]#
```

See the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide* for detailed information:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>

Disconnecting the Snapshot

You use the **snap disconnect** command to remove the mappings for storage entities in the Snapshot copy.

You can use this command to disconnect Snapshot copies that span multiple storage system volumes or multiple storage systems. The storage entities and volumes can reside on the same storage system or different storage systems.

When you disconnect a file system, SnapDrive for Linux and UNIX always removes the mount point. In Figure 8-6 on page 297, we disconnect the volume group db-1_SdDg_0.

Example 8-6 Disconnect the Snapshot with SnapDrive

```
[root@domino1 ~]# snapdrive snap disconnect -vg db-1_SdDg_0 -full
```

```
deleting disk group db-1_SdDg_0
- fs /notesdata/db_0 ... deleted
- hostvol db-1_SdDg_0/db_SdHv_0 ... deleted
- dg db-1_SdDg_0 ... deleted
- LUN itsotuc3:/vol/vol_DominoDB/db_SdLun_0 ... deleted
- LUN itsotuc3:/vol/vol_DominoDB/db-1_SdLun_0 ... deleted
```

```
[root@domino1 ~]#
```

- vg: The volume group name of the Snapshot you wish to disconnect.
- full: Use the -full option to tell SnapDrive for UNIX to disconnect the Snapshot copy even if a host-side entity is named that includes other entities; for example, a volume group is named that has one or more host volumes. As an best practice, always disconnect volume groups in junction with the -full option.

Renaming the Snapshot

You can use the **snapshot snap rename** command to change the name of an existing Snapshot copy. The administrator may need to change the Snapshot name if he is required to do so.

If you rename one of these Snapshot copies, you must also rename all the related Snapshot copies using the same name. This is because SnapDrive for Linux and UNIX uses a short name when it creates the Snapshot copy, even though it spans multiple storage systems or volumes. The **rename** command changes the name of the current Snapshot copy, but it does not change the name of the related Snapshot copies in the other locations.

In Example 8-7, we rename the Snapshot **snap1** at **itsotuc3:/vol/vol_DominoLog** to **snappi1**.

Example 8-7 Renaming the Snapshot with SnapDrive

```
[root@domino1 ~]# snapdrive snap rename itsotuc3:/vol/vol_DominoLog:snap1 snappi1
snap rename renamed Snapshot itsotuc3:/vol/vol_DominoLog:snap1 to new name snappi1
[root@domino1 ~]#
```

Note: If you rename a Snapshot in a Snapshot group, remember to rename all Snapshots in the group. This makes restoring the Snapshot group an easier task.

Deleting the Snapshot

The **snapdrive snap delete** command removes the Snapshot copies you specify from a storage system. This command does not perform any operations on the host. It only removes the Snapshot copy from a storage system, if you have permission to do so. Example 8-8 shows the deletion of our two created Snapshots with the name **snap1**.

Example 8-8 Deleting Snapshots with SnapDrive

```
[root@domino1 ~]# snapdrive snap delete -snapname
itsotuc3:/vol/vol_DominoLog:snap1 itsotuc3:/vol/vol_DominoDB:snap1
snap delete: deleted snapshot itsotuc3:/vol/vol_DominoLog:snap1
snap delete: deleted snapshot itsotuc3:/vol/vol_DominoDB:snap1
[root@domino1 ~]#
```

- `-snapname`: Long name of the Snapshot. This option is stackable, as shown in the example.

Note: For more options for the `snapdrive snap` command, see the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>

8.2.1 Restoring Snapshot copy with SnapDrive

The `snapdrive snap restore file_spec` command restores data from the Snapshot copy you specify on the command line to the storage system. This operation replaces the contents of the `file_spec` that you specified on the `snap restore` command line with the contents of the `file_spec` arguments found in the specified Snapshot copy.

You can also restore Snapshot copies for non-existent `file_spec` arguments. This happens when the value you specify no longer exists on the host, but existed when you took the Snapshot copy. For example, it might be a file system that you have now unmounted or a disk group that you have removed.

Normally, you restore Snapshot copies from the host where you took the Snapshot copies (in other words, the originating host). You can also restore Snapshot copies using a different, or non-originating, host. For more details, see the *IBM System Storage N series SnapDrive for UNIX 2.2 Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>.

Note: When you restore Snapshot copies for volume groups, or for host volumes and file systems that are created on them, SnapDrive for Linux and UNIX restores the whole volume group. If you specify part of a volume group, SnapDrive still restores the entire volume group. An error occurs if you enter only a subset of the host volumes or file systems in each volume group on the command line. You can include the `-force` option to override this error; however, SnapDrive then restores the entire disk group.

In Example 8-9, we restore Snapshot `snap1` of the Lotus Domino database and transactional log storage.

Example 8-9 Restoring the Lotus Domino database and transactional log storage from Snapshot

```
[root@domino1 /]# snapdrive snap restore -vg db_SdDg log_SdDg -snapname snap1
Starting to restore /dev/mapper/db_SdDg, /dev/mapper/log_SdDg
WARNING: This can take several minutes.
DO NOT CONTROL-C!
If snap restore is interrupted, the filespecs
being restored may have inconsistent or corrupted
data.
For detailed progress information, see the log file /var/log/sd-recovery.log
Importing db_SdDg, log_SdDg
Successfully restored snapshot snap1 on 2 filer volumes:
itsotuc3:/vol/vol_DominoLog
itsotuc3:/vol/vol_DominoDB

disk group db_SdDg containing host volumes
db_SdHv (filesystem: /notesdata/db)
disk group log_SdDg containing host volumes
log_SdHv (filesystem: /notesdata/log)
```

```
[root@domino1 /]#
```

- -vg: The volume group name of the Snapshot you wish to restore. This option is stackable.
- -snapname: Name of the Snapshot. You can use the short (for example: snap1) or long (for example: itsotuc3:/vol/vol_DominoLog:snap1) name. If the volume entity on your host no longer exists, you must specify the long name.

Reset the snap reserve

Every volume created contains a *snap reserve* for further usage. This snap reserve is exclusive to Snapshots and cannot be used as active data storage.

By default, the snap reserve option for Data ONTAP is 20%. When you use Data ONTAP in an FCP or iSCSI environment, we recommend that you reset the snap reserve option to 0% on all storage system volumes holding SnapDrive for UNIX LUNs. The volume can be increased or decreased based on the situation and business needs.

To reset the snap reserve option on the storage system, complete the following steps:

1. Log in to your IBM System Storage N series storage system with SSH (you have to start it first) or Telnet.
2. Enter the following command to list all the volumes and their Snapshot reserves:

```
snap reserve
```

Example 8-10 shows the output from our lab configuration.

Example 8-10 Command-line output of snap reserve

```
itsotuc3*> snap reserve
Volume vol0: current snapshot reserve is 20% or 4194304 k-bytes.
Volume vol_DominoDB: current snapshot reserve is 20% or 91855256 k-bytes.
Volume vol_DominoLog: current snapshot reserve is 20% or 7969176 k-bytes.
itsotuc3*>
```

3. Reset the snapshot reserve space to 0% by entering the following command:

```
snap reserve vol_name 0
```

vol_name is the name of the volume on which you want to set the snap reserve option (see Example 8-11).

Example 8-11 Reset the snapshot reserve to 0%

```
itsotuc3*> snap reserve vol_DominoDB 0
itsotuc3*> snap reserve vol_DominoLog 0
itsotuc3*> snap reserve
Volume vol0: current snapshot reserve is 20% or 4194304 k-bytes.
Volume vol_DominoDB: current snapshot reserve is 0% or 0 k-bytes.
Volume vol_DominoLog: current snapshot reserve is 0% or 0 k-bytes.
itsotuc3*>
```

8.2.2 Scheduling Snapshot Using SnapDrive

To schedule SnapDrive Snapshot copies, complete the following steps:

Note: All steps except Step 1 in the following procedure are performed using the Scheduled Task Wizard, a Windows task scheduling tool available on your Windows server.

Steps for scheduling Snapshot copies for Windows

- ▶ Create a batch file (a file with a .bat extension) containing the following command on the Windows host on which you are scheduling Snapshot copies:

```
sdcli snap create [-m MachineName] -s SnapshotName -D  
DriveLetterList [. . .] [-x]
```

MachineName is the name of the Windows host on which the command will be executed. If no machine name is specified, the command is executed on the local machine.

SnapshotName is the name of the Snapshot copy to be created.

DriveLetterList is a list of space-separated drive letters.

When the -x flag is specified, Snapshot copies are created only for the drives specified by the -D flag. Otherwise, Snapshot copies are created for all the disks on the storage system volumes used by the listed drives.

Example 8-12 shows how to create a Snapshot copy named Jun_13_03 for each volume containing one or more of the LUNs mapped to the specified drives (that is, J:, K:, and L:). The Snapshot copies created are consistent for all LUNs contained by those volumes.

Example 8-12 Snap Create example

```
sdcli snap create -s Jun_13_03 -D j k l
```

- ▶ Select **Start** → **Settings** → **Control Panel** → **Scheduled Tasks**.
- ▶ Double-click **Add Scheduled Task**. The Scheduled Task Wizard is launched. Click **Next**.
- ▶ After the next window appears, click **Browse** and navigate to the folder where the batch (.bat) file you created in Step 1 is located.
- ▶ Select the batch file.
- ▶ After the next window appears, select from the list of frequencies, and then click **Next**.
- ▶ After the next window appears, enter a start time and complete the detailed frequency parameters. The option details displayed in this window vary depending on the Snapshot copy frequency you picked in the previous window.
- ▶ In the next window, type the user name (the administrator account name and password, repeated for confirmation) and then click **Next**.

Steps to schedule Snapshot copies for UNIX

Snapshots can be scheduled using cron directories. Refer to Example 8-13

Example 8-13 Scheduling Snapshot for UNIX

```
#!/bin/sh
```

```
DATE=`date`  
DATEHOUR=`date +%I`  
LOGFILE="/tmp/cronsnap.log"
```

```
echo $DATE >> $LOGFILE
snapdrive snap create -vg log_SdDg db_SdDg -snapname snap_h_$DATEHOUR -force
-noprompt >> $LOGFILE
echo ----- >> $LOGFILE
```

8.2.3 Scheduling Snapshot Using SnapManager for Microsoft Exchange

When you specify a full database backup operation or a database verification operation for databases that reside on a volume that is a SnapMirror source volume, SnapManager automatically provides you with the option to automatically perform a SnapMirror update after the operation finishes successfully.

To schedule a one-time SnapManager backup with SnapMirror replication, complete the following steps:

- ▶ Ensure that SnapDrive is properly licensed for use with SnapMirror.
- ▶ Configure SnapMirror on both the source volume to be replicated and its destination volumes.
- ▶ Disable the SnapMirror replication schedule on the storage system so that SnapDrive will monitor when a Snapshot is taken, and then initiate a replication in response. To do this, modify the schedule parameter (in the `/etc/snapmirror.conf` file or in storage system view mode) to indicate to SnapMirror on the storage system that no schedule is set. Use “- - -” as the unscheduled time.
- ▶ Begin specifying the full database backup operation or the database verification operation.

Note: It is possible to set the SnapMirror update schedule elsewhere, that is, in the SnapMirror `/etc/snapmirror.conf` file on the storage system. However, when SnapMirror is used by SnapManager, you must rely on the SnapManager backup schedule to drive the SnapMirror replication updates.

- You can use either the Backup and Verification tab or the Backup Wizard to specify the details of the operation.
- Be sure to enable the SnapMirror option.
- Be sure to schedule the operation rather than start it immediately.
 - If you are using the Backup and Verification tab, click **Schedule** instead of Backup Now or Verify Now.
 - If you are using the Backup Wizard, deselect the **Now** option in the When to Run this Operation window.
- ▶ Use the Schedule dialog box to specify when you want this mirrored backup operation to be run.

8.3 Scheduling Snapshot using FilerView

Log on to the IBM System Storage N series storage system. Select **Volume** → **Snapshot** → **Configure**. Refer to Figure 8-14.

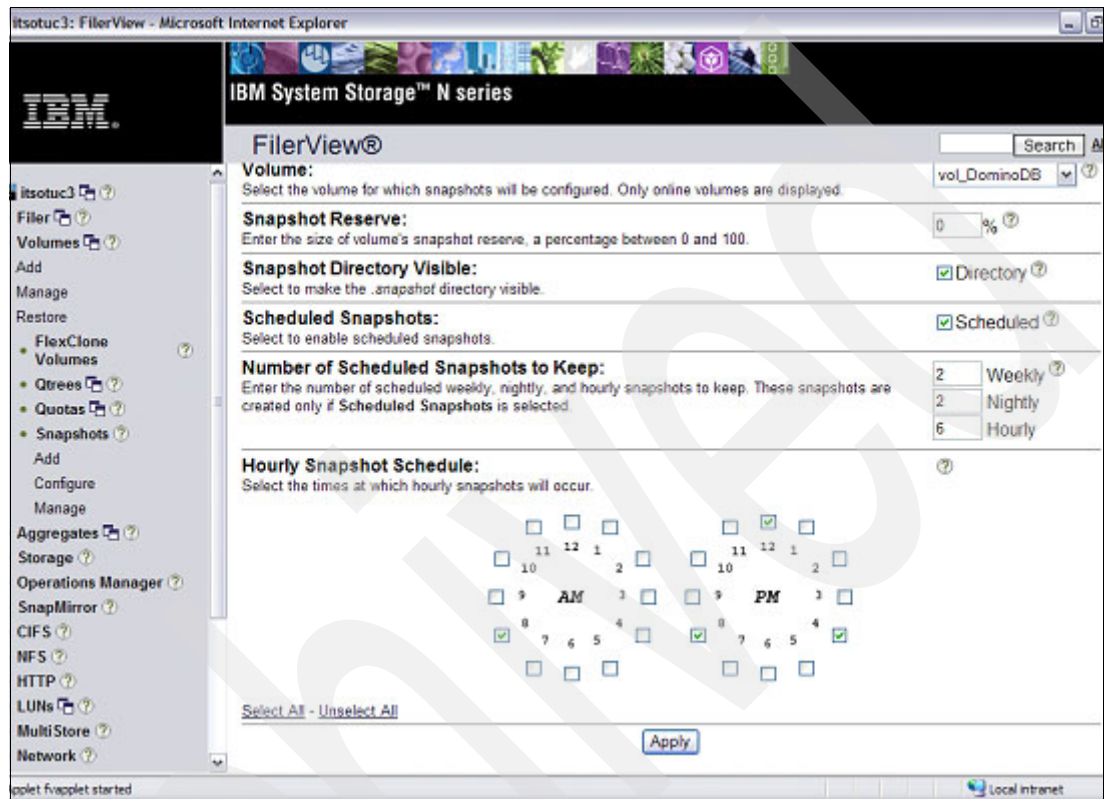


Figure 8-14 Configuring Snapshot

Select the schedule when Snapshots need to be taken. Refer to Figure 8-15 on page 303.

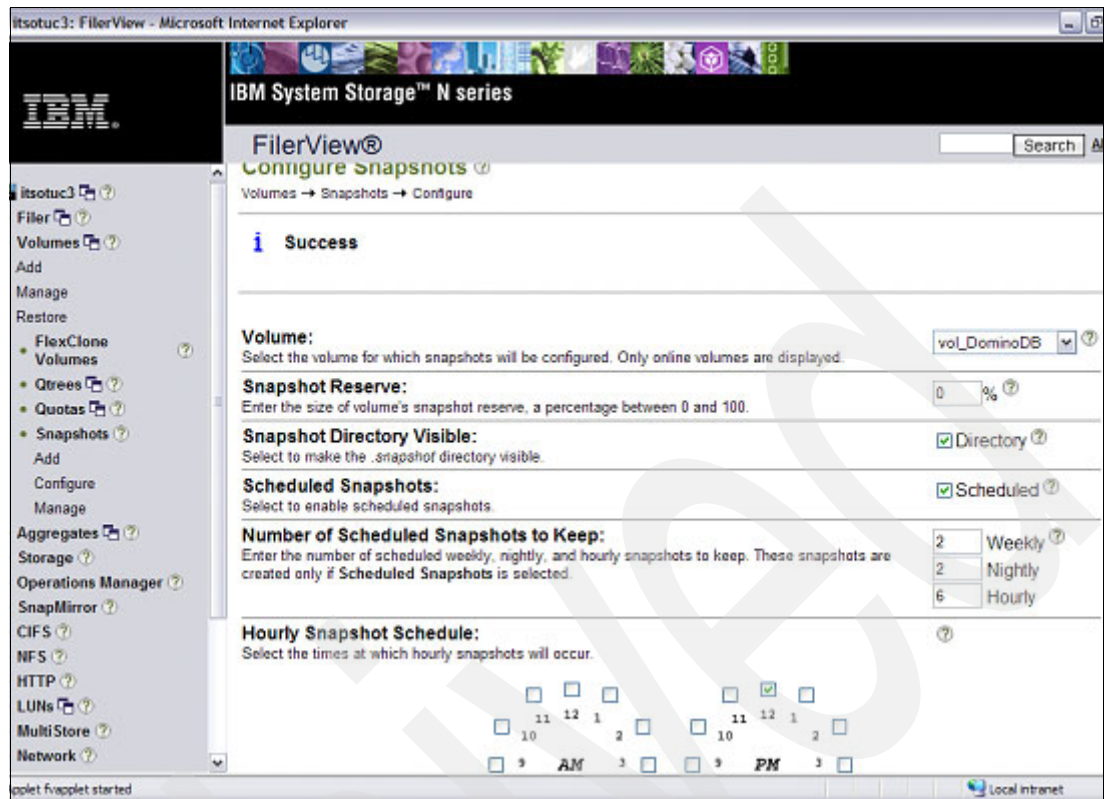


Figure 8-15 SnapShot Schedule

8.4 Performance commands and tools

The performance of the IBM System Storage N series can be measured with several tools and commands. There are N series console commands, such as the client using the **perfstat** script, using Operations Manager and SNMP. The performance of Lotus Domino and Microsoft Exchange will be affected if the N series storage that is used does not perform to the optimum level. Constant monitoring of storage is required to help the storage administrator fine-tune the storage so that the performance for applications like Lotus Domino and Microsoft Exchange does not deteriorate. The following N series commands will help you monitor the performance of the storage.

8.4.1 Data ONTAP commands

► Sysstat

The **sysstat** command reports aggregated N series performance statistics. Refer to Figure 8-16.

CPU	NFS	CIFS	HTTP	Total	Net kB/s		Disk kB/s		Tape kB/s		Cache age
					in	out	read	write	read	write	
43%	2470	0	0	2470	5218	7110	11123	48	0	0	4
70%	2873	0	0	2873	6887	11924	17474	16647	0	0	4
64%	3038	0	0	3038	7685	13129	14456	22148	0	0	4
48%	2954	0	0	2954	7757	13895	11972	168	0	0	4
54%	3511	0	0	3511	6606	12668	12352	0	0	0	4
54%	3910	0	0	3910	7327	14056	10004	24	0	0	4
68%	3094	0	0	3094	7151	13989	16020	14955	0	0	4
65%	3334	0	0	3334	6622	11755	13260	23165	0	0	4
76%	4207	0	0	4207	7330	16334	11738	628	0	0	4

Figure 8-16 *sysstat N series performance statistics*

► qtree stats

The **qtree stats** command shows how many NFS or CIFS operations are caused by user accesses into qtrees. Refer to Figure 8-17.

Volume	Tree	NFS ops	CIFS ops
-----	-----	-----	-----
vol0	unixhome	96414769	0
vol0	sims	610671	0
vol1	winhome	0	346978004
vol1	winapps	0	7036597
vol2	software	181653611	1757843416
vol2	tools	53109414	61724267
No qtrees are in use in volume vol3			

Figure 8-17 *qtree stats*

► statit

Gathers a set of performance statistics over an interval between the time **statit** is begun and ended for:

- CPU
- Network interfaces
- Disks
- System software

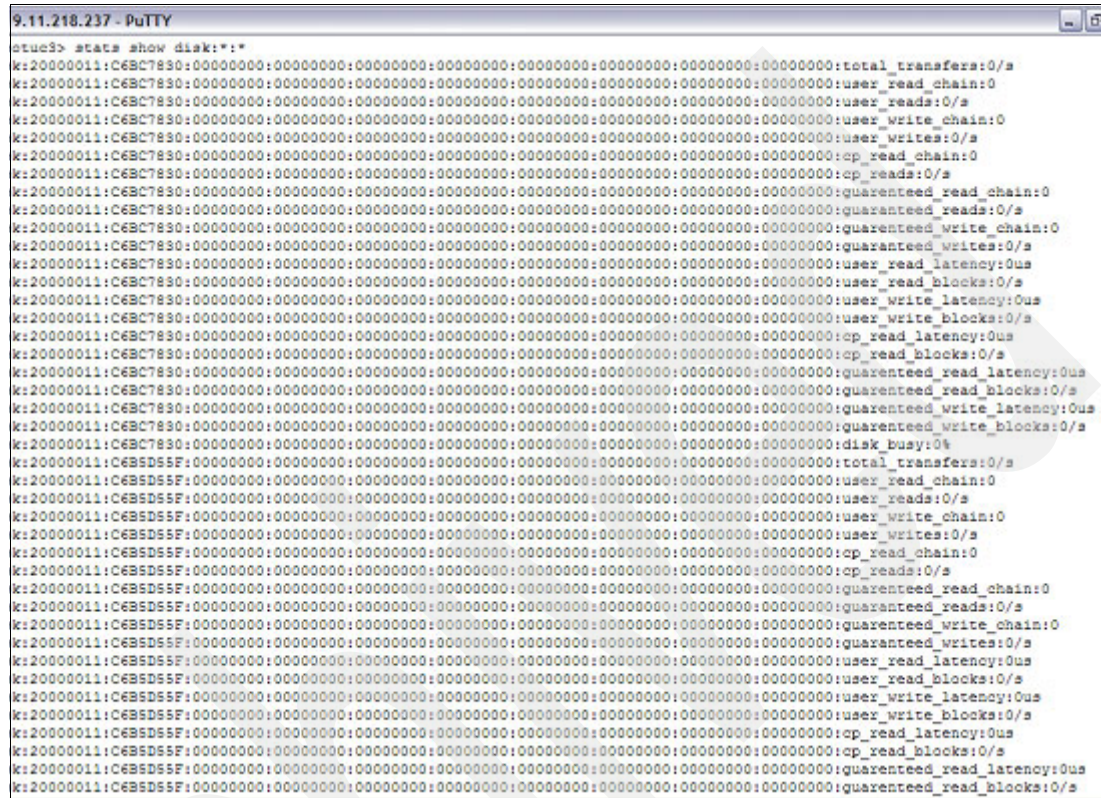
The **statit -b** command begins collecting data and the **statit -e** command ends the measurement interval and generates a report. Refer to Figure 8-18.

Multiprocessor Statistics (per second)			
	cpu0	cpu1	total
sk switches	88295.21	93934.69	182229.91
hard switches	10349.39	12997.44	23346.84
domain switches	1093.05	1120.04	2213.09
CP rupts	9730.45	975.84	10706.30
nonCP rupts	1502.80	35.94	1538.74
CP rupt usec	20213.94	10146.06	30360.01
nonCP rupt usec	2594.54	307.63	2902.17
idle	11631.81	14706.57	26338.38
kahuna	402646.60	353889.44	756536.10
network	249175.98	288404.39	537580.45
storage	57091.47	60831.06	117922.53
exempt	149351.86	155580.15	304932.03
raid	106830.52	115707.08	222537.42
target	247.80	232.26	480.06
netcache	0.00	0.00	0.00
netcache2	215.49	195.35	410.84
114.943430 seconds with one or more CPUs active			(99%)
1.842204 seconds with one CPU active			(2%)
113.101226 seconds with both CPUs active			(98%)

Figure 8-18 statit for multiprocessors

► **stats**

The **stats** command collects or views statistical data on Storage Systems. Refer to Figure 8-19.



```
9.11.218.237 - PuTTY
stuc3> stats show disk:*
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:total_transfers:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_chain:0
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_reads:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_chain:0
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_writes:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_chain:0
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_reads:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_chain:0
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_reads:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_write_chain:0
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_writes:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_latency:0us
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_blocks:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_latency:0us
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_blocks:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_latency:0us
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_blocks:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_latency:0us
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_blocks:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_write_latency:0us
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_write_blocks:0/s
k:20000011:C63C7830:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:disk_busy:0%
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:total_transfers:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_chain:0
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_reads:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_chain:0
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_writes:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_chain:0
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_reads:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_chain:0
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_reads:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_write_chain:0
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_writes:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_latency:0us
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_read_blocks:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_latency:0us
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:user_write_blocks:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_latency:0us
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:cp_read_blocks:0/s
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_latency:0us
k:20000011:C63D55F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000:guaranteed_read_blocks:0/s
```

Figure 8-19 stat show command

8.5 Operations Manager

The integrated management approach enables you to monitor and manage your IBM System Storage N series storage systems from a single location using the following management tools:

- Operations Manager, a Web-based user interface from which you can monitor and manage multiple storage systems and N series clusters.
- Operations Manager, which provides infrastructure services for various clients, for example, Performance Advisor.
- FileView software, which manages individual storage systems and clusters of storage systems.
- Performance Advisor, which runs on the Operations Manager Client (Java™ thick client) and provides a single location to view comprehensive storage system and vFiler unit performance information.
- Operations Manager Host Agent, an independent software application that resides on a host with which Operations Manager interacts.

In a environment where multiple mail servers are running on multiple IBM System Storage N series storage systems, Operations Manager provides a single point of view for monitoring N series.

8.5.1 Operations Manager

Operations Manager provides infrastructure services such as discovery, monitoring, role-based access control, auditing, and logging for products in the Storage and Data Suites. Operations Manager software runs on a separate workstation or server. It does not run on the storage systems.

Operations Manager has a command-line interface for scripting commands that might otherwise be performed through a Web-based user interface, known as Operations Manager.

Figure 8-20 shows the terminology and the architecture of the product, with respect to Operations Manager (the Web-based user interface), the server, and the clients. You can install Operations Manager on a separate workstation or on the same workstation as the server.

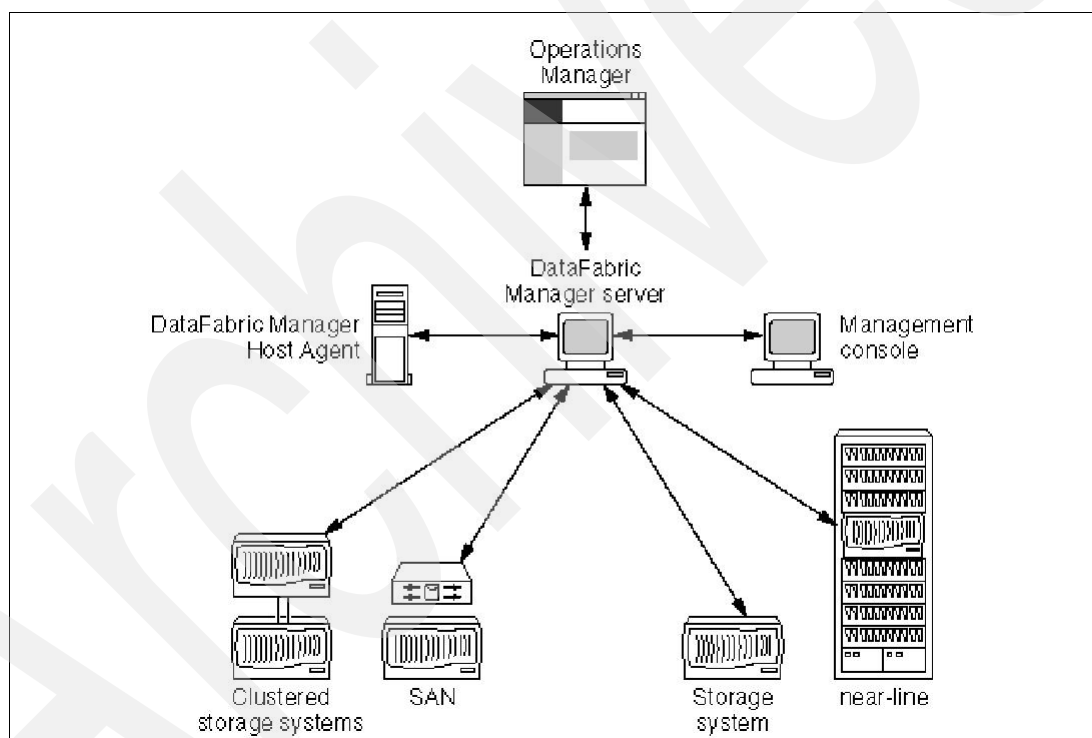


Figure 8-20 Operation Manager

Operations Manager is the Web interface used for day-to-day monitoring, alerting, and reporting on storage system infrastructure.

Feature of Operations Manager

- Discovery

You can configure Operations Manager for the discovery of storage systems, volumes, qtrees, LUNs, disks, quotas, and then view the results filtered by role-based access control and user-defined grouping.

- Monitoring and reporting

You can monitor device or object health, capacity utilization, and performance. You can also view or export reports with the relevant information and create custom reports.

- Alerting

You can configure alerts and thresholds for event management. Operations Manager issues alerts and Operations Manager generates event reports for monitored systems, volumes, qtrees. Operations Manager alerts are sent through e-mail, pager, or by generating SNMP traps to be sent to other monitoring applications.

- Management

Use Operations Manager to:

- Group devices, vFile units, host agents, volumes, qtrees, and LUNs into meaningful groups for ease of management. The groups are stored within Operations Manager and are shared with other client applications.
- Configure role-based access control settings.
- Define group configuration management templates and apply those templates to one or more systems.
- Edit volume, qtree, or user quotas.
- Run Data ONTAP CLI commands simultaneously on multiple systems.

Operation manager also add volumes. Refer to Figure 8-21.

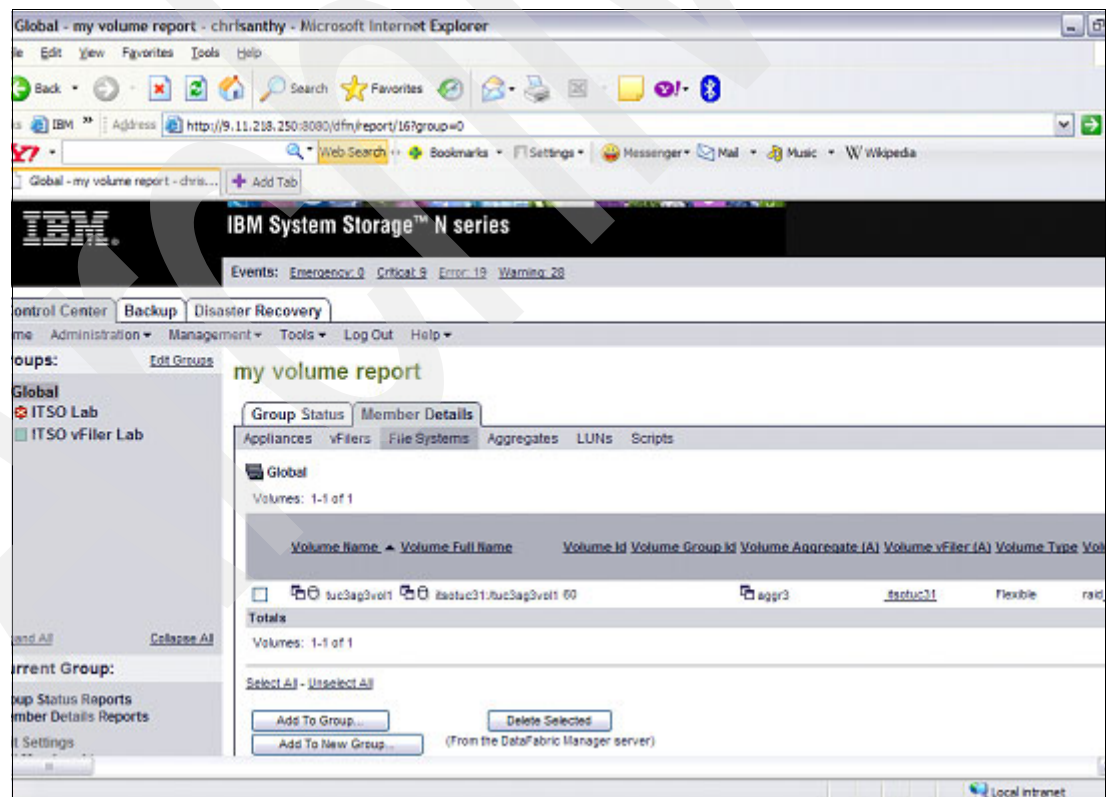


Figure 8-21 Addition of Volumes Operations Manager

Operations Manager back ups schedule operations. Refer to Figure 8-22.

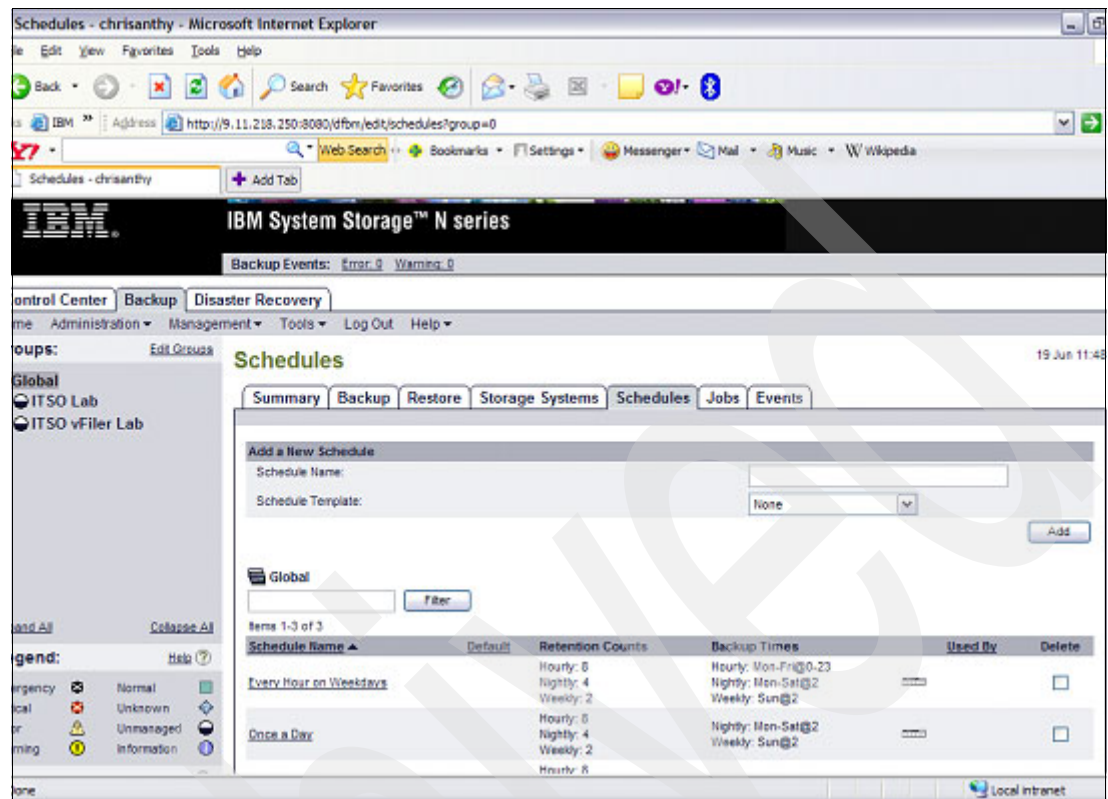


Figure 8-22 Scheduling backup operations

SnapMirroring is used for data recovery and data backup using the disaster recovery manager in Operations Manager. A volume that needs to be backed up also needs to be mirrored. The backup is done in a predetermined schedule. Refer to Figure 8-22.

Refer to Figure 8-23 for information about the Operation Manager Mirroring operation.

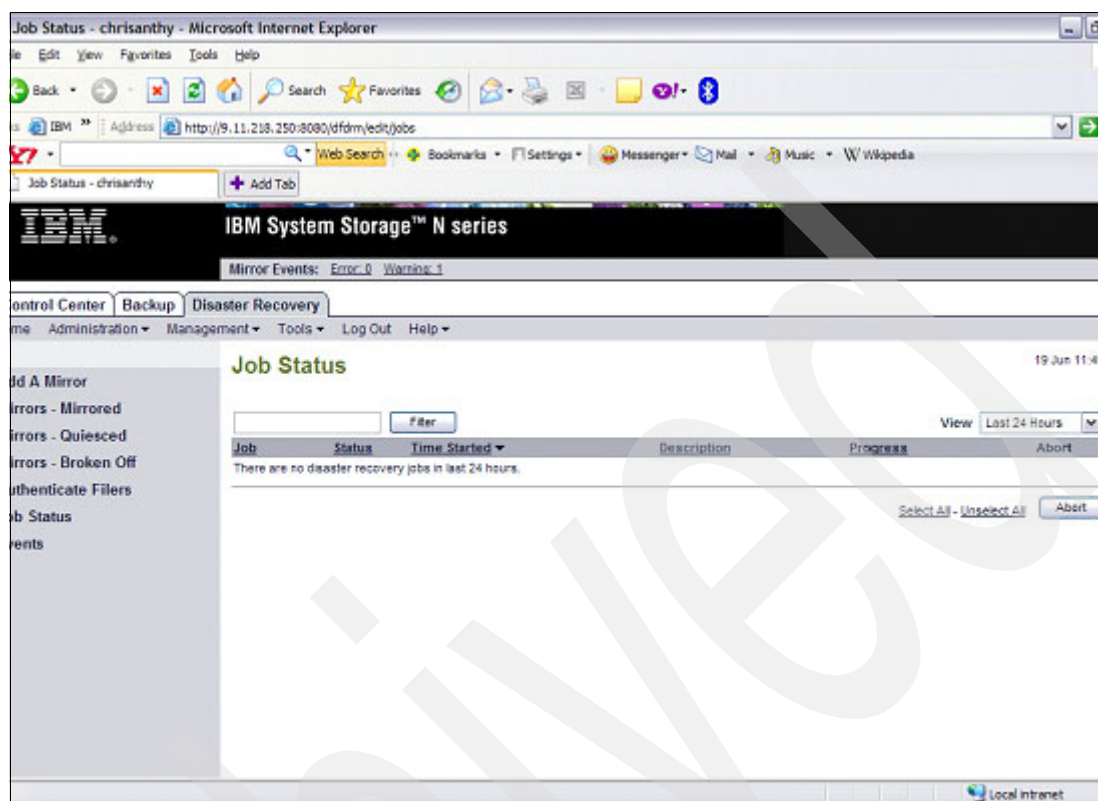


Figure 8-23 Operation Manager Mirroring Operation

8.6 FlexClone

A FlexClone volume is a writable point-in-time image of a flexible volume or another FlexClone volume. FlexClone volumes add a new level of agility and efficiency to storage operations. They take only a few seconds to create and are created without interrupting access to the parent flexible volume. FlexClone volumes use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and clone.

Lotus Domino and MS Exchange use FlexClone for application testing with live data, data mining, parallel processing, online backup, system deployment, and IT Operations.

8.6.1 Creating FlexClone

Example 8-14 shows how to create a new read-write clone named Widget, based on an existing flexible volume named Gadget. The command reuses an existing Snapshot copy in Gadget as the backing for the clone.

Example 8-14 Create Flex™ Clone

```
vol clone create Widget -s file -b Gadget GadgetSnap
```

8.6.2 Status of FlexClone

The command in Example 8-15 prints verbose status information for the volume named Widget. The volume can be a flexible volume, FlexClone volume, or traditional volume. It prints less information without the -v flag.

Example 8-15 FlexClone Status

```
vol status -v Widget
```

8.6.3 Splitting the FlexClone

The command in Example 8-16 starts the process of splitting the FlexClone volume named Widget from its parent flexible volume. The split will make a copy of all blocks associated with the FlexClone volume that currently reside in the parent flexible volume.

The process can take a while to complete and proceeds in the background. Afterwards, the volume will be an independent flexible volume with space guarantees reenabled.

Example 8-16 FlexClone initiation

```
vol clone split start Widget
```

The command checks the progress of the ongoing clone split for the Widget volume.

The command **vol clone split status Widget** stops the clone split process. Any blocks that have already been copied will continue to be associated with the clone volume. Another **split start** command will resume the split with any blocks that are still shared.

Example 8-17 Vol clone split status command

```
vol clone split stop Widget
```

8.6.4 Space utilization with FlexClone

The command in Example 8-18 displays disk usage associated with aggregates. The example limits it to the aggregate named aggr1.

Example 8-18 Disk Usage

```
df -A aggr1
```

The command in Example 8-19 shows the command to display the name of the aggregate containing the volume named Widget.

Example 8-19 Displays aggregate name

```
vol container Widget
```

8.7 Reallocation

Data ONTAP 7G introduces a new **reallocate** command to manage reallocation activities. Reallocation scans ensure that blocks in a LUN or a volume are laid out sequentially. This improves the performance of read and write commands, thereby improving the read/write performance of LUN-based applications that access data on the N series.

Lotus Domino and Microsoft Exchange uses reallocation to improve performance, as this feature leads to faster access to the disk. Some of the important uses of a reallocation scan for Lotus Domino and Microsoft Exchange are:

- ▶ Runs in the background so applications can continue to run.
- ▶ Evaluates the layout of blocks in a LUN to determine the current level of optimization.
- ▶ Rewrites blocks as necessary in order to achieve and maintain the desired level of optimization.
- ▶ Dynamically adjusts to avoid taking too many system resources while reallocating the volume or file.
- ▶ Running a LUN reallocation scan run before the creation of a Snapshot copy optimizes both the LUN and the Snapshot copy.

Example 8-20 shows how to reallocate and schedule a reallocation of a LUN.

Example 8-20 LUN reallocation

Check the LUN /vol/db1/lun1 allocation periodically.
`reallocate start /vol/db1/lun1`

Schedule a reallocation job at 11 pm every Saturday.
`reallocate schedule -s "0 23 * 6" /vol/db1/lun1`

8.8 User management

The N series supports role-based access control (RBAC). RBAC is a method for managing the set of actions that a user or administrator may perform in a computing environment. The feature is important for applications like Lotus Domino and Microsoft Exchange where data security is paramount. Only users who have authority should have access to modify the storage configurations.

Only certain users should be allowed certain tasks. RBAC solves this management problem by allowing you to define sets of capabilities (roles) that are not assigned to any particular user. Users are assigned to groups based on their job functions, and each group is granted the set of roles required to perform those functions. Using this method, the only configuration required for an individual administrator is to ensure that the administrator is a member of the appropriate groups; that administrator will inherit all the correct capabilities because of the group membership and the roles assigned to those groups.

First, we need to define the roles. There are some functions that all of the administrators should have access to, such as the login capabilities and certain reporting commands (for example, **df**). Rather than create a role for each of these groups and duplicate the common capabilities across all three groups, we will put the common capabilities in a role that will be assigned to all three groups. Then we can give each group a second, unique role that will

have the capabilities specific to that group. Refer to Example 8-21 on page 313 and Example 8-22 on page 313.

Example 8-21 Assigning roles

```
itsotuc> useradmin role add fileradmin -c "Capabilities common to all administrators" -a login-telnet, login-ssh, cli-df*, cli-help*, cli-man*, cli-sysstat*, cli-uptime*,cli-options*,cli-passwd*
```

Example 8-22 Storage Administrator role assignment

```
itsotuc> useradmin role add storageadmin -c "Storage administrators" -a cli-*,security-*, api-*, login-console, login-http-admin
```

For individual administrators, in order to provide centralized password management, we will use domain users instead of local users. Refer to Example 8-23 and Example 8-24.

Example 8-23 Adding individual users

```
itsotuc> useradmin domainuser add sven -g fileradmins
```

Example 8-24 Adding a storage administration user

```
itsotuc# useradmin domainuser add prashant -g storageadmin
```

Note: The user names above (prashant and sven) must be valid accounts within the Windows domain of which the N series system is a member.

8.9 Autosupport

Autosupport is a sophisticated, event-driven logging agent featured in the Data ONTAP operating system inside IBM System Storage N series storage systems that continuously monitors the health of your system. It keeps a watchful eye on a multitude of preset conditions. This feature is important for applications such as Lotus Domino and Microsoft Exchange, as the N series health is constantly monitored and e-mails are sent to the storage administrator about the N series' status. This will help the administrator solve the storage problems immediately and thus performance of Lotus Domino and Microsoft Exchange will not deteriorate.

The Autosupport feature triggers the automatic sending of notification messages to IBM Service and Support. Autosupport also has the ability to send notification messages to one or more customer-specified e-mail addresses, which can alert recipients to potential problems with the IBM System Storage N series storage system. As necessary, IBM Service and Support will contact customers based on the contact information in the customer's record for resolution of potential system problems.

8.9.1 Configuring Autosupport from FilerView

Open FilerView and select **Filer** → **Configure Autosupport**. Configure the mailhosts and mail address where notification needs to be sent. Refer to Figure 8-24.

The screenshot shows the IBM System Storage N series FilerView web interface in a Microsoft Internet Explorer browser window. The page title is "Configure Autosupport" and the breadcrumb is "Filer → Configure Autosupport". The left sidebar contains a navigation menu with options like "Show Status", "Manage Licenses", "Report", "Syslog Messages", "Use Command Line", "Configure Syslog", "Configure File System", "Configure Autosupport", "Test Autosupport", "Set Date/Time", "Configure Miscellaneous", "Shut Down and Reboot", "Show System Status", "Volumes", "Aggregates", "Storage", "DFM", "SnapMirror", "CIFS", "NFS", "HTTP", and "LUNs". The main content area contains the following configuration fields:

- Autosupport Enabled:** A dropdown menu set to "Yes".
- Mailhosts:** A table with one row containing the IP address "192.168.3.103".
- From:** A text field containing the email address "postmaster@itsa.tucson".
- To:** A table with one row containing the email address "administrator@itsa.tucson".
- Note To:** An empty text field.
- Buttons:** "Reset" and "Apply" buttons at the bottom.

The status bar at the bottom of the browser window shows "Local intranet".

Figure 8-24 Configuring Autosupport



Backup and restore of Microsoft Exchange servers using IBM System Storage N series

This chapter will help you configure and run backup and restore for Microsoft Exchange servers on IBM System Storage N series storage system using SnapDrive and SnapManager for Microsoft Exchange (SME) and how to recover single mailboxes using Single Mailbox Recovery (SMBR).

9.1 Backup and restore

Backup and restore operations are critical administrative tasks for Microsoft Exchange server systems. The backup should not only guarantee the recover of the data should a disaster occur, but also must guarantee the normal operation of the Microsoft Exchange environment by committing the transactional log files on the databases and deleting these log files after the operation.

Another point is the rapid growth of the Microsoft Exchange databases. In an enterprise, most customers utilize 20 databases (the maximum) to keep individual database sizes smaller for backup, recovery, and DR (repair/offline defrag). This large number of databases and mailboxes (and therefore the database and transaction log files) demands a considerable time for backup and for restore. Sometimes, the backup job can last more than 24 hours, which is useless.

The restore operation is also a concern in cases where you need to restore either the whole server or a single user mailbox. Depending on the size of the database, it can take some hours or days to get the data back on the Microsoft Exchange server.

SnapManager for Microsoft Exchange eases the backup and restore operations by using the N series Snapshot technology and providing quick and verified backup sets for the environment.

Single Mailbox Recovery for Microsoft Exchange provides the means to recover a single mailbox or even a single message on a specific mailbox, without recovering the whole Microsoft Exchange server or impacting the production performance.

Note: As a best practice, and in order to support up to the minute restore capabilities, the circular logging feature should not be enabled on the Microsoft Exchange Server.

9.1.1 Backup

Microsoft Exchange servers that have databases and transaction log files on the IBM System Storage N series storage system can be backed up on different ways depending on the environment and the planned resources.

Backup using SnapManager

SnapManager is the recommended application for backup, restore, and management of Microsoft Exchange servers installed and configured to use the IBM System Storage N series storage system.

SnapManager uses the N series Snapshot technology to enhance backup time and reliability.

Before installing SnapManager for Microsoft Exchange, verify that SnapDrive is already installed (with MPIO support if necessary). SnapManager must be installed on the Microsoft Exchange server itself so that the server can be configured, but can also be installed on a management server for the daily activities and verification process.

To configure backups using SnapManager, follow these steps:

1. On the Microsoft Exchange server, select **Start** → **All Programs** → **IBM** → **SnapManager for Exchange Management Console**. The SnapManager for Exchange window will be shown (as seen on Figure 9-1).

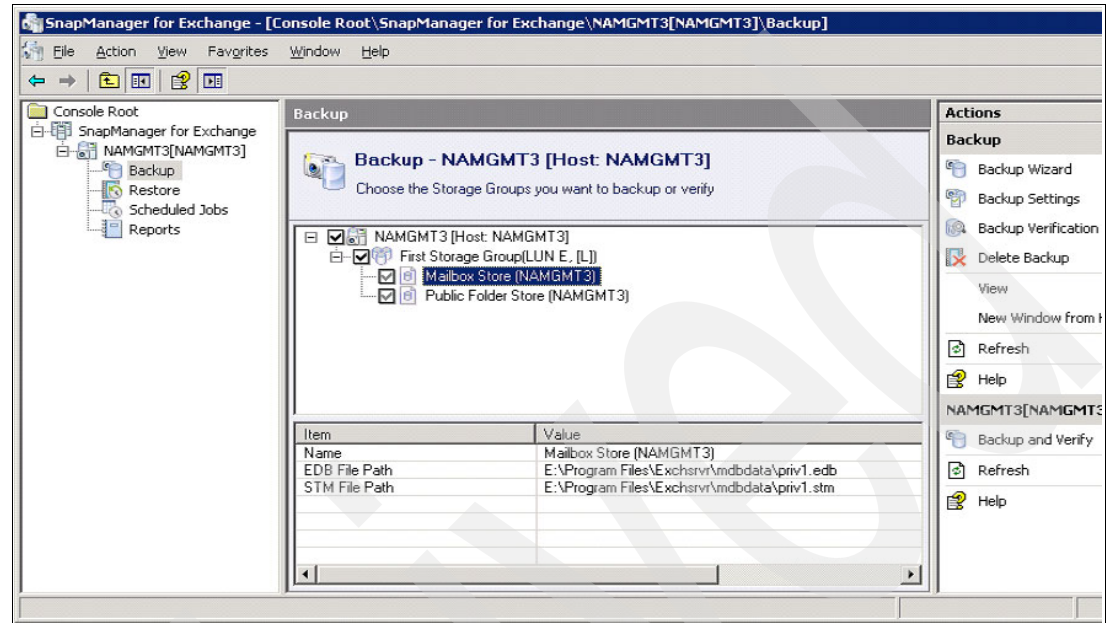


Figure 9-1 SnapManager for Exchange window

2. A backup operation can be manually configured or you can use the Backup Wizard to help you configure it. Select the target server on the left pane and on the Actions pane (right), click **Backup Wizard**. This will bring the Backup Wizard window (Figure 9-2). Click **Next** in this window.



Figure 9-2 Backup Wizard window

- [illegible]

4. In the Backup or Verify Databases and Transaction Logs window (Figure 9-4), select the operation to perform. If you select the option to **Verify the Databases** and transaction logs, you need to specify the number of backups to verify. At this point, the Backup Wizard will only verify the integrity and functionality of the existent backup sets. No backup job will be executed with this option. For backup, select the option **Back up databases and transaction logs** and click **Next**.

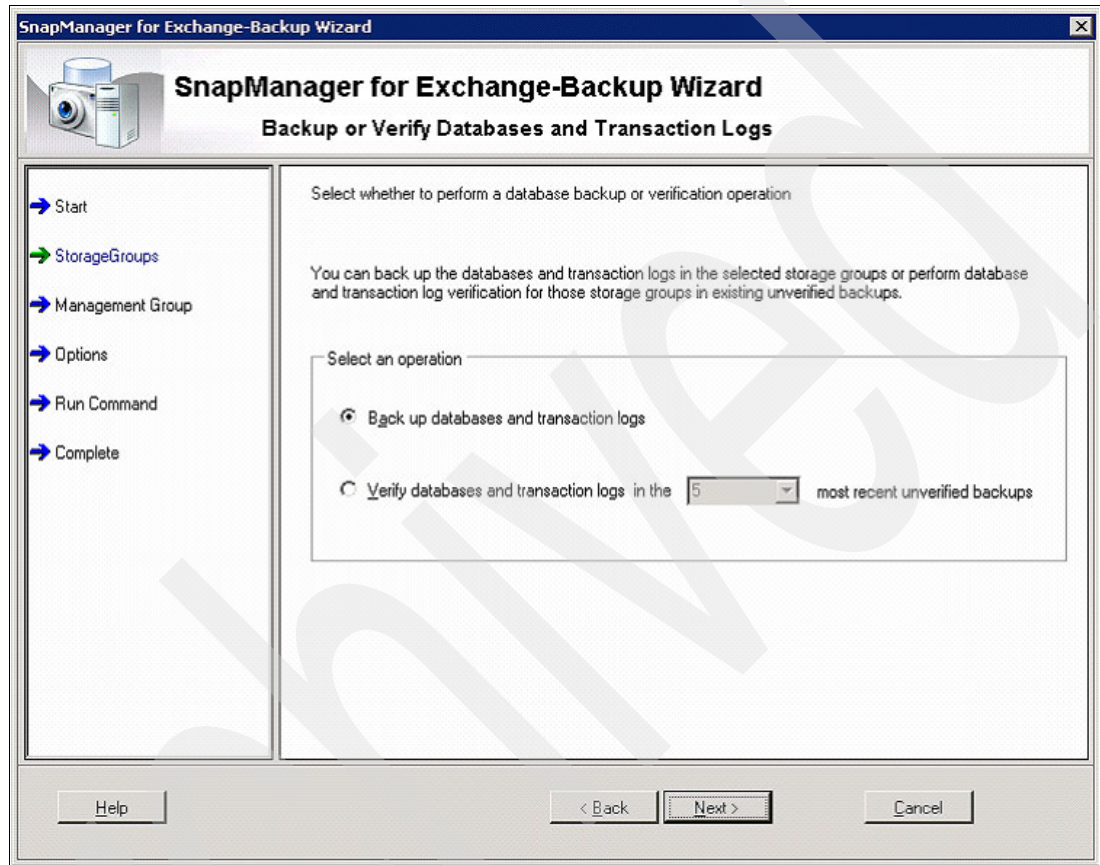


Figure 9-4 Backup or Verify window

5. The Backup Management Group window will be shown (Figure 9-5). Select the management group on which to include this backup. Click **Next**.

Note: The Backup Management Groups do not depend on or enforce how often the backup job is performed. It is rather a labelling convention for the Snapshots copies made by the SnapManager for Microsoft Exchange.

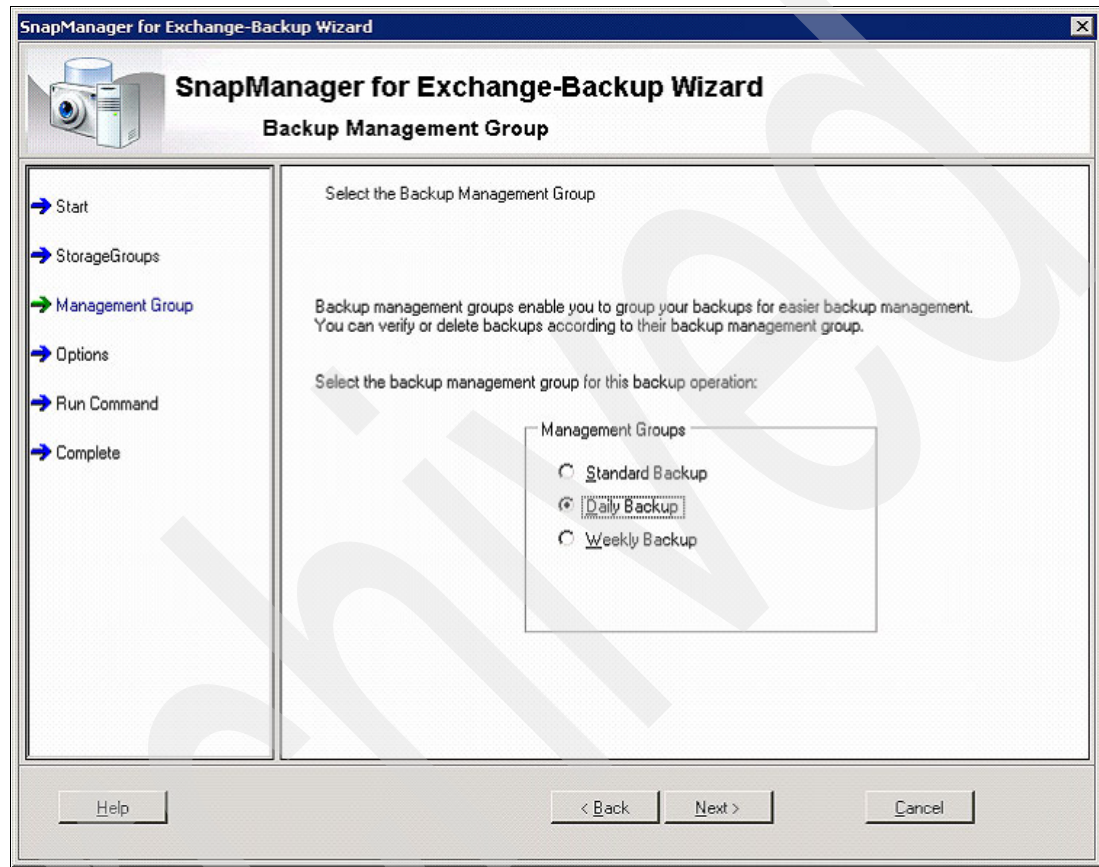


Figure 9-5 Backup Management Group window

6. On the Naming Convention window (Figure 9-6), select the convention to be used for naming the Snapshots copies created on the IBM System Storage N series storage system. If you select the generic naming convention, the newest Snapshot set will always have the `_recent` on the end of its name. If you select the unique time stamp, each Snapshot set will be named using the time stamp. Note that this configuration is to help you manage your backup set's names only. After select the desired name convention, click **Next**.

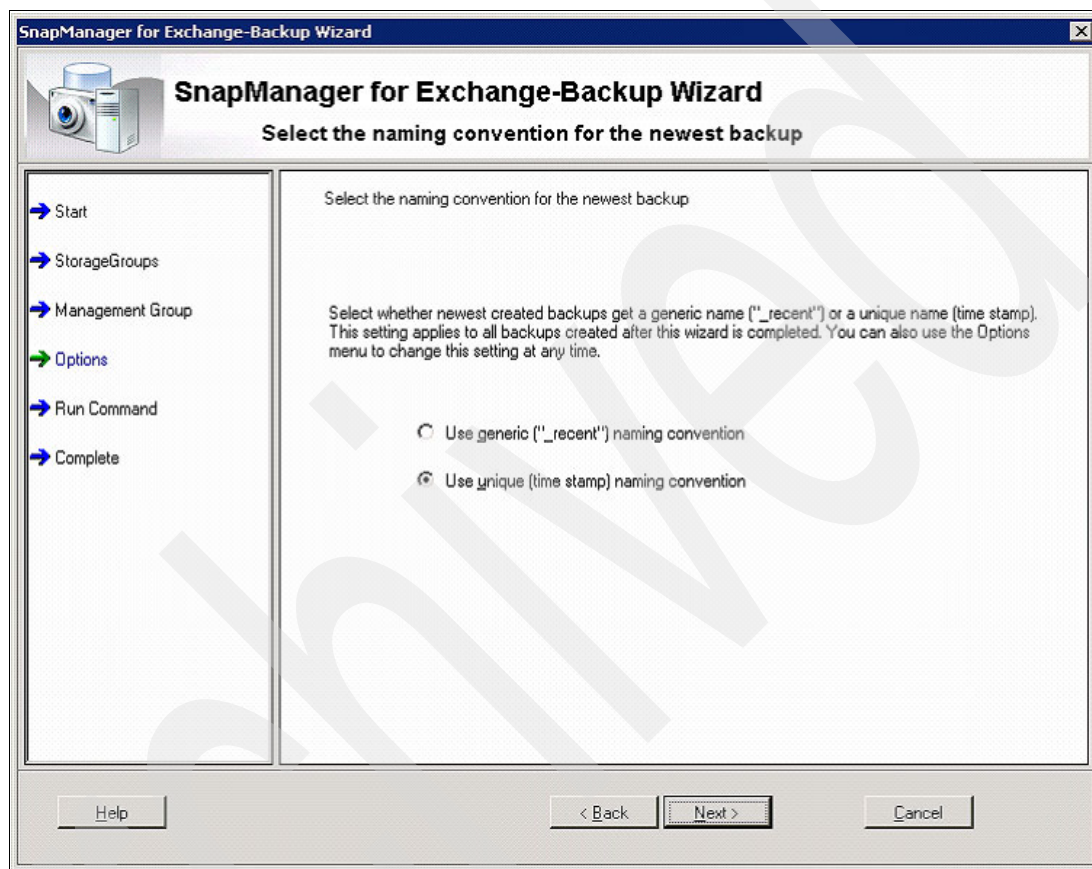


Figure 9-6 Naming Convention window

7. In the Delete Older Backups window (Figure 9-7), select the policy to be used to delete older backups. You can select to delete older backups (and then specify how many backup sets or backup days should be kept) or select to not delete older backups at all. This configuration should be done to meet your company's requirements for backup and restore. In the example scenario, we are keeping eight backup sets and deleting the oldest every time a backup job runs. After the configuration is done, click **Next**.

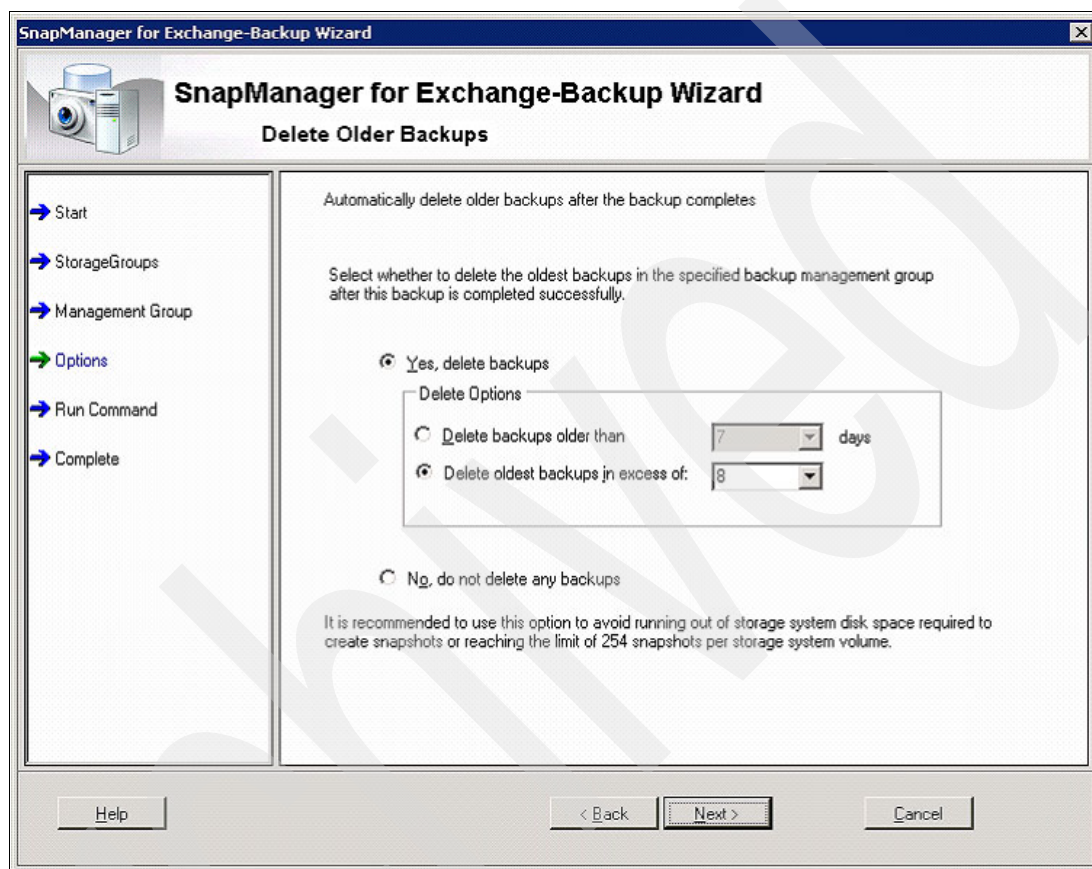


Figure 9-7 Delete Older Backups window

8. The next window will allow for the configuration of the restore capabilities (as seen in Figure 9-8). You can configure SnapManager for Microsoft Exchange to preserve all the transaction log files on the SnapInfo Directory to allow for up to the minute restore of the older backup sets. This directly impacts storage space usage. A different configuration is to remove the transaction log files on the SnapInfo Directory associated with the backup set when deleting that backup set. This can prevent backup sets older than the one that is being deleted to be restored up to the minute (because some transaction log files would be missing). This configuration can vary, depending on the business need, planned infrastructure, backup and restore policy, storage limits, and other variables. After selecting your configuration, click **Next**.

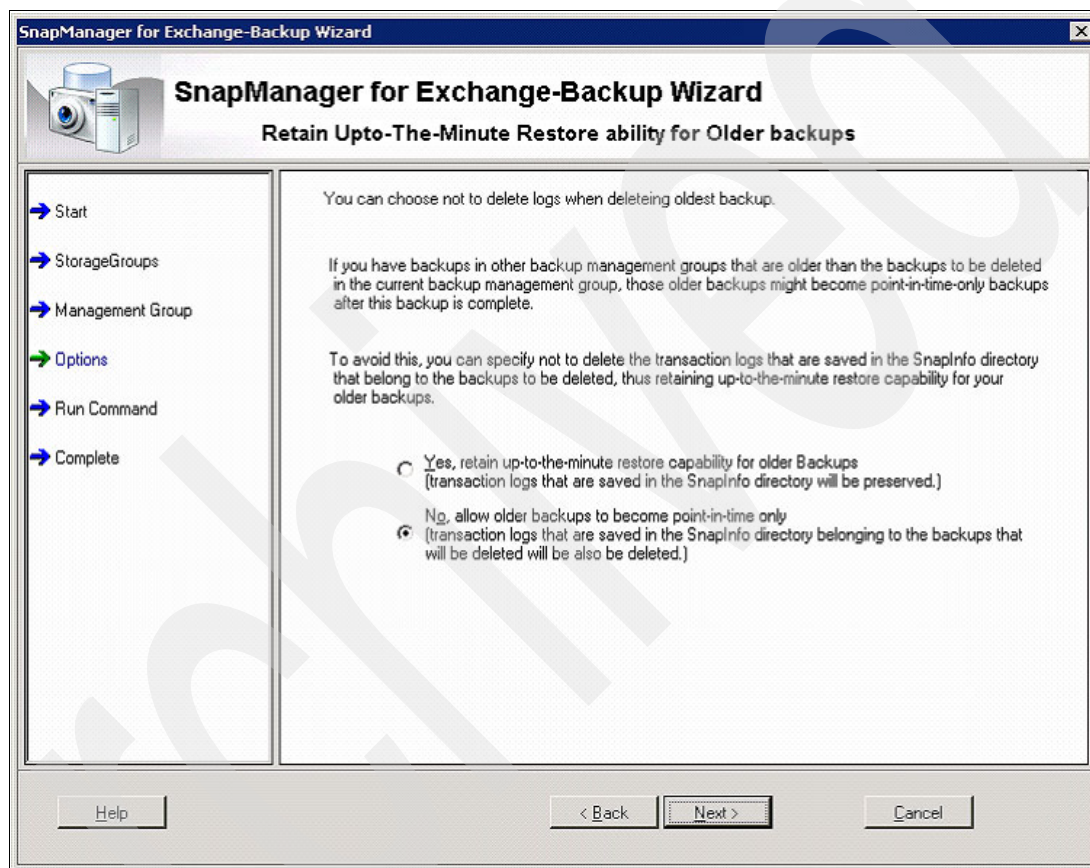


Figure 9-8 Retain Transaction Log files window

9. Because in our example scenario we configured the SnapInfo archived transaction log files to be deleted when removing the oldest backup sets, the SnapManager Warning Message window will be shown (as shown in Figure 9-9 on page 325) stating that older backups than the one being deleted will become point in time backups. Click **Yes** to confirm the configuration to remove the older transaction log files and click **Next** in the Retain Transaction Log files window (Figure 9-8).




Figure 9-9 SnapManager Warning Message window

10. In the Truncated Log Backup window (Figure 9-10), select whether the truncated log files (log files already committed to the database) should be backed up or not. Again, this configuration depends on the planned backup and restore policy and management. In the example scenario, we selected the truncated log files to be backed up so that any older backup set can be restored up to the minute. Click **Next**.

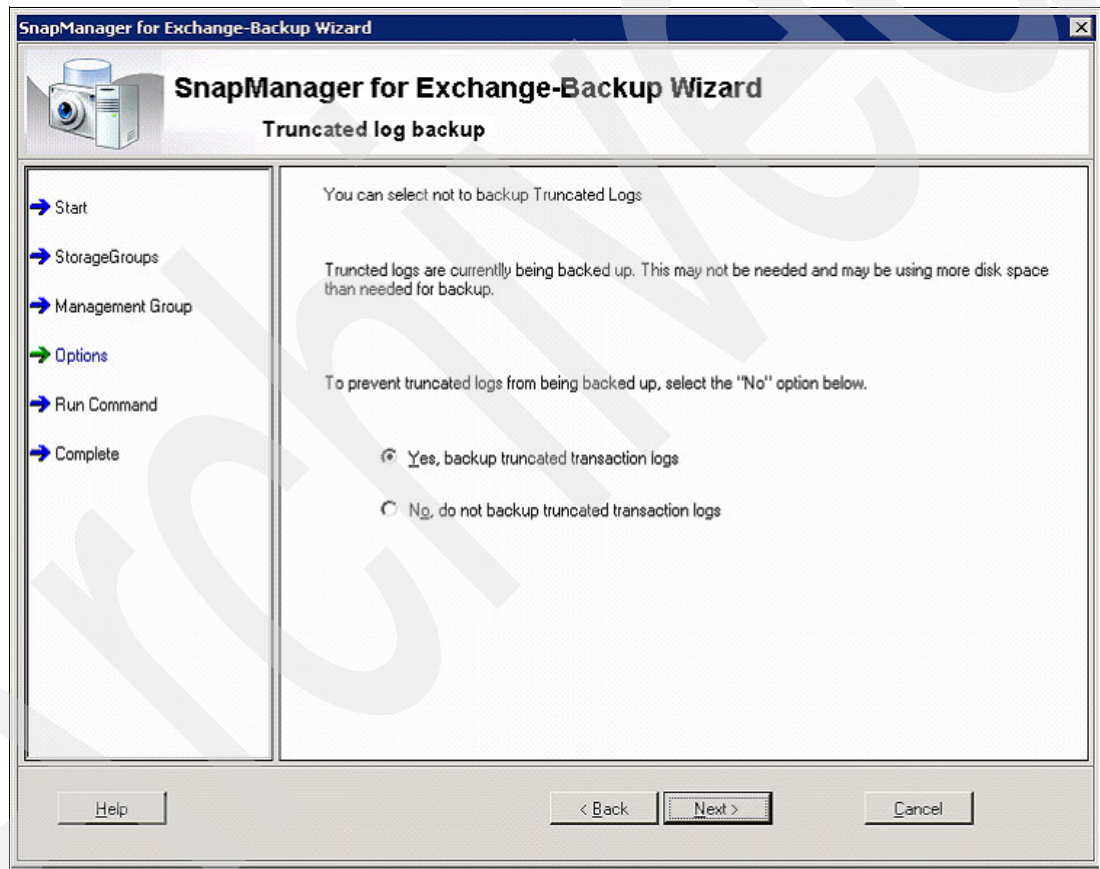


Figure 9-10 Truncated Log Backup window

11. In the Backup Verification window (Figure 9-11), select whether the backup will be verified after it is finished (and which server will be used for the verification process) or not. If you choose to not verify the backup now, you can still verify the backup later, using the verification option shown in the Backup or Verify window (Figure 9-4 on page 320). The recommendation is that every backup is verified to guarantee that you have a valid backup set should a disaster occurs. By default, the local server will be used to verify the backup set after it finishes. To select a remote verification server, make sure that the remote server has:

- a. SnapDrive installed
- b. The appropriate LUN driver installed (iSCSI or FCP)
- c. Connectivity to the IBM System Storage N series storage system (either using iSCSI or FCP)
- d. SnapManager installed and configured to use the same account used on the Microsoft Exchange production server
- e. The Exchange Management Tools installed

Note: The recommended configuration is to verify the backup sets as soon as they finish so that the Microsoft Exchange administrators have the guarantee of a valid backup set. Using remote validation server provides the best performance and does not impact the production Microsoft Exchange server.

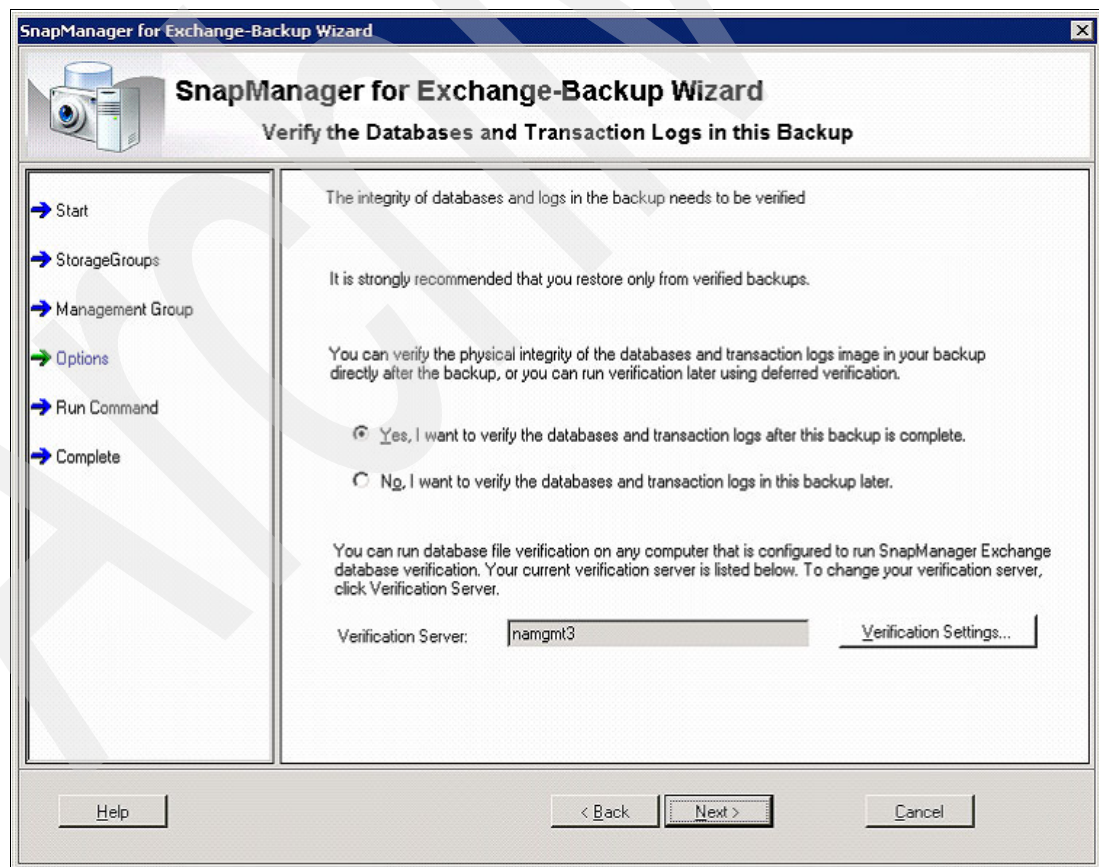


Figure 9-11 Backup Verification window

12. In the Run Command After Operation window (Figure 9-12), select whether you need to run a command after the backup operation is finished or not. In the example scenario, no command will be run after the backup job. Click **Next**.

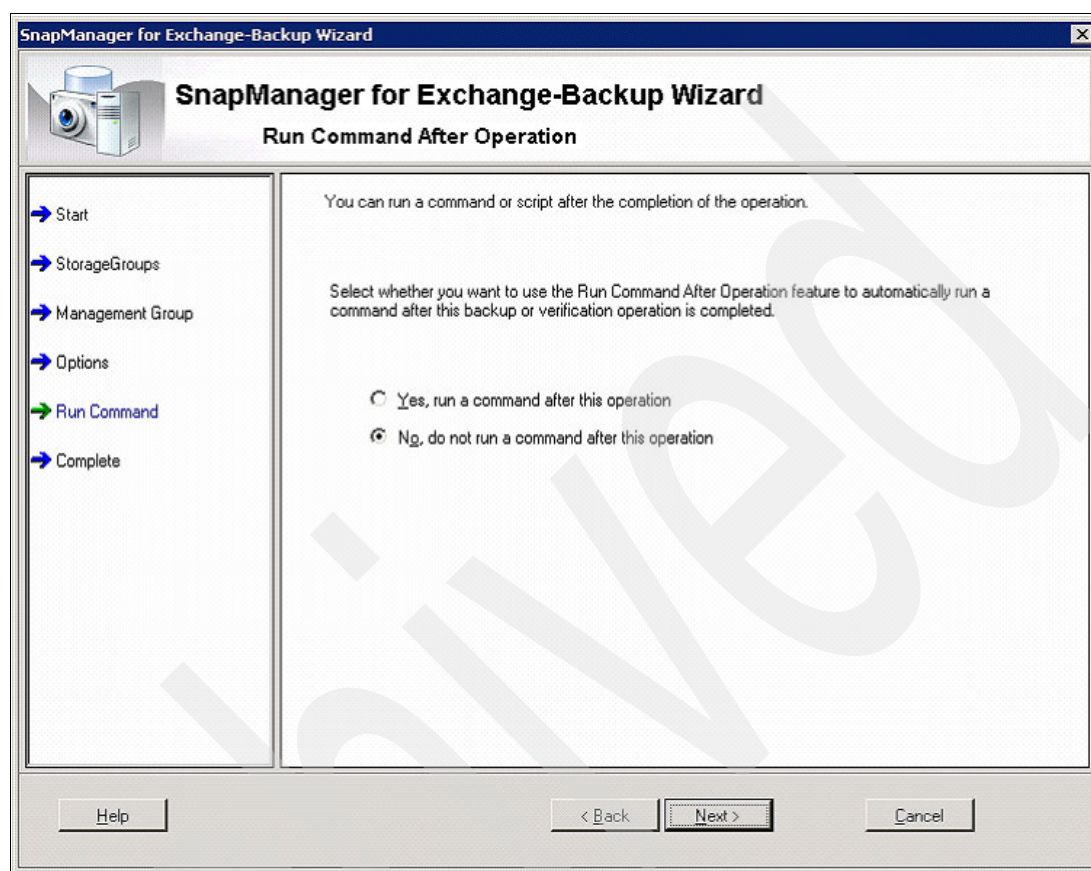
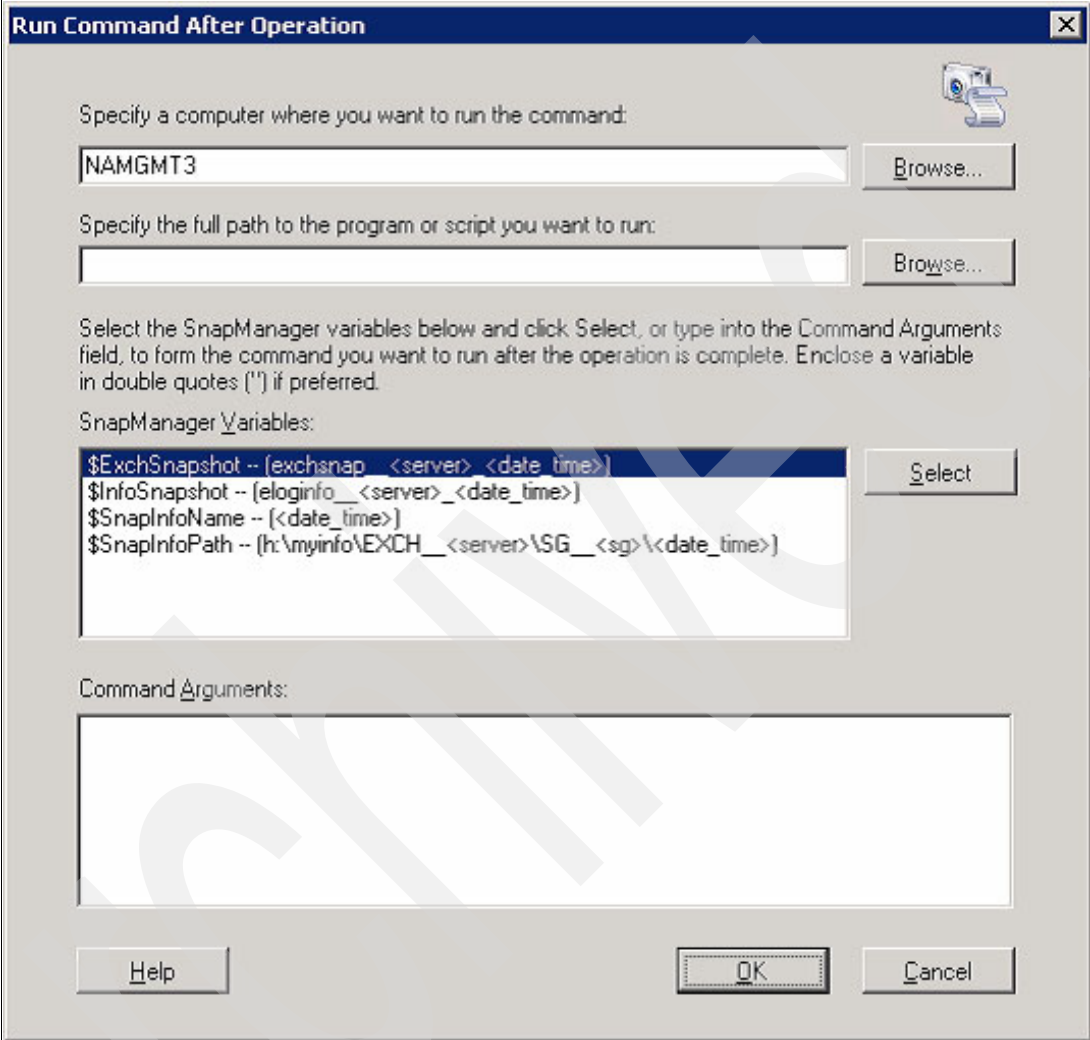


Figure 9-12 Run Command After Operation window

13. If you configured the Backup Wizard to run a command after the backup job is finished when you click **Next**, the Run Command After Operation detail window (Figure 9-13) will be shown. Type in the full path for the application to be run and any other argument to it. Click **OK** and then click **Next**.



The dialog box is titled "Run Command After Operation" and contains the following sections:

- Specify a computer where you want to run the command:** A text box containing "NAMGMT3" and a "Browse..." button.
- Specify the full path to the program or script you want to run:** An empty text box and a "Browse..." button.
- Select the SnapManager variables below and click Select, or type into the Command Arguments field, to form the command you want to run after the operation is complete. Enclose a variable in double quotes ("") if preferred.**
- SnapManager Variables:** A list box containing the following variables:
 - \$ExchSnapshot -- (exchsnap <server> <date_time>)
 - \$InfoSnapshot -- (eloinfo__<server>_<date_time>)
 - \$SnapInfoName -- (<date_time>)
 - \$SnapInfoPath -- (h:\myinfo\EXCH__<server>\SG__<sg>\<date_time>)A "Select" button is located to the right of the list box.
- Command Arguments:** A large empty text box for entering command arguments.
- Buttons:** "Help", "OK", and "Cancel" buttons are located at the bottom of the dialog.

Figure 9-13 Run Command After Operation details window

14. The Completing Backup Wizard window will be shown (Figure 9-14). If this backup is a one time only job, click **Finish**. If you are configuring a recurring backup job (such as a Daily backup job), click **Schedule**.

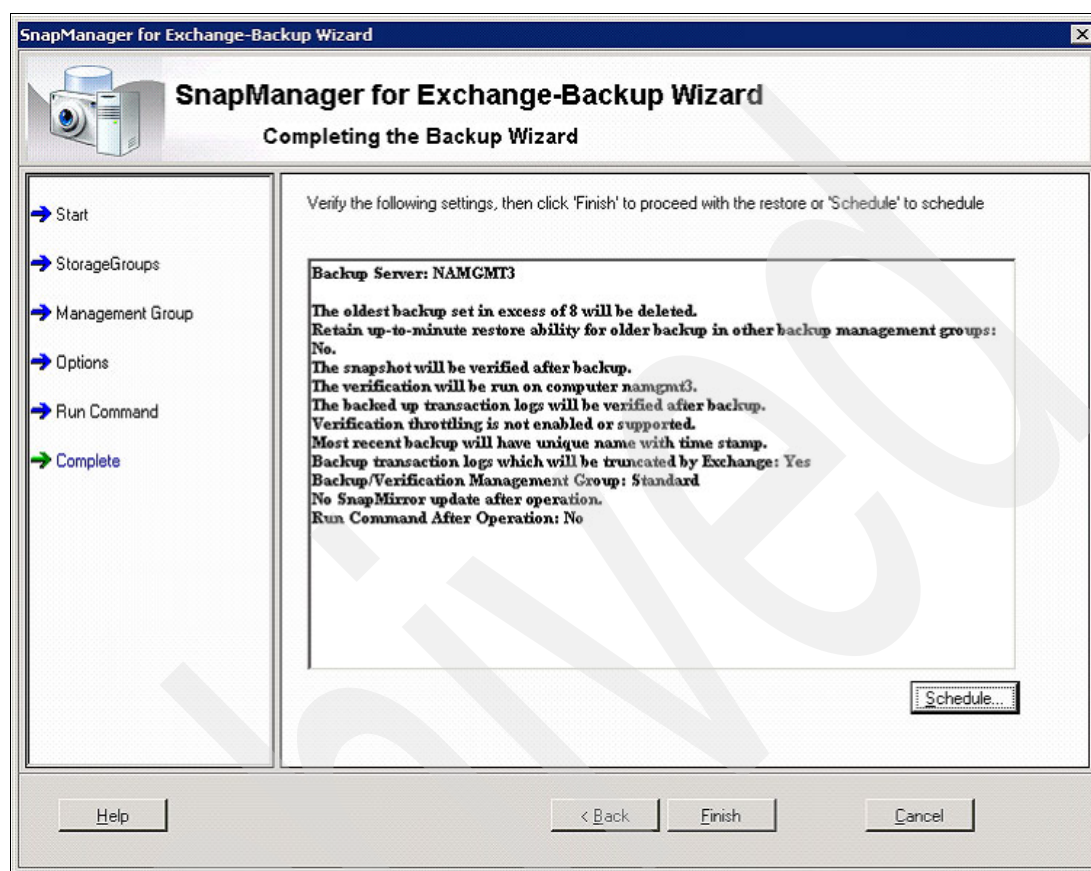


Figure 9-14 Completing Backup Wizard window

15. When you click **Finish**, the Backup Status window will appear, as shown in Figure 9-15, so that you can start the process. Click **Start Now** and the backup wizard will start the backup. The report is generated online and can be viewed by clicking the **Backup Report** tab.

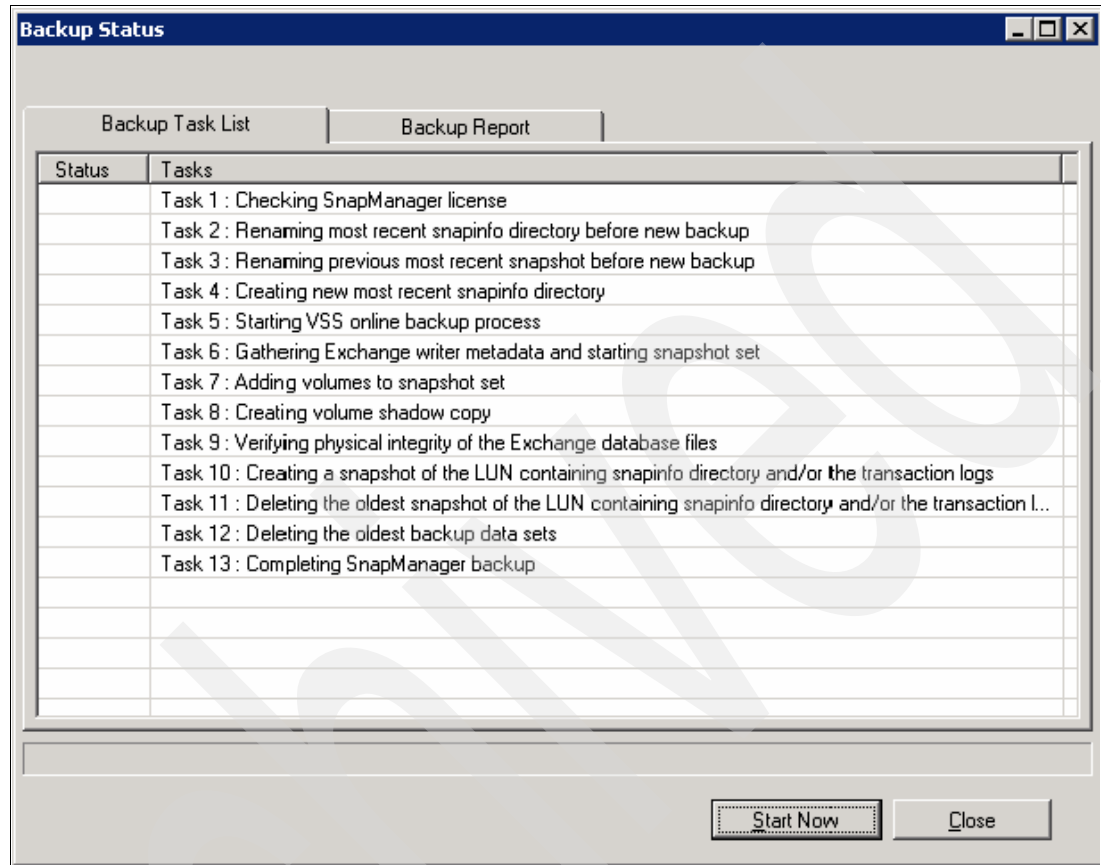


Figure 9-15 Backup Status window

16. If you click the **Schedule** button in the Completing Backup Wizard window (Figure 9-14 on page 329), the Backup Status window will not be shown. Instead, the Backup Wizard Schedule window will be shown (Figure 9-16). Type in a name for the Backup Job and the credentials to run the job. Click **OK**.

Schedule...

Schedule Job Name:

☐ Replace job if it exists

new-backup -Server 'NAMGMT3' -GenericNaming -ManagementGroup 'Standard' -NoTruncateLogs \$False -NoUTMRstore -RetainBackups 8 -StorageGroup 'First Storage Group' -Verify -VerificationServer 'namgmt3' -UseMountPoint -MountPointDir 'C:\Program Files\Exchsrvr\MDBDATA\Recovery\Mapping Point' -BackupCopyRemoteCCRNNode \$False

Run As:

Password:

Confirm password:

Figure 9-16 Backup Wizard Schedule window

17. The backup job will be scheduled using the Windows Task Scheduler. The Task Scheduler window will open (as seen on Figure 9-17) already completed with some information. Click the **Schedule** tab.

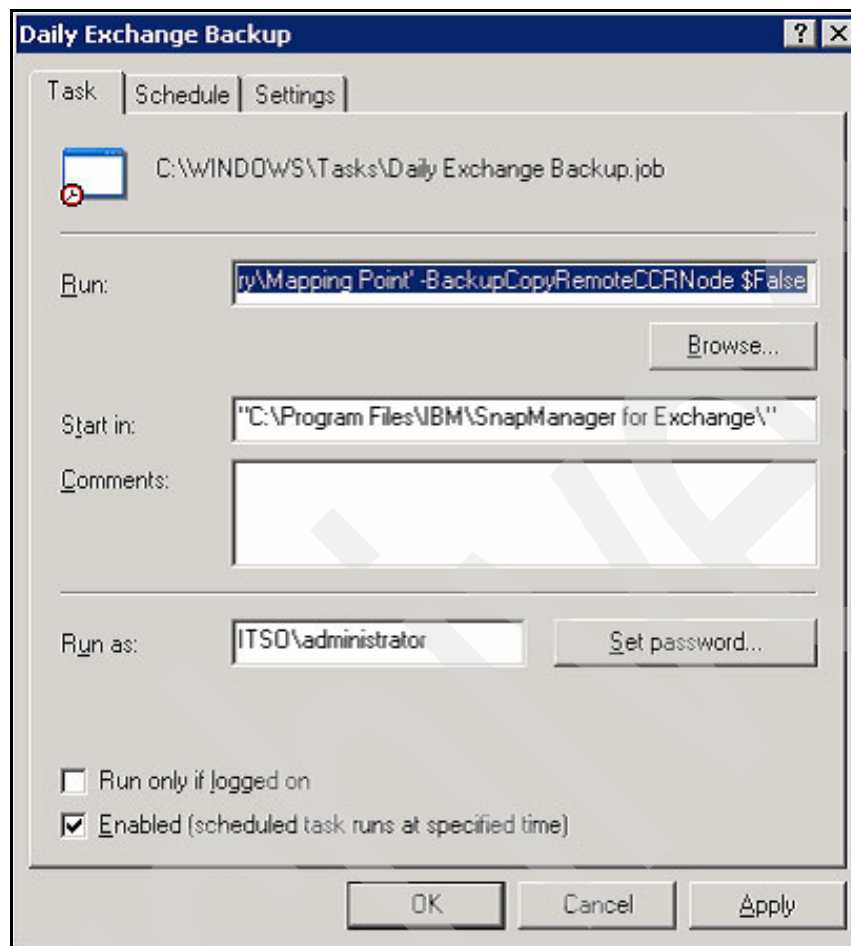


Figure 9-17 Task Scheduler window

18. In the Task Scheduler window (Figure 9-18 on page 333), select the schedule type (weekly, monthly, daily, and so on) and select the Start date for the task. After reviewing all the configurations, click **OK**.

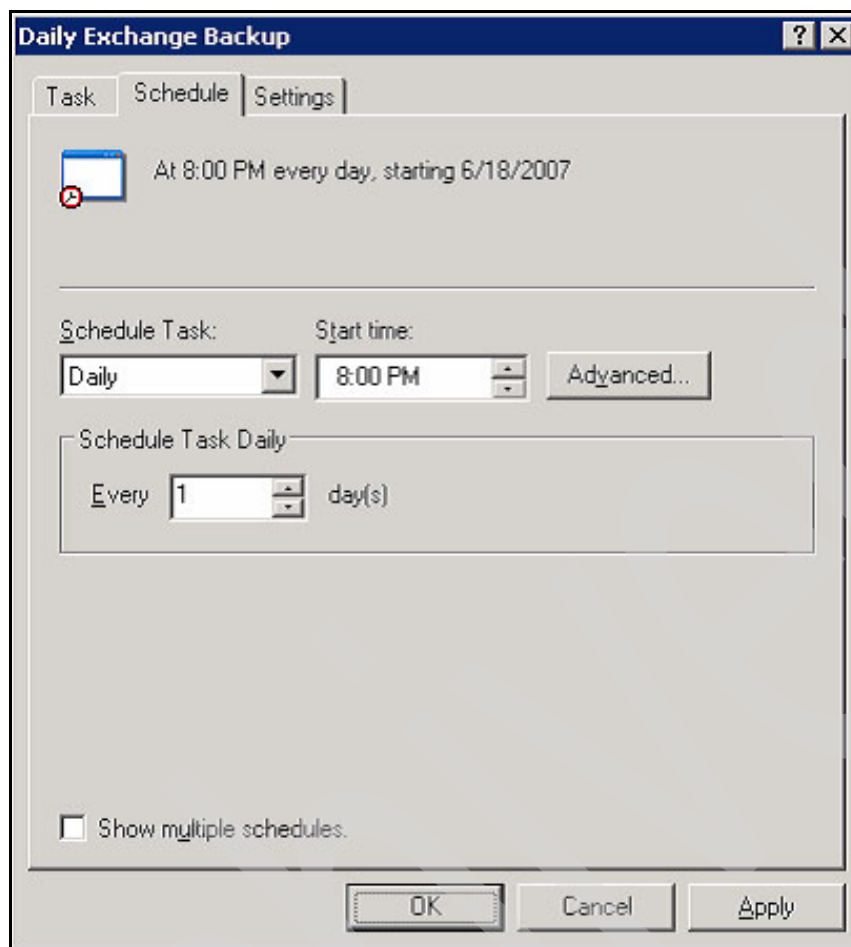


Figure 9-18 Schedule Task window

19. The task will be created in the Windows Task Scheduler and scheduled for the configured start time and repeat interval.

Note: SnapManager backups reside on the IBM System Storage N series storage system. For external backups on media tapes or site disaster protection, use additional backup software.

Backup using SnapDrive

Those environments that cannot take advantage of the manageability provided by the SnapManager for Microsoft Exchange still need to back up their Microsoft Exchange servers. In order to execute this backup, the SnapDrive software can be used as an additional layer of protection, generating Snapshots on the IBM System Storage N series storage system and restoring them later, if needed.

Note: Dismounting the databases should be done before backup

SnapDrive does not have the Advanced Programming Interface (API) to integrate with the Microsoft Exchange server and, for this reason, does not commit the Transaction Log files on the database. SnapDrive Snapshots are not a complete solution for backing up the Microsoft Exchange server environment. It is rather an additional protection in case a restore is needed.

When using the SnapDrive software to generate Snapshots, use additional backup software, such as IBM Tivoli Storage Manager (ITSM) that integrates with the Microsoft Exchange server and backs up the environment.

When creating Snapshots, always create a snapshot for the database LUNs and for the transactional log files LUNs so that the Mailbox Store can be restored using the log files.

Note: Snapshots generated by SnapDrive are not integrated to or managed by SnapManager. To delete them from the IBM System Storage N series storage system, use the FilerView, Data ONTAP command-line interface (CLI), or the SnapDrive MMC.

To create a Snapshot using SnapDrive:

1. Open the Computer Management MMC.
2. Expand **Storage**.
3. Expand **SnapDrive**.
4. Expand **Disks**.
5. Expand the disk to have a Snapshot created.
6. Right-click the Snapshots folder beneath the disk and click **Create Snapshot...**, as shown in Figure 9-19).

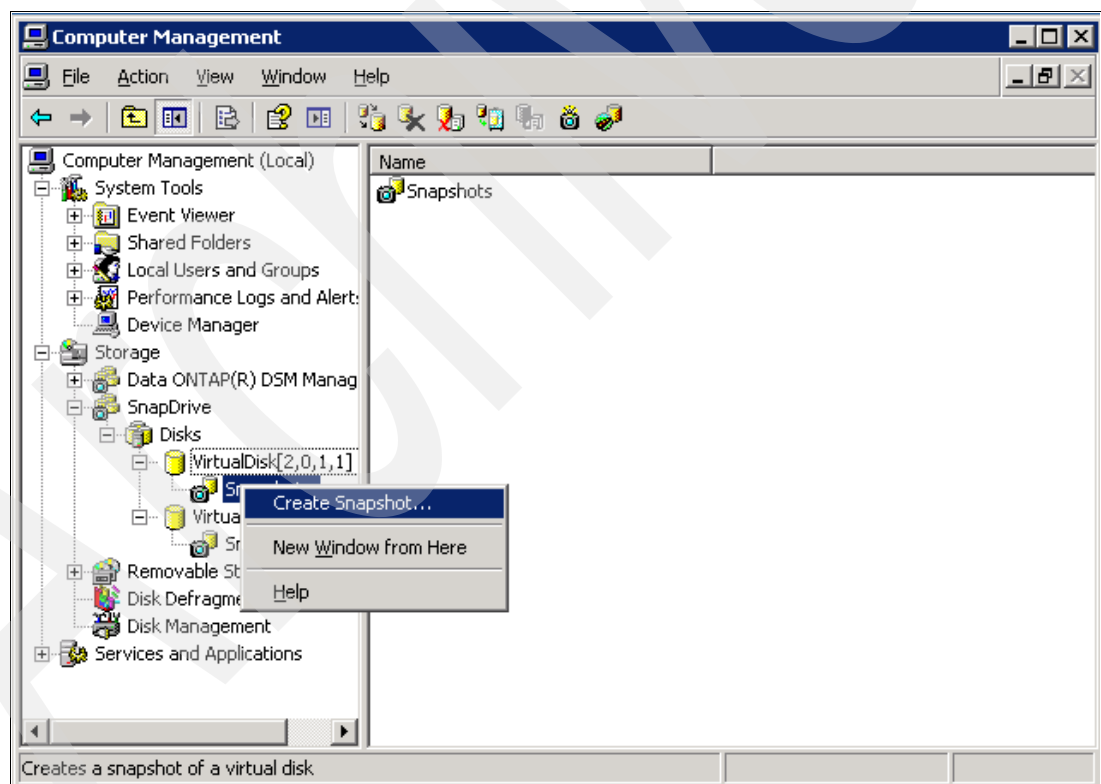


Figure 9-19 Create Snapshot window

7. Type in a name for the Snapshot in the Snapshot name window (Figure 9-20 on page 335). Notice that the name is the reference to the Snapshot. Remember to include any needed information about the Snapshot name (such as dates or type of Snapshot). Click **OK**.

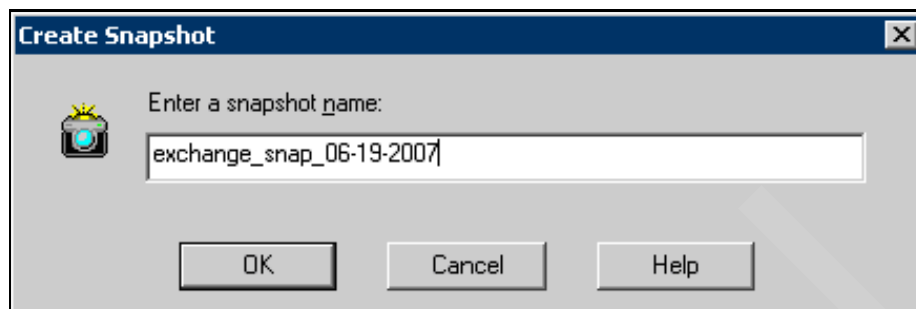


Figure 9-20 Snapshot name window

8. Snapshot will be created and shown in the Snapshots view on the SnapDrive MMC (Figure 9-21).

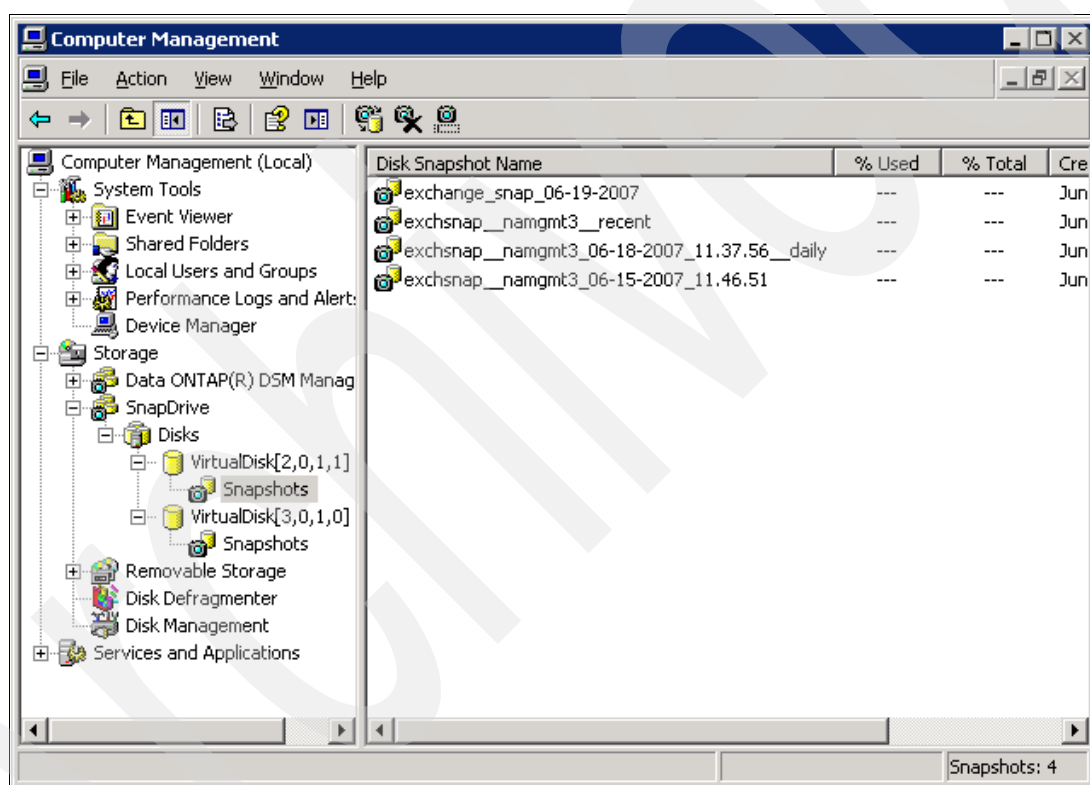


Figure 9-21 Snapshots view window

9.1.2 Restore

Depending on the way the Microsoft Exchange server was backed up, the restore operation can be more or less complex.

Restore from SnapManager based backups

If the backup solution for your Microsoft Exchange infrastructure is based on the SnapManager for Microsoft Exchange, the restore of services can be accomplished in a matter of minutes.

Because the backups from SnapManager are based on the Snapshot technology, the restores are much more fast and reliable.

Prior to starting the restore operation, there are some configurations that can be changed on the SnapManager for Microsoft Exchange for restore operations:

1. On the SnapManager for Exchange MMC, right-click the server name and select **Backup Verification Settings**.
2. The Database Verification Settings window will be shown (Figure 9-22). Click the **Access LUN in Snapshot** tab. This configuration controls how the Snapshot will be mounted on the Microsoft Exchange server in order to accomplish the restore operation. The default configuration is to mount the Snapshot in an empty NTFS directory. If necessary, change the configuration so that a drive letter will be used instead of a Volume Mount Point.

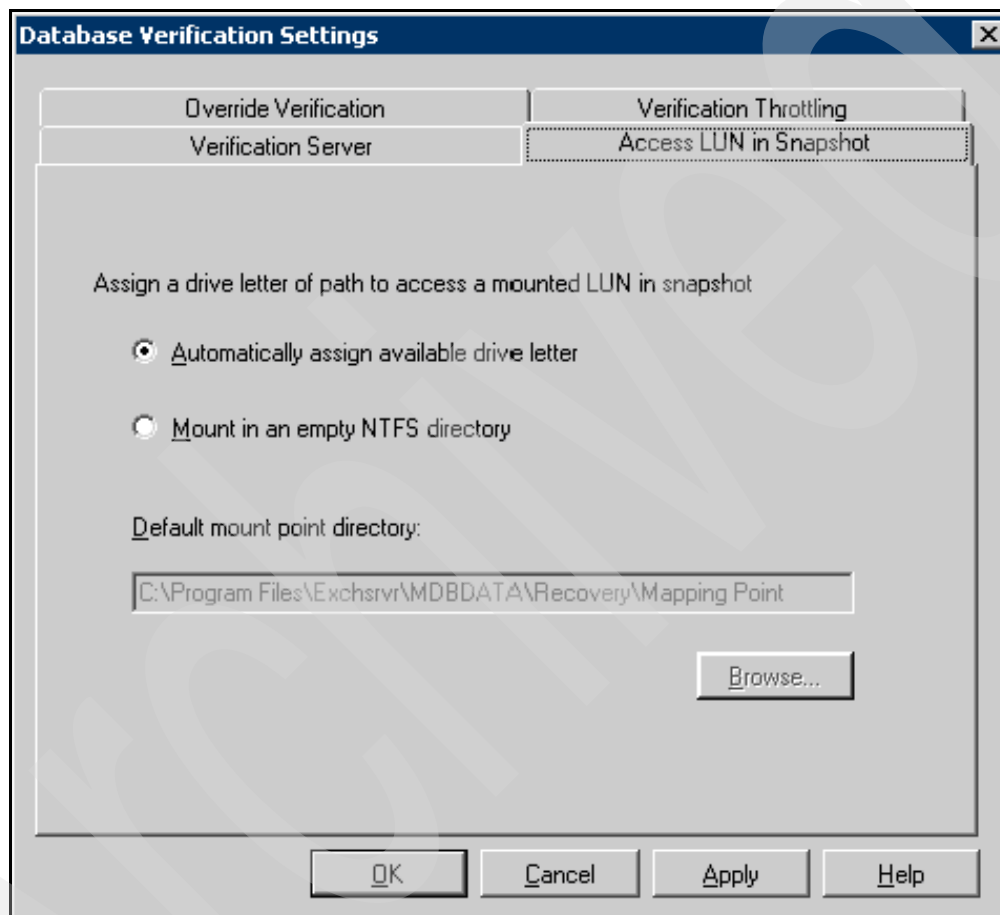


Figure 9-22 Database Verification Settings window

3. Click the **Override Verification** window (Figure 9-23). The check box **Override database verification requirement for restore** controls if the restore operation can be executed on a non-verified backup set or if the restore operation will only be allowed to run using a verified backup set. The recommended configuration is to leave the default configuration to only run restore on verified backup sets (check box unmarked).

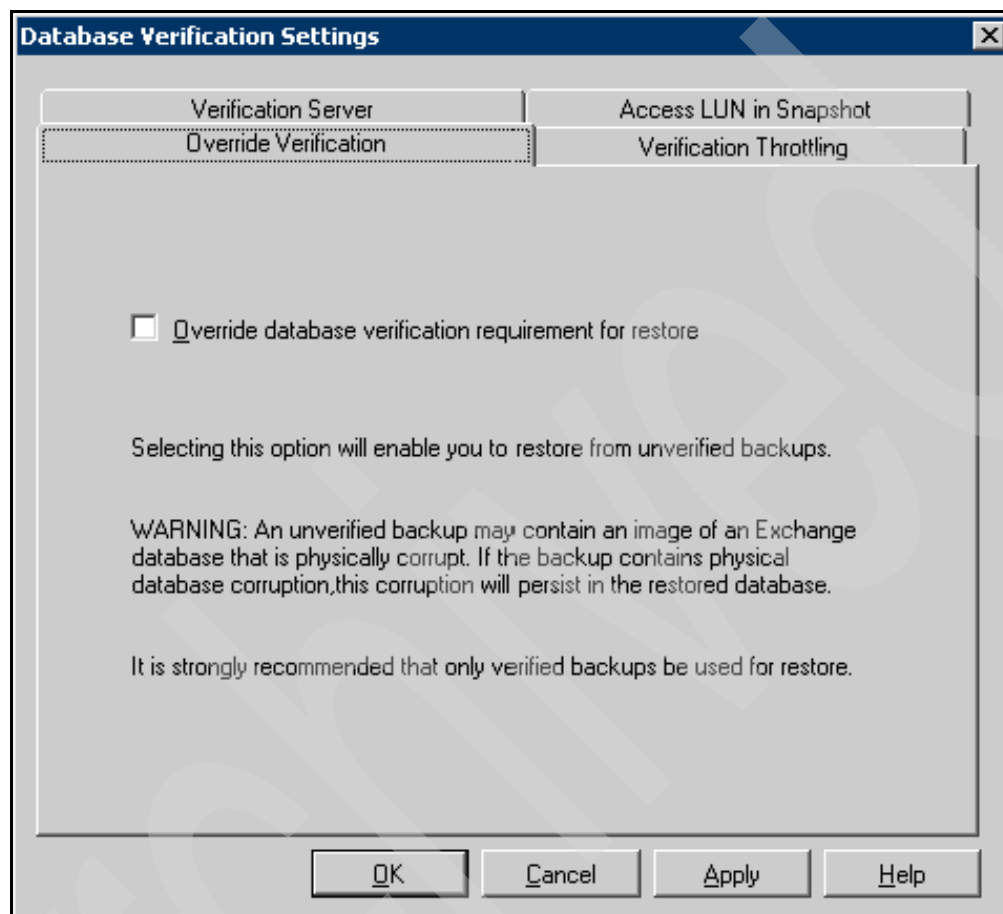


Figure 9-23 Override Verification window

4. Click the **Verification Throttling** window (Figure 9-24). The database verification process can impact the performance of either the Microsoft Exchange server or the IBM System Storage N series storage system because of its intense I/O use. If tuning is needed for the Verification Throttling, mark the check box and define how many I/O operations should be done between 1 second pauses.

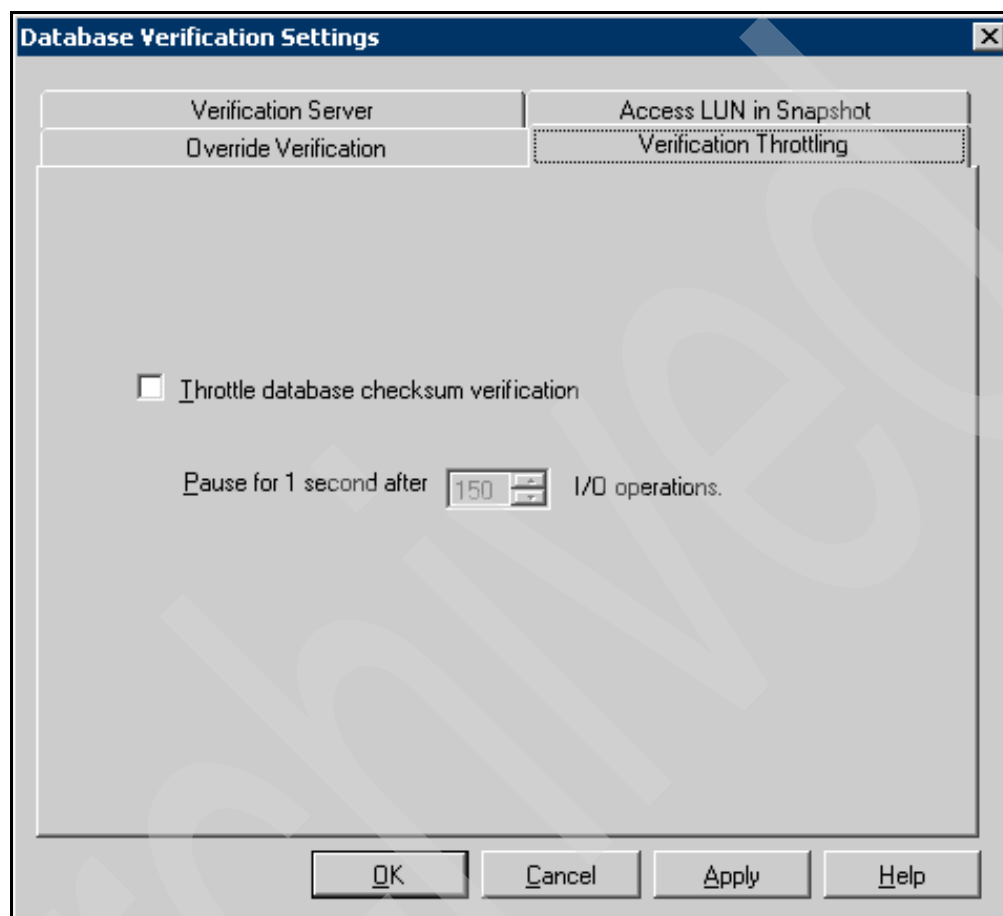


Figure 9-24 Verification Throttling window

5. Click the **Verification Server** window (Figure 9-25 on page 339). This configuration tab allows you to change the verification server to be used as the default verification server for Backup and Verify operations.

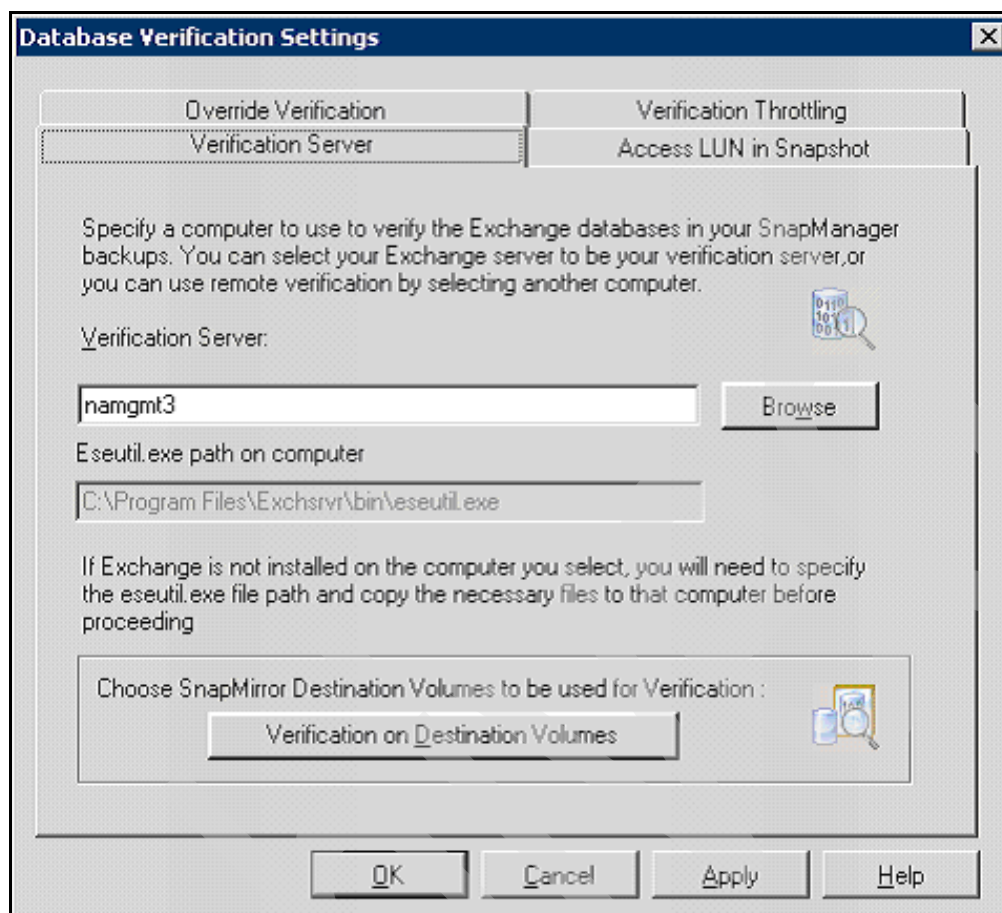


Figure 9-25 Verification Server window

After making configuration changes in the Database Verification Settings window, click **OK**.

These are the steps to restore the Exchange services on a Microsoft Exchange server that was backed up using the SnapManager for Microsoft Exchange application:

1. On the Microsoft Exchange server, select **Start** → **All Programs** → **IBM** → **SnapManager for Exchange Management Console**. Expand the server name and click **Restore**. Expand the server name and the available Storage Groups. The list of backup sets available will be shown (Figure 9-26).

Note: The way the backup sets are shown on SnapManager for Microsoft Exchange depends on the Microsoft Exchange infrastructure on the N series storage. If you have a LUN for each Mailbox Store, then the backup can be done on that specific Mailbox Store and so can the restore.

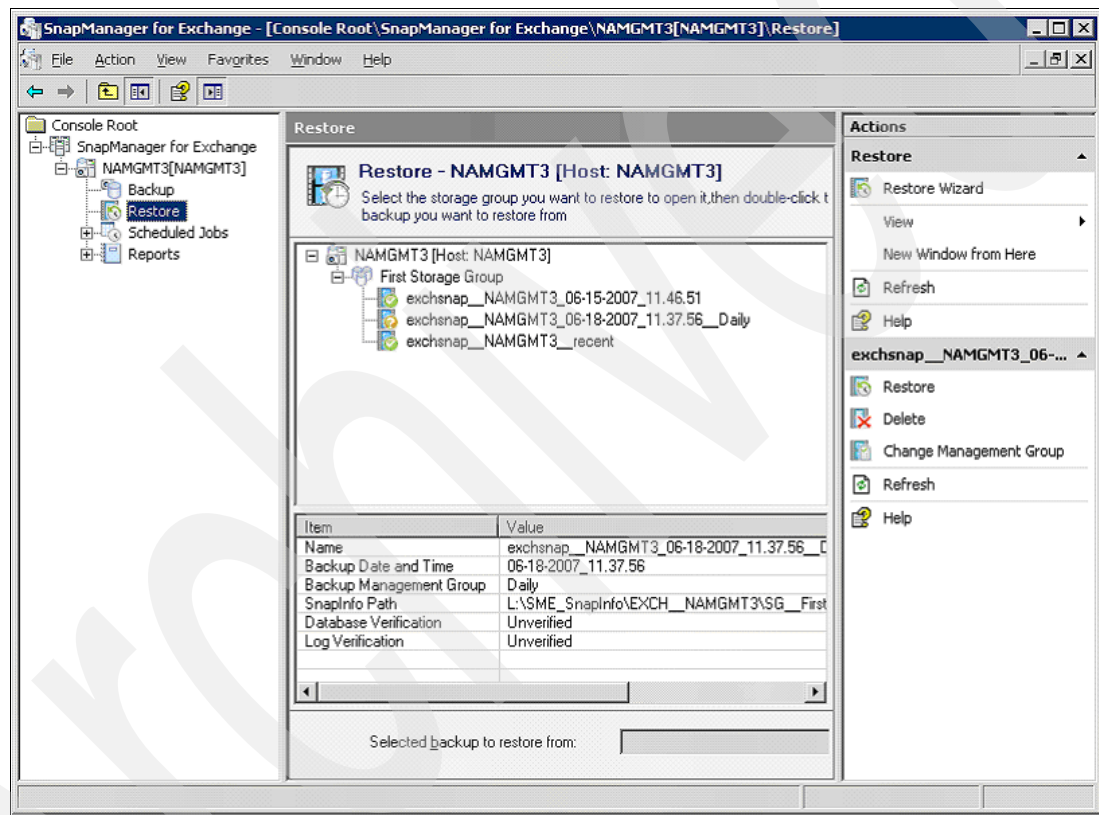


Figure 9-26 Backup sets available

2. Notice that some backup sets have the green check mark while others have the orange quest mark. This shows whether the backup sets are verified or not. By default, you can only restore from verified backup sets. If for any reason the backup had not been verified at the time of creation, the verification process can be started in order to verify it.

- a. To start the verification process, right-click the server name and start the Backup Wizard. Follow the steps described earlier in this chapter for the backup operation, but in the Backup or Verification window (Figure 9-27), select the verification option and choose the number of recent backups to verify. Click **Next**.

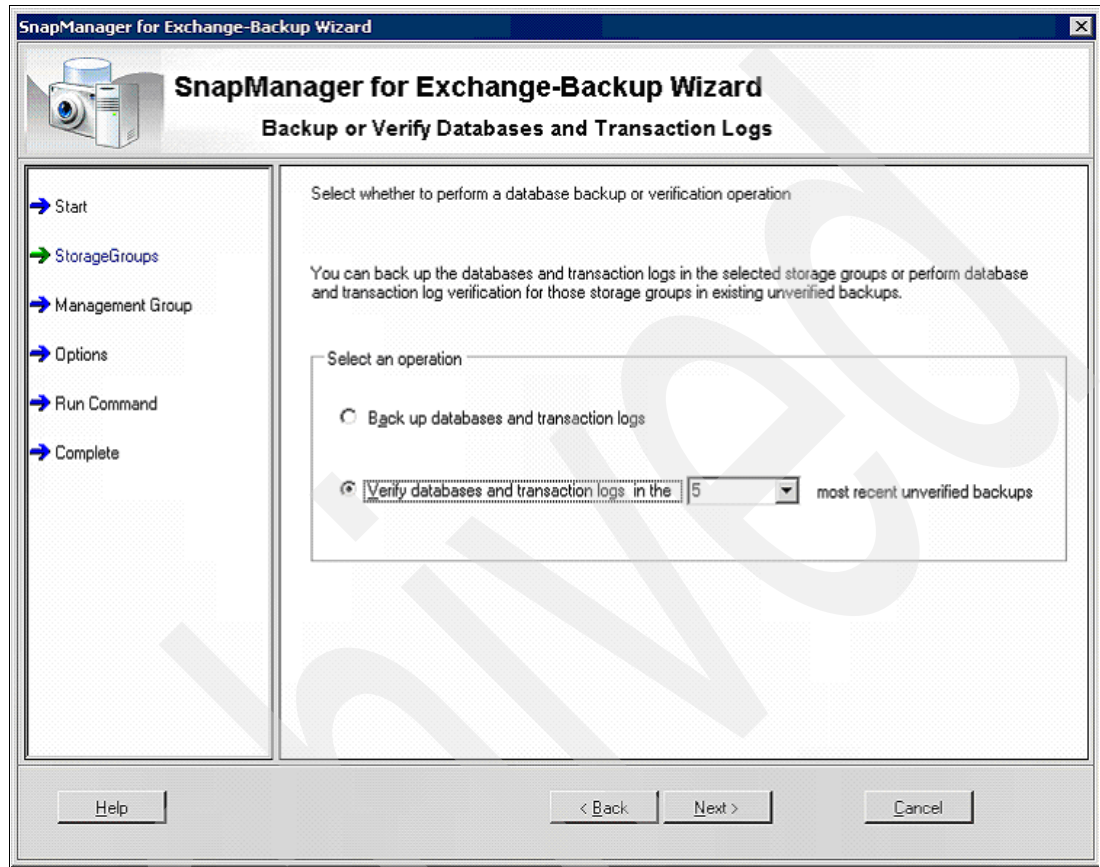


Figure 9-27 Backup or Verification window

- b. Select the Management Groups that should have the backup sets verified. You can select each Management Group independently or you can select **All** to have the number of backup sets defined on all the Management Groups verified (Figure 9-28). Click **Next**.

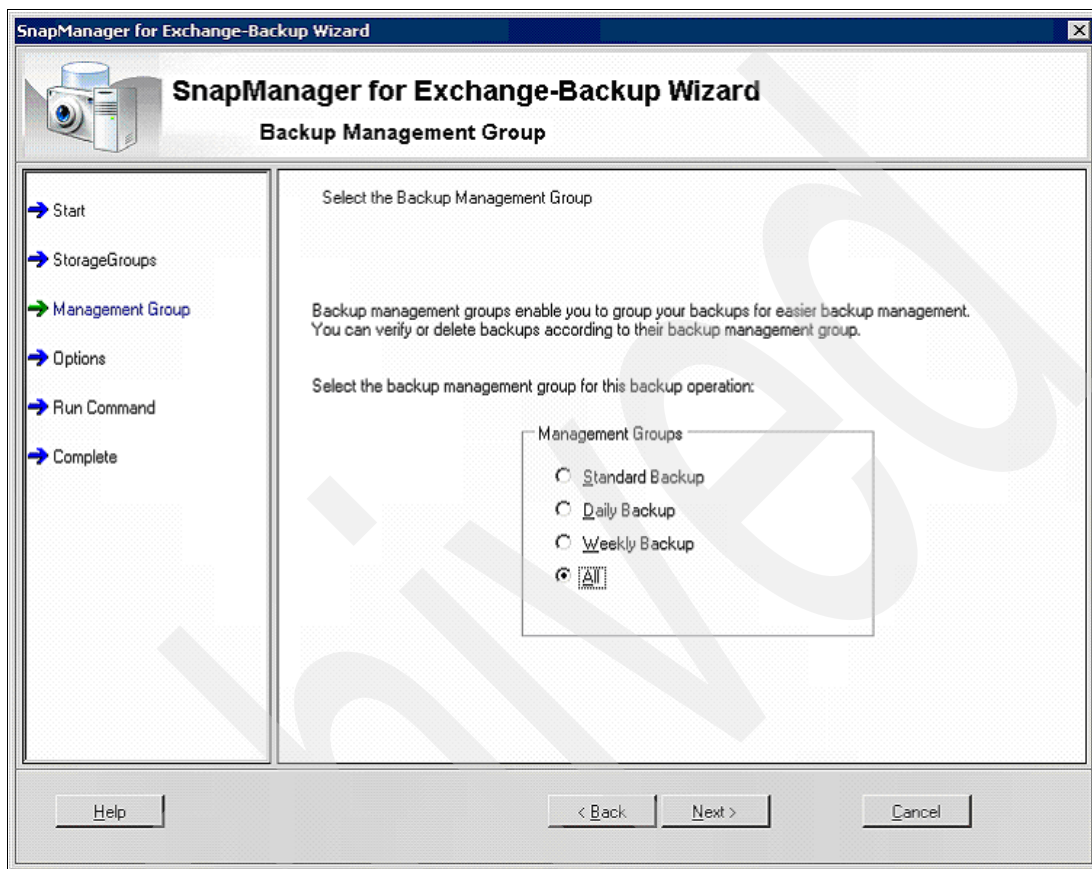


Figure 9-28 Management Group window

- c. In the Verification Server window (Figure 9-29), check the server that will be used for verification and click **Next**.

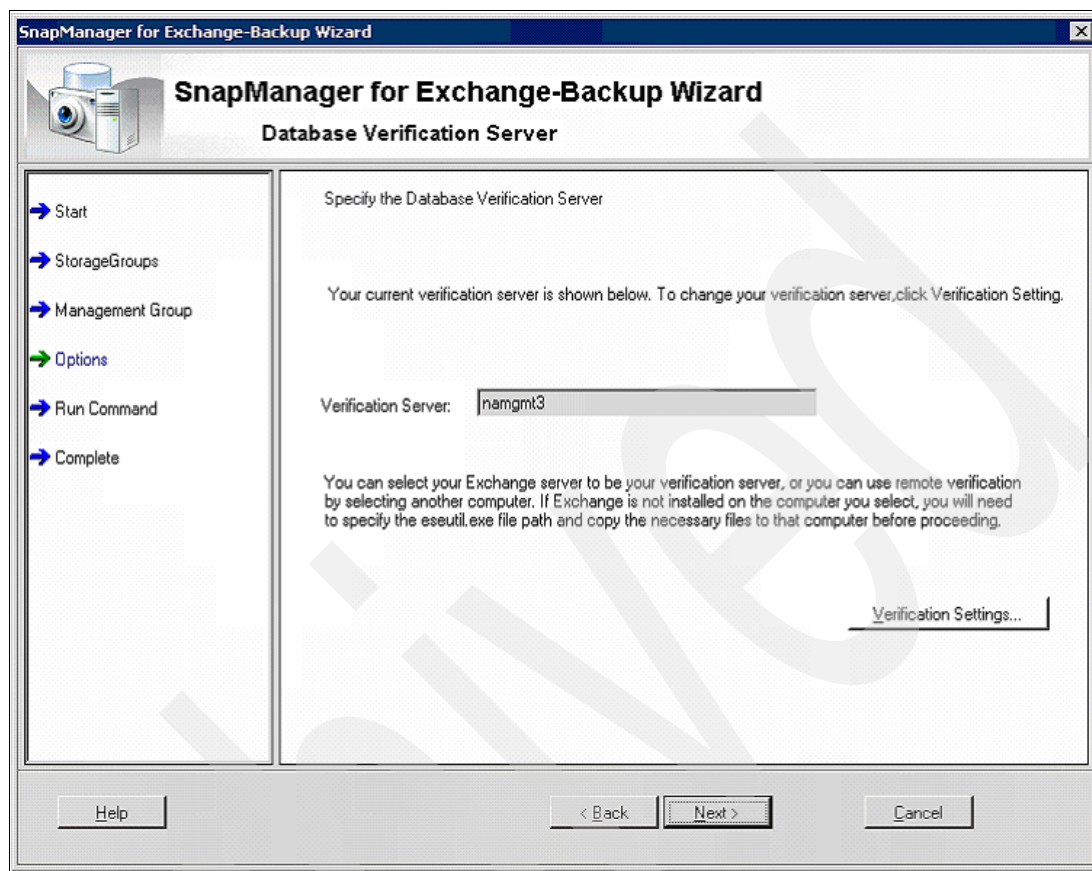


Figure 9-29 Verification Server window

- d. Select whether to run a command after the operation or not (Figure 9-30) and click **Next**.

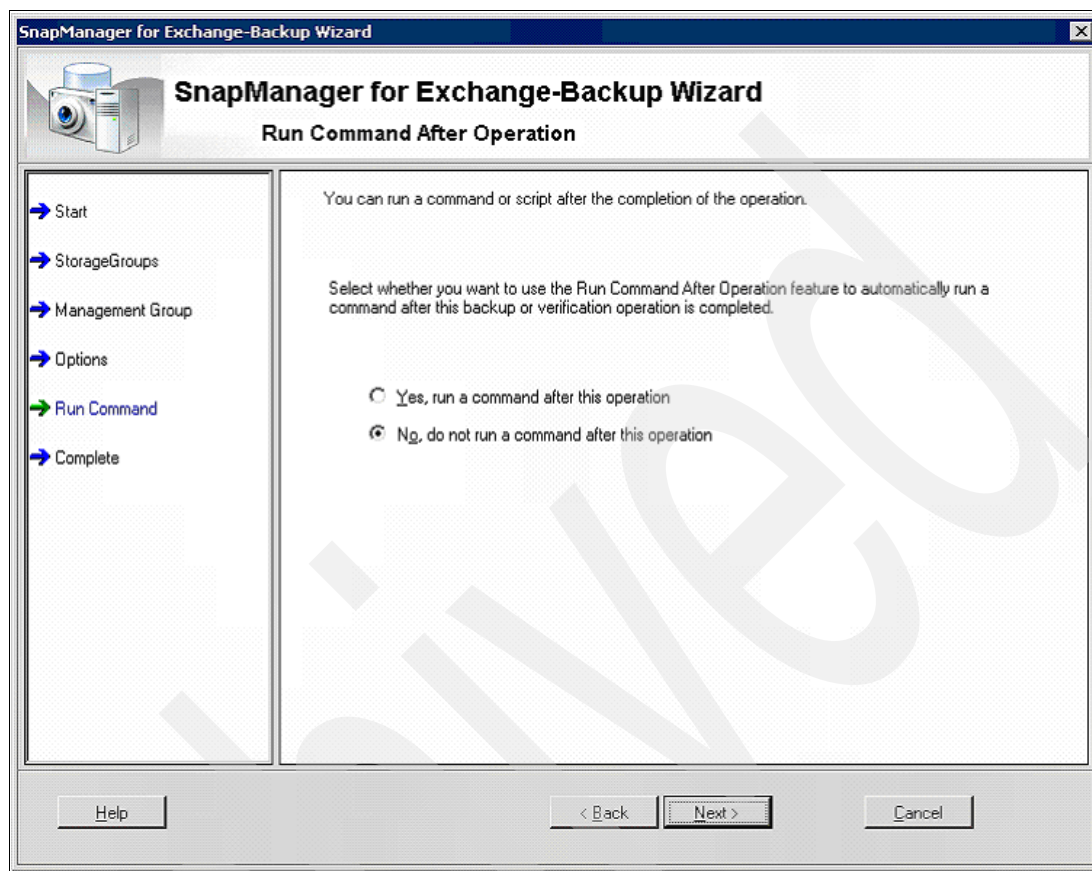


Figure 9-30 Run Command After Operation window

- e. In the Completing Backup Wizard window (Figure 9-31 on page 345), verify the summary. If any change is necessary, click **Back**. If the summary is correct, click **Finish** to run it now or **Schedule** to create a schedule for this verification process.

Note: The verification cannot be done on the volume that has a Snapshot operation in progress. If you are scheduling the verification to occur at a specific time and date, check if no Snapshot operations will be in use on that time and date.

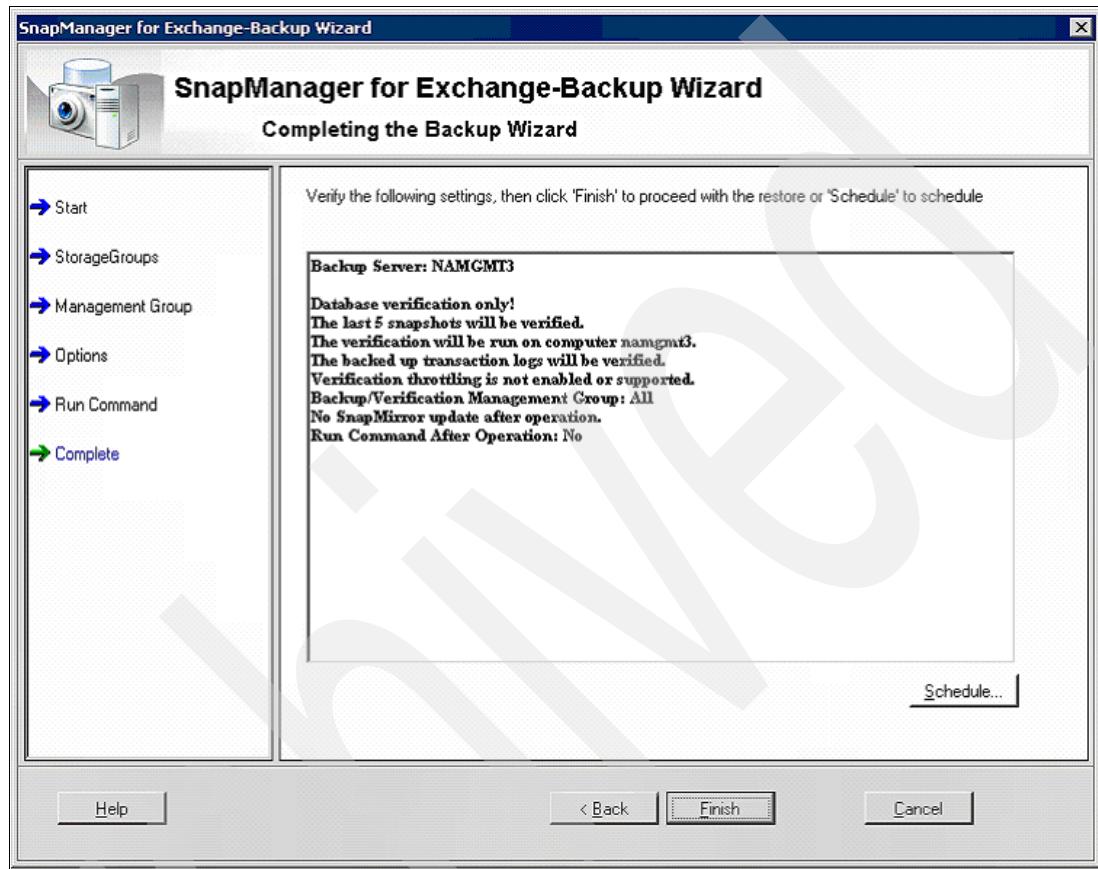


Figure 9-31 Completing Backup Wizard window

-
- The screenshot shows a window titled "Backup Status" with two tabs: "Backup Task List" and "Backup Report". The "Backup Task List" tab is active, displaying a table with two columns: "Status" and "Tasks". The table contains five rows of tasks, all of which are currently empty. At the bottom of the window, there are two buttons: "Start Now" and "Close".
- | Status | Tasks |
|--------|-------|
| | |
| | |
| | |
| | |
| | |
- Start Now Close

3. After the verification process is done, the backup sets list can be refreshed to show that all the backup sets now are verified (Figure 9-33).

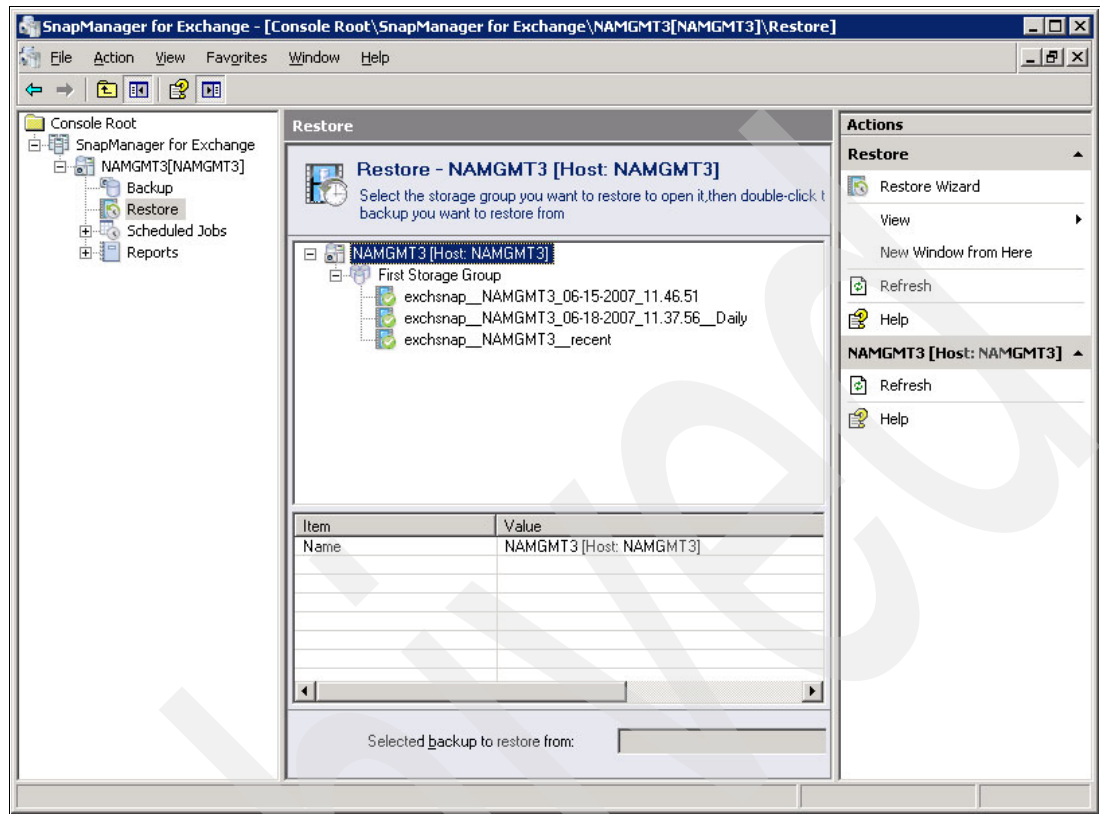


Figure 9-33 Backup sets available after verification

4. To start the restore process, right-click the server name and select **Restore Wizard**. In the Restore Wizard welcome window (Figure 9-34), click **Next**.



Figure 9-34 Restore Wizard Welcome window

5. In the Exchange Server selection window (Figure 9-35), select whether this restore is from a SnapManager backup set created by the same Exchange server or if it is from a SnapManager backup set created on a different Exchange server and Archived. The most common scenario is to restore a backup executed on the Microsoft Exchange server. Click **Next**.




Figure 9-35 Exchange server selection window

- a. If you are restoring from an archive, the Restore from Archive window (Figure 9-36) will be shown. Select the server that executed the backup and the path where the SnapInfo Directory Snapshot is mounted. Prior to starting the restore operation, you need to connect the SnapInfo Directory Snapshot to the path. When you click **Next**, a list of backups executed by the specified server will be shown and you will have to select which backup set is going to be restored.



Figure 9-36 Restore from Archive window

6. In the Choose backup window (Figure 9-37), expand the server name and expand the Storage Group that you plan to restore. Select the SnapManager backup set that you plan to restore. Notice that depending on the convention name selected at the time of the backup, the backup set names can be different. If you used the generic convention naming, the backup set with the *_recent* on the end of its name is the most up to date set. Select the backup set by double clicking it and click **Next**.

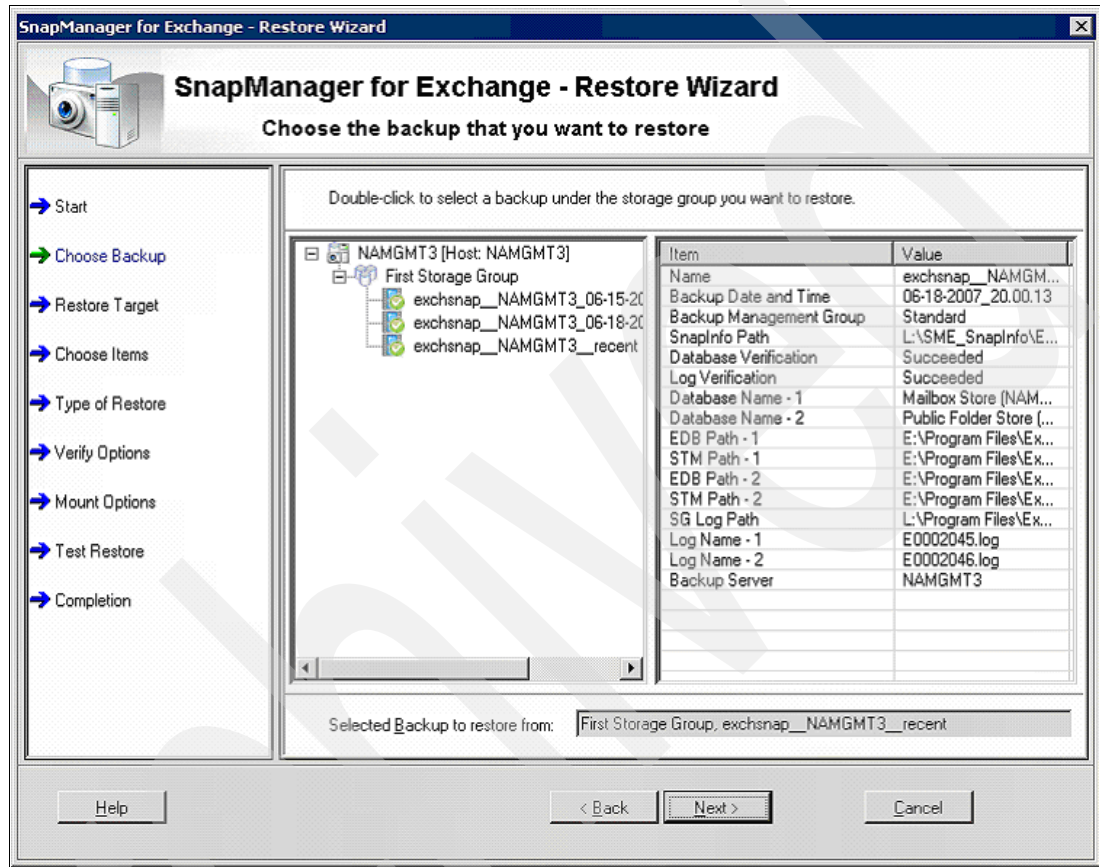


Figure 9-37 Choose backup window

7. In the Choose target window (Figure 9-38), select whether this backup set will be restored to the same Storage Group (which means it will overwrite your actual production environment) or if it will be restored to a Recovery Storage Group. The restore operation to a Recovery Storage Group is shown in “Restore to Recovery Storage Groups” on page 358. Click **Next**.

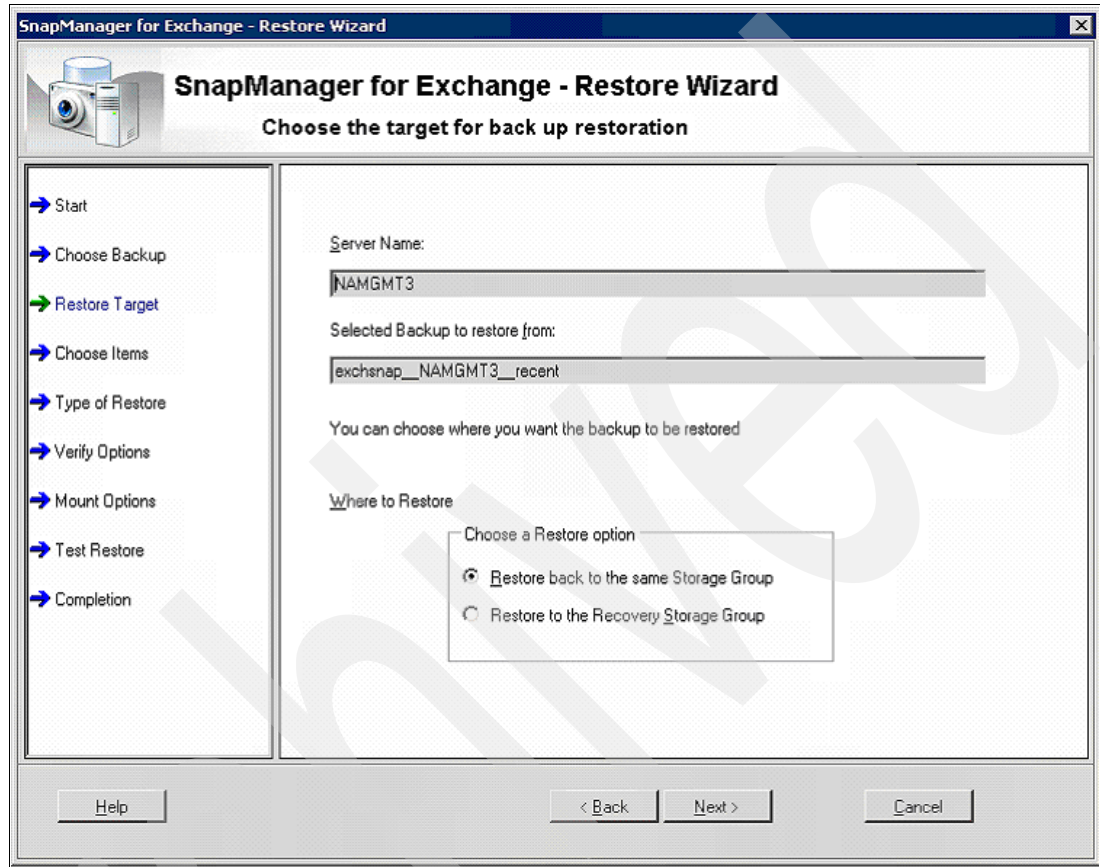


Figure 9-38 Choose target window

8. In the Choose items to be restored window (Figure 9-39), select which Mailbox Stores should be restored. Remember that all the Mailbox Stores stored on the same LUN have to be restored together. All Mailbox Store databases that resides on the same LUNs can not be selected separately. After the selection is done, click **Next**.

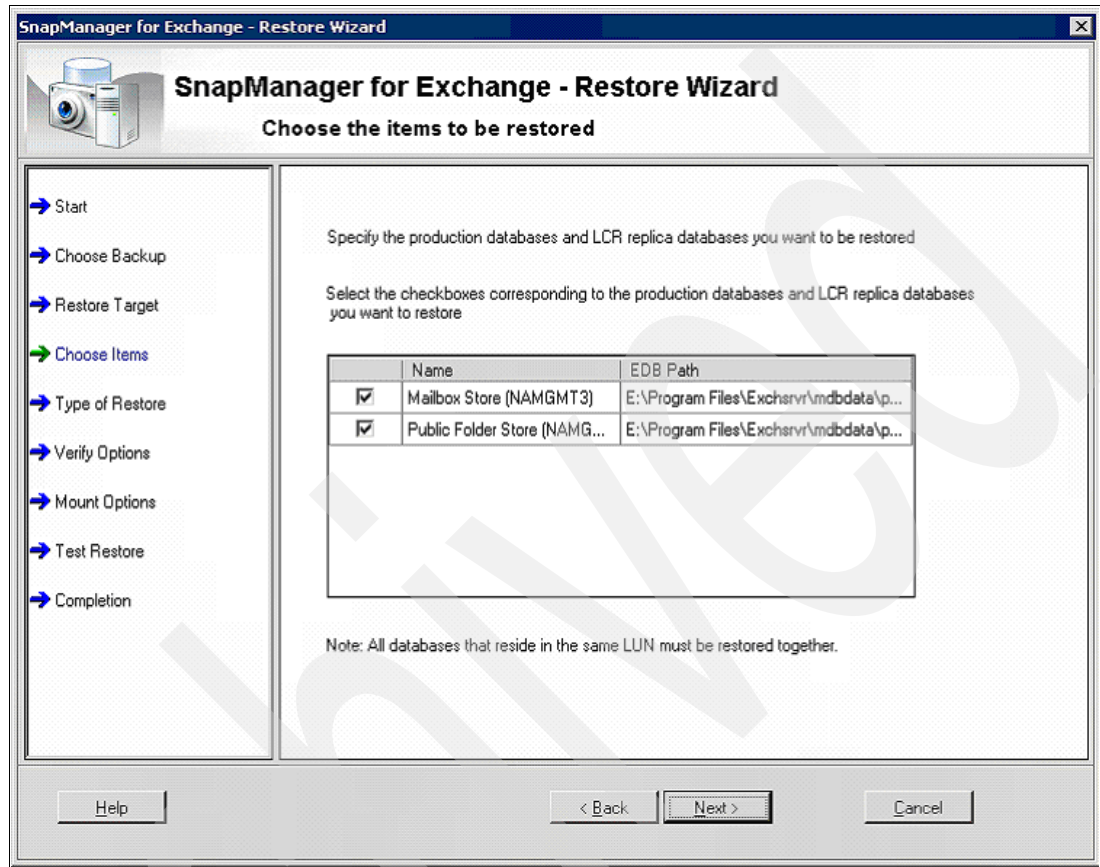


Figure 9-39 Choose items to be restored window

9. In the Type of Restore window (Figure 9-40), there are two available options: Point in time restore and Up to the minute restore. Point in time restore will restore only the Microsoft Exchange database on the SnapManager backup set, without applying any transaction log file to the database after the restore operation is done. That means that your Microsoft Exchange will be restored to the same point it was when that specific backup set was created. Up to the minute restore will restore the Microsoft Exchange database on the backup set and will apply all the transaction log files from the backup set creation date to the last transaction log file available. In order to do that, the SnapInfo Directory will be used to provide the older log files that are not on the Transaction Log files folder anymore. Select the desired type of restore and click **Next**.

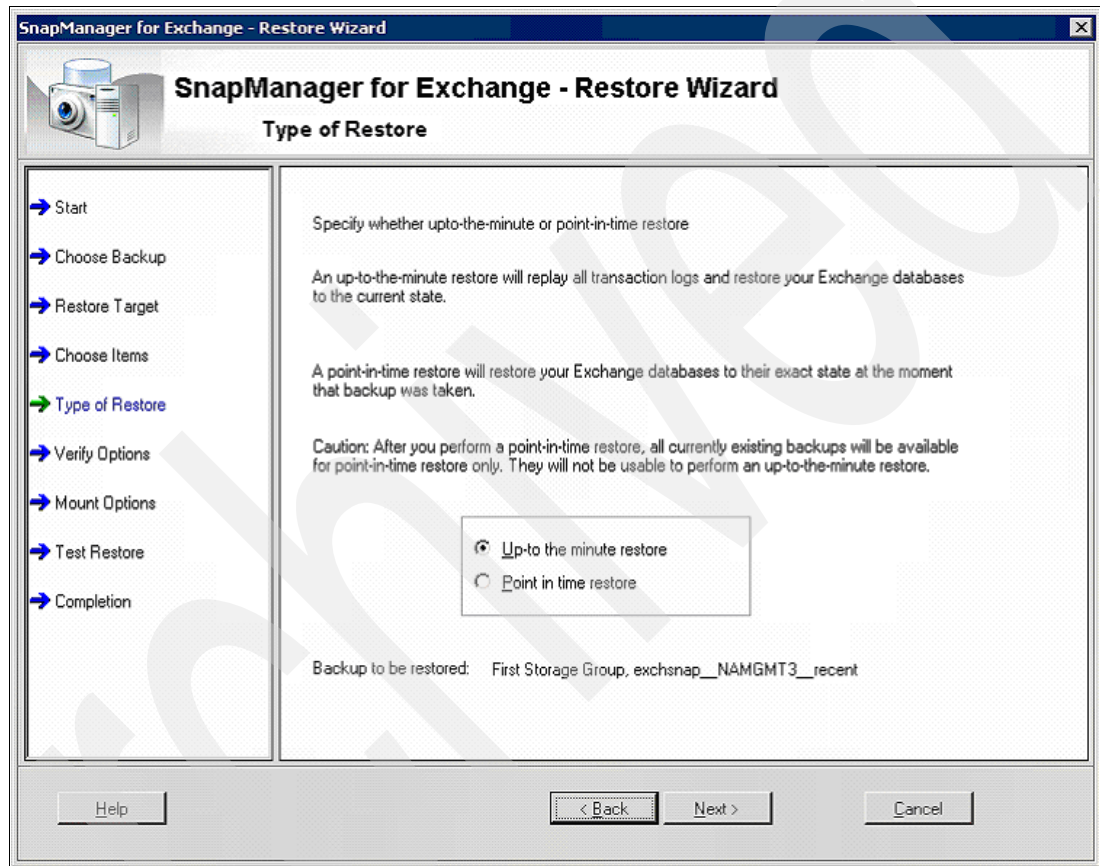


Figure 9-40 Type of Restore window

10. In the Test Restore window (Figure 9-41), specify whether this is an actual restore of the Microsoft Exchange server or if this is just a test restore to validate the process, files, Snapshots, databases, SnapInfo Directory, and so on. Click **Next**.

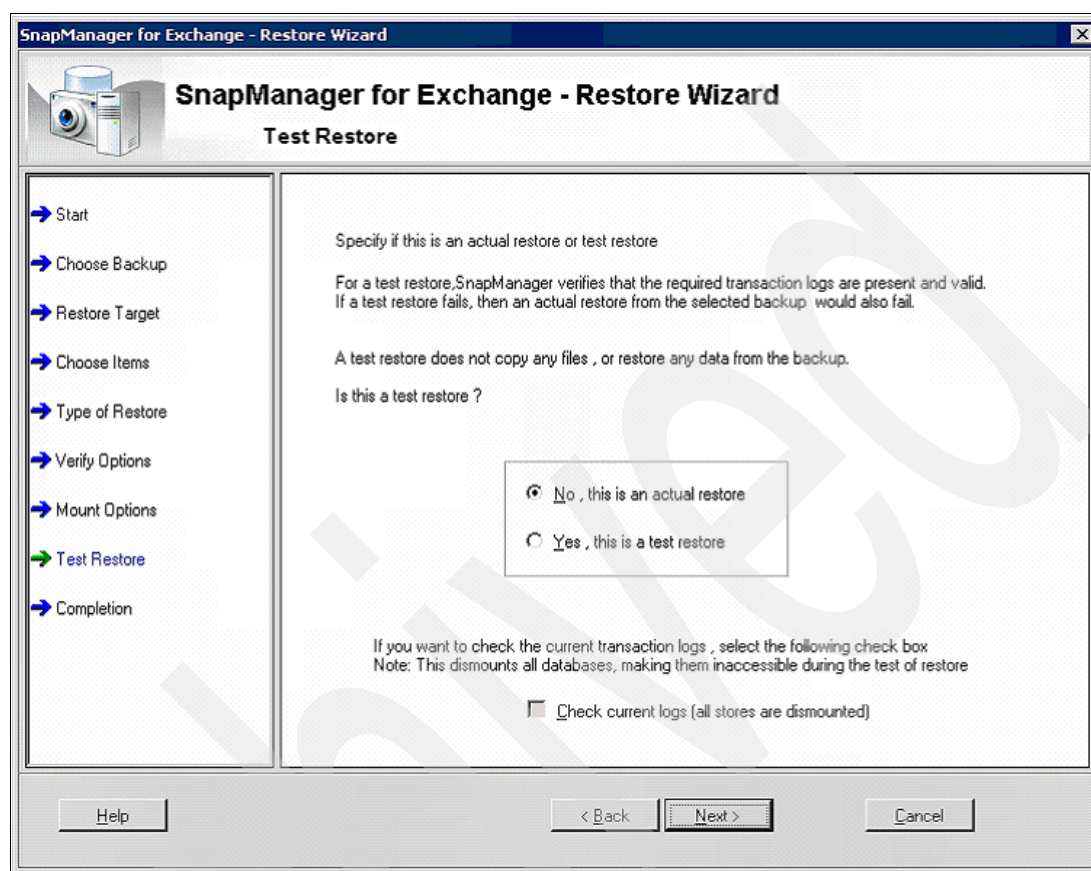


Figure 9-41 Test Restore window

11. In the Verify Options window (Figure 9-42), specify which validation tests should be done on the database and transaction log files before the actual restore operation starts. To verify only the database and log files sequence in order to make sure no log file is missing, uncheck the **Exhaustive Verification** check box. If you check the **Exhaustive Verification** check box, a complete test on each transaction log file will be done to guarantee that no problem will occur because of any single log file problem. This verification also makes the process more time and resource consuming. After selecting the desired verification, click **Next**.

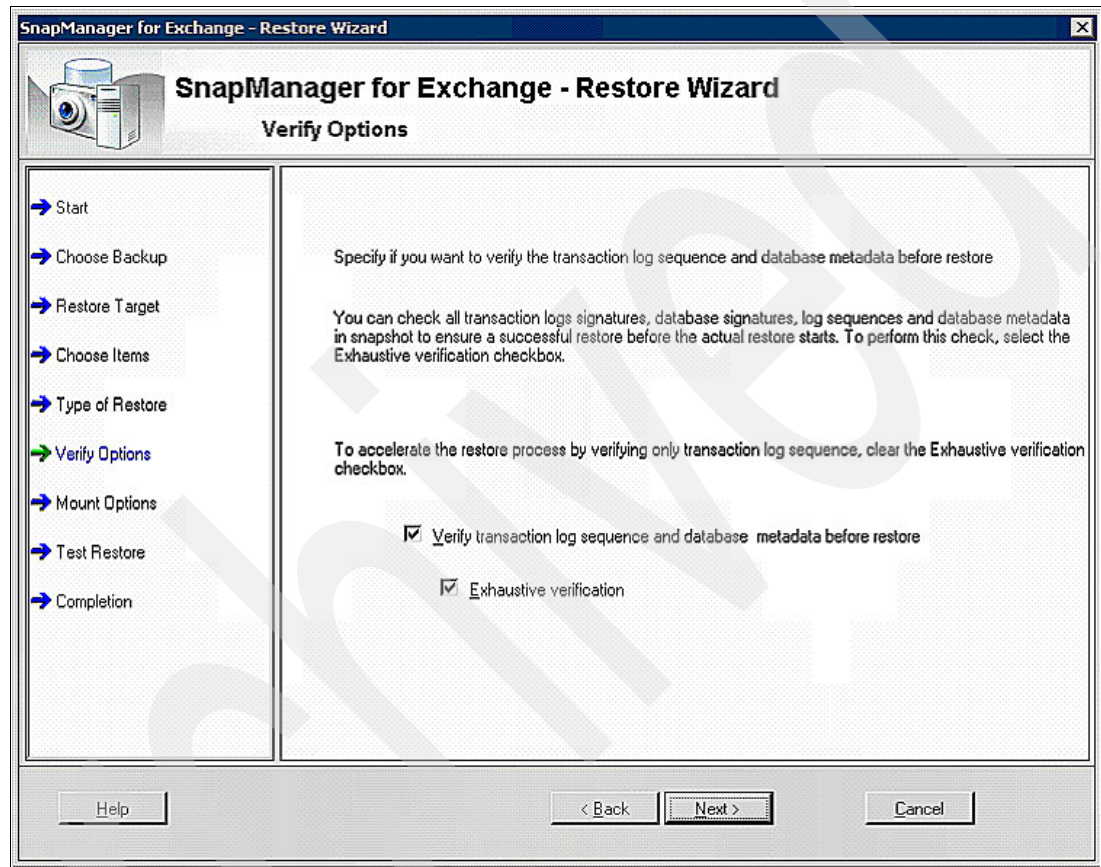


Figure 9-42 Verify Options window

12. In the Mount Options window (Figure 9-43), check the check box **Recover and Mount databases after restore** if you want SME to automatically perform a soft recovery on the Microsoft Exchange database just restored and mount the database after the process. If you plan to mount the database at a later time manually, just uncheck the check box and, at the end of the process, SME will not perform the soft recovery on the database. Click **Next**.

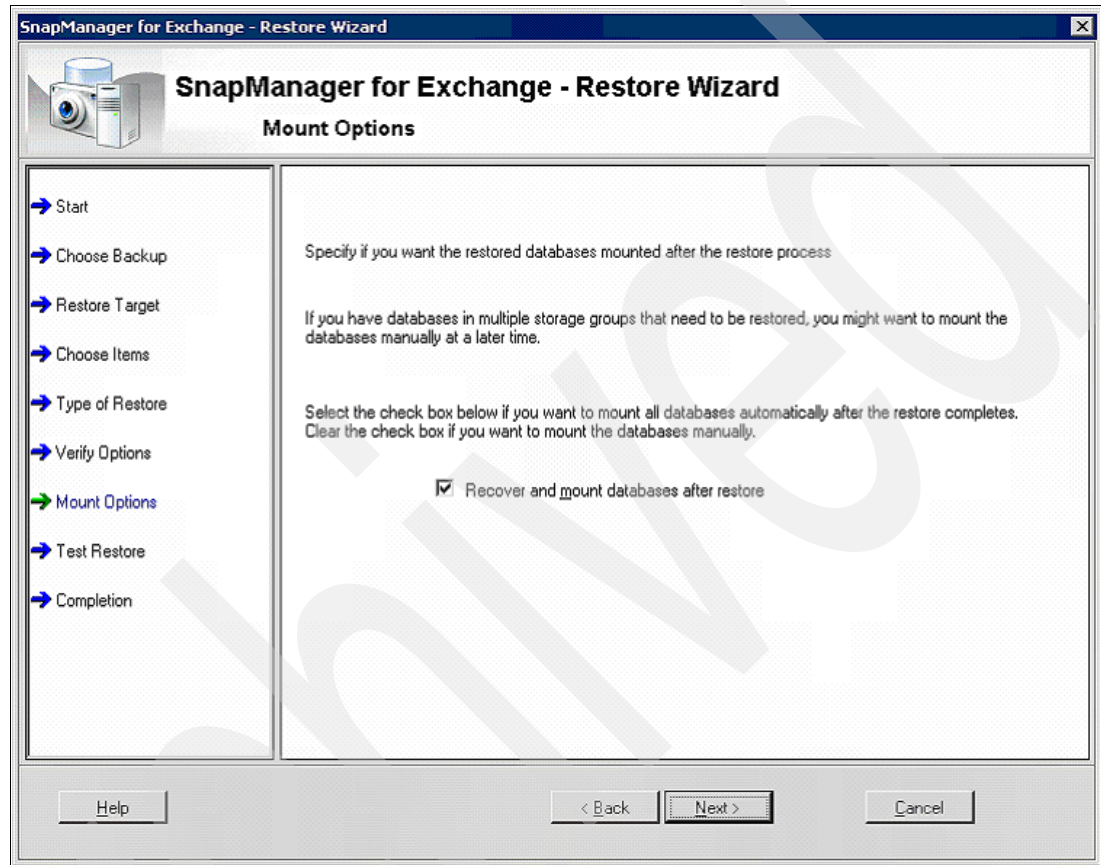


Figure 9-43 Mount Options window

13. In the Summary window (Figure 9-44), review the options selected and click **Finish**. The Status window will be shown. Click **Start Now** to process the restore and execute all the selected steps.

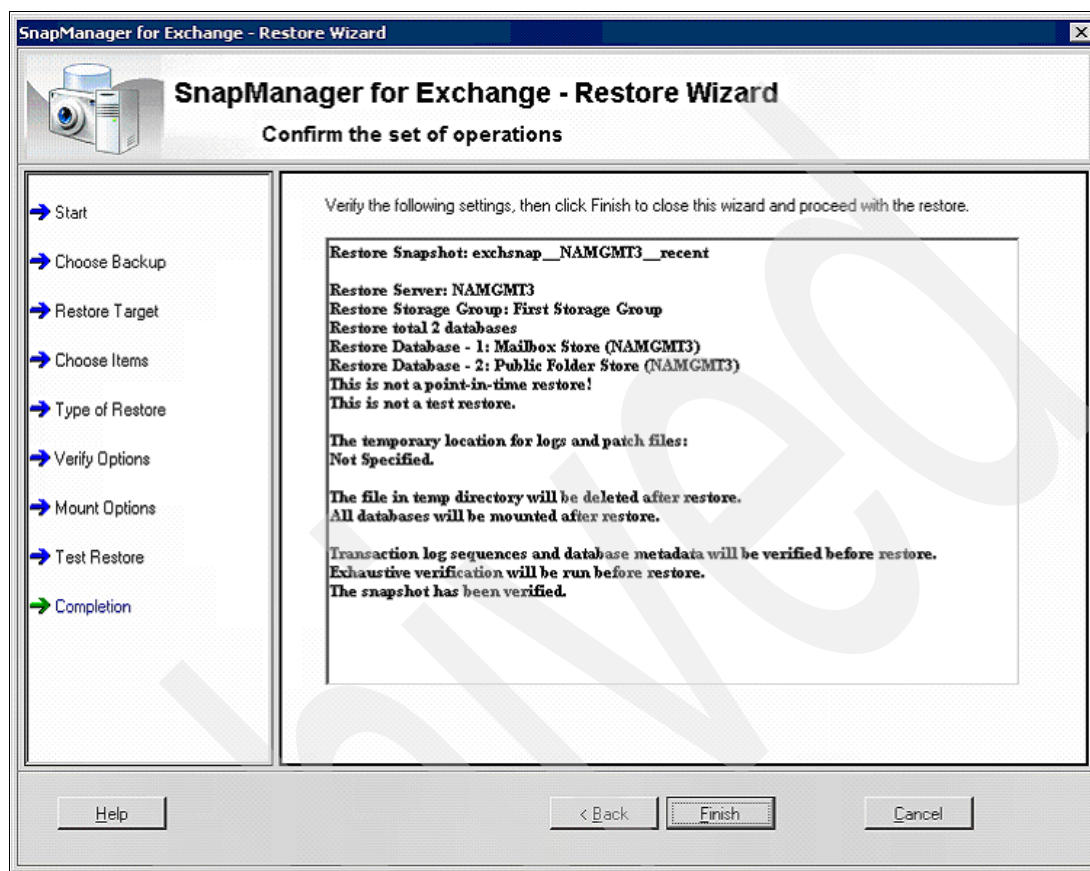


Figure 9-44 Summary window

When the restore operation is done, the Mailbox Stores will be mounted if you selected to mount them after the restore.

Restore to Recovery Storage Groups

There are some special situations when you need to restore the data to a different place than the production environment. For example, when the Microsoft Exchange server infrastructure on the IBM System Storage N series storage system is configured to have a single LUN for the entire Storage Group (and therefore for each and all of the Mailbox Stores on that Storage Group), you need to restore only one of the Mailbox Stores. If a restore operation is executed using the option to restore to the original Storage Group, all mailbox stores would be overwritten by the restore operation.

Another situation is when you need to restore a single mailbox without using the N series Single Mailbox Recovery software. In this case, the restore operation should be executed on a Recovery Storage Group and then have the specific mailbox exported from the Mailbox Store on the Recovery Storage Group and merged on the production environment.

In order to restore a SnapManager backup set to a Recovery Storage Group, some manual interventions are necessary on the Microsoft Exchange server.

Note: SnapManager for Microsoft Exchange does not support restoring SnapManager backup sets to the existing Recovery Storage Group. Before executing the restore operation, any Recovery Storage Group should be removed from the Exchange System Manager.

These are the steps to restore a SnapManager backup set to a Recovery Storage Group:

1. In the Microsoft Exchange server, select **Start → All Programs → IBM → SnapManager for Exchange Management Console**. To start the restore process, right-click the server name and select **Restore Wizard**. In the Restore Wizard Welcome window (Figure 9-45), click **Next**.



Figure 9-45 Restore Wizard Welcome window

2. In the Exchange Server selection window (Figure 9-46), select whether this restore is from a SnapManager backup set created by the same Exchange server or if it is from a SnapManager backup set created on a different Exchange server and archived. The most common scenario is to restore a backup executed on the Microsoft Exchange server. Click **Next**.

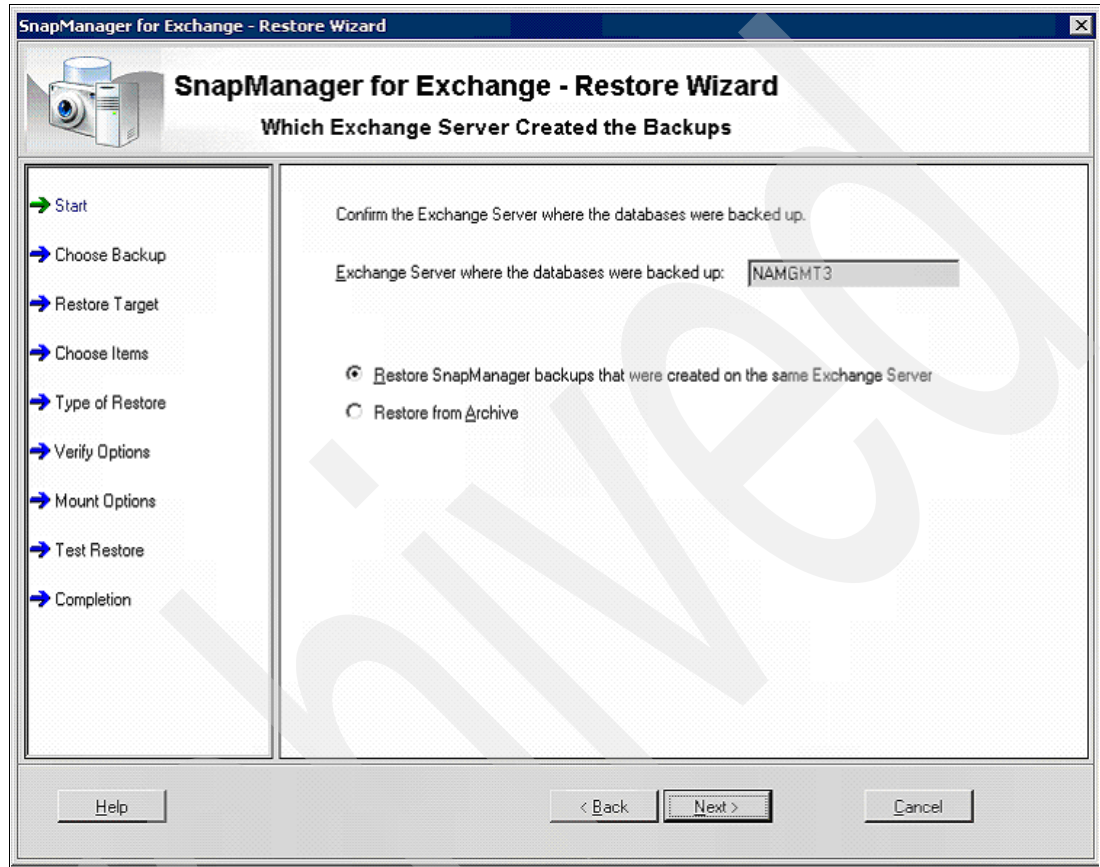


Figure 9-46 Exchange Server selection window

3. In the Choose backup window (Figure 9-47), expand the server name and expand the Storage Group that you plan to restore. Select the SnapManager backup set that you plan to restore. Notice that depending on the convention name selected at the time of the backup, the backup set names can be different. If you used the generic convention naming, the backup set with the *_recent* on the end of its name is the most up to date set. Select the backup set by double clicking it and click **Next**.

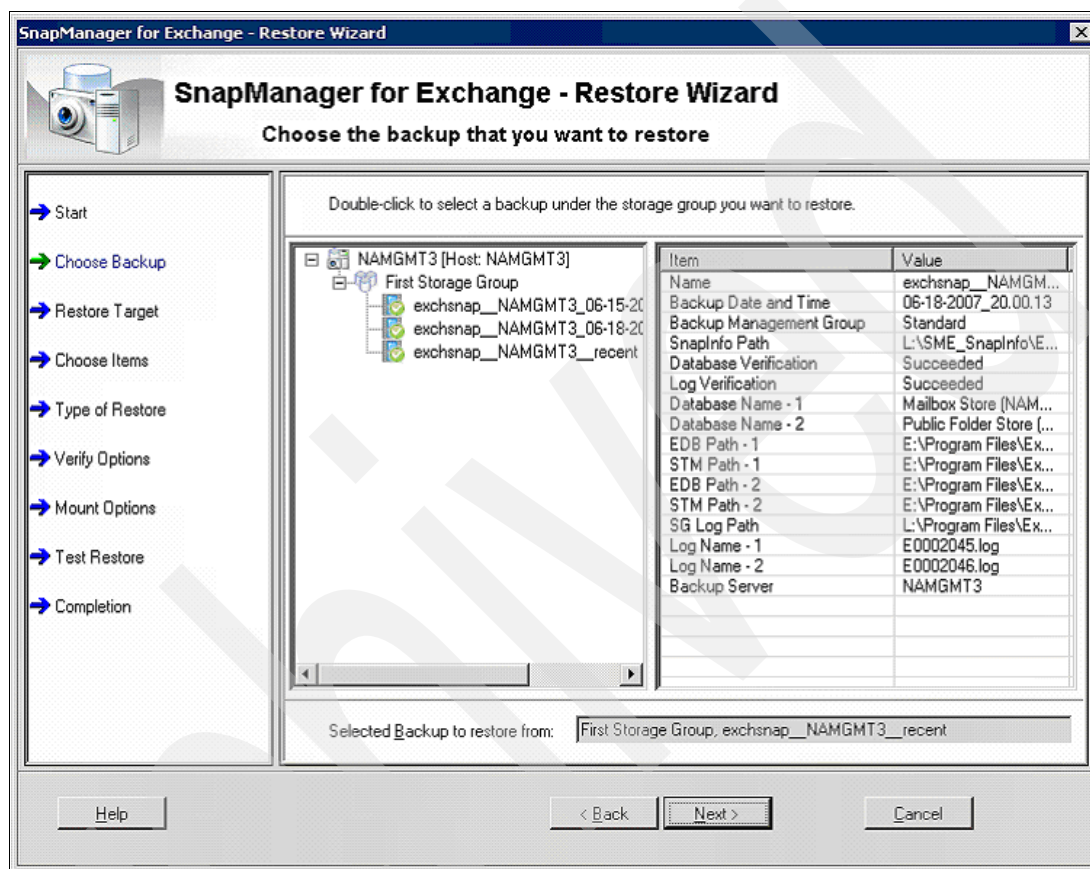


Figure 9-47 Choose backup window

4. In the Choose target window (Figure 9-48), select **Restore to the Recovery Storage Group**. Click **Next**.

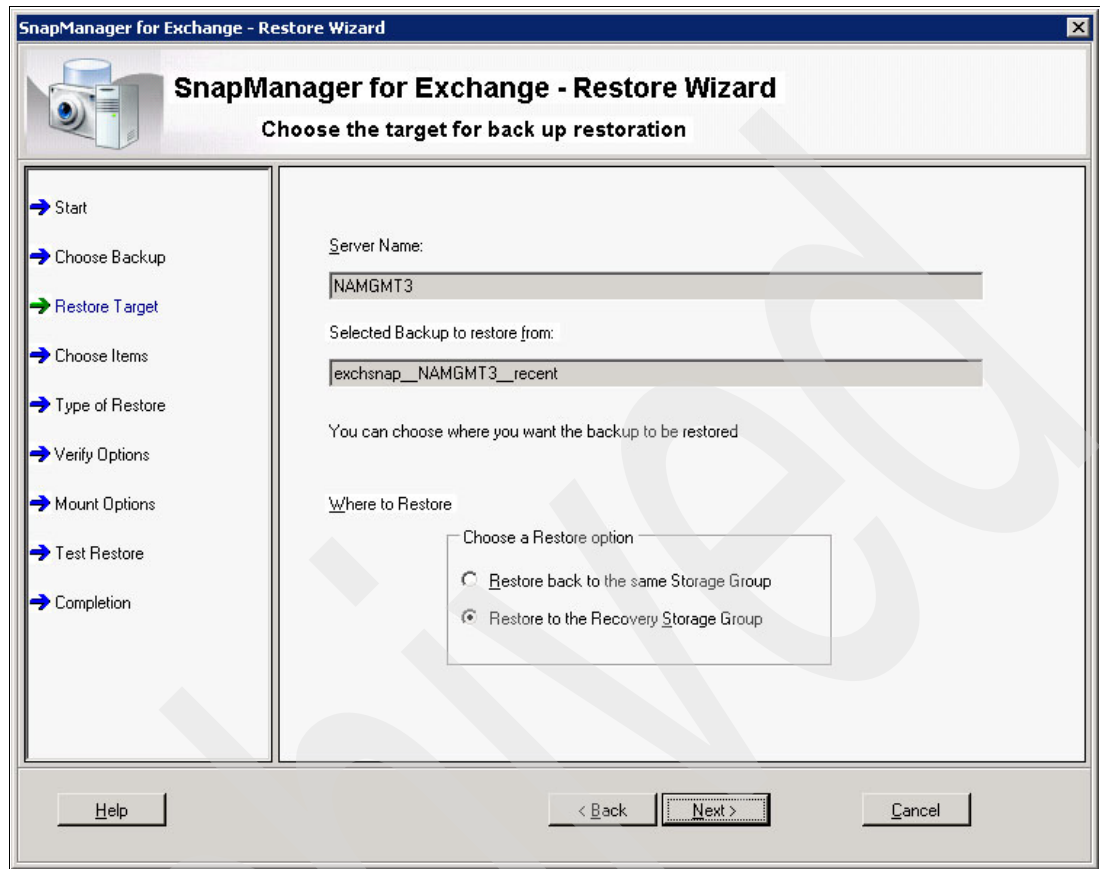


Figure 9-48 Choose target window

5. In the Choose Mailbox Store window (Figure 9-49), select the Mailbox Store that should be restored. Notice that if more than one Mailbox Store are on the same LUN, you cannot select them separately. Click **Next**.

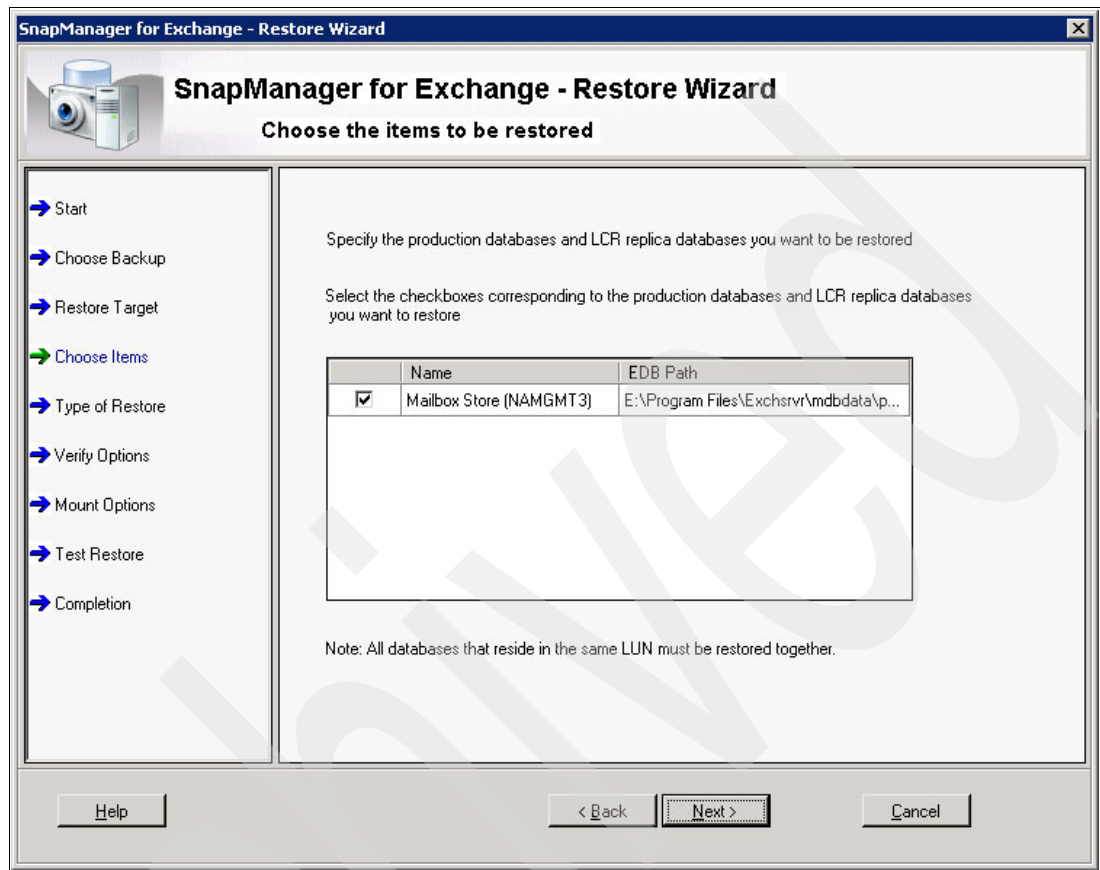


Figure 9-49 Choose Mailbox Store window

6. A verification at the Microsoft Exchange server will be done to check if any Recovery Storage Group exists. If SME finds any Recovery Storage Group configured for the server, you must first remove the Recovery Storage Group using the Exchange System Manager and rerun the Restore Wizard. If SME does not find any Recovery Storage Group at the server, then the Choose Recovery Storage Group window (Figure 9-50) will be shown. Check the Microsoft Exchange server name and click **Next**.

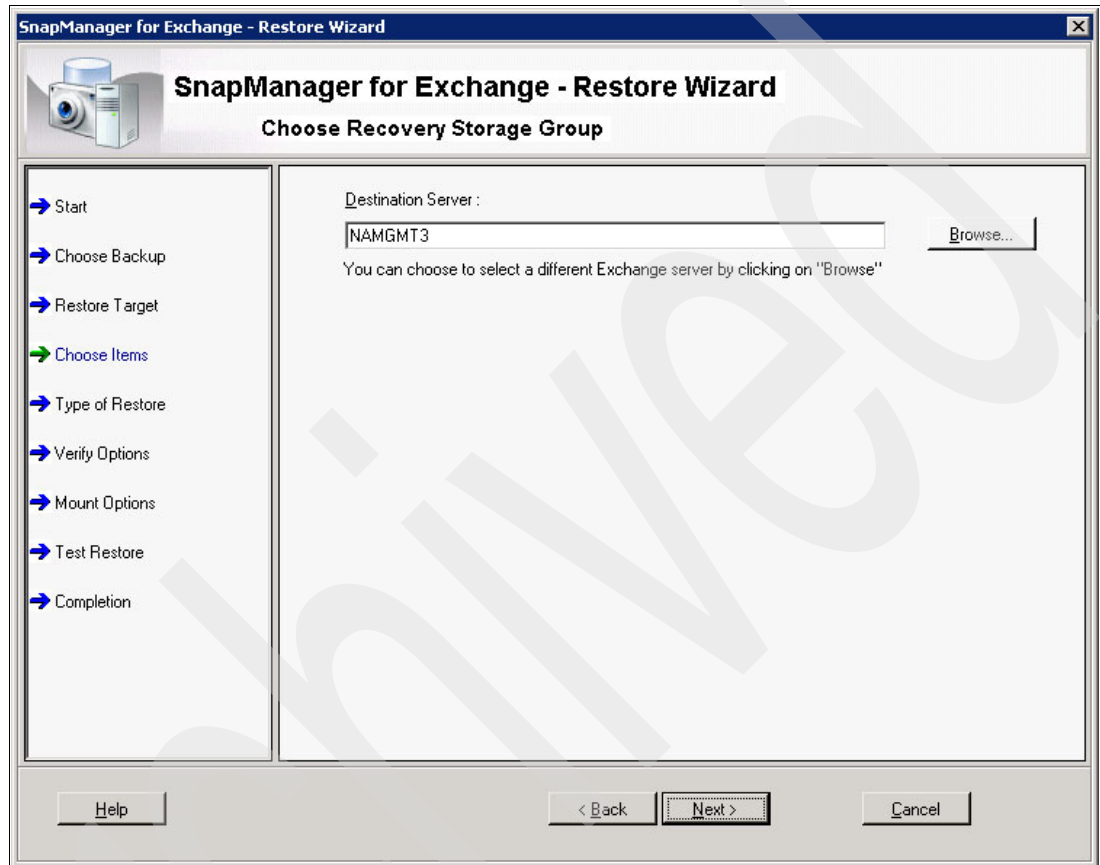


Figure 9-50 Choose Recovery Storage Group window

7. In the Verify Options window (Figure 9-51), specify which validation tests should be done on the database and transaction log files before the actual restore operation starts. To verify only the database and log files sequence in order to make sure no log file is missing, uncheck the **Exhaustive Verification** check box. If you check the **Exhaustive Verification** check box, a complete test on each transaction log file will be done to guarantee that no problem will occur because of any single log file problem. This verification also makes the process more time and resource consuming. After selecting the desired verification, click **Next**.



Figure 9-51 Verify Options window

8. In the Mount Options window (Figure 9-52), check the **Recover and Mount databases after restore** check box if you want SME to automatically perform a soft recovery on the Microsoft Exchange database just restored and mount the database after the process. If you plan to mount the database at a later time manually, just uncheck the check box and at the end of the process, SME will not perform the soft recovery on the database. Click **Next**.



Figure 9-52 Mount Options window

9. In the Summary window (Figure 9-53), review the options selected and click **Finish**. The Status window will be shown. Click **Start Now** to process the restore and execute all the selected steps.

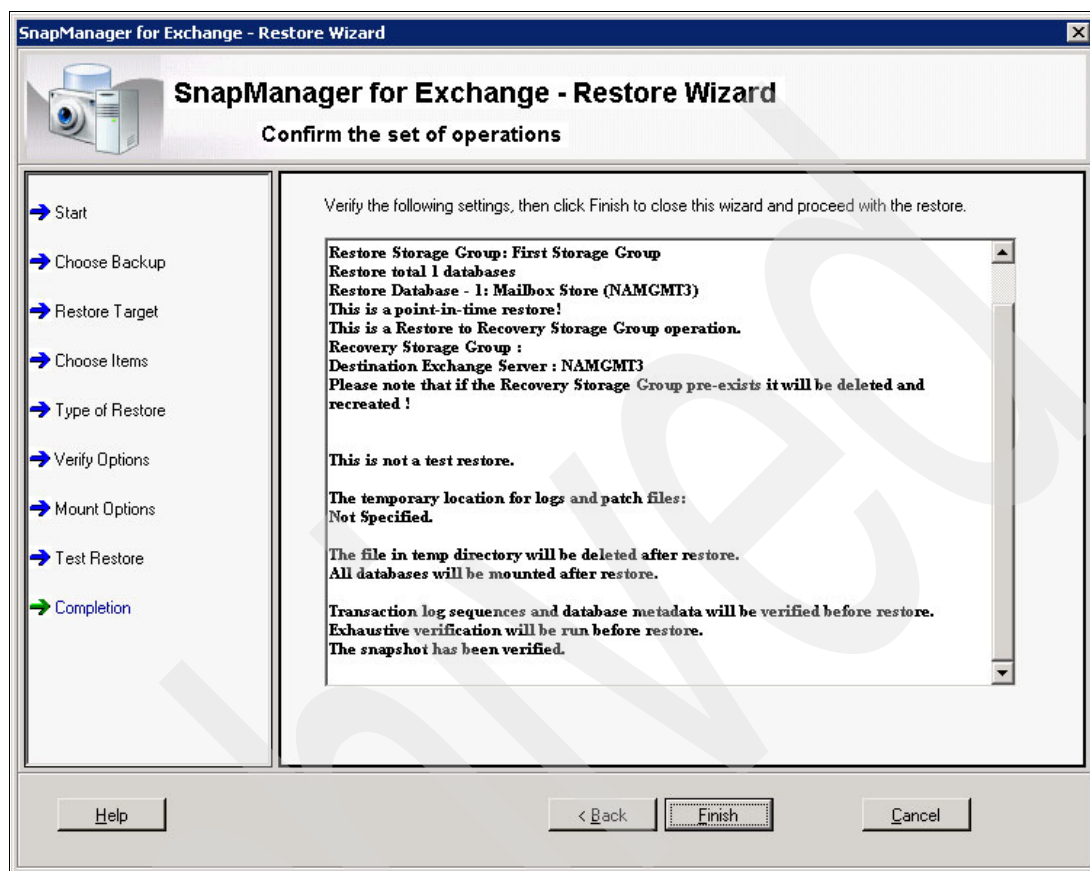


Figure 9-53 Summary window

At the end of the restore process, the Recovery Storage Group will not be created on the Exchange System Manager. A manual process to create and mount it is needed.

Notice that two new disks will be connected on the SnapDrive MMC (Figure 9-54). Depending on how the SME is configured to access LUNs on Snapshots, the drives can be mounted as drive letters (such as the example scenario case) or as a volume mounting point (creating a new folder on the configured path).




Figure 9-54 SnapDrive MMC

One of the drives will contain the database files while the other will contain the transaction log files. If you configured the SME restore operation to mount the database at the end of the process, a soft recovery will already be executed on the database, which means that any necessary transaction log files will already be replayed on the database. The database will be in a consistent state, ready to be mounted on Exchange System Manager.

To mount and enable the contents of the Mailbox Store recovered, follow these steps:

1. In Exchange System Manager, expand all the trees to get to the server being restored. Right-click the server name and select **New** → **Recovery Storage Group**, as shown in Figure 9-55.

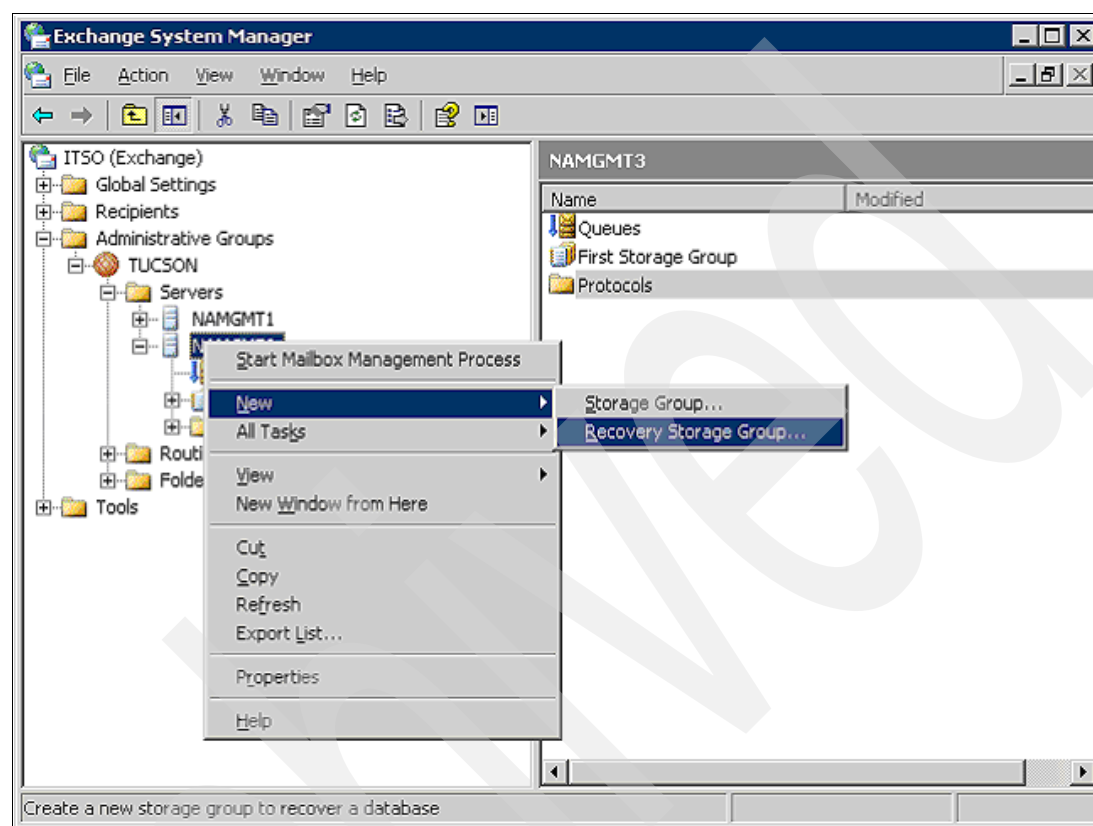


Figure 9-55 New Recovery Storage Group

2. In the Recovery Storage Group Properties window (Figure 9-56), type in a name for the Recovery Storage Group. The Transaction Log Location and System path location fields should point to the new location where the transaction log files were restored. In the example scenario, the path is Y:\Program Files\exchsrvr\MDBDATA. Click **OK**.




Figure 9-56 Recovery Storage Group Properties window

3. The Recovery Storage Group will be created using the information provided for the transaction log files. After that, the Mailbox Store to be restored should be configured in this Recovery Storage Group. In the Exchange System Manager, right-click the Recovery Storage Group and select Add Database to Recover, as shown in Figure 9-57.

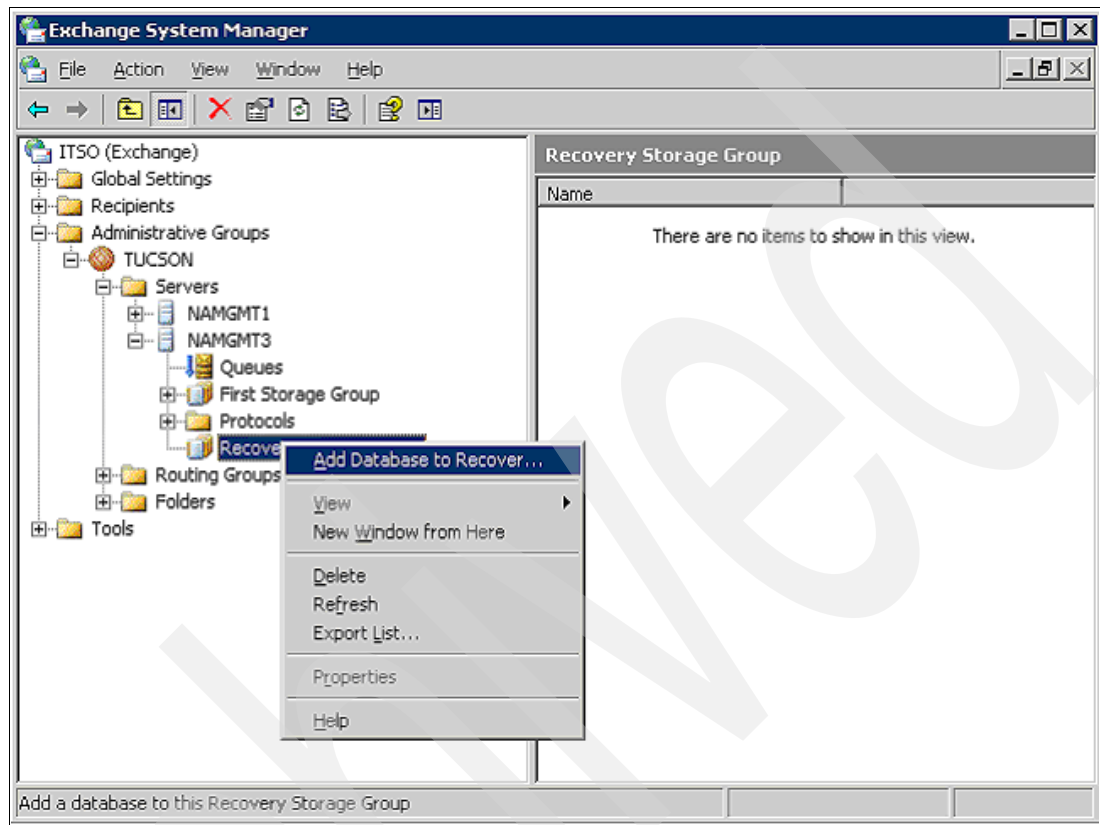


Figure 9-57 Add Database to Recover

4. The Exchange System Manager will automatically search for all available databases on that particular server and show a list of the Mailbox Stores on the server (see Figure 9-58). Select the Mailbox Store that you restored from the SnapManager backup set and click **OK**.

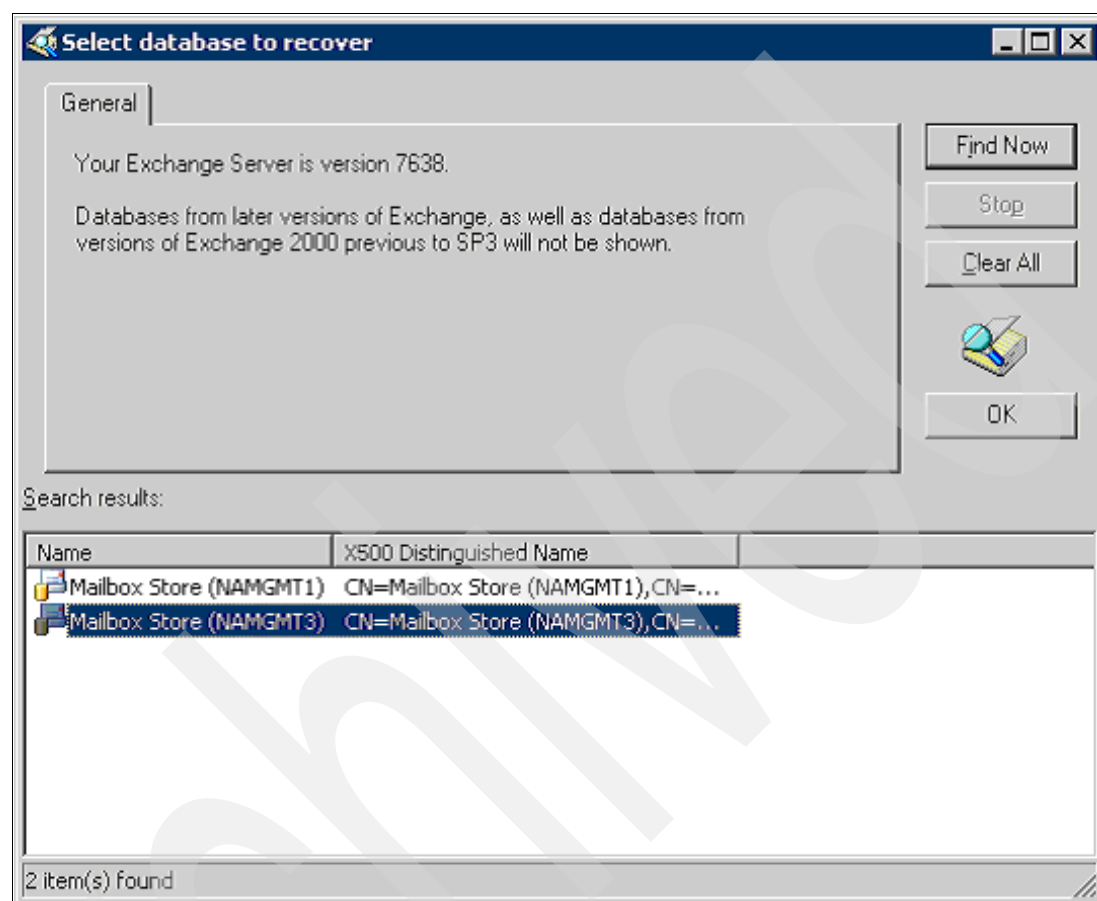


Figure 9-58 Select database to recover

5. The Mailbox Store Properties will be shown. Click the **Database** tab (Figure 9-59) and point the Exchange database and Exchange streaming database fields to the path for the databases you have just restored. Make sure that the **This database can be overwritten by a restore** check box is selected. Click **OK**.

Note: You cannot point to the database file itself at this time because the Recovery Storage Group will try to create a new file. Point the path to the same folder where the database resides and use a different name for the file.

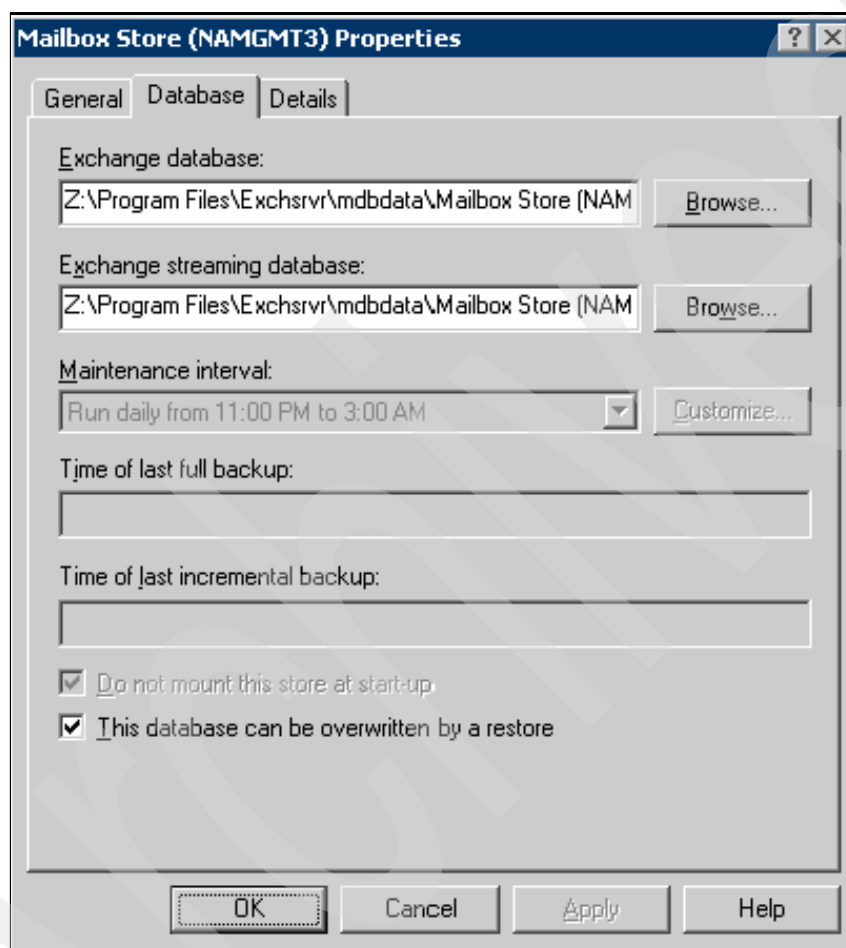


Figure 9-59 Mailbox Store Properties

6. Go to the database folder and rename the original database files (in the example scenario, priv1.edb and priv1.stm) to the names typed in the Mailbox Store Properties window (Figure 9-60).

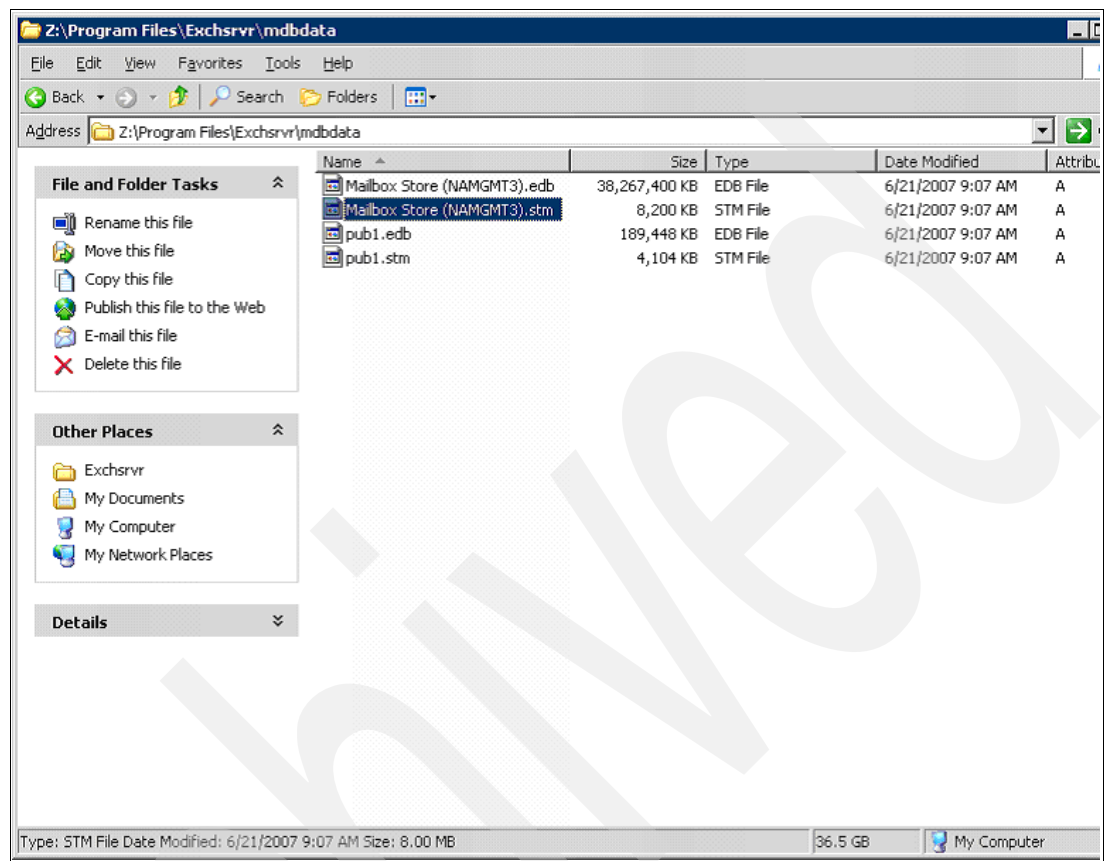


Figure 9-60 Rename database files

7. In the Exchange System Manager, right-click the Mailbox Store in the Recovery Storage Group and select **Mount Store**.

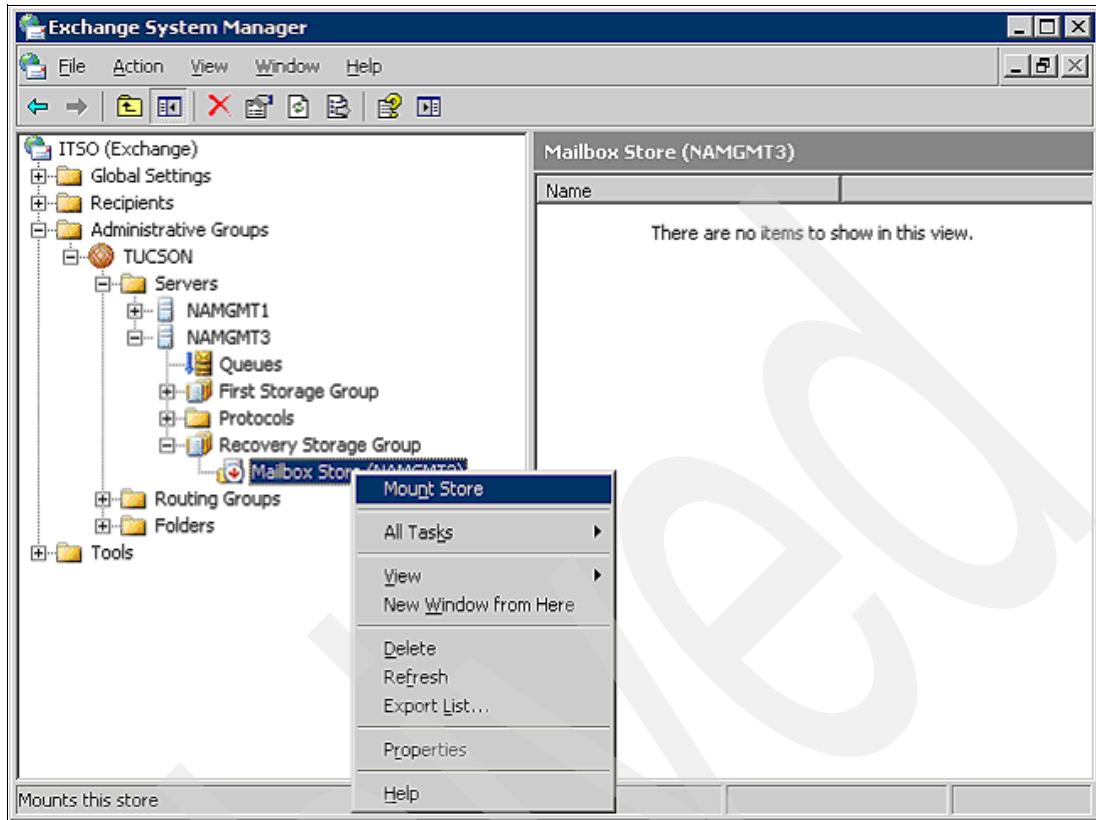


Figure 9-61 Mount Store

8. A warning message will be shown (Figure 9-62) stating that the Mailbox Store should be mounted only after the restore process is complete. Click **Yes** to have the Mailbox Store mounted.

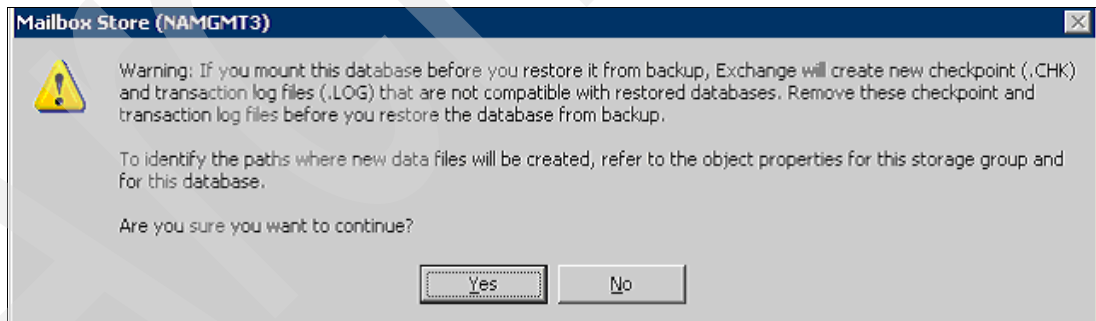


Figure 9-62 Mount Store warning window

9. A warning window will be shown to inform you that the Mailbox was successfully mounted. Click **OK**.

10. The restore process is now complete. In Exchange System Manager, expand the **Recovery Storage Group**, expand the **Mailbox Store** on the Recovery Storage Group, and expand the **Mailboxes** folder. You will notice that all mailboxes were restored (Figure 9-63).

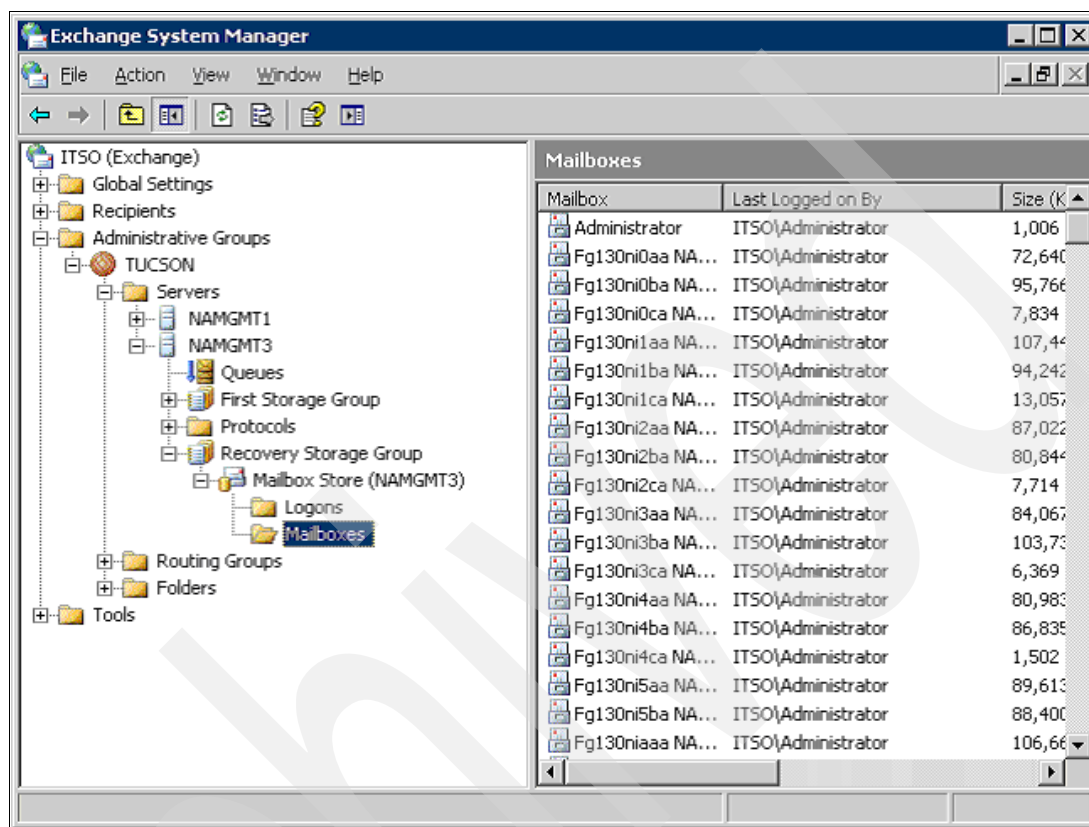


Figure 9-63 Mailbox Store restored

Restore from SnapDrive Snapshots

If a restore from a SnapDrive Snapshot backup is needed, there are some procedures to follow in order to accomplish it.

Most of the time, SnapDrive Snapshot restores will be point in time restores because when restoring the database LUN, the transaction log files LUN has to be restored together to bring the database to a consistent state. Afterwards, not all the transaction log files from that point in time to the last minute prior to the disaster will be available.

For these reasons and a lack of integration with the Microsoft Exchange server API, SnapDrive Snapshot restores are not the recommended way to recover your Microsoft Exchange environment.

To restore a Snapshot created on your Microsoft Exchange server, follow these steps:

1. In the Microsoft Exchange System Manager, expand the server name, expand the Storage Group that host the Mailbox Store to be restored, right-click the Mailbox Store to be restored, and select **Dismount Store**.
2. Dismount any other Mailbox Store on the Storage Group, because the transactional log files are related to the Storage Groups rather than to the Mailbox Stores.

Note: If more than one Microsoft Exchange server's database is configured on the specified LUN, all Mailbox Stores that have databases on that particular LUN will have to be restored together. Also, all Mailbox Stores on the same Storage Group have to be restored together because of the transactional log files.

3. In the Computer Management MMC, expand **Storage** and expand the **SnapDrive MMC**. Verify the disks that must have the Snapshots restored.
4. In the specified disks, expand the **Snapshots** folder. Right-click the desired Snapshot and select **Restore disk from Snapshot**.
5. Repeat the same operation for each of the disks for databases that reside on the Storage Group.
6. Repeat the same operation for the transaction log files LUN to restore it.
7. After all restores complete, right click each of the Mailbox Stores in the Exchange System Manager and select **Mount Store**.

9.1.3 Single Mailbox Recovery

Single Mailbox Recovery software (SMBR) helps administrators recover single mailboxes or even single messages from a Microsoft Exchange database, without the need to recover the whole Storage Group or Mailbox Store.

SMBR works with your existing backup and restore solution to ease the restore of single items when this is the case. In our example scenario, SMBR is being used with the SnapManager for Microsoft Exchange.

Without the use of SMBR, if you need to restore a single mailbox from a backup set, you have two options: Restore the whole backup set and export the mailbox to a Personal Folders file (.PST), or use a brick-level backup tool (which backs up each mailbox separately). Both solutions are time and resource consuming.

With the use of SMBR, the mailbox (single message, contact, or attachment) can be copied from the backup set directly to a production Microsoft Exchange server or to a Personal Folders file (.PST).

SMBR eliminates the need for a Microsoft Exchange recovery server in order to access the database file (.EDB). It directly access the file and makes the contents available to the administrator.

These are some of the benefits from using SMBR with your backup and restore solution:

- ▶ Minimize the time to restore an individual mailbox. Single Mailbox Recovery can slash restore time, making it possible to restore mail items from a previous full backup directly into your production Microsoft Exchange server, or directly into a new or existing PST file. This eliminates the need for a recovery server and the extra steps required to separately import mail back into Exchange or Outlook.
- ▶ Eliminate backups of an individual mailbox. Single Mailbox Recovery eliminates the need to back up individual mailboxes (brick-level backups) because they can be restored directly from an EDB file.
- ▶ Minimize the time to locate all mail matching specific criteria. Single Mailbox Recovery includes an Advanced Find feature that can search across all mailboxes in an archive EDB file, rather than searching one mailbox at a time or bringing an old backup back online for analysis.

- Minimize the time to back up all mailboxes. Single Mailbox Recovery eliminates the need to back up mailboxes individually.

For detailed information about Single Mailbox Recovery software, refer to the Web site:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskind=1&brandind=5000029&familyind=5329811>

In order to connect to the production Microsoft Exchange server and export messages, SMBR needs a MAPI connection to it. Prior to installing SMBR, Microsoft Outlook must be installed on the server to provide this MAPI connection. After installation, configure an mail account on the Microsoft Outlook and execute it to establish the MAPI connection.

Note: Do not install the Microsoft Outlook on the Microsoft Exchange server. Microsoft's KB266418 (<http://support.microsoft.com/kb/266418>) explains that this is not a supported environment because Outlook's MAPI overwrites Exchange's MAPI.

After Microsoft Outlook is installed and configured, install the Single Mailbox Recovery software on the same server.

Using Single Mailbox Recovery software

1. SMBR operations are based on the Microsoft Exchange database file (.EDB). To access the database file, first you need to identify the SnapManager backup set from which you need to restore the mailbox or message.
2. If you have SnapDrive installed on the SMBR server, directly connect to the Snapshot identified in the step above. If you are not using SnapDrive on the SMBR server, connect to the Snapshot using the SnapManager for Exchange server and share or copy the file to the SMBR. You need to access both the database Snapshot and the transaction log files Snapshot.
3. Open the Single Mailbox Recovery software. Figure 9-64 shows the software interface.

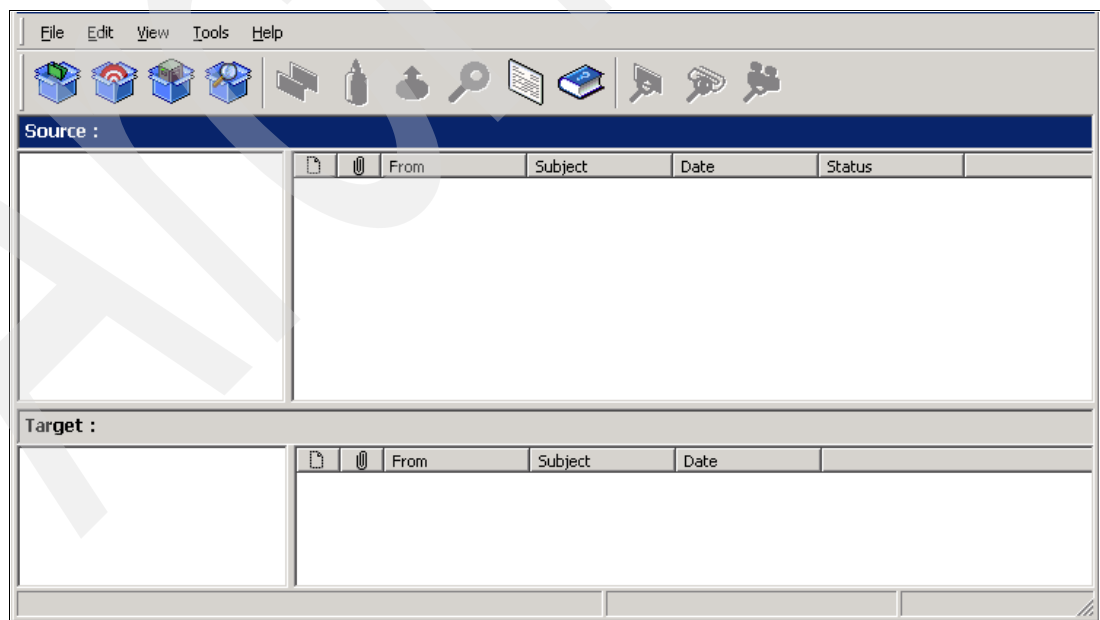


Figure 9-64 Single Mailbox Recovery software

4. To start a recovery process, select **File** → **Open Source** (Figure 9-65 on page 379).

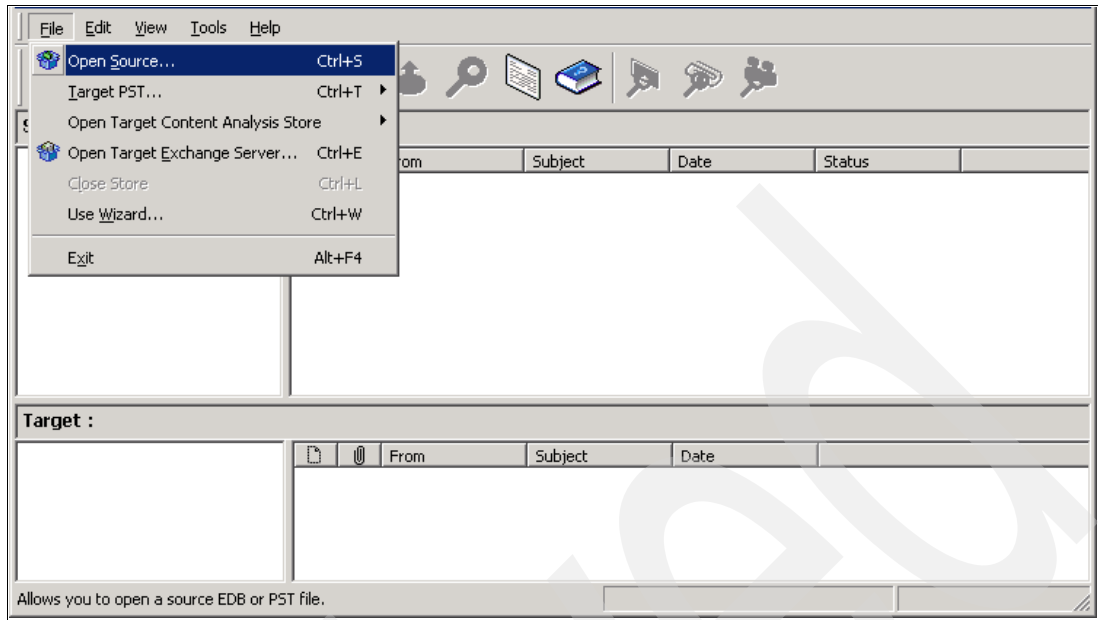


Figure 9-65 Open Source

5. The Select Source Files window will be shown (Figure 9-66). In the Source File field, specify the path to the Microsoft Exchange database on the connected SnapDrive. The path varies depending whether you mapped it as a drive letter or a volume mounting point. In the Log File Path field, specify the path to the transactional log files on the connected SnapDrive. Click **OK**.

Note: When connecting the drives on the SnapDrive, connect the database Snapshot and transactional log files Snapshot from the same SnapManager backup set.

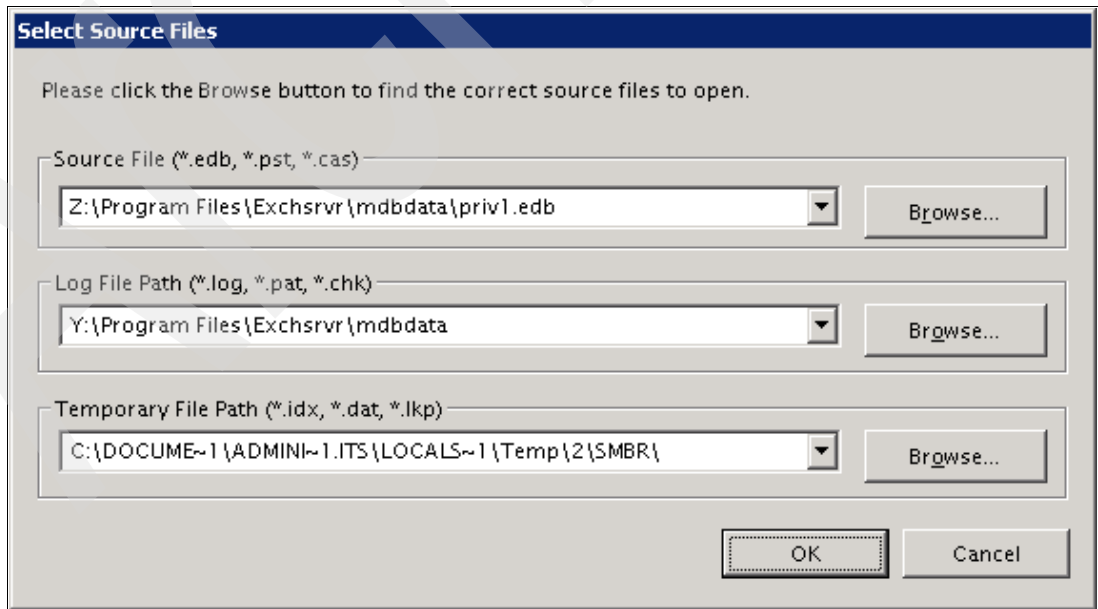


Figure 9-66 Select Source Files window

6. If the temporary folder still does not exist, a warning message will be shown asking for confirmation to create it. Click **Yes** (Figure 9-67).

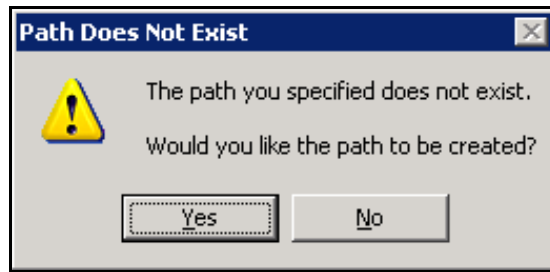


Figure 9-67 Create temporary folder

7. SMBR processes the database file and shows the content (see Figure 9-68) in the Source pane.

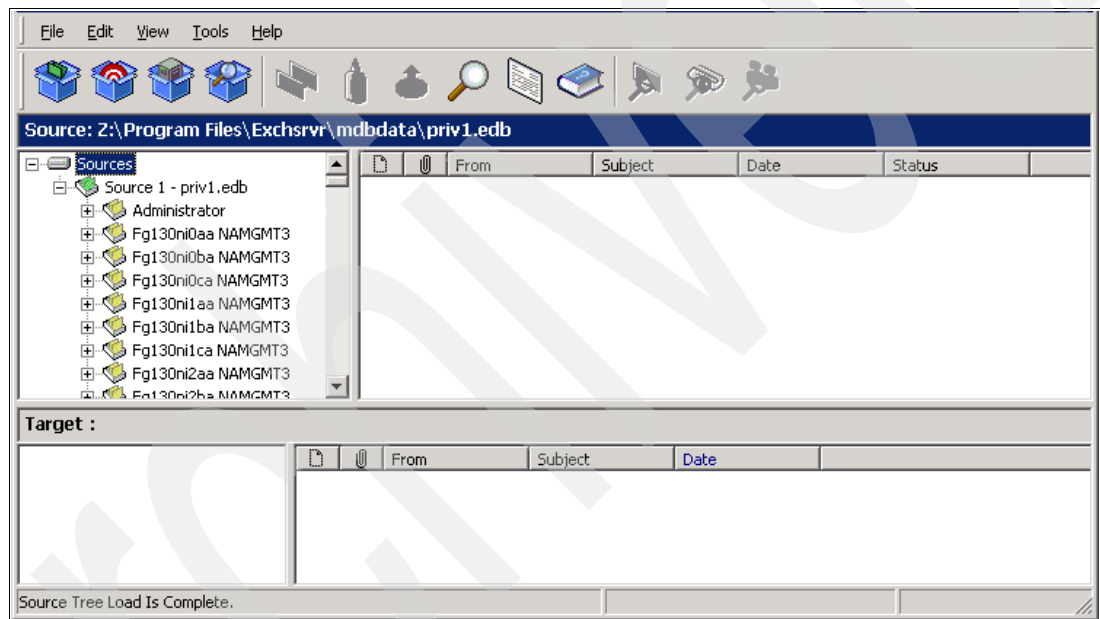


Figure 9-68 Source window

8. Now that the Microsoft Exchange database is open, a target source must be configured. Select **File → Target PST** to have the data exported to a PST. At this time, you can open an existing PST or create a new PST file. To target a Microsoft Exchange server, select **File → Open Target Exchange Server** (Figure 9-69 on page 381).

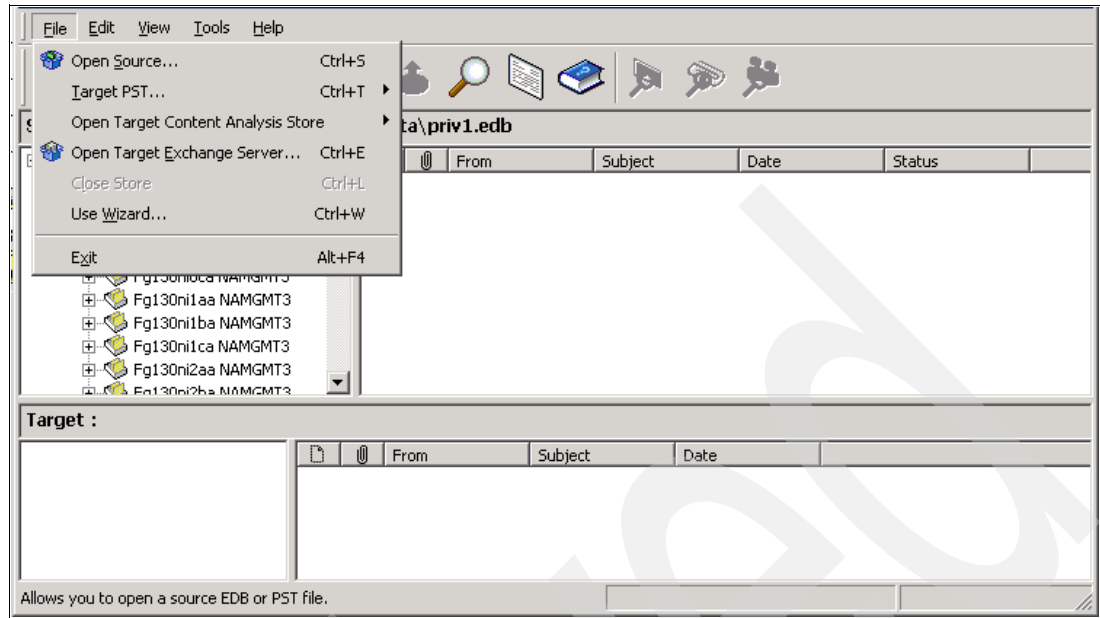


Figure 9-69 Open Target window

9. When connecting to an existing Microsoft Exchange Server, you need to specify which mailbox you are going to access and on which server this mailbox is located (Figure 9-70). To access this mailbox, the account being used has to have access permissions on the specified mailbox. Click **OK**.




Figure 9-70 Connect to Exchange Server window

10. The target Microsoft Exchange mailbox is shown in the Target Pane (Figure 9-71).

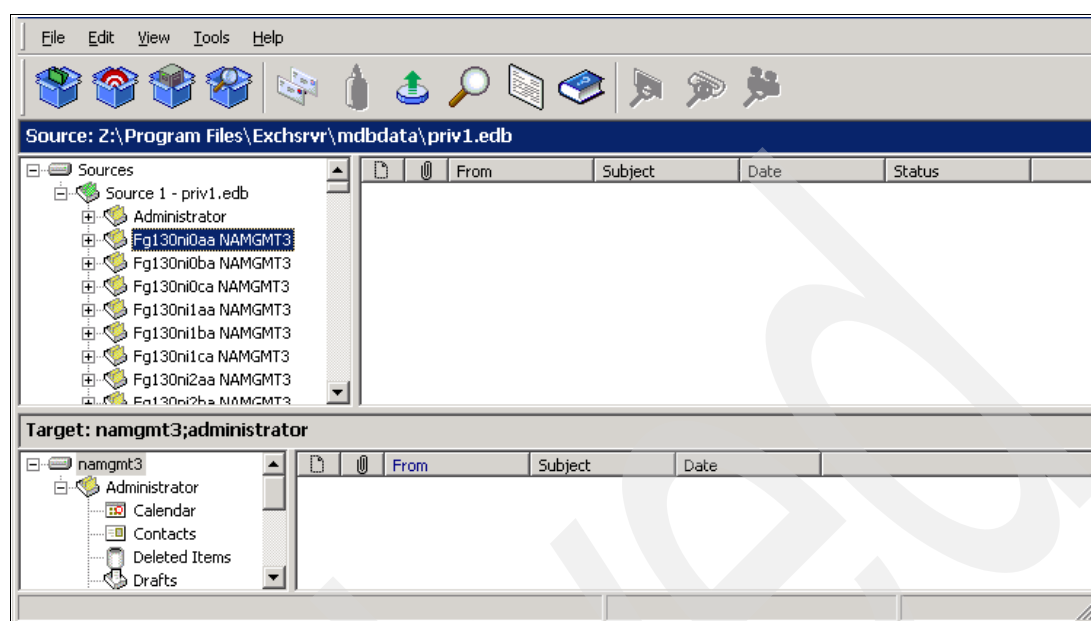


Figure 9-71 Target mailbox open

11. Find the data you need to export from the Source pane, and drag and drop it to the target pane on the folder to which you need the data copied.

12. After the copy process is done, the Copy Summary window (Figure 9-72) will be shown. Click **Close**.

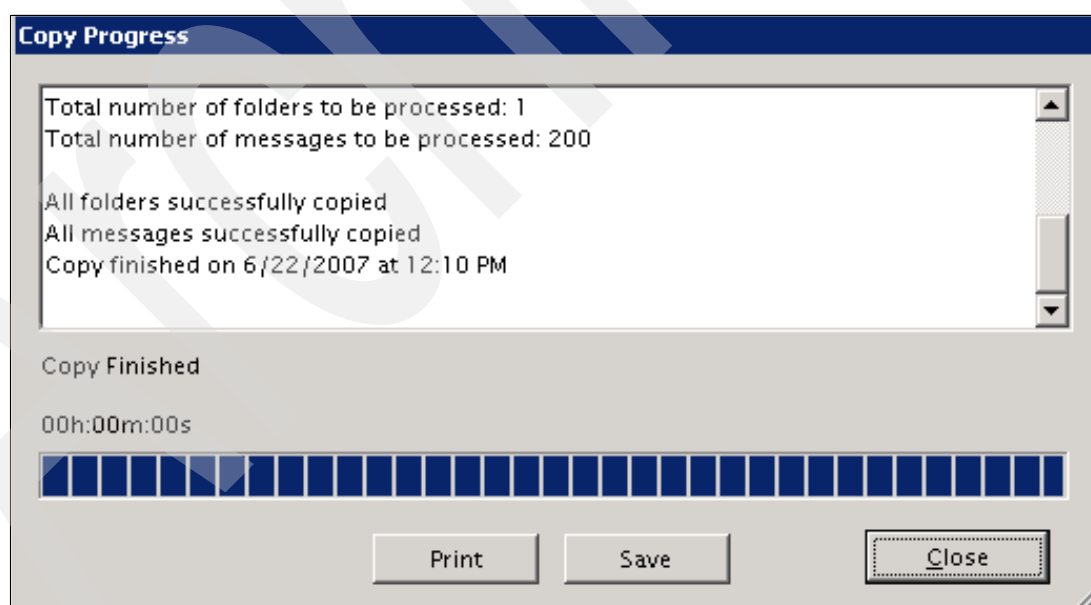


Figure 9-72 Copy summary window

13. If instead of copying data from a specific mailbox you want to copy a mailbox itself, use a Personal Folders file (.PST) as the target instead of a Microsoft Exchange server.

14. After the process is complete, disconnect each of the Snapshot LUNs from SnapDrive.

Backup and restore of Lotus Domino server

Backup and restore operations are critical administrative tasks for mail systems such as a Lotus Domino server. The backup should not only guarantee the recovery of data in case of a disaster but also must guarantee quick single database recovery.

Another important fact is the rapid growth of user and application databases. It is not unusual to see Lotus Domino servers with several thousand user databases. The large number and size of the single database files (and therefore the transaction log files) demand considerable time for backup and for restore. Sometimes the backup job can last more than 24 hours, becoming useless.

A restore operation is also a concern in cases where you need to restore either the whole server or a single database. Depending on the size of the database and location of your backup (disk, tape, or other media or location), it can take some hours get the data back on the Lotus Domino server.

SnapDrive for Linux and UNIX eases the backup and restore operations by using the Snapshot technology and providing quick and verified backup sets for the environment. In junction with SnapMirror, you benefit from the powerful replication mechanism and a disaster recovery solution.

10.1 Backup and restore

The IBM System Storage N series storage operating System Data ONTAP provides a variety of methods for protecting data in an iSCSI or Fibre Channel SAN. These methods, described in Table 10-1, are based on Snapshot technology in Data ONTAP, which enables you to maintain multiple read-only versions of LUNs online per volume.

Snapshot copies are a standard feature of Data ONTAP. A Snapshot copy is a frozen, read-only image of the entire Data ONTAP file system, or Write Anywhere File Layout (WAFL) volume, which reflects the state of the LUN or the file system at the time the Snapshot copy is created. The other data protection methods listed in Table 10-1 rely on Snapshot copies or create, use, and destroy Snapshot copies, as required.

Table 10-1 N series data protecting methods

Method	Used to...
Snapshot	Make point-in-time copies of a volume.
SnapRestore	<ul style="list-style-type: none">▶ Restore a LUN or file system to an earlier preserved state in less than a minute without rebooting the storage system, regardless of the size of the LUN or volume being restored.▶ Recover from a corrupted database or a damaged application, a file system, a LUN, or a volume by using an existing Snapshot copy.
SnapMirror	<ul style="list-style-type: none">▶ Replicate data or asynchronously mirror data from one storage system to another over local or wide area networks (LANs or WANs).▶ Transfer Snapshot copies taken at specific points in time to other storage systems. These replication targets can be in the same data center through a LAN or distributed across the globe connected through metropolitan area networks (MANs) or WANs. Because SnapMirror operates at the changed block level instead of transferring entire files or file systems, it generally reduces bandwidth and transfer time requirements for replication.
SnapVault	<ul style="list-style-type: none">▶ Back up data by using Snapshot copies on the storage system and transferring them on a scheduled basis to a destination storage system.▶ Store these Snapshot copies on the destination storage system for weeks or months, allowing recovery operations to occur nearly instantaneously from the destination storage system to the original storage system.
SnapDrive for Windows, Linux or UNIX	<ul style="list-style-type: none">▶ Manage storage system Snapshot copies directly from a Windows or UNIX host.▶ Manage storage (LUNs) directly from a host.▶ Configure access to storage directly from a host. <p>SnapDrive for Windows supports Windows 2000 Server and Windows Server 2003.</p> <p>SnapDrive for UNIX supports a number of UNIX environments. For a list of supported host environments and additional information about the SnapDrive feature, see http://www.ibm.com/storage/support/nas/</p>

Method	Used to...
Native tape backup and recovery	<p>Store and retrieve data on tape.</p> <p>Data ONTAP supports native tape backup and recovery from local, gigabit Ethernet, and Fibre Channel SAN-attached tape devices. Support for most existing tape drives is included, as well as a method for tape vendors to dynamically add support for new devices. In addition, Data ONTAP supports the Remote Magnetic Tape (RMT) protocol, allowing backup and recovery to any capable system. Backup images are written using a derivative of the BSD dump stream format, allowing full file-system backups as well as nine levels of differential backups.</p>
NDMP	<p>Control native backup and recovery facilities in storage systems and other file servers. Backup application vendors provide a common interface between backup applications and file servers.</p> <p>NDMP is an open standard for centralized control of enterprise-wide data management. For more information about how NDMP-based topologies can be used by storage systems to protect data, see the <i>IBM System Storage N series Data ONTAP Storage Management Guide</i> at: http://www.ibm.com/servers/storage/support/nas/dataontap/</p>

In this chapter, we discuss the Snapshot through SnapMirror backup procedure in conjunction with SnapDrive on Linux and Solaris, HP UNIX, and AIX, hereafter referred to simply as Linux and UNIX. We describe several features for backing up and restoring the Lotus Domino database and transactional log files. At the end of 10.1.1, “Backup” on page 385, we describe our recommended backup strategy for Lotus Domino on Linux and UNIX.

Note: For more information about the backup and recovery methods, refer to *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide* and *IBM System Storage N series Block Access Management Guide for FCP and iSCSI* at:
<http://www.ibm.com/servers/storage/support/nas/dataontap/>

10.1.1 Backup

In this chapter, we discuss several backup solutions with the N series and IBM Lotus Domino. If you use CIFS or NFS with N series and Lotus Domino, you might want to use the Backup to Disk SnapVault feature of N series and an external backup server to connect the N series share for backup. This provides a very easy backup and recovery scenario. When using LUNs, as we do in our case, there are some restrictions. This chapter will help you to design the backup strategies depending on your business needs.

Using Snapshot

A Snapshot copy is a frozen, read-only image of a traditional volume, a flexible volume, or an aggregate that reflects the state of the file system at the time the Snapshot copy was created. Snapshot copies are your first line of defense for backing up and restoring data.

Some facts about Snapshot copies:

- Data ONTAP can be configured to maintain a configurable Snapshot schedule that creates and deletes Snapshot copies automatically for each volume.

- Use SnapDrive to make Snapshot copies of LUNs. It takes care of flushing all host operating system buffers.
- You can store up to 255 Snapshot copies at one time on each volume.
- You can specify the percentage of disk space that Snapshot copies can occupy. The default setting is 20% of the total (both used and unused) space on the disk.

Note: For more information about requirements, hints, and commands for SnapDrive, we highly recommend the *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssglS7001600&aid=1>

When you create a Snapshot copy of a storage entity, such as a file system or volume group, SnapDrive creates a Snapshot copy that contains the image of all the storage system volumes that comprise the entity you specified using a *file_spec* argument. Use the **snapdrive snap create** command to create a new Snapshot:

```
snapdrive snap create file_spec -snapname snap_name
```

The *file_spec* argument specifies the storage entity, such as the file system or LUN that SnapDrive uses to create the Snapshot copy, while *snap_name* is the Snapshot name.

The **snapdrive snap create** command in Example 10-1 on page 387 creates a Snapshot named snap1 of the volume vol0.

SnapDrive makes consistent storage components that comprise the entity you requested in the Snapshot copy. This means that LUNs or directories being used outside those specified by the **snapdrive snap create** command's *file_spec* argument may not have consistent images in the Snapshot copy. SnapDrive for Linux and UNIX enables you to restore only the entities specified by the *file_spec* argument that are made consistent in the Snapshot copy.

Snapshot copies of entities contained on a single storage system volume are always crash consistent. SnapDrive for Linux and UNIX takes special steps to ensure that Snapshot copies that span multiple storage systems or storage system volumes are also crash consistent (for example, Lotus Domino database and transactional logging files).

For a complete list of supported host environments with the SnapDrive feature, refer to the following URL:

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

Data ONTAP Version 7.2 and greater provides support for consistency groups and storage system fencing. SnapDrive for UNIX uses these features to ensure that all Snapshot copies that span multiple volumes are crash consistent.

Note: When you create a Snapshot copy that spans multiple storage system volumes on storage systems prior to Version 7.2, SnapDrive for UNIX ensures consistency by freezing I/O to the requested LUNs. If a freeze is not provided by the host, for example, on a Linux host, SnapDrive for UNIX makes a “best effort” to create a consistent Snapshot copy by taking the Snapshot copy without freezing the target storage, and then checking for read-write I/Os that occurred to the storage entities when the Snapshot copy was taken. If SnapDrive can create a crash consistent Snapshot copy, the **snap create** command succeeds. If it cannot, SnapDrive discards the Snapshot copy and informs the user of the failure. SnapDrive will never create a Snapshot copy unless the data is crash consistent.

For creating Snapshots of Lotus Domino database and transactional log volumes with SnapDrive, we recommend the usage of Data ONTAP Version 7.2 because of its improved support for consistency groups.

To create a crash consistent Snapshot copy across multiple volumes, SnapDrive for UNIX:

- ▶ Freezes I/O to every volume that contains a storage entity.
- ▶ Makes a Snapshot copy of each volume.

The time it takes to fence the volume and create the Snapshot copy is limited, and is controlled by Data ONTAP.

To ensure that a Snapshot copy is “application-consistent,” you may need to stop or do whatever steps are required to quiesce the application before making the Snapshot copy. Note that database hot backup facilities depend on the methods used by the DBMS, and do not always quiesce I/O to database files.

If the application has not completed its transactions and written data to the storage system, the resulting Snapshot copy might not be application consistent.

Note: IBM Lotus Domino can recover its databases from a crash consistent Snapshot copy, so you do not need to stop it first. For this action, use transactional logging to be sure that no data inconsistency can occur.

Create Snapshots with SnapDrive

To create a Snapshot copy from the Lotus Domino database and transactional log LUN using SnapDrive for Linux and UNIX, complete the following steps:

1. Create the Snapshot with the **snapdrive snap create** command (Example 10-1).

Example 10-1 Creating the Snapshot of Lotus Domino database and transactional log

```
[root@domino1 /]# snapdrive snap create -vg db_SdDg log_SdDg -snapname snap1
Successfully created Snapshot snap1 on 2 filer volumes:
    itsotuc3:/vol/vol_DominoDB
    itsotuc3:/vol/vol_DominoLog
```

```
Snapshot snap1 contains:
disk group db_SdDg containing host volumes
    db_SdHv (filesystem: /notesdata/db)
disk group log_SdDg containing host volumes
    log_SdHv (filesystem: /notesdata/log)
```

```
[root@domino1 /]#
```

Where:

- vg: This switch follows the volume group from which a Snapshot should be created. It is allowed to stack more than one volume group for crash consistency.
- snapname: The short name of your Snapshot. In our case, it is snap1.

Note: Unless you specify otherwise, SnapDrive assumes that all entities that you specify on a given **snap create** command line are related, that is, that the validity of updates to one entity may depend on updates to the other entities specified. When storage entities have "dependent writes" in this way, SnapDrive takes steps to create a Snapshot copy that is crash consistent for all storage entities as a group.

List all created Snapshots

Use the **snapdrive snap list** command for a list of all Snapshot on the specified file system or volume (Example 10-2).

Example 10-2 List all Snapshots

```
[root@domino1 /]# snapdrive snap list -vg db_SdDg log_SdDg

snap name host date snapped
-----
itsotuc3:/vol/vol_DominoDB:snap1 domino1 Jun  8 16:24 db_SdDg log_SdDg
itsotuc3:/vol/vol_DominoLog:snap1 domino1 Jun  8 16:24 db_SdDg log_SdDg
[root@domino1 /]#
```

Where -vg is a list all Snapshots of specified volume groups.

Note: At the end of every row, the Snapshot group is quoted. During the time the Snapshot was created, SnapDrive took stopped all I/O operation on db_SdDg and log_SdDg.

Snapshots can now be made for further backup operations. For example, you can connect (with the **snapdrive snap connect** command) the Snapshot and make a regular file based backup from your host, or do a fast restore directly from the Snapshot with SnapRestore for recovering your whole LUN.

LUN cloning

A LUN clone is a point-in-time, writable copy of a LUN in a Snapshot copy. Changes made to the parent LUN after the clone is created are not reflected in the Snapshot copy.

A LUN clone shares space with the LUN in the Snapshot copy being made. The clone does not require additional disk space until changes are made to it. You cannot delete the Snapshot copy until you split the clone from it. When you split the clone from the Snapshot copy, you copy the data from the Snapshot copy to the clone. After the splitting operation, both the Snapshot copy and the clone occupy their own space.

The main difference between a LUN Snapshot and LUN cloning is that a cloned LUN is writable. As data is written to either the parent or the FlexClone volume, that data is no longer shared between the parent and FlexClone volumes, and the FlexClone volume starts to require more space from its containing aggregate, depending on the size of your changed data.

LUN cloning is not recommended as a main backup method, because all data still resides on the same volume and aggregate.

You might use LUN cloning for following reasons:

- ▶ You need to create a temporary copy of your Lotus Domino data for testing purposes.
- ▶ You need to make a copy of your data available to additional users without giving them access to the production data.
- ▶ You want to create a clone of the Lotus Domino databases for manipulation and projection operations, while preserving the original data in unaltered form.
- ▶ You want to access a specific subset of a LUN's data (a specific logical volume or file system in a volume group, or a specific file or set of files in a file system) and copy it to the original LUN, without restoring the rest of the data in the original LUN. This works on operating systems that support mounting a LUN and a clone of the LUN at the same time. SnapDrive for Linux and UNIX allows this with the **snap connect** command (this could be done with a Snapshot as well).

Note: Because the data still resist on the same physical storage, operations on the LUN, such as backup, will impact the performance of the productive Lotus Domino server.

Follow these steps to create a LUN clone:

1. Create a Snapshot of your LUN with SnapDrive on the Lotus Domino server.

Example 10-3 shows the creation of an new Snapshot from the Lotus Domino database and transactional log volume group with the name snap1.

Example 10-3 Creating a new Snapshot

```
[root@domino1 /]# snapdrive snap create -vg db_SdDg log_SdDg -snapname snap1
Successfully created Snapshot snap1 on 2 filer volumes:
    itsotuc3:/vol/vol_DominoDB
    itsotuc3:/vol/vol_DominoLog

Snapshot snap1 contains:
disk group db_SdDg containing host volumes
    db_SdHv (filesystem: /notesdata/db)
disk group log_SdDg containing host volumes
    log_SdHv (filesystem: /notesdata/log)
[root@domino1 /]#
```

2. Create the LUN clone with the Data ONTAP command-line interface. Use the following command:

```
lun clone create clone_lun_path -b parent_lun_path parent_snap
```

Where:

- clone_lun_path is the path to the clone you are creating.
- parent_lun_path is the path to the original LUN.
- parent_snap is the name of the Snapshot copy of the original LUN.

On Example 10-4, we create a new clone called db_SdLun_clone1 that is derived from our Snapshot snap1.

Example 10-4 Creating the clone

```
itsotuc3> lun clone create /vol/vol_DominoDB/db_SdLun_clone1 -o noreserve -b
/vol/vol_DominoDB/db_SdLun snap1
itsotuc3>
```

Note: We use the `-o noreserve` option because we do not change any data on our clone. If you create a clone without the `-o noreserve` option, you need at least enough free space on your volume as the size of your parent LUN.

3. Verify the created clone. Use the **lun show** command on the Data ONTAP command-line interface, as shown in Example 10-5.

Example 10-5 Verify the created clone

```
itsotuc3> lun show
/vol/vol_DominoDB/db_SdLun 146.0g (156786229248) (r/w, online, mapped)
/vol/vol_DominoDB/db_SdLun_clone1 146.0g (156786229248) (r/w, online)
/vol/vol_DominoLog/log_SdLun 5.0g (5377097728) (r/w, online, mapped)
itsotuc3>
```

The next steps are optional. We split the clone from the Snapshot, so we have a complete copy of the data. This process will take some time, depending on the amount of data that needs to be copied. After that, we delete the backing Snapshot copy on which the clone relied.

4. Split the clone from the backing Snapshot copy. After that, we have a full copy of our LUN. Take care that you have enough free space on your volume. Run the following command:

```
lun clone split start lun_path
```

where *lun_path* is the path to the cloned LUN.

Example 10-6 shows the command. We split the created `db_SdLun_clone1` clone.

Example 10-6 Splitting the clone

```
itsotuc3> lun clone split start /vol/vol_DominoDB/db_SdLun_clone1
itsotuc3>
```

You can check the status of the split with the **lun clone split status** command, as shown in Example 10-7.

Example 10-7 Check the status of splitting the clone

```
itsotuc3> lun clone split status /vol/vol_DominoDB/db_SdLun_clone1
lun clone split status: Done 11370706 of 38277888 blocks (29% complete).
itsotuc3>
```

Note: The split is completed when the status switches to LUN is not a clone:

```
itsotuc3> lun clone split status /vol/vol_DominoDB/db_SdLun_clone1
lun clone split status: /vol/vol_DominoDB/db_SdLun_clone1: LUN is not a clone
itsotuc3>
```

5. The Snapshot can be deleted when there are no clones relying on it. Example 10-8 on page 391 shows the usage of **snapdrive snap delete** command.

Example 10-8 Deleting the Snapshots

```
[root@domino1 ~]# snapdrive snap delete -snapname  
itsotuc3:/vol/vol_DominoLog:snap1 itsotuc3:/vol/vol_DominoDB:snap1  
snap delete: deleted Snapshot itsotuc3:/vol/vol_DominoLog:snap1  
snap delete: deleted Snapshot itsotuc3:/vol/vol_DominoDB:snap1  
[root@domino1 ~]#
```

Since you have a new LUN with the name `/vol/vol_DominoDB/db_SdLun_clone1`, this could be mapped to any other or to the same server. Use the **snapdrive storage connect** command on your Lotus Domino server on which you want to map the database and transactional log storage (see Example 10-9).

- ▶ Storage that includes LVM entities has special requirements. To use the **snapdrive storage connect** command to connect LVM entities, you must create the storage so that each entity in the storage hierarchy has exactly one instance of the next entity. For example, you can use the **snapdrive storage connect** command to connect a storage hierarchy that has one volume group (vg1) with one host volume (hostvol1) and one file system (fs1). However, you cannot use the **snapdrive storage connect** command to connect a hierarchy that has one volume group (vg1) with two host volumes (hostvol1 and hostvol2) and two file systems (fs1 and fs2).
- ▶ The **snapdrive storage connect** command connects a file system created directly on a LUN only when the underlying LUN is partitioned.

Note: For the complete guidelines for connecting Snapshot copies with SnapDrive, refer to *IBM System Storage N series SnapDrive for UNIX Installation and Administration Guide*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>

Example 10-9 Connect to the split clone

```
[root@domino1 ~]# snapdrive storage connect -fs /notesdata/db_0 -hostvol  
db_SdDg/db_SdHv -lun itsotuc3:/vol/vol_DominoDB/db_SdLun_clone1 -nopersist  
  
mapping lun(s) ... done  
discovering lun(s) ... done  
  
LUN itsotuc3:/vol/vol_DominoDB/db_SdLun_clone1 connected  
- device filename(s): /dev/sdd  
Importing db_SdDg  
Connected fs /notesdata/db_0  
[root@domino1 ~]#
```

Note: You will not be able to map an additional host volume within an existing name when using LVM. An error like this might occur:

0001-377 Command error: Disk group name db_SdDg is already in use or conflicts with another entity.

In this case, first disconnect the volume group with the following command:

```
[root@domino1 notesdata]# snapdrive storage disconnect -vg db_SdDg -full
```

Using disk

Working with Snapshots allows you to make very fast restores and save space because only the changed blocks are saved. However, Snapshot will not help you when a physical failure of your disks or aggregate occurs. In this situation, you have to back up your data to other physical disks. There are two ways to accomplish this task:

- ▶ SnapMirror
- ▶ Volume copy

Note: The Data ONTAP SnapVault feature is a disk based backup system, and because it uses qtrees as backup storage, it is not supported for volumes containing Data ONTAP LUNs, so we do not discuss it in this book. For more information about SnapVault, see *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

SnapMirror

The Data ONTAP SnapMirror feature enables an administrator to mirror Snapshot images either asynchronously or synchronously:

- ▶ When mirroring asynchronously, SnapMirror replicates Snapshot images from a source volume to a partner destination volume or qtree, thus replicating source object data on destination objects at regular intervals.
- ▶ When mirroring synchronously, SnapMirror replicates Snapshot images from a source volume to a partner destination volume at the same time it is written to the source volume.

Additionally, you can configure a synchronous SnapMirror replication to lag behind the source volume by a user-defined number of write operations or milliseconds. This option is useful if you are balancing the need for synchronous mirroring with the performance benefit of asynchronous mirroring.

You can access the information about the destination volume to:

- ▶ Provide users quick access to mirrored data in the event of a disaster that makes the source volume or qtree unavailable
- ▶ Update the source to recover from disaster or user error
- ▶ Archive the data to tape
- ▶ Back up or distribute the data to remote sites

SnapMirror is usually used to replicate Snapshot images from one IBM System Storage N series storage server to another IBM System Storage N series storage server. These are two common configurations:

1. Source to destination method

The SnapMirror destination volumes are read-only objects, usually on a separate system, to which the source volumes are replicated (see Figure 10-1 on page 393). The destination volumes are normally accessed by users only when a disaster takes down the source volumes and the administrator uses SnapMirror commands to make the replicated data at the destination accessible and writable.




Figure 10-1 SnapMirror source to destination method

2. Source to destination to tape method

A common variation to the basic SnapMirror backup deployment adds a tape backup of the destination volume (see Figure 10-2). By running a tape backup off the SnapMirror destination volume, you do not subject the heavily user-accessed source volume to the performance degradation, system unavailability, and complexity of a direct tape backup.




Figure 10-2 SnapMirror source to destination to tape method

To set up a asynchronous SnapMirror configuration on the Lotus Domino database and transactional log volume, follow these steps:

1. Make sure you have purchased a SnapMirror license for both the source and destination systems. Install the license on both the source and destination IBM System Storage N series storage system (see Example 10-10).

Example 10-10 Install the SnapMirror license

```
itsotuc4> license add agy2csd
A snapmirror site license has been installed.
    snapmirror enabled.
itsotuc4>
```

2. On the source system console, use the **options snapmirror.access** command to specify the host names of systems that are allowed to copy data directly from the source system (see Example 10-11).

Example 10-11 Set up security for SnapMirror operations

```
itsotuc3> options snapmirror.access host=itsotuc4
itsotuc3>
```

- Through the Data ONTAP FilerView, create or edit the `/etc/snapmirror.conf` file on the destination system to specify the volumes to be copied and the schedule (*minute hour day_of_month day_of_week* or *async*) on which the destination is updated. In Example 10-12, we create a asynchronous mirror of the Lotus Domino database and transactional log volume (both from the source N series storage server itsotuc3) to the IBM System Storage N series storage server itsotuc4 and its volumes vol1 and vol2 (see Example 10-14).

Note: For SnapMirror volume replication, you must create a restricted volume to be used as the destination volume.

Example 10-12 /etc/snapmirror.conf

```
itsotuc3:vol_DominoDB  itsotuc4:vol1 - async
itsotuc3:vol_DominoLog itsotuc4:vol2 - async
```

- Start the SnapMirror service on the source and destination IBM System Storage N series storage servers, as shown in Example 10-13.

Example 10-13 Starting the SnapMirror service

```
itsotuc3> snapmirror start
itsotuc3>
```

Note: For optimum SnapMirror volume replication performance, make sure that the SnapMirror source volume and destination volume contain disks of the same size, organized in the same RAID configuration.

- In the destination system console, use the **snapmirror initialize** command to create an initial complete (baseline) copy of the source on the destination and start the mirroring process.

In Example 10-14, we initialize the mirror from Lotus Domino transactional log volume vol_DominoLog on itsotuc3 to volume vol2 on the IBM System Storage N series storage server itsotuc4.

Example 10-14 Initialize the SnapMirror

```
itsotuc4> snapmirror initialize -S itsotuc3:vol_DominoLog itsotuc4:vol2
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.
itsotuc4>
```

- Verify the mirroring status with the **snapmirror status** command. The status In-sync means a successful replication status. Until now, the Transferring status was about the mirror initializing (see Example 10-15).

Example 10-15 Checking the mirror status

```
itsotuc3> snapmirror status
Snapmirror is on.
Source           Destination    State  Lag  Status
itsotuc3:vol_DominoDB  itsotuc4:vol1 Source  -    Transferring
itsotuc3:vol_DominoLog itsotuc4:vol2 Source  -    In-sync
itsotuc3>
```

The restricted volumes vol1 and vol2 on our destination IBM System Storage N series storage server itsotuc4 contains our Lotus Domino database and transactional log LUNs (see Example 10-16).

Example 10-16 lun show command at the mirror destination

```
itsotuc4> lun show
          /vol/vol1/db_SdLun          146.0g (156786229248) (r/o, online)
          /vol/vol2/log_SdLun         5.0g (5377097728) (r/o, online)
itsotuc4>
```

This volumes could be mounted on a different host for read-only backups, or, if a disaster takes the source volume down, you are able to mount the mirrored volume in a writable mode.

You can use the **snapmirror quiesce** and **snapmirror break** commands to convert a SnapMirror destination, with read-only status, to a writable volume. You might want to convert a destination to a writable volume to perform one of the following tasks:

- ▶ **Data migration:** You want to move your users working data from one volume (your current source volume) to another volume (your current destination volume) and make that data accessible and writable in its new location.
- ▶ **Disaster recovery:** In case your source volume or storage Server is suddenly unavailable, you want your current destination volume to serve as the substitute for your users retrieval and input source.
- ▶ **Application testing:** You want to make your current destination volume writable to test a new application on a mirrored replication of your current data rather than risk corruption of original data on the source volume.

Converting the destination to a writable volume lets you use data on the destination for these situations or in any other situation in which the original source is unavailable.

Example 10-17 SnapMirror break

```
itsotuc4> snapmirror quiesce vol2
snapmirror quiesce: in progress
This can be a long-running operation. Use Control - C (^C) to interrupt.
snapmirror quiesce: vol2 : Successfully quiesced
itsotuc4> snapmirror break vol2
snapmirror break: Destination vol2 is now writable.
Volume size is being retained for potential snapmirror resync. If you would lik
e to grow the volume and do not expect to resync, set vol option fs_size_fixed t
o off.
itsotuc4>
```

Note: For more information about SnapMirror and SnapMirror commands, see *IBM System Storage N series Data ONTAP 7.2 Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

Another method to access data at the SnapMirror location is to use the Data ONTAP SnapMirror feature in combination with the FlexClone volume.

Volume copy

You can use the **vol copy** command to copy LUNs; however, this requires that applications accessing the LUNs are quiesced and offline prior to the copy operation.

The **vol copy** command enables you to copy data from one WAFL volume to another, either within the same storage system or to a different storage system. The result of the **vol copy** command is a restricted volume containing the same data that was on the source storage system at the time you initiate the copy operation.

The following are requirements for copying a volume:

- ▶ The source and destination volumes must be of the same type: either both traditional or both flexible volumes.
- ▶ The capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- ▶ The source and destination on the N series must have a trust relationship with each other.
- ▶ The destination volume must exist, and must not be the root volume.
- ▶ Remote Shell access must be enabled.
- ▶ The destination volume must not contain data that you want to preserve.

Command usage:

```
vol copy start [ -S | -s Snapshot ] source destination
```

If the **-S** flag is used, the command copies all Snapshots in the source volume to the destination volume. To specify a particular Snapshot to copy, use the **-s** flag followed by the name of the Snapshot. If neither the **-S** nor **-s** flag is used in the command, the storage system automatically creates a distinctively-named Snapshot at the time the **vol copy start** command is executed and copies only that Snapshot to the destination volume.

Note: The source and destination volume must be the short volume name. If the source or destination is on a different N series, use **filer:** in front of the volume name.

Correct command example:

```
vol copy start itsotuc3:vol_DominoLog itsotuc4:vol1
```

Incorrect command example:

```
vol copy start itsotuc3:/vol/vol_DominoLog itsotuc4:/vol/vol1
```

Example 10-18 shows a copy process of the volume **vol_DominoLog** to another N series.

Example 10-18 Volume copy

```
itsotuc3> vol copy start -S vol_DominoLog itsotuc4:vol1
This can be a long-running operation. Use Control - C (^C) to interrupt.
Copy Volume: vol_DominoLog on machine 9.11.218.237 to volume: vol1
VOLCOPY: Starting on volume 1.
10:27:18 MST : vol copy restore 0 : begun, 249 MB to be copied.
[...]
10:27:28 MST : vol copy restore 0 : 100% done, 249 MB copied.
itsotuc3>
```

Using tape

You use tape backup and recovery to create tape archives or retrieve data from tape archives. Tape archives complement online backup methods, and you should plan to do tape backups regardless of the online backup and recovery methods you use.

With Data ONTAP, you can use the **dump** command for backup to tape. Additionally, you can use a third-party backup tool, such as IBM Tivoli Storage Manager (ITSM) and the Data ONTAP Network Data Management Protocol (NDMP) implementation.

Data ONTAP NDMP backup and recovery operations use the dump and restore services for data backup to tape and data restoration from tape. However, accessing this data protection services through backup applications that support NDMP offers a number of advantages:

- ▶ NDMP backup applications provide sophisticated scheduling of data protection operations across multiple storage systems. They also provide media management and tape inventory management services to eliminate or minimize manual tape handling during data protection operations.
- ▶ NDMP supports multiple topology configurations, allowing efficient sharing of secondary storage (tape library) resources through the use of three-way network data connections. NDMP backup applications typically provide user-friendly interfaces that simplify the management of data protection services.

Note: For more information about how to use NDMP with N series and ITSM, see the IBM Redbooks publication *Using the IBM System Storage N series with IBM Tivoli Storage Manager*, found at:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247243.pdf>

We recommend the usage of the SnapMirror source to destination to tape method for tape backup; however, the described method will work without SnapMirror as well.

Figure 10-3 shows a common SnapMirror source to destination to tape configuration.




Figure 10-3 SnapMirror source to destination to tape method

The **dump** command uses a Snapshot copy of your data for the backup, so you do not need to take the storage system or volume offline before initiating the backup.

The **dump** command names each Snapshot copy it creates `Snapshot_for_backup.n`, where:

- ▶ `n` is an integer starting at 0.
- ▶ Each time the **dump** command creates a Snapshot copy, it increments the integer by 1.
- ▶ The storage system resets the integer to 0 when it is rebooted.

When Data ONTAP executes multiple **dump** commands simultaneously, the **dump** commands create multiple Snapshot copies. For example, if Data ONTAP is running two **dump** commands simultaneously, you find these Snapshot copies in the volumes from which data is being backed up: Snapshot_for_backup.0 and Snapshot_for_backup.1.

Note: When you are backing up from a Snapshot copy, **dump** does not create an additional Snapshot copy. Because we are using Snapshots with the Data ONTAP scheduler and SnapDrive for Linux and Unix, we will take existing Snapshots for creating the dump.

Example 10-19 shows a simple usage of the **dump** command. The snap1 Snapshot was created on the SnapMirror source site itsotuc3 through SnapDrive on the Lotus Domino server. Now this Snapshot is dumped to tape on the SnapMirror destination site itsotuc4.

Example 10-19 Dump a mirrored Snapshot to a local tape

```
itsotuc4> dump 0f rst1m /vol/vol2/.Snapshot/snap1
DUMP: Using Full Volume Dump
DUMP: Dumping tape file 1 on rst1m
DUMP: Date of this level 0 dump: Thu Jun 14 21:32:02 2007.
DUMP: Date of last level 0 dump: the epoch.
DUMP: Dumping /vol/vol2/.Snapshot/snap1 to rst1m
DUMP: mapping (Pass I)[regular files]
DUMP: mapping (Pass II)[directories]
DUMP: estimated 127287 KB.
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
DUMP: dumping (Pass V) [ACLs]
DUMP: 127373 KB
DUMP: DUMP IS DONE
itsotuc4>
```

Where:

- ▶ 0f are the options of the **dump** command. The 0 means a full backup, while f specifies the tape file.
- ▶ rst1m is the argument of the f option. It is the tape device /dev/rst1m.
- ▶ /vol/vol2/.Snapshot/snap1 is the dump path.

This command backs up to tape all files and directories in the /vol/vol2/.Snapshot/snap1 volume. This includes our Lotus Domino LUNs, in this case the transactional log file LUN /vol/vol2/log_SdLun.

Note: For more information about the **dump** and **restore** commands, see *IBM System Storage N series Data ONTAP 7.2 Data Protection Tape Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

Recommended backup configuration

The N series Snapshot capability is the first protection step in our recommended backup configuration for the Lotus Domino server database and transactional logs.

A backup solution for Lotus Domino should meet the following requirements:

- ▶ Provide a fast backup with less impact on performance.
- ▶ Provide several point-in-time backups with minimal loss of data in case of a database failure or user fault.
- ▶ Provide fast recovery of single user databases.
- ▶ Have no downtime for backup and reduced downtime for recovery.
- ▶ Have a short downtime following a disaster on the storage server.
- ▶ Provide backup to tape for archiving.

We use a combination of the features of SnapDrive for Linux and UNIX and SnapMirror for our configuration. The process is described in these steps and in Figure 10-4 on page 400:

- ▶ Using Linux as an example, you would create Snapshots from the Lotus Domino database and transaction log volume through the Linux cron daemon and SnapDrive every hour.
 - A minimum amount of additional space is required (only for changed data blocks).
 - You can recover your databases at an exact point-in-time.
 - Recovery is very fast (the data relies on a disk).
- ▶ Using the N series servers itsotuc3 and itsotuc4, mirror the volumes from the IBM System Storage N series storage server itsotuc3 to a different IBM System Storage N series storage server itsotuc4.
 - In a disaster, you can use the mirrored data on the external storage as active data.
 - Tests, changes, or backups can be done on the mirrored side and off loads the production system.
- ▶ Mount the Lotus Domino database and transactional log Snapshots once a day at the backup server and take a full backup of all the files with a third-party backup tool, such as IBM Tivoli Storage Manager (ITSM).
 - Back up your Lotus Domino database and transactional log files as usual, but without a performance impact on your production system.

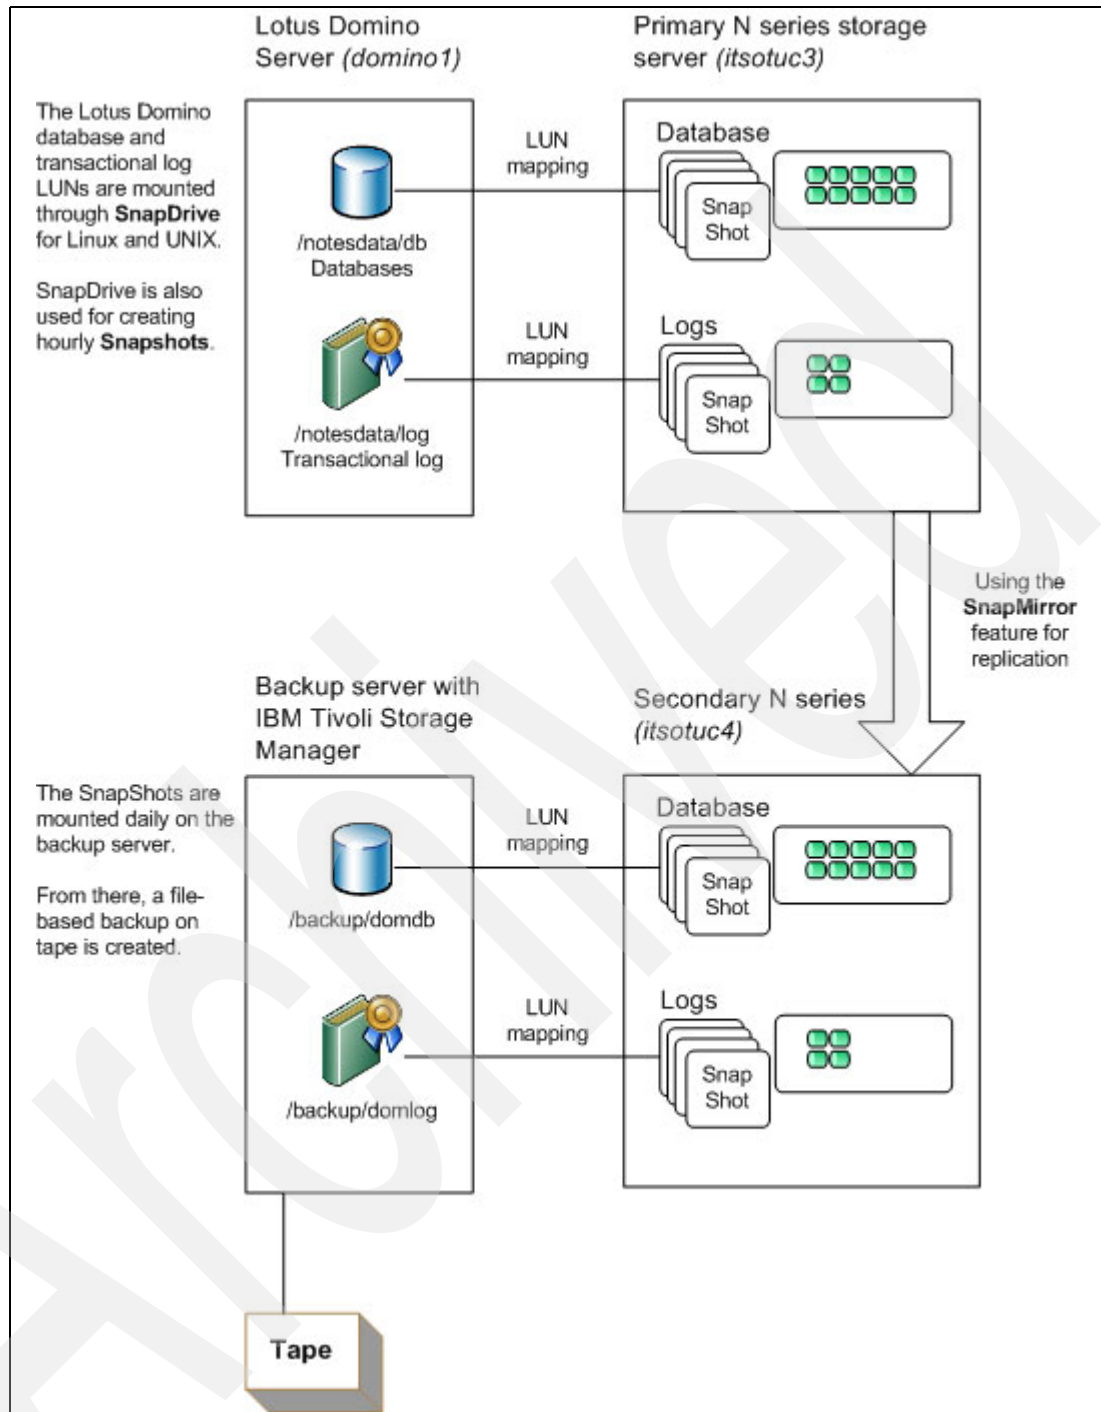


Figure 10-4 Recommended backup configuration

10.1.2 Restoring from Snapshot

If you need to restore a single user database, you can connect to a Snapshot copy and quickly restore the file you want by using the **cp** command. If you need to restore the whole LUN because of a failure in a change process, such as a template update, use the **snapdrive snap restore** command.

Connecting to the Snapshot

SnapDrive for UNIX lets you connect to a Snapshot copy from a different location on a host. This new location can be on the host where you took the Snapshot copy (the originating host) or on a different host (the non-originating host).

Being able to set up the Snapshot copies in a new location means you can back up a Snapshot copy to another medium, perform maintenance on a disk group, or simply access the Snapshot copy data without disrupting the original copy of the data.

Example 10-20 Connect to the Snapshot with SnapDrive

```
[root@domino1 /]# snapdrive snap connect -vg db_SdDg -autorename -snapname
itsotuc3:/vol/vol_DominoDB:snap2

connecting db_SdDg:
  LUN copy db_SdLun_0 ... created
    (original: itsotuc3:/vol/vol_DominoDB/db_SdLun)
  LUN copy db-1_SdLun_0 ... created
    (original: itsotuc3:/vol/vol_DominoDB/db-1_SdLun)

  mapping new lun(s) ... done
  discovering new lun(s) ... done
  Importing db-1_SdDg_0
Successfully connected to Snapshot itsotuc3:/vol/vol_DominoDB:snap2
  disk group db-1_SdDg_0 containing host volumes
    db_SdHv_0 (filesystem: /notesdata/db_0)
[root@domino1 /]#
```

Where:

- ▶ -vg: The volume group name to which you wish to connect.
- ▶ -autorename: This option tells SnapDrive for Linux and UNIX to generate a new, unused name for the destination entity if the default name is in use.
- ▶ -snapname: Long name of the Snapshot.

Restoring from the Snapshot

After you connected to the Snapshot, use your favorite copy command for copying the files from your Snapshot copy mount point to the active file system.

Example 10-21 Restoring from the Snapshot using cp

```
[root@domino1 /]# cp /notedata/db_0/sschaf.nsf /notesdata/db/sschaf.nsf
[root@domino1 /]#
```

If you need to restore the whole file system, use the **snapdrive snap restore** command. You do not need to make a connection with **snapdrive snap connect** first.

The **snapdrive snap restore file_spec** command restores data from the Snapshot copy you specify on the command line to the storage system. This operation replaces the contents of the *file_spec* that you specified on the **snap restore** command line with the contents of the *file_spec* arguments found in the specified Snapshot copy.

You can also restore Snapshot copies for non-existent *file_spec* arguments. This happens when the value you specify no longer exists on the host, but existed when you took the Snapshot copy. For example, it might be a file system that you have now unmounted or a disk group that you have removed.

Normally, you restore Snapshot copies from the host where you took the Snapshot copies (in other words, the originating host). You can also restore Snapshot copies using a different, or non-originating, host. For more details, see:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001600&aid=1>.

Note: When you restore Snapshot copies for volume groups, or for host volumes and file systems that are created on them, SnapDrive for Linux and UNIX restores the whole volume group. If you specify part of a volume group, SnapDrive still restores the entire volume group. An error occurs if you enter only a subset of the host volumes or file systems in each volume group on the command line. You can include the `-force` option to override this error; however, SnapDrive then restores the entire disk group.

In Example 10-22, we restore Snapshot `snap1` of the Lotus Domino database and transactional log storage.

Example 10-22 Restoring the Lotus Domino database and transactional log storage from Snapshot

```
[root@domino1 /]# snapdrive snap restore -vg db_SdDg log_SdDg -snapname snap1
Starting to restore /dev/mapper/db_SdDg, /dev/mapper/log_SdDg
WARNING: This can take several minutes.
DO NOT CONTROL-C!
If snap restore is interrupted, the filespecs
being restored may have inconsistent or corrupted
data.

For detailed progress information, see the log file /var/log/sd-recovery.log
Importing db_SdDg, log_SdDg
Successfully restored Snapshot snap1 on 2 filer volumes:
itsotuc3:/vol/vol_DominoLog
itsotuc3:/vol/vol_DominoDB

disk group db_SdDg containing host volumes
db_SdHv (filesystem: /notesdata/db)
disk group log_SdDg containing host volumes
log_SdHv (filesystem: /notesdata/log)
[root@domino1 /]#
```

Where:

- `-vg`: The volume group name of the Snapshot you wish to restore. This option is stackable.
- `-snapname`: Name of the Snapshot. You can use the short (for example: `snap1`) or long (for example: `itsotuc3:/vol/vol_DominoLog:snap1`) name. If the volume entity on your host no longer exists, you must specify the long name.

SnapRestore

SnapRestore enables you to quickly revert a local volume or file on a storage system to the state it was in when a particular Snapshot copy was taken. In most cases, reverting a file or volume is much faster than restoring files from tape or copying files from a Snapshot copy to the active file system.

Note: SnapRestore must be licensed on your N series.

You use SnapRestore to recover from data corruption. If a primary storage system application corrupts data files in a volume, you can revert the volume or specified files in the volume to a Snapshot copy taken before the data corruption.

SnapRestore carries out Snapshot restoration more quickly, using less disk space, than an administrator can achieve by manually copying volumes, directories, or large files to be restored from the Snapshot system to the active file system. A large volume directory restore can be carried out in a few seconds using the SnapRestore feature.

SnapRestore can restore large volumes or files even if space limitations would prevent restoration by copying from a Snapshot copy.

For our Lotus Domino server, which is using LUNs to store its database and transactional log files, we can use SnapRestore for instant recovery of our created Snapshots.

Note: As a best practice, use this feature with data stored directly on the volume, not in LUNs. SnapRestore will quickly restore data on a volume. Working with LUNs, you can use the SnapDrive **snapdrive snap restore** command for recovery of the whole LUN. SnapRestore does not enable you to restore several files in the LUN.

If you use Lotus Domino or Microsoft Exchange with a CIFS connection, this option is highly recommended.

Example 10-23 shows the restore of our Lotus Domino database volume from the Snapshot snap1:

```
snap restore -t vol -s Snapshot_name volume_name
```

Where:

- ▶ -t vol specifies the volume name to revert.
- ▶ -s Snapshot_name specifies the name of the Snapshot copy from which to revert the data. You can enter only one Snapshot copy name. In our case, it is snap1.
- ▶ volume_name is the name of the volume to be reverted. Enter the name only, not the complete path. You can enter only one volume name. In our case, it is /vol/vol_DominoDB.

Example 10-23 Restore a volume with SnapRestore

```
itsotuc3> snap restore -t vol -s snap1 /vol/vol_DominoDB
system> WARNING! This will restore a volume from a Snapshot into
the active filesystem. If the volume already exists in the active
filesystem, it will be overwritten with the contents from the
Snapshot.
Are you sure you want to do this? y
You have selected file /vol/vol_DominoDB, Snapshot snap1
Proceed with restore? y
itsotuc3>
```

Note: The volume to be reverted must not be a mirror used for data replication. If you revert to a Snapshot copy taken before a SnapMirror Snapshot copy, Data ONTAP can no longer perform an incremental update of the mirror; it must re-create the baseline version of the mirror.

You can even use **snap restore** to revert a single file to a selected Snapshot copy. This is practical when the file is so large that you cannot copy the previous file version from the Snapshot copy to the active file system. Remember, this is not essential for files they reside in LUNs. For more information about the SnapRestore feature, see the *IBM System Storage N series Data ONTAP Data Protection Online Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

Using tape

You restore data from tape under the following circumstances:

- ▶ Files were deleted from disk but backed up to tape.
For example, if you inadvertently delete a file and want to restore the file, you can recover the file from tape.
- ▶ Files are corrupted.
If some files are corrupted, you can restore the subtree containing the files.
- ▶ No disk slots are available for expansion. If the storage system runs out of storage space, you can do the following tasks:
 - Back up the entire storage system.
 - Replace the current disks with disks of greater capacity.
 - Set up volumes on new disks.
 - Restore the storage system from tapes.
- ▶ The entire storage system is damaged and unusable. Contact technical support to determine whether you can repair the storage system and restore data from tape.

You do not need to restore a deleted file from tape if you can recover the deleted file from a Snapshot copy on the storage system. If the file is in a Snapshot copy, copying the file to the active file system is faster than recovering the file from tape.

The **restore** command enables you to recover all the information that you backed up with the **dump** command so that it is the same as when you did the backup (see Example 10-19 on page 398). The restore command syntax is as follows:

```
restore options [arguments] [files ...]
```

The options flag can be one restore type with modifiers. We use the options **r** and **f**:

- ▶ Option **r** specifies restoring all the files in the dump tape.
- ▶ Option **f** specifies that there is a named device. Its argument is **rst1m**.

Example 10-24 Restore from tape


```
itsotuc4> restore rf rst1m
RESTORE: tape file #1 on device rst1m
RESTORE: Thu Jun 14 19:04:57 2007: Begin level 0 restore
RESTORE: Thu Jun 14 19:04:57 2007: Reading directories from the backup
RESTORE: Thu Jun 14 19:04:57 2007: Creating files and directories
RESTORE: Thu Jun 14 19:04:57 2007: Writing data to files
RESTORE: RESTORE IS DONE
itsotuc4>
```

Note: For more information about the **dump** and **restore** commands, see the *IBM System Storage N series Data ONTAP 7.2 Data Protection Tape Backup and Recovery Guide*, found at:

<http://www.ibm.com/servers/storage/support/nas/dataontap/>

Note: As a best practice, and in order to support up to the minute restore capabilities, the archive logging feature and usage of a third-party backup tool such as IBM Tivoli Storage Manager should be used with the Lotus Domino server.

Archived



Integrating Lotus Domino for Windows into the IBM System Storage N series

This appendix describes the installation of IBM Lotus Domino server on the Microsoft Windows operating system with the IBM N series storage system.

Introduction

The tight integration of SnapDrive with Windows and IBM System Storage N series storage systems offers the ease of storage management and high data availability required for Lotus Domino environments. An IBM System Storage N series storage system in conjunction with SnapDrive offers an enterprise-class solution for Lotus Domino. This chapter describes the steps necessary to integrate an IBM System Storage N series storage system with an IBM Lotus Domino server for Windows 2003 using SnapDrive over an iSCSI connection to the storage system.

The main topics covered in this appendix are:

- ▶ Preparing the IBM System Storage N series storage system for Lotus Domino
- ▶ Configuring Windows server and installing Lotus Domino software
- ▶ Migrating Lotus Domino server from local disks to LUNs on an IBM System Storage N series storage system

Configuration

In this scenario, we use Lotus Domino V7.0.2 enterprise server for Windows 2003. The N series supports Lotus Domino V6.5 and above on Microsoft Windows, UNIX, and Linux. The N series also supports Domino partitioning and Domino clustering for high availability and disaster recovery.

Ensure that the server meets the minimum requirements for running Lotus Domino and that the server hardware is on the Windows Hardware Compatibility List (HCL). For customers who plan to install Lotus Domino on an IBM System Storage N series storage system cluster, Lotus Domino clustering must be implemented for failover. In a cluster implementation, separate volumes are required for each Domino data directory. A list of supported configurations can be found at the following IBM Web site:

<http://www-1.ibm.com/support/docview.wss?rs=0&uid=swg27005236>

Microsoft iSCSI software initiator

For companies that do not have an FCP infrastructure in place or for those that want to access the storage using their Ethernet infrastructure and knowledge, iSCSI can be used as the access protocol for communication between the IBM Lotus Domino Server and IBM System Storage N series storage system.

The recommended configuration when using iSCSI protocol is to use iSCSI Initiator adapters. This adapters off load the largest amount of processing from the server's CPU, improving performance.

In case there are no iSCSI adapters on your planned environment, the iSCSI Initiator software can be used to provide the same connectivity to the IBM System Storage N series storage system. The use of multipaths is also recommended when using hardware or software based iSCSI solution, as shown on Figure A-1 on page 409.




Figure A-1 Multipathing configuration for Lotus Domino Server using iSCSI

In Figure A-1, there are two interfaces on the server (either iSCSI hardware based or Gigabit Ethernet cards) that connect to two different LAN switches. It is important, for performance and reliability reasons, that these LAN segments and switches are other than the public ones. The IBM System Storage N series storage system will have two of its adapters also connected to both switches.

Assuming that all the infrastructure is already in place and working and that you are not using HBA iSCSI-enabled adapters, Microsoft iSCSI Software Initiator has to be installed on the server. After installing and configuring it, SnapDrive should be installed and configured as well so that the LUNs can be created.

Note: Despite the fact that the LUNs can be created from the IBM System Storage N series storage system, the recommended procedure is to create the LUNs from the Lotus Domino Server using SnapDrive.

SnapDrive software

The IBM System Storage N series SnapDrive feature provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file, down through the volume manager, to the storage-system-side logical unit numbers (LUNs) providing the actual repository. In addition, it simplifies the backup of data and helps you decrease the recovery time.

SnapDrive provides a layer of abstraction between an application running on the host operating system and the underlying IBM System Storage N series storage systems (see Figure A-2). Applications that are running on a server with SnapDrive use virtual disks (or LUNs) on IBM System Storage N series storage systems, as through they were locally connected drives or mount points. This allows applications that require locally attached storage, such as IBM Lotus Domino and Microsoft Exchange, to benefit from the N series technologies, including Snapshot, flexible volumes, cloning, and space management technologies.

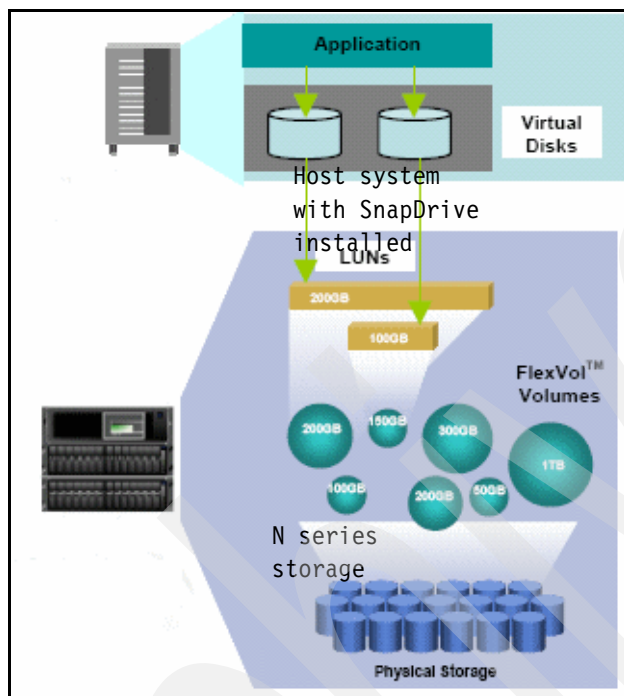


Figure A-2 Example of a typical SnapDrive deployment

SnapDrive includes all the necessary drivers and software to manage interfaces, protocols, storage, and Snapshot copies. Snapshot copies are nondisruptive to applications and functions on execution. Snapshot backups can also be mirrored across LAN or WAN links for centralized archiving and disaster recovery.

Benefits of SnapDrive

Most of today's enterprises use business-critical applications such as IBM Lotus Domino and their storage management team faces number of challenges. They must:

- ▶ Support new business initiatives with a minimal increase in operating budget.
- ▶ Protect data from corruption, disaster, and attacks.
- ▶ Back up data without any performance degradation, quickly and consistently, without any errors.

SnapDrive addresses these problems by providing simplified and intuitive storage management and data protection from a host/server perspective. The following list highlights some of the important benefits of SnapDrive:

- ▶ Allows hosts and applications administrators to quickly create virtual disks with a dynamic pool of storage that can be reallocated, scaled, and enlarged in real time, even while system are accessing data.

- ▶ Dynamic on-the-fly file system expansion; new disks are usable within seconds.
- ▶ Snapshot copies provide rapid backup and recovery capability with minimal resource and capacity requirements.
- ▶ Supports multipath technology for high performance. (Check the compatibility matrix first. At the time the book was written, SnapDrive for Linux did not support multipath technology.)
- ▶ Enables connections to existing Snapshot copies from the original host or different host.
- ▶ Independent of underlying storage access media and protocol; SnapDrive supports FCP, iSCSI, and NFS as the transport protocols (NFS supports only Snapshot management)
- ▶ Robust and easy-to-use data and storage management feature and software.

SnapDrive requirements

IBM System Storage N series SnapDrive is a licensed feature and is available by contacting IBM Support.

There are some general requirements for SnapDrive:

- ▶ Host operating system and appropriate patches
- ▶ Host file systems
- ▶ IP access between the host and storage system
- ▶ Storage system licenses
- ▶ FCP Host Utilities or iSCSI Host Utilities required software
- ▶ For security reasons, we recommend a separate user account on the IBM System Storage N series storage server. See 7.2, “Role-based access control of the IBM System Storage N series” on page 242 for more details.

The operating system requirements can be found in the *IBM System Storage N series SnapDrive for Windows 4.2.1 Release Notes*, found at:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7001650&aid=1>

For more information about SnapDrive, please refer to the *IBM System Storage N series SnapDrive 4.2 for Windows Installation and Administration Guide*, GC26-7880.

Network

For I/O traffic from the Windows server hosting the Domino binaries and the storage system housing the Lotus Domino databases and transaction logs, a private LAN or VLAN is mandatory. IBM will not support Domino I/O traffic over a public network. This can be accomplished using a crossover cable in a simple networking topology or network switch, with VLANs dedicated to Lotus Domino and auxiliary traffic for antivirus, and so on, and IBM System Storage N series storage system workloads. This can also be thought of as an IP SAN, since the Domino server will be connected to the storage system over a private network, and no other machines will have access to LUNs on the storage system. A gigabit IP SAN between the Lotus Domino server machine and the IBM System Storage N series storage system is required for the following reasons:

- ▶ Eliminates issues of contention or latency.
- ▶ Minimizes security issues and hacking threats.
- ▶ Provides storage capacity scalability.

We strongly recommend that Gigabit Ethernet be used to support your Lotus Domino database environment on the IBM System Storage N series storage system.

IBM System Storage N series storage system

An IBM System Storage N series storage system provides a virtual storage layer called *flexible volumes*, hereafter known as flexible volume or volumes. A flexible volume is created within an aggregate and is loosely coupled with its containing aggregate. The flexible volume provides greater performance than traditional volumes and can grow and shrink as needed by executing one simple command. Data ONTAP makes it easy to control the placement of related Domino file systems on a flexible volume (Figure A-3).

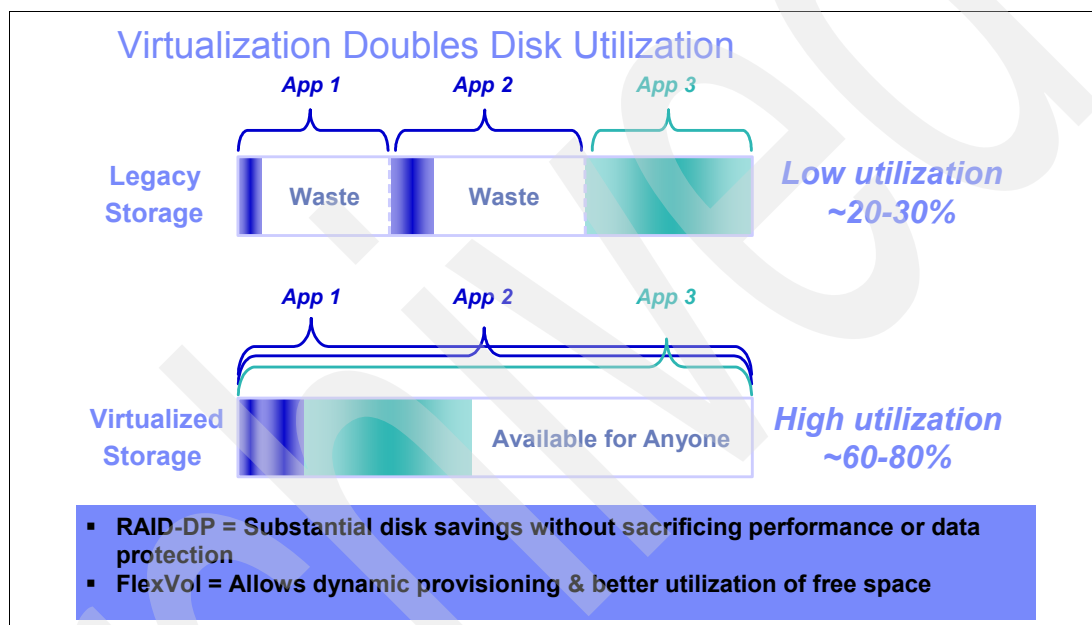


Figure A-3 FlexVol

A key to ease use and manageability of a Lotus Domino server housed on an IBM System Storage N series storage system is the notion that the entire Domino environment can be stored on one or more flexible volumes. This configuration requires minimal attention by Domino administrators and system administrators, allowing the storage system to manage the physical data storage, which ensures high performance and availability. However, there are a few storage system physical design considerations that need to be taken into consideration to ensure that these benefits are not compromised:

- The root volume should be its own volume. The root volume can be a traditional volume or a flexible volume. For extra reliability and flexibility, we recommend that you mirror the root volume. The root volume usually contains data that does not change much over time; a low change rate on the root volume would imply that it does not need to be backed up as often. And, should a data volume fail, having a separate, still functioning root volume will save valuable time in the recovery process.
- We strongly recommend that all the Lotus Domino files be stored on a volume on the IBM System Storage N series storage system that is *not* the root volume.
- We strongly recommend that a large aggregate be used whenever possible for performance reasons. Multiple flexible volumes required for a Domino environment can be created within a single large aggregate.

- ▶ Domino transaction log files should be kept on a different LUN from the LUN housing the Domino data directory. Also, Domino data and transaction logs should reside on separate flexible volumes. If database and transaction logs are housed in the LUNs on the same flexible volume, the recovery of the LUNs will return the logs to the same state they were in at the time the Snapshot copy was created. As a result, log data generated after the creation of the Snapshot copy will be lost, thus making roll-forward recovery for Domino impossible, since log file data needed would no longer be available.
- ▶ If NDMP is used to back up Domino residing on the IBM System Storage N series storage system, multiple backup tasks can be spawned. It is simpler and more efficient to create a volume, define your LUNs in the root of the volume, and let the storage system manage the storage.

FCP or iSCSI licenses must be activated on the storage system. The iSCSI and FCP licenses supplied with SnapDrive enable all the CIFS functionality necessary for using SnapDrive, including CIFS share creation. If you also want full-featured, direct CIFS access to a particular storage system, you must install a separate CIFS license on that storage system.

Using Lotus Domino with an IBM System Storage N series storage system

There are several advantages to storing Lotus Domino data and transaction logs on an IBM System Storage N series storage system. This section examines some of those advantages and explores some steps that must be taken to ensure the overall Domino performance.

Advantages of IBM System Storage N series storage systems

Running Lotus Domino with databases and transaction log files stored on an IBM System Storage N series storage system has several advantages:

- ▶ **Extremely Fast Backup:** Snapshot copies (see Figure A-4 on page 414) can be created in a matter of seconds, regardless of the size of the Domino database or the level of activity on the IBM System Storage N series storage system used. This reduces the Domino backup window from hours to seconds and allows Domino administrators to take frequent full backups without having to take the Domino server offline.

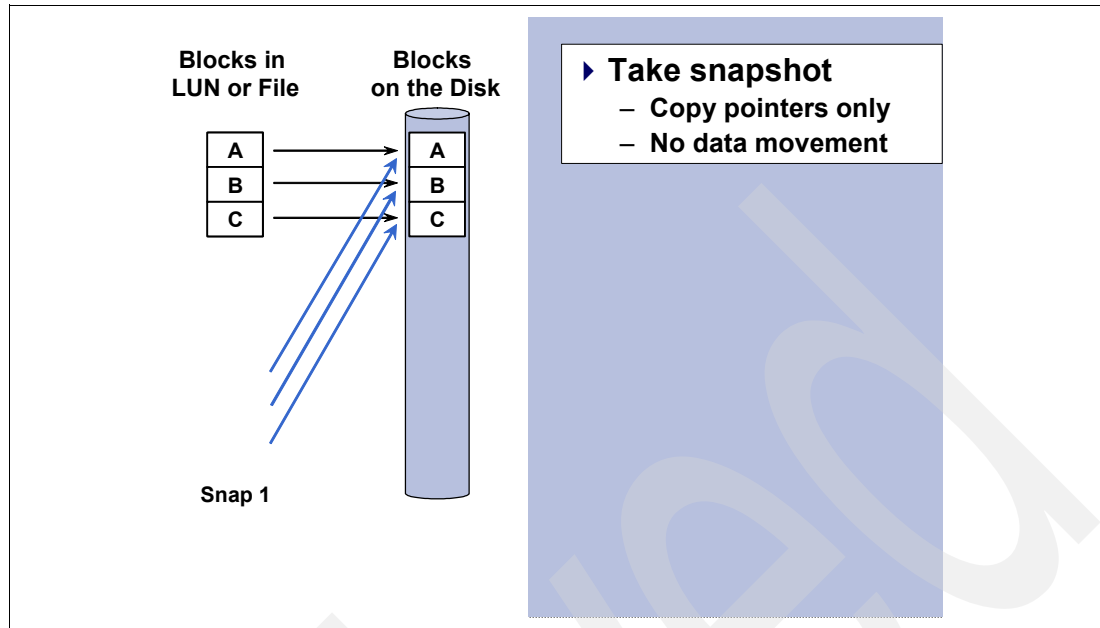


Figure A-4 Snapshot

- ▶ **Quick Recovery:** Using the Data ONTAP SnapRestore command, an entire Domino environment can be restored in a matter of seconds. Since there is no data copying involved, an incredible amount of time is saved as the file system is put back to the original state it was in at the time the Snapshot was created. Data ONTAP supports 255 Snapshot copies per WAFL volume. The ability to store a large number of low-impact, frequently created Snapshot copies brings the time needed to perform a roll-forward recovery operation down to minutes or seconds. In many circumstances, it allows the Domino administrator to restore the Domino server immediately without the need to restore from tape.
- ▶ **High Availability:** The need for 24x7 availability is fast becoming a reality for organizations of all sizes. Organizations cannot tolerate scheduled downtime or afford extended periods of slow system response caused by traditional Domino server backup methods. Snapshot technology that can create Domino server backups in a matter of seconds without bringing the server down can be used as complementary technology to ensure higher system uptime.
- ▶ **High Reliability:** The RAID architecture used for IBM System Storage N series storage systems is unique and provides greater reliability than many traditional RAID implementations. If a disk in an N series RAID group fails, it is reconstructed automatically without any user intervention. Additionally, N series supports the RAID-DP architecture. RAID-DP is considered approximately 4,000 times more reliable than traditional RAID.
- ▶ **No Impact on System Response Time during backup:** A Snapshot copy is simply a picture of the file system at a specific point-in-time. Therefore, creating a Domino server backup using Snapshot does not involve actual data movement (data I/O), so the backup process has virtually no performance impact on system response time.
- ▶ **Minimum Storage requirement:** Two Snapshot copies created in sequence differ from each other by the blocks added or changed in the time interval between their creation. This block-incremental behavior limits associated storage capacity consumption.
- ▶ **Load Balance:** Many of the tasks associated with load balancing between multiple Domino directories can be eliminated. Because of the high performance of the IBM System Storage N series storage system, only one volume needs to be defined for each directory used.

Lotus Domino transaction logging

Lotus Domino supports transaction logging. With this feature enabled, the system captures database changes and writes them to the transaction log. Then if a system or media failure occurs, you can use the transaction log and a previously created backup to recover your domino databases.

Transaction logging provides three main benefits:

- ▶ **No fix-up required:** In most cases, you no longer need to run the fix-up task to recover databases following a system failure, which results in faster server restarts. Fix-up must check every document in each database, while transaction log recovery applies or undoes only those transactions not written to disk at the time of the system failure.
- ▶ **Superior performance:** Transaction logging saves processing time because it allows Domino to defer database updates to disk during periods of high server activity. Transactions are recorded sequentially in the log files, which is much quicker than database updates to random, non-sequential parts of a disk. Because the transactions are already recorded, Domino can safely defer database updates until a period of low server activity.
- ▶ **Simplify the backup process:** Using transaction logging simplifies your daily backup procedure. You can use Snapshot to perform daily incremental backups of the transaction logs, rather than perform full database backups.

Lotus Domino's serialized log writes play to one of the strengths of the IBM System Storage N series storage system. WAFL is very efficient at writing data to a storage system volume. The storage system's cache effectively groups incoming data and writes it out to disk in an efficient manner.

Lotus Domino supports three types of transaction logging

Figure A-5 shows the default circular logging configuration in the Lotus Domino Administrator. You will find it at Configuration Server Current Server Document Transactional Logging.

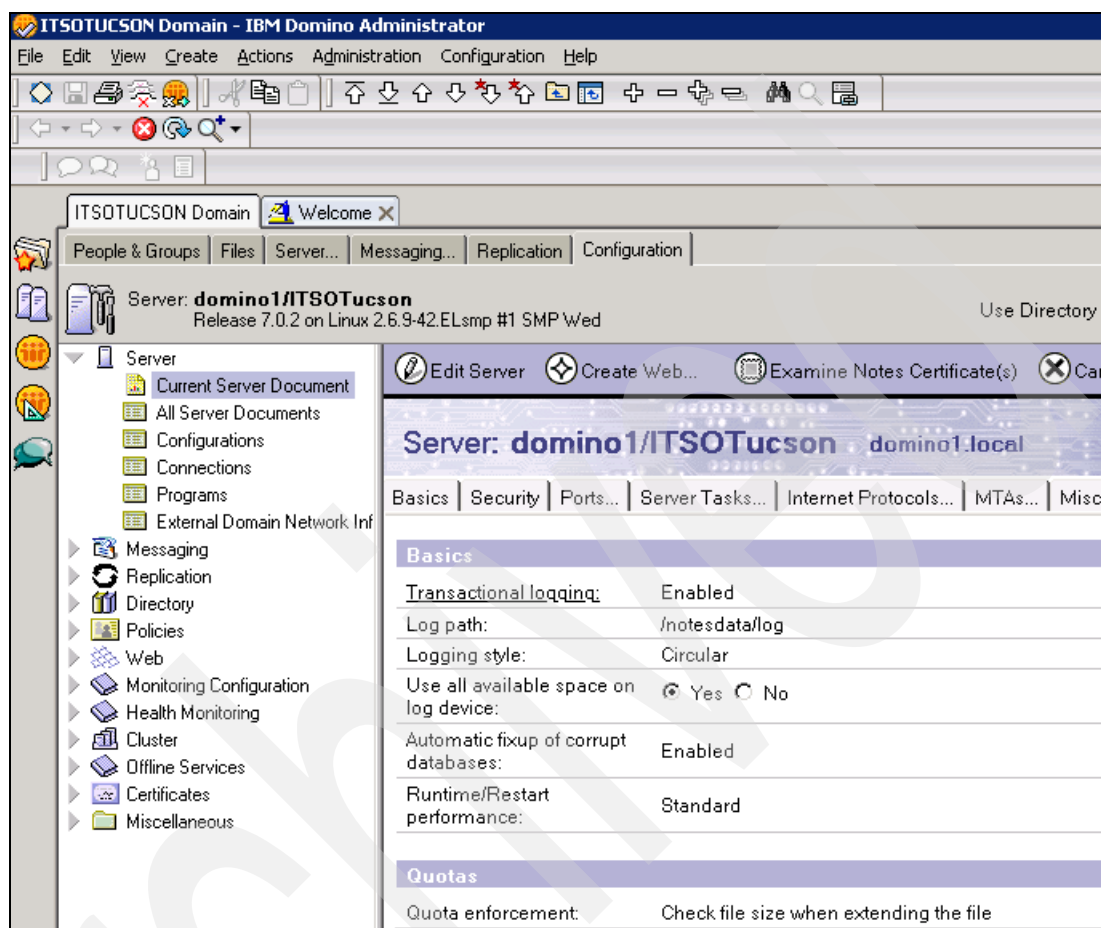


Figure A-5 Transactional logging configuration in the Lotus Domino Server Document

- ▶ **Circular:** Circular logging is the default transaction logging type. It uses one transaction log file of a predefined size (192 MB to 4 GB). The transaction log size is defined in the server document. The circular logging overwrites the old data once the defined size is reached with the new transaction. Domino writes the old transaction log data to the disk before it is overwritten. This method helps in server crash recovery, but limits the Domino administrator's ability to perform roll forward recovery because the old logs are overwritten. Circular logging does not support incremental backups.
- ▶ **Linear:** Similar to circular logging, but allows a log space larger than 4 GB.
- ▶ **Archival:** The method of transaction logging does not overwrite the transaction log file. After a file has reached its defined size, Domino creates a new transaction log file. After all the data is committed to the database, the old transaction log file becomes available for archiving. For archiving the transaction log file, you can use a third-party tool such as Tivoli or your custom scripts. Archive logging will enable rollback, media recovery, and point-in-time recovery.

Data and transaction logs *must* be stored on separate LUNs.

Configure the Domino environment

The following sections will explain how to prepare the IBM System Storage N series storage system to house the Lotus Domino database and log files.

Configuring the storage system

Before you create a Lotus Domino database on an IBM System Storage N series storage system, you need to configure the storage system by completing the following simple configuration steps:

1. Set up and install Data ONTAP V7.2 or above. Install Data ONTAP V7.2 on your storage system if it is not already installed.
2. Activate the license and start the appropriate services.

On the IBM System Storage N series storage system, file access protocols FCP and iSCSI are licensed services. You need to enable the following services by activating license keys for the protocol you intend to use:

- iSCSI, if you plan to use iSCSI-accessed virtual disks.
- FCP, if you plan to use FCP-accessed virtual disks.
- SnapRestore, which is required for restoring virtual disks from Snapshot copies.
- SnapMirror, if you plan to use the SnapMirror option for replicating the Domino database to another IBM System Storage N series storage system.

Note: To enable the iSCSI protocol and the iSCSI adapter on the IBM System Storage N series storage system, you will need an iSCSI license to be installed on the IBM System Storage N series storage system. In the FilerView, select **Filer** → **Manage Licenses** and scroll down the right pane until you see the iSCSI license. Add the iSCSI License and click **Apply**.

After activating the license key, you need to start the service by executing the following command on the Data ONTAP command-line interface (CLI) (see Example A-1):

```
[fcp | iscsi ] start
```

Example: A-1 Starting the iSCSI protocol

```
itsotuc3> iscsi start
iSCSI service started
itsotuc3>
```

3. Update the `/etc/host` file on the storage system.

The storage system must be able to communicate with the Windows server and vice versa. The storage system can communicate with the Windows server if there is an entry in its `/etc/hosts` file for it or, alternatively, if it uses some other host name resolution techniques like NIS or DNS. By default, the `/etc/hosts` file is checked first for host name resolution. The easiest way to update the `/etc/hosts` file on the storage system is by using FilerView. Entries made in the `/etc/hosts` file should have the following format:

```
[LDServerIP] [LDServerName]
```

Where:

- LDServerIP identifies the IP address assigned to the Windows server.
- LDServerName identifies the name assigned to the Windows server.

For example, to add an entry for a Windows server named ldsrv01 that has the IP address 172.17.32.112, you would add the following line to the /etc/hosts file on the IBM System Storage N series storage system:

```
172.17.32.112      ldsrv01
```

4. Enable rsh access

In order to use rsh (remote shell) from a Windows server, you must enable rsh by turning the rsh option on and editing the /etc/hosts.equiv files. You can turn the rsh option on by executing the command in Example A-2 on the storage system:.

Example: A-2 rsh enabling

```
itsotuc3> options rsh.enable on
itsotuc3>
```

Edit the /etc/hosts.equiv file and add an entry for each user on the Windows server that can issue the **rsh** command. Each entry should have the following format:

[LDServerName] [UserName]

Where

- LDServerName identifies the name assigned to the Windows server.
- UserName identifies the name assigned to the user on the Windows server.

For example, to enable rsh for a user named ldadmin on the Windows server named ldsrv01, you would add the following entry to the /etc/hosts.equiv file on the storage system:

```
ldsrv01 ldadmin
```

If a user is a domain user, then the entry should follow a format that looks similar to the following:

[LDServerName] [DomainName]\[UserName]

Where

- LDServerName identifies the name assigned to the Windows server.
- DomainName identifies the domain name the user belongs to.
- UserName identifies the name assigned to the user on the Windows server.

For example, to enable rsh from a Windows server named ldsrv01 for a user named ldadmin, who belongs to a domain named corp, you would add the following entry to the /etc/hosts.equiv file:

```
ldsrv1 corp\ldadmin
```

Create an aggregate

In order to create a Lotus Domino database on the storage system, you need to create aggregates and flexible volumes. An aggregate is a physical pool of storage, comprised of one or more RAID groups.

When creating the aggregate, a name should be defined. There are some naming conventions for the aggregate's name. It should:

- ▶ Begin with either a letter or an underscore.
- ▶ Contain only letters, digits, and underscores.
- ▶ Contain no more than 255 characters.

After you have the name, size, and disk configurations planned, these are the steps to create an aggregate:

1. Open the FilerView for the IBM System Storage N series storage system where you want to create the aggregate.
2. In the FilerView, select **Aggregates** → **Add**. This will bring up the Add New Aggregate window, as shown in Figure A-6. Click **Next**.



Figure A-6 Add New Aggregate window

3. The Aggregate Name window will be shown (Figure A-7). Type in the name for the aggregate you are creating. Select whether this aggregate will be a mirrored aggregate (by checking **Mirror**) or an unmirrored aggregate (by leaving **Mirror** unchecked). The parity should also be defined in this window. If you are creating a RAID-DP based aggregate, select the **Double Parity** check box. If the Double Parity check box is unchecked, the aggregate will be created using RAID 4. Click **Next**.

http://9.11.218.237/ - itsotuc3: Aggregate Wizard - Windows Internet Explorer

Aggregate Wizard - Aggregate Parameters

Aggregate Name:
Enter a name for the new aggregate. ?

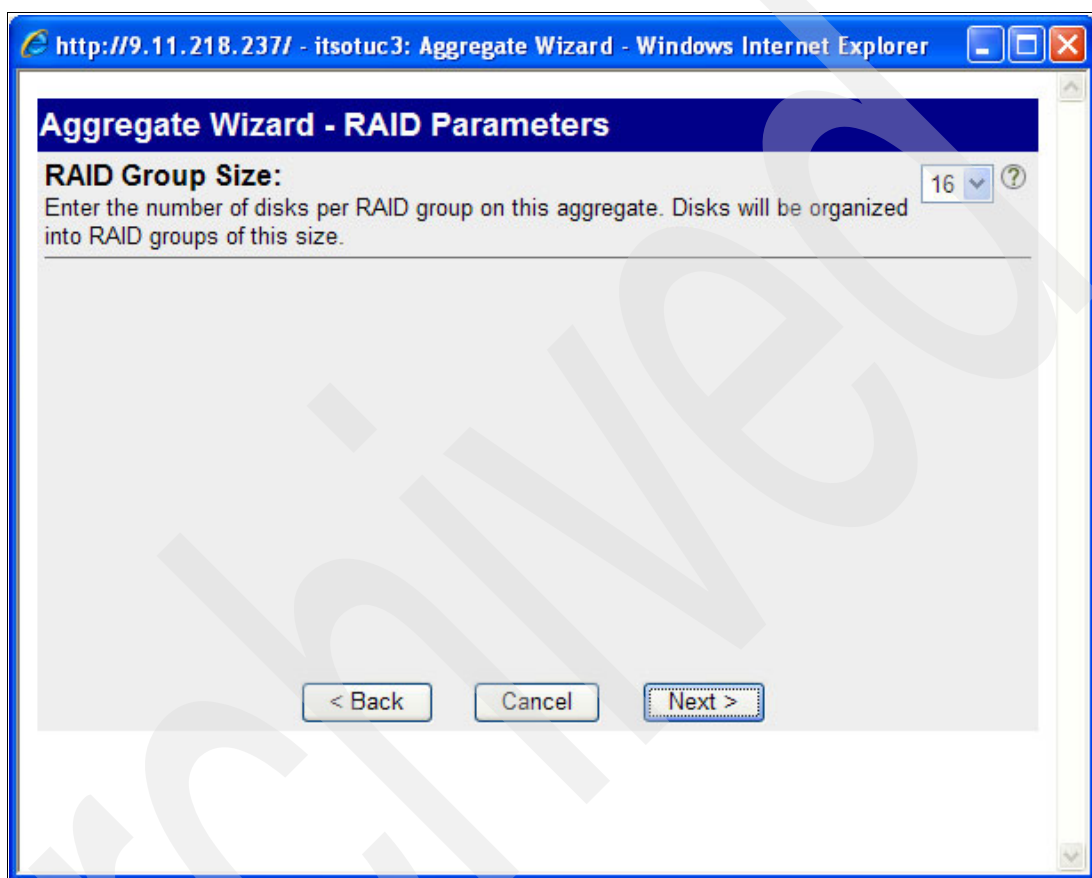
Synchronous Mirroring:
Select to enable local synchronous mirroring on this aggregate. Enabling this option requires twice as many disks. ☐ Mirror ?

Double Parity:
Select to enable double parity on this aggregate. Enabling this option requires an extra disk per RAID group. ☒ Double Parity ?

< Back Cancel Next >

Figure A-7 Aggregate Name window

4. In the RAID Parameters window (Figure A-8), select the number of disks that will be used on each RAID Group created for the aggregate. The recommended number of disks per RAID group is 16 disks. If you are using less than 16 disks per RAID Group, protection against disk failure is increased, but performance will decrease because there will be less disk spindles for accessing the data. If you are using more than 16 disks per RAID Group, performance will be increased (more disk spindles to access the data), but protection against disk failure will decrease. Click **Next**.



The screenshot shows a web browser window titled "http://9.11.218.237/ - itsotuc3: Aggregate Wizard - Windows Internet Explorer". The main content area is titled "Aggregate Wizard - RAID Parameters". Below the title, there is a section labeled "RAID Group Size:" with a text input field containing the number "16" and a help icon. Below this, there is a large, empty rectangular area. At the bottom of the window, there are three buttons: "< Back", "Cancel", and "Next >".

Figure A-8 RAID Parameters window

5. In the Disk Selection window (Figure A-9), select the method that should be used to identify the disks used on the aggregate. The default method is Automatic so that the IBM System Storage N series storage system will automatically select the disks based on your choices for size and number of disks from the next windows. If for any reason you need to select specific disks to compose the RAID Groups, click **Manual** and select the number and size of disks to be included on the aggregate. Click **Next**.




Figure A-9 Disk Selection window

6. If you select the Automatic method selection for the disks, the Disk Size window will be shown, as shown in Figure A-10. Select the disk size from the available options or select **Any Size** and click **Next**.




Figure A-10 Disk Size window

7. Select the number of disks of the selected size to be used on the aggregate in the Number of Disks window (Figure A-11). Click **Next**.



Figure A-11 Number of Disks window

8. Review your selection in the Summary window and click **Commit**.
9. The aggregate will be created. In the FilerView, select **Aggregates** → **Manage** and a list of the existing aggregates will be shown, along with their status, RAID level, size, available size, and other informations.

Creating a volume

Every volume on the IBM System Storage N series storage system must be created on an aggregate. For IBM Lotus Domino servers, the volumes should be created on different aggregates.

The volume name must follow these naming conventions:

- ▶ Begin with either a letter or an underscore.
- ▶ Contain only letters, digits, and underscores.
- ▶ Contain no more than 255 characters.

These are the steps to create the volume on the aggregate:

1. Open the FilerView for the IBM System Storage N series storage system where you want to create the aggregate.

2. In the FilerView, select **Volumes** → **Add**. This will bring up the Add New Volume window, as shown in Figure A-12. Click **Next**.




Figure A-12 Add New Volume window

3. The Volume Type Selection window will be shown (Figure A-13). Select **Flexible** for flexible volumes or **Traditional** for Traditional volumes. The recommended type for IBM Lotus Domino Server is FlexVol. Click **Next**.


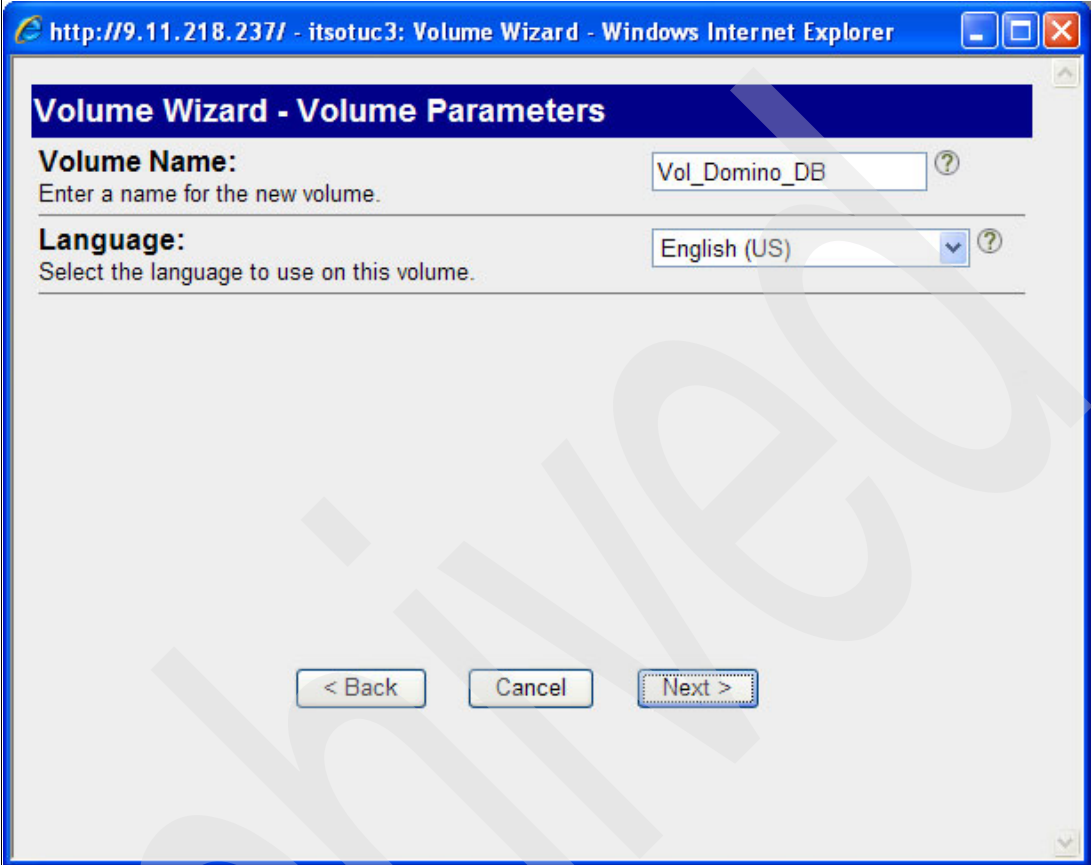


Figure A-13 Volume Type Selection window

4. In the Volume Parameters window (Figure A-14), type in the volume name and select the language used on the volume. By default, the root volume language will be selected. Click **Next**.



The screenshot shows a web browser window with the address bar displaying "http://9.11.218.237/ - itsotuc3: Volume Wizard - Windows Internet Explorer". The main content area has a title bar "Volume Wizard - Volume Parameters". Below the title bar, there are two sections: "Volume Name:" with a text input field containing "Vol_Domino_DB" and a help icon, and "Language:" with a dropdown menu showing "English (US)" and a help icon. Below these sections are three buttons: "< Back", "Cancel", and "Next >".

Figure A-14 Volume Parameters window

5. In the FlexVol Parameters window (Figure A-15), select the aggregate you want to create the volume on. Type in the size for the volume in KB, MB, GB, or TB. Select the type of Space Guarantee to be used. The default and recommended is volume. This option will pre-allocate the entire volume size on the aggregate. Other options are file space guarantee and none. Click **Next**.

Volume Wizard - Flexible Volume Parameters

Containing Aggregate
Select the aggregate to contain this volume. Only non-snaplock aggregates are displayed.

Aggr_Domino_DB (321 GB, raid_c)

Total Volume Size:
Enter the total amount of space for this volume. The total volume size includes space for the snapshot reserve and the file system overhead in addition to the usable space.

360 GB

Space Guarantee
Sets the space guarantee. Volume guarantees space for the entire the volume in the containing aggregate; File guarantees space for a file at file allocation time.

volume

< Back Cancel Next >

Figure A-15 FlexVol Parameters window

6. Review the selection in the Summary window and click **Commit**.
7. The volume will be created. In the FilerView, select **Volumes** → **Manage** and a list of the existing volumes will be shown, along with their status, RAID level, size, available size, and other information.

After the volumes are created for the Lotus Domino system, they must be shared. This is done by using the CIFS option on the FilerView.

1. In the FilerView, select **CIFS** → **Shares** → **Add**. The Add a CIFS Share window will be shown, as shown in Figure A-16 on page 429. Type in the following information:
 - a. Share Name: This is the name that will be used to access the volume for the LUN creation on the Lotus Domino server.
 - b. Mount Point: The path to connect to this volume on the IBM System Storage N series storage system, such as /vol/Vol_Domino_DB or /vol/Vol_Domino_Log.
 - c. Share Description: General description for the Share.
 - d. Max. Users: Maximum number of concurrent users at a time on the Share.
 - e. Force Group: Not used for volumes accessed by Windows hosts.
2. Click **Add**.

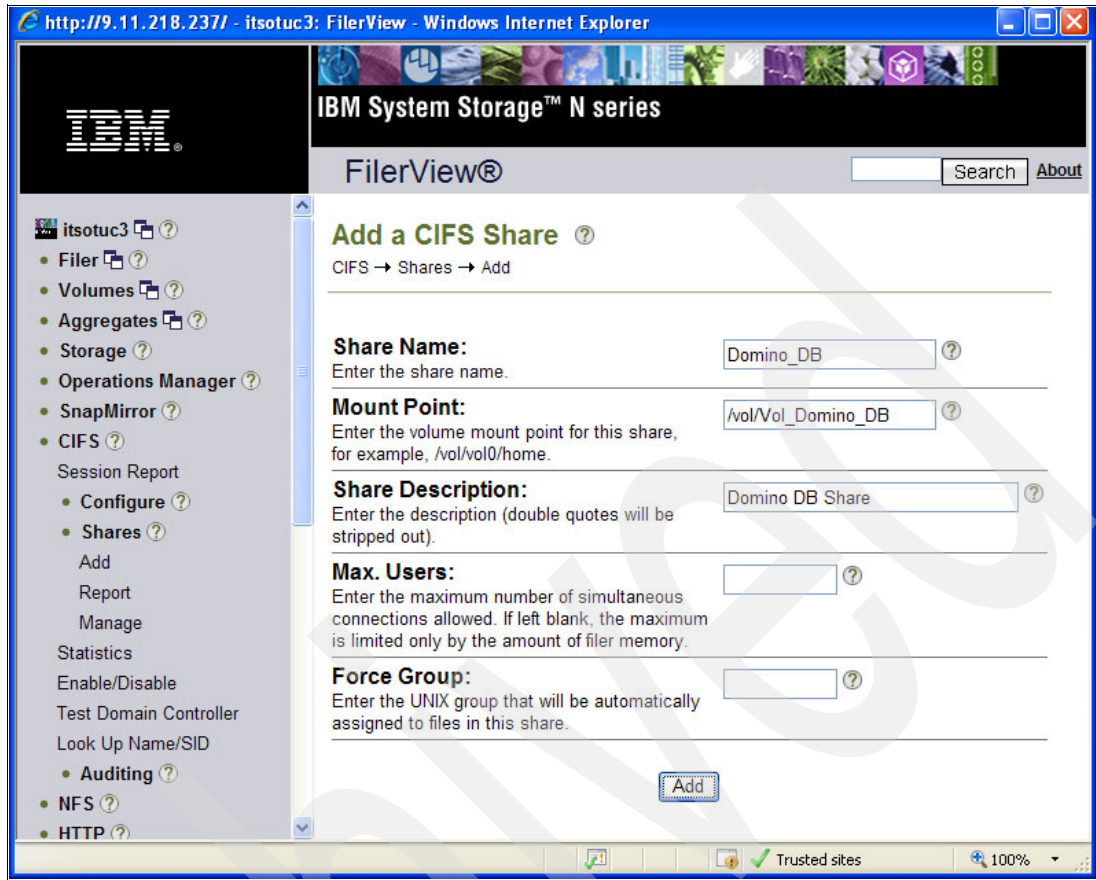


Figure A-16 Add a CIFS Share window

Specify the volume security settings to use

You must ensure that the volume(s) on the storage system that are to be used to house a Domino database support either the NTFS or the Mixed security style. The security setting of a volume can be set using the following command:

```
qtree security [VolName] [Unix | NTFS | mixed ]
```

where VolName identifies the name assigned to the new volume.

For example, to change the security style of a volume named dominodata to NTFS, you would execute the command in Example A-3 on the storage system:.

Example: A-3 Setting the volume security options

```
itsotuc3> qtree security /vol/vol_Domino_DB NTFS
itsotuc3>
```

Repeat this step and change the security style for all the volumes to be used for the database or transactional logs.

Disable the automatic Snapshot feature

Normally, a Lotus Domino database is backed up based on a user-defined schedule. Therefore, we recommend that you turn off the automatic Snapshot feature for all volumes that are used for the database and its transaction log files. The automatic Snapshot feature can be turned off by executing the following command on the storage system:

```
vol options [VolName] nosnap on
```

where VolName identifies the name assigned to the new volume.

For example, to turn the automatic Snapshot feature off for a volume named dominodata, you would execute the command in Example A-4 on the storage system:

Example: A-4 Disable the automatic Snapshot feature

```
itsotuc3> vol options vol_Domino_DB nosnap on  
itsotuc3>
```

Repeat this step and turn the auto Snapshot off for all the volumes that are used by the Database.

Setting the snap reserve option for the volume

By default, Data ONTAP sets the snap reserve option for new volumes to 20%. You should reset the snap reserve option to 0% on all volumes holding SnapDrive virtual disks. To reset the snap reserve options, you would execute the following command on the storage system:

```
snap reserve -V [VolName] [Percent]
```

Where:

- ▶ VolName identifies the name assigned to the new volume.
- ▶ Percent identifies the volume space reserved for the Snapshot copies. A valid value is between 0 and 100.

For example, to set **snap reserve** to 0% for a flexible volume named dominodata, you would execute the command in Example A-5 on the storage system.

Example: A-5 Setting the snap reserve option

```
itsotuc3> snap reserve -V vol_Domino_DB 0  
itsotuc3>
```

Obtain a node name for the storage system

Target IBM System Storage N series storage systems are identified by a unique node name that is used for creating persistent binding for the storage devices on the host. You can find out the node name by executing the following command:

```
[ fcp | iscsi ] nodename
```

For example, to find the node name assigned to an iSCSI target, you would execute the command in Example A-6 on page 431 on the IBM System Storage N series storage system.

Example: A-6 Obtain the node name for the storage system

```
itsotuc3> iscsi nodename  
iSCSI target nodename: iqn.1986-03.com.ibm:sn.101165597  
itsotuc3>
```

Make a note of the node name returned by this command; you will need it later when configuring the HBA on the Windows server.

Create a user with administrator privileges

Create a user that belongs to the local administrator group by entering the following command on the storage system:

```
useradmin useradd [UserName] -g [AdminGroup]
```

Where:

- ▶ `UserName` identifies the name assigned to the user.
- ▶ `AdminGroup` identifies the name assigned to the local administrator group.

For example, to create a user named `ldadmin` that belongs to the local administrator group named `administrators`, you would execute the command in Example A-7 on the storage system.

Example: A-7 Create a user with administrator privileges

```
itsotuc3> useradmin user add ldadmin -g administrators  
New password:  
Retype new password:  
User <ldadmin> added.  
itsotuc3>
```

Configuring the Windows server

The following sections will explain how to configure the Windows server to access the disks on the IBM System Storage N series storage system.

Installing the Microsoft iSCSI software

Microsoft iSCSI Software Initiator is the software installed on the server that allows the SCSI communication over TCP/IP. This software is required if you are using SCSI protocol to communicate with the IBM System Storage N series storage system but do not have the hardware based iSCSI adapters.

Microsoft iSCSI Software Initiator will create additional layers in the network communication so that a layer for the iSCSI protocol and for SCSI drivers be present. In this manner, the regular Network Interface Card (NIC) could be used to communicate with the IBM System Storage N series storage system.

As a best practice, always use the latest version of software and drivers on your environment, unless there are compatibility issues. For information about Microsoft iSCSI Software Initiator and the latest versions, go to:

<http://www.microsoft.com/windowsserver2003/technologies/storage/iscsi/msfiSCSI.msp>
x

The following steps outline the Microsoft iSCSI Software Initiator software installation:

1. In the Welcome window (Figure A-17), click **Next**.



Figure A-17 Welcome window

2. In the Installation Options window (Figure A-18), select the following options and click **Next**:
 - a. Initiator Service
 - b. Software Initiator
 - c. Microsoft MPIO Multipathing Support for iSCSI

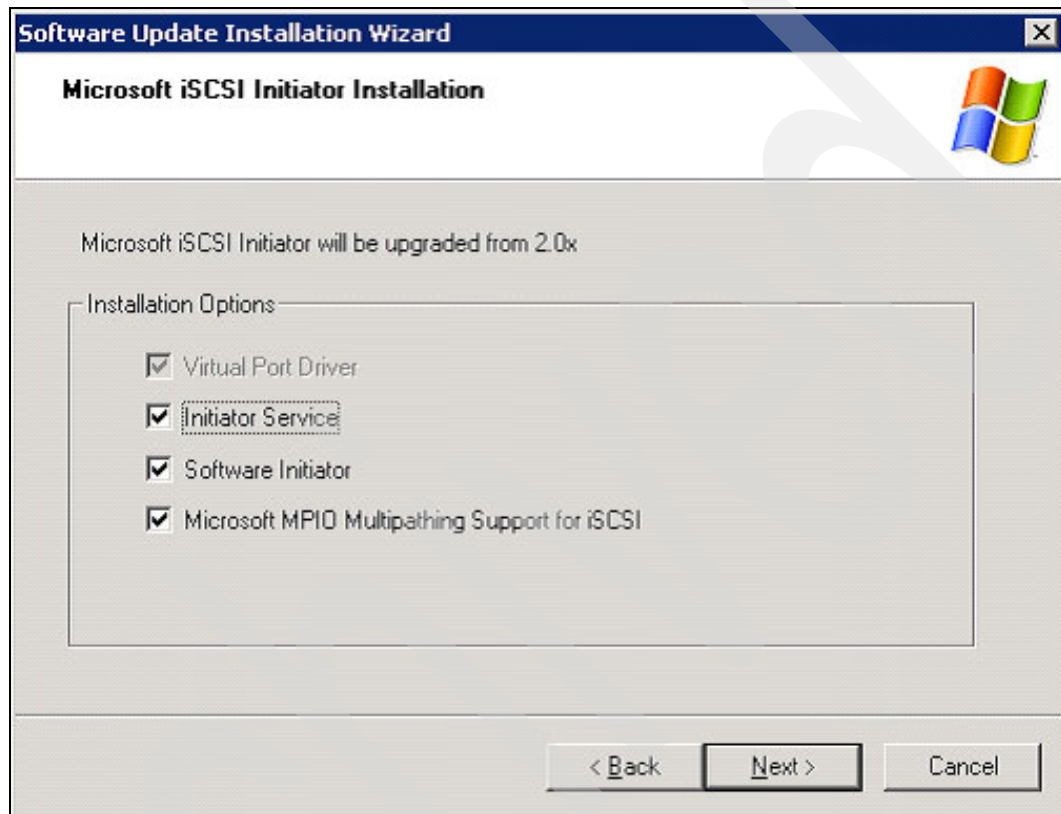


Figure A-18 Installation options window

3. In the License Agreement window (Figure A-19), agree to the terms and click **Next**.




Figure A-19 License Agreement window

4. The installation starts. At the end of the installation, in the Finish window (Figure A-20 on page 435), click **Finish**. If you do not want your server to reboot now, select the check box **Do not restart now**. Otherwise, your server will reboot immediately.



Figure A-20 Finish window

Configuring the Microsoft iSCSI software

1. Start Microsoft iSCSI Initiator by clicking the desktop icon or by selecting **Start** → **All Programs** → **Microsoft iSCSI Initiator** → **Microsoft iSCSI Initiator**. This will bring up the iSCSI Initiator Properties window.

2. Copy the Initiator Node Name from the iSCSI Initiator Properties window (Figure A-21). In this example, it would be `iqn.1991-05.com.microsoft:namgmt1.itso.tucson`.

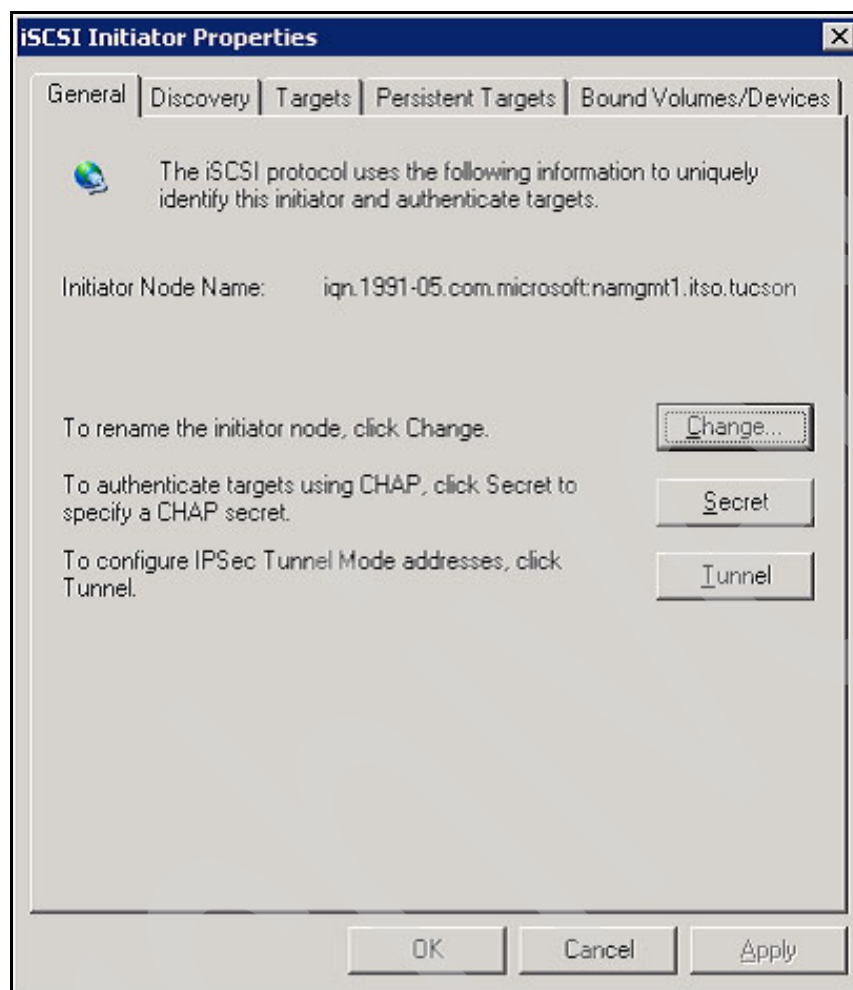


Figure A-21 iSCSI Initiator Properties window

3. In the FilerView, select **LUNs** → **Initiator Groups** → **Add**. In the Add Initiator Group window (Figure A-22), type in a name for the group you are creating, select **iSCSI** as the protocol for the group, select **Windows** as the operating system, and paste in the iSCSI initiator node name you copied in the previous step.

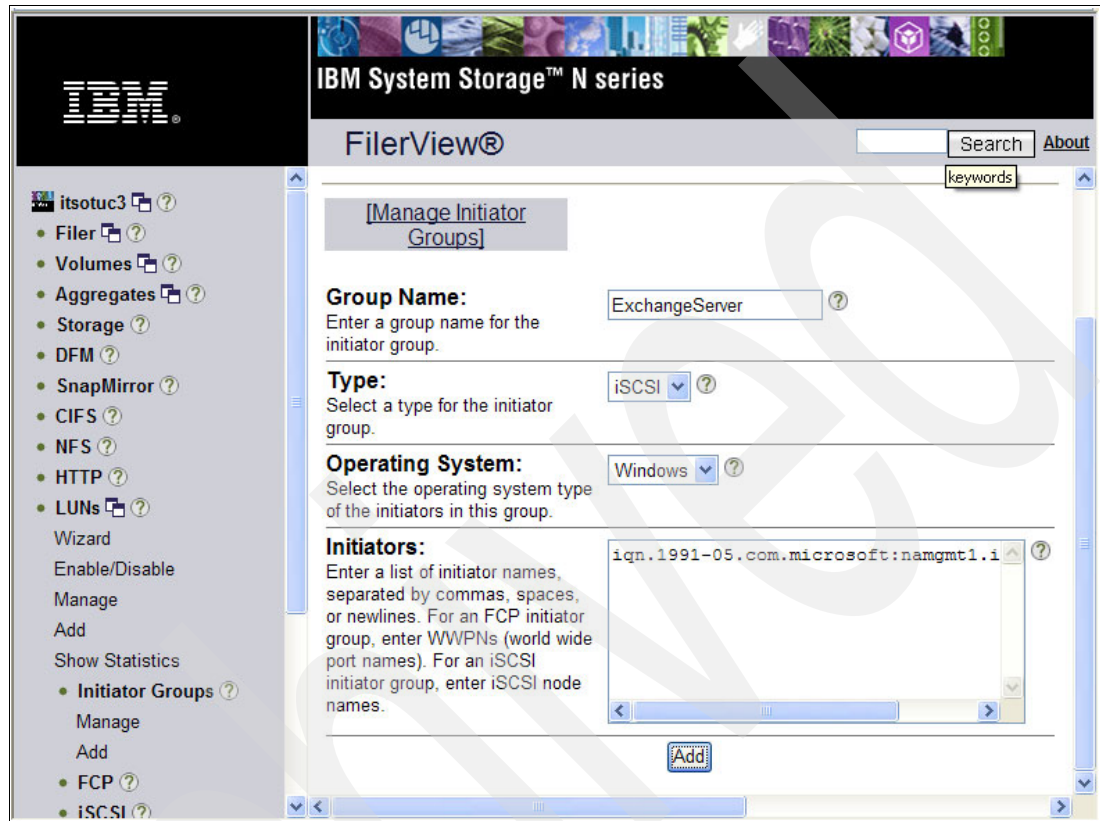


Figure A-22 Add Initiator Group window

4. In the Lotus Domino Server, click the **Discovery** tab and the iSCSI Initiator Discovery window will be shown (Figure A-23). Click **Add** in the Target Portals session.

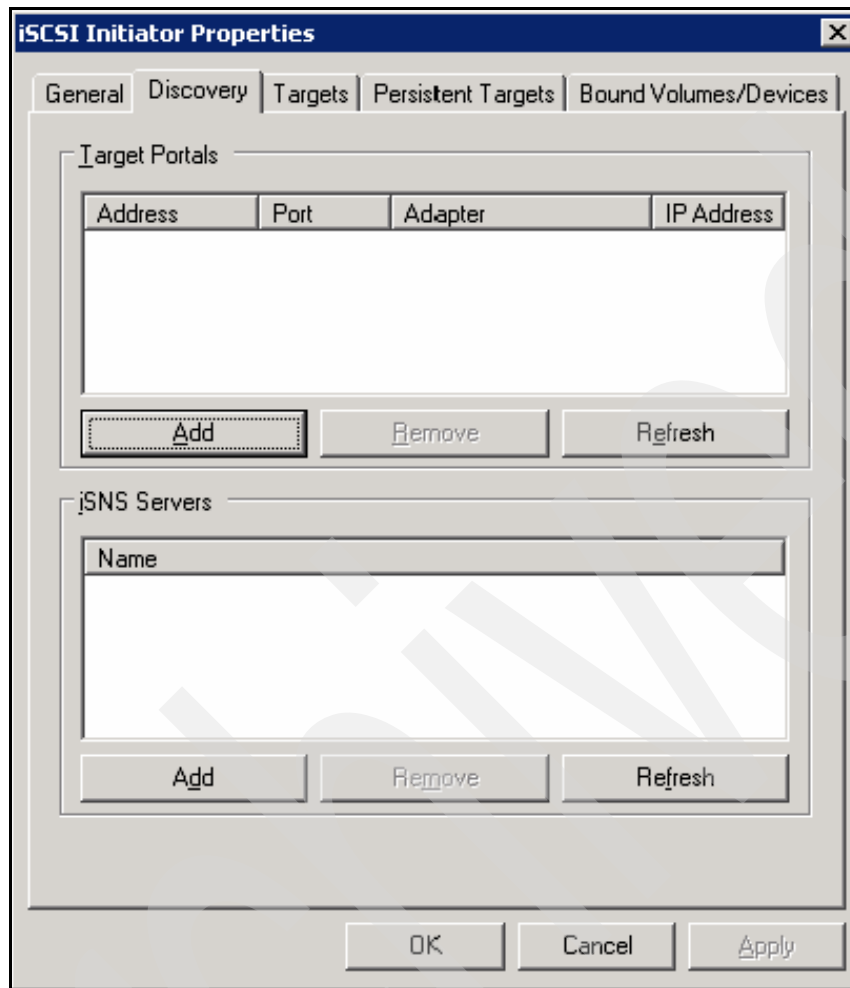


Figure A-23 iSCSI Initiator Discovery window

5. In the Add Target Portal window (Figure A-24), type in the IP address or DNS name of the filer and click **Advanced**.

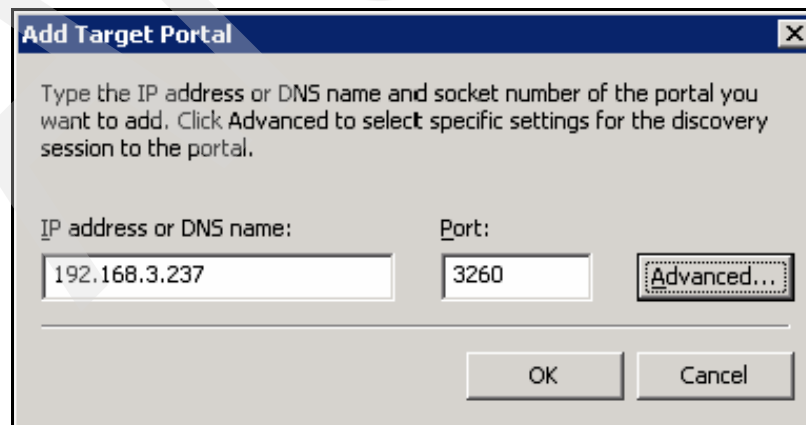


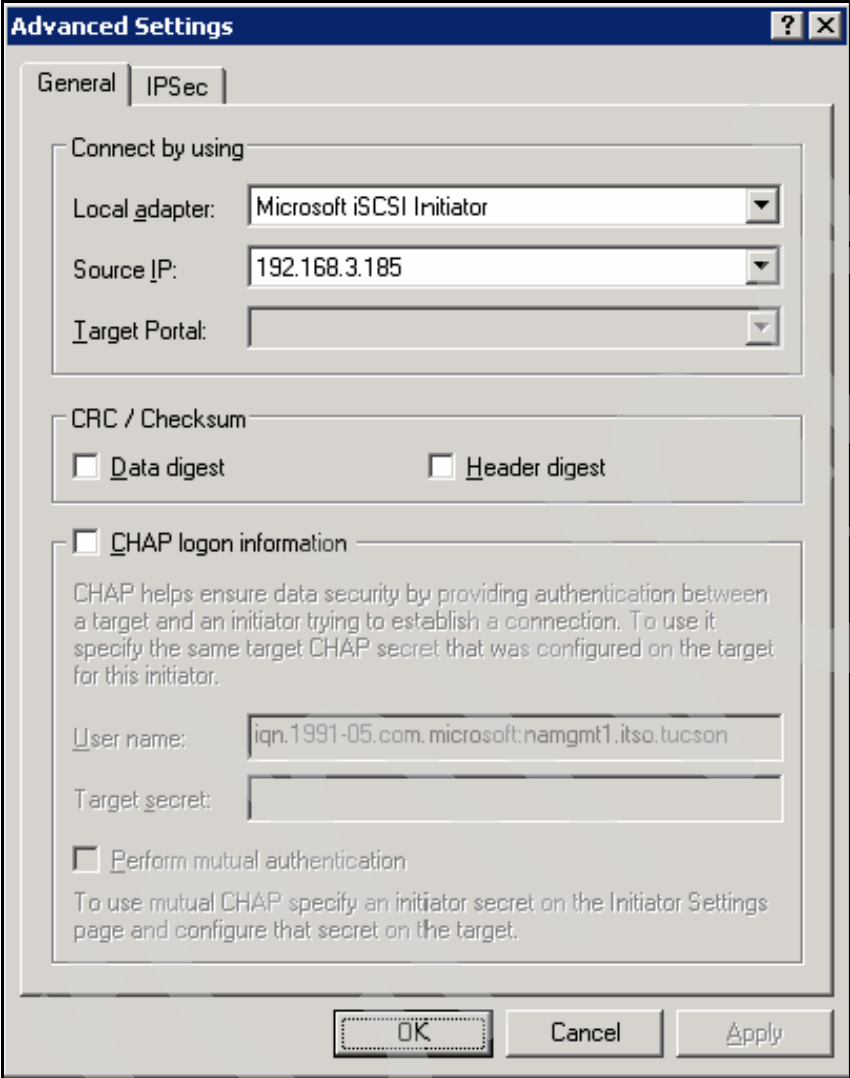
Figure A-24 Add Target Portal window

6. In the Advanced Settings window (Figure A-25), select **Microsoft iSCSI Initiator** as the Local Adapter and then select one of the IP addresses as the Source IP. At this time, the CHAP Authentication protocol may be configured, if defined on the IBM System Storage N series storage system configuration. For this scenario, we are not going to use CHAP. Click **OK**.

The image shows a Windows-style dialog box titled "Advanced Settings" with a blue header bar containing a question mark and a close button. The dialog has two tabs: "General" and "IPSec", with "IPSec" currently selected. The "Connect by using" section contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.181", and "Target Portal:" which is empty. Below this is a "CRC / Checksum" section with two unchecked checkboxes: "Data digest" and "Header digest". The "CHAP logon information" section is also unchecked and contains a text box for "User name:" with the value "iqn.1991-05.com.microsoft:namgmt1.itso.tucson", an empty "Target secret:" text box, and an unchecked "Perform mutual authentication" checkbox. A descriptive paragraph explains that CHAP helps ensure data security by providing authentication between a target and an initiator. At the bottom of the dialog are three buttons: "OK", "Cancel", and "Apply".

Figure A-25 Advanced Settings window

7. Repeat steps 4 on page 438, 5 on page 438, and 6 on page 439 for the additional IP address, as shown in Figure A-26. Click **OK**.



The image shows a Windows-style dialog box titled "Advanced Settings" with a question mark icon and a close button. It has two tabs: "General" and "IPSec", with "IPSec" currently selected. The dialog is divided into several sections. The first section, "Connect by using", contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.185", and "Target Portal:" which is empty. The second section, "CRC / Checksum", contains two checkboxes: "Data digest" and "Header digest", both of which are unchecked. The third section, "CHAP logon information", starts with an unchecked checkbox. Below it is a text box with the text "CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it specify the same target CHAP secret that was configured on the target for this initiator." This is followed by a "User name:" label and a text box containing "iqn.1991-05.com.microsoft:namgmt1.itso.tucson", and then a "Target secret:" label and an empty text box. The final section contains an unchecked checkbox labeled "Perform mutual authentication" and a text box with the text "To use mutual CHAP specify an initiator secret on the Initiator Settings page and configure that secret on the target." At the bottom of the dialog are three buttons: "OK", "Cancel", and "Apply".

Advanced Settings

General | **IPSec**

Connect by using

Local adapter: Microsoft iSCSI Initiator

Source IP: 192.168.3.185

Target Portal:

CRC / Checksum

☐ Data digest ☐ Header digest

☐ CHAP logon information

CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it specify the same target CHAP secret that was configured on the target for this initiator.

User name: iqn.1991-05.com.microsoft:namgmt1.itso.tucson

Target secret:

☐ Perform mutual authentication

To use mutual CHAP specify an initiator secret on the Initiator Settings page and configure that secret on the target.

OK Cancel Apply

Figure A-26 Advanced Settings window

8. In the iSCSI Initiator Discovery window (Figure A-27), you will notice that two paths were created for the same Target Portal, one for each interface. Depending on your infrastructure configuration, this may change due to more Target Portals configured on the IBM System Storage N series storage system.




Figure A-27 iSCSI Initiator Discovery window

9. Click the **Targets** tab and a list of the targets iSCSI storage devices will be shown (Figure A-28). Click **Log On** to configure the paths to the IBM System Storage N series storage system.

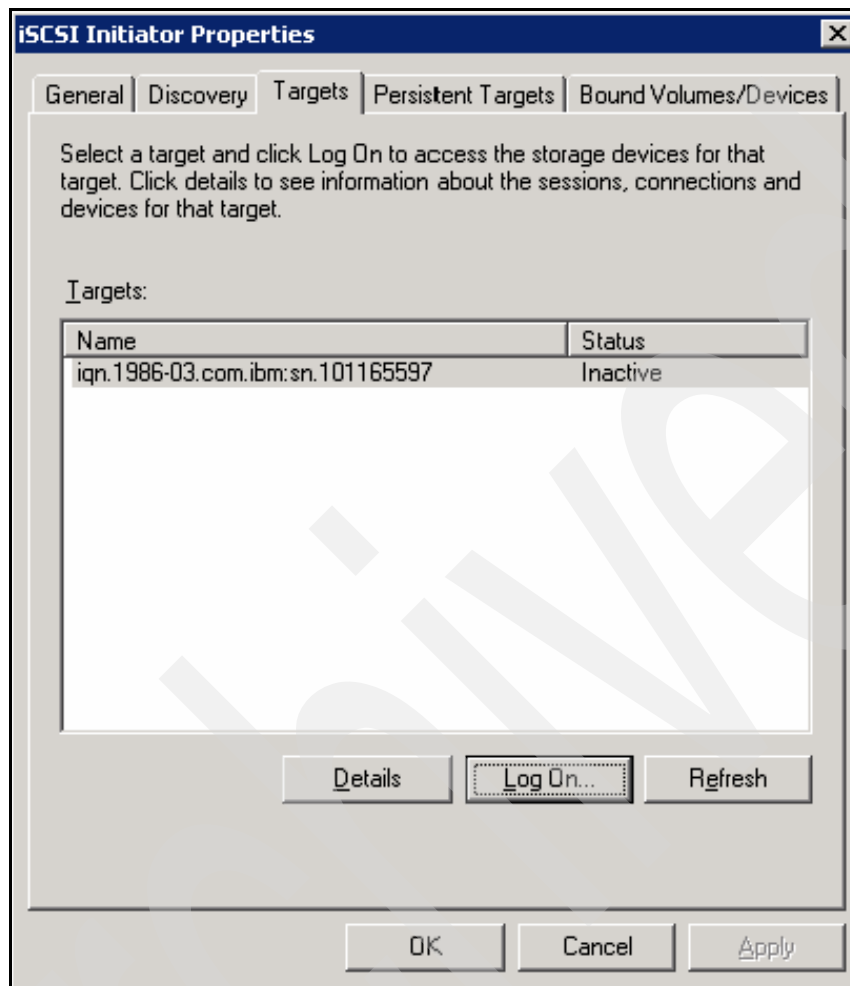


Figure A-28 iSCSI Initiator Targets window

10. In the Log On to Target window, select both check boxes, as shown in Figure A-29. This will configure the path be persistent and will enable the multipathing. Click **Advanced**.


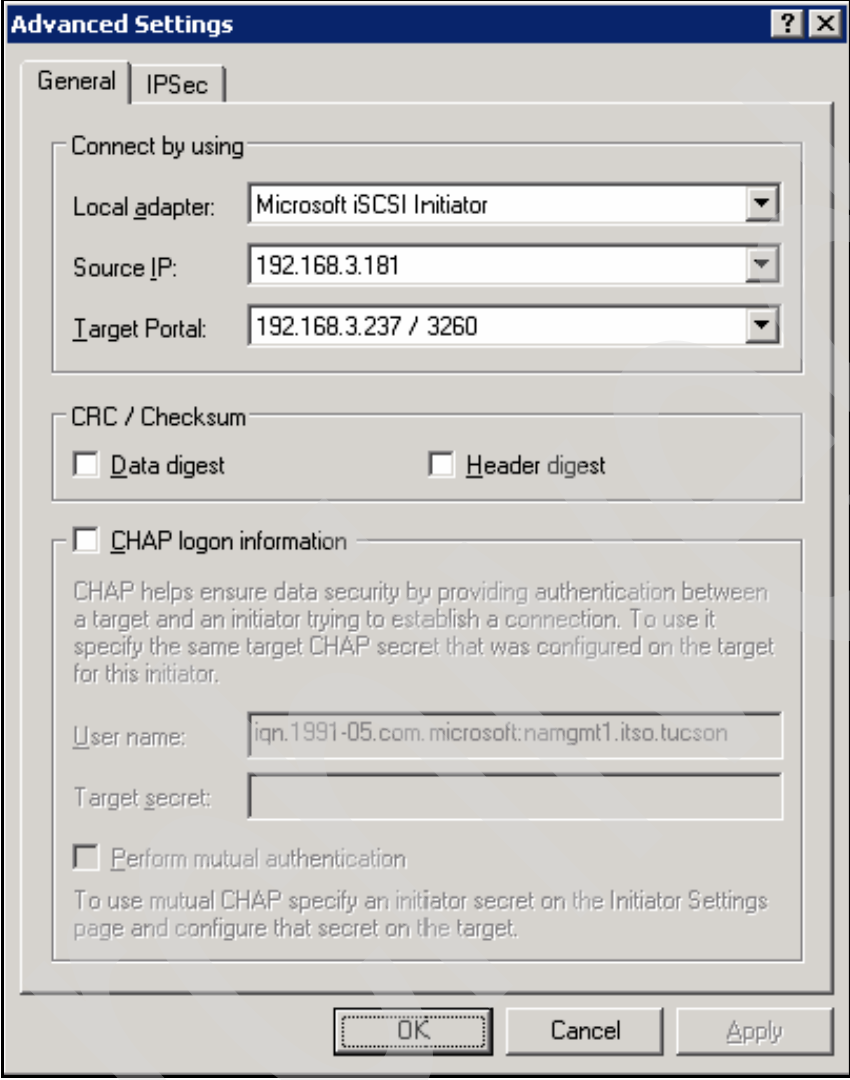


Figure A-29 Log On to Target window

11. In the Advanced Settings window (Figure A-30), select **Microsoft iSCSI Initiator** as the Local Adapter, select the first IP address as the Source IP, and select the proper Target Portal's combination of IP Address and port number. Click **OK**.



The image shows a Windows-style dialog box titled "Advanced Settings" with a question mark and close button in the title bar. It has two tabs: "General" and "IPSec", with "IPSec" currently selected. The "Connect by using" section contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.181", and "Target Portal:" set to "192.168.3.237 / 3260". Below this is a "CRC / Checksum" section with two unchecked checkboxes: "Data digest" and "Header digest". The "CHAP logon information" section is also unchecked and contains a text box for "User name:" with the value "iqn.1991-05.com.microsoft:namgmt1.itso.tucson", an empty "Target secret:" text box, and an unchecked "Perform mutual authentication" checkbox. A descriptive paragraph explains CHAP's role in data security. At the bottom are "OK", "Cancel", and "Apply" buttons.

Advanced Settings

General | **IPSec**

Connect by using

Local adapter: Microsoft iSCSI Initiator

Source IP: 192.168.3.181

Target Portal: 192.168.3.237 / 3260

CRC / Checksum

☐ Data digest ☐ Header digest

☐ CHAP logon information

CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it specify the same target CHAP secret that was configured on the target for this initiator.

User name: iqn.1991-05.com.microsoft:namgmt1.itso.tucson

Target secret:

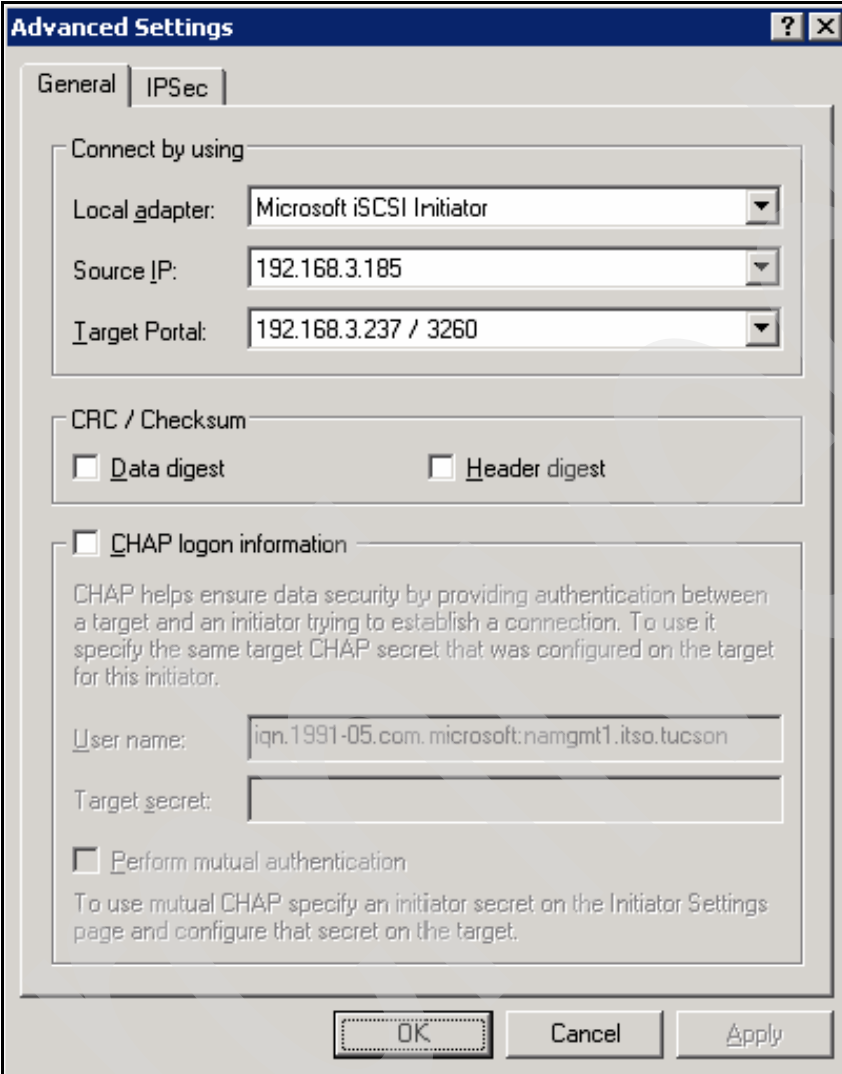
☐ Perform mutual authentication

To use mutual CHAP specify an initiator secret on the Initiator Settings page and configure that secret on the target.

OK Cancel Apply

Figure A-30 Advanced Settings window

12.Repeat steps 9 on page 442, 10 on page 442, and 11 on page 443 for the additional IP address on the Lotus Domino Server. Select this additional IP address as the Source IP, as shown in Figure A-31. Click **OK**.

The image shows a Windows-style dialog box titled "Advanced Settings" with a blue header bar containing a question mark and a close button. It has two tabs: "General" and "IPSec", with "IPSec" being the active tab. The "Connect by using" section contains three dropdown menus: "Local adapter:" set to "Microsoft iSCSI Initiator", "Source IP:" set to "192.168.3.185", and "Target Portal:" set to "192.168.3.237 / 3260". Below this is a "CRC / Checksum" section with two unchecked checkboxes: "Data digest" and "Header digest". Further down is a "CHAP logon information" section, also with an unchecked checkbox. It includes a text box for "User name:" containing "iqn.1991-05.com.microsoft:namgmt1.itso.tucson" and an empty "Target secret:" text box. A final unchecked checkbox is labeled "Perform mutual authentication". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Apply".

Advanced Settings

General | **IPSec**

Connect by using

Local adapter: Microsoft iSCSI Initiator

Source IP: 192.168.3.185

Target Portal: 192.168.3.237 / 3260

CRC / Checksum

☐ Data digest ☐ Header digest

☐ CHAP logon information

CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it, specify the same target CHAP secret that was configured on the target for this initiator.

User name: iqn.1991-05.com.microsoft:namgmt1.itso.tucson

Target secret:

☐ Perform mutual authentication

To use mutual CHAP, specify an initiator secret on the Initiator Settings page and configure that secret on the target.

OK Cancel Apply

Figure A-31 Advanced Settings window

13.Click **OK** to close the iSCSI Initiator Properties window.

At this point, the communication between the IBM Lotus Domino Server and the IBM System Storage N series storage system is established. The SnapDrive installation can take place.

Installing SnapDrive for Windows

In order to simplify storage management, you need to install a SnapDrive on Windows server that will house the Domino software. SnapDrive for Windows is a licensed product and can be obtained by contacting IBM Support.

These are the steps to install SnapDrive on the Lotus Domino Server:

1. In the SnapDrive installation welcome window (Figure A-32), click **Next**.

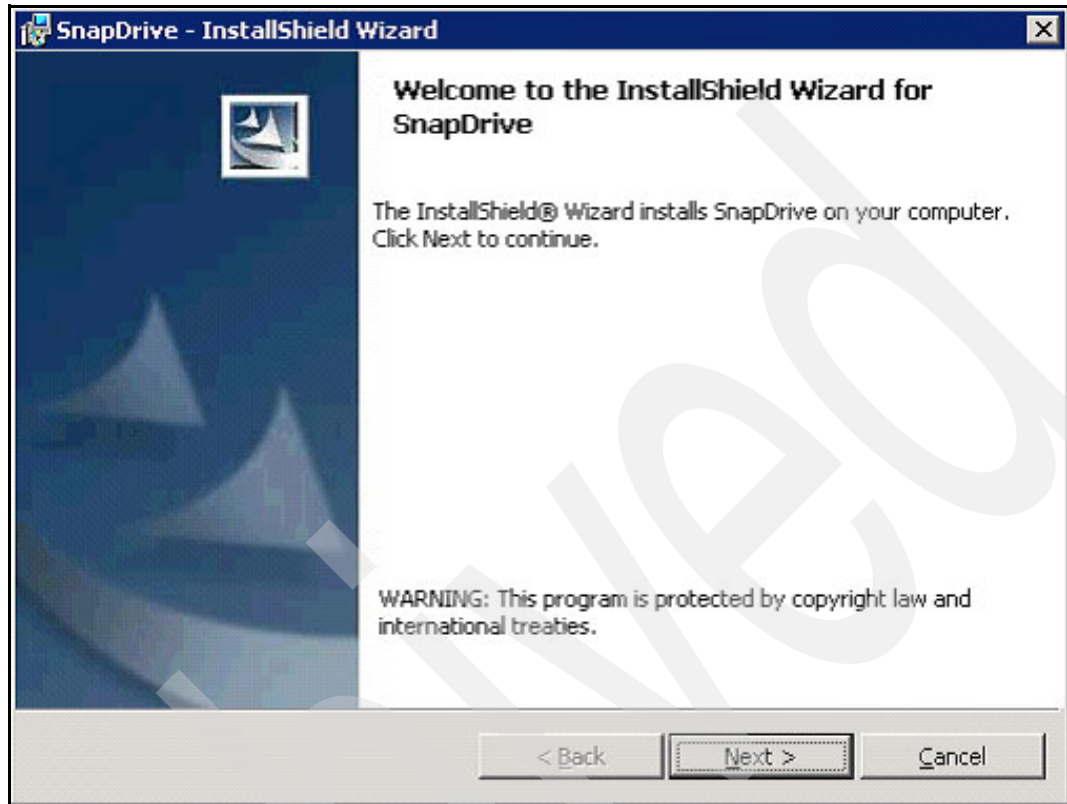


Figure A-32 SnapDrive installation welcome window

2. In the License Agreement window (Figure A-33), accept the terms of the License Agreement and click **Next**.

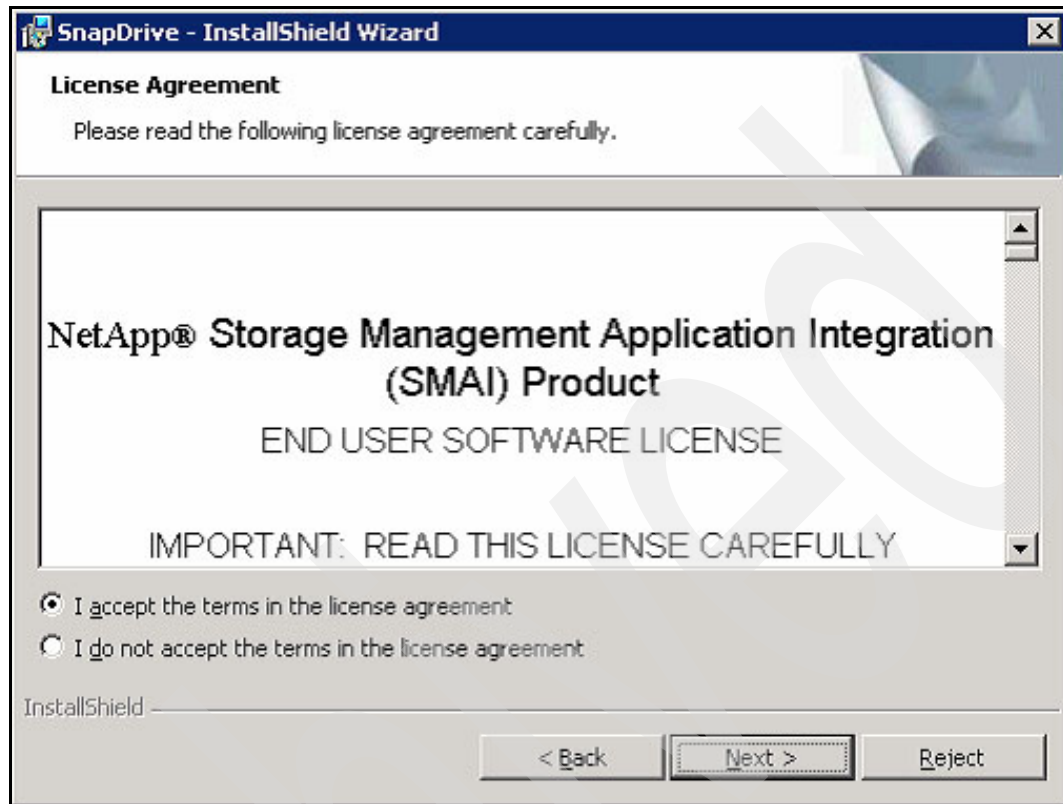


Figure A-33 License Agreement

3. In the License key window (Figure A-34), type in the License key for SnapDrive and click **Next**.

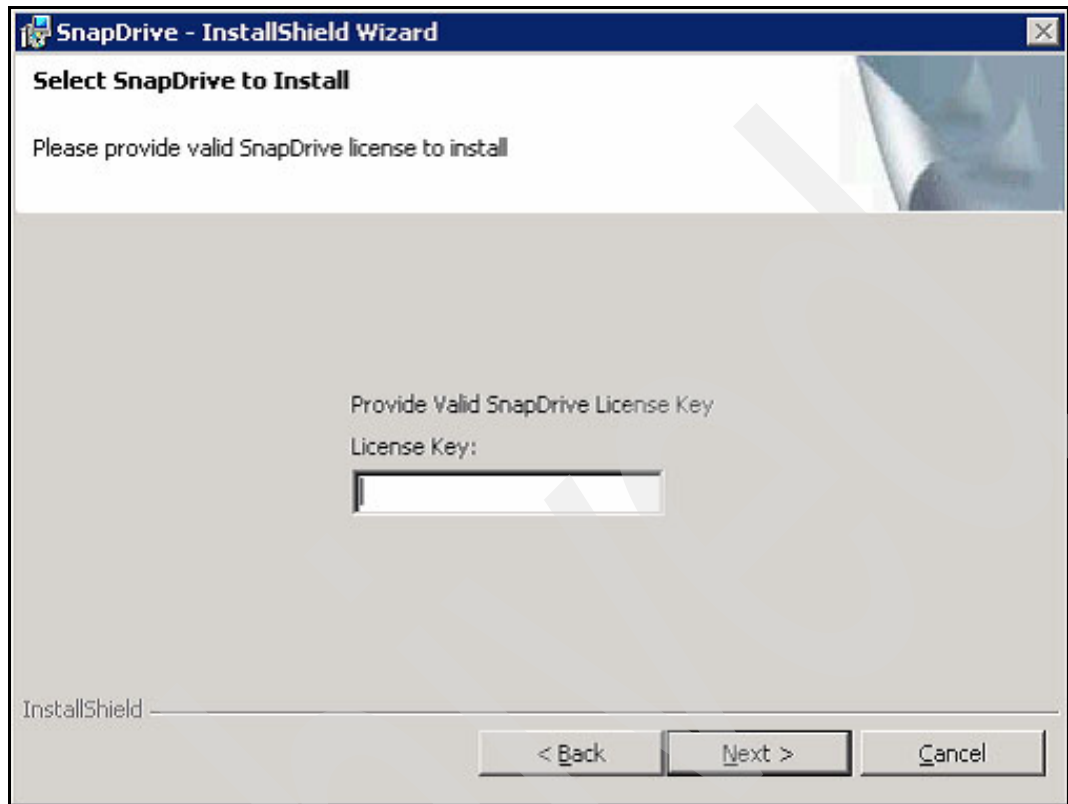


Figure A-34 License Key

4. SnapDrive will check for the minimum requirements for the iSCSI driver version, as shown in Figure A-35. Click **Next**.

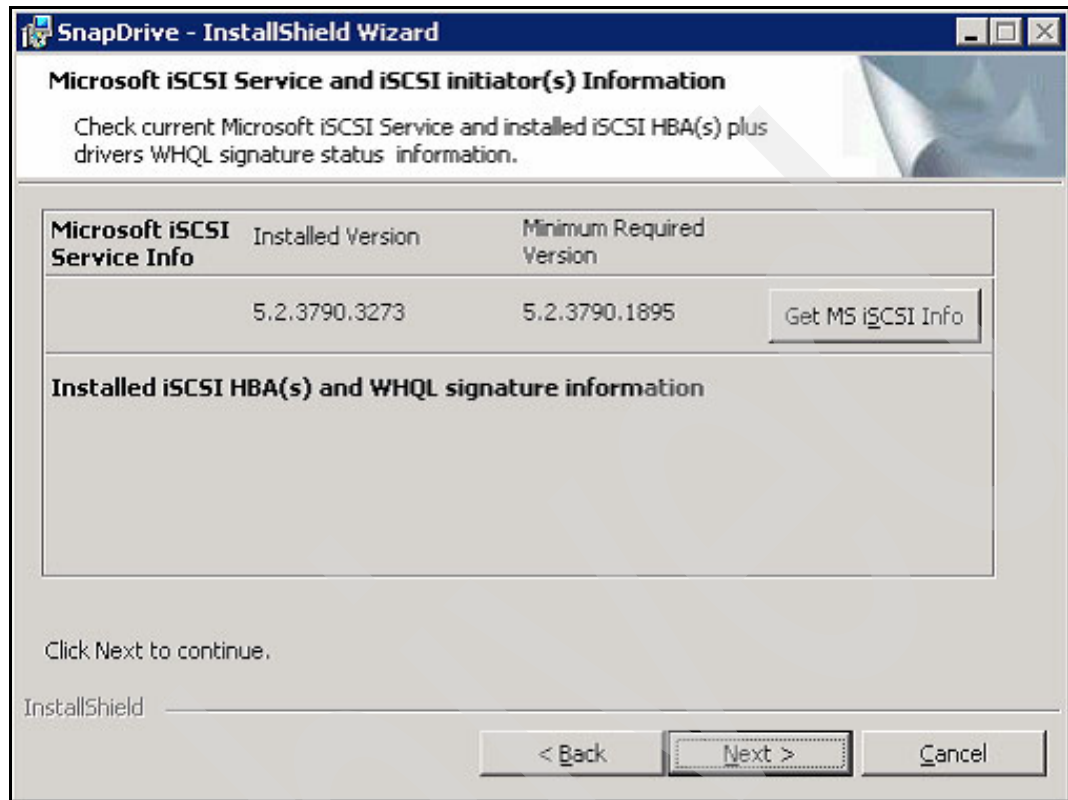
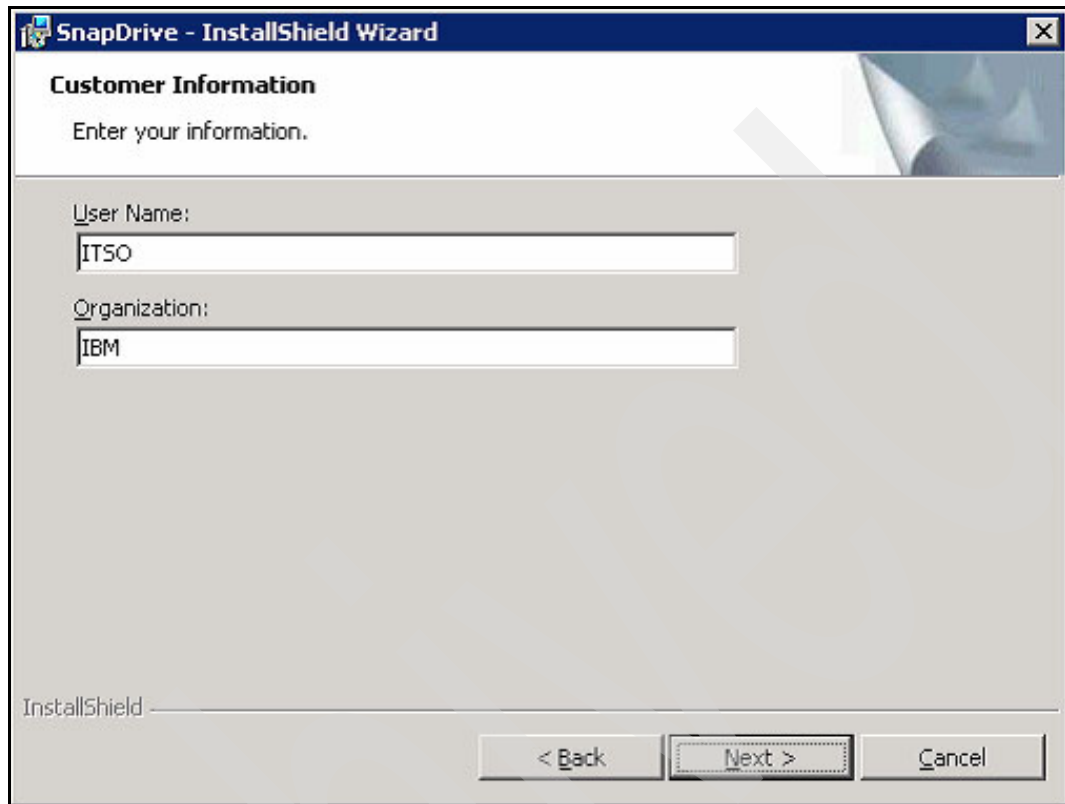


Figure A-35 iSCSI Service and iSCSI Initiator Information window

5. In the Customer information window (Figure A-36), type in the User Name and Organization information and click **Next**.



The image shows a Windows-style dialog box titled "SnapDrive - InstallShield Wizard". The window has a blue title bar with a close button (X) in the top right corner. Below the title bar, the text "Customer Information" is displayed in bold, followed by the instruction "Enter your information." in a smaller font. The main area of the window contains two text input fields. The first field is labeled "User Name:" and contains the text "ITSO". The second field is labeled "Organization:" and contains the text "IBM". At the bottom of the window, there is a status bar that says "InstallShield". To the right of the status bar are three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

Figure A-36 Customer information

6. In the Destination Folder window (Figure A-37), confirm or change the destination folder for installation files and click **Next**.

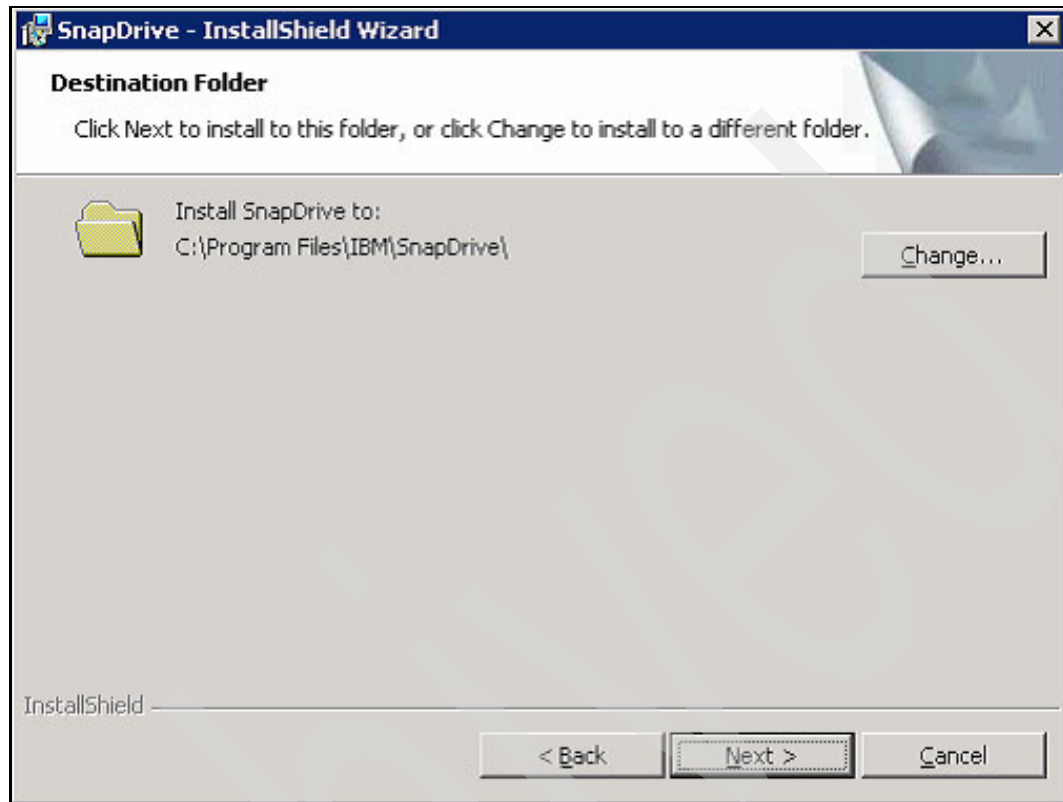
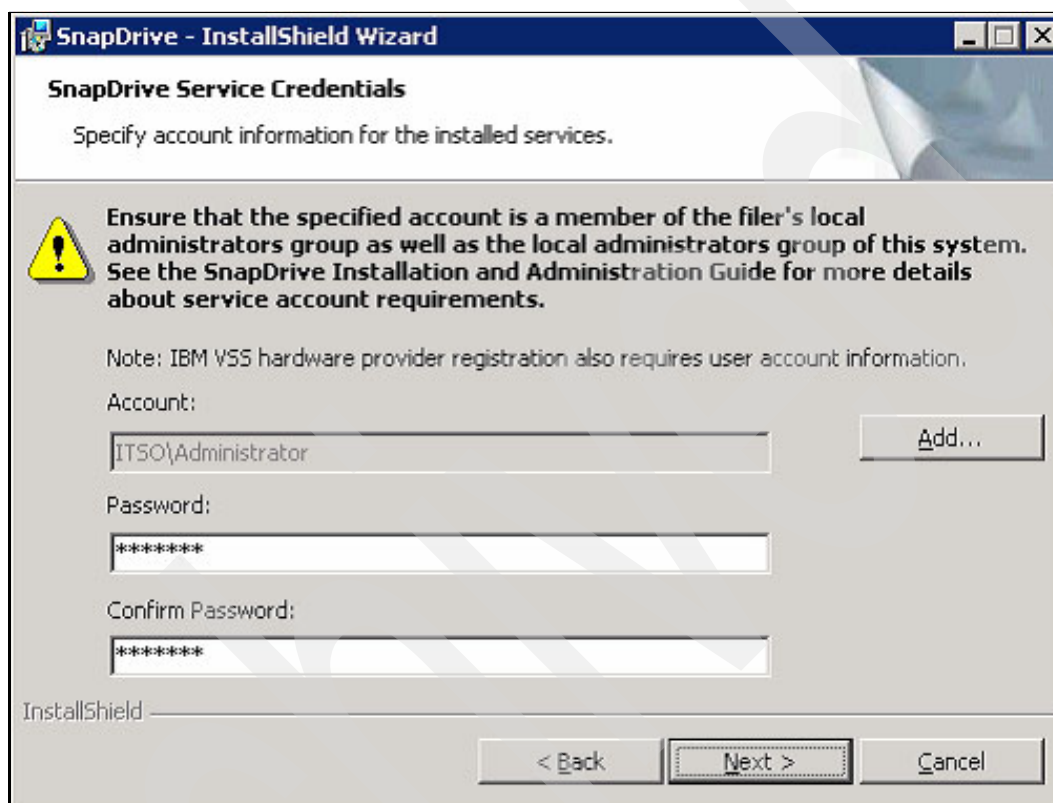


Figure A-37 Destination Folder

7. In the SnapDrive Service Credentials window (Figure A-38), type in the account and password for the user account to be used to start the SnapDrive service and click **Next**.

Note: The SnapDrive service user account should be a member of the Active Directory's Domain Admins group and be a member of the filer's local administrators group.



The image shows a Windows-style dialog box titled "SnapDrive - InstallShield Wizard". The subtitle is "SnapDrive Service Credentials". Below the subtitle, it says "Specify account information for the installed services." There is a yellow warning icon with an exclamation mark. The text next to the icon reads: "Ensure that the specified account is a member of the filer's local administrators group as well as the local administrators group of this system. See the SnapDrive Installation and Administration Guide for more details about service account requirements." Below this, a note states: "Note: IBM V55 hardware provider registration also requires user account information." There are three input fields: "Account:" with the text "ITSO\Administrator" and an "Add..." button to its right; "Password:" with a masked password "*****"; and "Confirm Password:" with a masked password "*****". At the bottom left, it says "InstallShield". At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

Figure A-38 SnapDrive Service Credentials

8. Click **Finish**.

After the installation is done, no restart is needed in order to get the SnapDrive for Windows working. When accessing the Computer Management MMC, you will notice the SnapDrive snap-in and the iSCSI Management snap-in showing the already configured connections to the IBM System Storage N series storage system (see Figure A-39).

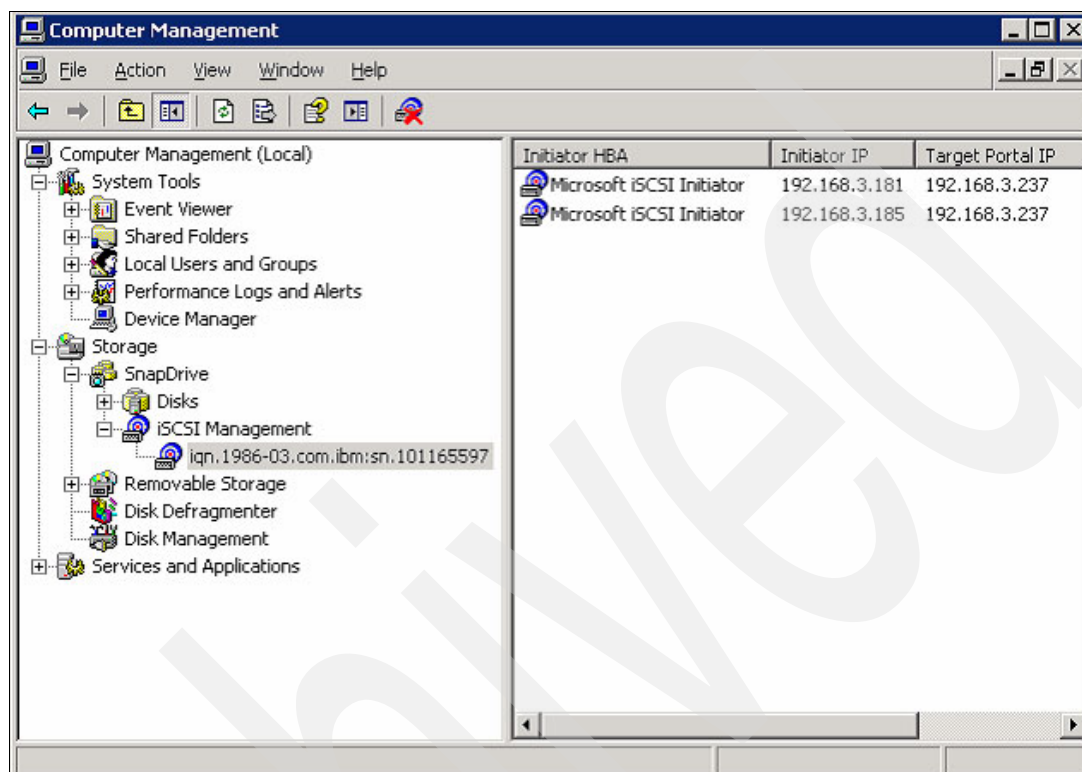


Figure A-39 Computer Management MMC with SnapDrive

Creating Virtual disks

Now that the Data ONTAP DSM for Windows MPIO, Microsoft iSCSI Initiator Software, and SnapDrive are installed, the disk drives can be added using the SnapDrive software.

A virtual disk created on the Windows server is in fact a LUN on the IBM System Storage N series storage system. You should not create LUNs using FilerView or the command line. Instead, you should use SnapDrive to create LUNs and virtual disks on the Windows server.

In order to create the disks from SnapDrive, we are assuming that:

- ▶ An aggregate has been created on the filer.
- ▶ A volume has been created on the aggregate.
- ▶ A CIFS share has been created mapping the path to the volume.
- ▶ iSCSI and networking infrastructures are in place and working.

These are the steps to create the LUN from the SnapDrive:

1. Access Computer Management MMC.
2. Expand **Storage** and then expand **SnapDrive**.
3. Right-click **Disks** and click **Create Disk**, as shown in Figure A-40 on page 453.

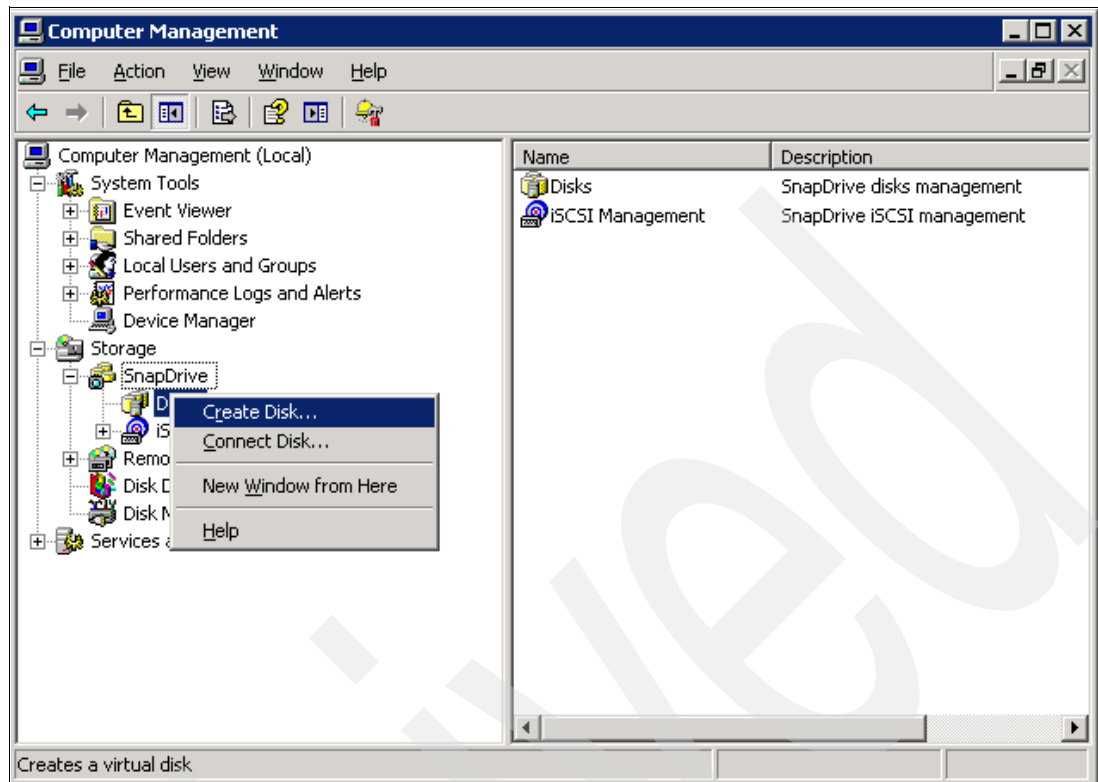


Figure A-40 Disk creation using SnapDrive

4. In the Create Disk welcome window (Figure A-41), click **Next**.

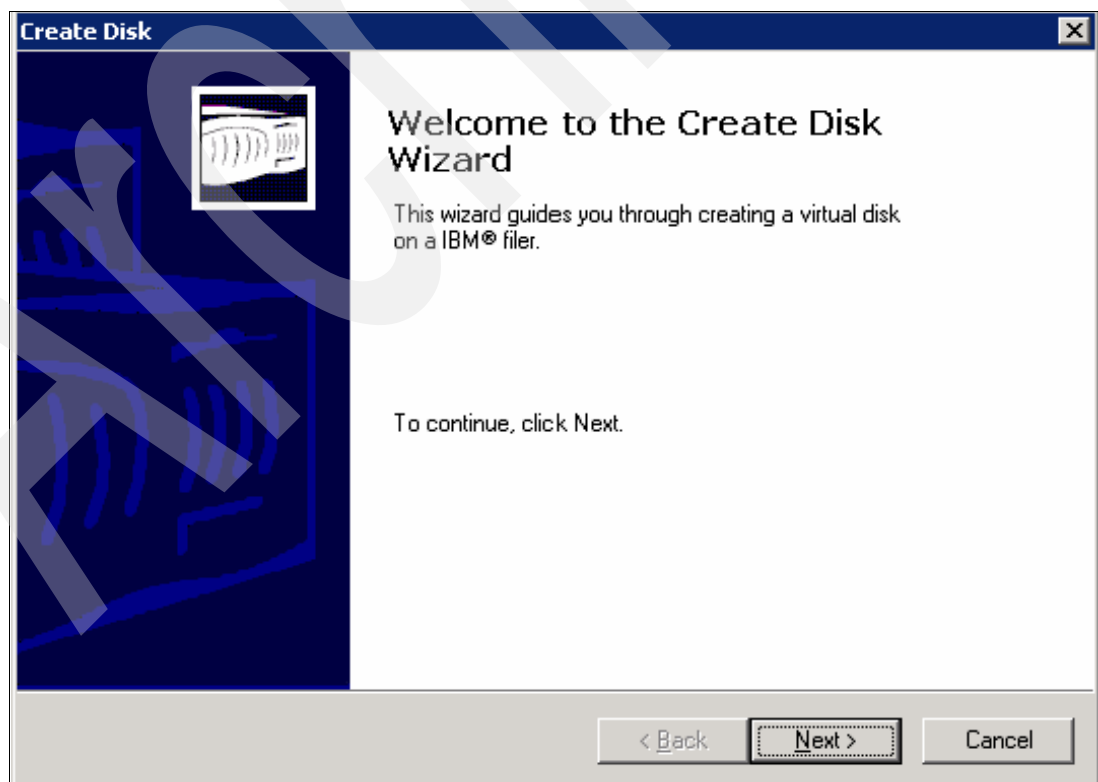
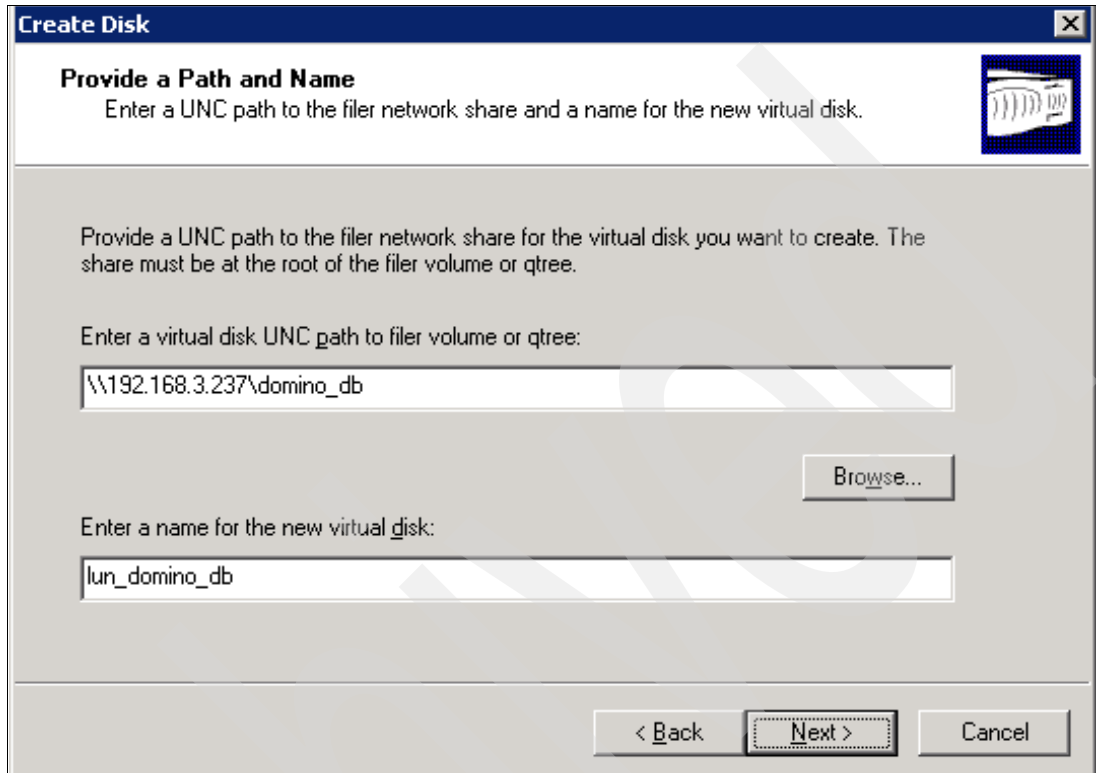


Figure A-41 Create Disk welcome window

5. In the Provide a Path and Name window (Figure A-42), enter the information needed. For the UNC path, use \\Filer IP\CIFS Share on the filer. Remember that this CIFS share points to the volume you created. For the name for the new virtual disk, type in the name you want to assign to the LUN that will be created. Click **Next**.



The screenshot shows a Windows-style dialog box titled "Create Disk". The main heading is "Provide a Path and Name". Below this, a subtitle reads: "Enter a UNC path to the filer network share and a name for the new virtual disk." There is a small icon of a disk in the top right corner. The main area contains two text input fields. The first field is labeled "Enter a virtual disk UNC path to filer volume or qtree:" and contains the text "\\192.168.3.237\domino_db". To the right of this field is a "Browse..." button. The second field is labeled "Enter a name for the new virtual disk:" and contains the text "lun_domino_db". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel". The "Next >" button is highlighted with a red border.

Create Disk

Provide a Path and Name
Enter a UNC path to the filer network share and a name for the new virtual disk.

Provide a UNC path to the filer network share for the virtual disk you want to create. The share must be at the root of the filer volume or qtree.

Enter a virtual disk UNC path to filer volume or qtree:

\\192.168.3.237\domino_db

Browse...

Enter a name for the new virtual disk:

lun_domino_db

< Back Next > Cancel

Figure A-42 Provide path and name window

6. In the Select a virtual disk type window (Figure A-43), select **Dedicated** if this disk will be accessed by only one server. Select **Shared** if this disk will be accessed by a Microsoft Cluster Service. Click **Next**.

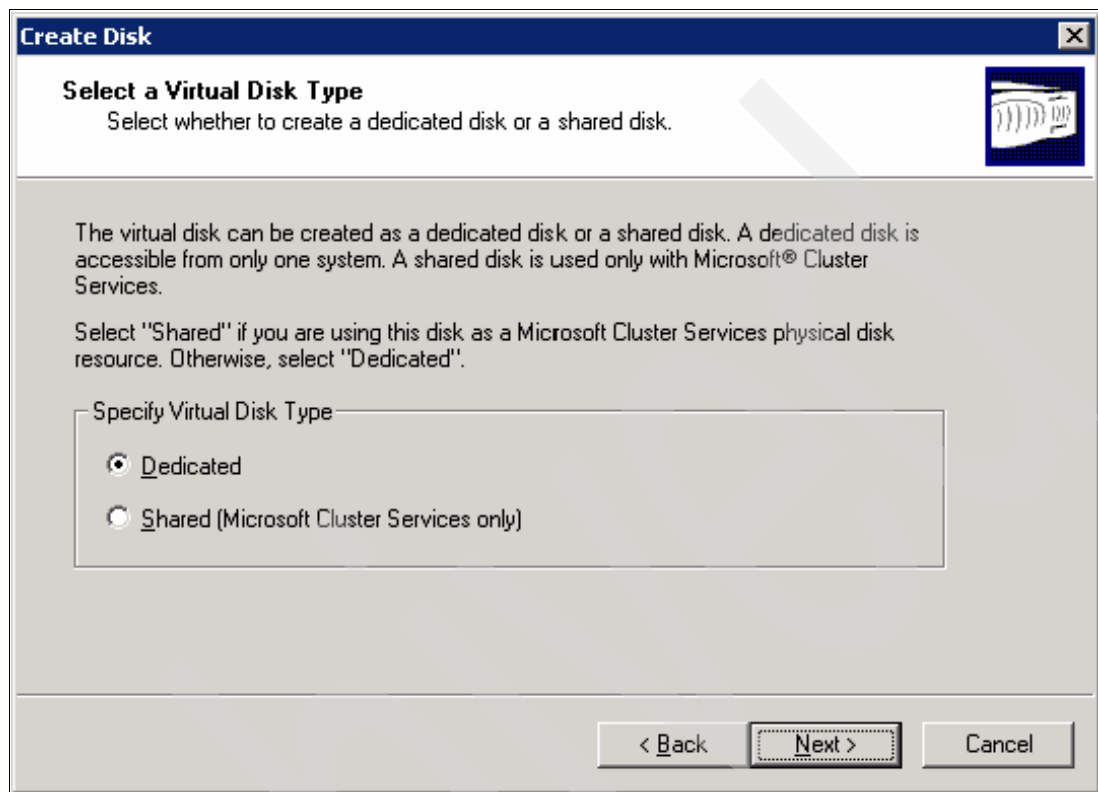


Figure A-43 Select a virtual disk type window

7. In the Select Virtual Disk Properties window (Figure A-44), select if you want to assign a drive letter (and which drive letter) for the disk being created or if you want to assign a Volume Mount Point. The next option will impact the size that will be available for the LUN creation. You need to select if you will reserve space for at least one snapshot of this LUN on the volume. Then, enter the size you want for this LUN. Click **Next**.

Create Disk

Select Virtual Disk Properties
Provide the drive letter and the size of the virtual disk to create.

Provide a Drive Letter or Volume Mount Point

☒ Assign a Drive Letter: L

☐ Use Volume Mount Point:

Do you want to limit the maximum disk size to accommodate at least one snapshot on the volume?

☐ Yes ☒ No

Enter a virtual disk size that is equal to or less than the maximum size, but greater than or equal to the minimum size.

Maximum Virtual Disk Size: 44 GB

Minimum Virtual Disk Size: 32 MB

Enter or Select Virtual Disk Size: 42 GB

< Back Next > Cancel

Figure A-44 Select Virtual Disk Properties window

8. In the Select Initiators window (Figure A-45 on page 457), select the initiators from the Available Initiators column on the left and click the arrow to move them to the Selected Initiators on the right. Because we are using iSCSI, the initiators will be listed as the Initiator node name on the Microsoft iSCSI Initiator software configuration. Click **Next**.

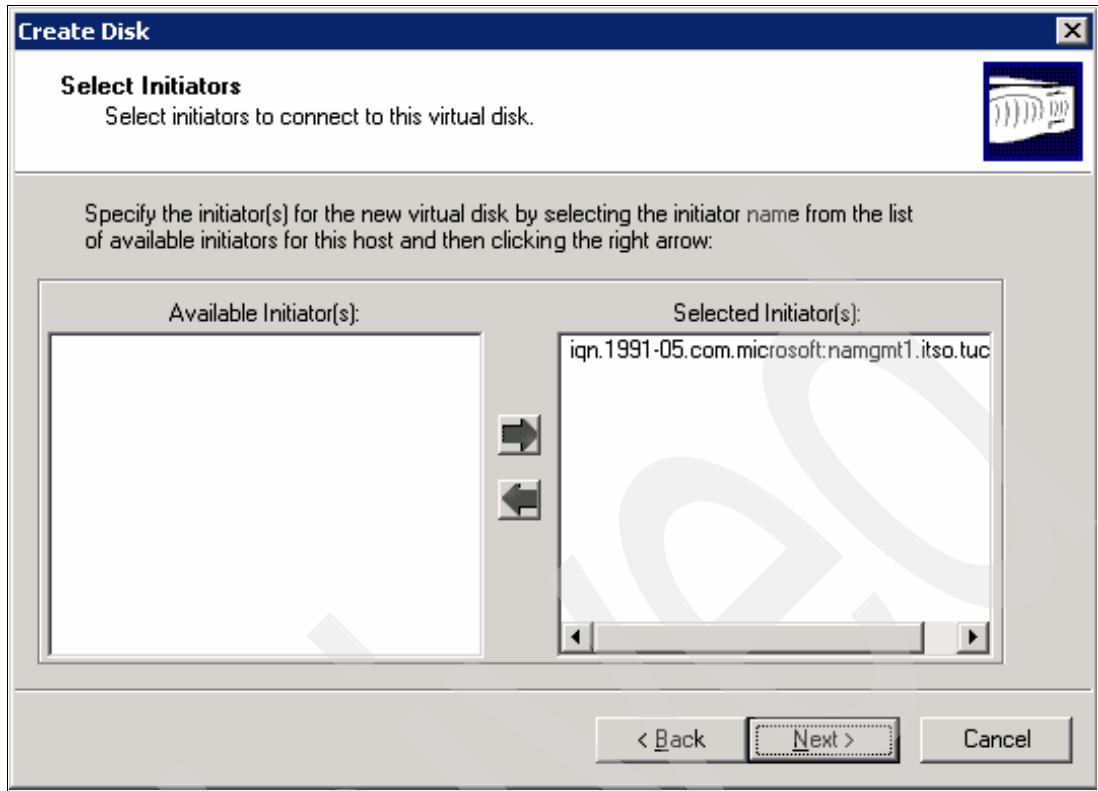


Figure A-45 Select Initiators window

9. In the Summary window, verify if all the information is right and click **Finish**. This will start the LUN creation process on the filer.
10. SnapDrive will format the drive and a Disks window will appear (Figure A-46). Click **OK**.




Figure A-46 Disks window

11. Now the disk is available for the Lotus Domino Server by using the drive letter you selected or by using the Volume Mount Point you created.

Important: When using iSCSI to map the LUNs from the IBM System Storage N series storage system to the Lotus Domino Server, you should guarantee that the Server service do not start *before* the Microsoft iSCSI Initiator service. If that happens, the shares on the LUNs will be unavailable until you restart the Server service.

To resolve this problem, create a dependency between the Server service and the Microsoft iSCSI Initiator service by following these steps:

1. Select **Start** → **Run** and type regedt32.exe to start the Registry Editor.
2. Navigate to the key HKLM\SYSTEM\CurrentControlSet\Services\LanManServer.
3. Select LanManServer and select **Edit** → **New** → **Multi-String Value**.
4. Type DependOnService as the value name. Notice the lack of space between the words and capitalization of letters.
5. Double-click the DependOnService value and type in MSiSCSI.
6. Click **OK**.
7. In the iSCSI Initiator Properties window, shown in Figure A-21 on page 436, click the **Bound Volumes/Devices** tab.
8. Click **Add** and type in the drive letters for all the LUN drives available to the IBM Lotus Domino Server. This configuration will guarantee that the Microsoft iSCSI Initiator service will only be started after all resources listed on this tab are online as well. Click **OK**.
9. Restart the server.

Note: When you create a virtual disk using SnapDrive, SnapDrive creates an igroup for the initiator and LUNs created for the initiator are mapped to this igroup automatically.

Installing and configuring a Lotus Domino server

The following steps are necessary to install and configure a Lotus Domino server that uses N series storage for its data and transaction logs:

1. Launch the Setup program for Lotus Domino server.
2. After accepting the license agreement, the next window will prompt you to enter the program file location, as shown in Figure A-47 on page 459. You can store the program file on the virtual disk or the local disk on the Windows server. For this write-up, we installed the Domino program files on a local disk (C:\Program files\lotus\domino).

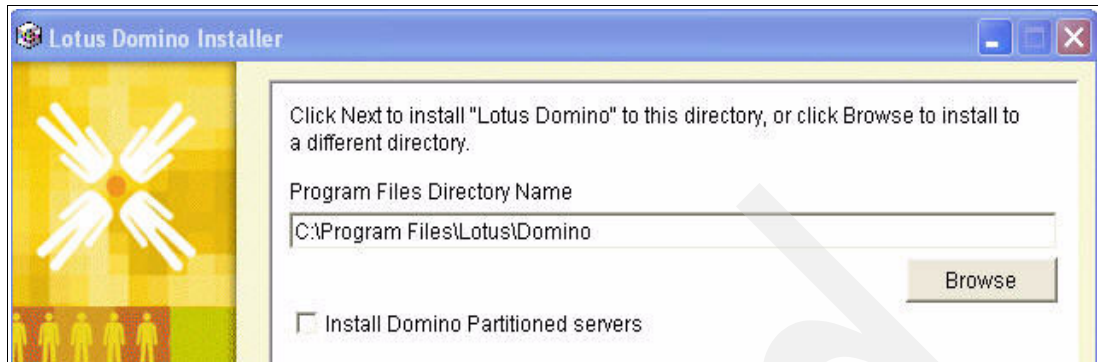


Figure A-47 Domino program file location

3. The next window will prompt you to enter the location for the Lotus Domino database data files, as shown in Figure A-48. You need to specify the network drive that corresponds to the virtual disk on the storage system for the Domino data folder (H:\lotus\domino\data).

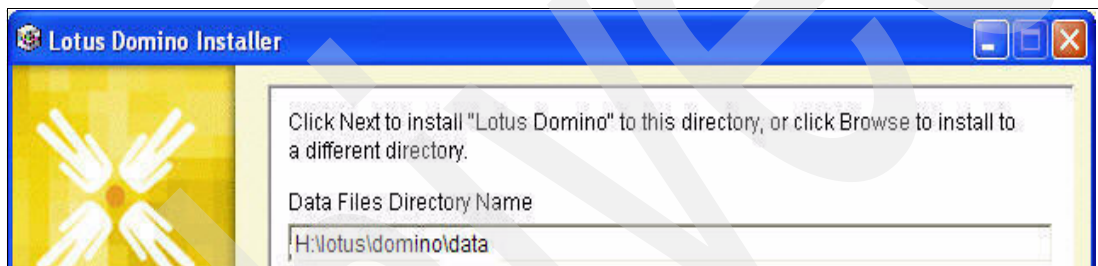


Figure A-48 Figure 16: Domino data files location

4. The next window prompts you to select the install type. Select the **Domino Enterprise Server**, as shown in Figure A-49, and click the **Next** button.

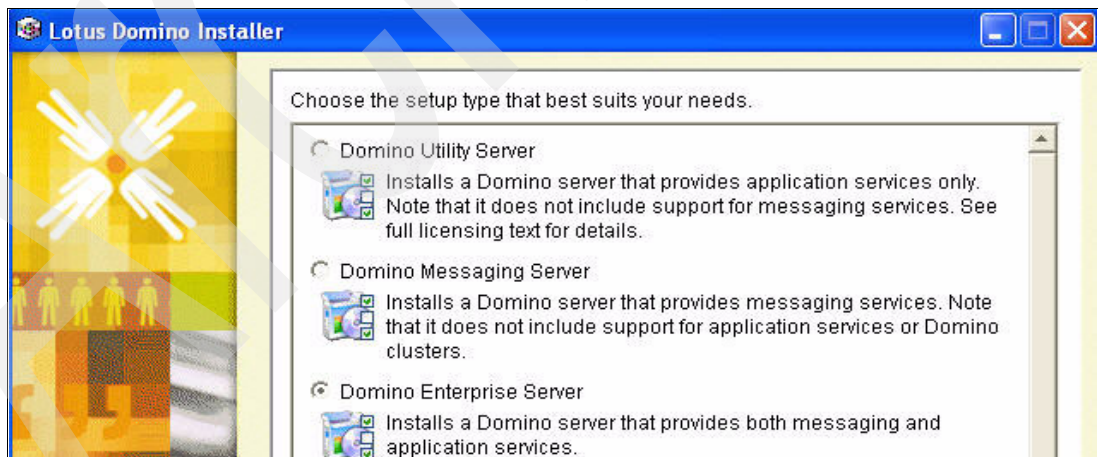


Figure A-49 Figure 17: Domino install type

5. The next window will present you with the installation summary, as shown in Figure A-50. After reviewing the installation summary, click **Next**.

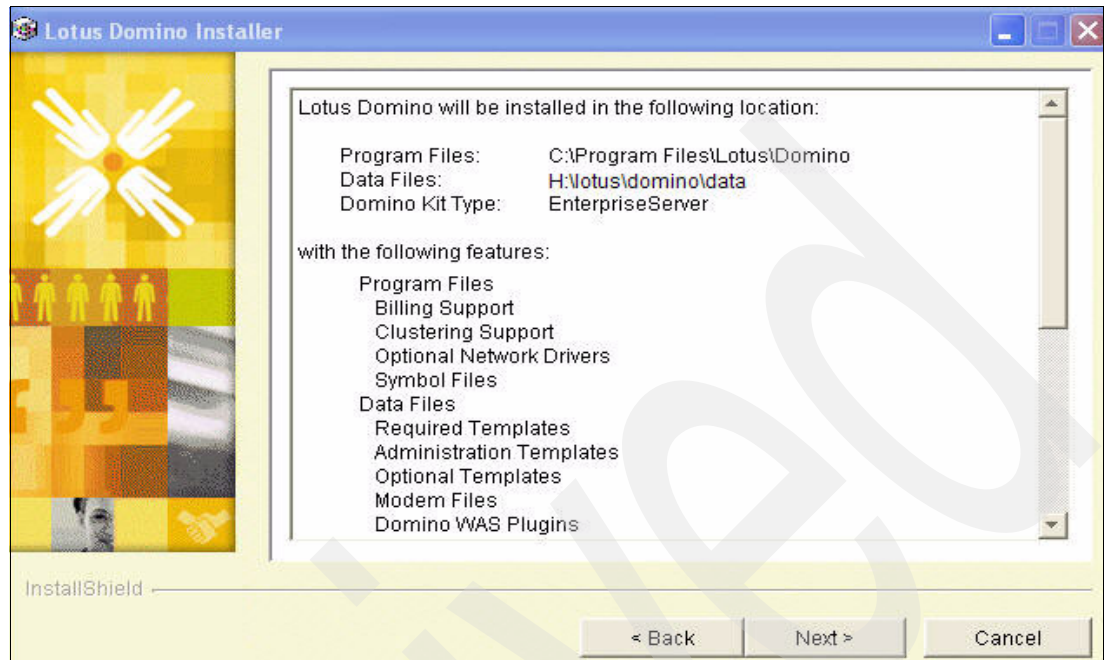


Figure A-50 Domino installation

6. When the Installation process is complete, click **Finish**
7. Next, start the Lotus Domino server setup process by selecting **Start** → **Program** → **Lotus Applications** → **Lotus Domino Server**.
8. In the next window, select **Start Domino as a Regular Application**, check the **Don't ask me again** check box, and click **Next**.
9. Click the **Next** button in the Welcome window.
10. In the next window, select the **Set up the first server or a stand alone server** radio button and click **OK**.
11. In the next window, specify the host details name and host description, as shown in Figure A-51 on page 461, and click the **Next** button.

Server setup

Provide a server name and title

You must provide a unique name for your new Domino server. Carefully choose the server name; you cannot easily change it later. By default, Setup recommends that you use the computer's host name as the server name.

Server name:
For example: Sales1

Optional: Provide a short title which describes the purpose or function of this server. (You can always change this information later in the Domino Directory)

Server title:
For example: Corporate Sales Server 1

☐ I want to use an existing server ID file:

Figure A-51 Domino first server configuration

12. The next window prompts you to specify the organization's name and the certificate's password. Enter the values for the organization's name and the certificate's password, as shown in Figure A-52, and click the **Next** button. If you desire to set up at the organization unit level, click the **Customize** button and proceed with the configuration.

Server setup

Choose your organization name

The organization name is usually your company name. It becomes part of each server and user name. Do not choose a long organization name. For example, instead of Acme Corporation, use Acme.

Organization name:
Minimum of 3 characters

This server's final name will be: fuj15/netapp
A typical user name will be: db2inst1/netapp

Organization Certifier password:
Minimum of 5 characters

Confirm password:

☐ I want to use an existing certifier ID file:
F:\Lotus\Domino\data\cert.id

To specify additional organization settings click Customize.

Figure A-52 Domino first server configuration

13. In the next window, enter the Domino domain name and click the **Next** button.
14. In the next window, specify the administration account details and click the **Next** button.
15. In the next step, you need to configure the internet services you desire for the Domino server (Figure A-53).

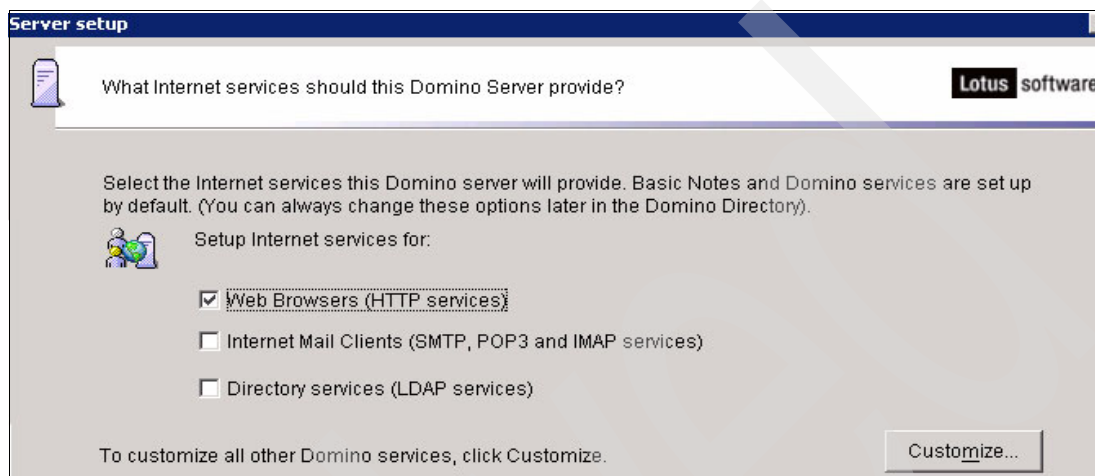


Figure A-53 Internet service configuration

16. By default, the installation will configure two network settings. You can click the **Customize** button and clear the port that will not be used. Figure A-54 illustrates keeping TCP/IP as the network setting and clearing the others.

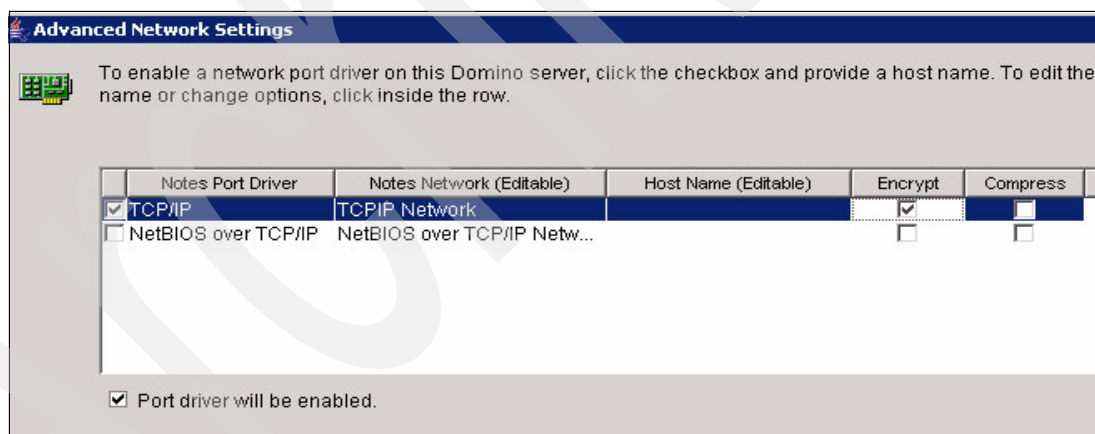


Figure A-54 Configure Network setting

17. After setting and reviewing the security settings, review the configuration summary in the next window and click the **Setup** button. When setup is complete, click the **Finish** button.

Enabling Lotus Domino Transaction Logging

Transaction logging is an optional feature in Lotus Domino, but based on integration work with Lotus Professional Services and close collaboration with Lotus Engineering, we recommend that you use the transaction logging feature in Domino. Transaction logging has many benefits, including increased performance for high Domino workloads. However, the enhanced data availability associated with much shorter recovery (consistency check) times

is the primary reason transaction logging should be seriously considered by all Domino server administrators. In order to enable transaction logging, you need to complete the following steps:

1. In the Domino Administrator, select the **Domino domain** tab.
2. Select **Server** → **Current Server Document** from the tree and click the **Configuration** tab.
3. Click the **Transactional Logging** tab.
4. Complete the fields shown in Table A-1, and then save the document.

Table A-1 Transaction logging fields

Field	Value
Transactional Logging	Choose Enabled . The default is Disabled.
Log path	The path name location of the transaction log. The default path name is \LOGDIR in the Domino data directory, although it is strongly recommended to store the log on a separate virtual disk residing on the IBM System Storage N series storage system. The separate device should have at least 1GB of disk space for the transaction log.
Maximum log space	The default is 192 MB. The maximum is 4096 MB (4 GB). Domino formats at least three and up to 64 log files, depending on the maximum log space allocated.
Use all available space on log device	Choose one: <ul style="list-style-type: none"> ► Yes to use all available space on the device for the transaction log. This is recommended if you use a separate device dedicated to storing the log. If you choose Yes, you do not need to enter a value in the "Maximum log space" field. ► No to use the default or specified value in the "Maximum log space" field.
Automatic fixup of corrupt databases	Choose one: <ul style="list-style-type: none"> ► Enabled (default). If a database is corrupt and Domino cannot use the transaction log to recover it, Domino runs the Fixup task, assigns a new DBIID, and notifies the administrator that a new database backup is required. ► Disabled. Domino does not run the Fixup task automatically and notifies the administrator to run the Fixup task with the -J parameter on corrupt logged databases.

Field	Value
Runtime/ Restart performance	<p>This field controls how often Domino records a recovery checkpoint in the transaction log, which affects server performance. To record a recovery checkpoint, Domino evaluates each active logged database to determine how many transactions would be necessary to recover each database after a system failure. When Domino completes this evaluation, it:</p> <ul style="list-style-type: none"> ► Creates a recovery checkpoint record in the transaction log, listing each open database and the starting point transaction needed for recovery. ► Forces database changes to be saved to disk if they have not been saved already. <p>Choose one:</p> <ul style="list-style-type: none"> ► Standard (default and recommended). Checkpoints occur regularly. ► Favor runtime. Domino records fewer checkpoints, which requires fewer system resources and improves server runtime performance. ► Favor restart recovery time. Domino records more checkpoints, which improves restart recovery time because fewer transactions are required for recovery.
Logging style	<p>Choose one:</p> <ul style="list-style-type: none"> ► Circular (default) to continuously re-use the log files and overwrite old transactions. You are limited to restoring only the transactions stored in the transaction log. ► Archive (recommended) to not re-use the log files until they are archived. A log file can be archived when it is inactive, which means that it does not contain any transactions necessary for a restart recovery.

Migrating an existing Lotus Domino server from local disk to a IBM System Storage N series storage system

A Domino administrator can migrate an existing Domino server from local disk to a IBM System Storage N series storage system copying the Domino data directory and transaction log files to the virtual disks created on the storage system. This section describes the manual process to migrate the Domino server.

In order to migrate the Lotus Domino server to virtual disks residing on the IBM System Storage N series storage system, you need to complete the following steps:

1. Stop the Domino server that is being migrated.
2. Make a complete backup of the original Lotus Domino server.
3. Identify the location of the program directory by looking at the line NotesProgram = in the NOTES.INI file.

4. Create the \Lotus\domino\data and \Lotus\domino\logs directories on the virtual disks you create for the Domino server.
5. Move the old \Lotus\domino\data old directories (the location of the data directory can be identified from the line `Directory=` in the `notes.ini`) to the new directory \Lotus\domino\data on the virtual disk.
6. Move the old \Lotus\domino\logs old directories (the location of the data directory can be identified from the line `Directory=` in the `notes.ini`) to the new directory \Lotus\domino\logs on the virtual disk.
7. The `notes.ini` file should have references to the new data and transaction log file locations. Go to the Domino directory and update the parameters shown in Table A-2 in the `notes.ini` file.

Table A-2 Changes for Domino migration

Old value	New value
<code>Directory=C:\Lotus\Domino\data</code>	<code>Directory=C:\Lotus\Domino\data</code>
<code>KeyFilename=C:\Lotus\Domino\data\server.id</code>	<code>KeyFilename=F:\Lotus\Domino\data\server.id</code>
<code>CertifierIDFile=C:\Lotus\Domino\data\cert.id</code>	<code>CertifierIDFile=F:\Lotus\Domino\data\cert.id</code>
<code>TRANSLOG_Path=C:\lotus\domino\logs</code>	<code>TRANSLOG_Path=G:\lotus\domino\logs</code>
<code>Previous_TRANSLOG_Path=C:\lotus\domino\logs\</code>	<code>Previous_TRANSLOG_Path=G:\lotus\domino\logs\</code>

In Table A-2 as an example, the Lotus Domino data and transaction log files are moved from directories on a local disk to the directories on virtual Disks H and G. Only storage is changed, not the Domino server, therefore the program files are not moved and kept in the same directory.

8. Start the Lotus Domino server to ensure that it was migrated correctly and test the migration by sending some e-mails from a Notes e-mail client. Also, open a couple of the databases and make sure that there are no issues.

Note: Before you start migrating your Lotus Domino server from local disks to virtual disks residing on an IBM System Storage N series storage system, you should pay attention to the following additional migration information:

- ▶ Make sure the Server document is updated if the Network Configuration section changes. Some customers may have a server IP address in the Net Address field. This must be updated if the new box being installed has a different IP address.
- ▶ Check for Directory Links when performing an upgrade/move.
- ▶ If the server name is to be changed, ensure that all encrypted databases are decrypted before copying the files to the new server.
- ▶ When moving servers between different operating system platforms, use FTP to relocate the databases or mail files. This will ensure the code pages are not corrupted. In some instances, using the OS copy has caused some database corruption.
- ▶ If moving from one machine to another and the drive mapping is different (for example, from C drive to D drive), change the following parameters in the NOTES.INI:

```
Directory=d:\Lotus\Domino\Data  
NotesProgram=d:\Lotus\Domino
```

Perform a search in the server's NOTES.INI file to ensure that these parameters are changed appropriately. Additionally, perform a "find" (Ctrl + F) on the old drive references and change appropriately.

If the server's name is changed when it is moved to the new hardware, the administrator can send users a mail message with a button having LotusScript® behind it to change their Location documents to reflect the new name of the server.

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 publications

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 468. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM N Series Storage Systems in a Microsoft Windows Environment*, REDP-4083
- ▶ *IBM System Storage N series Fibre Channel and iSCSI Configuration Guide*, SG24-7496
- ▶ *N Series SnapManager with Microsoft Exchange*, REDP-4160
- ▶ *Using the IBM System Storage N Series with IBM Tivoli Storage Manager*, SG24-7243

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM System Storage N series Data ONTAP 7.2 File Access and Protocols Management Guide*, GC26-7965
- ▶ *IBM System Storage N series Data ONTAP 7.2 Software Setup Guide*, GC26-7975
- ▶ *IBM System Storage N series Data ONTAP DSM 3.0 for Windows MPIO Installation and Administration Guide*, GC27-2053

Online resources

These Web sites are also relevant as further information sources:

- ▶ Support for Network attached storage (NAS) and iSCSI
<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/storageselectproduct?brandind=5000029&familyind=0&oldfamily=0&continue.x=14&continue.y=12>
- ▶ Support for Data ONTAP
<https://www-304.ibm.com/jct01004c/systems/support/myview/supportsite.wss/supportresources?brandind=5000029&familyind=5329797&taskind=1>
- ▶ Support for DSM for Windows MPIO
<https://www-304.ibm.com/jct01004c/systems/support/myview/supportsite.wss/supportresources?brandind=5000029&familyind=5355915&taskind=1>

How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks publications, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy IBM Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

2864-A10 24
2864-A20 24
2865-A10 24
2865-A20 24
2866-A10 29
2866-A20 29
2867-A10 29
2867-A20 29
2868 - A10 24
2868 -A20 24

A

A10 17, 22
Access Control List 57
Add Target Portal window 438
-addlun 270, 293
AdminGroup 431
administration account 462
administrator account 300
Advanced Settings window 439, 443
aggregate 83, 96, 125–128, 136–137, 177, 209, 287,
412, 419, 424, 452
 volumes
 introduction 3
aggregate failure 271
Aggregate Name 130, 420
aggregates 6, 127–128, 419
aggregates. 137
AIX 71, 107, 262, 292
application design elements 53
application testing 163, 395
Archival 416
archiving 399
ASCII 20
asynchronous mode 160
authenticated 148
Automatic method selection 423
-autorename 401
AutoSupport 264
availability 3

B

Backup 301
backup 301, 413, 464
Backup data 158
backup process 415
Backup Wizard 301
backup/recovery 3
backups 123–124
bezel 32
Block I/O 105–106, 145
block level 3

BM Lotus Domino 7 57
BM System Storage N series Gateway 40
Boot device 268
business continuity 66

C

cache 120, 123
cache percentage 122
capacity 120, 288
Cascading Mirrors 161
CIFS 105, 177, 452
CIFS share 209
circular logging 416
CLI 89
client/server 54
clone 164, 166
cloned volumes 163
clones 162
cloning 163
cluster configuration 56
Cluster Failover 21
Cluster Service 213, 455
clustering technology 56
Common Internet File System (CIFS) 144
compact flash 20
compatibility matrix 156
compliance 63
Computer Management MMC 209
concurrent users 143
Configuration Report window 238
Configuration Status 237
console port 34
copper 20
corporate data 161
corruption 410
CPU 19
CPU tray 20
Create the Lotusdominodata and Lotusdominologs direc-
tories on the 465

D

DAS 62–64
data availability 5
data directory 465
data integrity 5
data management 5
data migration 395
Data Mining 164
Data ONTAP 5, 72, 146–147, 154–155, 159, 171,
242–243, 299, 385–387, 392, 398, 414, 430, 452
Data ONTAP 7.2 417
Data ONTAP DSM 171
Data ONTAP DSM for Windows 177
Data ONTAP DSM for Windows MPIO 156, 209

- Data ONTAP FilerView 85, 394
- Data ONTAP NDMP 397
- Data ONTAP, 397
- Data ONTAP® 5
- data protection 40
- database 58
- database LUN 123–125
- Database LUN size 122, 124
- database maintenance 120
- database placement 272
- database verification 301
- Database Volume 124–125
- Database Volume size 124
- DBMS 105
- default-transport 264
- Destination Folder window 450
- destination volume 396
- Diagnostic 21
- disaster recovery 62, 111, 161, 395
- disk configurations 128, 419
- disk sanitization 5
- Disk Selection window 422
- disk space 59, 155
- Distributed 58
- DMP backup applications 397
- domain 418
- domain users 148
- DomainName 418
- Domino 72
- Domino Administrator 463
- Domino administrator 464
- Domino Directory 57
- Domino domain 462
- Domino server 464
- DR 63
- drive flexibility 14
- dual-node 32
- dump command 397–398
- dynamic 411
- dynamic pool 158

E

- e-mail 62
- e-mail archival 67
- e-mails 120
- encryption 57
- Enterprise Server 53
- ERP 161
- Eseutil.exe 226
- ESH2 21
- Ethernet 20
- Ethernet ports 35
- Exchange Server 227
- Express 53
- external storage 399

F

- fabric 100
- FC 177

- FCP 5, 97, 105, 107, 109, 156, 170, 187, 240, 408, 413, 417
- FCP connectivity 159
- FCP Host 411
- FCP protocol 156, 264
- Fibre Channel 20, 99, 170, 240
- Fibre Channel Protocol (FCP) 144
- Fibre Channel SAN 152
- Fibre Channel switch model 241
- Fibre Channel switches 241
- Fibre Channel zone 149
- File I/O 105
- file system 270
- File Transfer Protocol (FTP) 144
- file_spec 386
- FilerView 128–129, 137–138, 142–143, 194, 280, 419, 424–425, 428
- filervol option 266
- Finish window 434
- fix-up 415
- FlexClone 111, 163–164, 395
- FlexClone feature 163, 165
- FlexClone volume 165–166
- FlexClone volumes 164–165
- flexible storage 3
- flexible volume 91, 136–137, 166
- flexible volumes 396
- FlexVol 81, 136, 139, 428
- FlexVol Parameters 141
- Force Group 428
- Fractional space reservation 80
- frequency 300
- fs argument 266
- FTP 105

G

- Gateways 40
- Gigabit 412
- group 147–148
- groups 147
- growby 293
- growth 120

H

- Hard zoning 150
- Hardware 21
- HBA 82, 101, 106, 241, 409, 431
- High Availability 414
- High Reliability 414
- Host bus adapter (HBA) 241
- host details name 460
- Host Type 98
- host volume group, 266
- hot site 161
- hot swappable 33
- HTTP 105
- Hyper Text Transfer Protocol (HTTP) 144

I

- I/O 272, 387, 411
- I/O performance 120
- IBM 32
- IBM Lotus Domino 52–54, 57–58, 158, 292, 410
- IBM Lotus Domino 7. 54
- IBM Lotus Domino Access for Microsoft Outlook 54
- IBM Lotus Domino application files 59
- IBM Lotus Domino data directories 59
- IBM Lotus Domino database 272
- IBM Lotus Domino database volume 270
- IBM Lotus Domino program files 59
- IBM Lotus Domino Server 444
- IBM Lotus Domino server 241, 407
- IBM Lotus Domino transactional log files 59
- IBM Lotus Domino user database files 59
- IBM Lotus Domino Web Access 54
- IBM Lotus Domino WebMail 54
- IBM Lotus Notes 54
- IBM N series 13
 - Gateway versus IBM N series storage systems 5
 - hardware 4
 - hardware quick reference
 - A models 6
 - A& G models 7
 - standard software features 8
 - storage systems A models 13
- IBM N series Gateway 5, 39
- IBM N series storage system 126, 132
- IBM N series storage systems 13
- IBM Storage System N series 3
- IBM System Storage 16
 - an introduction to N3700 15, 24, 29
- IBM System Storage N series 3, 14, 119, 159, 279, 409
- IBM System Storage N series Gateways 38
- IBM System Storage N series SnapDrive 411
- IBM System Storage N series SnapDrive for Linux 259
- IBM System Storage N series SnapDrive for Windows 4.2.1 411
- license key 417
- ID file 57
- iGroup 97
- infrastructure 24, 188
- Initiator 184
- Initiator Group 186–187
- Initiator Node Name 193, 436
- Initiator Service 190, 433
- initiators 183, 215
- Install and Update Software 262
- Install Software 262
- Install Software menu 262
- install type. 459
- installation 263
- Installation Options 190
- installation summary 460
- internet 54
- Internet SCSI Protocol (iSCSI) 144
- Internet services 462
- interoperability 241
- IP 63, 411

- IP address 161, 440, 444
- IP connectivity 159
- iSCSI 97, 105, 109, 145, 187, 209, 408, 452
- iSCSI adapters 188
- iSCSI hardware 188, 409
- iSCSI Initiator 198
- iSCSI Initiator Discovery window 441
- iSCSI Initiator Properties 201, 444
- iSCSI protocol 156
- iSCSI storage devices 198, 442
- iSCSI target 430
- iSCSI, 417

L

- latency 411
- LDAP 54
- LDServerIP 418
- LDServerName 418
- LEDs 37
- license 53, 107, 111, 156, 173, 417
- license agreement 191, 203, 434, 446, 458
- License Agreement window 219
- License key 174
- License key window 204, 447
- Linear 416
- Linux 107, 259, 399
- Linux root file system 272
- Load Balance 414
- local administrator group 431
- LockVault Compliance 5
- Logical Unit Numbers (LUNs) 144
- logical volume 270
- Lotus Applications 460
- Lotus Domino 52, 54, 71–73, 91, 107–108, 292, 294–295, 389, 399, 408, 412–413, 415, 430
- Lotus Domino 7.0.2 408
- Lotus Domino data 413, 465
- Lotus Domino database 267, 273, 395, 399, 402, 417–418
- Lotus Domino database data 459
- Lotus Domino databases 389
- Lotus Domino Enterprise Server 53
- Lotus Domino Express 53
- Lotus Domino LUNs 398
- Lotus Domino Messaging Server 53
- Lotus Domino program files 274
- Lotus Domino Server 438, 457
- Lotus Domino server 57, 59, 294, 383, 412, 458, 464–465
- Lotus Domino Server I 53
- Lotus Domino server setup 460
- Lotus Domino software 408
- Lotus Domino transactional log 273
- Lotus Domino transactional log Volume 394
- Lotus Domino transactional logging path 274
- Lotus Domino transactional logs 274
- Lotus Domino Utility Server 53
- Lotus Domino. 51
- Lotus Domino's serialized log writes 415
- lspp -l 264

- LUN 6, 76, 122, 124–125, 143, 177, 184, 209, 237, 266, 292, 389–391, 452
- LUN clone 388–389
- lun clone 390
- LUN cloning 388–389
- LUN creation 216
- LUN Map 186–187
- LUN mapping 98
- lun show 390
- LUN size 124
- LUN Snapshot 388
- LUNs 126, 144, 152, 187, 241, 265, 269, 272, 292, 396
- LUNs, 264
- lvextend 293
- lvextend command 270
- LVM 391
- lvm lvdisplay 270

M

- Mail enabled 58
- Mailbox 120, 122–123
- mail-enable 54
- Microsoft Exchange 61–62, 120, 126, 137, 156, 184
- Microsoft Exchange 2003 62
- Microsoft Exchange architecture 68
- Microsoft Exchange backup 62
- Microsoft Exchange databases 120–121, 166
- Microsoft Exchange deployment 120
- Microsoft Exchange environment 63, 145
- Microsoft Exchange message store 64
- Microsoft Exchange Operations 164
- Microsoft Exchange Server 194, 200–201, 216, 224
- Microsoft Exchange server 143, 145, 153
- Microsoft Exchange Server 2003 123–125
- Microsoft Exchange server infrastructure 145
- Microsoft Exchange Server software installation 217
- Microsoft Exchange server's 137
- Microsoft Exchange servers 120
- Microsoft Exchange transaction logs 121
- Microsoft Exchange. 61
- Microsoft iSCSI Initiator 192, 196, 200, 435, 443
- Microsoft iSCSI Software Initiator 189, 431–432
- Microsoft iSCSI Software Initiator software installation 189
- Microsoft MPIO Multipathing Support for iSCSI 190, 433
- Microsoft Windows 63
- Mirrored aggregates 128
- mirrored backup 301
- mirroring 392
- mirroring synchronously 392
- MMC 177, 223, 452
- mobile devices 120
- Model 26
- Mount Point 428
- MPIO 103, 153–154
- Multi-disk 28
- multipath technology 158, 411
- multipathing 152
- multipathing-type 264

N

- N series 63, 65, 147–148, 156, 161, 166, 169, 239, 243, 299, 385, 396, 399, 408
- N series server 399
- N series SnapDrive feature installation 260
- N series storage 259, 411
- N series storage operating System Data ONTAP 384
- N series storage server 159
- N series storage system 72, 136, 143–144, 153, 240–241, 244, 264, 412, 417, 430
- N3700 4, 16–20
- N5000 5, 24, 26, 30
- N5000 series 4
- N5200 24, 28
- N5500 24, 27
- N7000 29, 31, 36
- N7000 series 4
- N7600 29
- N7800 29, 36
- Name resolution infrastructure 159
- NAS 5
 - IBM TotalStorage NAS 13
- NDMP 397, 413
- NDMP backup 397
- near-line Feature 5
- near-line storage 14
- network 259
- Network File System (NFS) 144
- NFS 39, 105, 109
- non-disruptively 3
- Notes database 53
- Notes Storage File 58
- NOTES.INI 464
- notes.ini 465
- NSF 58, 74
- NVRAM 30
- NVRAM6 36

O

- Object oriented 58
- on-the-fly 158
- operating budget 158
- operating system 241, 411
- optical 20
- optional software 9
- options snapmirror.access 393
- organization name 461

P

- Parallel Processing 164
- patches 159
- PCI 36
- PCI slots 34
- PCI-Express 36
- Performance 127
- performance 137, 410, 415
- point-in-time 399
- port based 149
- port zoning 99, 151

- prerequisites 161
- privileges 146
- production data 164
- production volume 161
- program file placement 272
- protect data 158
- protocol 105, 170, 240
- provisioning 63
- Public Folders 120

Q

- quota limits 120

R

- RAID 6, 74, 127, 131, 421
- RAID 4 5, 90
- RAID DP 86
- RAID Group 86
- RAID-DP 5
- RBAC 146
- real-time 58
- reboot 192, 222
- Red Hat Linux Enterprise Server 272
- Redbooks Web site 468
 - Contact us xiii
- redundant 20
- redundant cooling 20
- Remote Shell access 396
- Replicated mirrors 161
- reports 161
- requencies 300
- resize operation 269
- resize2fs 271
- restore command 404
- Role-based access controls, or RBAC 242
- root 261
- Root disk groups 268
- root volume 140, 396, 412
- RPM 85
- rpm command 259
- rsh 418
- rsh (remote shell) 418
- rst1m 398

S

- s Snapshot 403
- SAN 5, 39, 64, 83, 102, 241
- SAN, 152
- SATA 14, 24, 29
- SCSI 5
- sdcli snap create 300
- security 411
- security settings 462
- security style 429
- Select Initiators window 456
- Semi-Synchronous Mode 160
- server 120
- serviceability 5

- Share 143
- Share Description 428
- Share Name 428
- Shared 58, 181
- single fabric 151–152
- Single instance 122
- Single Mailbox Recovery for Microsoft Exchange (SMBR) 64
- single user databases 399
- SME 217, 221, 224, 234–235
- SME configuration 228
- SMIT 261–262
- SMIT log 264
- SMTP 54, 122, 125, 232
- snap connect 389
- snap disconnect 296
- snap reserve 299, 430
- snap reserve option 299
- snap reserve vol_name 299
- snap restore 404
- snap restore -t 403
- SnapDrive 97, 103–104, 107–109, 146–147, 156–159, 172, 177, 184, 202, 205, 209, 241–243, 264, 292, 297, 301, 386, 408, 410–411, 444–445, 452
- SnapDrive for Linux 399
- SnapDrive for Linux and UNIX 265, 269, 292, 296, 383, 387
- SnapDrive for Microsoft Windows 156
- SnapDrive for UNIX 262, 264, 267–269, 292, 295, 387, 401
- SnapDrive for Window 177, 452
- SnapDrive installation 172, 445
- SnapDrive MMC 158
- SnapDrive on AIX server 261
- SnapDrive on Red Hat Linux 259
- SnapDrive Service 176, 451
- SnapDrive Service Credentials 208
- snapdrive snap create 267–268, 386
- snapdrive snap create file_spec 386
- snapdrive snap delete 297
- snapdrive snap list 388
- snapdrive snap restore 298, 401
- snapdrive snap restore file_spec 401
- SnapDrive SnapShot copies 300
- snapdrive storage connect 391
- snapdrive storage create 266
- snapdrive storage resize 269, 271, 292
- snapdriveadmin 244
- SnapInfo Directory 123–124, 233
- SnapInfo Directory size 123
- SnapManager 121, 301
- SnapManager for Exchange 166, 217–218
- SnapManager for Exchange (SME), 64
- SnapManager for Microsoft Exchange 166, 220
- SnapManager for Microsoft Exchange, 223
- SnapManager Server 222
- SnapManager version 4.0 166
- SnapMirror 110, 156, 161, 301, 392, 417
- SnapMirror backup 393
- SnapMirror destination 392

- SnapMirror feature. 160
- snapmirror initialize 394
- SnapMirror license 393
- snapmirror quiesce 395
- SnapMirror replication 301
- SnapMirror service 394
- SnapMirror source 397
- snapmirror status 394
- SnapMirror technology 66
- SnapMirror volume 161
- SnapMirror volume replication 161
- snapname 268, 298, 388, 401–402
- SnapRestore 402–403, 417
- SnapShot 77, 155
- Snapshot 64, 104, 137, 154–155, 158, 163, 268, 294–295, 297–299, 385, 387, 389–390, 411, 430
- SnapShot copies 155, 158, 267, 296–298, 300, 384–386, 401–402, 411
- SnapShot copy 267–268, 385–387, 390, 400, 404
- snapshot reserve space 299
- snapshot snap rename 297
- SnapShots 79, 269, 295, 300, 388, 396
- SnapVault 162
- SnapVault feature 162
- soft zoning 100, 150
- Software Initiator 190, 433
- Software Installation and Maintenance 261
- software package 262
- software quick reference 11
- space guarantee 93
- SSH 243
- stand-alone server 460
- storage 124
- storage access media 411
- storage capacity 411
- storage expansion 20
- Storage Group 122–123
- storage infrastructure 120
- storage management 40, 159, 411
- storage partitioning 271
- storage provisioning 40
- Storage requirement 414
- storage segmentation 59
- storage space 125
- storage system 20, 241, 411
- Storage system licenses 159
- storage systems 5, 40
- storage utilization 40
- Summary window 424, 428, 457
- synchronous mode 160
- synchronous SnapMirror 392–393
- synchronous SnapMirror feature 160
- Synchronous SnapMirror mode 160
- synchronously 160
- SyncMirror 86, 104
- System Deployment 164
- system initialization 5
- System Response Time 414
- System Storage N series
 - introduction 3

T

- t vol 403
- tape 16, 64
- tape backup 161
- Target Portal 195
- Target window 199, 442
- TCP/IP 462
- thin server 13
- third-party 120
- Tivoli Storage Manager 77
- TOE 106
- Traditional Volumes 91, 136
- transaction log 230, 413
- transaction log files 231
- Transaction Log LUN size 125
- transaction logging 415, 462–463
- Transaction Logs 124–125
- transaction logs 123–125, 416
- Transaction Logs LUN 123–124
- Transaction Logs LUN size 123
- Transaction Logs LUN sizing 123
- transactional log 58, 273
- transactional log files 59
- transactional log placement 272
- transactional logging 77
- Transactional Logging tab 463

U

- UNC 180, 454
- UNC path 212
- Unmirrored aggregates 127
- User ID files 57
- useradmin 147–148
- useradmin command 243–244
- useradmin user add 146, 242
- UserName 418, 431
- utilization 39

V

- Verification Server 226
- vg 388, 401–402
 - 268
- vg db_SdDg 293
- virtual disk 182, 410, 452
- Virtual Disk Properties 214, 456
- virtualization 38
- vol copy 396
- vol copy start 396
- vol create 96
- vol status 97
- VolName 429–430
- Volume 92, 137, 177, 424, 428, 452
- Volume cloning 163
- Volume copy 392
- volume group 270
- Volume name 424
- Volume Parameters window 427
- volume SnapMirror 161
- Volume Type Selection window 426

volume(s) 429
volume_name 403
volumes 428

W

WAFL 5, 83, 104
warm backup site 161
Web-based Distributed Authoring and Versioning (Web-DAV) 144
Welcome panel 460
Windows hosts 143
Windows server 417–418, 431
Wizard 225, 300
World Wide Node Name zoning 153
WorldWide Port Name 150
WORM 67
WWN spoofing 99
WWPNs 150, 185

Z

zeroing 88
zoning 99, 150, 152
zoning configuration 150

Archived



Redbooks

Using the IBM System Storage N series with Mail Servers

(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages



Using the IBM System Storage N series with Mail Servers



Using N series with Lotus Domino

Using N series with Microsoft Exchange

Using N series software with mail servers

This IBM Redbooks publication gives mail database administrators a introduction to the IBM System Storage N series. It also gives both storage administrators and mail database administrators the information to prepare, install, and administer the IBM System Storage N series when used in conjunction with Microsoft Exchange and IBM Lotus Domino. Often the lines of responsibility with regards to these tasks are blurred or crossed over, so some sections of this book may be redundant for experienced N series administrators but of importance to mail database administrators. Conversely, some sections that may be redundant for mail database administrators may help storage administrators.

**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

SG24-7483-00

ISBN 0738486574