



# IBM System Storage N series and Digital Video Surveillance

Meeting the Digital Video Surveillance  
Storage Capacity demands with N series

Protecting Digital Video Data  
with N series

Implementing Digital Video  
Surveillance with N series



Alex Osuna  
Haijiang Sha  
Marc Tu-Duy-Khiem  
Marco Aurélio de Mello Santos





International Technical Support Organization

**IBM System Storage N series and Digital Video  
Surveillance**

October 2008

Archived

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

## **First Edition (October 2008)**

This edition applies to Data ONTAP Version 7.1 and above

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

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



# Contents

<b>Notices</b> .....	ix
Trademarks .....	x
<b>Preface</b> .....	xi
The team that wrote this book .....	xi
Become a published author .....	xii
Comments welcome .....	xiii
<b>Chapter 1. Introduction to IBM System Storage N series</b> .....	1
1.1 IBM N series hardware .....	2
1.2 Comparing N series Gateway to N series storage systems .....	3
1.2.1 IBM N series A models hardware quick reference .....	4
1.2.2 IBM N series A and G models hardware quick reference .....	5
1.3 The IBM N series standard software features .....	7
1.4 Optional software .....	8
1.5 Software quick reference .....	10
1.6 IBM System Storage N series A models .....	12
1.6.1 Drive flexibility .....	14
1.6.2 Near-line storage .....	14
1.6.3 IBM System Storage N3000 introduction .....	15
1.6.4 IBM System Storage N5000 introduction .....	27
1.6.5 IBM System Storage N7000 introduction .....	33
1.7 IBM N series gateways (G models) .....	47
1.7.1 Two halves to set up .....	51
1.7.2 IBM N series Gateway highlights .....	52
1.7.3 Gateway RAID .....	52
1.7.4 IBM N5200, N5300, N5500, and N5600 Gateway models .....	54
1.7.5 IBM Gateway models N7600, N7700, N7800, and N7900 .....	56
1.7.6 LUN sizing .....	56
1.7.7 LUN mapping .....	57
1.8 Interoperability between G and A models .....	58
1.9 The N series expansion units .....	60
1.9.1 Intermixing EXN units with N series A models .....	61
1.9.2 EXN2000 .....	62
1.9.3 EXN1000 .....	65
1.9.4 EXN4000 .....	66
<b>Chapter 2. Introduction to Digital Video Surveillance</b> .....	69
2.1 Video as information: the rules have changed .....	70

2.2	Digital Video Surveillance	72
2.3	Analog video surveillance: do not throw out your video cameras	74
2.4	Building a digital video strategy: choosing future-ready technology	76
2.4.1	Open standard	76
2.4.2	Converged network	76
2.4.3	Video information	77
2.5	Converged network infrastructure	77
2.6	Advanced system capabilities	79
2.7	Systems, processes, and policies all contribute to the solution	82
2.8	The IBM Digital Video Surveillance solution	84
2.8.1	Data capture	85
2.8.2	Data transfer	86
2.8.3	Storage	87
2.8.4	Application and data integration	87
2.9	Components of a DVS solution	89
2.10	IBM Smart Surveillance Solution (S3)	90
2.11	Industries that use Digital Video Surveillance	94
2.12	IBM digital video surveillance solution business partners	95
2.13	The Benefits of IBM System Storage N series in the Digital Video Surveillance solution	97
<b>Chapter 3. Benefits of the IBM System Storage N series in the digital video surveillance solution</b>		
3.1	More flexible, lower TCO, unified storage	102
3.2	Enterprise vendor and solution	104
3.3	Performance and scalability	105
3.4	Flexible storage for different applications	106
3.5	SnapMirror	107
3.6	Clustering	108
3.6.1	Benefits of clustering for DVS solutions	109
3.7	FlexVol	110
3.8	MultiStore	110
<b>Chapter 4. Cisco Video Surveillance Media Server software</b>		
4.1	Introduction	114
4.1.1	Cisco Video Surveillance Media Server software	115
4.1.2	Media Platform features	116
<b>Chapter 5. Our environment</b>		
5.1	Hardware	120
5.2	Software	121
5.3	Operating system	122
5.4	SAN configuration	122
5.5	Network	123

5.6 Camera .....	125
<b>Chapter 6. Preparing the N series for Digital Video Surveillance .....</b>	<b>127</b>
6.1 The physical environment .....	128
6.2 Sizing N series storage to DVS .....	129
6.2.1 Calculating the storage needs <sup>1</sup> .....	132
6.2.2 Sample DVS storage calculations <sup>2</sup> .....	134
6.2.3 Factors in determining capacity hardware and license needs .....	135
6.3 N series calculations .....	136
6.4 Hardware considerations .....	137
6.4.1 Factors to consider .....	137
6.4.2 Mirroring data .....	137
6.4.3 RAID groups .....	139
6.5 SnapMirror .....	142
6.5.1 Criticality to business .....	143
6.5.2 System performance on the secondary storage device .....	143
6.5.3 Network bandwidth considerations .....	144
6.5.4 Asynchronous mode .....	144
6.5.5 Synchronous mode .....	146
6.5.6 Semi-synchronous mode .....	152
6.6 Why DVS needs SnapLock .....	154
6.6.1 License requirement .....	154
6.6.2 Network effects .....	154
6.6.3 SnapLock Compliance and SnapLock Enterprise .....	154
6.7 Fibre Channel Protocol topologies .....	156
6.7.1 FCP topology recommendations .....	159
6.8 FCP environment .....	160
6.8.1 Operating system requirements .....	160
6.8.2 Setting up N series .....	161
6.8.3 Setting up the host .....	165
6.8.4 Installing the FCP Linux Host Utilities .....	166
6.8.5 Installing HBAs .....	168
6.8.6 Setting the required HBA and driver parameters .....	173
6.8.7 Recording the World Wide Port Names of the host bus adapters. .	176
6.8.8 Loading the host bus adapter driver on the host .....	177
6.8.9 Editing the host's /etc/multipath.conf file .....	178
6.8.10 Starting the multipath service .....	180
6.8.11 Configuring the multipath service to start automatically .....	181
6.9 iSCSI topologies .....	181
6.10 iSCSI environment .....	184
6.10.1 Operating system requirements .....	184
6.10.2 N series requirements .....	185
6.10.3 Ethernet requirements .....	185

6.10.4	Host setup	185
6.10.5	N series setup	191
6.11	Cluster configuration	193
6.11.1	Cluster failover	194
6.11.2	CFMODE types	194
6.11.3	CFMODE set up	199
<b>Chapter 7.</b>	<b>Setting up the Digital Video Surveillance Solution</b>	<b>201</b>
7.1	N series configuration for a Fibre Channel Protocol environment	202
7.1.1	Initial configuration	204
7.1.2	Creating the aggregate	206
7.1.3	Creating the volumes	211
7.1.4	Creating the LUN	217
7.1.5	Mapping LUNs to initiator groups	219
7.2	Configuring Linux for the FCP environment	224
7.3	Configuring N series for the iSCSI environment	230
7.3.1	Initial Configuration	230
7.3.2	Creating the aggregate	233
7.3.3	Creating the volumes	233
7.3.4	Creating the LUN	234
7.3.5	Mapping the LUNs to the initiator groups	235
7.4	Configuring Linux for the iSCSI environment	237
7.5	Installing the Cisco Video Surveillance software	243
7.6	Installing the camera	246
<b>Chapter 8.</b>	<b>DVS operations with N series</b>	<b>247</b>
8.1	Capturing data	247
8.1.1	Starting the video proxy	248
8.1.2	Updating a proxy	256
8.1.3	Stopping a proxy	258
8.1.4	Viewing a proxy	259
8.1.5	Listing all proxies	262
8.2	Retaining and archiving data	263
8.2.1	The archive commands	264
8.2.2	List all archives command	279
8.2.3	List all running archives command	281
8.2.4	Archive details command	281
8.3	Clip management	283
8.3.1	Creating a clip	283
8.3.2	Extracting a clip with the AX Client	287
8.4	Migrating data	291
8.5	Expiring data	291
<b>Chapter 9.</b>	<b>Exploiting N series features and DVS</b>	<b>293</b>

9.1 FlexVol .....	294
9.1.1 Adding space to the aggregate (optional) .....	294
9.1.2 Changing the size of a volume .....	298
9.1.3 Having more space at the server level .....	304
9.2 Multistore .....	309
9.2.1 Adding a Multistore license .....	309
9.2.2 Preparing the node for Vfiler creation .....	310
9.2.3 Creating a Vfiler .....	313
9.2.4 Using Vfiler .....	320
9.3 SnapMirror .....	323
9.4 SnapLock .....	332
9.4.1 Adding the license .....	332
9.4.2 Creating SnapLock volumes .....	332
9.4.3 Managing WORM files .....	344
<b>Chapter 10. Administering DVS data with N series</b> .....	349
10.1 FilerView .....	350
10.1.1 Before you use FilerView .....	351
10.1.2 Accessing storage system using FilerView .....	351
10.1.3 Managing storage system with FilerView .....	354
10.1.4 Commonly used items in FilerView .....	356
10.2 Operations Manager .....	359
10.2.1 Before you use the Operations Manager .....	360
10.2.2 Accessing storage system using the Operations Manager .....	361
10.2.3 Setting up the Operations Manager .....	362
10.2.4 Managing the storage system with Operations Manager .....	363
10.3 Secure Shell .....	374
10.3.1 Before you use Secure Shell .....	375
10.3.2 Using Secure Shell .....	377
10.4 Telnet .....	384
10.4.1 Before you use Telnet .....	384
10.4.2 Using Telnet .....	385
10.5 Serial Console Access .....	387
<b>Appendix A. SnapDrive for UNIX on a Linux host</b> .....	391
SnapDrive overview .....	391
Installing SnapDrive for UNIX .....	393
Using SnapDrive .....	397
Summary .....	414
<b>Related publications</b> .....	415
IBM Redbooks .....	415
Other publications .....	415
Online resources .....	416

How to get Redbooks..... 416

Help from IBM ..... 416

**Index** ..... 417

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

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

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

AIX®  
Alerts®  
IBM®

Redbooks®  
Redbooks (logo) ®  
System Storage™

System x™  
Tivoli®  
TotalStorage®

The following terms are trademarks of other companies:

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

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

SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

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

QLogic, SANsurfer, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

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

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

Internet Explorer, Microsoft, SharePoint, SQL Server, Windows NT, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

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

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

“These materials have been reproduced by IBM with the permission of Cisco Systems Inc. COPYRIGHT © 1992 - 2008 CISCO SYSTEMS, INC. ALL RIGHTS RESERVED



# Preface

In this IBM® Redbooks® publication, we introduce Digital Video Surveillance (DVS) and the role that IBM plays in this industry. We provide the benefits of and define the role of the IBM System Storage™ N series in the DVS solution. We also include an example DVS application software installation and how it operates with the N series storage system.



*Figure 1 Hajiang Sha, Alex Osuna, Marco AurÉlio de Mello Santos, Marc Tu-Duy-Khiem*

## 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 at the International Technical Support Organization, Tucson Center. He writes extensively and teaches IBM classes worldwide on all areas of Storage. Before joining the ITSO three years ago, Alex was a Tivoli® Principal Systems Engineer specializing in Storage. He has over 30 years in the IT industry mainly focused on storage hardware and software. He has over 10 certifications from IBM, RedHat, and Microsoft®.

**Haijiang Sha** is an Advisory IT Specialist in IBM China. He has 10 years of experience in the network support field. He has 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.

**Marc Tu-Duy-Khiem** is an IT specialist in the EMEA Product and Solution Support Customer Center, Montpellier France. He was part of the EMEA/IBM/Oracle® Joint Solutions Center that supports Oracle and System x™ sales through architecture definition, benchmarks, proof-of-concept, publications, briefings, and education. He now supports N series sales.

**Marco AurÉlio de Mello Santos** is an IT Specialist in IBM Global Technology Services in Brazil where he works in a worldwide project as L2 support. He has more than seven years of experience in IT infrastructure. His areas of expertise include Microsoft infrastructure environments and network-attached storage (NAS) support.

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

Sangam Rancherla  
International Technical Support Organization, San Jose Center

Kent Breau  
Cisco Technologies

Syed Amin  
Product & Partner Engineer 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](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

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

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

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

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

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

- ▶ Mail your comments to:

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



# Introduction to IBM System Storage N series

In this chapter, we introduce the IBM System Storage N series and describe the hardware and software.

The IBM Storage System N series provides a range of reliable, scalable storage solutions for a variety of storage requirements. These capabilities are achieved using network access protocols, such as NFS, CIFS, HTTP, iSCSI, and 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 installing, administrating, and troubleshooting uncomplicated and straightforward.

The IBM System Storage N series is designed from the ground up as a standalone 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 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, you need to be able to expand it quickly and

non-disruptively. When existing storage exists, but is deployed incorrectly, the capability to reallocate available storage from one application to another quickly and simply cannot occur.

- ▶ To maintain availability and productivity during upgrades. If outages are required, N series keeps them to the shortest time possible.
- ▶ Create effortless backup and recovery solutions that operate commonly across all data access methods.
- ▶ File and block level services in a single system that help 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 resource can change non-disruptively, easily, and quickly. Online storage resource redeployment is possible.
- ▶ Upgrade process is easily and quickly. Non-disruptive upgrade is possible.
- ▶ Strong data protection solutions and support online for backup and recovery.

## 1.1 IBM N series hardware

Table 1-1 lists the available N series models:

- ▶ N3000 series
- ▶ N5000 series
- ▶ N7000 series

*Table 1-1 IBM N series storage systems*

A20 models	Max # of drives	Max capacity in TB
N3700	56	16.8
N3300	68	68
N3600	104	104
N5200	168	84
N5300	336	336
N5500	336	168
N5600	504	504
N7600	840	840

A20 models	Max # of drives	Max capacity in TB
N7700	840	840
N7800	1008	1008
N7900	1176	1176

Figure 1-1 illustrates the N series Gateway models.

Use	Model	Capacity
Midrange	N5200 G10 & G20 	84TB
Midrange	N5300 G10 & G20 	336TB
Midrange	N5500 G10&G20 	80TB
Midrange	N5600 G10&G20 	504TB
Enterprise class	N7600 G10&G20 	840TB
Enterprise class	N7700 G10 & G20 	840TB
Enterprise class	N7800 G10 & G20 	1008TB
Enterprise class	N7900 G10 & G20 	1176TB

Figure 1-1 IBM N series Gateway models

## 1.2 Comparing N series Gateway to N series storage systems

In this section, we compare N series Gateway to N series storage systems:

- ▶ 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 same version of Data ONTAP®
- ▶ An N5000 and N7000 series Gateway physical attributes are the same as the N5000 and N7000 models A10 and A20 storage systems
- ▶ Differences exist in system initialization and storage expansion
- ▶ The N series Gateway does not use the following features of Data ONTAP:
  - SnapLock® Compliance
  - LockVault™ Compliance
  - Nearstore option
- ▶ N series storage systems only use disk storage that IBM provides. The Gateway models support heterogeneous storage and IBM expansion units
- ▶ Data ONTAP is enhanced to enable the Gateway Series solution
- ▶ A RAID array from a separate storage system can provide LUNs to the Gateway:
  - Each LUN is equivalent to an IBM disk
  - LUNs are assembled into aggregates/volumes, and then formatted with the WAFL file system just like the IBM N series storage systems

## 1.2.1 IBM N series A models hardware quick reference

Table 1-2 is a quick hardware reference for the N series A models.

*Table 1-2 A models Hardware Quick Reference*

Function	N3700	N3300	N3600	N5200	N5500	N5300	N5600	N7600	N7700	N7800	N7900
Maximum Raw Capacity in TB A10 models	16	68	104	84	168	336	420	672	840	672	1176
Maximum Raw Capacity in TB A20 models	16	68	104	84	168	336	504	840	840	1008	1176
Fibre Channel Disk drives	144 GB 10K RPM, 144 GB 15K RPM, 300 GB 10K RPM - EXN2000 144 GB 15K, 300 GB 10K - EXN4000										
SATA Disk Drives	250 GB 7.2K RPM., 320 GB 7.2K RPM., 500 GB 7.2K RPM, 750 GB 7.2 KRPM, 1 TB 7.2 KRPM										
Maximum number of disks	56	68	104	168	336	336	420(A10) 504(A20)	672(A10) 840(A20)	840	672 (A10) 1008 (A20)	1176
Expansion units supported	EXN1000 (SATA), EXN2000 (FC), EXN4000 4 Gbps (FC)										



Table 1-3 is a hardware quick reference for the N series G models.

*Table 1-3 N series G models quick reference*

Function	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
Maximum Raw Capacity in TB G10 models	84	336	84	504	672	840	672	1176
Maximum Raw Capacity in TB G20 models	84	336	84	504	840	840	1008	1176
Max. number of – Logical Units (LUNs) on back-end disk storage array	168	252	336	420 for A10 and 504 for A20	840	840	1008	1176
Max LUN size in GB	1000	1000	500	1000	1000	1000	1000	1000
Maximum Volume size in TB	16	16	16	16	16	16	16	16

## 1.2.2 IBM N series A and G models hardware quick reference

*Table 1-4 Storage System Reference*

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
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										

Function	N3700	N3300	N3600	N5200	N5300	N5500	N5600	N7600	N7700	N7800	N7900
Onboard I/O ports per node	2 X GbE 2 X Optica I FC	4 x GbE 4 x Optica I FC	4 x GbE 4 x Optica I 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 (4 Gbps)	6 X GbE 8 X FC	6 X GbE 8X FC	6 X GbE 8 X FC	6 X GbE 8X 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	3 X PCI-E 6 X PCI-E	5 X PCI-E, 3 X PCI-X	3 X PCI-E 6 X PCI-E
NVRAM in MB per node	128	128	256	512	512	512	512	1024	1024	4096	4096
Memory in GB per node	1	1	2	2	4	4	8	16	32	32	64
Redundancy/ High Availability	CompactFlash, dual-redundant hot-plug integrated cooling fans, hot-swappable autoranging power supplies, clustered storage controllers, hot-swappable disk bays**										
Required rack space	3U	2U	4U	3U per node	3U per node	3U per node	3U per node	6U per node	6U 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 1.8 GHz dual-core	two 2.6 GHz AMD Opteron	two 2.6 GHz AMD Opteron	four 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 1.8 GHz dual-core	four 2.6 GHz AMD Opteron	four 2.6 GHz AMD Opteron	eight 2.6 GHz AMD Opteron	eight 2.6 GHz AMD Opteron

## 1.3 The IBM N series standard software features

Table on page 8 contains licensed, no charge features that are available with IBM N series.

Table 1-5 Standard software features

Feature	Description
Data ONTAP	Operating system software that optimizes data serving and allows multiple protocol data access.
File Transfer Protocol (FTP)	A standard Internet protocol, which 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, and provides 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. Dynamic, nondisruptive (thin) storage provisioning, space- and time-efficiency. These flexible volumes can span multiple physical volumes without regard to size.
FlexShare™	FlexShare gives administrators the ability to leverage existing infrastructure and increase processing utilization without sacrificing the performance of critical business needs. Using FlexShare, administrators can confidently consolidate different applications and data sets on a single storage system. FlexShare gives administrators the control to prioritize applications based on how critical they are to the business.
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. 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 that is featured in the Data ONTAP operating software and inside each N series system, which 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-mail

Feature	Description
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"> <li>▶ Ensures high data availability for business-critical requirements by eliminating a single point-of-failure.</li> <li>▶ Must be ordered for A20 clustered configurations or upgrades from A10 to A20.</li> <li>▶ Active-active pairing delivers even more “nines to right of the decimal point”.</li> </ul>
NIS	N series does provide NIS client support and can participate in NIS domain authentication.
Integrated automatic RAID manager	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 that the customer needs to diagnose problems.

## 1.4 Optional software

Table 1-6 contains the optional software. For more information, refer to Section 1.5, “Software quick reference” on page 10, which contains additional optional software that is available with IBM N series.

Table 1-6 *Optional software*

Feature	Description
CIFS	Provides file system access for Microsoft Windows environments.
NFS	Provides file system access for UNIX® and Linux environments.
Hypertext Transfer Protocol (HTTP)	HTTP allows you to transfer displayable Web pages and related files.
FlexClone®	Designed to provide instant replication of data volumes and sets without requiring additional storage space at the time of creation

Feature	Description
Multistore	<ul style="list-style-type: none"> <li>▶ Permits an enterprise to consolidate a large number of Windows, Linux, or UNIX file servers onto a single storage system</li> <li>▶ Many “virtual filters” on one physical appliance ease migration and multi-domain failover scenarios</li> </ul>
SnapLock	Provides non-erasable and non-rewritable data protection that helps enable compliance with government and industry records retention regulations
SnapMirror®	<ul style="list-style-type: none"> <li>▶ Remote mirroring software that provides automatic block-level incremental file system replication between sites.</li> <li>▶ Available in synchronous, asynchronous, and semi-synchronous modes of operation</li> </ul>
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®	Enables Windows and UNIX applications to access storage resources on N series storage systems, which are presented to the Windows 2000 or later operation system as locally attached disks. For UNIX users, 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, SQL Server®, and SAP® backup and restore. SnapManager software simplifies Exchange data protection by automating processes to provide hands-off, worry-free data management. A new integrated GUI-based administration, SnapManager Console, was introduced for the MS Office SharePoint® Server. It unifies the entire stack from applications to servers to storage for backup, recovery, and extraction of critical information.
SnapValidator®	For Oracle deployments, SnapValidator provides an additional layer of integrity that checks between the application and N series storage. SnapValidator allows Oracle to create checksums on data that is transmitted to N series storage for writes to disk and include the checksum as part of the transmission.
SyncMirror®	SyncMirror is synchronous mirror of a volume. It maintains a strict physical separation between the two copies of your mirrored data. In case of an error in one copy, the data is still accessible without any manual intervention.
Single Mailbox Recovery for Exchange (SMBR)	SMBR is a software option from SnapManager that takes near-instantaneous online backups of Exchange databases, verifies that the backups are consistent, and rapidly recovers 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.

Feature	Description
Operations Manager	Provides remote, centralized management of IBM N series data storage infrastructure, which includes global enterprise, storage network, and so on.
MetroCluster	MetroCluster software provides an enterprise solution for high availability over wide area networks.
NearStore® option	A disk-based, secondary storage device for enterprise applications.
Advanced Single Instance Storage	Significantly improves physical storage efficiency and network efficiency by enabling the sharing of duplicate data blocks.
Virtual File Manager (VFM®)	IBM System Storage N series Virtual File Manager (VFM) software is a comprehensive solution for managing unstructured file data. It is designed to provide data management functionality for server and storage consolidation, migration, remote office data management, and disaster recovery features while avoiding disruption to users. It provides all of this functionality through automated policy-based data management, leveraging a global namespace.

## 1.5 Software quick reference

Table 1-7 is a software quick reference list that includes both the optional and included software.

Table 1-7 Software quick reference

Product, feature, or function	Included or 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

<b>Product, feature, or function</b>	<b>Included or optional</b>	<b>N3000 A10 &amp; A20</b>	<b>N5X00 A10 &amp; A20</b>	<b>N5x00 G10 &amp; G20</b>	<b>N7x00 A10 &amp; A20</b>	<b>N7x00 G10 &amp; G20</b>
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(A20)	X(G20)	X(A20)	X(G20)
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
SnapDrive for Windows and UNIX: AIX, Solaris™, HP-UX, 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						
SnapManager for Exchange	Optional	X	X	X	X	X

Product, feature, or function	Included or optional	N3000 A10 & A20	N5X00 A10 & A20	N5x00 G10 & G20	N7x00 A10 & A20	N7x00 G10 & G20
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
MetroCluster A2X Models only	Optional		X	X	X	X
Disk Sanitization	Included	X	X		X	X
SnapLock Compliance	Optional	X	X		X	
LockVault Compliance	Optional	X	X		X	
NearStore Option Bundle	Optional	X	X	X	X	X
RAID4, RAID-DP™	Included	X	X		X	
Advanced Single Instance Storage	Optional		X	X	X	X
VFM	Optional	X	X	X	X	X

## 1.6 IBM System Storage N series A models

The A models of the IBM System Storage N series offer multi protocol connectivity using internal storage or storage that expansion units provide, as shown in Figure 1-2 on page 13. The IBM System Storage N series systems are designed to provide integrated block- and file-level data access, which allows concurrent operation in IP SAN (iSCSI), FC SAN, NFS, and CIFS environments. Other storage vendors might 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, which helps you gain a competitive advantage.

The 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 system, 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 that are normally found in the standard operating systems.

The N series comes with pre-configured software and hardware and with no monitor or keyboard for user access, which 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 an N series storage systems product is its ability to be installed rapidly using minimal time and effort to configure the system. It is integrated seamlessly into the network, which makes IBM N series products especially attractive when lack of time and skills are elements in the decision process.

Figure 1-2 shows the IBM N series A model.

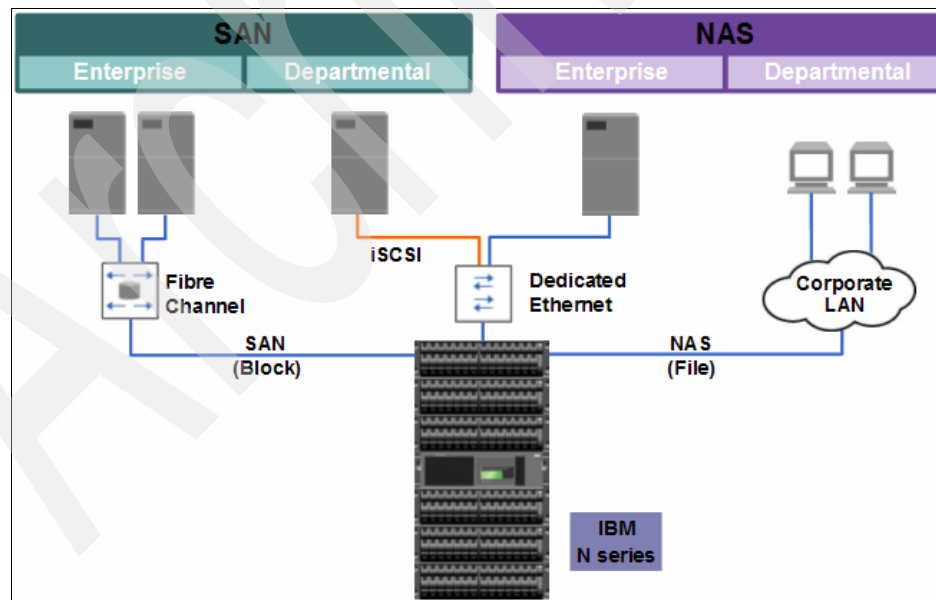


Figure 1-2 IBM N series A models

### 1.6.1 Drive flexibility

The IBM System Storage N series products provide network attached storage for environments where you need to utilize your storage investment in a multifaceted environment. The IBM System Storage N series storage systems also provide a tremendous amount of versatility by allowing this solution to be populated with both Fibre Channel disk drives and serial advanced technology attachment (SATA) disk drives. An N series that is populated with Fibre Channel disk drives might be suitable for mission-critical high-performance data transaction environments; whereas, an N series that is populated with SATA disk drives is attractive to those who want to 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 1-8 provides information about drive positioning.

Table 1-8 Drive positioning

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

### 1.6.2 Near-line storage

The IBM N series with SATA Drives offers Near-line Storage. Figure 1-3 on page 15 shows 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 were traditionally 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, which enables organizations to do daily backups to disk and backup 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 that Near-line storage provides is that you can leverage your existing investment in primary storage, your backup application, and tape libraries.

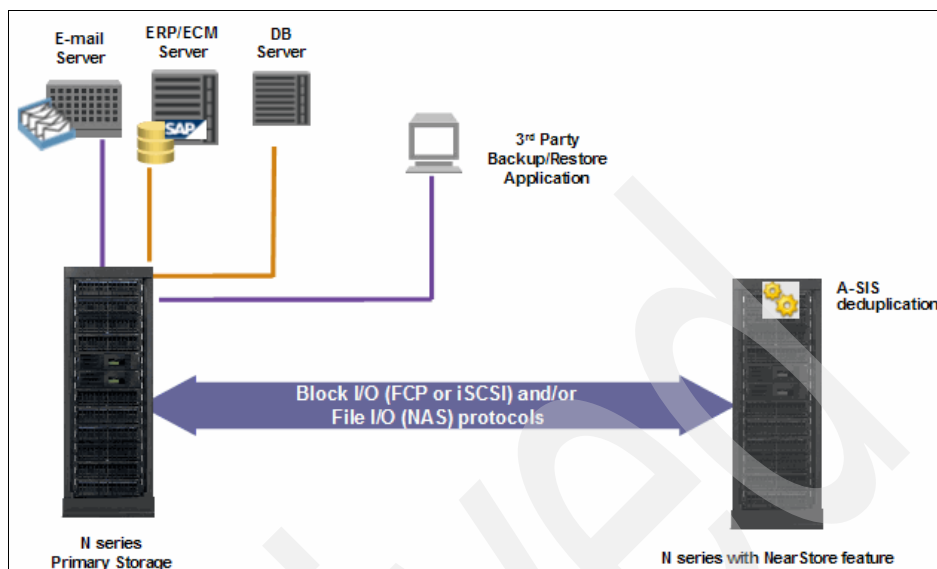


Figure 1-3 Near-line Storage

### 1.6.3 IBM System Storage N3000 introduction

The N3000 systems 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, intelligent management software, and data protection capabilities, such as higher-end N series systems, in a cost-effective package. N series innovations include Serial-Attached SCSI (SAS) drive support, expandable I/O connectivity, and on board remote management.

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

- ▶ High availability, leverages proven features including a high performing and scalable operating system, data management software, and redundancy features
- ▶ Backup and recovery features that support disk-based backup with file or application-level recovery with Snapshot and SnapRestore software features
- ▶ Simple replication and Disaster Recovery that provide easy-to-deploy mirroring solution that is highly tolerant of WAN interruptions
- ▶ Management simplicity, self-diagnosing systems that enable on-the-fly provisioning

- ▶ Versatile, single, integrated architecture that supports concurrent block I/O and file serving over Ethernet and Fibre Channel SAN infrastructures

The 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 that addresses 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 series supports Ethernet and Fibre Channel environments, which enables economical NAS, FC, and iSCSI deployments. The N3000 system functions as a “unification engine,” which is designed to allow you to simultaneously serve 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.

There are no PCI adapter slots on the N3300 and N3700 systems, no additional adapter options are supported for the N3300 and N3700 systems. There is one available PCIe adapter slot per node on the N3600 storage system. For an A20 model, you must add adapters in pairs, one per node, so that both nodes are populated with one of the same type of PCIe adapter.

### **N3700**

The N3700 storage system, shown in Figure 1-4 on page 17, is a 3U solution that provides 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, supports failover and failback functions to maximize reliability. The N3700 storage system can support 14 internal hot-plug disk drives with

scalability provided through attachment to up to three expansion units, each with a maximum of 14 drives. The N3700 can also connect to a Fibre Channel tape for backup.

A list of supported Tape drives are at:

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

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

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



*Figure 1-4 N3700*

The type of controller defines the model. Figure 1-5 on page 18 shows a single control unit. The single node A10 uses a single control unit with dual node clustered A20 using two control units, as shown in Figure 1-6 on page 18.



Figure 1-5 N3700 A10



Figure 1-6 N3700 A20

The N3700 comes with redundant power supplies for higher reliability, as shown in Figure 1-7.




Figure 1-7 Redundant Power supplies

From the rear of the N3700, on the power supply, you can see the Diagnostic and Operational LEDs, which we show in Figure 1-8 on page 19. Table 1-9 on page 19 explains the LEDs and what the possible configurations mean.




Figure 1-8 Diagnostic and Operational LEDs

Table 1-9 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, temperature fault	Off	
PSU Status Normal	Off	Power Supply failure
AC missing for this PSU	Off	
Fan Fault	Off	
Output voltage, current, temperature fault	On	
PSU Status Normal	Off	Fan failure
AC missing for this PSU	Off	
Fan Fault	On	
Output voltage, current, temperature fault	Off	



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

Figure 1-9 shows the CPU module, which 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, both cooling fans will continue to operate.





Figure 1-9 CPU Tray module front view

The N3700 is based around a MIPS dual core processor. It has 1 GB of system memory of which 128 MB is defined as non-volatile due to having a battery back up. The battery is a three cell Li-Ion, as shown in Figure 1-10 on page 21.



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



*Figure 1-10 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, shown in Figure 1-11 on page 22. The flash card contains a copy of the Data ONTAP operating system and firmware. The operating system is also stored on each disk drive.




Figure 1-11 Bottom of CPU tray module showing Compact Flash card

Figure 1-12 shows rear ports on the N3700. 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 Plugable SFPs.

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

The second port, channel B, is copper and is exclusively used for connection of the expansion unit. A special copper cable (option x6531-C) is used for this connection.




Figure 1-12 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 a RJ45 to DB-9 cable.

### **N3700 hardware features**

The N3700 hardware features are:

- ▶ 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

Although ESH2 modules in the EXN2000 would support up to 84 drives per loop or one N3700 and five expansion shelves, only 56 drives, 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), which is not available on the N series.

### **Optional N3700 hardware**


The second CPU tray supports cluster failover. Table 1-10 shows additional N3700 specifications.

*Table 1-10 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 and N3600, Figure 1-13 on page 24 and Figure 1-14 on page 24, systems provide multiple I/O connectivity options, a small footprint to hold high-density SAS drives, external expansion using low-cost SATA drives and Fibre Channel disks for production applications. The N3300/3600 also use 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, firmware upgrades, console commands, and diagnostics to help maintain the reliability of the system and your business-critical data. Figure 1-13 on page 24 is an example of the N3300.



*Figure 1-13 N3300*

Figure 1-14 is an example of the N3600.



*Figure 1-14 N3600*

Figure 1-15 shows the rear panel of the N3300, and Figure 1-16 on page 25 shows the rear panel of the N3600. The single node A10 uses a single control unit with dual node clustered A20 using two control units (see Figure 1-6 on page 18).



*Figure 1-15 N3300 rear view*



Figure 1-16 N3600 rear view

The N3300 has redundant power supplies, as shown in Figure 1-17.

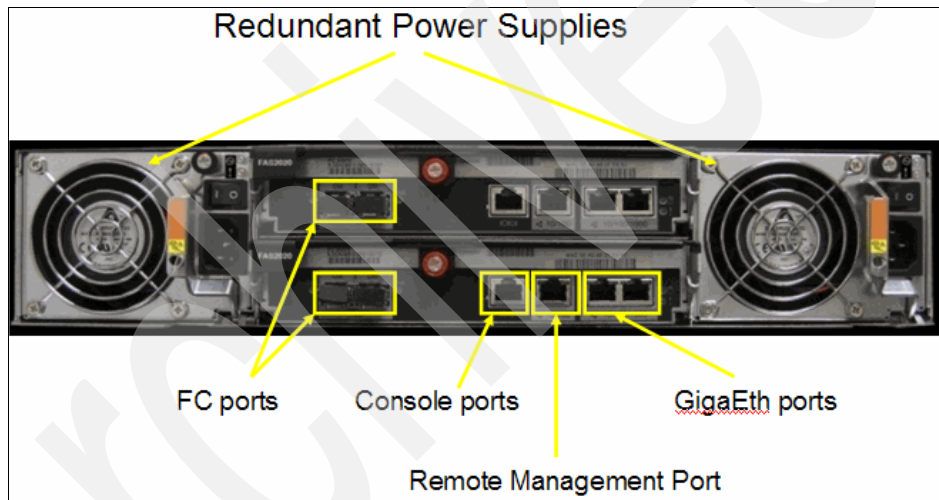


Figure 1-17 external ports on N3300

N3300 is a 2U high device. It has 12 internal SAS drive bays and can support up to two external disk expansion units. Each Controller has dual Gigabit Ethernet ports and dual 4 Gbps Fibre Channel ports (see Figure 1-17). The N3300 also has one console port and one Remote Management port.

**Tips:** The N3300 series supports SAS, FC, and SATA disk technologies. 12 SAS disk drives are supported in the controller chassis. You can configure the N3300 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, as shown in Figure 1-18.




Figure 1-18 N3600 power supplies and expansion slot

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, as shown in Figure 1-19. It also has one console port and one Remote Management port. N3600 has a PCIe Slot on each controller.




Figure 1-19 External ports on N3600

**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***

The key specifications for the N3300 and N3600 are:

- ▶ 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, Fibre Channel
- ▶ Each controller: Dual Gigabit Ethernet Ports and dual 4 Gbps Fibre Channel Ports
- ▶ Onboard Remote Platform Management
- ▶ Internal SAS Drive Bays

The N3000 series 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.

## **1.6.4 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 midrange to high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help to 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 can 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 an EXN1000, EXN2000, and EXN4000 storage expansion unit (see Figure 1-22 on page 30 and Figure 1-21 on page 29). 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 1-3 on page 5.


Figure 1-20 on page 29 shows the N5200. From the Front the N5600 and N5300 also looks very similar to the N5200 and N5500. Some of the differences are on the rear (Figure 1-24 on page 31), especially the absence of a LVD iSCSI connector. The N5600 and N5300 also uses a BIOS prompt upon boot rather than a Common Firmware Environment (CFE) prompt.





*Figure 1-20 N5200*


Figure 1-21 shows the N5500 with EXN2000 shelves.



*Figure 1-21 N5500 with EXN2000 shelves*

Depending on the model that you have, you will see one N5200 or two N5500 modules internally.

Figure 1-22 shows the N5200 One CPU model.



*Figure 1-22 N5200 One CPU module*

The easily accessible rear of the N5000 series provides I/O connectivity and power supply access and status indications, as shown in Figure 1-23 on page 31, which shows the rear view of the N5500 and N5200.

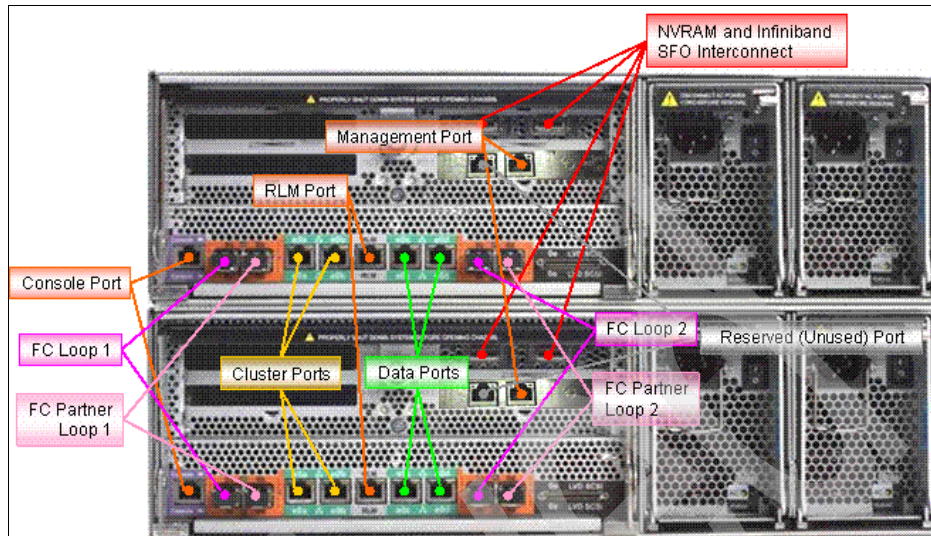



Figure 1-23 Rear view of N5500 and N5200

Figure 1-24 shows the rear view of the N5600 and N5300 and points out the differences between the N5200/N5500 and the N5300/N5800.



*Figure 1-24 Rear view of N5600 and N5300*

Figure 1-25 on page 32 shows the top view of the N5200 or N5500, specifically the modular design and field replaceable unit capabilities.




Figure 1-25 Top view of N5000

Figure 1-26 shows the motherboard of the N5200, which is a self containing unit that holds components such as memory, CPU, and interfaces.




Figure 1-26 N5200 and N5500 motherboard

### RAID group sizes

Table 1-11 shows the RAID group size in drive types.

Table 1-11 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

N5200, N5300, N5500, and N5600 multi-disk drive options offer mission critical, Near-line, and compliance storage solutions. Figure 1-27 shows the multi-storage options.

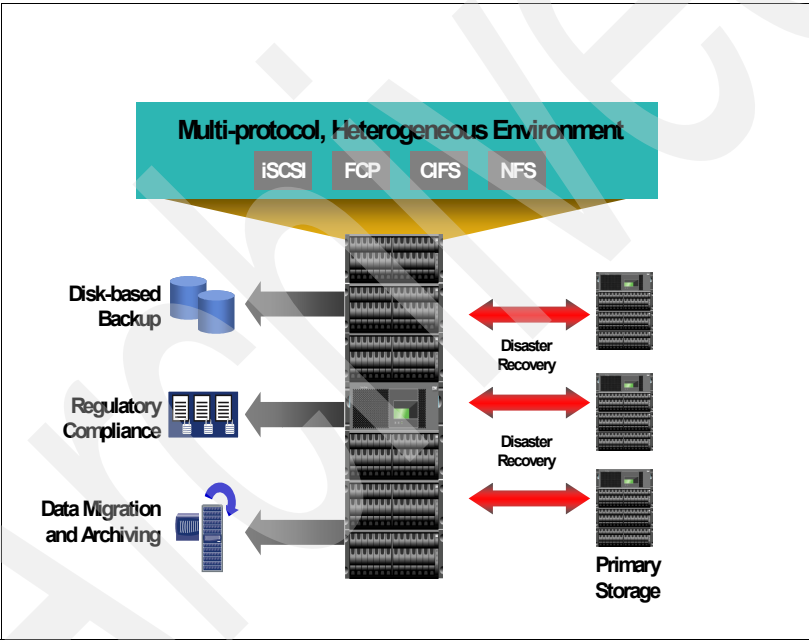


Figure 1-27 Multi storage options

### 1.6.5 IBM System Storage N7000 introduction

The IBM System Storage N7000 series offers additional choices to organizations that face the challenges of enterprise data management. The IBM System Storage N7000 series delivers high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help to support your efforts to increase reliability,



simplify and unify storage infrastructure and maintenance, and deliver exceptional economy.

The IBM N7000 A series comes in four models:

- ▶ N7600:
  - 2866-A10 Single Node
  - 2866-A20 Clustered
- ▶ N7700:
  - 2866-A11 Single Node
  - 2866-A21 Clustered
- ▶ N7800:
  - 2867-A10 Single Node
  - 2867-A20 Clustered
- ▶ N7900:
  - 2867-A11 Single Node
  - 2867-A21 Clustered

FC or SATA (both can be used behind a single controller but not in the same drawer).

Like its N5000 predecessor, the front of the N7000 series unit, Figure 1-28 on page 35, has the LCD display and the standard three LEDs that indicate system activity, status, and power. Externally the N7600 and N7800 appear the same. When compared to the N7600, the differences are internal with the increased CPU, memory, and NVRAM capability of the N7800.



*Figure 1-28 Front view of N7000*

Figure 1-29 on page 36 shows the rear of the N7000. Notice the redundant power supplies, the NVRAM card, the Gigabit Ethernet interfaces, and the Fibre Channel interfaces. The console port and RLM port are also located on the rear.




Figure 1-29 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 1-29).

Figure 1-30 on page 37 shows a racked N7000.





*Figure 1-30 N7000 racked*

A dual-node N7600 supports a maximum of 60 storage expansion units (EXN1000 and EXN2000). A dual-node N7800 supports a maximum of 72 storage expansion units (EXN1000 and EXN2000). See Figure 1-32 on page 39. A dual-node N7700 and N7900 supports a maximum of 60 and 84 storage expansion units respectively. Each rack holds a maximum of 12 expansion units. The N7000 products are installed by an IBM service or qualified IBM Business partner and are not customer setup.

Figure 1-31 on page 38 shows a clustered N7000 with multiple expansion units.



*Figure 1-31 Clustered N7000 with multiple expansion units*

When you remove the bezel, Figure 1-32 on page 39, you see the CompactFlash card reader, and directly below it, the Remote LAN Module (RLM). The RLM is required in all N7000 series systems. The systems will not boot unless the card is present. Also, notice the five fan units. The fans are hot swappable and are numbered in Figure 1-32 on page 39 for your reference.




Figure 1-32 Front of N7000 with bezel removed

On the side of the system, Figure 1-33 on page 40, notice that there are two handles on each side to help you lift the system, which is very heavy. Fully loaded, it weighs 120 pounds. IBM recommends that before you lift the system, you remove the fan units and the two power supplies, which reduces the weight to slightly over 90 pounds. We recommend 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 1-33 Caution with lifting the N7000

As shown in Figure 1-34 on page 41, the two hot swappable power supplies are visible and removable from the rear of the N7000.





Figure 1-34 N7000 power supplies

In the middle, Figure 1-35 on page 42, there are 9-PCI slots, which are numbered 1-9, from left to right, and all of the slots are slot-specific. When installing any adapter cards, always use the System Configuration Guide, which is on the IBM site for reference:

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





*Figure 1-35 PCI slots*

Moving below the PCI slots, Figure 1-36 on page 43 shows the console port and the RLM port.



Figure 1-36 RLM and console ports

Below the Ethernet ports, Figure 1-37 on page 44 shows the Fibre Channel Tray or “FC Tray”, with eight onboard Fibre Channel ports. This tray is actually a Field Replaceable Unit.




Figure 1-37 Fibre Channel ports


From inside the N7000 looking from the top, shown in Figure 1-38 on page 45, the PCI slots and system memory are visible but the processors are not. They are on the other side of the motherboard tray. Recall that the N7900 and N7800 have eight CPUs, and the N7700 and N7600 have four processors.

From this perspective, you can see the nine PCI slots. Slots 3, 4, and 9 are black and represent PCI-X.

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

Notice that the NVRAM6 adapter resides in slot 2 on this standalone 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 1-38 Top of N7000*

Keep in mind that the N7800 and N7900 use an NVRAM6 adapter with 4 Gigabytes of memory, and the N7600 and N7700 use an NVRAM6 adapter with 1024 Megabytes of memory, as shown in Figure 1-39 on page 46.




Figure 1-39 NVRAM

There are some new LEDs that you should be aware of. The fan units, PCI slots, and memory DIMMs have LEDs to indicate a failed component, as shown in Figure 1-40.




Figure 1-40 New LEDs

# RAID Group size

Table 1-12 shows the RAID Group size in drive type.

Table 1-12 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
N7700 RAID 4	8	14	7	7
N7700 RAID DP	16	28	14	16
N7800 RAID 4	8	14	7	7
N7800 RAID DP	16	28	14	16
N7900 RAID 4	8	14	7	7
N7900 RAID DP	16	28	14	16

## 1.7 IBM N series gateways (G models)

The IBM System Storage N series Gateways, an evolution of the N5000 series product line, are a network-based virtualization solution that virtualizes tiered, heterogeneous storage arrays that allow you to leverage the dynamic virtualization capabilities that are available in Data ONTAP across multiple tiers of IBM and vendor-acquired storage (Figure 1-41 on page 49). Like all IBM N series storage systems, the IBM 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.

The industry's most comprehensive virtualization solution, the N series Gateways provide 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, as shown in Figure 1-41 on page 49.

The IBM N series Gateway presents 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. N series Gateways can take storage array LUNs (which are treated as disks) and virtualize them through Data ONTAP, which

presents a unified management interface or uses expansion units, such as the EXN4000.

The IBM N series Gateway, Figure 1-42 on page 50, offers you new levels of performance, scalability, and a robust portfolio of proven data management software. IBM N series storage systems seamlessly integrate into mission-critical SAN environments and provide a simple, elegant data management solution that decreases management complexity, improves asset utilization, and streamlines operations to increase business agility and to reduce total cost-of-ownership for organizations that are looking for ways to leverage SAN-attached storage to 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 that require NAS, SAN, or IP SAN data access.

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, iSCSI) and file (NFS, CIFS) access
- ▶ Increasing local or remote shared file services and file access workloads.

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

With Data ONTAP, the N series Gateway now supports attachment of heterogeneous storage systems and IBM expansion units of the type that are used with N series storage systems, shown in Figure 1-43 on page 51.

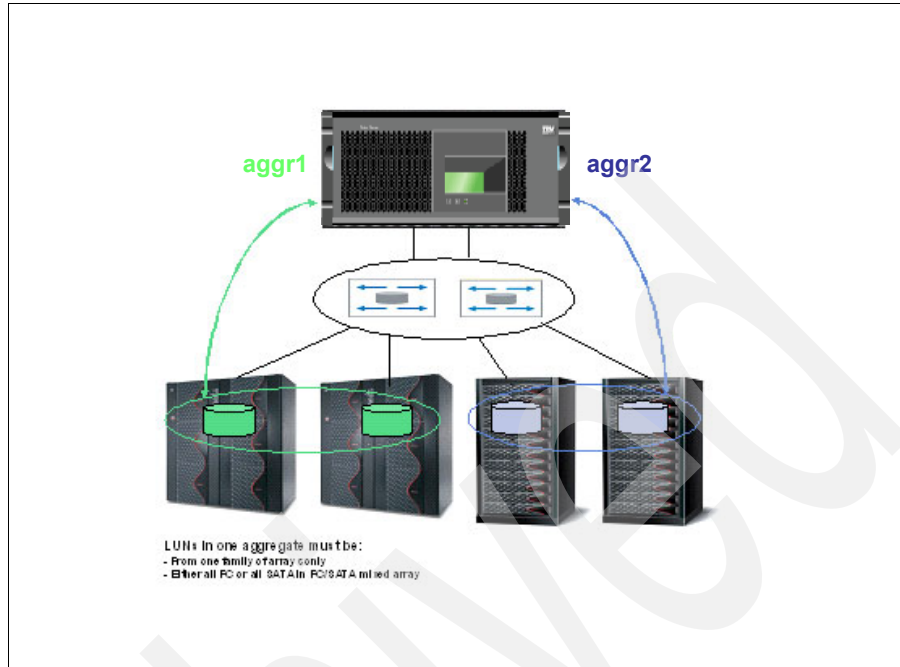


Figure 1-41 Heterogeneous storage

Figure 1-42 on page 50 illustrates the Gateway topology.

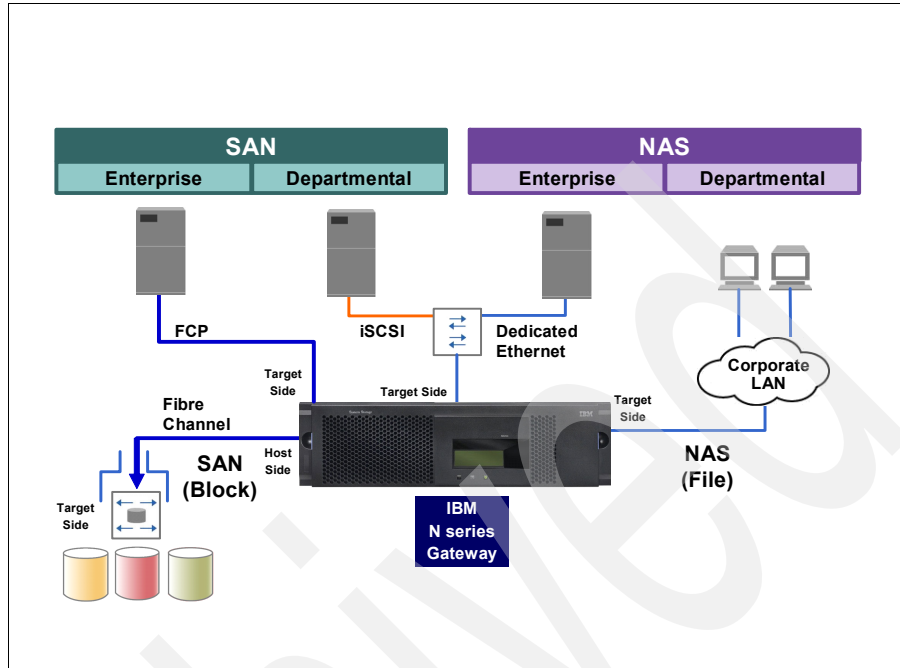


Figure 1-42 Gateway topology

Figure 1-43 on page 51 shows storage systems with expansion units.




Figure 1-43 Storage systems with expansion units

### 1.7.1 Two halves to set up

Think of an N series Gateway implementation 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, etc.)
- ▶ Setting up backup including NDMP dumps to tapes.

The back-end implementation includes all tasks that are required to set up the N series Gateway's system up to the point where it is ready for Data ONTAP installation:

- ▶ Array LUN formatting
- ▶ Port assignment
- ▶ Cabling
- ▶ Switch zoning
- ▶ Assigning LUNs to the V-Series system
- ▶ Creating aggregates
- ▶ Loading Data ONTAP



## 1.7.2 IBM N series Gateway highlights

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

- ▶ 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, as shown in Figure 1-44.




Figure 1-44 Tiered heterogeneous storage

## 1.7.3 Gateway RAID

Gateways use RAID0 on top of RAID1, RAID5, or RAID10 on RAID storage subsystems, as shown in Figure 1-45 on page 53. Physical disk operations, such as scrubbing, is disabled.






Figure 1-45 RAID configuration

RAID0 is used to write data. Example 1-1 is an example of volume status with the Gateway, which looks very similar to what you would see on a 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

---

## 1.7.4 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 that support high-demand operations. It houses hot swappable, redundant power supplies and fans and supports multi path failover protection and host dual pathing between the unit and its SAN-attached storage device. Additionally, 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 are in the spectrum of Data ONTAP features that are supported and enabled:

- ▶ 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 1-13 shows the Gateway capacity.

*Table 1-13 Gateway capacity*

Model	Maximum capacity
2864-G10	50TB
2864-G20	50TB per node
2869-G10	126TB

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

**Important:** If you plan 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 on failure. (The `cf.takeover.on_panic` option controls whether a cluster partner immediately takes over for a panicked partner.)

Table 1-14 shows the maximum number of LUNs for Gateways.

*Table 1-14 Maximum number of LUNs*

Model	Maximum number of LUNs
N5200 2864-G10 (non cluster model)	168
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
N5600 2868-G20 (clustermodel)	For each node, single node N5600 2868-G10 values apply

## 1.7.5 IBM Gateway models N7600, N7700, N7800, and N7900

The IBM System Storage N7000 series gateway models offer additional choice to organizations that face the challenges of enterprise data management. The IBM System Storage N7000 series delivers high-end enterprise storage and data management value with midrange affordability. Built-in enterprise serviceability and manageability features help to support your efforts to increase reliability, simplify and unify storage infrastructure and maintenance, and deliver exceptional economy. The IBM N series Gateway models N7600 and N7800 deliver all of 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. 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 G series comes in four models:

- ▶ N7600:
  - 2866-G10 Single Node
  - 2866-G20 Clustered
- ▶ N7700:
  - 2866-G11 Single Node
  - 2866-G21 Clustered
- ▶ N7800:
  - 2867-G10 Single Node
  - 2867-G20 Clustered
- ▶ N7900:
  - 2867-G11 Single Node
  - 2867-G21 Clustered

## 1.7.6 LUN sizing

The maximum and minimum sizes that Gateway supports for LUN sizes are:

- ▶ Maximum LUN size: 1024 GB
- ▶ Minimum LUN size: 1001 MB

**Note:** The Data ONTAP definition of a GB is: one GB is equal to  $1000 \times 1024 \times 1024$  bytes. Therefore, the maximum LUN size that Data ONTAP supports means  $1024 \times 1000 \times 1024 \times 1024 = 1,048,576,000,000$  bytes.

## 1.7.7 LUN mapping

Storage Subsystem LUNs are converted to disks for the IBM N series Gateway. The equivalent disk count is the Gateway disk count = LUN count.

Figure 1-46 is an example of an Array LUN that is 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 (Figure 1-47 on page 58) that we use on the N series A models.

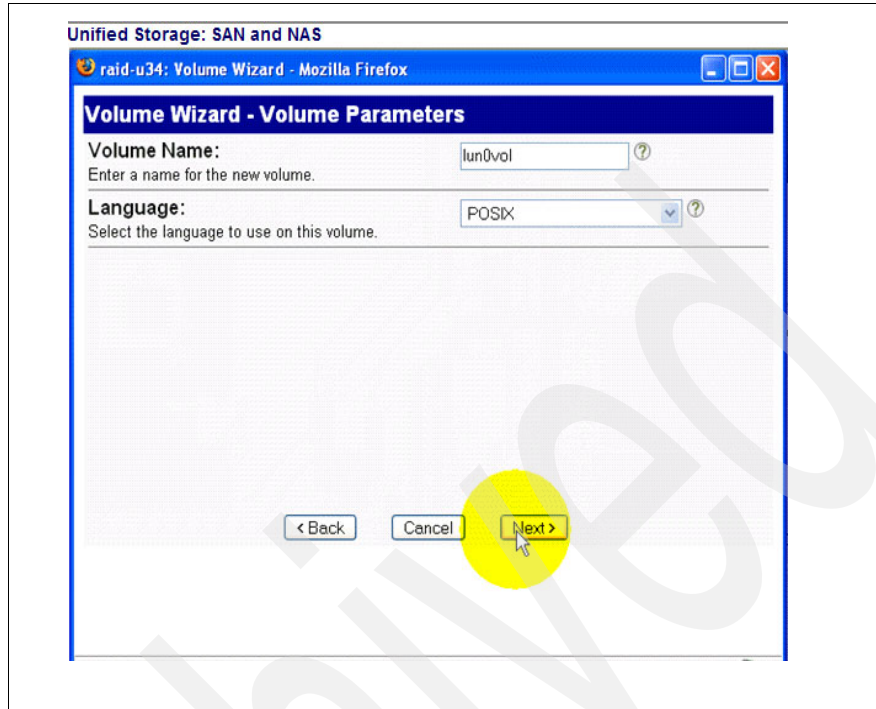


Figure 1-47 Volume wizard

**Note:** Do not map LUN 0 to gateway systems, even if LUN 0 is a storage LUN.

## 1.8 Interoperability between G and A models

In this section, we provide the items that are not interoperable between the G and A models:

- Replication between SnapMirror on G model SnapMirror on A model, as shown in Figure 1-48 on page 59. Includes async, semisync, and synchronous




Figure 1-48 SnapMirror interoperability

- ▶ Disk-to-disk backup from SnapVault primary on G model to SnapVault secondary on A model, as shown in Figure 1-49 on page 60.
- ▶ Disk-to-disk backup from SnapVault primary on A model SnapVault secondary on G model.




Figure 1-49 SnapVault interoperability

## 1.9 The N series expansion units

There are currently three disk storage expansion units that are specifically designed for the IBM N series filers:

- ▶ IBM EXN4000 Fibre Channel disk storage expansion unit
- ▶ IBM EXN2000 Fibre Channel disk storage expansion unit
- ▶ IBM EXN1000 serial advanced technology attachment (SATA) storage expansion unit

**Note:** EXN expansion units are not intended for attachment to a Gateway.

Multiple EXN1000s, each having different SATA disk drive feature codes, might be attached to the same N series filer on the same Fibre Channel loop. Multiple EXN2000s and EXN4000s, each having different Fibre Channel disk drive feature codes, might be attached to the same N series filer on the same Fibre Channel loop. For the latest storage expansion unit support information, visit:

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



## 1.9.1 Intermixing EXN units with N series A models

EXN4000s and EXN2000s are both Fibre Channel disk storage expansion units. You can mix EXN4000 and EXN2000 expansion units within the same loop, but you must set the speed switches on all EXN4000s and EXN2000s to the same speed, which is either 1 Gbps or 2 Gbps, as shown in Figure 1-51 on page 62. Intermixing Fibre Channel and SATA disk drives in a supported N series filer configuration is supported as follows:

- ▶ Intermixing Fibre Channel disk expansion units with SATA disk expansion units on the same loop is not supported.
- ▶ EXN4000s or EXN2000s (Fibre Channel disk drives) and EXN1000s (SATA disk drives) may be attached to the same N series filer only if the Fibre Channel disk expansion units (EXN4000s or EXN2000s) are on separate loops than the SATA disk expansion units (EXN1000s).



Figure 1-50 Speed switches

**Note:** Intermixing Fibre Channel and SATA disk drives in an N3700 configuration is not supported. Only N3300, N3600, N5000, and N7000 series models support Fibre Channel and SATA disk drives intermixing in a configuration.

## 1.9.2 EXN2000

The EXN2000 is a fibre expansion unit for the N series that looks very similar to the N3700; however, unlike the N3700, which has the CPU modules, the EXN2000 supports only the disk modules and the connectivity to them. See Figure 1-51, which is an image of the EXN2000.



*Figure 1-51 EXN2000*

The EXN2000 is identical to the N3700 chassis except that the slot that holds 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, which is illustrated in Figure 1-52 on page 63.

The maximum number of drives per shelf is unaffected by the capacity of the individual drive modules. We do not recommend that you mix drives of different capacity in the same shelf 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 1-52 Rear of EXN2000 schematic of arbitrated loop versus switch hub

### Switched Hub architecture

Switched Hub architecture, shown in Figure 1-53 on page 64, has the benefit of additional availability, boosted performance in high I/O environments, and more powerful diagnostic abilities. Figure 1-54 on page 64 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 was 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 because that traffic no longer flows through each disk drive. The hubs are capable of local communication to the disk drives and then more efficiently conveying 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 and are hot swappable, which allows maximum availability.

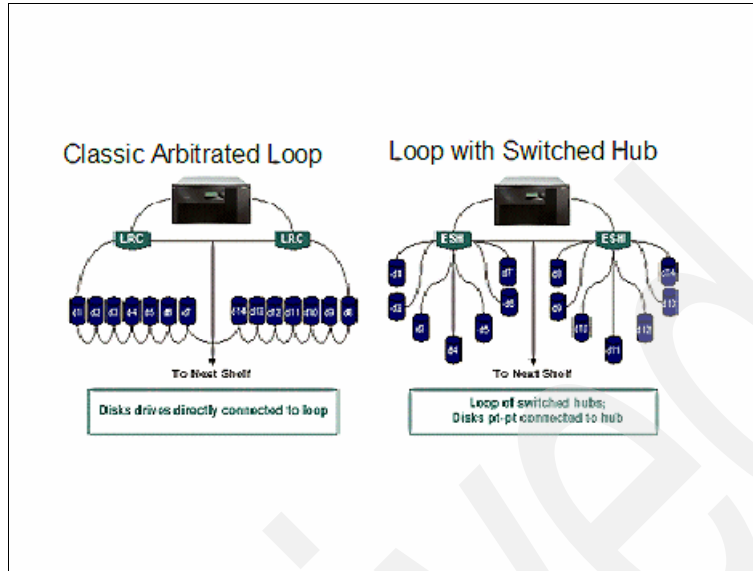


Figure 1-53 Switched Hub architecture

Figure 1-54 shows the external ports on ESH2 modules.




Figure 1-54 External ports on ESH2 module

EXN2000 was withdrawn from marketing in May 22, 2007.

### 1.9.3 EXN1000

The EXN1000 uses the same shelf and hardware that the EXN2000 and EXN4000 uses, so it has the same dimensions. The EXN1000 also supports the same number of disks per shelf 14. The main differences are:

- ▶ Drive type SATA versus Fibre Channel
- ▶ Interface module the AT-FCX versus the ESH2

Figure 1-55 shows the EXN1000 expansion unit.



Figure 1-55 EXN1000 expansion unit

AT-FCX refers to the controller module, shown in Figure 1-56, of the SATA storage expansion unit.



Figure 1-56 AT-FCX module



Data ONTAP supports up to 400 RAID groups per storage system or cluster. When you configure 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.

## 1.9.4 EXN4000

The EXN4000 uses the same shelf and hardware that the EXN2000 uses, 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, 4 Gb loop speed when it works with EXN4000. The ESH4 has LEDs that indicate whether the module is functioning normally (Figure 1-59 on page 67), whether there are any problems with the hardware, and the loop speed operation of the EXN4000. The main differences in the EXN4000 and the EXN2000 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

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 1-57 shows the EXN4000 expansion unit.



Figure 1-57 EXN4000 expansion unit

Figure 1-58 on page 67 shows the 2xESH4 and 2xPSU fans.



Figure 1-58 2xES4, 2xPSU/Fans

Figure 1-59 shows the location of the LEDs for an ES4.

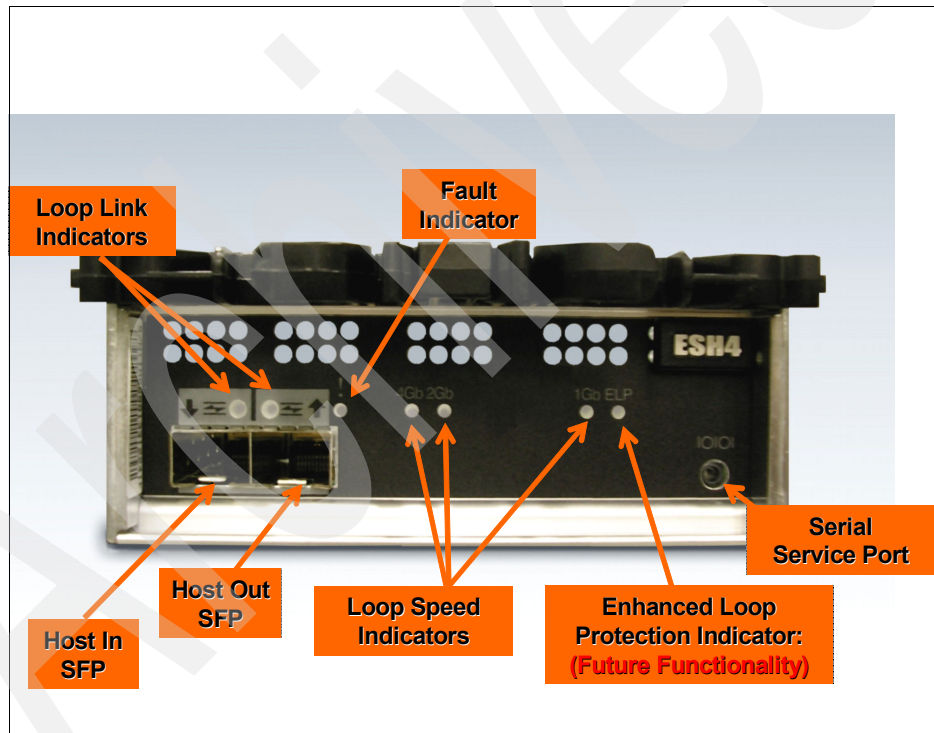


Figure 1-59 Location of the LEDs for an ES4

EXN4000 is the replacement for the EXN2000 FC Storage Expansion Unit.





# Introduction to Digital Video Surveillance

Given a choice, most law-abiding citizens probably prefer not to regard physical security as their top priority. Today, however, virtually every municipality, agency, educational institution, transportation center (including airports and seaports), utility plant, and medical center must plan for threats and disasters. They also must set procedures in place to help safeguard the lives and property that are entrusted to their care.

Video surveillance has long been available as a physical security measure. Placing video cameras in sensitive or strategic areas of their premises, organizations have utilized closed-circuit television (CCTV) to monitor activities, both as a deterrent to crime and as a record of the movements of people and property. Many have also utilized mobile methods of video surveillance, such as using cameras mounted in patrol cars or other security vehicles, to record events. Today, video surveillance remains as vital as ever, but it assumes a new role. As a consequence of both advances in technology and increasingly complex security requirements, video surveillance must be viewed as a key component of a comprehensive security program, rather than as an end unto itself.

Figure 2-1 on page 70 shows Digital Video Security (DVS) as a key component of a security structure.

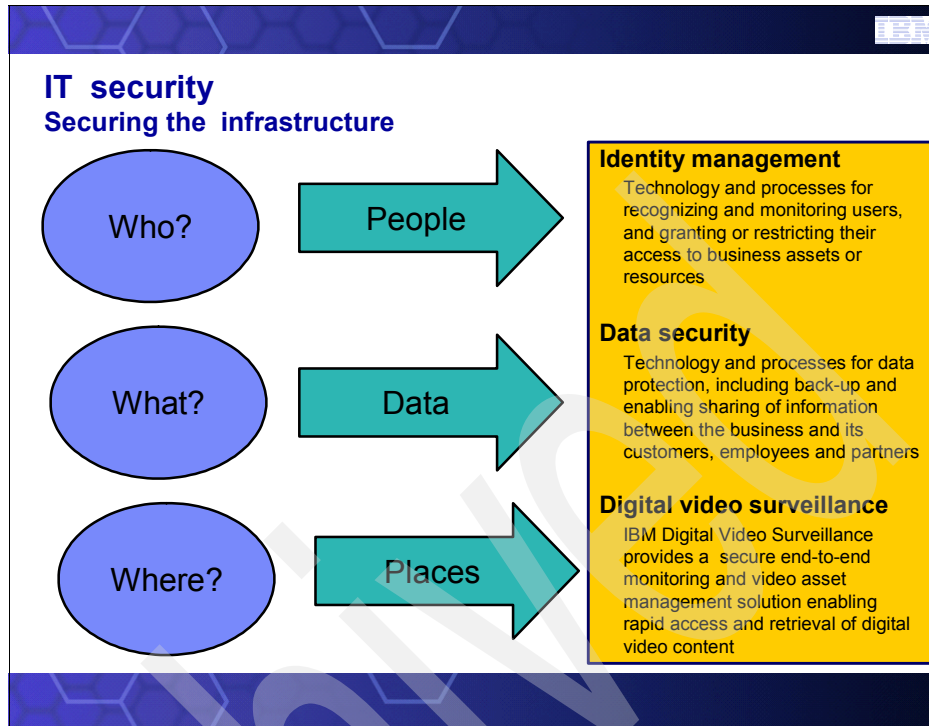


Figure 2-1 DVS as a key component of a security structure

## 2.1 Video as information: the rules have changed

In a free but litigious society, the principles of risk management dictate the need to develop multifaceted strategies for incident prevention, mitigation of damage, and proactive management of events and their aftermath. Both man-made situations and natural disasters require these strategies. What has become evident, however, is not so much the specific need for video footage, but a growing need for a wide range of decision-support information — in the squad car, the security guard's office, the administration building, the firefighters' or SWAT team's command center, the ambulance, and the courtroom.

How big a role that captured video plays in an organization's security program depends, today, upon how well the video capture and management system can bring information to the right people, at the right time, by answering questions, such as:

- ▶ How much ancillary data can be captured with each video frame?
- ▶ How well do the data capture and data management functions integrate with other systems and agencies?
- ▶ What level of advanced features and controls is implemented to help capture the most meaningful data?

It should be evident by now that video surveillance is becoming an information technology (IT) issue. As Figure 2-2 shows, physical security measure can be integrated into IT infrastructure, and we can gain more valuable information from the video images.




Figure 2-2 Integrate the physical security and IT technology

The sooner that the IT staff is engaged in the process of selecting an organization's video capture/video surveillance solution, the more value can be

designed into the system. Indeed, even the highest level of management in an organization might recognize the need to become involved.

## 2.2 Digital Video Surveillance

Current users of video surveillance understand the potential evidentiary benefits of video footage in legal proceedings. The main attraction of digital data collection for these current users is that the image quality is superior, as shown in Figure 2-3, in comparison to traditional videotape with digital frame resolution at 500 lines (versus 240 for videotape). The fidelity of digital images generally remains constant as data is transferred and copied, as well, without the deterioration in quality that often comes with traditional videotape, which is typically reused many times. The expectation is that users of digital surveillance will ultimately find that digital video capture can help improve conviction rates and speed disposition of court cases, it can help deter illegal activity by its very presence, and might help inhibit false claims from being presented. Indeed, digital video surveillance holds the promise of potentially reducing court backlogs because high-quality video evidence might increase the number of cases that are settled out of court and also the promise of potentially reducing the risk of lawsuits in the first place.

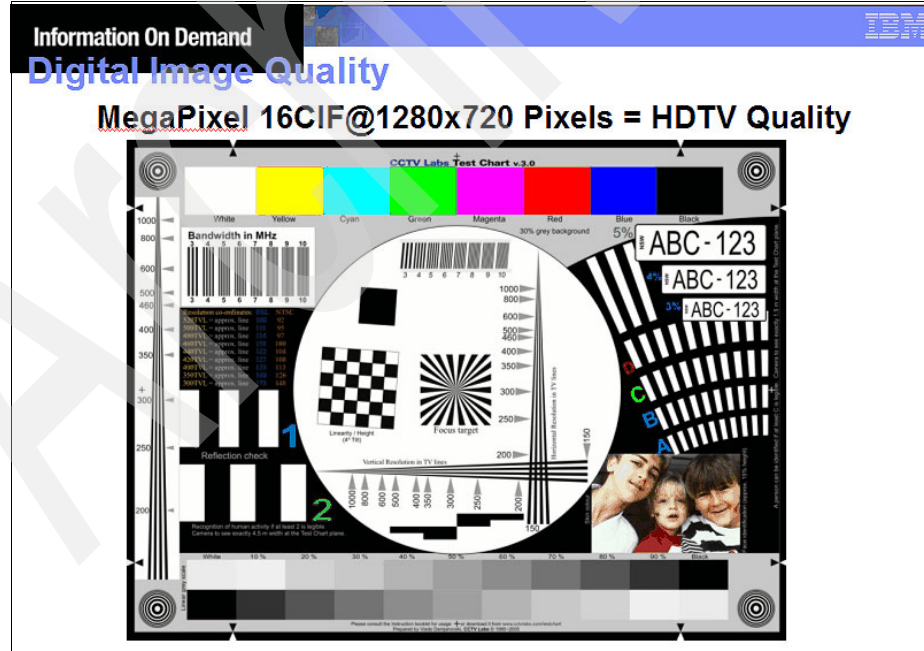


Figure 2-3 Digital image quality

Digital data collection can provide users with numerous tools to help build a complete picture of the environment on which events have occurred or are taking place. In a networked environment, it also might provide the opportunity for a team response. A logically integrated digital environment can enable premises guards and local law enforcement agencies to simultaneously view data. It permits watermarking and the capture of ancillary data, such as biometric information, along with each frame. Saved in centralized storage or archive libraries, digital data can be easily located using simple search methods. Digital surveillance data can also be smoothly integrated with access control and screening systems to provide a unified “command center” view of a single location or complex facilities network.

Using an integrated digital video surveillance system, you can create a virtually seamless community that is responsible for monitoring premises, responding quickly to crises or security breaches in order to better mitigate damage and help save lives, and follow through with the necessary evidentiary documentation of incidents.

An important point to note is that an integrated digital solution can take advantage of wireless technologies to transmit data, even from remote areas. In fact, wireless data transmission is sometimes the only feasible means of reaching the field. The incident of massive wildfires in 2003, for instance, showed the value of gathering footage on the scene and transmitting it live, along with additional critical environmental data, to the personnel in command of the firefighting efforts. The system enabled remotely dispersed emergency workers to view activities across the site, from different points-of-view, and to use that visual data to plan, coordinate, and allocate the use of resources.

Although streaming live (real time) video requires high bandwidth, and might prove too costly or difficult in certain situations, you can use wireless transmission routinely in mobile digital video surveillance systems. Where the infrastructure enables, you can use wireless technology to transfer data from a hard disk that is deployed in a vehicle (such as in a school bus or police squad car) directly to a main storage network, which frees up the mobile disk space for reuse. Enabling wireless capability in the design phase of a digital video surveillance system is a future-ready strategy. Increases in bandwidth that are expected to occur over time, which includes capacity that third parties add, can improve the cost structure for wireless transmission and therefore increase the range of cost-effective uses.

Figure 2-4 on page 74 illustrates wireless IP cameras.

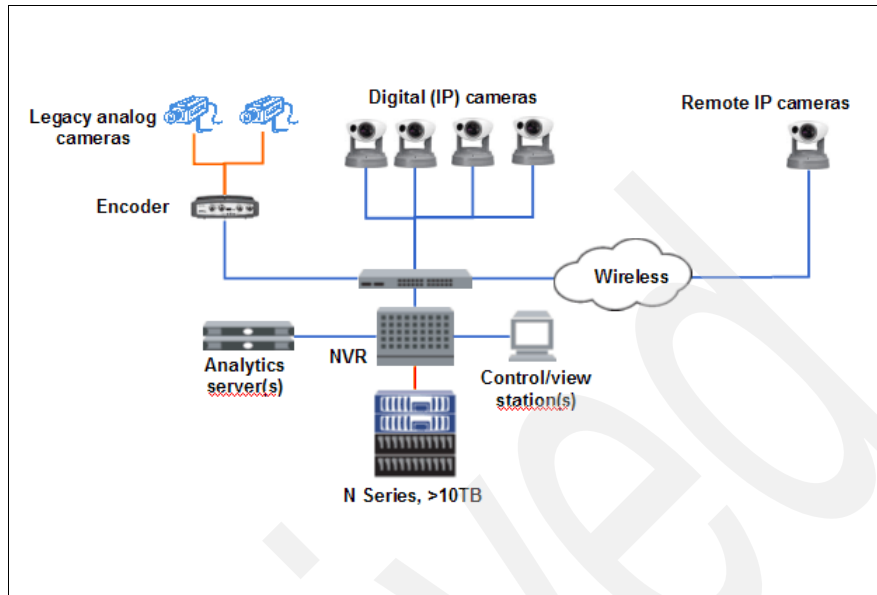


Figure 2-4 Wireless IP cameras

## 2.3 Analog video surveillance: do not throw out your video cameras

Traditionally, videotape has been the medium that we used for capturing video surveillance footage, in both mobile and stationary environments. Using videotape involves familiar equipment, such as fixed mounted analog (nondigital) video cameras that each feed to a standard video cassette recorder (VCR). The VCR receives video input and captures data onto analog videotape. The VCR uses a tape format called video home system (VHS) and has a screen resolution at 240 lines. Figure 2-5 on page 75 shows a traditional video surveillance model. Analog video systems typically use a standard set of procedures, which includes removing and replacing the tape in each camera at specific intervals, storing and cataloguing an ever-growing volume of used tapes, and possibly reusing those same tapes after a prescribed retention period elapses.

## Traditional Video Surveillance System

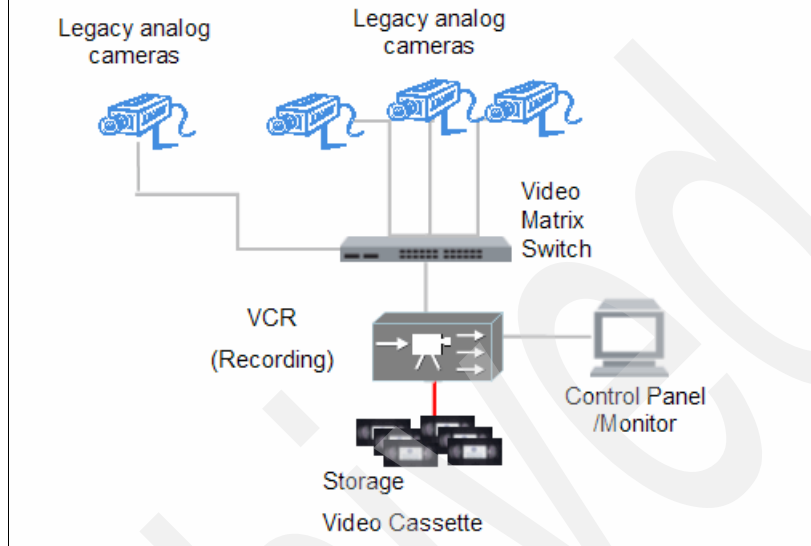


Figure 2-5 Traditional Video Surveillance System

Current users are very familiar with the time-consuming and labor-intensive tasks that are involved in locating specific footage in a videotape archive. Depending on the size of the videotape library, maintenance can involve several people, and once the desired tapes are located for use in legal proceedings, other staffers might be assigned to review the footage and identify the appropriate clips. Users have also learned that the quality of videotape images deteriorates over time and with tape reuse. In spite of its shortcomings, however, traditional videotape-based surveillance has been a valued and valid mainstream tool for many law-enforcement and other agencies.

The good news for current users of analog video surveillance systems is that you can incorporate your existing cabling and analog cameras into a digital surveillance solution. An automated, seamless process is available to convert analog input into a digital format that enables you to take advantage of virtually all of the features of a digital system after the digital surveillance infrastructure is in place.

## 2.4 Building a digital video strategy: choosing future-ready technology

As the capabilities of digital hardware, software, and networks continue to grow at a rapid pace, keeping your systems from becoming obsolete remains an important challenge. The keys to a “future-ready” video surveillance system are:

- ▶ Purchasing components based on open standards, for example, avoiding proprietary hardware
- ▶ Deploying all communications (sending and receiving data) over a converged network that can handle voice, video, and data in a single, high-bandwidth environment
- ▶ Integrating new and heritage systems to deliver the right information to the right users in a timely, streamlined fashion

### 2.4.1 Open standard

Hardware that is proprietary or limited to a single use might hinder your ability to scale your system to changing circumstances. Many video surveillance systems users find, over time, that their constituents want more areas or events to be monitored. Mayors, police chiefs, town and school administrators, airport and seaport executives and others, seeing the value of initial surveillance efforts, typically put forth requests for additional coverage after the systems are implemented. In an ideal video surveillance system, it should be easy to replace or upgrade system components because using components that are built on open standards facilitates such system changes.

### 2.4.2 Converged network

A converged network indicates that lots of applications and systems are integrated together, such as data access system, video surveillance system, and telephony system, and so on. All of these applications and systems share network resources. Currently, wide area network (WAN), such as Internet, and local area network (LAN) have been dominated by IP network. In the future, carrier-class IP will become a fundamental requirement in the development of IP Next-Generation Network. In other words, most delivery vehicles that service providers use today are subject to change, which also includes video service, telephony service, data service, and even storage.



### 2.4.3 Video information

In closed circuit television (CCTV) system, video content (images) is recorded in videotape. If we want to locate specific footage in a videotape archive, we must have people in a videotape library place to do a lot of searching, which depends on the size of the videotape library and its index system. It is time-consuming and labor-intensive work. But if we convert video content into video information in IP streams, then we can take advantage of this video information. We cannot let just the surveillance software help us to search; instead, we must also integrate other applications to gain more value from the original video content.

## 2.5 Converged network infrastructure

Let us discuss converged network in more detail. Converged network, also called “unified network” or “all-IP (Internet Protocol) network,” is a converged network infrastructure that carries all transmissions, including digitally encoded voice and video, as data that can be transported over one network. Converged network represents an evolution of communications technology where everything is treated as data to help deliver greater cost effectiveness, increased capability, and better resource management. Here are some characteristics and benefits of a converged network:

- ▶ A converged network is built of hardware and software that conforms to open standards and is readily expandable and compatible with a multitude of devices.
- ▶ A variety of security-related data streams, including the digital video itself, data from security sensors, and command messages can share a common IP-based network infrastructure.
- ▶ With convergence, system administration can be consolidated and centralized, which makes it more cost-effective, with reduced ongoing maintenance requirements compared to separate infrastructures for voice, video, and data transmission.
- ▶ A converged network infrastructure can enable the delivery of information through wired or wireless wide area networks, LANs, and over the Internet or intranets, to and from desktop computers, laptops, personal digital assistants (PDAs), mobile command units, cameras and other special-purpose devices, which includes those that might run business or security processes, such as access control.

Figure 2-6 on page 78 illustrates a converged network.

## Converged network in DVS system

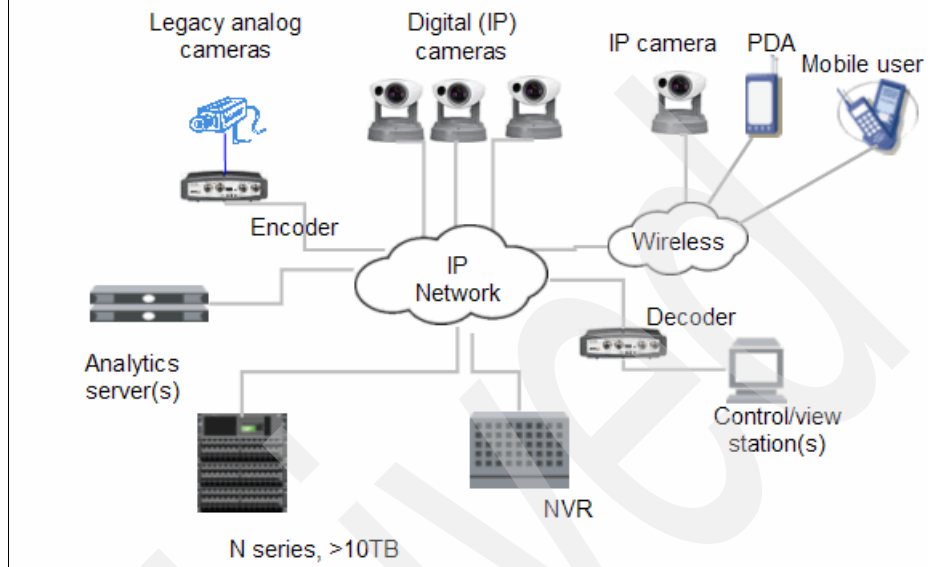


Figure 2-6 Converged network

In many cases, an organization's existing data network components can be deployed in the implementation of digital video surveillance.

### System and process integration

Knowing what information you want to derive from captured data is the key to understanding your requirements for which systems and processes need to be integrated with video surveillance in your organization for the greatest value. Converged networks are highly robust and enable you to gather information from numerous disparate systems and concentrate that information in a single view for decision-making. At the same time, digital video surveillance over converged networks can provide the decentralization and collaboration capabilities needed to enable you to view the same activity simultaneously so that you can take appropriate coordinated actions. Finally, a converged network can be sufficiently high-bandwidth and robust to support a full range of unrelated activities for thousands of users.

## 2.6 Advanced system capabilities

The minimum capability that is expected from a video surveillance system, analog or digital, is the passive collection of video surveillance footage, which you can store, retrieve, and potentially use as evidence, as needed, for investigation or prosecution when an illegal or questionable event occurs.

However, some systems offer capabilities that far exceed those expectations. In addition to capturing the name and serial number of the duty officer or guard, the date, time, and other environmental data that is pertinent to an incident, the end-to-end digital video surveillance solution that is available from IBM (Figure 2-7 on page 80), for instance, can be enabled to support the following advanced capabilities:

- ▶ Motion detection: detects and alerts staff to motion in user-defined regions
- ▶ Directional motion: detects and alerts staff to the presence of objects moving in a user-defined direction, such as cars moving the wrong way on a one way street
- ▶ Abandoned object: detects and alerts staff to the presence of abandoned objects in a user-defined region, such as unattended packages that were left in the monitored area
- ▶ Object removal: detects and alerts staff if one or more objects are removed in a user-defined region and can monitor high-value assets, such as art exhibits
- ▶ People count: counts the number of people in a user-defined region. Since certain crowd conditions can signify trouble or the need for increased personnel on the spot, it alerts staff to inappropriate people count
- ▶ Camera move/blind: detects and alerts staff if a camera is blinded or moved to a different direction, which might indicate tampering
- ▶ License-plate recognition: locates and reads a vehicle's license plate in order to grant or deny premises access to specific vehicles and to track their movement

Figure 2-7 on page 80 illustrates the IBM solution.

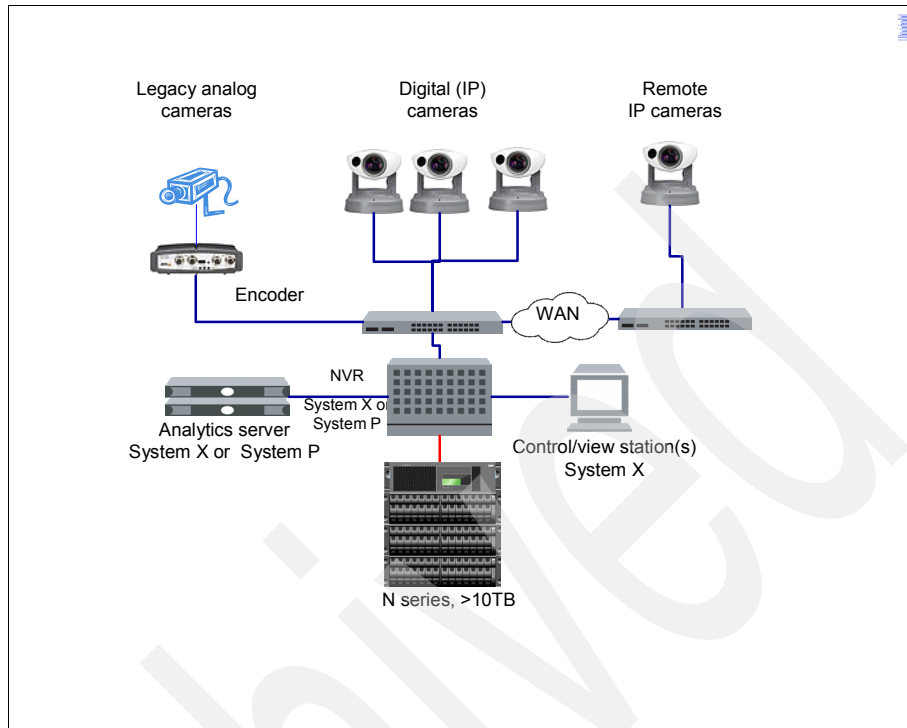


Figure 2-7 IBM solution

These advanced capabilities can reduce the amount of personnel that you need to monitor premises. This is called *smart monitoring*, which is surveillance that can trigger alerts or alarm conditions where it might otherwise take security longer to discover that something is amiss. Figure 2-8 on page 81 shows the improvement we get from video surveillance's new capability.




Figure 2-8 Video Analytics-enabled Surveillance System

Object recognition, as reflected in the advanced features that we previously described above, is a close cousin of biometrics as we know it today. Biometric values are stored sets of unique identifiers, such as fingerprints, voiceprints, irises, facial features, and other identifiers that are likely to be added and gain acceptance over time. Object recognition is possible by databases that contain the size, shape, and mass parameters of a large variety of objects, such as humans, animals, vehicles of different makes and models, packages, specific paintings on the walls of a museum, and so on.

Integration with these object and biometric databases can help make a video surveillance system even more useful— both as a deterrent to illegal activity and as a proactive tool for physical security and premises access. Integration, in this case, implies that the surveillance system:

- ▶ Captures the object or biometric data while recording the activity
- ▶ Automatically synchronizes this data with the activity footage
- ▶ Shares this data among agencies that monitor or review the activity
- ▶ Responds proactively with feedback that indicates a match with the database

All four functions must be enabled for the biometric system to contribute to the capture and detention of wanted persons.

## 2.7 Systems, processes, and policies all contribute to the solution

With an understanding of the range of capabilities of today's sophisticated digital video surveillance systems, you can make your own choices concerning system configuration and the policies that govern system use, for example, when discussing storage options for a digital video surveillance system, organizations must consider the policies that govern their data retention:

- ▶ Is there a minimum video data retention period required by law or regulation?
- ▶ Beyond the regulatory concerns, what are the organization's own needs for retention? From a practical viewpoint, how long should data remain in short-term storage and how long in the archives?
- ▶ What are the cost considerations that apply to library management personnel versus more full-featured, automated storage systems (that is, those that transparently move the data from short-term to long-term storage and delete data at the expiration date as opposed to performing these tasks manually)?
- ▶ What happens to data volume when more data capture or data integration features are deployed? How much additional storage is required if a screening or biometric database is also deployed? Can some or all of the biometric data be served using a third-party or an Internet application, removing it from local storage?
- ▶ How can the appropriate configuration, including people, storage hardware, software, and policy be achieved?

Actually no two video surveillance deployments are identical. A different set of questions apply for each aspect of the solution design. Generally, video surveillance customers have some common concerns. Understanding these concerns will give you a lot of help in video surveillance design and deployment:

- ▶ The high maintenance and high costs of analog VCR system.

A typical analog video cassette recorder, industrial-based cassette recorder solution is very high maintenance. You have probably seen pictures of the back rooms with guards and an entire wall full of cassette recorders. The tapes have to be stored. You cannot really transmit the video like you could in a digital format. But the main issue here is high maintenance. We are seeing customer transition out of analog into digital, and high costs for the equipment, the upkeep, and the labor drive this transition.

- ▶ Improve physical security measures.

After September 11, 2001, securing your physical assets became just as important as securing your IT assets—people, facilities, and infrastructure. Security is only as good as the weakest link. We need to protect key sites,

such as petrochemical facilities and utility plants, from outside intruders; therefore, customers are looking to improve the security of their sight. How can they improve security, from the physical sense, to provide more safety?

- Can video data be shared across a converged network?

With an analog video cassette recorder solution, you cannot really share your video surveillance data; however, with an IT-based network solution, anybody, anywhere can access that video, which is a key issue.

Organizations want to be able to view their sites across multiple states. In a metro environment, the police want to be able to see what is happening in the schools, on the buses, or in the subways. So the pain point is that the authorized remote user knows there is a video record, but that user cannot access it remotely.

- Will video analytics be more effective than analog video system?

Organizations, the security and business side, want to turn video data into useful information that they can use to manage their business. So the question is how can you mine this terabyte of data to provide useful information? Can you use it to better manage shelves in the store? Can you use it to provide better service by looking at the number of people in line and how long they have been there? The queue length and providing automated alerts that might bring people from other parts of the store to the cash registers to alleviate the line or in a bank to provide quicker service to high level, high executive, high-powered individuals who have money to spend. So companies want to turn video data into useful information.

- Who will provide the video surveillance service to them?

We talked about on demand, open system, open architecture, and open standards. The physical security side of the house and the IT side of the house want to bring the same standards to video surveillance. Why? Because the IT department is going to be supporting these applications. They do not want to have 20 different solutions, different protocols, and different standards on their network because then they cannot appropriately and cost effectively support the network. And now it will not scale. So companies want an open and standards-based approach to video surveillance.

- Is the network performance being impacted severely?

Ensure that when you start transporting video across the network, the administrator can manage the impact in the network and can ensure the quality of service for the critical business applications and the performance for real-time video. One of the issues early in the transition from analog to IP-based network video was latency. IP applications have latency, and you must really design the network appropriately so that you do not get flutter and jitter and all of the issues that degrade the video, which are the same issues

on an IP-based phone. So managing the impact in the network is another pain point, ensuring quality of service.

We just discussed several important factors to consider in designing a system and processes to meet your organization's specific requirements. The elements that require review and input generally fall into the categories of data capture, data transfer (also called data transport and aggregation), data storage, application and data integration, and data management (or user interface, which includes generating reports from the system).

## 2.8 The IBM Digital Video Surveillance solution

As a leading IT provider, IBM developed an end-to-end solution for digital video surveillance that is worth examination by any organization struggling with the question, "What should we do to implement a future-ready video surveillance program?" The IBM Digital Video Surveillance solution provides a secure, cost effective, end-to-end monitoring and video asset management solution that focuses on:

- ▶ Networked video surveillance: Current IP network deployed globally. IBM DVS solution will take advantage of it. Video surveillance system operates over the IP network (internet).
- ▶ Physical and logical security convergence: Video information can be centrally managed and can be shared in control.
- ▶ Making the video more useful: You can use the video information for live surveillance and for interactive functions, such as biometric identification.
- ▶ Back end infrastructure optimization: Provides effective and cost effective video usage and storage solution. You can easily access and share the video information.
- ▶ Standards-based approach: Proprietary protocol locks you into a single-vendor solution and decreases the video surveillance system's scalability.
- ▶ Open architecture and design: Open, standards-based solution that can scale and adapt to new technologies.
- ▶ Enterprise Security Framework: Integrates physical and logical security to effectively protect people, facilities, data, and applications.

Figure 2-9 on page 85 illustrates an IBM DVS solution.



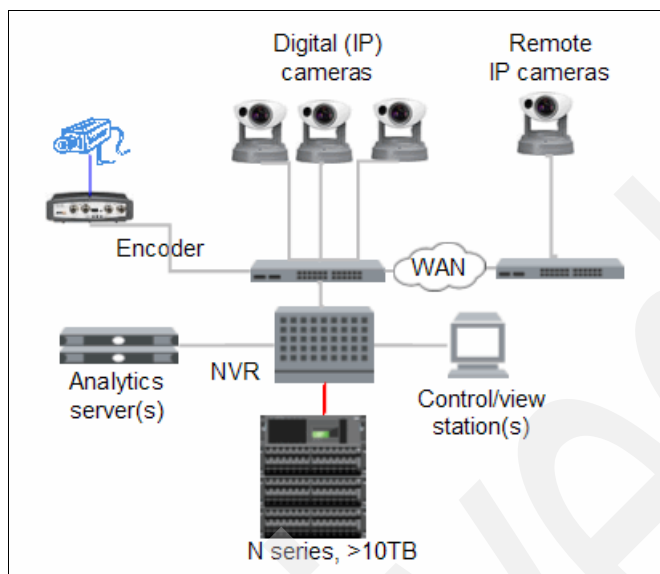


Figure 2-9 IBM DVS solution

The basics of the digital video capture and management solution from IBM, for either mobile or stationary installations or both, are:

- ▶ Data capture
- ▶ Data transfer
- ▶ Storage
- ▶ Application and data integration

### 2.8.1 Data capture

Data capture indicates how to capture video data, which includes the heritage analog camera and digital (IP) camera:

- ▶ The analog camera has an encoder server that converts analog video data into IP traffic. The video IP stream is then transferred to the Network Video Recorder (NVR) server through an IP network.
- ▶ The digital camera records video in disks. Remote IP cameras transfer recorded video-to-video storage system through an IP network.

Table 2-1 on page 86 shows the difference between the analog camera and the digital camera. They have different scalability and different storage methods.

Table 2-1 Analog camera versus digital camera

Video type, cabling	All analog	Mostly analog	Mostly digital	All digital
<b>Cameras</b>	<10	10-50	30-200	100+
<b>Storages</b>	VCR <sup>a</sup>	DVR <sup>b</sup> /NVR <sup>c</sup>	NVR	NVR
<b>Retention</b>	7-30 days	10 ~ 90 days	30 ~ 180 days	> 90 days
<b>Maintainer</b>	Security Dept	IT and Sec.	Usually IT	IT

a. VCR indicates Video Cassette Recorder.

b. DVR indicates Digital Video Recorder.

c. NVR indicates Network Video Recorder.

Small scale video surveillance systems have limited cameras, a low requirement of retention time, and the Security department maintains the video surveillance system. Larger scale systems, especially with more digital cameras, can support more cameras and more retention days.

In the DVS solution, you can utilize the analog camera, but the video data must be digitally encoded for integration by the encoder device. You can also analyze the digital video data and transfer it directly to a NVR server using digital cameras or “Pan-Tilt-Zoom” cameras.

## 2.8.2 Data transfer

After you capture the video data, you must transfer it to a centralized storage device over the IP network. After the video data is captured, it needs to be transferred to a centralized storage device over the IP network, which includes the switch and router. Can the existing network infrastructure support the transport of video content, which includes bandwidth, utilization, burst capacity provisioning, and QoS of enterprise applications.

Currently, 100 Mbps links to the user desktop is very common. The effective throughput of each of these links is approximately 80 Mbps. Normally MPEG-2 creates about 4 to 5 Mbps of video stream from a 30 frames per second D1 720 x 480 resolution video camera. One 100 Mbps link can support about 20 high-quality (720x480) video surveillance streams. Many customers are now equipped with switches that support 10 Gbps to the desktop. In these cases, it could be about 2000 high-quality (720x480) video surveillance streams. Presently, most video surveillance deployments only carry hundreds of video streams; therefore, it is unlikely that video surveillance would put a strain on a LAN environment. But with WAN connectivity, there are some challenges, which we can correct by putting surveillance policies and QoS policy in the network to minimize the bandwidth consumption.

### 2.8.3 Storage

For analog video data, we put the data on videotape and we need a lot of maintenance on these tapes, and for this reason we cannot keep a long retention of the video data. However, if video data is stored in digital format, we can put them into disks. With the development of new compression technology and the decreasing of disk price, digital video data could be more and more popular. VHS machines that are used to play videotape are not being made anymore. So in the future, DVR/NVR would eventually replace VCR. In other words, the video data would be stored in the hard disk with digital format.

As a container of video information, video storage plays a crucial role in the IBM Digital Surveillance Solution. You can store digital video data in direct attached disk (DAS), NAS disk, or SAN disk. The DAS method has a limitation of its disk capacity. DAS also lacks redundancy and is hard to manage in large deployments. The SAN method has good performance, but it is expensive to deploy and requires special SAN-skill maintainers. The NAS disk, like the N series, can leverage current IP resources to store video data, and it is also easy to maintain. The IBM System Storage N series can provide a flexible, lower TCO, unified storage platform:

- ▶ N series storage appliance that is capable of high-performance FC SAN and cost-efficient IP SAN (iSCSI) attachment. Buy what you need and grow on-demand
- ▶ N series storage appliance can be a centralized storage system that is designed for short-term retrieval by applications
- ▶ N series storage appliance can be a solution for long-term retention or archival

### 2.8.4 Application and data integration

A surveillance system application includes a suite of server-based technologies for management, distribution, and storage of video surveillance data in a network environment. A network surveillance system is comprised of up to thousands of cameras and viewers. The application is the foundation for digital video surveillance because it enables distribution of video in the right format for multiple users throughout an extended enterprise and enables authorized users to access, view, and control cameras. It also gives managers limitless scalability and options for storage and retrieval of archived video data. To archive, we need hardware, software and middleware to make the captured video and environmental data more readily available to the desired applications.

The video surveillance solution includes:

- ▶ **Media Transcoder:** converts the analog video data into digital video stream.
- ▶ **Application Server:** authenticates and manages access to video feeds and gives integrators an out of the box component to quickly deploy video monitoring applications.
- ▶ **Command Server:** an application for monitoring video feeds in the command center. The command center must have coordinated viewing of multiple camera feeds at one time and the ability to switch manually or automatically between feeds and display patterns.
- ▶ **Media Server:** the core component in video surveillance applications is the Media Server, which enables you to distribute, archive, and manage video feeds. It offers the power and flexibility to meet a diverse range of video surveillance requirements and coexist on an IP network with other IT applications.

Figure 2-10 shows the different servers that are in the DVS system working together to fulfill the Video Management function.

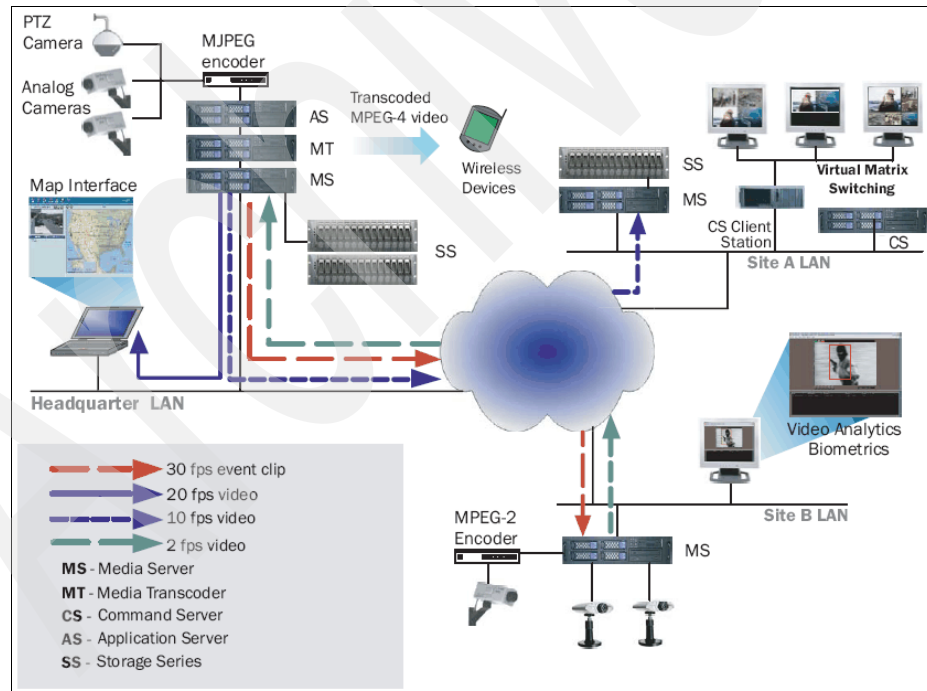


Figure 2-10 Applications servers in the Video Surveillance System

## 2.9 Components of a DVS solution

The Digital Video Surveillance System is a complex system that includes hundreds of cameras, a data transfer system, a video monitor center, a data management subsystem, and a data storage subsystem. Figure 2-11 illustrates the components in DVS system:

- ▶ Camera device
- ▶ Network connection
- ▶ Encoder device
- ▶ Servers
- ▶ Storage device
- ▶ Monitor device

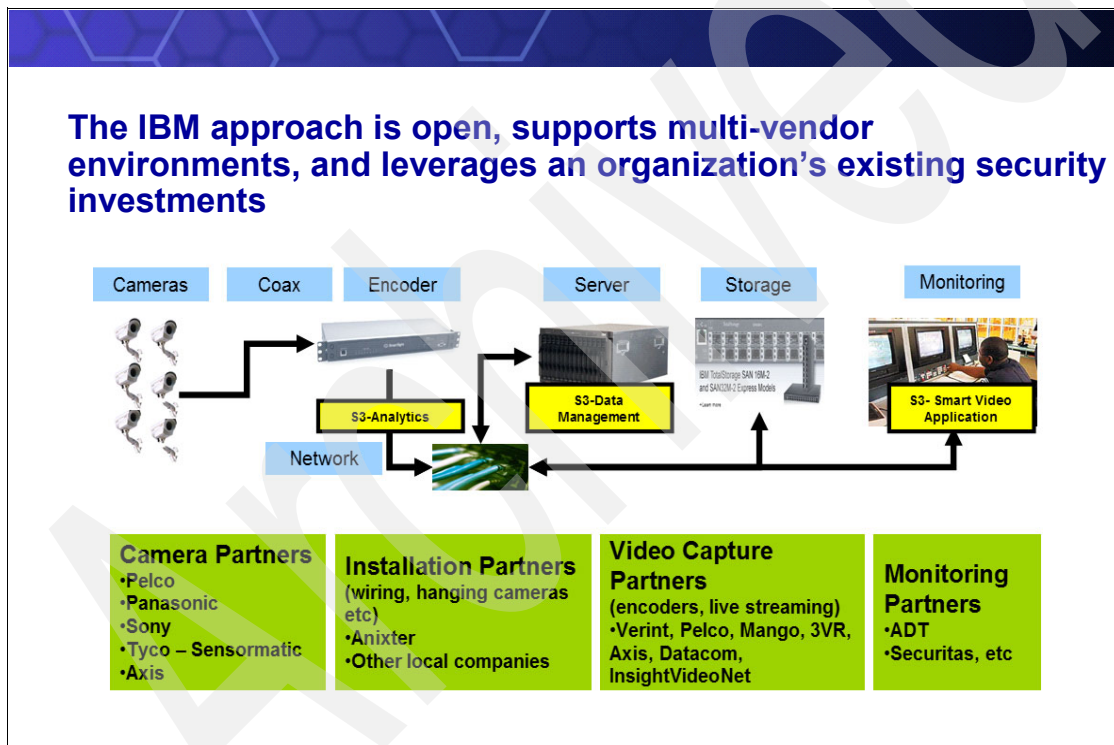


Figure 2-11 Components of the DVS solution

*Cameras*, work as DVS solution's input device. Cameras do video capture work and can determine the video image quality. With better image quality, we can get more valuable information with the help of IT technology. Cameras can be divided into analog cameras, digital (IP) cameras, and remote IP cameras.

Currently, analog cameras are still being used popularly in Video Surveillance systems.

*Network connections* include the cable system and the data switch. Multiple parallel cable plants are necessary to deploy the video surveillance solution:

- ▶ Coaxial cables for traditional analog video transmission
- ▶ Unshielded twisted pair cable and fibre cable for IP network deployment
- ▶ Data switch to provide connectivity between cameras and IP network
- ▶ Router device if the DVS solution is separated between long distance sites

*Encoder device*, encoder/decoder device, helps to convert the analog data into digital IP stream. It supports multiple video codecs, which includes MJPEG and MPEG-4. Also, at the monitor side we need a decoder device to convert digital IP stream into analog video image.

*Servers* are platforms that combine powerful video and audio servers with an open API to allow systems integrators to create customer solutions. A server is a suite of servers that are in charge of management, analytics, distribution, and storage of video surveillance data in a network environment. The server includes an application server, command server, and media sever.

*Storage device* records digital video information. Physically, it is connected to the IP network, but logically it is behind the video surveillance server. The server gets the video information input from the IP network and then records the video information into the storage device. The storage device can also be connected to a sole SAN network, if the server is SAN-enabled or if there is a IP-SAN gateway to help the server to communicate with the storage device.

*Monitor device* is used to view live video. The operator selects the desired video input and specifies where the video is to be displayed. In a DVS system, the monitor device can be separated across the IP network, and the authorized user can monitor video data locally.

## 2.10 IBM Smart Surveillance Solution (S3)

S3 stands for Smart Surveillance Solution. The IBM Smart Surveillance Solution (S3), shown in Figure 2-12 on page 91, turns Video into useful information, which allows customers to realize greater value from their video investment. The IBM S3 solution helps you integrate data from a variety of monitoring devices and apply advanced analytics to improve security, reduce security overhead, and increase organizational flexibility and business intelligence.

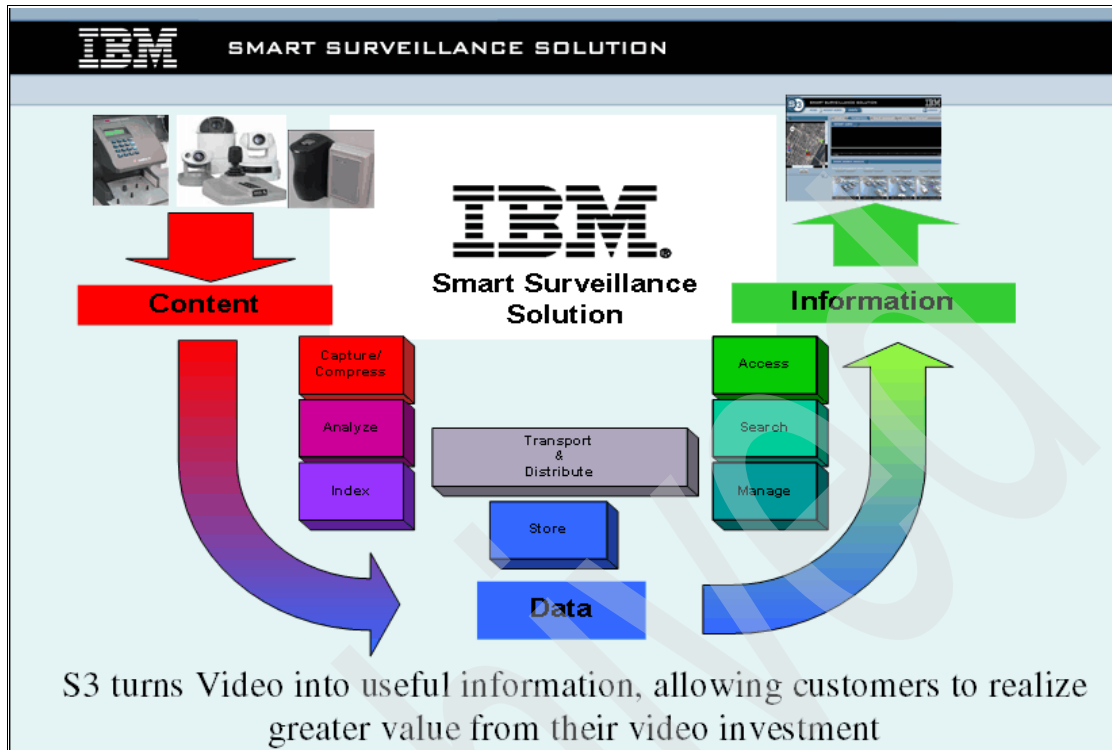


Figure 2-12 IBM S3 solution

The IBM Smart Surveillance Solution helps companies that are migrating from tape-based monitoring to digital video surveillance to get more value out of the video content that they capture and store. The solution can help lower the total cost of ownership for physical security, speed response time, and increase organizational flexibility by transforming footage into valuable business information.

The IBM smart surveillance solution is designed to help you:

- ▶ *Integrate* data from multiple sensors and analog and digital cameras into one searchable system
- ▶ *Search* your data according to a broad range of parameters that include time, date, alert, object, size, location, and color
- ▶ *Share* data across agencies, counties, and departments
- ▶ *Access and view* information remotely from a Web browser
- ▶ *Prioritize* critical events and store them long term

- ▶ *Analyze* your footage for perpetrator identification, theft prevention, damage detection, customer behaviors, and business intelligence
- ▶ *Use* your security data in a variety of contexts, which includes litigation preparation, regulatory compliance documentation, and sales conversion analysis.

Because you do not have to manually search through piles of videotapes and because you can target your search more specifically, you are more likely to detect threats and apprehend perpetrators, which helps to improve security and operations. Regardless of whether your business is located in the public or private sector, the ability to reduce your need for manual review and for redeploying armed guards can help reduce your costs.

Event-prioritization capabilities can help you lower your data management costs. The added flexibility that comes from being able to add new technology or analytical capabilities to the open-architecture solution framework can enhance the agility of your business. And the ability to use your data in a variety of ways to support business decision making can help you get more return from your security investments.

As one of the most experienced systems integrators in the world, IBM leverages its worldwide consultants and a network of established digital video surveillance technology providers to bring you integrated solutions that include cameras, sensors, servers, networking technology, software, services, and analytics. Our proprietary smart surveillance technology allows a wider range of search capabilities and helps you to leverage your data in more ways than virtually any other offering on the market today.

Figure 2-13 on page 93 summarizes the IBM Smart Surveillance Solution functions as:

- ▶ Behavior analysis  
By analyzing video data, you can generate alerts when configurable tripwires are touched. You can also search video images by event attributes and object appearance.
- ▶ License plate recognition  
You can match license plates that are captured by video recorder against a license plate watch-list.



► Face analysis

You can match faces against a face watch-list or capture the face view of people for identification. The IBM S3 solution also considers people's privacy in the video surveillance environment. It has a function that limits access to camera and redact information from video, such as fuzzy metadata representation.

► Event integration

You can generate an alert when configurable events occur, such as sensor events, transaction logs, 911 call logs, RFID events, and GPS metadata.

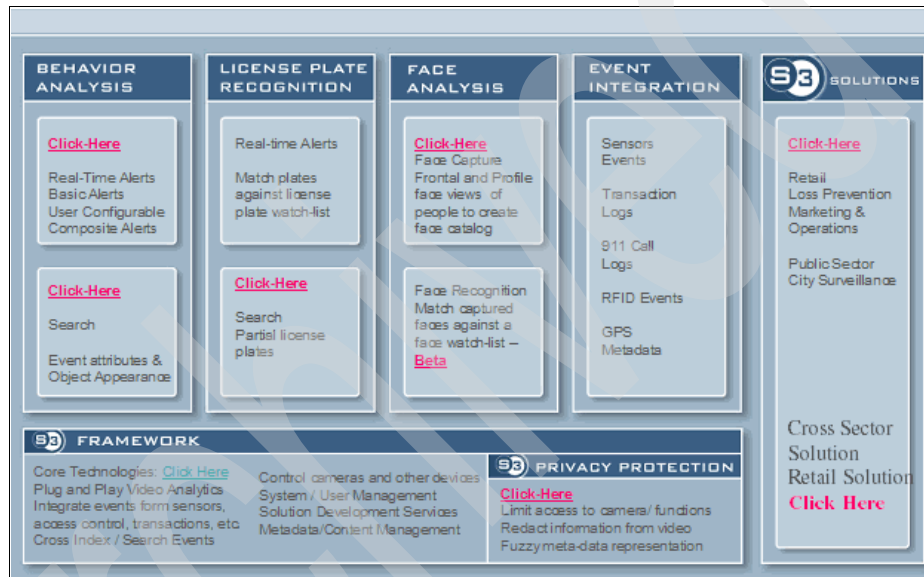


Figure 2-13 Functions of the IBM Smart Surveillance Solution

Figure 2-14 on page 94 gives you an example of the behavior analysis function of the IBM Smart Surveillance Solution.



Figure 2-14 IBM Smart Surveillance Solution Real Time Alerts®

## 2.11 Industries that use Digital Video Surveillance

The IBM DVS solution is a comprehensive, open, standards-based, scalable, highly-available physical security monitoring solution. It enables real-time access and analytics of critical video information and provides secure, cost-effective access to and storage of diverse digital video media.

Who is suitable for the IBM DVS solution? To answer that question, we must first answer some other questions:

- ▶ Who needs a video surveillance system?
- ▶ Do they have an aging and limited video surveillance system already?
- ▶ Do they want to use new technologies to get information from the video data?
- ▶ Do they want to lower the cost while improving the security capabilities?

Yes, after going over these questions, you should have an idea about who needs the IBM DVS solution.

These are industries who have a DVS requirement:

- ▶ *Financial institutions* can use DVS to reduce ATM and in-branch fraud.
- ▶ *Retailers* can add directly to their bottom line by reducing shrinkage from employee theft and shoplifting.
- ▶ *Central Government* can use DVS to revise perception of the two key dimensions of danger and risk: threat and vulnerability.

- ▶ *Manufacturers* can monitor Total Quality Management for complex processes.
- ▶ *Transportation companies* can reduce risk by monitoring hazardous supply transit and handling.
- ▶ *Entertainment venues* can reduce employee or customer fraud and facility vandalism.
- ▶ *Educational institutions* can monitor student facilities to provide a secure environment.
- ▶ *Telecommunication* can use DVS to leverage its facility resource.

## 2.12 IBM digital video surveillance solution business partners

Currently IBM does not manufacturer all of the components for the DVS solutions; however, we team with partners such as, Cisco, Anixter, ADT, Pelco, Bosch, and Panasonic, just to name a few, and we have very good relationships, alliances, and partnerships with the right leading edge, best-of-breed digital video surveillance solution vendors in the market to provide this core technology. Table 2-2 shows the relationship between IBM and our partners in specific areas and in the IBM customer field.

Table 2-2 IBM DVS partners

Capture	Video management	Analytics	Service
AXIS	Cisco	Guardian Analytics	ANIXTER
BOSCH	Dedicated Micros	Cernium	ADT
PELCO	INSIGHTVIDEONET	-----	-----
Panasonic	-----	-----	-----

Table 2-3 on page 96 contains Web sites of some major partners in the IBM DVS solution.

*Table 2-3 Web sites of some major IBM DVS solution partners*

<b>Partner</b>	<b>Web site</b>
AXIS	<a href="http://www.axis.com/">http://www.axis.com/</a>
BOSCH	<a href="http://www.boshsecurity.com/">http://www.boshsecurity.com/</a>
PELCO	<a href="http://www.pelco.com/">http://www.pelco.com/</a>
Panasonic	<a href="http://www.panasonic.com/">http://www.panasonic.com/</a>
Cisco	<a href="http://www.cisco.com">http://www.cisco.com</a>
Dedicated Micros	<a href="http://www.dedicatedmicros.com/">http://www.dedicatedmicros.com/</a>
Guardian Analytics	<a href="http://www.guardiananalytics.com/">http://www.guardiananalytics.com/</a>
Cernium	<a href="http://www.cernium.com/">http://www.cernium.com/</a>
Anixter	<a href="http://www.anixter.com/">http://www.anixter.com/</a>
ADT	<a href="http://www.adt.com">http://www.adt.com</a>
Safety Vision	<a href="http://www.safetyvision.com/">http://www.safetyvision.com/</a>
iMove	<a href="http://www.imove.com/">http://www.imove.com/</a>
DEFENDERtech	<a href="http://www.defendertech.com/">http://www.defendertech.com/</a>
RETAIL EXPERT	<a href="http://www.retailexpert.com/">http://www.retailexpert.com/</a>
ESRI	<a href="http://www.esri.com/">http://www.esri.com/</a>
Qwest	<a href="http://www.qwest.com/">http://www.qwest.com/</a>
VectorMAX	<a href="http://www.vectormax.com/">http://www.vectormax.com/</a>
INSIGHTVIDEONET	<a href="http://www.insightvideonet.com/">http://www.insightvideonet.com/</a>

## 2.13 The Benefits of IBM System Storage N series in the Digital Video Surveillance solution

In the DVS system, storage plays a key role. It records video data, which is the most valuable action in the video surveillance system. Previously, traditional video surveillance systems recorded video on video tape, which needs lots of space to store these tapes and lots of maintainers to maintain tape, index video, and locate data. The analog tape will eventually phase out because video data can be recorded in digital format on disk.

IBM System Storage effectively meets the needs of both security and IT as those needs converge to a more network-centric, digital video surveillance architecture.

When improving surveillance operations is your goal, IBM System Storage is the answer. It has enough power and reliability to support more simultaneous camera feeds, higher resolution video, longer retention periods, more robust analytics, and applications beyond surveillance, all while remaining easy enough for one administrator to manage terabytes of footage.

In video systems, the image quality is expressed by resolution, for example, VGA indicates 640 x 480 pixels, CIF indicates 352 x 288 pixels, 4CIF indicates 704 x 576 pixels, and QVGA indicates 1280 x 960 pixels.

We assume that a camera uses QVGA format to record video images to get high quality of the image, and the system gathers seven frames per second: each frame needs about 55 KB bandwidth to transfer and record and then for 1 second it needs 386 KB. For one month, one camera needs 1.01 TB of space to record video data. If we consider record colorful images, even with new compression technologies (10:1) we still need 300 GB of space to record images for one month for only one camera.

Figure 2-15 on page 98 shows the disk space that DVS consumes.

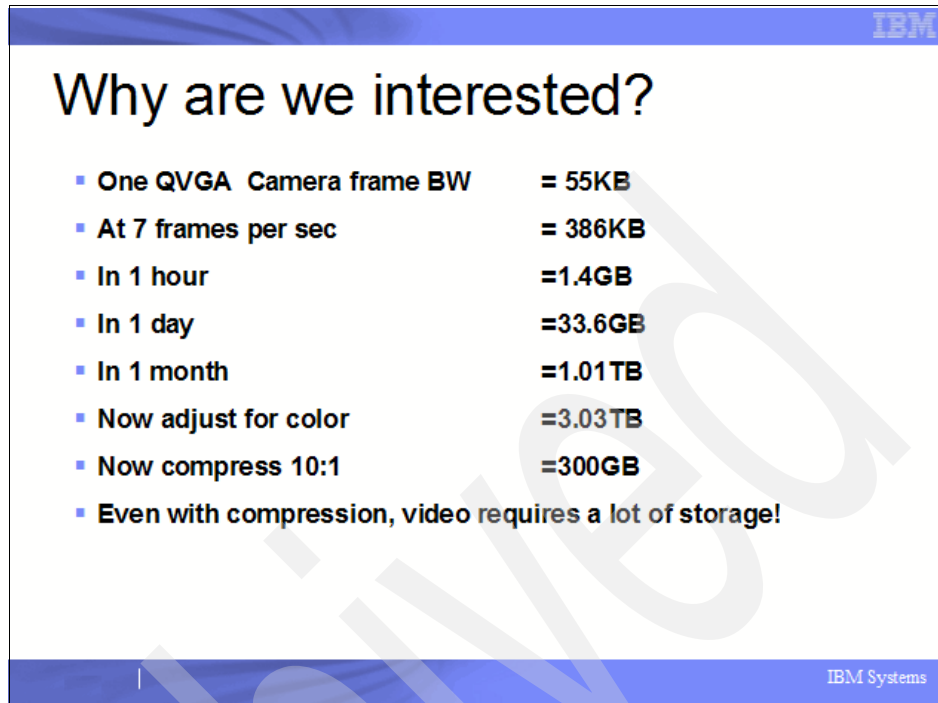


Figure 2-15 Disk space consumed in DVS

Based on the analysis of disk space in the DVS application, we think such growing expectations are being met today with IBM System Storage N series solutions, which is proven to withstand the varied data storage demands of leading regional and worldwide enterprises. IBM N series continues to win high marks in even the most demanding security application environments.

IBM System Storage N series as part of a total DVS solution:

- ▶ More flexible, Lower TCO, Unified Storage:
  - A storage appliance that is capable of both high performance FC SAN and cost efficient IP SAN (iSCSI) attachment
  - Buy what you need, grow on-demand
  - Store other applications, DVS index (database)
- ▶ Enterprise Vendor/Solution:
  - Higher availability = less downtime, lower risk
  - RAID-DP protection against inexpensive disk drive technology
  - Storage leader, Global sales, support and service

- ▶ Performance & Scalability:
  - A storage system that can scale up to 500 TB provides a smaller number of storage footprints to manage a greater number of cameras per server/storage
  - Fast record and overwrite without fragmentation
  - Clustering & redundancy enables customers to expand/upgrade with zero/minimal downtime

The latest evolution in surveillance architectures now builds upon the use of network video recorders that run on standard Windows or Linux servers, in conjunction with analog-to-IP encoders and digital/IP cameras. Using existing Ethernet backbones, NVRs also bring with them the ability to decouple NVR application servers from their underlying storage. In addition to offering better support and more rapid response to emerging events, NVR environments can also reap the benefits from IBM enterprise class system storage: higher availability, better performance, higher capacities to support higher quantities and higher quality video, the ability to centralize and share video among multiple users, and the ability to perform high-speed recording and playback of multiple video streams simultaneously. Used with NVRs, IBM System Storage N series allows video footage to be retrieved and viewed easily by other authorized users in the network while providing robust, steady performance.





## **Benefits of the IBM System Storage N series in the digital video surveillance solution**

Video's quality has a relationship with its storage capacity. Higher quality digital video requires higher storage resources. So extending the storage system is a better idea. IBM System Storage effectively meets the needs of both security and IT as those needs converge to a more network-centric, digital video surveillance (DVS) architecture.

When improving surveillance operations is the goal, IBM System Storage N series is the answer because it has enough power and reliability to support more simultaneous camera feeds, higher resolution video, longer retention periods, and more robust analytics and applications beyond surveillance, all while remaining easy enough for one administrator to manage terabytes of footage.

Network-centric DVS systems are coming of age in the enterprise. Orchestrating the move to DVS is often the Chief Information Officer (CIO), Chief Technology Officer (CTO), or Information Technology Director, who must balance the security or loss prevention teams' needs to record, retain, share, and analyze critical video over the network against ITs own need for low total cost of

ownership (TCO) and enterprise-class availability, reliability, and performance. When every second counts, world-class physical security and loss-prevention teams must rely on quick access to surveillance footage. They expect network-centric digital video to be readily available, of excellent viewing quality, and able to be analyzed by local or remote users. IT organizations that are faced with ensuring a smooth transition to DVS must choose their video repository carefully. Critical to success is a network storage platform that is powerful enough to meet enterprise requirements for flexibility, reliability, and ease of management, while causing no hiccups in continuous recording and video availability. The storage must also scale easily within and across systems as stored footage grows to accommodate higher resolution cameras, use of video analytics, and the corporate reuse of digital video for other groups, such as human resources, marketing, and sales. Such growing expectations are being met today with IBM System Storage N series solutions. Proven to withstand the varied data storage demands of leading regional and worldwide enterprises, N series continues to win high marks in even the most demanding security application environments.

### **3.1 More flexible, lower TCO, unified storage**

In this section, we discuss the benefits of using IBM N series:

- ▶ A storage system that is capable of both high performance FC SAN and cost efficient IP SAN (iSCSI) attachment:
  - The IBM N series connectivity options allow you to connect using NFS or CIFS. This flexibility enables you to adapt to remote and local environments, multivendor environments, and small to large installations.

Figure 3-2 on page 104 shows the N series' multiprotocols.

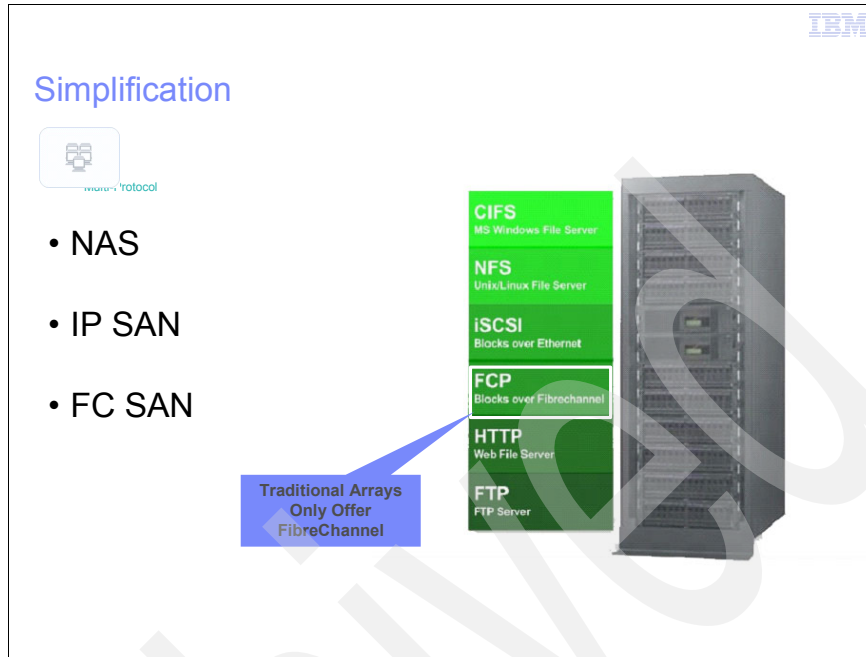


Figure 3-1 N series multiprotocols

- Buy what you need and grow on-demand.
- The N series is not just for DVS. You can use it for storage for other applications, databases, mail servers, file servers, and DVS index (database) file servers. Additionally, it can operate in small-to-large installations with its array of models and storage capacities. Its protocol and block options, such as NFS, CIFS, iSCSI, and FCP also make it adaptable to varying application environments. Given these capabilities, you do not need different storage for different applications because you can use the N series for all of your storage needs.

Figure 3-2 on page 104 illustrates the unified storage.

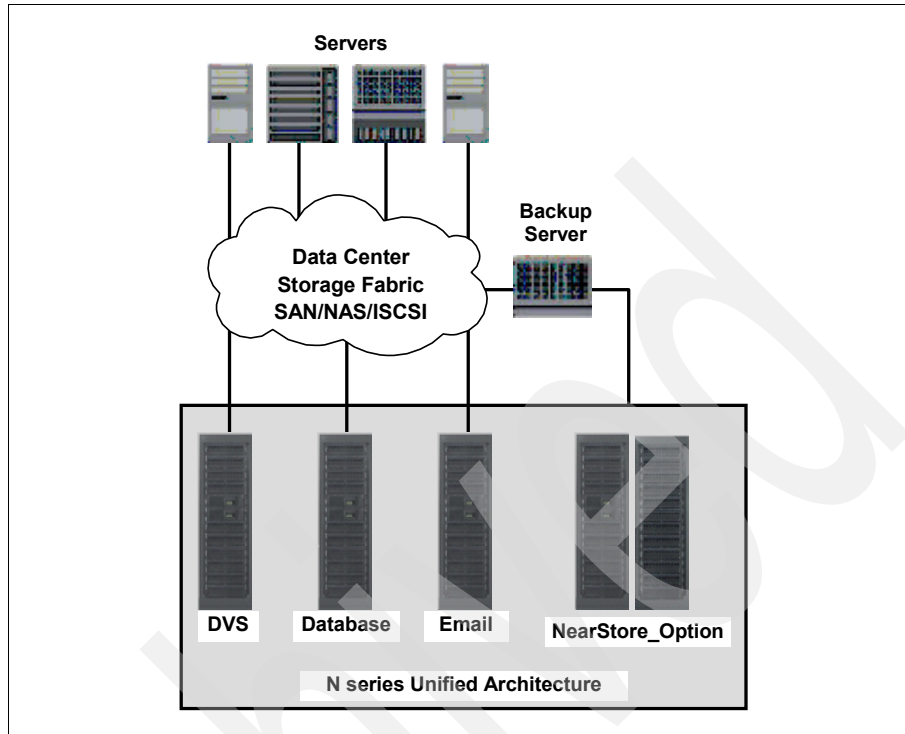


Figure 3-2 Unified storage

### 3.2 Enterprise vendor and solution

The high availability design of the N series and Data ONTAP qualifies it as an Enterprise solution. Some of the highlights are:

- ▶ Higher availability = less downtime and lower risk.
- ▶ RAID-DP protection against inexpensive disk drive technology RAID-DP stands for RAID Double Parity, as shown in Figure 3-3 on page 105, and it significantly increases the fault tolerance from failed disk drives over traditional RAID. When all relevant numbers are plugged into the standard mean time to data loss (MTTDL) formula for RAID-DP versus single-parity RAID, RAID-DP is on the order of 10,000 times more reliable on the same underlying disk drives.

## High Reliability



Figure 3-3 RAID-DP

### 3.3 Performance and scalability

N series is a storage system that can scale up to 500 TB and provides a smaller number of storage footprints to manage a greater number of cameras per server/storage:

- ▶ Given the calculations in the following list, notice how rapidly your storage capacity requirements might grow for Digital Video Surveillance. N series can meet these requirements with its array of products:
  - One QVGA Camera frame BW = 55 KB
  - At 7 frames per sec = 386 KB
  - In 1 hour = 1.4 GB
  - In 1 day 33.6 GB
  - In 1 month 1.01TB
  - Now adjust for color = 3.03 TB
  - Now compress 10:1 = 300 GB
  - Even with compression, video requires a lot of storage

Figure 3-4 on page 106 illustrates N series' scalability.




Figure 3-4 N series scalability

- Fast record and overwrite without fragmentation

### 3.4 Flexible storage for different applications

Where as archived data or video clips might never be accessed and can reside on SATA drives, recently captured data is more subject to recall and analysis, which requires higher response times of fibre drives. With N series and its SATA expansion units, EXN1000 and EXN4000, both requirements are filled with the same storage system, which is illustrated in Figure 3-5 on page 107.

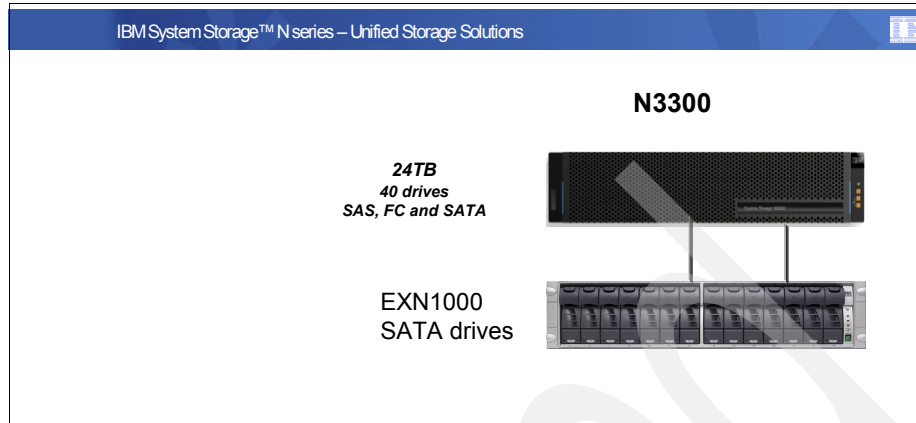


Figure 3-5 N3300

### 3.5 SnapMirror

SnapMirror is a feature of N series that allows DVS data to be replicated between N series storage systems over a network for backup or disaster recovery purposes. After an initial baseline transfer (Figure 3-6 on page 108) of the entire DVS volume or file system, subsequent updates only transfer new and changed data blocks from the source to the destination, which makes SnapMirror highly efficient in terms of network bandwidth utilization. The destination file system is available for read-only access or the mirror can be “broken” to enable writes to occur on the destination. After breaking the mirror, synchronizing the changes made to the destination back onto the source file system reestablishes the mirror.

In the traditional asynchronous mode of operation, updates of new and changed data from the source to the destination occur on a schedule that the storage administrator defines. These updates could be as frequent as once per minute or as infrequent as once per week to match the DVS installation policies. Synchronous mode is also available, which sends updates from the source to the destination as they occur, rather than on a schedule. If configured correctly, synchronous mode can guarantee that data written on the source system is protected on the destination even if the entire source system fails because of natural or human-caused disaster. A semi-synchronous mode is also provided, which can minimize loss-of-data in a disaster while also minimizing the performance impact of replication on the source system. To maintain consistency and ease of use, the asynchronous and synchronous interfaces are identical with the exception of a few additional parameters in the configuration file.

Figure 3-6 on page 108 shows the SnapMirror setup.

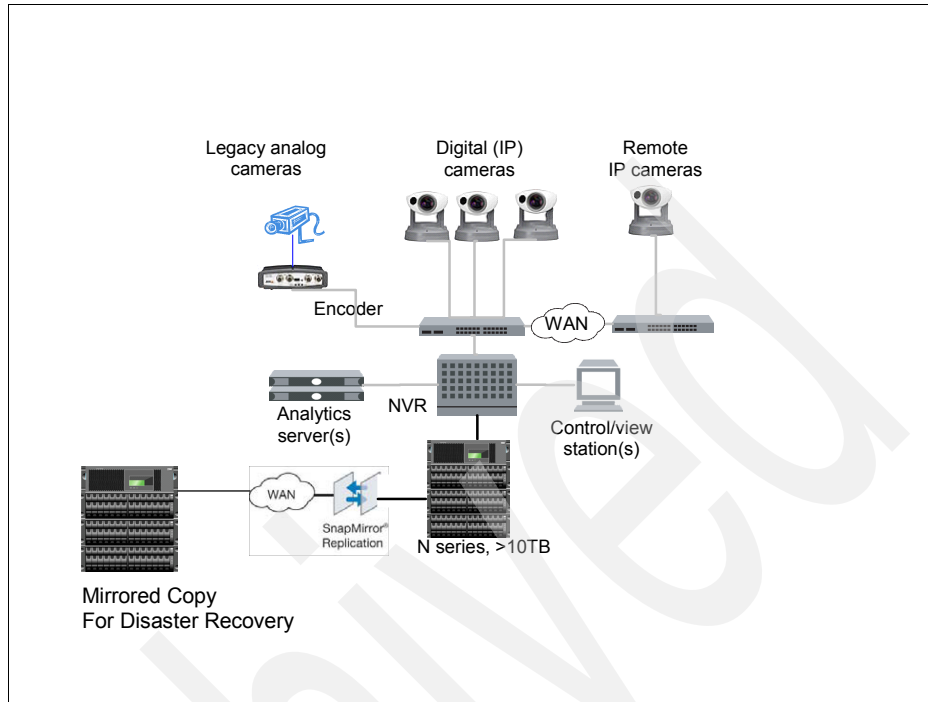


Figure 3-6 SnapMirror

### 3.6 Clustering

The basics of an IBM System Storage N series cluster consists of two nodes that can takeover or failover their resources or services to the associated counterpart nodes. This functionality assumes that each node can access all resources, which that means both nodes must have access to all disks physically (cabling) and logically, as shown in Figure 3-7 on page 109. The N3700 Model A20 and N3300/3600 combines both cluster nodes in one shelf and does not require interconnect cables.






Figure 3-7 Basic cluster configuration

### 3.6.1 Benefits of clustering for DVS solutions

The benefits of clustering for DVS solutions are:

- ▶ Increased capacity for growth.
- ▶ Fault tolerance.
- ▶ When one node fails or becomes impaired, a takeover occurs, and the partner node continues to serve the failed node filesystem data, which allows the DVS application to continue to stream data.
- ▶ Online/non disruptive software upgrades.
- ▶ When you halt one node and allow takeover, the partner node continues to serve data for the halted node while you upgrade the node you halted, or in the case of a cluster failover event, the partner node automatically takes over, as illustrated in Figure 3-7.
- ▶ Non disruptive storage system and disk maintenance.
- ▶ Clustered systems eliminate all single points-of-failure.

### 3.7 FlexVol

FlexVol provides flexible volumes and is a ground breaking new technology. The flexible volumes, shown in Figure 3-8, are logical data containers that are managed separately from their underlying physical storage that can be sized, resized, managed, and moved independently and non disruptively. This is key to DVS where storage growth spikes and unplanned capacity requirements often occur. Flexible volumes are file systems that hold user data that is accessible using one or more of the access protocols that Data ONTAP supports, which includes NFS, CIFS, HTTP, FTP, FCP, and iSCSI. Because each flexible volume is a separate file system, you can create one or more Snapshots of the data in a volume so that multiple, space-efficient, point-in-time images of the data can be maintained for such purposes as backup and error recovery.




Figure 3-8 FlexVol

### 3.8 MultiStore

With the MultiStore® feature, N series can be virtualized to appear as three separate N series storage subsystems, each one having their own virtualized addressable network connections and storage but physically utilizing one set of network connections and pool of storage. In the DVS environment, depending on industry, application, policy, or security requirements, you might want to separate your video data this way. Figure 3-9 on page 111 is an example of MultiStore,

where the external cameras view the outside, parking lot, back lot, and so on, and the department store cameras view the cash registers and merchandise. The back office cameras view the employee operations.

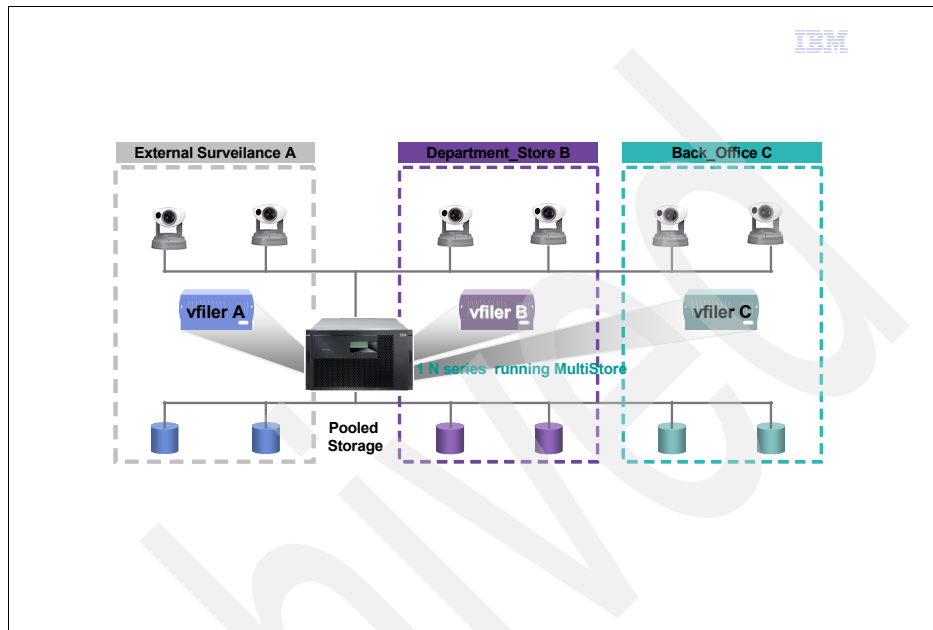


Figure 3-9 MultiStore example

Archived

## Cisco Video Surveillance Media Server software

There are several IBM business partners that provide the video surveillance software for the IBM DVS solution. For this IBM Redbooks publication, it was not possible to use all of them in the time frame given to complete this book; therefore, we used the Cisco Video Surveillance Management System (VSMS) software; however, our selection is not a recommendation of one IBM Business Partner over another. Figure 4-1 on page 114 illustrates the role of Cisco in the IBM DVS solution.

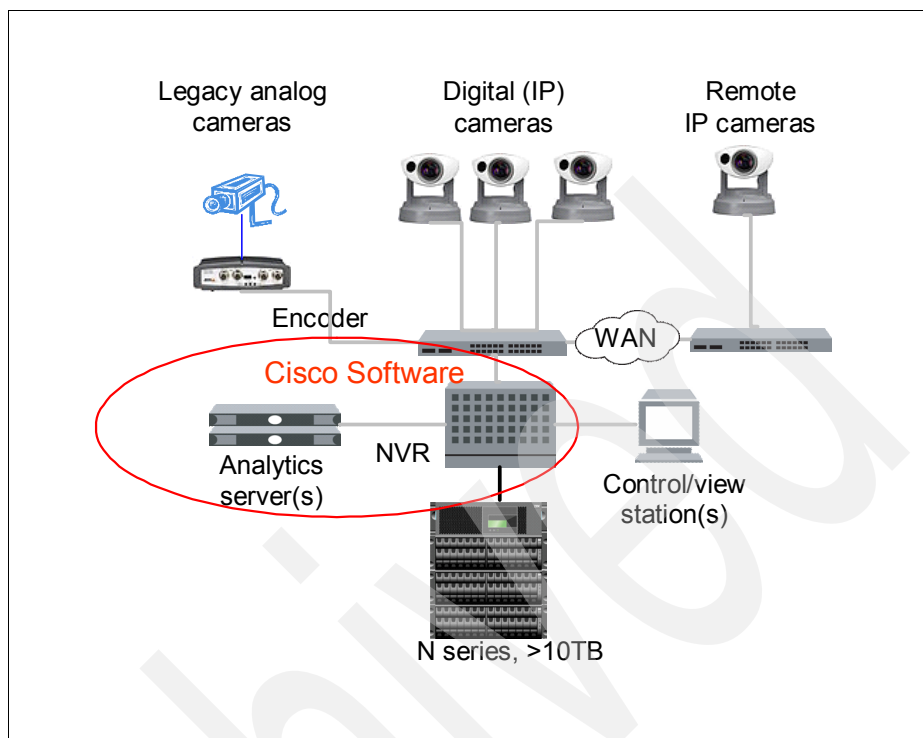


Figure 4-1 Cisco VSMS

## 4.1 Introduction

Cisco offers network-centric video surveillance software and hardware that supports video transmission, monitoring, recording, and managing capabilities. Cisco video surveillance solutions work in unison with the advanced features and functions of the IP network infrastructure—switches, routers, and other network security devices—to enable secure, policy-based access to live or recorded video. With support for many third-party video surveillance cameras, encoders, and applications, Cisco video surveillance solutions allow you to build best-in-class video surveillance systems that optimize cost, performance, and capability. Cisco video surveillance products are deployed within the Cisco Intelligent Converged Environment architecture. Through this architecture, you can access video at any time from any place, which enables real-time incident response, investigation, and resolution. As an extension of the Cisco Self-Defending Network, the Cisco Intelligent Converged Environment

enables you to use existing investments in video surveillance and physical security while enhancing the safety of people and protecting assets. The open, standards-based Cisco infrastructure enables you to deploy and control new security applications and maximizes the value of live and recorded video.

### 4.1.1 Cisco Video Surveillance Media Server software

The Cisco VSMS performs the following networked video surveillance system functions:

- ▶ Collects and routes video from a wide range of third-party cameras and video encoders over an IP network
- ▶ Event-tags and records video for review and archival purposes
- ▶ Secure local, remote, and redundant video archive capabilities
- ▶ Bandwidth management over LANs and WANs

By using the power and advanced capabilities of today's IP networks, Cisco VSMS software allows you to add new third-party applications, additional users, cameras, and storage over time. As a result, the Cisco VSMS software provides unparalleled video surveillance system flexibility and scalability to support:

- ▶ Small systems to those with thousands of cameras (video feeds)
- ▶ Hundreds of simultaneous users viewing live or recorded video
- ▶ Standard video compression algorithms, such as MJPEG, MPEG-2, and MPEG-4 simultaneously in a single Media Platform or system
- ▶ Conservation of storage using events or loop-based archival options
- ▶ Integration with other security and IT applications
- ▶ IT-caliber fault-tolerant storage for greater efficiency and easier maintenance

The Cisco VSMS is complemented by other Cisco Video Surveillance Manager applications that provide video display control and distribution (virtual matrix switching), customizable Web-based user interface for roles-based operation and management, system configuration, and options to support storage area networks (SANs) and network- and direct-attached storage (NAS and DAS). Unlike many other video surveillance offerings that use proprietary hardware, the Media Platform and other Cisco Video Surveillance software run on Linux-based servers. As a result, a wide diverse range of deployment scenarios are supported.

## 4.1.2 Media Platform features

Cisco's VSMS offers video, archiving, and system tool features.

### Powerful video management

The video management feature offers:

- ▶ Standards-based architecture, which provides the flexibility to use a broad range of cameras, codecs, viewing platforms, and network topologies
- ▶ Low-latency video with high-quality images that include megapixel camera video
- ▶ Simultaneous support for MJPEG, MPEG-2, and MPEG-4
- ▶ Unparalleled scalability in terms of number of sites, cameras, viewers, and storage

### Flexible archiving

The archiving feature offers:

- ▶ Archives at different frame rates, durations, and locations
- ▶ Efficient redundant multisite archiving for bandwidth conservation
- ▶ Loop- and event-based video and audio recording and clipping capabilities

### Sophisticated system tools

The system tools offer:

- ▶ Enhanced diagnostic tools that provide notification and API support for failure of proxies and archives
- ▶ Simplified configuration that allows you to import data, in mass, from a spreadsheet
- ▶ Support for redundancy configurations, which includes failover and complex high-availability scenarios
- ▶ Backup utility that supports configurations of devices, archives, events, and pan tilt zoom (PTZ) for a quick and easy restore to a secondary server
- ▶ Minimized load on video servers' platforms by streaming only the active video channels
- ▶ Integration with third-party Electronic Access Control (Lenel)
- ▶ Allows you to sync existing media platforms with Cisco Operations Manager with the click of a button



In this IBM Redbooks publication, we did not use all of the features and functions of the Cisco VSMS software; instead, we focus on the basics of preparation, set up, installation, video capture, and archiving using the N series.

Archived



## Our environment

In this chapter, we describe the environment that we use for this IBM Redbooks publication. We include details about hardware, software, the operating system of the server, SAN configuration, network configuration, and N series configuration.

## 5.1 Hardware

In this section, we provide specifics about the hardware that we use:

- For storage, we use a N series N5200-G20 with integrated FC and Ethernet ports for connections. This is just one of the possible configurations that are available from the N series model product line. Data ONTAP is the same on all N series storage systems and is virtually the same Data ONTAP versions for the N series Gateways, so this book is relevant to all Models of the N series.

Figure 5-1 shows the cluster Gateway model.



*Figure 5-1 N5200 Gateway*

- Behind the N series Gateway is a DS4700 with expansion units, as shown in Figure 5-2 on page 121.



Figure 5-2 DS4700

- ▶ We use two X series servers to host the Digital Video Surveillance (DVS) application:
  - One with Fibre Channel connection using QLogic® QLA2340 HBA cards to connect to the N series
  - One with a iSCSI gigabit Ethernet connection to the N series. We use the software initiator for iSCSI connection that the operating system provides.

## 5.2 Software

The DVS application that we use is Cisco Video Surveillance Media Platform, as shown in Figure 5-3 on page 122.




Figure 5-3 Cisco DVS application

## 5.3 Operating system

The operating system that hosts the DVS application is Suse Linux 9 Enterprise Server. Refer to Chapter 6, “Preparing the N series for Digital Video Surveillance” on page 127 for more information about packages and installation options. You can also refer to:

<http://www.novell.com/linux/>

For the client server, we use Windows XP.

<http://www.microsoft.com/windows/products/windowsxp/default.mspx>

## 5.4 SAN configuration

Each node of the N series Gateway has two paths to the SAN switches and is zoned to the two paths of the DS4700. We configure these two ports as initiators. Each node also has an additional path to the SAN switches to serve as a target to the LINUX host. Figure 5-4 on page 123 illustrates our SAN configuration.

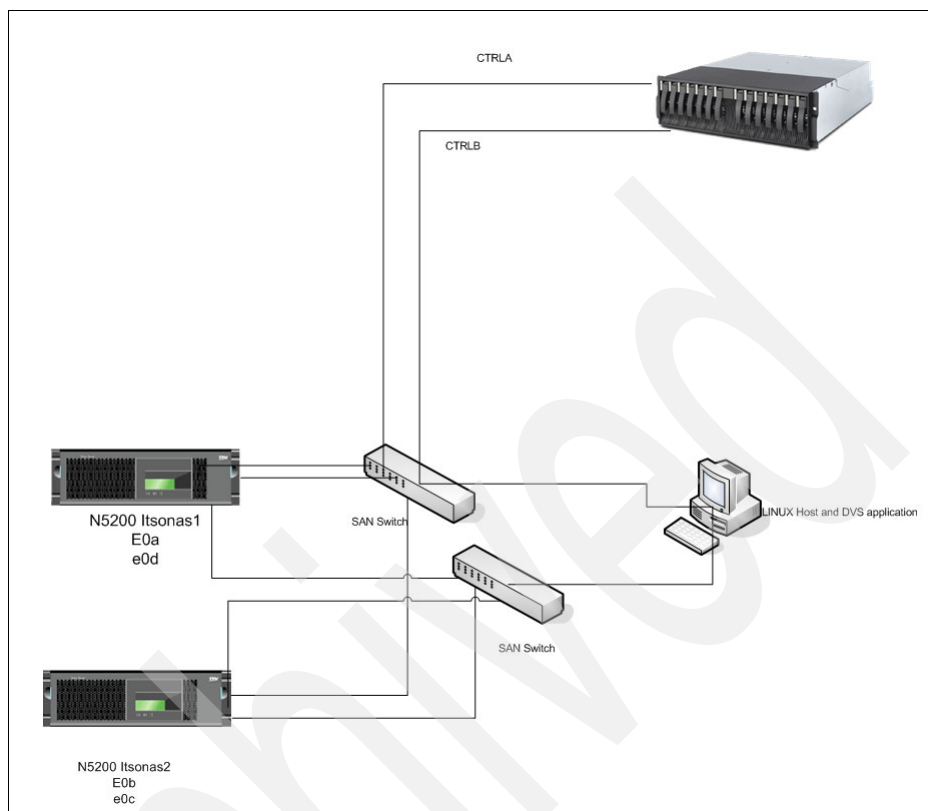


Figure 5-4 SAN configuration

## 5.5 Network

For the iSCSI environment, we use the Ethernet network, with two connections from all of the elements of the configuration, the two N series nodes, and the Linux Host, as shown in Figure 5-5 on page 124.




Figure 5-5 Network configuration



## 5.6 Camera

In our environment, we have two AXIS 207 Network cameras, which we show in Figure 5-6 and Figure 5-7 on page 126, with the following specifications:

- ▶ Up to 30 frames per second for all resolutions
- ▶ Resolution: 640 x 480, 480 x 360, 352 x 288, 320 x 240, 240 x 180, 176 x 144, and 160 x 120
- ▶ Video compression: Motion JPEG and MPEG-4



*Figure 5-6* AXIS 207 camera

Figure 5-7 on page 126 shows our network with the AXIS Network camera.




Figure 5-7 Network with AXIS Network camera

# Preparing the N series for Digital Video Surveillance

In this chapter, we discuss:

- ▶ Preparing the IBM System Storage N series for Digital Video Surveillance Software (DVS)
- ▶ Sizing the N series to DVS, SAN environment and the Fibre Channel Protocol (FCP) topologies, Ethernet environment, and the iSCSI (internet SCSI) topologies
- ▶ Setting up FCP and iSCSI on host and N series storage
- ▶ Implementing cluster and cluster failover modes (CFMODE)
- ▶ Setting up RAID-DP and RAID4
- ▶ Mirroring data with N series SnapMirror
- ▶ SUSE® Linux Enterprise Server 9 (SLES9) requirements and configuration
- ▶ Setting up DVS for the N series

## 6.1 The physical environment

For this book, we created two environments with different topologies, FCP and iSCSI because some of your smaller customers, remote installations, or those customers that do not want to invest in SAN technologies, might be more suited to the iSCSI environment whereas your large customers or large installations might require the connectivity of the SAN environment or want to utilize their SAN investment. We discuss how to prepare both environments in this chapter.

Figure 6-1 shows a DVS environment that uses a FCP topology.




Figure 6-1 DVS environment using FCP topology

In a SAN environment, we use a FCP topology and the media platform is attached to a N series storage system through HBA and fibre switches. We work with some high availability (HA) resources and we discuss those resources through this chapter. Figure 6-2 shows a DVS environment using iSCSI topology.




Figure 6-2 DVS environment using iSCSI topology

In the iSCSI environment, the media platform is attached to N series storage systems through an Ethernet card and Ethernet switches that work as a single network configuration.

## 6.2 Sizing N series storage to DVS

Determining the exact amount of storage you will need for your DVS installation depends on a number of factors. The following list offers some general guidelines. For help in determining more specific networked DVS storage needs for your organization, contact your IBM System Storage N series representative.

The amount of storage needed can be a factor of:

- ▶ The number of surveillance cameras in use
- ▶ How many hours per day each camera records footage

Some cameras record 24 x 7, while others might record only when motion is automatically detected. Still others might be scheduled to record only during business hours.

- ▶ The type of camera resolution being used

Most IP/digital cameras transmit at either CIF, 2CIF, or 4CIF resolution. 4CIF is 640 x 480 or 704 x 480 pixels, 0.3 Megapixels, or akin to VGA monitor or DVD movie quality. Some IP cameras go up to eight Megapixels today, which allows coverage of a large area or high quality digital zoom. The camera resolution used can relate to the type of surveillance that is likely to be performed (that is, facial recognition or license-plate tracking might require cameras with higher pixel counts).

Figure 6-3 provides a resolution sample.




Figure 6-3 Resolution sample

Table 6-1 guides you through some of the common resolution settings (in pixels) and how they correspond to different market standards. The earlier standards are from the analog world and are based on TV-line resolutions.

*Table 6-1 Common resolution settings*

Resolution	Pixels	Typical market	Comments
CIF	352 x 240	US, Japan	NTSC standard
CIF	352 x 288	Europe	PAL standard
4CIF	704 x 480	US, Japan	NTSC standard
4CIF	704 x 576	Europe	PAL standard
QVGA	320 x 240	Global	Digital standard
VGA	640 x 480	Global	Digital standard
XVGA	1024 x 768	Global	Digital standard
MegaPixel	1280 x 960	Global	Digital standard

- The frame rate (frames per second, or fps)
 

In terms of frame rate, 30 fps for NTSC or 25 fps for PAL is television-quality, which is an excellent frame rate for viewing live video and usually more than enough for performing post-event review and analytics. In fast-moving environments with large crowds of people (such as casinos, some retail settings, etc.) higher frames per second is advised.
- The use of real-time monitoring
 

To minimize network bandwidth or storage consumption, some cameras might also be set up to transmit at 1-5 fps unless they detect some abnormality, such as unexpected motion or an alarm being triggered. After the abnormality is detected, they might automatically switch to recording at 30 fps until the abnormality is no longer detected. Overall, use of this capability can significantly cut down on storage consumption.
- Type of video compression:<sup>1</sup>

Two of the most common types of video compressions are MPEG-4 and Motion JPEG. MPEG-4 and Motion JPEG each employ a different technique to reduce the amount of data that is transferred and stored in a network video system. Each format has its advantages and disadvantages. MPEG-4 transmits only parts of an image that differ from an earlier referenced image. At high frame rates and particularly with scenes that have static areas, MPEG-4 requires less bandwidth and storage than with Motion JPEG. MPEG-4 is a licensed technology, so if a network camera supports MPEG-4,

<sup>1</sup> From AXIS IP-Surveillance Design Guide

be sure to obtain if the MPEG-4 license fee is already included in the product's purchase price. MPEG-4 provides support for synchronized audio whereas Motion JPEG does not.

Figure 6-4 compares MPEG-4 and JPEG.

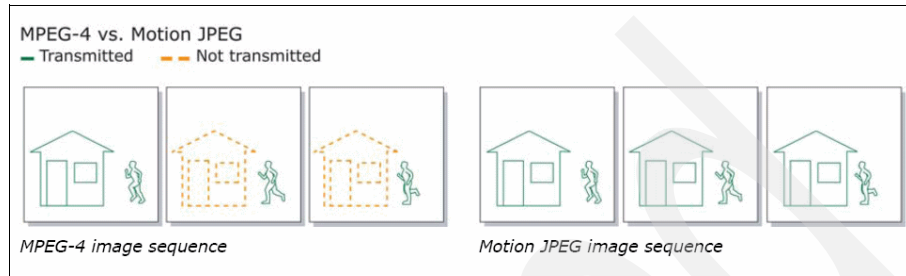


Figure 6-4 MPEG-4 versus JPEG

With Motion JPEG, each image is a complete JPEG compressed image and is simple to encode and decode. One of the advantages of Motion JPEG is that it guarantees the image quality that you set (either high or low) regardless of movement or image complexity. Motion JPEG has low latency and is license free. However, Motion JPEG files are usually larger than those compressed with the MPEG-4 standard. When looking at video compression, it is important to select the compression that best suits your application. One of the best ways to maximize the benefits of both standards is to look for network video products that can deliver simultaneous MPEG-4 and Motion JPEG streams. AXIS network video products, for example, provide the two video compression formats, which gives users the flexibility to both maximize image quality for recording and reduce bandwidth needs for live viewing. With limited bandwidth, you might want to view at full frame rate or 30/25 (NTSC/PAL) frames per second (fps) with MPEG-4 and record with guaranteed quality using Motion JPEG.

#### ► The retention time

How long do you plan to store the recorded footage? Retention times can be anywhere from 30 days to three years or more, depending on how the organization uses the footage. Retention times are not an exact science, although there are norms emerging in vertical use case.

## 6.2.1 Calculating the storage needs<sup>1</sup>

To appropriately calculate the storage requirements of a network surveillance system, there are a number of elements to factor in, such as the number of cameras that are required in your installation, the number of hours a day each camera records, how long the data is stored, and whether the system uses



motion detection or continuous recording. Additional parameters to consider include frame rate, compression, image quality, and complexity.

The type of video compression that is employed also effects storage calculations. Systems employing Motion JPEG compression vary storage requirements by changing the frame rate, resolution, and compression. If MPEG compression is used, bit rate is the key factor that determines the corresponding storage requirements.

Fortunately, there are formulas for calculating the amount of storage to buy. These formulas are different for Motion JPEG and MPEG-4 compressions because Motion JPEG consists of one individual file for each image, while MPEG-4 is a stream of data that is measured in bits per second. We discuss these formulas in the next sections.

### Motion JPEG calculation<sup>2</sup>

Image size x frames per second x 3600s = Kilobyte (KB) per hour/1000 = Megabyte (MB) per hour.

MB per hour x hours of operation per day / 1000 = Gigabyte (GB) per day.

GB per day x requested period of storage = Storage need.

Table 6-2 illustrates the Motion JPEG calculations for three cameras.

Table 6-2 Total for three cameras and 30 days of storage = 1002 GB

Camera	Resolution	Image size (KB)	Frames per second	MB/hour	Hours of operation	GB/day
No.1	CIF	13	5	234	8	1.9
No.2	CIF	13	15	702	8	5.6
No.3	4CIF	40	15	2160	12	26

### MPEG-4 calculation

Bit rate / 8(bits in a byte) x 3600s = KB per hour / 1000 = MB per hour.

MB per hour x hours of operation per day / 1000 = GB per day.

GB per day x requested period of storage = Storage need.

**Note:** The formula does not take into account the complexity of the image, which is an important factor that can influence the size of the storage required.

<sup>2</sup> Calculations provided by Cisco

Table 6-3 shows the MPEG-4 calculation for three cameras.

*Table 6-3 Total for three cameras and 30 days of storage = 204 GB*

Camera	Resolution	Bit Rate (Kbit/s)	Frames per second	MB/hour	Hours of operation	GB/day
No.1	CIF	170	5	76.5	8	0.6
No.2	CIF	400	15	180	8	1.4
No.3	4CIF	880	15	396	12	5

## 6.2.2 Sample DVS storage calculations<sup>2</sup>

To calculate DVS storage:

1. Calculate the bit rate per second per camera.

This is the amount of network bandwidth that is required per camera. (In general, cameras transmitting 30 fps at 2CIF or 4CIF resolution tend to require from 2 Mbits/second to 6 Mbits/second of bandwidth.) You can calculate the bit rate as follows:

Camera Resolution x Frame Rate / Compression = Bit Rate per Camera  
(typically in Mbits/second)

2. Determine the amount of usable storage that you need, as shown in Table 6-4.

Bit Rate per Camera x Number of Cameras x Retention Time = Usable Storage

**Note:** DVS industry commentator Fredrik Nilsson offers other storage calculations for use with Motion JPEG or MPEG compression.

*Table 6-4 Determining the amount of usable storage*

Example	Calculation	Usable storage required
Example #1: ▶ One camera ▶ Continuous operation 24x7 ▶ 90-day retention time ▶ Bandwidth needed: 6 Mbits/second	6 Mbits per sec. x 3600 sec./hour x 24 hrs. /day x 90 days/8 bits/byte	5.8TB for 90 days (storage needed for one camera only)

Example	Calculation	Usable storage required
Example #2: ▶ 40 cameras ▶ Continuous operation 24 x 7 ▶ 30-day retention time ▶ 4CIF resolution at 30 fps with bandwidth needed of 4 Mbits/second	Step 1: 4 Mbits per sec. x 40 cameras = 160 Mbits per second (or 20 Mbytes per second) Step 2: 20 MB per sec. x 3600 sec. /hour x 24 hrs. /day x 30 days	5.1TB for 30 days (storage needed for 40 cameras)

### 6.2.3 Factors in determining capacity hardware and license needs

This section contains the most common factors in determining DVS requirements.

#### Individual remote offices or centralized data capture

How your surveillance network is set up, bandwidth, and remote office connectivity determines how your DVS solution is configured.

#### Centralized data capture

If you are implementing a smaller surveillance system that involves eight to ten cameras, you should be able to use a basic 100 megabit (Mbit) network switch without having to consider bandwidth limitations. Most companies can implement a surveillance system of this size using their existing network.

If you implement 10 cameras or more, estimate the load in the network using a few rules of thumb:

- ▶ A camera uses approximately 2 to 6 megabits of bandwidth when configured to deliver high-quality images at a high frame rate.
- ▶ With more than 12 to 15 cameras, consider using a switch with a gigabit (Gbit) back bone. If a gigabit-supporting switch is used, the server that runs the video management software should have a gigabit network adapter installed.

#### Remote offices

If you are implementing a remote office surveillance system, you need to have the same considerations of the centralized data capture that we previously mentioned, and if you want to do data replication, mirroring data, and disaster recovery, you need to have more than one N series and associated licenses. For help in determining more remote offices surveillance system needs for your organization, contact your IBM System Storage N series representative.

## 6.3 N series calculations

After you calculate the capacity that is needed for DVS, look at the N series models and determine which model would best suit your needs. After you determine the best model for your environment, go through the following factors in determining your N series storage configuration:

- ▶ Clustering: This requires the A20 or G20 models of the N series and, in effect, gives you two nodes for disaster recovery.
- ▶ Location of Root volume: Will you isolate the root volume or share it with DVS. If standalone, you have to subtract its volume capacity from the total usable capacity for DVS.
- ▶ Snapshot %: Normally we do not recommend using more than 20% of the raw capacity for Snapshot. In DVS, you most likely will not use Snapshots of the video streams, but use Snapshots of the video clips instead. If you are using snapshots of the LUNs, you will need a minimum of 100% space for the first Snapshot.
- ▶ Number of Spares: The more capacity, the more spares, as shown in Figure 6-5.
- ▶ Mirroring: Effectively doubles the amount of capacity needed.




Figure 6-5 Hot spares

## 6.4 Hardware considerations

Depending on your business requirements, different factors come into play in determining the best solution for you. We discuss those factors in this section.

### 6.4.1 Factors to consider

In this section, we discuss the factors for you to consider as you determine the best solution for you.

#### Archiving and retention archive

N series supports both Fibre Channel (FC) and Serial ATA disks (SATA) concurrently, which allows you to implement different policies for treatment of video data, for example:

- ▶ You might use Fibre Channel disk for video streams captured in the last seven days.
- ▶ For older video captures or clips, you might want to archive to SATA drives.
- ▶ For a large amount of cameras, streaming simultaneously higher performing FC drives is better.

#### Drive recommendations

We recommend the FC drives to the stream capture and SATA drives to the clips and archive system on this environment.

### 6.4.2 Mirroring data

Depending on the regulations or laws that your business has to comply with, the absolute protection of data might require mirroring and that you use the N series SnapMirror feature.

#### Overview of aggregates

Before we discuss mirroring and SnapMirror, you must understand the basic building blocks of N series space organization, starting with an aggregate. An *aggregate* is a collection of physical disk space that is used as a container, depending on whether you want to take advantage of RAID-level mirroring and physical layer. If the aggregate is unmirrored, it only contains a plex. 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. If the SyncMirror feature is licensed and enabled, Data ONTAP adds a second plex to

the aggregate, which serves as a RAID-level mirror for the first plex in the aggregate.

When you create an aggregate, Data ONTAP assigns data disks and parity disks to RAID groups depending on the options that you choose, such as the size of the RAID group or the level of RAID protection.

Each aggregate possesses its own RAID configuration and set of assigned disks. Within each aggregate, you can create one or more FlexVol volumes. You can increase the usable space in an aggregate by adding disks to existing RAID group, or by adding new RAID groups.

Table 6-5 lists the limits of aggregates and volumes.

*Table 6-5 Limits of aggregates and volumes*

Items	Number or size
Maximum number of aggregates	100
Maximum aggregate size	16 TB
Minimum aggregate size	10 Gb
Maximum number of RAID groups in an aggregate	150

### **Mirrored aggregate**

A mirrored aggregate consists of two plexes, which provide an even higher level of data redundancy through RAID-level mirroring. For an aggregate to be enabled for mirroring, the storage system must have a SyncMirror license for `syncmirror_local` or `cluster_remote` installed, and the storage system's disk configuration must support RAID-level mirroring.

When you enable SyncMirror, Data ONTAP divides all of the hot spot spare disks into two disk pools to ensure that a single failure does not affect disks in both pools. This division allows the creation of mirrored aggregates. Data ONTAP uses disks from one pool to create the first plex and another pool to create the second plex. A failure that affects one plex will not affect the other plex.

The plexes are physically separated and are updated simultaneously during normal operation. After the plex that had a problem is fixed, you can resynchronize the two plexes and reestablish the mirror relationship.

Figure 6-6 on page 139 shows that SyncMirror is enabled and that plex0 and plex1 contain copies of one or more file systems. There are also hot spare disks in disk shelves and a pool for each sub-container waiting to be assigned.

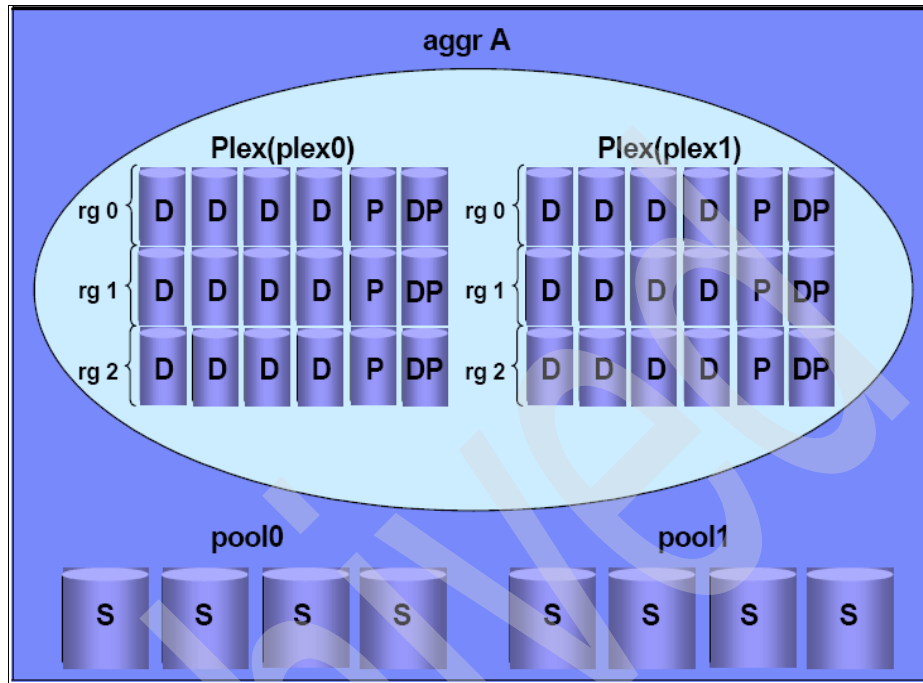


Figure 6-6 Mirrored aggregate

### 6.4.3 RAID groups

A RAID group consists of one or more data disks, across which client data is striped and stored, plus one or two parity disks. The purpose of a RAID group is to provide parity protection from data loss across its included disks. RAID4 uses one parity disk to ensure data recoverability, if one disk fails within the RAID group. RAID-DP uses two parity disks to ensure data recoverability, even if two disks within the RAID group fail.

#### Levels of RAID protection

Data ONTAP supports three levels of RAID protection: RAID4 and RAID-DP for the N series storage systems, which you can assign on a per-aggregate basis, and RAID 0 for the N series Gateways:

- ▶ If an aggregate is configured for RAID4 protection, Data ONTAP reconstructs the data from a single failed disk within a RAID group and transfers that reconstructed data to a spare disk.
- ▶ If an aggregate is configured for RAID-DP protection, Data ONTAP reconstructs the data from one or two failed disks within a RAID group and transfers that reconstructed data to one or two spare disks as necessary.

Figure 6-7 shows the RAID levels.




Figure 6-7 RAID levels

### ***RAID4 protection***

RAID4, as shown in Figure 6-8 on page 141, provides single-parity disk protection against single-disk failure within a RAID group. The minimum number of disks in a RAID4 group is two: at least one data disk and one parity disk. If there is a single data or parity disk failure in a RAID4 group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the failed disk's data on the replacement disk. If no spare disks are available, Data ONTAP goes into degraded mode and alerts you of this condition.






Figure 6-8 RAID4

**Attention:** With RAID4, if there is a second disk failure before data can be reconstructed from the data on the first failed disk, there is no data loss. To avoid data loss when two disks fail, you can select RAID-DP, which provides two parity disks to protect you from data loss when two disk failures occur in the same RAID group before the first failed disk is reconstructed.

### ***RAID-DP protection***

RAID-DP provides double-parity disk protection when the following conditions occur:

- ▶ There are media errors on a block when Data ONTAP is attempting to reconstruct a failed disk.
- ▶ There is a single- or double-disk failure within a RAID group. The minimum number of disks in a RAID-DP group is three: at least one data disk, one regular parity disk, and one double-parity (or dParity) disk.

If there is a data-disk or parity-disk failure in a RAID-DP group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the data of the failed disk on the replacement disk. If there is a double-disk failure, Data ONTAP replaces the failed disks in the RAID group with two spare disks and uses the double-parity data to reconstruct the data of the failed disks on the replacement disks. Figure 6-9 on page 142 shows a traditional volume that is configured for RAID-DP protection.




Figure 6-9 RAID-DP

**Tip:** The recommended RAID group setting is to use RAID-DP on all N series.

## 6.5 SnapMirror

SnapMirror is a feature of IBM N series that allows a dataset to be replicated between IBM N series storage systems over a network for backup or disaster recovery purposes. After an initial baseline transfer of the entire dataset, as shown in Figure 6-10 on page 143, subsequent updates only transfer new and changed data blocks from the source to the destination, which makes SnapMirror highly efficient in terms of network bandwidth utilization.

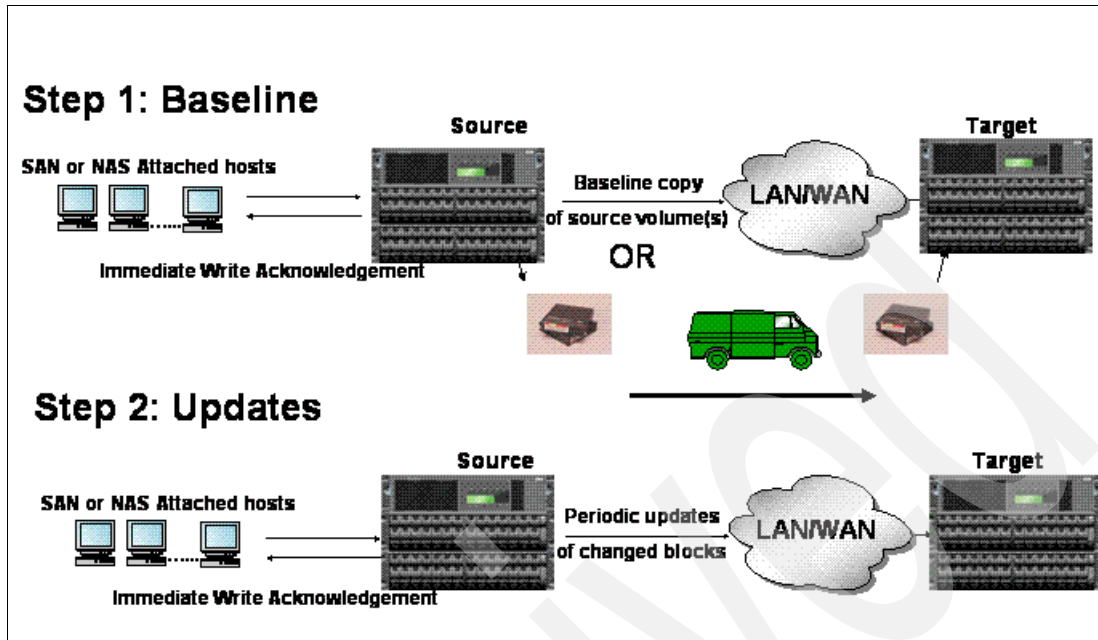


Figure 6-10 Baseline creation

### 6.5.1 Criticality to business

On DVS environments, we recommend this feature because if your video surveillance information is critical to your organization, you can use the SnapMirror to mirror or replicate your archive videos in another location.

### 6.5.2 System performance on the secondary storage device

Because the secondary storage system must write all of the same data that the primary storage system writes, it is important that the secondary system can maintain the write throughput that is expected on the primary system, for example, if a low-end filer, such as the N3700, is configured as a secondary for a high-end filer, such as a N7600, the write performance on the N7600 is limited to what can be achieved on the N3700. For this reason, the best practice is to always configure the same model of filer for both primary and secondary systems or a higher-end filer on the secondary system.

### 6.5.3 Network bandwidth considerations

Because all of the data that is written to the primary storage must be replicated to the secondary storage as it is written, write throughput to the primary storage cannot generally exceed the bandwidth that is available between the primary and secondary storage devices. Because SnapMirror transfers can be performed over standard Ethernet networks and over Fibre Channel networks, there is a choice for transport. This choice is, most likely, determined by preference or by existing infrastructure rather than by performance needs.

In general, the configuration guideline is to configure the network between the primary and secondary storage with at least as much bandwidth as the network between the clients and the primary storage.

#### Required licenses

To work with SnapMirror you must have the following licenses:

- ▶ SnapMirror
- ▶ SnapMirror\_sync

#### SnapMirror modes

You can use SnapMirror in three different modes, which we discuss in the next sections:

- ▶ Asynchronous
- ▶ Synchronous
- ▶ Semi-synchronous

### 6.5.4 Asynchronous mode

In asynchronous mode, as shown in Figure 6-11 on page 145, SnapMirror performs incremental, block-based replication as frequently as once per minute. Consult your technical team for the best plan for your environment or to determine whether synchronous SnapMirror is a better match for you. The performance impact on the source IBM System Storage N is minimal as long as the system is configured with sufficient CPU and disk I/O resources.

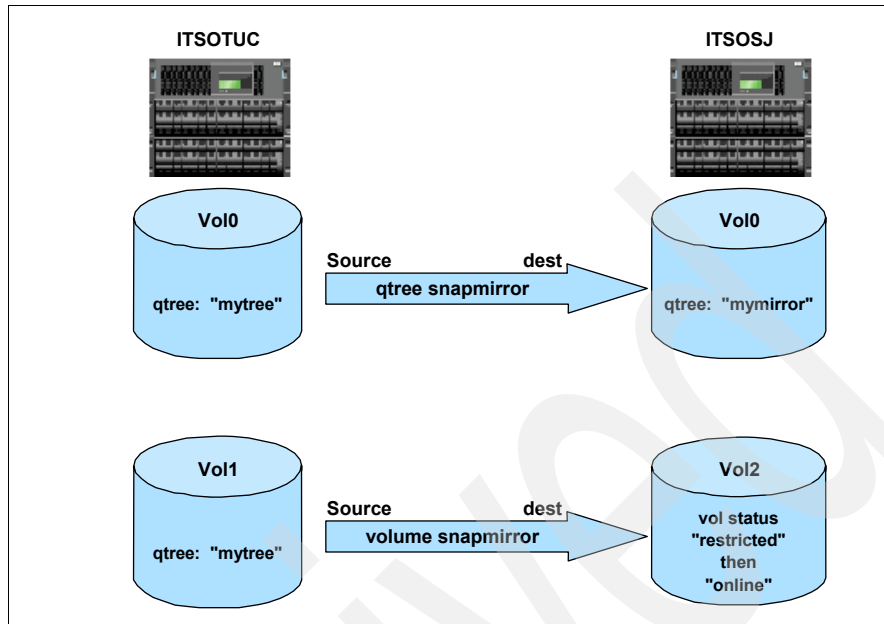


Figure 6-11 Asynchronous SnapMirror options

## Asynchronous mode initialization

The first and most important step in asynchronous mode involves creating a one-time, baseline transfer of the entire dataset, which is required before incremental updates can be performed. The process for asynchronous mode initialization is:

1. The primary storage system takes a Snapshot copy (a read-only, point-in-time image of the file system).

Referring to Figure 6-10 on page 143, this Snapshot copy is called the "baseline" copy.

2. All data blocks referenced by this Snapshot copy and any previous Snapshot copies are transferred and written to the secondary file system.
3. After initialization is complete, the primary and secondary file systems have at least one Snapshot copy in common.

## Asynchronous mode updates

After initialization, both scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the primary to the secondary file system.

The process for asynchronous mode updates is:

1. The primary storage system takes a Snapshot copy.
2. The new Snapshot copy is compared to the baseline Snapshot copy to determine which blocks changed.
3. The changed blocks are sent to the secondary and written to the file system.
4. After the update is complete, both file systems have the new Snapshot copy, which becomes the baseline Snapshot copy for the next update.

Because asynchronous replication is periodic, SnapMirror can consolidate writes and conserve network bandwidth.

### 6.5.5 Synchronous mode

Synchronous SnapMirror 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. This ensures that data that is written on the source system is protected on the destination even if the entire source system fails. It guarantees zero data loss in the event of a failure, but can have a significant impact on performance. It is not necessary or appropriate for all applications.

Synchronous SnapMirror, shown in Figure 6-12 on page 147, replicates data between single storage systems or clustered storage systems that are located at remote sites using IP or FCP infrastructure with no special converters required. Synchronous SnapMirror is simply a mode of operation or a feature that was recently added to the SnapMirror software.

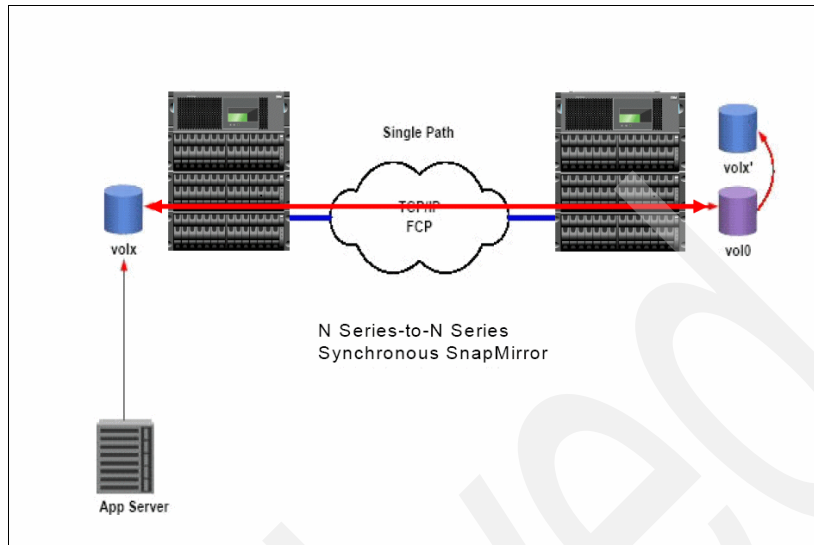


Figure 6-12 Single path SnapMirror

**Note:** Synchronous SnapMirror is only supported in a configuration using FlexVols and FCP through RPQ.

## Terminology

To avoid any potential confusion, in this section we review exactly what is meant by the term *synchronous* in this context. The best way to do this is to examine a scenario where the primary data storage device fails completely and then examine the disaster's impact on an application.

In a typical application environment, the following steps occur:

1. A user saves some information in the application.
2. The client software communicates with a server and transmits the information.
3. The server software processes the information and transmits it to the operating system on the server.
4. The operating system software sends the information to the storage.
5. The storage acknowledges receipt of the data.
6. The operating system tells the application server that the write is complete.
7. The application server tells the client that the write is complete.
8. The client software tells the user that the write is complete.

In most cases, these steps take only tiny fractions of a second to complete. If the storage system fails in such a way that all data on it is lost (for example, as a result of a fire or flood that destroys all of the storage media), the impact to an individual transaction varies based on when the failure occurs, as explained here:

- ▶ If the failure occurs before step 5, the storage never acknowledges receipt of the data, which results in the user receiving an error message from the application saying that it failed to save the transaction.
- ▶ If the failure occurs after step 5, the user sees client behavior that indicates correct operation (at least until the following transaction is attempted). Despite the indication by the client software (in step 8) that the write was successful, the data is lost.

The first case is obviously preferable to the second because it provides the user or application with knowledge of the failure and the opportunity to preserve the data until the transaction can be attempted again. In the second case, the data can be discarded based on the belief that it is already safely stored.

With traditional asynchronous SnapMirror, data is replicated from the primary storage to a secondary or destination storage device on a schedule. If this schedule were configured to cause updates once per hour, for example, it is possible for a full hour of transactions to be written to the primary storage and acknowledged by the application, only to be lost when a failure occurs before the next update. For this reason, many customers attempt to minimize the time between transfers. Some customers replicate as frequently as once per minute, which significantly reduces the amount of data that could be lost in a disaster.

This level of flexibility is good enough for the vast majority of applications and users. In most real-world environments, loss of one minute or five minutes of data is of trivial concern compared to the downtime that is incurred during such an event. Any disaster that completely destroys the data on the IBM System Storage N series would most likely also destroy the relevant application servers, critical network infrastructure, and so on.

However, there are some customers and applications that have a zero data loss requirement even in the event of a complete failure at the primary site, as illustrated in Figure 6-13 on page 149.



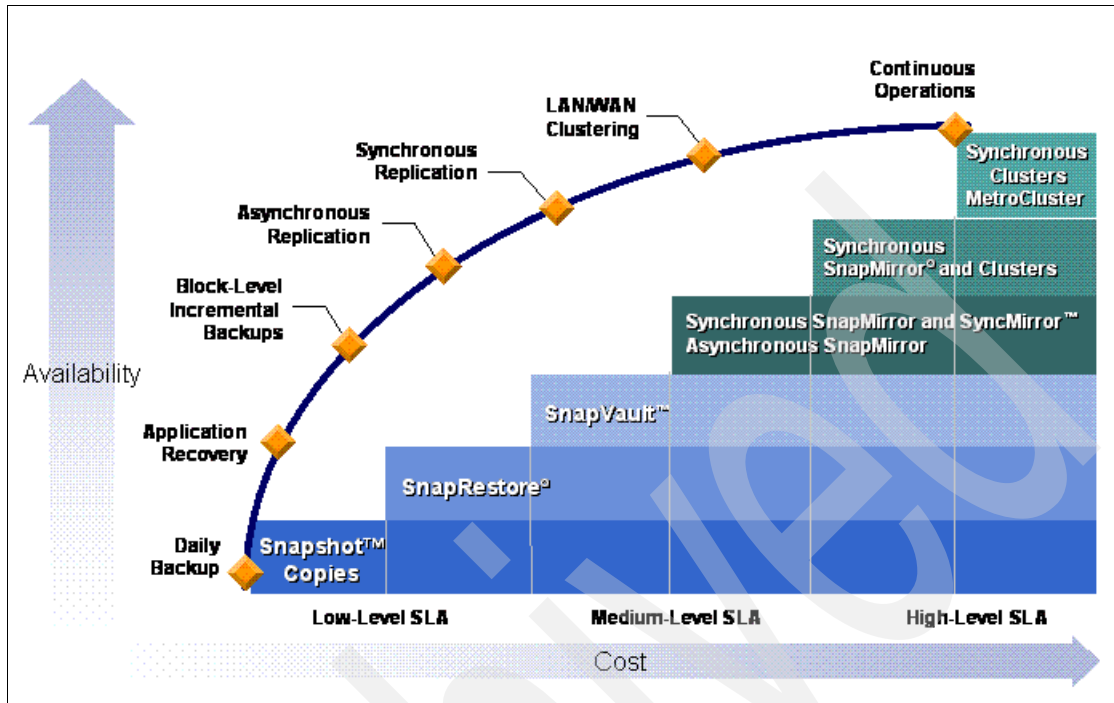


Figure 6-13 Availability

For these situations, synchronous mode is appropriate because it modifies the application environment such that data replication to the secondary storage occurs with *each* transaction, as in the following process example:

1. A user saves some information in the application.
2. The client software communicates with a server and transmits the information.
3. The server software processes the information and transmits it to the operating system on the server.
4. The operating system software sends the information to the primary storage.
5. The primary storage sends the information to the secondary storage.
6. The secondary storage acknowledges receipt of the data.
7. The primary storage acknowledges receipt of the data.
8. The operating system tells the application server that the write is complete.
9. The application server tells the client that the write is complete.
10. The client software tells the user that the write is complete.

The key difference, from the application's point-of-view, is that the storage does not acknowledge the write until the data is written to both the primary and the secondary storage. This has some performance impact, which we discussed later, but modifies the failure scenario in the following beneficial ways:

- ▶ If the failure occurs *before* step 7, the storage never acknowledges receipt of the data, which results in the user receiving an error message from the application that says it failed to save the transaction. This causes inconvenience, but no data loss.
- ▶ If the failure occurs during or after step 7, the data is safely preserved on the secondary storage system despite the failure of the primary.

**Note:** Regardless of what technology is used, it is always possible to lose data. The key point is that with synchronous mode, loss of data that was acknowledged is prevented.

## Operation

The first step involved in synchronous replication is a one-time, baseline transfer of the entire dataset. The following steps delineate the process:

**Note:** SnapMirror must be licensed before synchronous SnapMirror.

After the baseline transfer is complete, SnapMirror can change to synchronous mode, as follows:

1. Asynchronous updates occur until the primary and secondary file systems are very close to being synchronized.
2. NVLOG forwarding begins, which is a method for transferring updates as they occur.
3. Consistency point (CP) synchronization begins, which is a method for ensuring that writes of data from memory to disk storage are synchronized on the primary and secondary systems.
4. New writes from clients or hosts on the primary file system are blocked until acknowledgment of those writes are received from the secondary system.
5. A final update occurs using the same method as asynchronous updates.

After SnapMirror determines that all data acknowledged by primary is safely stored on the secondary, the system is in synchronous mode. At this point, the output of a SnapMirror status query shows that the relationship is "In Sync."

**Note:** If the environment cannot maintain synchronous mode (because of networking or destination issues), SnapMirror drops to asynchronous mode. When the connection is reestablished, the source filer asynchronously replicates data to the destination once each minute until synchronous replication is reestablished. After this occurs, a message is logged of the change of status (“into” or “out of” synchronous status). This “safety net” is known as *fail-safe synchronous*.

## Synchronous mode paths

More than one physical path might be required for a synchronous mirror. Synchronous SnapMirror supports up to two paths for a particular relationship. These paths can be Ethernet, Fibre Channel, or a combination of the two.

Multipath support allows synchronous and semi-synchronous traffic to be load-balanced between these paths and provides for failover in the event of a network outage. There are two modes of multipath operation:

- ▶ Multiplexing mode, as illustrated in Figure 6-14, is when both paths are used simultaneously and load-balancing transfers across the two. When a failure occurs, the load from both transfers moves to the remaining path.
- ▶ Failover mode, is when one path is specified as the primary path in the configuration file. This path is the desired path and is used until a failure occurs. The second path is then used.




Figure 6-14 SnapMirror Multipath

## 6.5.6 Semi-synchronous mode

SnapMirror also provides a semi-synchronous mode, sometimes called *semi-sync*. Synchronous SnapMirror can be configured to lag behind the source volume by a user-defined number of write operations or milliseconds.

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—Of course, for this very reason it is possible to lose acknowledged data.

This mode is also like synchronous mode in that updates from the primary storage to the secondary storage occur right away, rather than waiting for scheduled transfers, which makes the potential amount of data that is lost in a disaster very small. Semi-synchronous mode minimizes data loss in a disaster, while also minimizing the extent to which replication impacts the performance of the source system.

Semi-synchronous mode provides a middle ground that keeps the primary and secondary file systems more closely synchronized than asynchronous mode. The configuration of semi-synchronous mode is identical to the configuration of synchronous mode with the addition of an option that specifies how many writes can be outstanding (unacknowledged by the secondary system) before the primary system delays acknowledging writes from the clients.

Internally, semi-synchronous mode works identically to synchronous mode, in most cases. The only difference lies in how quickly client writes are acknowledged; however, the replication methods that are used are the same.

However, you can configure semi-synchronous mode in a way that changes the replication strategy. A CP is triggered when NVRAM is one-half full, or every 10 seconds, whichever occurs sooner.

If you configure semi-synchronous mode to allow unacknowledged transactions that are greater than 10 seconds old, SnapMirror falls back to performing CP synchronization only. NVLOG forwarding is halted because a CP synchronization is sufficiently frequent to meet the service level requested.

When a CP synchronization occurs under such circumstances, the tetris that is sent to the secondary filer includes not just the list of data blocks to be written but also the content of those data blocks because with NVLOG forwarding disabled, the secondary system does not have a copy of the data until the CP synchronization occurs.

For the vast majority of configurations, NVLOG forwarding is desirable. Therefore, we do not recommend that you configure SnapMirror to allow more than 10 seconds of outstanding data if you want higher synchronicity levels.

However, if NVLOG forwarding is not required, specifying a large time value for outstanding data might reduce the overall CPU usage on the primary storage system, which can allow for significant increases in overall throughput, if CPU usage is a limiting factor.

**Note:** Unlike asynchronous mode that can replicate either volumes or quota trees, synchronous and semi-synchronous modes work only with volumes.

### **Semi-synchronous mode scenario**

The scenario for using semi-sync mode is:

1. A user saves some information in the application.
2. The client software communicates with a server and transmits the information.
3. The server software processes the information and transmits it to the operating system on the server.
4. The operating system software sends the information to the primary storage.
5. The primary storage sends the information to the secondary storage. The primary storage simultaneously acknowledges receipt of the data.
6. The operating system tells the application server that the write is complete.
7. The application server tells the client that the write is complete.
8. The client software tells the user that the write is complete.
9. At some point after step 5, the secondary acknowledges receipt of the data. (Note that step 9 could potentially occur before or simultaneously with step 6.)

If the secondary storage system is slow or unavailable, it is possible that the primary storage system acknowledged a large number of transactions that are not yet protected on the secondary. These transactions represent a window of vulnerability to loss of acknowledged data.

For a window of zero size, you can use fully synchronous mode rather than semi-sync. If you use semi-sync mode, you can customize the size of this window based on your needs and the applications' needs. It can be specified as a number of operations, milliseconds, or seconds.

If the number of outstanding operations equals or exceeds the number of operations that you specify, the primary storage system does not acknowledge further write operations until the secondary acknowledges some.

Likewise, if the secondary has not acknowledged the oldest outstanding transaction within the amount of time that the user specified, the primary storage system does not acknowledge further write operations until all responses from the secondary are being received within that time frame.

## **6.6 Why DVS needs SnapLock**

The SnapLock feature is important on environments that need to do retention of archives and clips. There are two types of Snaplock products, SnapLock Compliance and SnapLock Enterprise, which we discuss in this chapter. SnapLock works with WORM (Write once, ready many).

### **6.6.1 License requirement**

The SnapLock feature is an optional license feature. When you use SnapLock, you must have more space available on your N series storage system because of the retention period requirements.

### **6.6.2 Network effects**

Using SnapLock does not increase your bandwidth consumption.

### **6.6.3 SnapLock Compliance and SnapLock Enterprise**

Both SnapLock Compliance and SnapLock Enterprise software products provide nonerasable, nonrewritable WORM (write once, read many) data permanence functionality that utilizes high-throughput magnetic disk drives in a cost-efficient, highly available RAID configuration. From a data protection perspective, the process of committing data to WORM status on either SnapLock product can be thought of in the same manner as storing data on an optical platter.

Like an optical platter that is “burned” with data, both SnapLock software products protect data that is committed to WORM status from any possible alteration or deletion until their retention period expires. Although SnapLock Compliance and SnapLock Enterprise data permanence is analogous to traditional optical WORM media, the similarities end there. SnapLock offers performance and reliability improvements over traditional WORM storage, while reducing both maintenance overhead and TCO.

SnapLock Compliance and SnapLock Enterprise are implemented as add-on licenses to Data ONTAP. Both SnapLock software products run on the N series

storage system and Gateway, which features lower cost, higher capacity ATA-based drives, the EXN1000, and on higher performance EXN2000 or EXN4000 expansion units that feature fiber attached disk drives.

This flexibility allows you to buy the amount of storage that fits your business needs for SnapLock WORM storage, whether it is a few hundred gigabytes of data or hundreds of terabytes. Additionally, to highlight this flexibility further, SnapLock Compliance or SnapLock Enterprise can be combined on the same storage system with traditional read-write volumes. In the next sections, we explain each product in more detail.

### **SnapLock Compliance**

SnapLock Compliance assists organizations in implementing a comprehensive archival solution for meeting Security Exchange Commission (SEC) or governmental regulations for data retention. Records and files that are committed to WORM storage on a SnapLock Compliance volume cannot be altered or deleted before the expiration of their retention period. Moreover, a SnapLock Compliance volume cannot be deleted until all data that is stored on it passes its retention period and is deleted by the archival application or some other process.

### **SnapLock Enterprise**

SnapLock Enterprise is geared towards assisting organizations with meeting self-regulated and best practice guidelines for protecting digital assets with WORM-type data storage. Data stored as WORM on a SnapLock Enterprise volume is protected from alteration or modification with one main difference from SnapLock Compliance: as the data being stored is not for regulatory compliance, an administrator can delete a SnapLock Enterprise volume, including the data it contains.

Figure 6-15 on page 156 shows a comparison of SnapLock Compliance and SnapLock Enterprise.

SnapLock™ Compliance	SnapLock™ Enterprise
<ul style="list-style-type: none"> <li>• “Strict” SnapLock <ul style="list-style-type: none"> <li>– Trust nobody</li> </ul> </li> <li>• Permanently non-erasable, non-rewritable disk storage (WORM) <ul style="list-style-type: none"> <li>– Until file expiration</li> <li>– Safe from any keyboard attack</li> </ul> </li> <li>• Complies w/ SEC Regulations <ul style="list-style-type: none"> <li>– Meets SEC 17a-4 requirements</li> <li>– Easy WORM-to-WORM replication</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• “Flexible” SnapLock <ul style="list-style-type: none"> <li>– Trust administrator</li> </ul> </li> <li>• Revision-safe, long-term storage solution <ul style="list-style-type: none"> <li>– Virus and application bug-proof</li> <li>– Enables best practices business records retention</li> </ul> </li> <li>• Partial storage admin control <ul style="list-style-type: none"> <li>– Admin can destroy volumes</li> <li>– Cannot modify/delete individual records</li> </ul> </li> </ul>

SnapLock is available on all FAS and NearStore systems

Figure 6-15 Comparison of SnapLock Compliance and SnapLock Enterprise


## 6.7 Fibre Channel Protocol topologies

There are three basic storage area network (SAN) topologies that are possible when connecting IBM System Storage N series storage systems and server systems with FC:

- ▶ *Direct-attached*: The servers (or hosts) are directly attached to the N series storage controller, as illustrated in Figure 6-16 on page 157.
- ▶ *Single fabric*: The servers are attached to N series storage controllers through a single FC fabric, as illustrated in Figure 6-17 on page 158. This fabric can consist of multiple FC switches.
- ▶ *Dual fabric*: Each server is attached to two physically independent fabrics that are connected to the N series storage controllers, which we illustrate in Figure 6-18 on page 159.

Figure 6-16 on page 157 shows the direct-attached topology.





*Figure 6-16 Direct-attached: one FC target port per controller single\_image topology*

With the direct-attached topology, the servers (or hosts) are directly attached to the N series storage controller, which is supported for all controllers in a non-high availability (non-HA) configuration. Single system\_image (SSI) CFMODE is required for direct-attach in a high-availability (HA) configuration.

Figure 6-17 on page 158 shows the single fabric topology.




Figure 6-17 Single fabric: High-availability storage controller connect topology

With the single fabric topology, the servers are attached to N series storage controllers through a single FC fabric that can consist of multiple FC switches. This configuration is supported in single\_image and standby CFMODEs. When a host has multiple paths to a single LUN that is configured on a storage controller, it is necessary for the host to have multi-pathing software installed. Although in this specific topology with two FC target ports per controller, single-attached hosts do not require multi-pathing software when the controllers are running in standby CFMODE. In single\_image CFMODE, a host with a single FC connection can still have multiple logical paths to the LUN, if there are multiple connections from the N series storage controllers to the fabric.

Figure 6-18 on page 159 shows the dual fabric topology.




Figure 6-18 Dual fabric: two fabrics per storage controller single\_image topology

With the dual fabric topology, each server is attached to two physically independent fabrics that are connected to the N series storage controllers. This topology is possible only with a high-availability storage controller configuration. When properly configured with multi-pathing software, dual-attached hosts are fully redundant from the storage perspective because the host bus adapter (HBA), wiring, fabrics, storage controller, and disks are all redundantly configured. For multiple host configurations, the hosts can be heterogeneous (for example, Windows and UNIX). This configuration is supported with single\_image CFMODE only.

### 6.7.1 FCP topology recommendations

The recommended FCP topology is Dual Fabric, which is supported for all IBM System Storage N series storage systems.

## 6.8 FCP environment

In our environment, we used a N series N5200 G20. The requirements for CPU, storage capacity, and features determine the best model for you, for instance, in the retail environments that have many remote department store locations, the N3300 using iSCSI is a common configuration.

Our environment consists of an IBM N series 5200-G20, Linux server, and an IP Camera. We are working with multipath, and the FCP topology is a dual fabric, Figure 6-18 on page 159, CFMODE that is configured as a single\_image, Figure 6-22 on page 196. SnapDrive is not supported when using multipath at this time. Refer to the interoperability matrix:

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

### 6.8.1 Operating system requirements

The DVS software that we use in this IBM Redbooks publication is the Cisco surveillance software, which uses SUSE Linux Enterprise Server 9 (SLES9) with Service Pack 3. We installed the version with the standard requirements and the following packages:

- ▶ db1, mysql: Cisco video surveillance software requirements
- ▶ gcc: SANsurfer® Linux Driver Installer
- ▶ multipath-tools: multipath service

**Note:** The packages are available on installation CDs. When the package's installation finishes, the next step must be to update the Service Pack 3 upgrade according to the interoperability matrix, and then restart the server.

#### N series requirements

The N series requirements are:

- ▶ Licensing and starting the Cluster service.
- ▶ Licensing and starting the FCP service.

## SAN requirements

The SAN requirements are:

- ▶ At least two switches with enough ports to support your environment. See the compatibility matrix at:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/storageselectproduct?brandind=5000029&familyind=0&continue.x=9&continue.y=7&oldfamily=0>

You also need:

- Two ports per N series node for host access.
  - Two ports per node for back-end storage in loop mode.
  - For a Metrocluster and only on the N series A models, at this time each expansion unit requires two ports for point-to-point connectivity. One or two ports per host system.
  - Two ports per storage subsystem, if you use N series Gateway.
- ▶ In our environment, we required two HBAs per Linux server.
  - ▶ Our IBM N series 5200-G20 had eight FC Ports x 4 Gb (4 ports per node).

## 6.8.2 Setting up N series

The following tasks are common for N series set up and use with DVS.

### Licenses

To verify which licenses are installed on the system, complete the following steps:

1. Enter the command `license`, as shown in Example 6-1.
2. If the licenses are properly added, go to 6.11, “Cluster configuration” on page 193.
3. If the licenses are not properly added, type the following command: `license add license number`, as shown in Example 6-2 on page 162.

#### *Example 6-1 License verification*

---

```
itsonas2> license
      a_sis not licensed
      cifs XXXXXXX
      cluster not licensed
      cluster_remote not licensed
      disk_sanitization not licensed
      fcp not licensed
```

```

flex_cache not licensed
flex_clone not licensed
gateway XXXXXXX
gateway_hitachi not licensed
http site XXXXXXX
iscsi not licensed
multistore not licensed
nearstore_option not licensed
nfs not licensed
smdomino not licensed
smsql not licensed
snaplock not licensed
snaplock_enterprise not licensed
snapmanagerexchange not licensed
snapmirror not licensed
snapmirror_sync not licensed
snapmover not licensed
snaprestore not licensed
snapvalidator not licensed
sv_linux_pri not licensed
sv_ontap_pri not licensed
sv_ontap_sec not licensed
sv_unix_pri not licensed
sv_windows_ofm_pri not licensed
sv_windows_pri not licensed
syncmirror_local not licensed
vld not licensed

itsonas2*>

```

---

Example 6-2 shows the command to add cluster and FCP licenses.

*Example 6-2 Add cluster and FCP licenses*

---

```

itsonas2*> license add XXXXXXX
A cluster license has been installed.
Clustered Failover will be enabled upon reboot.
Make sure that each individual service is licensed
on both nodes or on neither node. Remember to configure
the network interfaces for the other node.
itsonas2*> Fri Sep 7 22:16:21 GMT [itsonas2: rc:notice]: cluster
licensed
itsonas2*> license add XXXXXXX
A fcp site license has been installed.
cf.takeover.on_panic is changed to on
Run 'fcp start' to start the FCP service.

```

```
Also run 'lun setup' if necessary to configure LUNs.  
FCP enabled.  
itsonas2*> Fri Sep 7 22:22:02 GMT [itsonas2: rc:notice]: fcp licensed
```

---

**Configuring the FCP port**

We have two different modes to configure the IBM N series FCP port:

- ▶ Initiator: This is the default mode of the FCP port. You do not have to configure it. The initiator connects the IBM N series to the hosts on the SAN environment.
- ▶ Target: The target mode requires configuration, and the target ports connect the IBM N series to the disk shelves.

To define the onboard FCP port as target:

1. To verify the current status, type **fcadmin config**, as shown in Example 6-3.
2. To configure the FC port as initiator, type **fcadmin config -e target 0b**, as shown in Example 6-4.
3. To verify the new status, type **fcadmin config**, as shown in Example 6-5 on page 164.

*Example 6-3 FC port status initiator*

```
itsonas1*> fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED.	online
0b	initiator	CONFIGURED	online
0c	initiator	CONFIGURED.	offline
0d	initiator	CONFIGURED.	offline

```
itsonas1*>
```

---

Example 6-4 shows the command for FC port configuration.

*Example 6-4 FC port configuration*

```
itsonas1*> fcadmin config -e -t target 0b  
Wed Aug 22 01:21:16 GMT [itsonas1: fci.config.state:info]: Fibre  
channel initiator adapter 0b is in the PENDING (target) state.  
A reboot is required for the new adapter configuration to take effect.  
Error: adapter 0b failed to come online
```

```
itsonas1*> Wed Aug 22 01:21:16 GMT [itsonas1: fci.config.offline:info]:
Fibre channel adapter 0b is offline because it is in the PENDING
(target) state.

itsonas1*>
```

---

Example 6-5 shows the command for the FC port status target.

*Example 6-5 FC port status target*

```
itsonas1*> fccadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED.	online
0b	target	CONFIGURED	online
0c	initiator	CONFIGURED.	offline
0d	initiator	CONFIGURED.	offline

```
itsonas1*>
```

---

**Configuring FCP services**

To verify the FCP services:

1. Type **fcv status**, as shown in Example 6-6.
2. If the fcv services are not running, type **fcv start**, as shown in Example 6-7.
3. To verify the world wide port names (WWPN), type the command **fcv show adapters**, as shown in Example 6-8 on page 165.

*Example 6-6 fcv status*

```
itsonas2*> fcv status
FCV service is not running.
itsonas2*>
```

---

Example 6-7 shows the command for starting FCP services.

*Example 6-7 Starting FCP services*

```
itsonas2*> fcv start
Fri Sep 7 23:26:31 GMT [itsonas2: fcv.service.startup:info]: FCP
service startup
Fri Sep 7 23:26:31 GMT [itsonas2: scsitarget.ispfct.onlining:notice]:
Onlining Fibre Channel target adapter 0a.
```



```
itsonas2*> Fri Sep 7 23:26:32 GMT [itsonas2:
scsitarget.ispfct.linkUp:notice]: Link up on Fibre Channel target
adapter 0a.
```

---

Example 6-8 shows the syntax for the command **fcv show adapters**.

*Example 6-8 fcv show adapters*

---

```
fcv show adapters
Slot:                0a
Description:         Fibre Channel Target Adapter 0a (Dual-channel,
QLogic 2322 (2362) rev. 3)
Adapter Type:       Local
Status:             ONLINE
FC Nodename:        50:0a:09:80:86:57:f3:ac (500a09808657f3ac)
FC Portname:        50:0a:09:81:86:57:f3:ac (500a09818657f3ac)
Standby:            No
itsonas2*>
```

---

### 6.8.3 Setting up the host

In this section, we explain how to set up the host.

#### HBA Driver verification

The Host Bus Adapters (HBAs) that are connected on the Linux server is the QLogic QLA2340 2 Gb fibre Channel HBAs. We work with two HBAs because we recommend working with multipath capability in this SAN environment. The first step is to verify the HBAs BIOS and drivers, which you can do during the POST screen, and you can verify the driver version can be using a command line. We extracted the information about the drivers and BIOS requirements from the interoperability matrix:

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

To verify a version of the HBAs cards driver attached on the host, run the command **cat /proc/scsi/qla2xxx/0**, which gives us all of the information about the HBAs card, including the driver version. Example 6-9 shows the driver version 8.01.02-sles and that an update is required.

*Example 6-9 HBA driver verification*

---

```
tonga:~ # cat /proc/scsi/qla2xxx/0
QLogic PCI to Fibre Channel Host Adapter for QLA2340:
Firmware version 3.03.18 IPX, Driver version 8.01.02-sles
```

```
ISP: ISP2312, Serial# Q08447
Request Queue = 0x37080000, Response Queue = 0x37070000
Request Queue count = 2048, Response Queue count = 512
Total number of active commands = 0
Total number of interrupts = 75
    Device queue depth = 0x10
Number of free request entries = 2046
Number of mailbox timeouts = 0
Number of ISP aborts = 0
Number of loop resyncs = 0
Number of retries for empty slots = 0
Number of reqs in pending_q= 0, retry_q= 0, done_q= 0, scsi_retry_q= 0
Host adapter:loop state = <READY>, flags = 0x1a03
Dpc flags = 0x4000000
MBX flags = 0x0
Link down Timeout = 000
Port down retry = 030
Login retry count = 030
Commands retried with dropped frame(s) = 0
Product ID = 4953 5020 2020 0002
```

SCSI Device Information:

```
scsi-qla0-adapter-node=200000e08b18ff8a;
scsi-qla0-adapter-port=210000e08b18ff8a;
```

FC Port Information:

SCSI LUN Information:

```
(Id:Lun) * - indicates lun is not registered with the OS.
tonga:~ #
```

---

## 6.8.4 Installing the FCP Linux Host Utilities

Verify that you have the correct FCP Host Utilities for your version of Linux. See the appropriate interoperability matrix for your IBM N series product, which is available on the IBM support Web site at:

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

## Downloading the Host Utilities software

To download the iSCSI Linux Host Utilities file:

1. Go to:  
<http://www.ibm.com/servers/storage/support/index.html>
2. In the Select your product box, select **Network attached storage (NAS)**. In the Product family field, select **iSCSI**. In the Product field, select **FCP Host Utilities**.
3. Click **Go**.
4. On the Support for FCP Host Utilities page, select the Download tab, and then perform the proper registration to download the FCP Host Utilities

## Installing the FCP Host Utilities software

To install the Host Utilities software, you must have performed the “HBA Driver verification” on page 165, and the proper version installed on your system, which you can check using the instructions at “Verifying or installing the required RPMs” on page 187.

To install the Host Utilities software:

1. Remove any previous Linux Host attach kit by changing to the directory where the previous version is installed (default is `/opt/sanlun/bin`), and type the `/uninstall` command.
2. Change the working directory to the directory that you copied the Host Utilities file.
3. Type the following command to decompress the file:  

```
tar -xvf ibm_linux_host_utils_3_0.tar
```
4. Change to the `ibm_linux_host_utils_3_0` directory. By default, this directory is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.
5. Type `/install`.

The diagnostic scripts are installed to the `/opt/ontap/santools` directory. This directory is different from the directory used by the previous version of the support kit.

For detailed information about running the diagnostic scripts, see the main pages in the `/opt/ontap/man/man1` directory.

## Removing existing QLogic HBA drivers

You need to remove any existing HBA driver modules.

To remove the QLogic driver modules:

1. Login to the Linux Host as root.
2. Stop qlremote by typing the command **kill qlremote**.
3. Change to the directory where the driver installer files are located.
4. Run the command **/qlinstall --uninstall**.

## Verifying or installing the required RPMs

To verify that the correct multipath RPMs are installed on your Linux host:

1. If you plan to use dm-multipath support, type the following commands, as shown in Example 6-10:

```
rpm -q device-mapper
rpm -q multipath-tools
```

The **rpm -q** commands return the name and version of the installed RPMs. See the Release Notes for the correct RPM version for your specific version of Linux.

2. If you do not have the required RPMs, install them using the appropriate method for your version of Linux.

### *Example 6-10 RPMs verification*

---

```
tonga:~ # rpm -q device-mapper
device-mapper-1.01.01-1.6
tonga:~ # rpm -q multipath-tools
multipath-tools-0.4.5-0.11
tonga:~ #
```

---

## 6.8.5 Installing HBAs

You must verify that your HBA model is supported. See the appropriate interoperability matrix for your N series product, which is available on the IBM support Web site at:

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

If you install more than one HBA, be sure that all of the HBAs are the same brand and bus type (PCI-X or PCI Express). Mixing HBA brands on one host can lead to driver conflicts.

Follow the HBA and host hardware vendors' instructions for installing the HBAs in the Linux host. If the HBAs are already installed, skip this step.

**Note:** Do not cable the HBAs to the switch or storage system yet. If existing HBAs are cabled, remove the Fibre Channel cables from the HBAs at this time.

## Downloading and installing the HBA driver and applications

We use the SANsurfer Linux Driver installer package and the HBA drivers are included on this package. You must download the SANsufer Linux driver installer package from the Qlogic Web site:

[http://support.qlogic.com/support/drivers\\_software.asp](http://support.qlogic.com/support/drivers_software.asp)

Follow the instructions on the Qlogic Web site or in the README file that is included with the driver for installing the driver software on your Linux host.

### Installing the SANsurfer

Be sure that the Fibre Channel cables are not connected to the HBA before you install the driver. A reboot is required after you install the HBA driver package.

To install the SANsufer software, type `./qlinstall`. Example 6-11 shows the SANsurfer installation.

#### Example 6-11 SANsurfer installation

```
tonga:/usr/src/qlafc-linux-8.01.04-3-install # ./qlinstall

#####
#          SANsurfer Driver Installer for Linux          #
#          Installer Version:  1.01.00pre6              #
#####

Kernel version: 2.6.5-7.244-bigsm
Distribution: SUSE LINUX Enterprise Server 9 (i586)

Found QLogic Fibre Channel Adapter in the system
  1. QLA2340
Installation will begin for following driver
  1. qla2xxx version: v8.01.04

Preparing...
#####
qla2xxx
#####
```

QLA2XXX -- Building the qla2xxx driver...

\

QLA2XXX -- Installing the qla2xxx modules to  
/lib/modules/2.6.5-7.244-bigsmpt/kernel/drivers/scsi/qla2xxx/...

Setting up QLogic HBA API library...

Please make sure the /usr/lib/libqlsdp.so file is not in use.

Installing ia32 api binary.

Done.

Unloading any loaded drivers

Unloaded module qla2300

Loading module qla2xxx\_conf version: v8.01.04....

Loaded module qla2xxx\_conf

Loading module qla2xxx version: v8.01.04....

Loaded module qla2xxx

Loading module qla2300 version: v8.01.04....

Loaded module qla2300

Building default persistent binding using SCL

Configuration saved on HBA port 0. Changes have been saved to  
persistent storage.

Please reload the QLA driver module/rebuild the RAM disk for the saved  
configuration to take effect.

Configuration saved on HBA port 0. Changes have been saved to  
persistent storage.

Please reload the QLA driver module/rebuild the RAM disk for the saved  
configuration to take effect.

Configuration saved on HBA port 1. Changes have been saved to  
persistent storage.

Please reload the QLA driver module/rebuild the RAM disk for the saved  
configuration to take effect.

Configuration saved on HBA port 1. Changes have been saved to  
persistent storage.

Please reload the QLA driver module/rebuild the RAM disk for the saved  
configuration to take effect.

Saved copy of /etc/sysconfig/kernel as

/usr/src/qlogic/v8.01.04-3/backup/kernel-2.6.5-7.244-bigsmpt-090507-0928  
16.bak

Saved copy of /etc/modprobe.conf.local as

```
/usr/src/qlogic/v8.01.04-3/backup/modprobe.conf.local-2.6.5-7.244-bigsm  
p-090507-092816.bak
```

```
Saved copy of /boot/initrd-2.6.5-7.244-bigsm as  
/usr/src/qlogic/v8.01.04-3/backup/initrd-2.6.5-7.244-bigsm-090507-0928  
16.bak
```

```
QLA2XXX -- Rebuilding ramdisk image...  
Ramdisk created.
```

```
Reloading the QLogic FC HBA drivers....  
Unloaded module qla2300  
Loading module qla2xxx_conf version: v8.01.04....  
Loaded module qla2xxx_conf  
Loading module qla2xxx version: v8.01.04....  
Loaded module qla2xxx  
Loading module qla2300 version: v8.01.04....  
Loaded module qla2300
```

```
Target Information on all HBAs:
```

```
=====
```

```
-----
```

```
HBA Port 0 - QLA2340 Port Name: 21-00-00-E0-8B-18-FF-8A Port ID:  
01-1C-00
```

```
-----
```

```
-----  
Path : 0  
Target : 0  
Device ID : 0x82  
Port ID : 01-13-00  
Product Vendor : NETAPP  
Product ID : LUN  
Product Revision : 0.2  
Node Name : 50-0A-09-80-86-57-F3-AC  
Port Name : 50-0A-09-81-86-57-F3-AC  
Product Type : Device  
Number of LUN(s) : 0  
Status : Online
```

```
-----
```

```
-----  
Path : 0  
Target : 1  
Device ID : 0x83  
Port ID : 01-11-DA
```

```
Product Vendor      : NETAPP
Product ID          : LUN
Product Revision    : 0.2
Node Name           : 50-0A-09-80-86-57-F3-AC
Port Name           : 50-0A-09-81-96-57-F3-AC
Product Type        : Device
Number of LUN(s)    : 0
Status              : Online
```

```
-----
-----
-----
HBA Port 1 - QLA2340 Port Name: 21-00-00-E0-8B-18-D4-8F Port ID:
01-1D-00
```

```
-----
Path                : 0
Target              : 0
Device ID           : 0x82
Port ID             : 01-13-00
Product Vendor      : NETAPP
Product ID          : LUN
Product Revision    : 0.2
Node Name           : 50-0A-09-80-86-57-F3-AC
Port Name           : 50-0A-09-81-86-57-F3-AC
Product Type        : Device
Number of LUN(s)    : 0
Status              : Online
```

```
-----
Path                : 0
Target              : 1
Device ID           : 0x83
Port ID             : 01-11-DA
Product Vendor      : NETAPP
Product ID          : LUN
Product Revision    : 0.2
Node Name           : 50-0A-09-80-86-57-F3-AC
Port Name           : 50-0A-09-81-96-57-F3-AC
Product Type        : Device
Number of LUN(s)    : 0
Status              : Online
```



```
#####
#               INSTALLATION SUCCESSFUL!!               #
#   SANSurfer Driver installation for Linux completed   #
#####
tonga:/usr/src/qlafc-linux-8.01.04-3-install #
```

---

## 6.8.6 Setting the required HBA and driver parameters

You need to set both HBA and driver parameters to ensure that multipathing and storage system failover work correctly.

### QLogic HBA parameters

To set the HBA parameters, complete the following steps on the Linux host using the **QLogic scli** command:

1. Start the **scli** command in interactive mode by typing **scli**, as shown in Example 6-12.
2. From the main menu, select the **Configure HBA Settings** option.
3. Select the first HBA port that is listed.
4. Select the option for Port Down Retry Count, and if you are using dm-multipath, set the value to **30**. If you are not using dm-multipath, set the value to **180**.
5. Select the option for Link Down Timeout, and set the value to **20**.
6. Repeat the settings for each HBA port that is listed.
7. Return to the main menu and quit.

#### Example 6-12 CLI

---

```
tonga:/usr/src/qlafc-linux-8.01.04-3-install # scli
Scanning QLogic FC HBA(s) and device(s) ...      -

SANSurfer FC HBA CLI

v1.06.16 Build 49

Main Menu

1: Display System Information
2: Display HBA Settings
3: Display HBA Information
4: Display Device List
5: Display LUN List
```

- 6: Configure HBA Settings
- 7: Target Persistent Binding
- 8: Selective LUNs
- 9: Boot Device
- 10: Driver Settings
- 11: HBA Utilities
- 12: Flash Beacon
- 13: Diagnostics
- 14: Statistics
- 15: Help
- 16: Quit

Enter Selection:

---

## QLogic driver parameters

To set the driver parameters, complete the following steps on the Linux host. You must edit either `/etc/modprobe.conf` or `/etc/modprobe.conf.local`, depending on where the QLogic installer put its options:

1. Unload the QLogic HBA driver by typing the appropriate command for your HBA model, as shown in Example 6-13 on page 175:  
`./qlinstall -ul qla2300.`
2. Open the `/etc/modprobe.conf` file with a text editor.
3. Locate the options `qla2xxx` statement that the QLogic installer added. If there is no options `qla2xxx` statement in `/etc/modprobe.conf`, close that file and edit `/etc/modprobe.conf.local` instead.
4. Modify the options `qla2xxx` statement with the following settings on a single line. Only one option `qla2xxx` statement is allowed.

```
ql2xsuspendcount=6 MaxRetriesPerPath=10
MaxRetriesPerIo=10 ql2xfailover=0
```

5. Save file and exit.
6. Unload the QLogic HBA driver using the appropriate command for your HBA mode, as shown in Example 6-13 on page 175:  
`./qlinstall -l qla2300`
7. Reload the QLogic HBA driver using the appropriate command for your HBA mode, as shown in Example 6-14 on page 175:  
`./qlinstall -l qla2300`

8. Enter the appropriate command for your HBA model to rebuild the initial ramdisk (initrd) image with the new parameter values, as show in Example 6-15:

```
./qlinstall -br -in qla2300
```

9. Reboot the Linux host using the updated image.

Example 6-13 shows the syntax for unloading the QLogic HBA driver.

*Example 6-13 Unload the QLogic HBA driver*

---

```
tonga:/usr/src/qlafc-linux-8.01.04-3-install # ./qlinstall -ul qla2300
Unloaded module qla2300
tonga:/usr/src/qlafc-linux-8.01.04-3-install #
```

---

Example 6-14 shows the syntax for reloading the QLogic HBA driver.

*Example 6-14 Reload the QLogic HBA driver*

---

```
tonga:/usr/src/qlafc-linux-8.01.04-3-install # ./qlinstall -l qla2300
Loading module qla2300 version: v8.01.04....
Loaded module qla2300
tonga:/usr/src/qlafc-linux-8.01.04-3-install #
```

---

Example 6-15 shows the syntax for rebuilding the initial ramdisk image.

*Example 6-15 rebuild the initial ramdisk (initrd) image*

---

```
tonga:/usr/src/qlafc-linux-8.01.04-3-install # ./qlinstall -br -in
qla2300
```

```
Saved copy of /etc/sysconfig/kernel as
/usr/src/qlogic/v8.01.04-3/backup/kernel-2.6.5-7.244-bigsm
p-090507-1011
14.bak
```

```
Saved copy of /etc/modprobe.conf.local as
/usr/src/qlogic/v8.01.04-3/backup/modprobe.conf.local-2.6.5-7.244-bigsm
p-090507-101114.bak
```

```
Saved copy of /boot/initrd-2.6.5-7.244-bigsm
p as
/usr/src/qlogic/v8.01.04-3/backup/initrd-2.6.5-7.244-bigsm
p-090507-1011
14.bak
```

```
QLA2XXX -- Rebuilding ramdisk image...
Ramdisk created.
```

## 6.8.7 Recording the World Wide Port Names of the host bus adapters

The storage system uses the World Wide Port Name (WWPN) of the HBA ports to identify them. The WWPN is added to an initiator group (igroup) to enable the initiator to access a particular LUN on the storage system.

The WWPN consists of 16 hexadecimal characters. When you enter the characters to create an igroup, separate each pair of characters with a colon, for example, 21:01:00:E0:8B:3A:BB:30.

Each physical port has its own WWPN. A two-port HBA has two unique WWPNs.

The Fibre Channel Protocol also defines a World Wide Node Name (WWNN).

The WWNN is not used when creating igroups on the storage system.

### Obtaining the World Wide Port Name for QLogic host bus adapters

To obtain the World Wide Port names of QLogic HBAs using the SANsurfer FC HBA CLI client:

1. On the Linux host console, enter the `sccli -i` command, as shown in Example 6-16.
2. Record the WWPN for each port that is listed.

For more information about using the SANsurfer FC HBA CLI client, see the *SANsurfer FC HBA CLI Application User's Guide*. To download the guide, point your browser to the following Web location, and then click your HBA model:

[http://support.qlogic.com/support/drivers\\_software.asp](http://support.qlogic.com/support/drivers_software.asp)

The WWPNs are also printed on a sticker on the bottom of most QLogic HBAs.

#### *Example 6-16 Obtaining the WWPNs*

---

```
tonga:~ # sccli -i
```

```
-----  
-----  
Host Name           : tonga  
HBA Model           : QLA2340  
Port                 : 0  
Node Name            : 20-00-00-E0-8B-18-FF-8A  
Port Name            : 21-00-00-E0-8B-18-FF-8A
```

```
Port ID                : 01-1C-00
Serial Number          : Q08447
Driver Version         : 8.01.04
BIOS Version           : 1.47
Firmware Version       : 3.03.19
Actual Connection Mode : Point to Point
Actual Data Rate       : 2 Gbps
PortType (Topology)    : FPort
Device Target Count    : 2
HBA Status             : Online
```

```
-----
Host Name              : tonga
HBA Model              : QLA2340
Port                   : 1
Node Name              : 20-00-00-E0-8B-18-D4-8F
Port Name              : 21-00-00-E0-8B-18-D4-8F
Port ID                : 01-1D-00
Serial Number          : Q09684
Driver Version         : 8.01.04
BIOS Version           : 1.47
Firmware Version       : 3.03.19
Actual Connection Mode : Point to Point
Actual Data Rate       : 2 Gbps
PortType (Topology)    : FPort
Device Target Count    : 2
HBA Status             : Online
```

```
-----
tonga:~ #
```

## 6.8.8 Loading the host bus adapter driver on the host

Before you load the driver, be sure that you have at least one LUN mapped to the Linux host as LUN 0 before you load the HBA driver.

### Starting Fibre Channel Protocol on a Linux host

To start the Fibre Channel Protocol on a Linux host, load the driver for your HBAs.

If you are using dm-multipath, be sure that you configured multipath before you start FCP.

To load the QLogic HBA driver, on the Linux host, enter the `/sbin/modprobe -v qla2300` command, as shown in Example 6-17.

*Example 6-17 Loading FCP driver*

---

```
tonga:~ # /sbin/modprobe -v qla2300
insmod
/lib/modules/2.6.5-7.244-bigsmp/kernel/drivers/scsi/qla2xxx/qla2300.ko
tonga:~ #
```

---

## 6.8.9 Editing the host's `/etc/multipath.conf` file

The `multipath.conf` file contains configuration settings for multipathing. A sample `multipath.conf` file is copied to your host when you install the `multipath-tools` RPM. The file is:

`/usr/share/doc/packages/multipath-tools/multipath.conf.annotated`

Copy this file to `/etc/`, and rename it `multipath.conf`.

### Editing the file

You need to edit the `/etc/multipath.conf` file to exclude ("blacklist") the local hard drives and other resources that should not be included in the multipathing configuration.

To edit the `/etc/multipath.conf` file to blacklist local drives:

1. Open `/etc/multipath.conf` with a text editor.
2. Locate the following statement:  

```
prio_callout "/sbin/mpath_prio_ontap/dev/%n"
```

Replace it with the following statement:

```
prio_callout "/opt/ontap/santools/mpath_prio_ontap/dev/%n"
```

If the statement is not found, add it to the `/etc/multipath.conf` file as shown in Step 8.
3. Locate the section labeled `## name: blacklist`.
4. Delete the comment characters (`#`) to enable the code in the blacklist section.
5. If the command in the blacklist section is `blacklist`, rename the command **`devnode_blacklist`**.
6. The default command excludes IDE hard drives `/dev/hdx` devices from `/dev/hda` through `/dev/hdz`. However, there might be a syntax error in the

sample /etc/multipath.conf file that prevents local drives from being properly blacklisted.

If the line devnode "**^hd[a-z] [[0-9]\*]**" has two extra bracket characters (bold), delete the bracket so that the line reads as follows:

```
devnode "^hd[a-z] [0-9]*"
```

If the line devnode "**^cciss!c[0-9]d[0-9]\*[p[0-9]\*]**" has two extra bracket characters (bold), delete the brackets so that the line reads as follows:

```
devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"
```

7. If your local drives are IDE drives in the above range, no further changes are needed; save the file and quit.

If your local drives are SCSI drives, add a new line within the **devnode\_blacklist** command to exclude the specific /dev/sdx devices. Be sure to blacklist only your local drives. Each path to an FCP LUN is also a /dev/sdx device that should be configured for multipath, for example, if you have two local SCSI drives, /dev/sda and /dev/sdb, add the following line:

```
devnode "sd[a-b] $"
```

**Note:** Be sure to include the "\$" in the devnode statement when you exclude the local SCSI drives. Without the \$, the multipath code excludes devices, which includes /dev/sdaa, /dev/sdab, and /dev/sdbb, which prevents multipathing from working, if you have many LUN paths.

8. Create a device-specific section at the end of the file for the storage system as follows. These settings apply only to IBM N series storage systems:

```
devices {
  device {
    vendor "NETAPP"
    product "LUN"
    path_grouping_policy group_by_prio
    getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout "/opt/ontap/santools/mpath_
prio_ontap /dev/%n"
    features "1 queue_if_no_path"
    path_checker readsector0
    failback immediate
  }
}
```

9. Save the changes.

Example 6-18 shows the required changes to the `/etc/multipath.conf` file for a system with two local SCSI drives.

*Example 6-18 Editing `/etc/multipath.conf`*

---

```
##
## name : blacklist
## scope : multipath & multipathd
## desc : list of device names to discard as not multipath
##          candidates
## default : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
devnode_blacklist {
devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
devnode "^hd[a-z][0-9]*"
devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"
devnode "sd[a-b]$"
}
#
##
devices {
device {
vendor "NETAPP"
product "LUN"
path_grouping_policy group_by_prio
getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
prio_callout "/opt/ontap/santools/mpath_
prio_ontap /dev/%n"
features "1 queue_if_no_path"
path_checker readsector0
failback immediate
}
}
```

---

### 6.8.10 Starting the multipath service

After you start the HBA drivers, on the Linux host, start the multipath service. To start the multipath service, on the Linux host console:

1. Start the scan for multipath LUNs by entering the `/etc/init.d/boot.multipath start` command.
2. After the command in the step 1 completes, start the multipath daemon by entering the `/etc/int.d/multipathd start` command.

Example 6-19 on page 181 shows the syntax for starting the multipath services.



#### *Example 6-19 Starting multipath services*

---

```
tonga:~ # /etc/init.d/boot.multipath start
Creating multipath targetsdm names    N

done
tonga:~ # /etc/init.d/multipathd start
Starting multipathd
done
tonga:~ #
```

---

### **6.8.11 Configuring the multipath service to start automatically**

After you start the multipath service on the Linux host, you can configure it to start automatically after reboot. To integrate the multipath setup into the boot sequence, on the Linux host console, add the multipath service to the boot sequence by entering the following commands, as shown in Example 6-20:

```
insserv boot.multipath multipathd
chkconfig --add multipathd
chkconfig multipathd on
```

#### *Example 6-20 Adding the multipath service to the boot sequence*

---

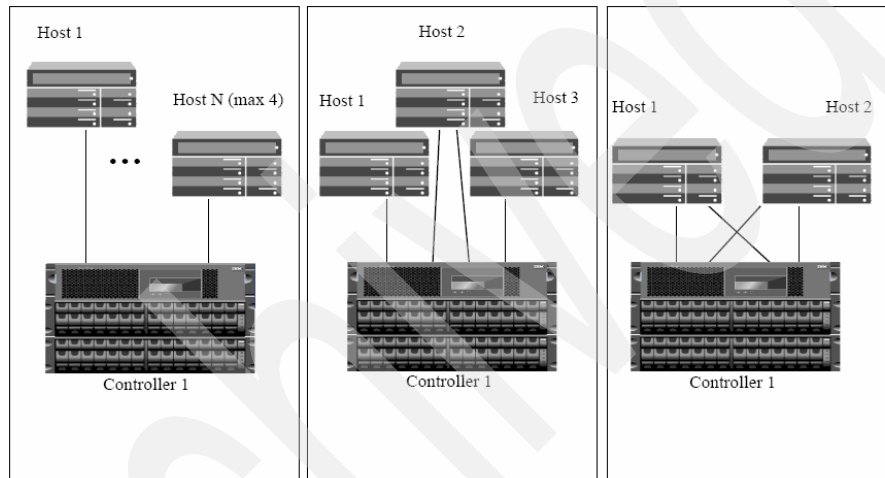
```
tonga:~ # insserv boot.multipath multipathd
tonga:~ # chkconfig --add multipathd
multipathd          0:off 1:off 2:off 3:on  4:off 5:on
6:off
tonga:~ # chkconfig multipathd on
tonga:~ #
```

---

## **6.9 iSCSI topologies**

Figure 6-19 on page 182, Figure 6-20 on page 183, and Figure 6-21 on page 184 show the supported topologies for all IBM System Storage N series storage systems, excluding the N3700 and N3300 that has a limitation of two Ethernet ports. Cases that require more than two Ethernet connections (ports) do not apply to the N3700 and N3300.

Figure 6-19 shows the direct-attached configurations. Because of the non redundant components of the single network and the single N series storage controller, none of these variants can be considered fully redundant. When a host has multiple paths to a single LUN configured on a storage controller (as with Host 2 in the middle configuration or the hosts in the configuration on the right of Figure 6-19), it is necessary for the host to have multi-pathing software installed. For multiple host configurations, the hosts can be heterogeneous (for example, Windows and UNIX) because each storage controller port can customize responses based on the type of operating system that is accessing it. Direct-attached configurations are not supported with HA storage system configurations.



*Figure 6-19 Direct-attached: Single storage controller connect topology*

Figure 6-20 on page 183 shows an IBM System Storage N series storage system that is attached to a single network. The network can consist of one or multiple switches. In fact, the storage controllers can be attached to multiple switches. The storage system can have anywhere from four to eight connections. Because of the non redundant nature of the single network, this configuration cannot be considered fully redundant.

When a host has multiple paths to a single logical unit number that is configured on a storage controller, the host must have multi-pathing software installed. In Figure 6-20 on page 183, all of the hosts have multiple paths because each controller has multiple connections to the network. As an example, Figure 6-20 on page 183 shows two IP connections to the network per storage controller; however, anywhere from one to the maximum number of connections is possible, within the storage controller that is used, are supported.




Figure 6-20 Single network: High availability storage controller connect topology

Figure 6-21 on page 184 shows a configuration with no single points-of-failure for the IBM System Storage N series storage system where two Ethernet ports are used on each storage controller. This configuration is possible only with an HA storage controller configuration. When properly configured with multi-pathing software, the hosts are fully redundant from a storage perspective because the host bus adapters, network interface controllers (NICs), wiring, networks, storage controller, and disks are all redundantly configured. For multiple host configurations, the hosts can be heterogeneous (that is, Windows and UNIX) because each storage controller port can connect to multiple operating systems. A minimum of two connections per storage controller is needed to protect against HBA, network, wiring, or controller failure. As an example, Figure 6-21 on page 184 shows two IP connections to the network; however, anywhere from two to the maximum number of connections possible within the particular storage controller that is used are supported.




Figure 6-21 Dual network: One or more Ethernet ports per storage controller topology

## 6.10 iSCSI environment

For our iSCSI environment, we use IBM N series 5200-G20, Linux server, and an IP Camera. This is an uncommon environment when using iSCSI because a N3300 is a more commonly used N series model for this environment, but generally all of the features and functions of Data ONTAP and N series are the same on the N5200 G20 as on the N3300.

We also work with SnapDrive and the iSCSI topology, which is a single network, as shown in Figure 6-20 on page 183, CFMODE, and it is configured as a single\_image, as shown in Figure 6-22 on page 196, in our environment. SnapDrive for Linux does not currently support multipathing. See the interoperability matrix, for more information:

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

### 6.10.1 Operating system requirements

The Cisco video surveillance software, which we use in this IBM Redbooks publication, uses SUSE Linux Enterprise Server 9 (SLES9) with Service Pack 3.

This version was installed with the standard requirements and the following packages:

- ▶ db1, mysql: Cisco video surveillance software requirements
- ▶ Linux-iSCSI: iSCSI initiator services

**Note:** The packages are available on installation CDs. When the package's installation finishes, the next step is to update the Service Pack 3 upgrade accord to the interoperability matrix, and then restart the server.

## 6.10.2 N series requirements

The N series requirements are:

- ▶ Licensing and starting the Cluster service.
- ▶ Licensing and starting the iSCSI service.

## 6.10.3 Ethernet requirements

The Ethernet requirements are:

- ▶ At least two Ethernet switches and enough ports to support two Ethernet ports for data transfer and communication on the Linux server
- ▶ IBM N series 5200-G20 has eight Ethernet ports x 1 Gb (4ports per node)

## 6.10.4 Host setup

### iSCSI Linux Host Utilities installation

Verify that you have the correct iSCSI Host Utilities for your version of Linux. See the appropriate interoperability matrix for your IBM N series product, which is available on the IBM support Web site at:

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

### Downloading the Host Utilities software

To download the iSCSI Linux Host Utilities file:

1. Go to:

<http://www.ibm.com/servers/storage/support/index.html>

2. In the Select your product box, in the Product family field, select **Network attached storage (NAS) & iSCSI**, and in the Product field, select **iSCSI Host Utilities**. Click **Go**.

3. On the Support for iSCSI Host Utilities page, click the **Download** tab, and then perform the proper registration to download the iSCSI Host Utilities.

## Installing the Host Utilities software

Use the following steps to install the Host Utilities software, as shown in Figure 6-21 on page 184:

1. Remove any previous version of the Host Utilities. Change to the directory where the previous version is installed (default is /opt/ontap/santools), and enter the `./uninstall` command.
2. Change to the working directory to which you saved the Host Utilities file.
3. Enter the `gunzip ibm_linux_host_utils_3_0.tar.gz` command to uncompress the file.
4. Enter the `tar -xvf ibm_linux_host_utils_3_0.tar.gz` command to uncompress the file.
5. Change to the `ibm_linux_host_utils_3_0` directory. By default, this directory is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.
6. Type the `./install` command.

The diagnostic scripts are installed to the /opt/ontap/santools directory. This directory is different from the directory that the previous version of the Host Utilities used.

For detailed information about running the diagnostic scripts, see the main pages in the /opt/ontap/man/man1 directory.

### *Example 6-21 iSCSI Host Utilities installation*

---

```
lochnese:/opt/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
lochnese:/opt/ibm_linux_host_utils_3_0 #
```

---

## Verifying or installing the required RPMs

To verify that the correct iSCSI and multipathing RPMs are installed on your Linux host:

Enter the `rpm -q linux-iscsi` command, as shown in Example 6-22. The `rpm -q` command returns the name and version of the iSCSI RPMs. See the *Release Notes* for the correct RPM version for your specific version of Linux.

If you do not have the correct RPMs for your version of SLES, install the required RPM versions from the media for your version of SLES. You can use the `rpm` or `yast` command.

### Example 6-22 RPMs verification

---

```
lochnese:~ # rpm -q linux-iscsi
linux-iscsi-4.0.1-88.26
lochnese:~ #
```

---

## Recording or changing the initiator node name on the host

The initiator node name is required to create igroups on the storage system. You map igroups to specific LUNs. Only the hosts in the igroup can discover the LUNs as local devices.

The default initiator node name uses the following format:

```
iqn.1987-05.com.cisco:RandomNumber
```

We recommend that you set the initiator node name to the following, where host-name is the host name of your Linux host:

```
iqn.1987-05.com.cisco: host-name
```

You can change the initiator node name or use the default node name.

To view or change the default node name, complete the following steps:

1. With a text editor, open the host's `/etc/initiatorname.iscsi` file.
2. Use the following table to decide your next step.

If ...	Then ...
If a default initiator name was not generated	Replace the text on the host's <code>/etc/initiatorname.iscsi</code> file with the following line: <code>iqn.1987-05.com.cisco:RandomNumber</code> The recommended value for <code>RandomNumber</code> is the host name of the Linux host. Record the node name and proceed to 6.8.9, "Editing the host's <code>/etc/multipath.conf</code> file" on page 178.
You want to use the default initiator node name	Record the node name, and proceed to 6.8.9, "Editing the host's <code>/etc/multipath.conf</code> file" on page 178.
You want to change the initiator node name	<ul style="list-style-type: none"><li>► Modify the <code>RandomNumber</code> part of the initiator node name. The following line shows an example node name. <code>iqn.1987-05.com.cisco:linux-host1</code>, as shown in Example 6-23.</li><li>► Record the node name and proceed to 6.8.9, "Editing the host's <code>/etc/multipath.conf</code> file" on page 178.</li></ul> The <code>RandomNumber</code> is the only component of the node name that you modify. You do not modify the other components of the node name. We recommend that you use the host name in place of the default random number.

Example 6-23 iSCSI Initiator node name

```
## DO NOT EDIT OR REMOVE THIS FILE!  
## If you remove this file, the iSCSI daemon will not start.  
## If you change the InitiatorName, existing access control lists  
## may reject this initiator. The InitiatorName must be unique  
## for each iSCSI initiator. Do NOT duplicate iSCSI InitiatorNames.  
InitiatorName=iqn.1987-05.com.cisco:lochnese  
~  
~  
~
```

6,44

A11



## Node name rules

If you change the host's initiator node name, be sure that the new name follows all of these rules:

- ▶ A node name can be up to 223 bytes.
- ▶ Uppercase characters are always mapped to lowercase characters.
- ▶ A node name can contain alphabetic characters (a to z), numbers (0 to 9) and three special characters:
  - Period (“.”)
  - Hyphen (“-”)
  - Colon (“:”)
- ▶ The underscore character (“\_”) is not supported.

## Editing the host's /etc/iscsi.conf file

You edit the /etc/iscsi.conf file to configure:

- ▶ Required iSCSI parameter settings
- ▶ The storage system's IP address

To edit the /etc/iscsi.conf file:

1. With a text editor, open the host's /etc/iscsi.conf file.
2. At the beginning of the file, add the following lines, and make sure that you add these lines before any other settings:

```
Continuous=no
HeaderDigest=never
DataDigest=never
```

These lines are required for proper operation. If you do not add these lines at the beginning of the file, you will experience problems with connectivity and performance.

3. Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you will use for iSCSI LUNs:

```
DiscoveryAddress=storage_system_IPaddress
```

storage\_system\_IPaddress is the IP address of an Ethernet interface on the storage system. Specify an interface that will be used for iSCSI communication. Gigabit Ethernet interfaces are strongly recommended.

**Example:** The following syntax is a sample DiscoveryAddress entries using storage system IP addresses:

```
DiscoveryAddress=192.168.10.100
DiscoveryAddress=10.61.208.200
```

### Starting the iSCSI service on the host

To start the iSCSI service, at the Linux host command prompt, start the iSCSI service by entering the `/etc/init.d/iscsi start` command, as shown in Example 6-24.

*Example 6-24 Starting the iSCSI services*

```
lochnese:~ # /etc/init.d/iscsi start
Starting iSCSI: iscsi iscsid fsck/mount
done
lochnese:~ #
```

### Configuring the iSCSI service to start automatically

To configure the iSCSI service to start automatically at system bootup:

1. Verify the status of the iSCSI service by typing the **chkconfig iscsi** command.
2. Use the following table to help you decide on your next step.

If...	Then...
The <b>chkconfig</b> command indicates that the iSCSI service is enabled, as follows: iscsi on	No action is needed.
The <b>chkconfig</b> command indicates that the iSCSI service is not enabled as follows: iscsi off	Start the iSCSI service by entering the <b>chkconfig iscsi on</b> command.

Example 6-25 shows the syntax to configure the iSCSI services to start automatically.

*Example 6-25 Configuring the iSCSI services*

```
lochnese:~ # chkconfig iscsi
iscsi off
lochnese:~ # chkconfig iscsi on
lochnese:~ # chkconfig iscsi
iscsi on
lochnese:~ #
```

## 6.10.5 N series setup

For N series set up, you must verify which licenses are installed on the system. To verify the license, type the **license** command, as shown in Example 6-26. If the license is not properly added, type the **license add license\_number** command. Run the **add** command until all of the necessary license are added on your system.

*Example 6-26 License verification*

---

```
itsonas2> license
      a_sis not licensed
      cifs XXXXXXXX
      cluster not licensed
      cluster_remote not licensed
      disk_sanitization not licensed
      fcp not licensed
      flex_cache not licensed
      flex_clone not licensed
      gateway XXXXXXXX
      gateway_hitachi not licensed
      http site XXXXXXXX
      iscsi not licensed
      multistore not licensed
      nearstore_option not licensed
      nfs not licensed
      smdomino not licensed
      smsql not licensed
      snaplock not licensed
      snaplock_enterprise not licensed
      snapmanagerexchange not licensed
      snapmirror not licensed
      snapmirror_sync not licensed
      snapmover not licensed
      snaprestore not licensed
      snapvalidator not licensed
      sv_linux_pri not licensed
      sv_ontap_pri not licensed
      sv_ontap_sec not licensed
      sv_unix_pri not licensed
      sv_windows_ofm_pri not licensed
      sv_windows_pri not licensed
      syncmirror_local not licensed
      vld not licensed

itsonas2*>
```

---

Example 6-27 is the syntax for adding clusters and iSCSI licenses.

*Example 6-27 Add cluster and iSCSI licenses*

---

```
itsonas2*> license add XXXXXX
A cluster license has been installed.
    Clustered Failover will be enabled upon reboot.
    Make sure that each individual service is licensed
    on both nodes or on neither node. Remember to configure
    the network interfaces for the other node.
itsonas2*> Fri Sep 7 22:16:21 GMT [itsonas2: rc:notice]: cluster
licensed
itsonas2*> license add XXXXXX
A iscsi site license has been installed.
cf.takeover.on_panic is changed to on
Run 'iscsi start' to start the iSCSI service.
Also run 'lun setup' if necessary to configure LUNs.
    iSCSI enabled.
itsonas2*> Tue Sep 11 23:06:47 GMT [itsonas2: rc:notice]: iscsi
```

---

## Configuring iSCSI services

To verify the iSCSI services:

1. Type **iscsi status**, as shown in Example 6-28. If the iSCSI service is not running, type **iscsi start**, as shown in Example 6-29 on page 193.
2. To add the iSCSI node name, type **iscsi nodename iscsi\_node\_name**.
3. To verify that the iSCSI node name was successfully added, type **iscsi initiator show**. This command gives you information about the iSCSI node name, as shown in Example 6-30 on page 193.
4. To verify that the iSCSI session connected on N series, type **iscsi session show**, as shown in Example 6-31 on page 193.

*Example 6-28 iscsi status*

---

```
itsonas2*> iscsi status
iSCSI service is not running
itsonas2*>
```

---

*Example 6-29 iscsi start*

---

```
itsonas2*> iscsi start
Fri Sep 7 23:35:00 GMT [itsonas2: iscsi.service.startup:info]: iSCSI
service startup
iSCSI service started
itsonas2*>
```

---

*Example 6-30 iscsi initiator show*

---

```
itsonas2*> iscsi initiator show
Initiators connected:
  TSIH  TPGroup  Initiator
    1    1000    lochnese (iqn.1987-05.com.cisco:lochnese /
00:02:3d:00:00:01)
itsonas2*>
```

---

*Example 6-31 iscsi session show*

---

```
itsonas2*> iscsi session show
Session 1
  Initiator Information
    Initiator Name: iqn.1987-05.com.cisco:lochnese
    ISID: 00:02:3d:00:00:01
    Initiator Alias: lochnese

itsonas2*>
```

---

## 6.11 Cluster configuration

To enable and verify the cluster status:

1. To verify the cluster status, type **cf status** and **cf monitor**, as shown in Example 6-32.
2. If the cluster services are not running, type the **cf enable** command, as shown in Example 6-33 on page 194.

*Example 6-32 cf status and cf monitor*

---

```
itsonas1> cf status
Cluster enabled, itsonas2 is up.
itsonas1*> cf monitor
```

```
current time: 05Sep2007 21:43:16
UP 4+23:21:23, partner 'itsonas2', cluster monitor enabled
VIA Interconnect is up (link 0 up, link 1 up), takeover capability
on-line
partner update TAKEOVER_ENABLED (05Sep2007 21:43:15)
itsonas1*>
```

---

#### *Example 6-33 cf enable*

---

```
itsonas1*> cf enable
itsonas1*> Wed Sep  5 02:33:35 GMT [itsonas1:
cf.misc.operatorEnable:warning]: Cluster monitor: operator initiated
enabling of cluster
Wed Sep  5 02:33:35 GMT [itsonas1:
cf.fsm.takeoverOfPartnerEnabled:notice]: Cluster monitor: takeover of
itsonas2 enabled
Wed Sep  5 02:33:35 GMT [itsonas1:
cf.fsm.takeoverByPartnerEnabled:notice]: Cluster monitor: takeover of
itsonas1 by itsonas2 enabled
itsonas1*> cf status
Cluster enabled, itsonas2 is up.
itsonas1*> Wed Sep  5 02:34:00 GMT [itsonas1:
monitor.globalStatus.ok:info]: The system's global status is normal.
itsonas1*>
```

---

### 6.11.1 Cluster failover

Cluster failover mode (CFMODE) is a storage system Fibre Channel Protocol failover mode setting that controls how the system handles Fibre Channel before and after cluster failover events. CFMODE is not available and is irrelevant on non-clustered storage controllers. You can change the mode using the **fcp set** command from the console after setting the privileges to advanced. On this environment, we use the single system image (single\_image or SSI), which has the following features:

- ▶ Supports all host OS types
- ▶ Supports all IBM System Storage N series and switches
- ▶ Efficient use of FC target ports
- ▶ All logical unit numbers (LUNS) are visible on all ports

### 6.11.2 CFMODE types

In this section, we discuss the types of CFMODE.

## Single system image CFMODE

Single system image, also called `single_image` or SSI, CFMODE is available on all IBM clustered N series storage systems with ONTAP 7.1 and later. SSI also allows all ports on the local controller to be usable for local LUN access and eliminates the port burn issue. SSI requires only a single port per controller to protect against head and fabric failure:

- ▶ Features:
  - Supports all IBM System Storage N series and switches
  - Efficient use of FC target ports
  - All logical unit numbers (LUNS) are visible on all ports
- ▶ Considerations:
  - None

Figure 6-22 on page 196 shows a dual fabric-attached configuration for the N7000 storage system, where dual-ported 4 Gb FCP target mode cards are used in each storage controller. In Figure 6-22 on page 196, the 4 Gb FC target port numbers (5a and 5b) are examples and, in reality, vary depending on the expansion slot where the FC target cards are installed in the controllers. This configuration is possible only with an active/active storage controller configuration. When properly configured with multi-pathing software, dual-attached hosts are fully redundant from the storage perspective because the HBA, wiring, fabrics, storage controller, and disks are all redundantly configured. For multiple host configurations, the hosts can be heterogeneous (for example, Windows and UNIX). In this topology, the 2 Gb onboard FC ports (0a and 0c) are used for back-end disk shelf storage connections. This configuration is supported with the `single_image` (SSI) CFMODE.




Figure 6-22 Single\_image CFMODE

## Dual fabric CFMODE

Dual fabric CFMODE is available only on the N3700 storage system. In this mode, LUNs that are served by a storage controller are always accessible from the FC ports on either of the storage controllers. For performance reasons, we recommend that you access the LUNs primarily through the FC ports on the storage controller that services that LUN. Supported multi-pathing software solutions take care of this selection automatically. Because this CFMODE requires a fabric loop connection type, it is not supported on some switches, such as the older model enterprise class McData switches:

- ▶ Features:
  - Supports all operating systems
  - Requires fewer switch ports
- ▶ Considerations:
  - Requires loop mode



## Partner CFMODE

Partner CFMODE, shown in Figure 6-23, is the default CFMODE for Data ONTAP 7.1. This mode is available on N5000 series clustered storage systems. In this mode, LUNs served by a storage controller are always accessible from the “A” FC ports on the storage controller that serve the LUN and on the “B” FC ports on the partner storage controller. For performance reasons, we recommend that you access the LUNs primarily through the FC ports on the controller that services that LUN. Supported multi-pathing software solutions take care of this selection automatically.

Partner CFMODE uses fabric logins and is supported on all switches:

- ▶ Features:
  - Supports all operating systems
  - Supports all switches
  - Easy to manage
- ▶ Considerations:
  - Requires more switch ports and wiring because both the A and B target FC ports must be connected

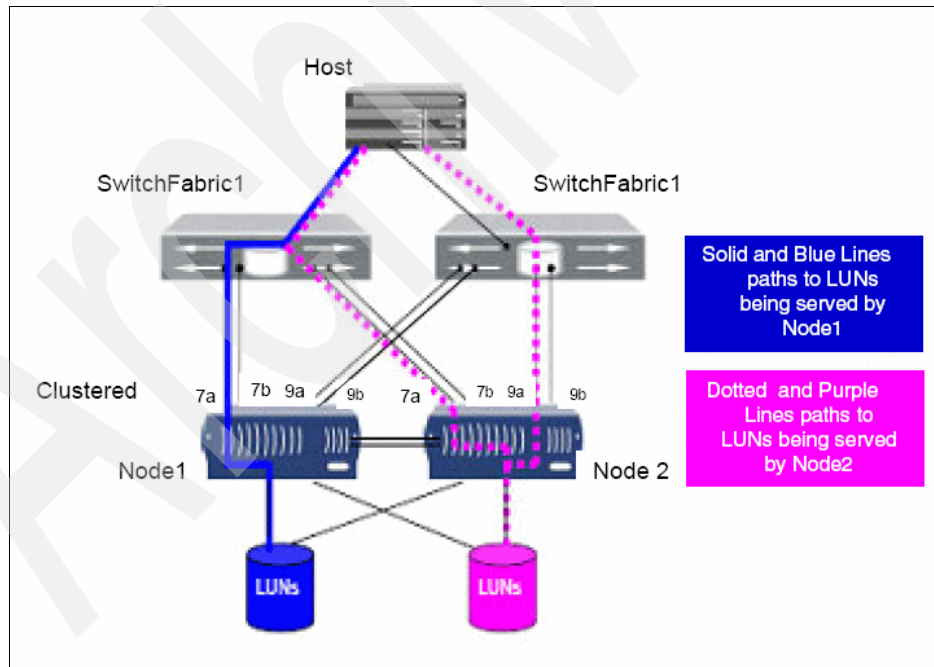


Figure 6-23 Partner CFMODE

## Standby CFMODE

Standby mode is available on N5000 series clustered storage systems. In this mode, LUNs served by a storage controller are always accessible from the “A” FC ports on the storage controller serving the LUN and are first made available on the “B” FC ports on the partner storage controller during a failover event. This CFMODE uses fabric logins and is supported on all switches:

- ▶ Features:
  - Supports all switches
  - Allows multiple active ports with the ASL 2.0 for Veritas
  - Windows Multi Path Input Output (MPIO) Version 3.2 will round-robin LUNs across the available primary paths
- ▶ Considerations:
  - Does not support all operating systems
  - Requires more switch ports and wiring, because both the A and B target FC ports must be connected
  - Requires a minimum of two FC target cards per storage controller or four target ports per controller on an N5000 series system

## Mixed CFMODE

Mixed CFMODE, which is shown in Figure 6-24 on page 199, is available on N5000 series clustered storage systems. This CFMODE is equivalent to standby mode for Windows and Solaris hosts and partner mode for AIX and Hewlett-Packard UNIX (HP-UX) hosts except that instead of using the physical A and B FC ports, this functionality is available on each physical FC port through virtual ports. Because the mixed CFMODE requires a fabric loop connection type, it is not supported on some switches, such as the older model enterprise class McData switches:

- ▶ Features:
  - Supports all operating systems
  - Requires fewer ports
- ▶ Considerations:
  - Requires loop mode

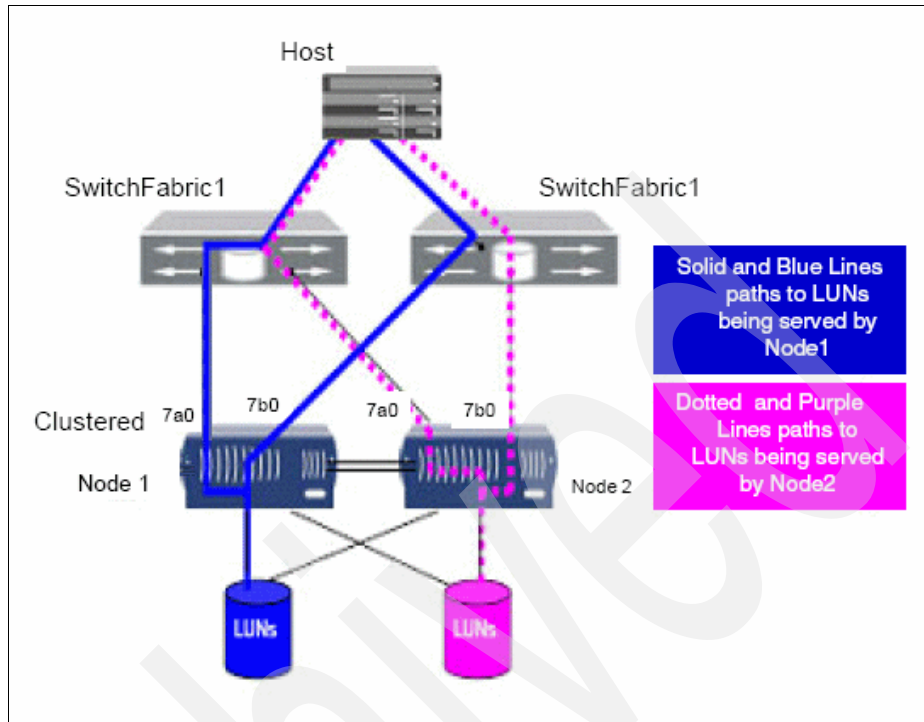


Figure 6-24 Mixed CFMODE

### CFMODE recommendations

The recommended CFMODE setting is to use single\_image CFMODE for N3700 storage system, N3000 series, N5000 series, and N7000 series clustered storage systems.

#### 6.11.3 CFMODE set up

To verify the cfmode status:

1. Verify the which cfmode you are working in by typing **fcpl show cfmode**, as shown in Example 6-34 on page 200.
2. Execute the **fcpl set cfmode** command with the proper parameters, as shown in Example 6-35 on page 200:

```
fcpl set cfmode [ -f ] { dual_fabric | mixed | partner | standby |
single_image }
```

```
fcpl set cfmode single_image
```

*Example 6-34 fcp show cfmode*

---

```
itsonas1*> fcp show cfmode  
fcp show cfmode: standbyfcp
```

---

*Example 6-35 fcp set mode*

---

```
itsonas1*> fcp set cfmode single_image  
Are you sure you want to change the cfmode?  
The cfmode setting must be the same on both filers in a cluster.  
Improper setting of this option can cause host outages, continue? yes  
Thu Sep 6 16:26:03 GMT [itsonas1:  
scsitarget.ispfct.cfmodeChanged:warning]: FCP cfmode changed from  
standby to single_image.  
itsonas1*>
```

---

# Setting up the Digital Video Surveillance Solution

In this chapter, we describe how to set up the DVS solution for the:

- ▶ N series
- ▶ Linux server
- ▶ DVS application
- ▶ Camera

## 7.1 N series configuration for a Fibre Channel Protocol environment

To administer the N series:

1. Connect with Internet Explorer® at the IP address of the node1:  
[http://www.ip-address.com/na\\_admin/](http://www.ip-address.com/na_admin/)
2. Enter credentials, as shown in Figure 7-1.




Figure 7-1 Credentials

3. Open the FilerView as, shown in Figure 7-2 on page 203 and Figure 7-3 on page 203.

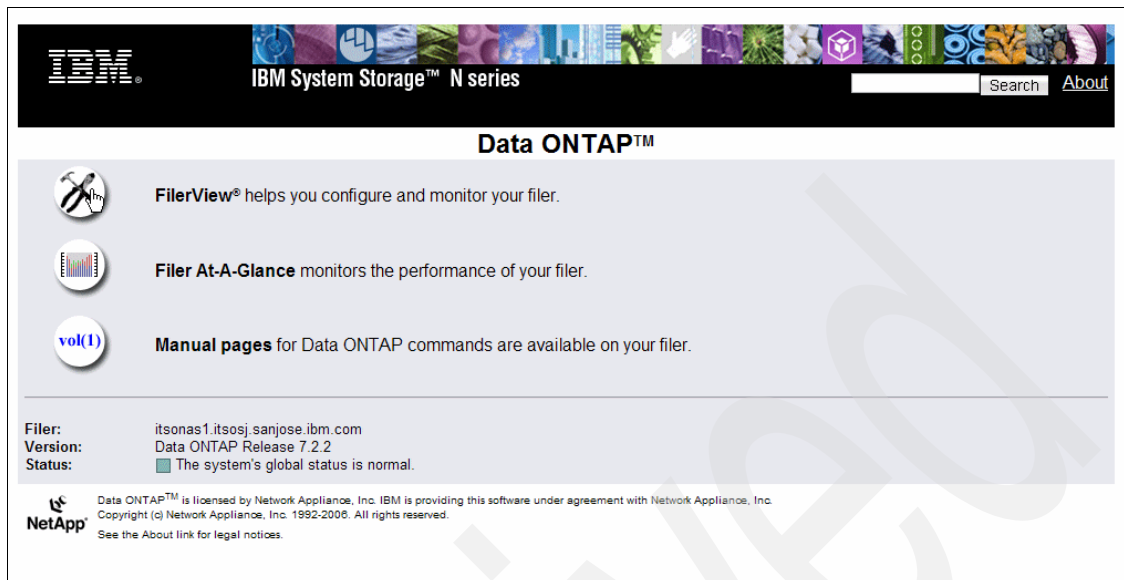


Figure 7-2 Welcome window

- Click **FilerView**. The first window of the FilerView with system status is displayed, as shown in Figure 7-3.

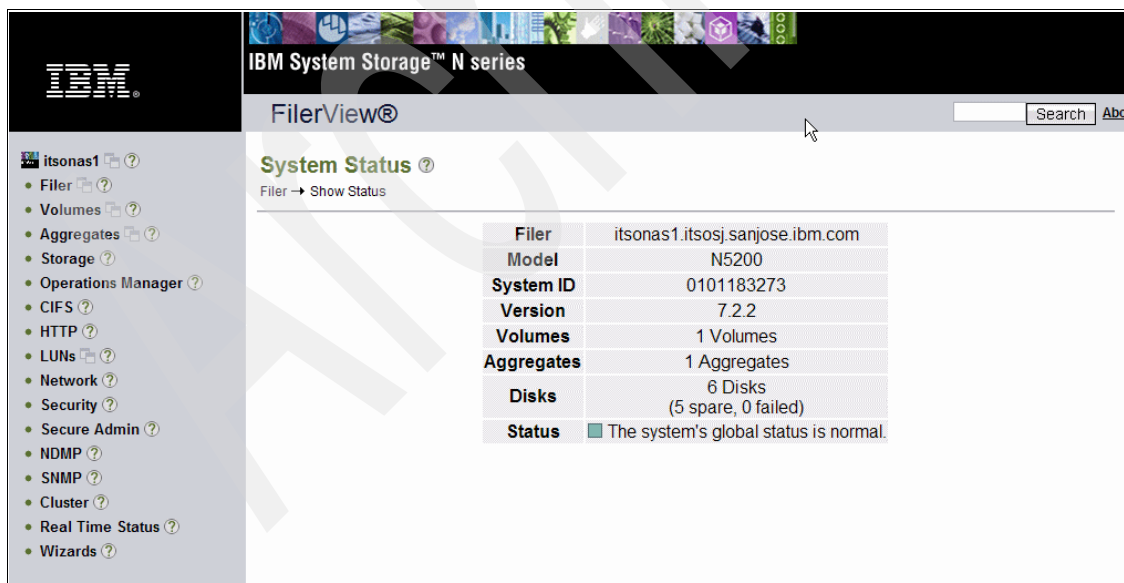


Figure 7-3 FilerView - System Status

## 7.1.1 Initial configuration

We used a N5200-G20 with a DS4700 to provide the LUNS. Initial configuration includes one aggregate per node, with root volume in this aggregate. Root volume is where Data ONTAP is installed as well as configuration files.

Figure 7-4 shows all of the disks as seen by itsonas1. To get to Figure 7-4, we connected to FilerView using the IP address of the node1 in our Internet browser and selected **Storage** → **Disks** → **Manage**:

- ▶ Disk itsosan02:24:126:L10 was used by itsonas1 as indicated in column type with mention data, and one aggregate, aggr0\_node1, was defined on this disk in the column aggregate.
- ▶ Disks itsosan02:24:126:L0 to itsosan02:24:126:L4 were available on this node, as indicated in column type with mention spare.
- ▶ Disks itsosan02:24:126:L5 to itsosan02:24:126:L9 and itsosan02:24:126:L11 belonged to itsonas2, as indicated in column type with mention partner.

IBM System Storage™ N series

FilerView®

Manage Disks ?

Storage → Disks → Manage

View Type: All Disks View

	Disk	Type	Checksum	Shelf	Bay	Chan	Used	Physical	Pool	Aggregate
<input type="checkbox"/>	itsosan02:24:126L0	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L1	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L2	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L3	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L4	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L5	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L6	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L7	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L8	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L9	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/>	itsosan02:24:126L10	data	zoned/block	?	?		17 GB	18 GB	Pool0	aggr0_node1
<input type="checkbox"/>	itsosan02:24:126L11	partner	zoned/block	?	?		0 MB	18 GB	Pool0	

Select All - Unselect All

Disks: 1-12 of 12

Refresh

Fail Remove

Figure 7-4 Node 1 - Storage / Disks / Manage - Initial Aggregate Configuration



Figure 7-5 displays information about the initial aggregates of our configuration that belong to itsonas1. Under **Aggregates** → **Manage**, there is one aggregate for the root volume aggr0\_node1. Figure 7-5 provides the following information:

- ▶ Size
- ▶ Number of disks inside
- ▶ The root volume of the itsonas1 is defined on this aggregate with the column Root checked



Figure 7-5 Node1 - Aggregates / Manage - Initial Aggregate Configuration

Figure 7-6 on page 206 shows the volumes in the initial configuration for itsonas1:

- ▶ There was one unique volume. Initial name was vol0, but we renamed it vol0\_node1 for more clarity.
- ▶ This volume was in aggregate aggr0\_node1.
- ▶ This volume was root volume as indicated by column root checked.



Figure 7-6 Node 1- Volumes / Manage - Initial Volume Configuration

## 7.1.2 Creating the aggregate

To create the aggregate:

**Important:** In our configuration with FCP, we created one aggregate where we put all of the volumes that are associated with the camera we use. This aggregate is created on itsonas1. By doing this, only resources of the itsonas1 (memory and processors) are used to provide data. itsonas2 remains available in a cluster configuration in case of failure. In a production environment with several cameras, files of the different cameras are distributed in a minimum of two aggregates, each one being associated to one node, to ensure good performance with both nodes serving data.

1. In FilerView, as shown in Figure 7-7 on page 207, under Aggregates, click **Add**.
2. A wizard is displayed, as shown in Figure 7-8 on page 207. Click **Next**.

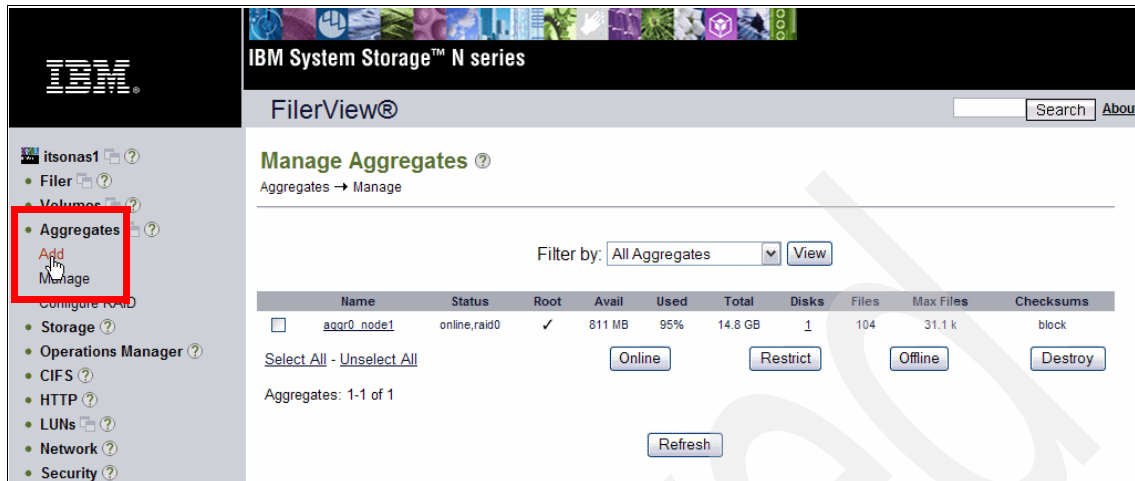


Figure 7-7 Add an Aggregate

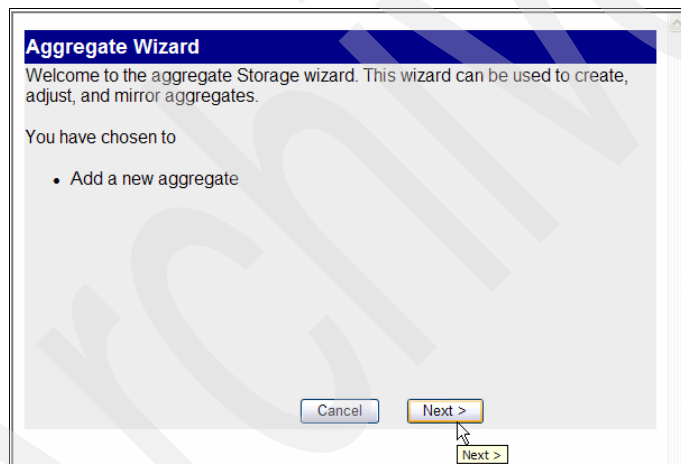


Figure 7-8 Aggregate Wizard: Add an Aggregate

3. In the next wizard, Figure 7-9 on page 208, define the name of the aggregate.

**Note:** We are using a Gateway model, so we do not have the choice for the RAID level. It is RAID0 over all the disks that we include in the aggregate, with data striped over all of the devices. Using a Gateway, if you want to implement RAID policy with parity disks or mirroring, you must implement it at the level of the storage system whose LUNs are presented to the Gateway.



Figure 7-9 Aggregate Wizard: Aggregate Name

4. In Figure 7-10, define the number of disks in the RAID groups. If you want to use RAID-DP, this is where specify RAID-DP groups of 14 disks, for example, with 12 disks for data and two disks for parity. In our case, we want to have one disk in the aggregate, so we chose one disk per RAID group.



Figure 7-10 Aggregate Wizard: RAID Group Size

5. As shown in Figure 7-11 on page 209, you can choose automatic or manual disk selection. Automatic disk selection means that the disks are automatically selected for you. Manual disk selection means that you can choose the disks yourself.

We selected manual disk selection.

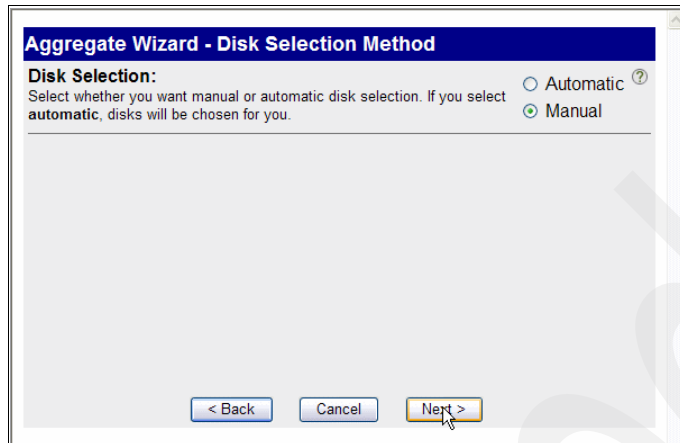


Figure 7-11 Aggregate Wizard: Disk Selection

6. In Figure 7-12, you select the disk, and click **Next**. We chose the ID and disk size and selected one disk: itsosan02:24.126L0. The disks we select from the list will form the aggregate.




Figure 7-12 Aggregate Wizard: Disks to add

A window appears, Figure 7-13 on page 210, that provides a summary of the new aggregate. Verify the settings, and click **Commit**.




Figure 7-13 Aggregate Wizard: Summary

A status message is displayed about the creation of the new aggregate.

7. In the FilerView option, under Aggregates → Manage, verify the successful creation of the new aggregate, as show in Figure 7-14.



Figure 7-14 Aggregates / Manage - Aggregate created

### 7.1.3 Creating the volumes

In the aggregate we used two volumes: one of 15 GB for the archives, and one of 2 GB for the clips. To create the volumes:

1. Open the FilerView in your browser, and click Volumes → Add, as shown in Figure 7-15.



Figure 7-15 Volumes / Add

2. Click **Next** to create a new volume on the filer, as shown in Figure 7-16.




Figure 7-16 Volume Wizard: Add a new volume

3. The FilerView Add Volume window, Figure 7-17, choose the type of volume that you want to create: traditional volumes or flexible volumes. We created a flexible volume.



Figure 7-17 Volume Wizard: Volume Type Selection

Figure 7-18 shows the Volume Name and the Language setting. For our environment, we chose US English.

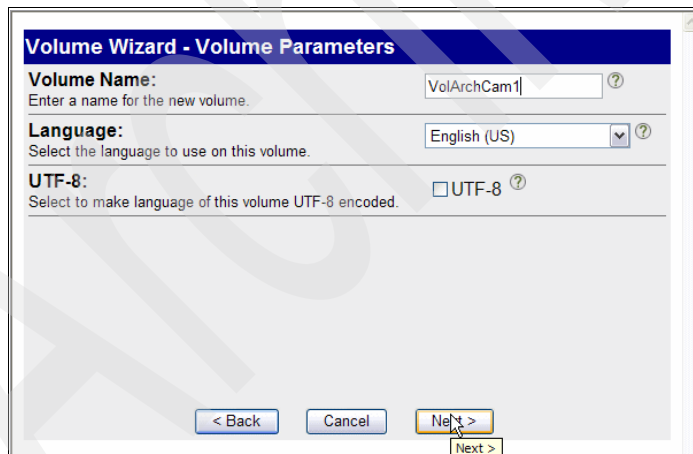


Figure 7-18 Volume Wizard: Volume Parameters



4. In Figure 7-19, choose the aggregate that your volume will reside in and the space guarantee.

There are three types of space guarantee:

- A space guarantee of **volume** preallocates space in the aggregate for the volume. The preallocated space cannot be allocated to any other volume in that aggregate. The space management for a FlexVol volume that has a space guarantee of volume is equivalent to a traditional volume.
- A space guarantee of **file** preallocates space in the aggregate so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are pinned for a Snapshot copy.
- A FlexVol volume that has a space guarantee of **none** reserves no extra space. Writes to LUNs or files that are contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

We chose **volume** as the space guarantee type for our volumes.



Figure 7-19 Volume Wizard: Volume Parameters 2

5. In Figure 7-20 on page 214, you define how much space you want for the volume and the percentage of space that will be allocated for snapshot, included in the size you enter (Total size option) or added to the size which will be really usable (Usable size option). Choose the size of your volume. We have 18.5 GB available.

For the snapshot reserve, we chose only 5% (default is 20%) because DVS data is not normally modified and users do not often need to come back at a previous state. In DVS context, recovery can be needed but more in a

situation of disaster recovery and to prevent this type of failure to impact the infrastructure, a strategy, such as SnapMirror, is more suited than Snapshot.

The screenshot shows a web browser window titled "http://9.43.86.31 - itsonas1: Volume Wizard - Microsoft Internet Explorer". The main content area is titled "Volume Wizard - Flexible Volume Size". It contains the following sections:

- Volume Size Type:** Two radio buttons are present: "Total Size" (selected) and "Usable Size". A help icon (?) is next to "Total Size".
- Volume Size:** A text input field contains "15", followed by a "GB" dropdown menu. A help icon (?) is next to the dropdown. Below the input field, it says "Enter the desired volume size. The containing aggregate, **aggr1\_node1** has a maximum of 18.5 GB space available." and "18.5 GB (Max)".
- Snapshot Reserve:** A text input field contains "5", followed by a "%" dropdown menu. A help icon (?) is next to the dropdown. Below the input field, it says "Enter the snapshot reserve for volume 'VolArchCam1'. The range is between 0% and 50%. The default is 20%."

At the bottom of the dialog, there are three buttons: "< Back", "Cancel", and "Next >".

Figure 7-20 Volume Wizard: Flexible Volume Size

6. The FilerView displays a summary, as shown in Figure 7-21, about the volume that will be created. Select **Commit**.

The screenshot shows a web browser window titled "Volume Wizard - Commit". The main content area is titled "Volume Wizard - Commit" and contains the text "Below is a summary of your changes." followed by a text box with the following summary:

```
Create New Volume
Volume Name: VolArchCam1
Aggregate Container: aggr1_node1 (18.5 GB, raid0)
Volume Size: 15 GB
Snapshot Reserve: 5%
Language: English (US) (en_US)
Space Guarantee: volume
```

At the bottom of the dialog, there are three buttons: "< Back", "Cancel", and "Commit". A mouse cursor is pointing at the "Commit" button.

Figure 7-21 Volume Wizard: Summary

7. A message is displayed, as shown in Figure 7-22 on page 215, that tells you that the Volume updated successfully. Click **Close** to exit the wizard.




Figure 7-22 Confirmation of the volume creation

8. Under **Volumes** → **Manage**, verify that the new volume is now ready for use, as shown in Figure 7-23.



Figure 7-23 Volumes / Manage - Volume created

9. Repeat steps 1 through 9 to create the second volume for the clips. The total size of this volume is 2 GB; otherwise, choose the same options and parameters that you chose for the first volume.

In Figure 7-24 on page 216, under Volumes → Manage, final configuration of the volumes for itsonas 1 is listed.

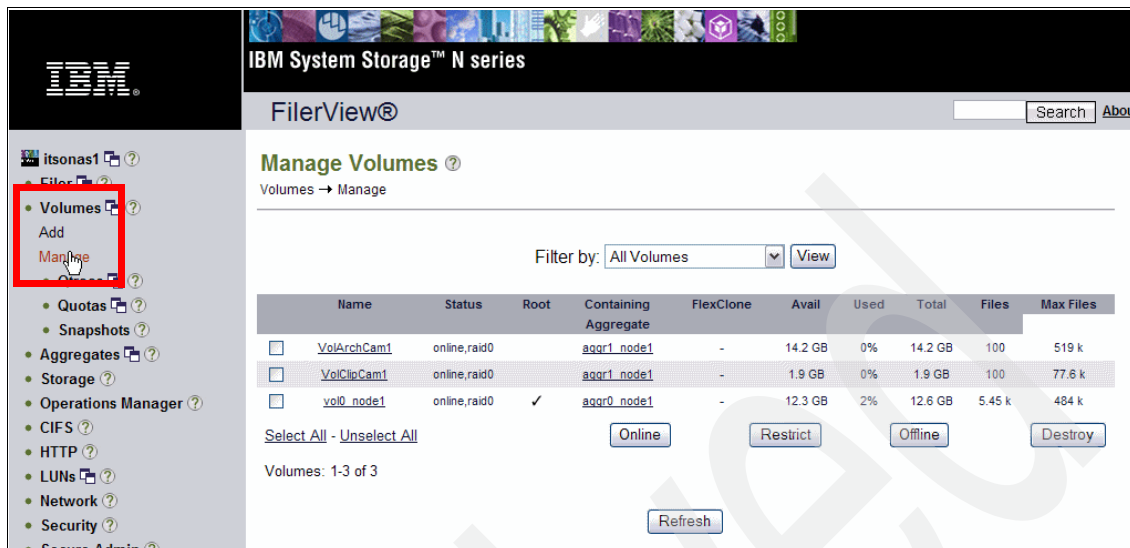


Figure 7-24 Volumes / Manage - Volumes created

### About Fractional Reserve:

Fractional reserve enables you to tune the amount of space that is reserved for overwrites based on application requirements and the rate of change of your data. You define fractional reserve settings per volume, for example, you can group LUNs with a high rate of change in one volume and leave the fractional reserve setting of the volume at the default setting of 100 percent. You can group LUNs with a low rate of change in a separate volume with a lower fractional reserve setting and therefore make better use of available volume space.

In our configuration, we put fractional reserve at zero with the following command:

```
vol options VolArchCam1 fractional_reserve 0
```

## 7.1.4 Creating the LUN

We have two volumes available, and we created one LUN in each volume using the following steps:

1. In FilerView, select **LUNs** → **Add**, to create the LUN for the archives, as shown in Figure 7-25:
  - Path: indicates the name of the LUN, LUNArchCam1 here, and the volume in which the LUN was created, VolArchCam1 here. All of these parameters are given in the path: /vol/VolArchCam1/LUNArchCam1
  - LUN Protocol Type: the type of operating system that will access the LUN. In our case, it was Linux.
  - Description: is an optional parameter. We put LUN for archives.
  - The volume VolArchCam1 had 14.2 GB available. We chose 14 GB for the size of the LUN.
  - By default, keep the Space Reserved option to avoid problems with allocated space.

Click **Add** at the bottom of the page, and a message is displayed:

LUN create: succeeded



Figure 7-25 Add a LUN

2. Under LUNs, select **Manage**, verify that the LUN is created, as shown in Figure 7-26.

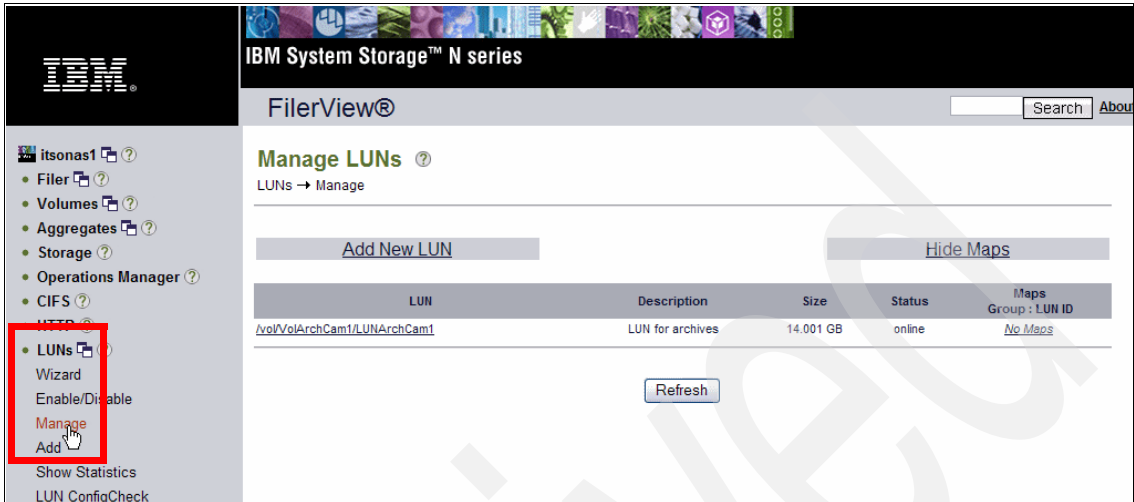


Figure 7-26 LUNs /Manage - LUN created

3. Repeat steps 1 and 2 to create the LUN of 1GB for the Clips, LUNClipCam1, in the volume VolClipCam1. Figure 7-27 shows output under **LUNs** → **Manage**, after the second LUN is created.




Figure 7-27 LUNs / Manage - Listing of the LUNs

## 7.1.5 Mapping LUNs to initiator groups

We must define which HBA cards can access which LUNs. Mapping takes two steps.

### Defining the initiator group

One initiator group can be:

- ▶ The worldwide name of an HBA card
- ▶ Several worldwide names of several HBA cards that are grouped together

Define initiator groups to indicate which LUN is associated to which initiator group and so, which LUN is seen by this HBA card or this group of HBA cards:

1. Select **LUN → Initiator Groups → Add**. Figure 7-28 on page 220 is displayed, where you perform the following actions:
  - a. Name this initiator group.
  - b. Indicate which protocol (FCP or iSCSI) will access the LUNs that are associated to this initiator group.
  - c. Indicate which operating system was installed on the server that will access the LUNs.
  - d. Indicate worldwide names of the HBA cards of the server that will access the LUNs.




Figure 7-28 Initiator Group Definition

2. Click **Add** at the bottom of the page. The following message is displayed:  
Initiator group create: succeeded
3. Under LUNs, select **Initiator Groups** → **Manage**, as shown in Figure 7-29 on page 221, verify that the initiator group was created.






Figure 7-29 LUNs / Initiator Groups / Manage - Listing of initiator groups

## Mapping the LUN

Indicate which LUNs to associate to the initiator group that we created:

- Under **LUNs** → **Manage**, Figure 7-30 is displayed. Look in the column Maps Group: LUN ID, and see that the value in the field is **no maps**, which means that no mapping was defined. Click the **No Maps** link.

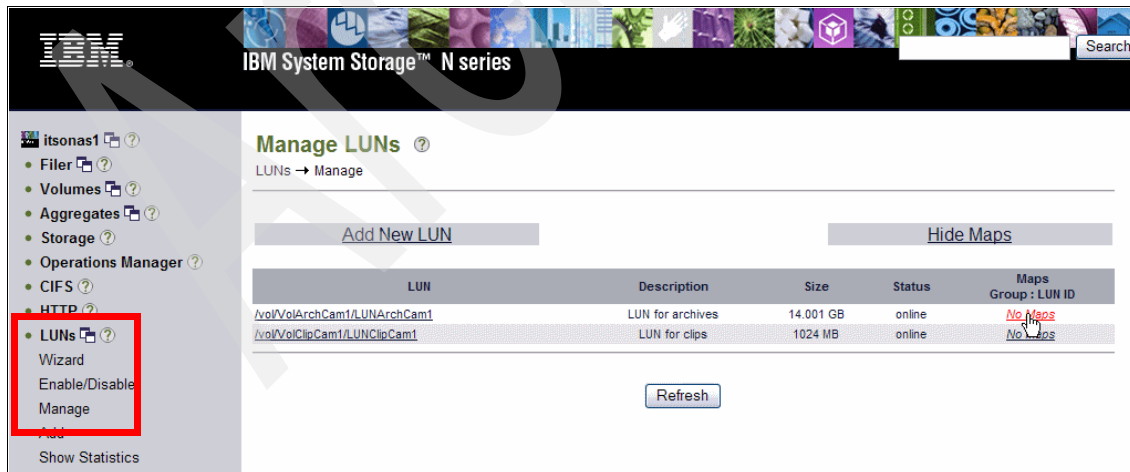


Figure 7-30 Mapping definition

The the current map that is associated with this LUN is displayed, as shown in Figure 7-31, which means that there is no map at the moment.

2. Click **Add Groups to Map**.



Figure 7-31 Initial Empty Mapping

3. As shown in Figure 7-32, select which initiator group to associate with this LUN. We selected the initiator group that we previously created, **DVS\_CAM1**. Click the **Add** button at the bottom of the page.



Figure 7-32 Initiator Group Selection

- As shown in Figure 7-33, define the LUN ID, which is how we differentiate the different LUNs that are associated to a single initiator group. The best thing is to begin with LUN ID 0, and increment one by one. In our case, it was the first LUN defined for this initiator so we chose **LUN ID 0**.

Click **Apply**.

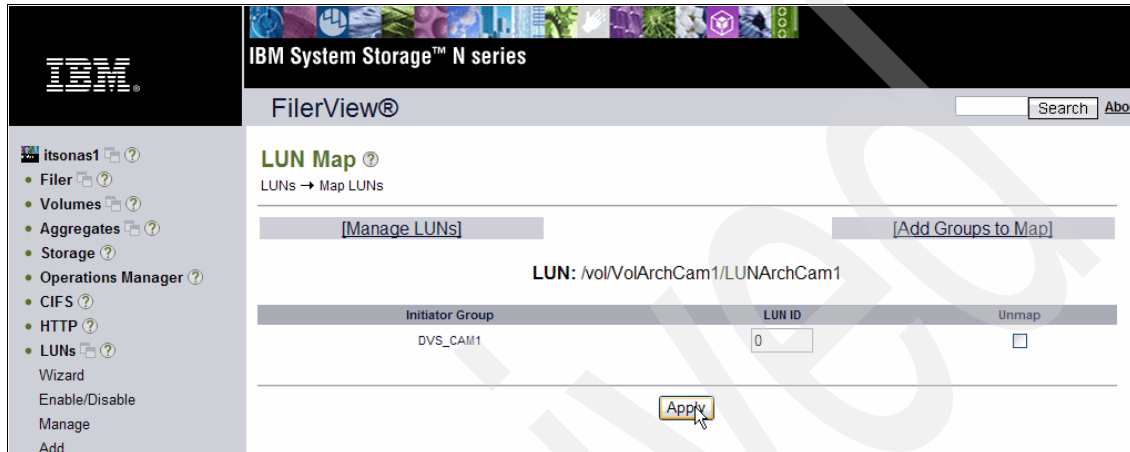


Figure 7-33 LUN ID selection

A message indicates that the LUN mapping was a success.

- Select **LUNs → Manage** to verify that the map is now defined for this LUN, as shown in Figure 7-34. LUNArchCam1 is now mapped with the initiator group DVS\_CAM1, with LUN ID 0.




Figure 7-34 Verification of the map

- Repeat the same steps for LUNClipCam1 with the same initiator group corresponding to the HBA cards of our Linux DVS server. It was the second LUN associated to this initiator group, so we chose LUN ID 1.

In FilerView, select LUNs → Manage, and the final map for both LUNs is displayed, as shown in Figure 7-35.

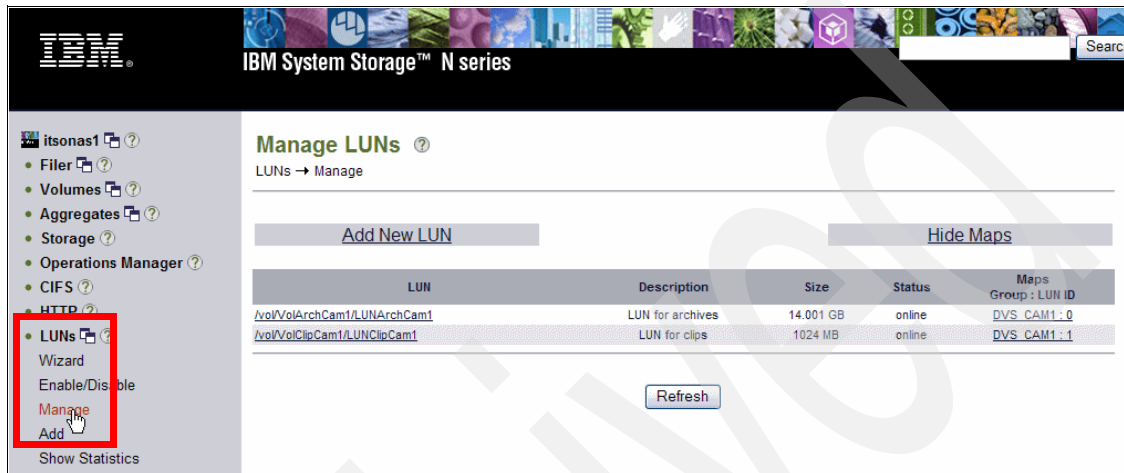


Figure 7-35 Map for the LUNs

## 7.2 Configuring Linux for the FCP environment

Complete these steps to use the LUNs on the Linux server that we created in N series:

- The new LUNs that we created in 7.3.4, “Creating the LUN” on page 234, are not discovered by the OS, as shown in Figure 7-36. Run the **sanlun lun show all** command:

```
tonga:~ # sanlun lun show all
no Filer LUNs available
```

Figure 7-36 List of the LUNs: no LUNs at the beginning

- On the Linux server, re-scan what was presented to the HBA cards to discover the LUNs that we just created. Use the following command, as shown in Figure 7-37 on page 225:

```
/opt/ontap/santools/qla2xxx_lun_rescan all
```

```
tonga:~ # /opt/ontap/santools/qla2xxx_lun_rescan all
performing scan on existing LUNs ...
scanning /proc/scsi/qla2xxx/0 ...
performing scan on existing LUNs ...
scanning /proc/scsi/qla2xxx/1 ...
```

Figure 7-37 Discovering New LUNs

3. Verify that the LUNs were discovered and are available, as shown in Figure 7-38. We recognized the LUNs that we created by running the **sanlun lun show all** command.

**Important:** The output in Figure 7-38 shows one LUN that is mapped to the Linux host from each of the two N series storage systems. There are two paths from the host to each storage system, which is why each LUN appears twice.

```
tonga:~ # sanlun lun show all
filer:          lun-pathname          device filename  adapter
protocol      lun size              lun state
itsonas1: /vol/VolArchCam1/LUNArchCam1 /dev/sda         host0
FCP          14.0g (15033434112)    GOOD
itsonas1: /vol/VolClipCam1/LUNClipCam1 /dev/sdb         host0
FCP          1g (1073741824)       GOOD
itsonas1: /vol/VolArchCam1/LUNArchCam1 /dev/sdc         host0
FCP          14.0g (15033434112)    GOOD
itsonas1: /vol/VolClipCam1/LUNClipCam1 /dev/sdd         host0
FCP          1g (1073741824)       GOOD
itsonas1: /vol/VolArchCam1/LUNArchCam1 /dev/sde         host1
FCP          14.0g (15033434112)    GOOD
itsonas1: /vol/VolClipCam1/LUNClipCam1 /dev/sdf         host1
FCP          1g (1073741824)       GOOD
itsonas1: /vol/VolArchCam1/LUNArchCam1 /dev/sdg         host1
FCP          14.0g (15033434112)    GOOD
itsonas1: /vol/VolClipCam1/LUNClipCam1 /dev/sdh         host1
FCP          1g (1073741824)       GOOD
```

Figure 7-38 Listing LUNs

**Note:** If have the graphical interface started, you might see a notice that a new LUN is available and proposing that you graphically configure it. In our installation, we went on with the manual configuration.

4. Create multipath devices that correspond to your LUNs multipath devices to work on. Use the **mult ipath** command.

Figure 7-39 on page 227 shows the result of this command with the multipath devices created:

- 360a980004334682f636f4457756f2f72 for disks `sda`, `sd c`, `sde`, and `sdg`, which correspond to `/vol/VolArchCam1/LUNArchCam1`, as shown in Figure 7-40 on page 228.
- 360a980004334682f636f4457766a7131 for disks `sdb`, `sdd`, `sdf`, and `sdh`, which correspond to `/vol/VolClipCam1/LUNClipCam1`, as shown in Figure 7-40 on page 228.

```

tonga:~ # multipath
dm names      N
dm info 360a980004334682f636f4457756f2f72  N
dm create 360a980004334682f636f4457756f2f72
360a980004334682f636f4457756f2f72 0
dm reload 360a980004334682f636f4457756f2f72 0
dm resume 360a980004334682f636f4457756f2f72  N
dm message 360a980004334682f636f4457756f2f72  N  switch_group 1

create: 360a980004334682f636f4457756f2f72
[size=14 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [prio=8]
\_ 0:0:1:0 sdc 8:32 [ready]
\_ 1:0:1:0 sdg 8:96 [ready]
\_ round-robin 0 [prio=2]
\_ 0:0:0:0 sda 8:0 [ready]
\_ 1:0:0:0 sde 8:64 [ready]

dm info 360a980004334682f636f4457766a7131  N
dm create 360a980004334682f636f4457766a7131
360a980004334682f636f4457766a7131 0
dm reload 360a980004334682f636f4457766a7131 0
dm resume 360a980004334682f636f4457766a7131  N
dm message 360a980004334682f636f4457766a7131  N  switch_group 1

create: 360a980004334682f636f4457766a7131
[size=1 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [prio=8]
\_ 0:0:1:1 sdd 8:48 [ready]
\_ 1:0:1:1 sdh 8:112 [ready]
\_ round-robin 0 [prio=2]
\_ 0:0:0:1 sdb 8:16 [ready]
\_ 1:0:0:1 sdf 8:80 [ready]

```

Figure 7-39 Create multipath devices

5. Look under /dev/disk/by-name/ to verify that the symbolic links were created, as shown in Figure 7-40 on page 228.

```
tonga:~ # ls -l /dev/disk/by-name/*  
lrwxrwxrwx 1 root root 10 Sep 14 14:26  
/dev/disk/by-name/360a980004334682f636f4457756f2f72 -> ../../dm-0  
lrwxrwxrwx 1 root root 10 Sep 14 14:26  
/dev/disk/by-name/360a980004334682f636f4457766a7131 -> ../../dm-1
```

*Figure 7-40 List of the multipath devices*

6. Create a file system of xfs type, as required for DVS installation in the new multipath devices. We used the following commands, as shown in Figure 7-41 on page 229:

```
mkfs -t xfs /dev/dm-0  
mkfs -t xfs /dev/dm-1
```



```

tonga:~ # mkfs -t xfs /dev/dm-0
dm_task_set_name: Device /dev/dm-0 not found
Command failed
dm_task_set_name: Device /dev/dm-0 not found
Command failed
meta-data=/dev/dm-0            isize=256    agcount=16,
agsize=229392 blks
        =                       sectsz=512
data      =                     bsize=4096    blocks=3670272,
imaxpct=25
        =                       sunit=0      swidth=0 blks,
unwritten=1
naming    =version 2           bsize=4096
log        =internal log       bsize=4096    blocks=2560,
version=1
        =                       sectsz=512    sunit=0 blks
realtime  =none                extsz=65536   blocks=0, rtextents=0
tonga:~ # mkfs -t xfs /dev/dm-1
dm_task_set_name: Device /dev/dm-1 not found
Command failed
dm_task_set_name: Device /dev/dm-1 not found
Command failed
meta-data=/dev/dm-1            isize=256    agcount=8,
agsize=32768 blks
        =                       sectsz=512
data      =                     bsize=4096    blocks=262144,
imaxpct=25
        =                       sunit=0      swidth=0 blks,
unwritten=1
naming    =version 2           bsize=4096
log        =internal log       bsize=4096    blocks=2560,
version=1
        =                       sectsz=512    sunit=0 blks
realtime  =none                extsz=65536   blocks=0, rtextents=0

```

Figure 7-41 Creation of the File System

7. Put the available disks at the Linux level by performing the following steps:
  - a. Create the following directories:
 

```

mkdir /Archives
mkdir /Clips

```

b. Add the following lines in /etc/fstab.

```
/dev/dm-0 /Archives xfs defaults 0 0  
/dev/dm-1 /Clips xfs defaults 0 0
```

```
/dev/dm-0    /Archives    xfs    defaults 0 0  
/dev/dm-1    /Clips       xfs    defaults 0 0
```

c. Mount the file systems with the **mount -a** command.

8. Using the **df** command, review the available LUNs, as shown in Figure 7-42.

```
tonga:~ # mount -a  
tonga:~ # df
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/hda2	73911032	3127980	70783052	5%	/
tmpfs	1876552	12	1876540	1%	/dev/shm
/dev/dm-0	14670848	272	14670576	1%	/Archives
/dev/dm-1	1038336	144	1038192	1%	/Clips

Figure 7-42 List of file system available

## 7.3 Configuring N series for the iSCSI environment

To work with itsonas2, we used an Internet browser to connect to FilerView with IP address of the itsonas2:

[http://www.ip-address/na\\_admin](http://www.ip-address/na_admin)

For more information about this process, refer to the beginning of 7.1, “N series configuration for a Fibre Channel Protocol environment” on page 202.

### 7.3.1 Initial Configuration

The initial configuration is very similar to node1 configuration because node2 used disks that were created on the same DS4500.

The aggregate configuration with utilization of the disks for each node is under:

**Storage → Disks → Manage**

**Aggregates → Manage**

The volume configuration for each node is under:

**Volumes → Manage**

As shown in Figure 7-43, all of the disks as seen by node2 are visible. To get information about them, select **Storage → Disks → Manage**.

Node2 used disk itsosan02:24:126:L11, which is indicated in the Type column with mention data, and aggr0\_node2 was defined on this disk, which we can see in the Aggregate column:

- ▶ Disks itsosan02:24:126:L5 to itsosan02:24:126:L9 are available on this node, as indicated in the Type column with mention spare.
- ▶ Disks itsosan02:24:126:L0 to itsosan02:24:126:L4 and itsosan02:24:126:L10 belong to node1, as indicated in the Type column with mention partner.

IBM System Storage™ N series  
FilerView®

Manage Disks ?  
Storage → Disks → Manage

View Type: All Disks View

Disk	Type	Checksum	Shelf	Bay	Chan	Used	Physical	Pool	Aggregate
<input type="checkbox"/> itsosan02:24:126L0	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L1	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L2	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L3	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L4	partner	zoned/block	?	?		0 MB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L5	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L6	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L7	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L8	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L9	spare	zoned/block	?	?		22 GB	22 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L10	partner	zoned/block	?	?		0 MB	18 GB	Pool0	
<input type="checkbox"/> itsosan02:24:126L11	data	zoned/block	?	?		17 GB	18 GB	Pool0	aggr0_node2

Select All - Unselect All Fail Remove

Disks: 1-12 of 12 Refresh

Figure 7-43 Node2 - Storage/Disks/Manage - Initial Aggregate Configuration

As shown in Figure 7-44 on page 232, information about the initial aggregates of our configuration that belong to node 2 are visible.

Under **Aggregates** → **Manage**, there is one root aggregate to aggr0\_node2. The following information was provided:

- ▶ Size
- ▶ Number of disks inside
- ▶ The fact that the root volume of node 1 is defined on this aggregate with the column Root checked



Figure 7-44 Node2 - Aggregates / Manage - Initial Aggregate Configuration

As shown in Figure 7-45 on page 233, the volumes are present in the initial configuration for node2:

- ▶ There was one unique volume vol0\_node2.
- ▶ This volume was in aggregate aggr0\_node2.
- ▶ This volume was root volume, as indicated by the root checked column.

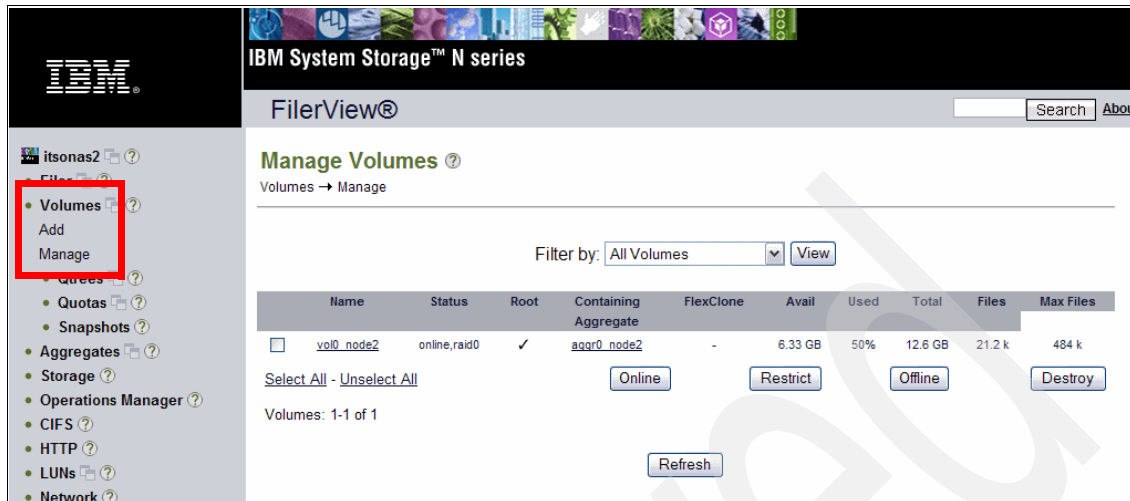


Figure 7-45 Node2 - Volumes / Manage - Initial Volume configuration

### 7.3.2 Creating the aggregate

Because we previously provided the process of creating an aggregate in 7.1.2, “Creating the aggregate” on page 206, and configuration of itsonas2 is very similar to the configuration of the node1, we do not repeat those instructions here.

We created one aggregate aggr0\_node2 in node2 with the same options as aggr0\_node1 in node1:

- ▶ Name: aggr0\_node2
- ▶ RAID group size: 1
- ▶ Disk Selection: Manual
- ▶ Disk used: itsonas02:24.126L9

### 7.3.3 Creating the volumes

Because we previously provided the process of creating the volumes in 7.1.3, “Creating the volumes” on page 211, and configuration of node2 is very similar to the configuration of the node1, we do not repeat those instructions here.

We created two volumes in aggr0\_node2, one of 15 GB for the archives and one of 2 GB for the clips:

- ▶ Volume VolArchCam2:
  - Volume Type: Flexible
  - Volume Name: VolArchCam1
  - Language: English US
  - Containing Aggregate: aggr1\_node2
  - Space Guarantee: volume
  - Volume Size Type: Total Size
  - Volume Size: 15 GB
  - Snapshot Reserve: 5%
- ▶ Volume VolClipCam2:
  - Volume Type: Flexible
  - Volume Name: VolClipCam1
  - Language: English US
  - Containing Aggregate: aggr1\_node2
  - Space Guarantee: volume
  - Volume Size Type: Total Size
  - Volume Size: 2GB
  - Snapshot Reserve: 5%

### 7.3.4 Creating the LUN

Because we previously provided the process for creating the LUNs in 7.1.4, “Creating the LUN” on page 217, and the configuration of node2 is very similar to the configuration of the node1, we do not repeat those instructions here.

We created two LUNs:

- ▶ LUN for archives:
  - Path: /vol/VolArchCam2/LUNArchCam2
  - LUN protocol type: Linux
  - Size: 14 GB
  - Space reserved: yes
- ▶ LUN for clips:
  - Path: /vol/VolClipCam2/LUNClipCam2
  - LUN protocol type: Linux
  - Size: 1 GB
  - Space reserved: yes

### 7.3.5 Mapping the LUNs to the initiator groups

In this section, we associate the LUNs with the iSCSI initiator, that is which LUNs is accessible by which servers. The iSCSI initiators define these servers, which can be of two different types:

- ▶ iSCSI cards
- ▶ iSCSI third-party software initiators

As we explained in a previous section (5.1, “Hardware” on page 120), in our environment we used iSCSI third-party software initiator, which is included by default in the SUSE Linux Distribution.

#### Defining the initiator group

We defined the iSCSI initiator name in a previous chapter, (6.8.3, “Setting up the host” on page 165, Example 6-23 on page 188). We must add it as a definition at the N series level, as shown in Figure 7-46 on page 236.

1. Under LUN, select **Initiator Groups** → **Add**:
  - Group name (name of the server): lochnese
  - Type: iSCSI
  - Operating System: Linux
  - Initiators: the name we defined in `/etc/initiatorname.iscsi:iqn.1987-05.com.cisco:lochnese` (6.8.3, “Setting up the host” on page 165, and more precisely Example 6-23 on page 188)

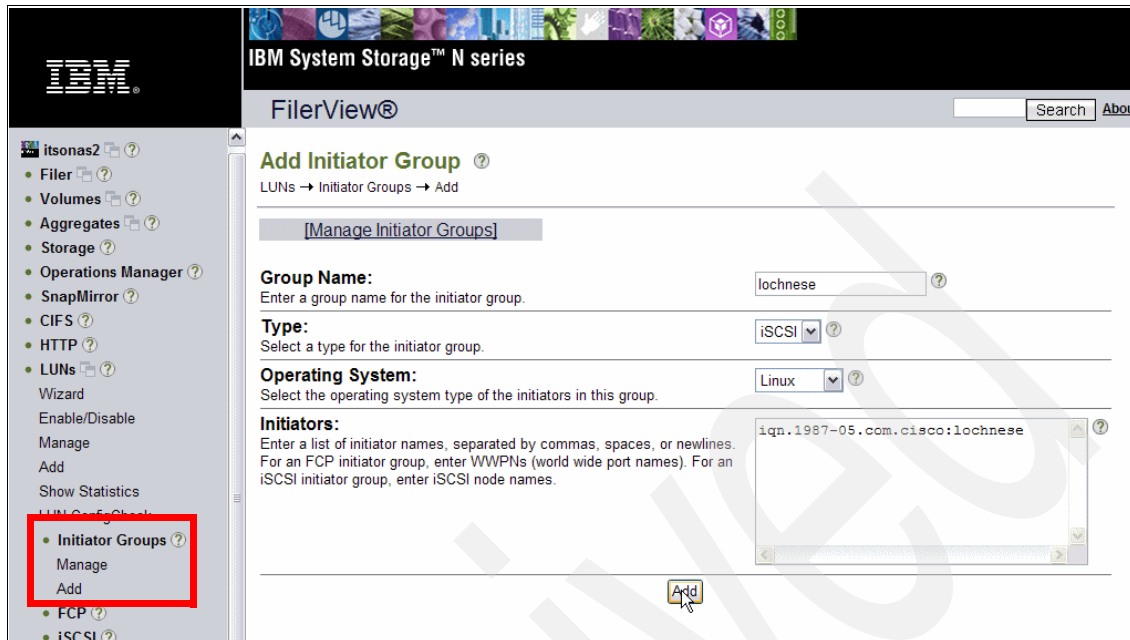


Figure 7-46 iSCSI initiator Definition

2. Click **Add** at the bottom of the page, and the following message is displayed:  
Initiator Group create: succeeded
3. Under LUNs, select Initiator Groups → Manage, verify that the initiator group was created, as shown in Figure 7-47 on page 237.






Figure 7-47 Verification of the iSCSI initiator creation

## LUN Mapping

Because we explained the process of mapping the LUNs in 7.1.5, “Mapping LUNs to initiator groups” on page 219, we do not repeat those steps here.

Both LUNs are mapped to the initiator group, lochnese, with LUN ID 0 for LUNArchCam2 and LUN ID 1 for LUNClipCam2.

## 7.4 Configuring Linux for the iSCSI environment

To configure Linux for the iSCSI environment:

At this point, we do not have any disks listed, as shown in Figure 7-48 on page 238, with the `ls -l /dev/disk/by-id/*` command.

```
lochnese:~ # ls -l /dev/disk/by-id/*
total 0
drwxr-xr-x  2 root root 25 Sep 17 11:28 .
drwxr-xr-x  3 root root 24 Sep 17 11:46 ..
lrwxrwxrwx  1 root root 12 Sep 17 11:28 DVD_GCC-4480 -> ../../../../hdc
```

*Figure 7-48 List of the LUNs at the beginning*

1. Restart the iSCSI service to check for the current status of the disks, as shown in Figure 7-49, using the command **hereunder**, after mapping is defined at the N series level.

```
/etc/init.d/iscsi restart
```

```
lochnese:~ # /etc/init.d/iscsi restart
Stopping iSCSI: sync umount sync iscsid
done
Starting iSCSI: iscsi iscsid fsck/mount
done
```

*Figure 7-49 iSCSI service restart for LUN discovery*

2. Verify that the devices were discovered, as shown in Figure 7-50 on page 239 and Figure 7-51 on page 240 with the output of the commands:

```
ls -l /dev/disk/by-id/*
```

```
fdisk -l
```

There are two new devices:

- iscsi-iqn.1986-03.com.ibm:sn.101184428-0 associated to /dev/sda
- iscsi-iqn.1986-03.com.ibm:sn.101184428-1 associated to /dev/sdb

```

lochnese:~ # ls -l /dev/disk/by-id/*
lrwxrwxrwx 1 root root 9 Sep 17 11:59
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-0 ->
.././sda
lrwxrwxrwx 1 root root 9 Sep 17 11:59
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-0-generic ->
.././sg0
lrwxrwxrwx 1 root root 9 Sep 17 11:59
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-1 ->
.././sdb
lrwxrwxrwx 1 root root 9 Sep 17 11:59
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-1-generic ->
.././sg1

/dev/disk/by-id/HL-DT-ST_RW:
total 4
drwxr-xr-x 2 root root 25 Sep 17 11:28 .
drwxr-xr-x 3 root root 4096 Sep 17 11:59 ..
lrwxrwxrwx 1 root root 12 Sep 17 11:28 DVD_GCC-4480 ->
../././hdc

```

Figure 7-50 Discovery of the new devices - part 1

```
lochnese:~ # fdisk -l
```

Disk /dev/hda: 80.0 GB, 80032038912 bytes  
 255 heads, 63 sectors/track, 9730 cylinders  
 Units = cylinders of 16065 \* 512 = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1		1	3813	30619890	7	HPFS/NTFS
/dev/hda2		3814	4335	4192965	82	Linux swap
/dev/hda3	*	4336	9729	43327305	83	Linux

Disk /dev/sda: 15.0 GB, 15033434112 bytes  
 64 heads, 32 sectors/track, 14337 cylinders  
 Units = cylinders of 2048 \* 512 = 1048576 bytes

Disk /dev/sda doesn't contain a valid partition table

Disk /dev/sdb: 1073 MB, 1073741824 bytes  
 34 heads, 61 sectors/track, 1011 cylinders  
 Units = cylinders of 2074 \* 512 = 1061888 bytes

Disk /dev/sdb doesn't contain a valid partition table

Figure 7-51 Discovery of the new devices - part 2

3. Create partitions on the devices /dev/sda and /dev/sdb using **fdisk**, as shown in Figure 7-52 on page 241:
  - a. Type the command **fdisk /dev/sda**.
  - b. Type **n** to create a new partition.
  - c. Choose to create a primary partition.
  - d. Choose 1 as the partition number.
  - e. Create a partition from the first cylinder to the last cylinder.
  - f. Type **w** to write your changes.

Repeat these steps to create the partition /dev/sdb1 for the device /dev/sdb with the command **fdisk /dev/sdb**.

```

lochnese:~ # fdisk /dev/sda
Device contains neither a valid DOS partition table, nor Sun, SGI
or OSF disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

The number of cylinders for this disk is set to 14337.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Warning: invalid flag 0x0000 of partition table 4 will be corrected
by w(rite)

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-14337, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-14337, default
14337):
Using default value 14337

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```

*Figure 7-52 Creation of the partition*

4. Create a file system of xfs type, which is required for DVS installation in the new partitions. Use the following commands, as shown in Figure 7-53 on page 242:

```

mkfs -t xfs /dev/sda1
mkfs -t xfs /dev/sdb1

```

```

lochnese:~ # mkfs -t xfs /dev/sda1
meta-data=/dev/sda1          isize=256    agcount=16,
agsize=229391 blks
           =                  sectsz=512
data      =                  bsize=4096    blocks=3670256,
imaxpct=25
           =                  sunit=0      swidth=0 blks,
unwritten=1
naming    =version 2         bsize=4096
log        =internal log     bsize=4096    blocks=2560,
version=1
           =                  sectsz=512    sunit=0 blks
realtime   =none             extsz=65536   blocks=0, rtextents=0

lochnese:~ # mkfs -t xfs /dev/sdb1
meta-data=/dev/sdb1          isize=256    agcount=8,
agsize=32761 blks
           =                  sectsz=512
data      =                  bsize=4096    blocks=262088,
imaxpct=25
           =                  sunit=0      swidth=0 blks,
unwritten=1
naming    =version 2         bsize=4096
log        =internal log     bsize=4096    blocks=1200,
version=1
           =                  sectsz=512    sunit=0 blks
realtime   =none             extsz=65536   blocks=0, rtextents=0

```

*Figure 7-53 Creation of the File System*

5. Put the available disks at the Linux level by performing the following actions:

a. Create the following directories:

```

mkdir /Archives
mkdir /Clips

```

b. Add the following lines in /etc/fstab:

```

/dev/sda1 /Archives xfs defaults 0 0
/dev/sdb1 /Clips xfs defaults 0 0

```

c. Mount the file systems with the **mount -a** command.

Following these actions, LUNs are available, as shown in Figure 7-54 on page 243, with the output of the **df** command.

```
lochnese:~ # mount -a
lochnese:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda3              43306144    2631716  40674428   7% /
tmpfs                  1876552         12   1876540   1% /dev/shm
/dev/sda1             14670784      272  14670512   1% /Archives
/dev/sdb1             1043552       144   1043408   1% /Clips
```

Figure 7-54 List of File System available

## 7.5 Installing the Cisco Video Surveillance software

To install the Cisco Video Surveillance software:

1. Verify that the date and time zone are correctly set.
2. Download and install the Java™ SDK 1.4.2 package, which is at:  
<http://java.sun.com/j2se/1.4.2/download.html>
3. Install the Cisco VSM software by first unzipping it, as shown in Figure 7-55.
4. Install the different RPMs, as shown in Figure 7-56 on page 244.

During the rest of this book, we use CMS to designate the Cisco VSM Platform.

```
tonga:~/installed_software # unzip
BroadWare-BMS-4.8.0-4-sles9-sp3.zip
Archive:  BroadWare-BMS-4.8.0-4-sles9-sp3.zip
  inflating: BroadWare_Base-4.8.0-4d-sles9-sp3.i586.rpm
  inflating: BroadWare_BWT-4.8.0-4d-sles9-sp3.i586.rpm
  inflating: BroadWare_Docs-4.8.0-4d-sles9-sp3.i586.rpm
  inflating: BroadWare_Drivers-4.8.0-4d-sles9-sp3.i586.rpm
  inflating: BroadWare_Server-4.8.0-4d-sles9-sp3.i586.rpm
  inflating: BroadWare_Tools-1.1-1.i586.rpm
  inflating: BMS_4_8.pdf
  inflating: BMS-Install-Server-4.8.txt
  inflating: BMS-Upgrade-to-4.8.txt
  inflating: BMS-Whats-New-4.8.txt
```

Figure 7-55 Unzip of the CMS file

```
rpm -ivh BroadWare_Base-4.8.0-4d-sles9-sp3.i586.rpm
rpm -ivh BroadWare_BWT-4.8.0-4d-sles9-sp3.i586.rpm
rpm -ivh BroadWare_Server-4.8.0-4d-sles9-sp3.i586.rpm
rpm -ivh BroadWare_Docs-4.8.0-4d-sles9-sp3.i586.rpm
rpm -ivh BroadWare_Drivers-4.8.0-4d-sles9-sp3.i586.rpm
rpm -ivh BroadWare_Tools-1.1-1.i586.rpm
```

Figure 7-56 Installation of the RPMs

5. Connect to the Cisco management console with a browser window, as shown in Figure 7-57, at the following address:

<http://IP address of the CMS server/broadware.html>

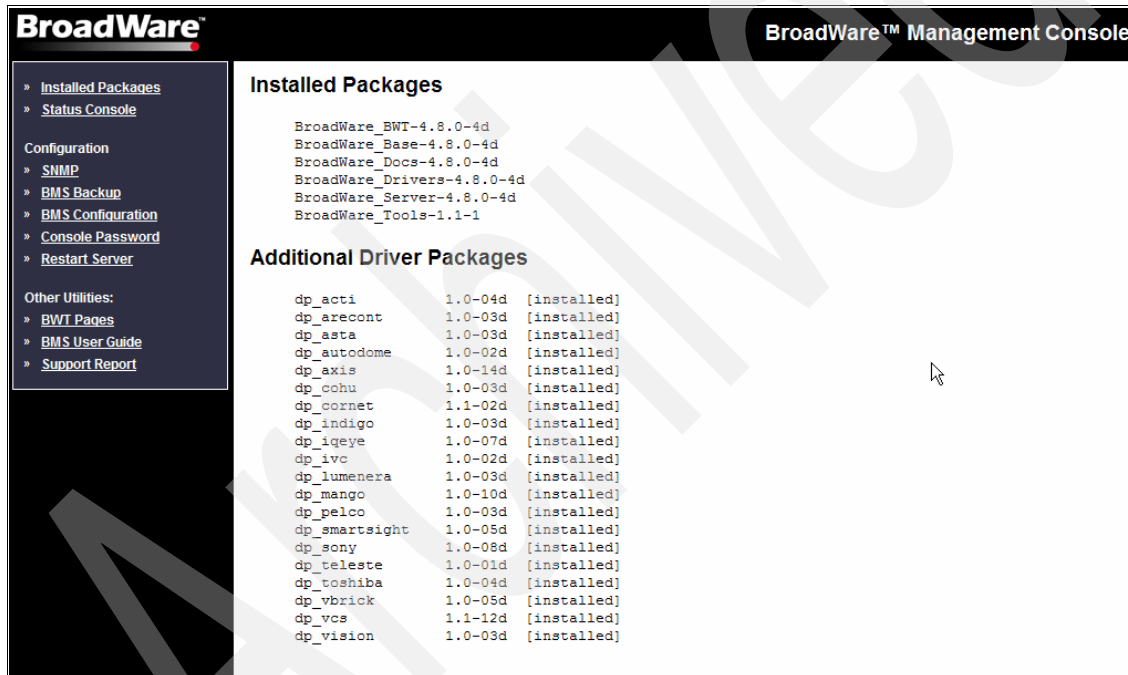


Figure 7-57 Cisco Management Console

6. In the configuration file, manually define which devices and repositories to use to store the archives and the clips. Parameters are in the file:

`/usr/BWhttpd/conf/system.cfg`



We defined the following lines in `/usr/BWhttpd/conf/system.cfg`, as shown in Figure 7-58 or Figure 7-59, depending on the server we were working on:

- Figure 7-58 is for the Fibre Channel server with multipath devices
- Figure 7-59 is for the iSCSI server

**Note:** The Cisco documentation has explanations about how to use the graphical interface to configure the devices to be used. We recommend that you manually define the devices to be used because the server does not always automatically discover the devices.

```
REPOS:/Archives
CLIP_REPOS:/Clips
EVENT_REPOS:/Clips
BWM_REPOS:/Clips
PARTITION:/dev/dm-1
PARTITION:/dev/dm-0
```

*Figure 7-58 CMS repositories and devices definition for FCP server*

```
REPOS:/Archives
CLIP_REPOS:/Clips
EVENT_REPOS:/Clips
BWM_REPOS:/Clips
PARTITION:/dev/sda1
PARTITION:/dev/sdb1
```

*Figure 7-59 CMS repositories and devices definition for iSCSI server*

7. CMS files are written by user nobody of the group nobody. Change the permissions of the devices and of the repositories to put the adapted authorizations. For devices, we used the `/etc/udev/udev.permissions` file:

- For the FCP server, add the following line:  
`dm*:root:root:777`
- For the iSCSI server, modify the following lines from:  
`sda*:root:root:660`  
`sdb*:root:root:660`  
to  
`sda*:root:root:777`  
`sdb*:root:root:777`

8. For the good authorizations to be applied to directories, change the permissions with the following commands:  

```
chmod -R 777 /Archives  
chmod -R 777 /Clips
```
9. Add these same two lines in the `/etc/rc.d/boot.local` file for these changes to remain after reboot.
10. Reboot the server for the configuration to be applied.

## 7.6 Installing the camera

To install the camera:

1. Configure the IP address of the camera to include it in your network by downloading an AXIS software, IP Utility on a Windows server that is connected to your network.
2. After the IP address is configured, verify the configuration by connecting with a Web Browser to:  
`http://ip-address-of-the-camera/`
3. Accept the installation of some additional AXIS modules to be able to see the video stream. For more information, refer to the documentation that is provided with the camera.

Figure 7-60 shows an example of the camera.



Figure 7-60 AXIS 207 camera

## DVS operations with N series

In this chapter, we describe how to use and configure the Cisco Surveillance Software to perform data capture, configure a proxy, archive, manage clips, migrate data, and expire data.

### 8.1 Capturing data

In this section, we describe how to capture data stream of the IP camera connected to the Cisco VSM Platform (CMS) hosts and the N series. The IP camera generates a stream, and we use the CMS to capture this stream and store it on N series, which the CMS uses a proxy to do. A proxy runs on a specific device that acts as a source for an encoder or an IP camera, which allows multiple other sources to view and record a single encoder or IP camera source. At least one proxy is required per a video source and supports other sources, such as clients, child proxies, and hardcovers.

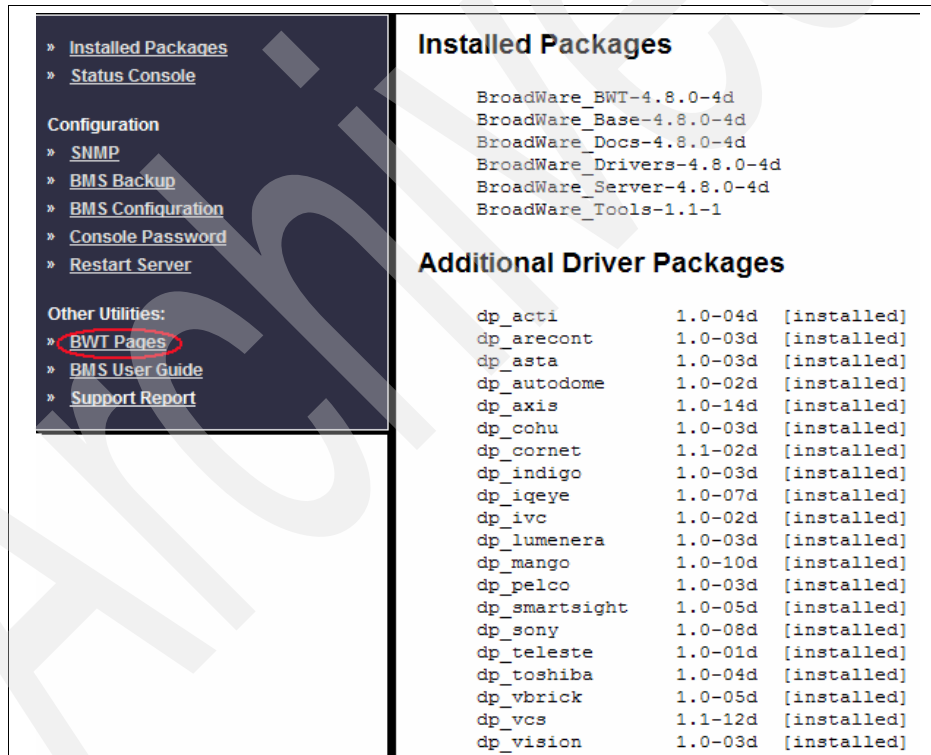
Child proxies run on local or other proxies and have the same resolution, quality, and media type of its host. They can also have a lower frame rate for MJPEG. CMS supports nine types of media proxies.

## 8.1.1 Starting the video proxy

The **start proxy** command starts a proxy for a MJPEG, MPEG, or video source. The source can be the original encoding device or another CMS proxy. Parent-child proxies can be nested indefinitely as resources permit.

To start the video proxy:

1. To access your CMS open your browser, and type the following Web address:  
**`http://ip_address/broadware.html`** (ip\_address is the IP address of the server where CMS is installed)
2. On the management console page, Figure 8-1, in the left menu, click **BWT Pages**.



The screenshot displays the CMS main menu page. On the left is a dark navigation menu with the following items: **Installed Packages**, **Status Console**, **Configuration** (with sub-items: **SNMP**, **BMS Backup**, **BMS Configuration**, **Console Password**, **Restart Server**), and **Other Utilities:** (with sub-items: **BWT Pages**, **BMS User Guide**, **Support Report**). The **BWT Pages** item is circled in red. The main content area on the right is titled **Installed Packages** and lists the following: **BroadWare\_BWT-4.8.0-4d**, **BroadWare\_Base-4.8.0-4d**, **BroadWare\_Docs-4.8.0-4d**, **BroadWare\_Drivers-4.8.0-4d**, **BroadWare\_Server-4.8.0-4d**, and **BroadWare\_Tools-1.1-1**. Below this is a section titled **Additional Driver Packages** containing a table of installed drivers.

Additional Driver Packages		
dp_acti	1.0-04d	[installed]
dp_arecont	1.0-03d	[installed]
dp_asta	1.0-03d	[installed]
dp_autodome	1.0-02d	[installed]
dp_axis	1.0-14d	[installed]
dp_cohu	1.0-03d	[installed]
dp_cornet	1.1-02d	[installed]
dp_indigo	1.0-03d	[installed]
dp_iqeye	1.0-07d	[installed]
dp_ivc	1.0-02d	[installed]
dp_lumenera	1.0-03d	[installed]
dp_mango	1.0-10d	[installed]
dp_pelco	1.0-03d	[installed]
dp_smartsight	1.0-05d	[installed]
dp_sony	1.0-08d	[installed]
dp_teleste	1.0-01d	[installed]
dp_toshiba	1.0-04d	[installed]
dp_vbrick	1.0-05d	[installed]
dp_vcs	1.1-12d	[installed]
dp_vision	1.0-03d	[installed]

Figure 8-1 CMS main menu page

3. You are redirected to a Media Platform page, Figure 8-2. Click **Proxy Commands** to open the Proxy Commands menu.

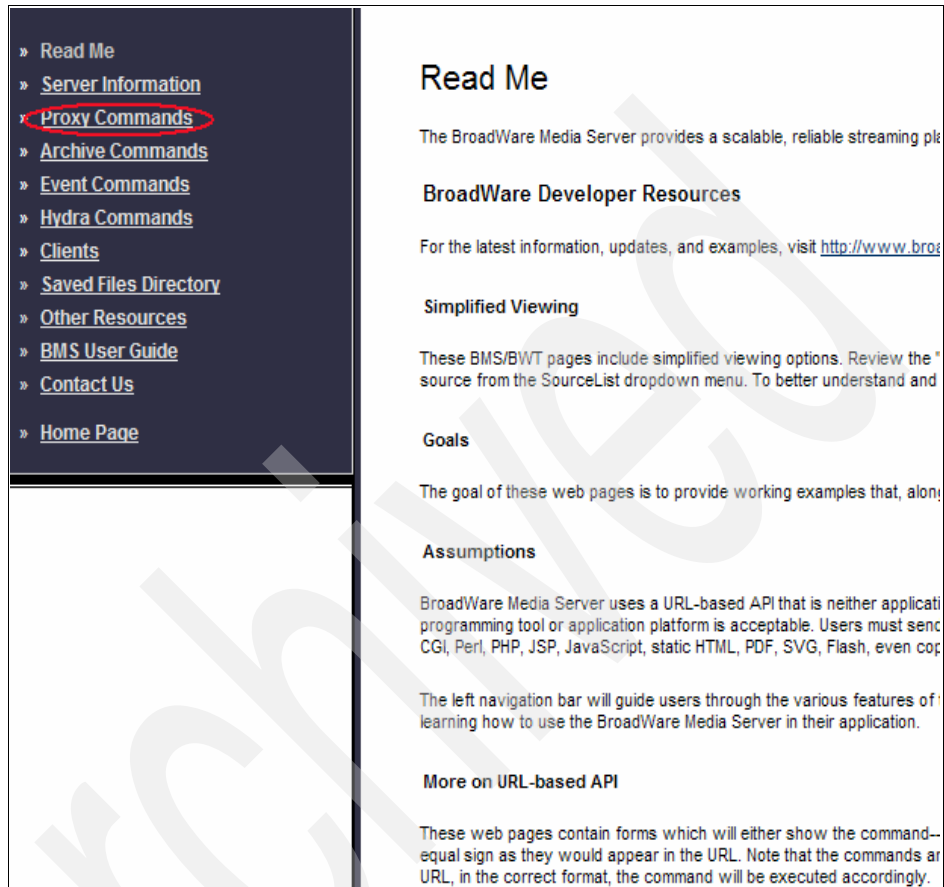


Figure 8-2 VSM

4. On the Proxy Commands page, Figure 8-3 on page 250, you can choose to start, update, list, view or stop an existing proxy. To create a proxy, from the navigation menu, click **Start Video Proxy** or under the Proxy Commands section, go to the Start Proxy section, and click **Video**.

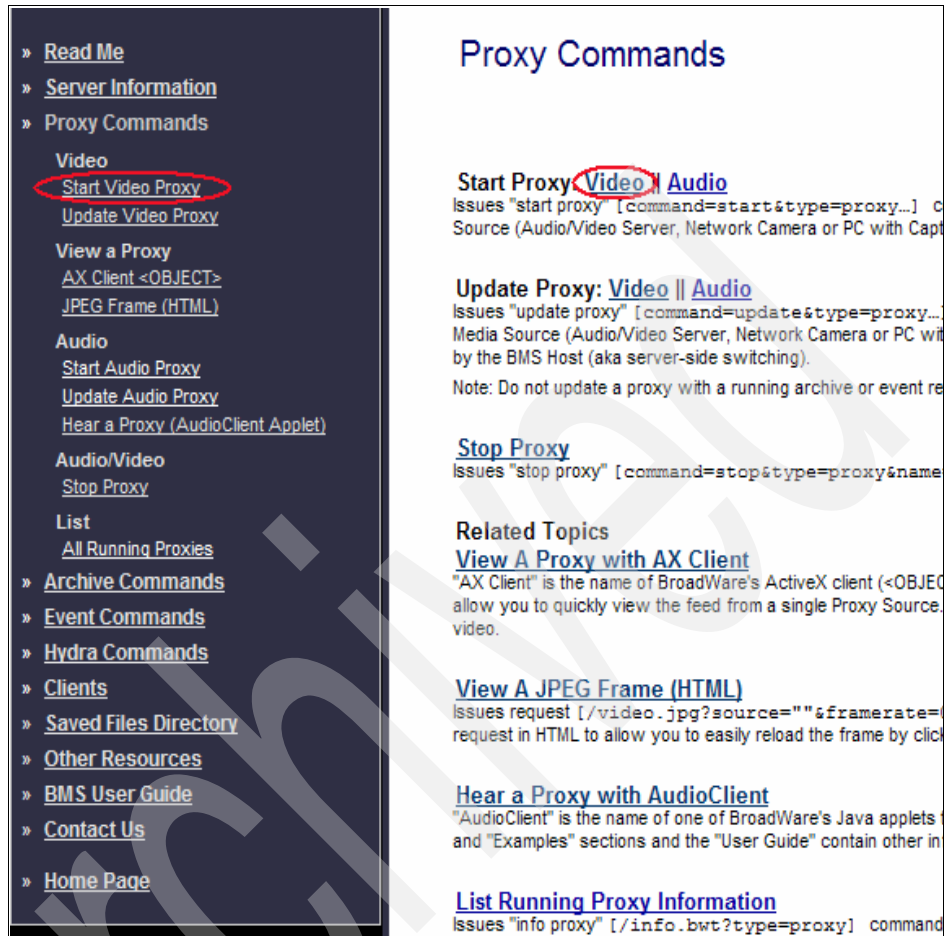


Figure 8-3 Proxy Commands

The Start Video Proxy page, Figure 8-4 on page 251, is displayed.

Start Video Proxy

9.43.86.85☐ Use 24h

change BMS host

name=REQUIRED

source=REQUIREDRunning Proxies (for child-proxies)

srcype=REQUIRED

username=(e.g. root) for supported devices.

password=(e.g. root) for supported devices.

mediatype=REQUIRED

udp=optional0 (TCP)

multicast=optional(address)

format=ntsc

resolution=cifThis parameter replaces width and height

width=optionalOver-ride resolution and format

height=optional

quality=50 (Medium)

framerate=optional5

Choose from typical framerates for JPEG feeds.

bitrate=optional56 (MPEG-4)

Choose from typical bitrates for MPEG feeds.

config=optional

execute commandshow commandreset form

Figure 8-4 Start Video Proxy

5. Populate the REQUIRED fields to start the video proxy. Table 8-1 provides a description of the fields that you are required to complete on the Start Video Proxy page.

Table 8-1 Start Video Proxy page required fields

Field	Description
host	Format: [host name.domain   IP address] Web address of host where CMS is running. CMS runs on port 80 by default. Specify additional ports here (for example, cms_host.broadware.com:8080).
command	Reserved value: [start]
type	Reserved value: [proxy]
name	Character class: [0-9   A-Z   a-z   _   -]; Reserved: (-1); Length: (1-64) Name for this proxy. This is the value that would be used as a source for a child proxy. Each proxy must have a unique name on a given CMS host.

Field	Description
source	<ul style="list-style-type: none"> <li>► Format: [#@address:port]</li> <li>► # (required): The source input number. This is the camera input number on the source.</li> <li>► address (optional): IP address/host name.domain of the source, if no address given, the address defaults to localhost.</li> <li>► port (optional): Port number, if no port is given the port defaults to 80.</li> </ul> <p>For SmartSight devices -</p> <ul style="list-style-type: none"> <li>► Format: [#_#@address:port]</li> <li>► #_# (required): The device input number on the device and the configuration number for the input.</li> <li>► address (optional): IP address/host name.domain of the device, if no address given, the address defaults to localhost.</li> <li>► port (optional): Port number, if no port is given the port defaults to 80.</li> </ul>
mediatype	<p>Reserved values: Supported Devices Codec used to encode the media.</p> <p>Example: jpeg, mpeg2-v (video and audio), mpeg2-e (video elementary), mpeg2-t (video transport), mpeg4-v (video only), and mpeg4-v.h (video only).</p>
srctype	<p>Reserved values: Supported Devices (for example, CMS proxy, video server, network camera, etc.) for this proxy.</p>

Table 8-2 on page 253 provides descriptions for the fields on the Start Video Proxy page that are optional.

**Fields that affect the storage size of N series:** The quality, framerate, width, height, and bitrate fields affect the storage size of N series. To obtain more information about sizing, go to 6.2, “Sizing N series storage to DVS” on page 129.



Table 8-2 Start Video Proxy page optional fields

Field	Description
quality	<p>Range: [1-100](default=50) Quality of the MJPEG feed, where 100 is the highest possible quality.</p> <p>If the quality parameter is set to less than 50, the framerate has priority and the requested framerate is as per the proxy_axis_mpeg4.xml file. If the quality parameter is set greater than 50, the generated image quality has priority (while maintaining the bitrate) and lower framerates are returned.</p> <p>Any number between 1 and 49 indicates the same priority for the framerate. Any number between 50 and 99 indicates the same priority for the image quality. The higher the bitrate, the higher the image quality.</p> <p>A child proxy must have the same quality value that the parent proxy has.</p>
framerate	<p>Range: [0.001-30](5) Maximum number of MJPEG frames per second transmitted.</p> <p>A child proxy cannot have a higher framerate value than the parent proxy has.</p>
width	<p>Range: [160-720](320) Width of the video feed in pixels.</p> <p>A child proxy must have the same width value that the parent proxy has.</p>
height	<p>Range: [112-576](240) Height of the video feed in pixels.</p> <p>A child proxy must have the same height value that the parent proxy has.</p>
bitrate	<p>Range: [28.8-5000](640) Kilobytes per second transmitted for MPEG feed.</p> <p>If the quality parameter is set to less than 50, the framerate has priority and the requested framerate is as per the proxy_axis_mpeg4.xml file. If the quality parameter is set to greater than 50, the generated image quality has priority (while maintaining the bitrate) and lower framerates are returned.</p> <p>Any number between 1 and 49 indicates the same priority for the framerate. Any number between 50 and 99 indicates the same priority for the image quality. The higher the bitrate, the higher the image quality.</p>
username	<p>Format: [user name] Administration user name for encoding device.</p> <p>If a username and password are enabled for a device or to do Camera Controls and Events Setup, this field is required.</p>

Field	Description
password	Format: [password] Administration password for encoding device.  If a username and password are enabled for a device or to do Camera Controls and Events Setup, this field is required.
resolution	Reserved values: [qcif   cif   2cif   4cif](cif) Resolution is used to start proxy using look-up values for width and height of video.
format	Reserved values: [ntsc   pal] (ntsc) Video standard name for the format that is used to start proxy using look-up values for width and height of video.  The width and height parameters take precedence over format.
udp	Reserved values: [0   1](0) Protocol for MPEG-4 multicast streams. Use 1 to turn on.  To start a multicast stream an address must be provided.
multicast	Address: [IP address   host name] To start or join a multicast stream, the address of the device must be included. When upgrading IMC, replace the previous CLASSID with the new CLASSID.
srctype	Reserved values: Supported Devices Type of device for this proxy.

**Start Video Proxy**

9.43.86.85 ☐ Use 24h

[change BMS host](#)

name= DVS0001

source= 1@9.43.86.100 Running Proxies (for child-proxies) ▼

srctype= axis207 [ AXIS 207 Network Camera ] ▼

username= root (e.g. root) for supported devices.

password= ••••• (e.g. root) for supported devices.

mediatype= jpeg ▼

udp= optional 0 (TCP) ▼

multicast= optional (address)

format= ntsc ▼

resolution= cif ▼ This parameter replaces width and height

width= optional Over-ride resolution and format ▼

height= optional

quality= 50 (Medium) ▼

framerate= optional 5 ▼

Choose from typical framerates for JPEG feeds.

bitrate= optional 56 (MPEG-4) ▼

Choose from typical bitrates for MPEG feeds.

config= optional

[execute command](#) [show command](#) [reset form](#)

Figure 8-5 Start Video Proxy

We populated the fields in Figure 8-5 as follows:

- Name= proxy name
- Source=ip address of the camera
- srctype= model of the camera
- Username= camera user ID
- Password= camera password
- Mediatype=jpeg

The other fields in Figure 8-5 were populated by default, which is the basic configuration that is required if you want to have a different option.

To finish the process click the **execute command** button, and you should receive the message proxy started successfully, as shown in Figure 8-6 on page 256.



Figure 8-6 Started Successfully

### 8.1.2 Updating a proxy

Use the following steps to update an existing proxy with different parameter values:

**Note:** If an archive is running against the proxy, updating the proxy source causes the archive to be unplayable. In this case, stop the archive first, update the proxy, and then start a new archive.

1. To access the Update Video Proxy page, in the left navigation, click **update video proxy**, as shown in Figure 8-7 on page 257. A similar page of the start video proxy is displayed, but you do not have the option to change the media type, so if you have a running proxy with JPEG media type and you want to change to MPEG media type, you need to create a new proxy.

**Update Video Proxy** 9.43.86.85 [change](#)

name= Running Proxies-REQUIRED

source= REQUIRED Running Proxies (for child-proxies)

srctype= REQUIRED

username= (e.g. root) for supported devices.

password= (e.g. root) for supported devices.

mediatype= Media Type is inherited from original proxy.

udp= optional 0 (TCP)

multicast= optional (address)

format= ntsc

resolution= cif This parameter replaces width and height

width= optional Over-ride resolution and format

height= optional

quality= 50 (Medium)

framerate= optional 5

Choose from typical framerates for JPEG feeds.

bitrate= optional 56 (MPEG-4)

Choose from typical bitrates for MPEG feeds.

config= optional

[execute command](#) [show command](#) [reset form](#)

Figure 8-7 Update Video Proxy

- To update the video proxy, from the name field, shown in Figure 8-8, choose an existing proxy. Update the proxy with the new values, and click the **execute command** button to save the changes.

**Update Recording Frame Rate** 9.43.86.85 ☐ Use 24H [change BMS host](#)

name= DVS\_Arcvhive

framerate= 7

NOTE: New frame rate cannot exceed [motion JPEG source proxy](#) frame rate.

[execute command](#) [show command](#) [reset form](#)

Figure 8-8 Update Video Proxy

A message is displayed that lets you know that the proxy reset successfully, as shown in Figure 8-9 on page 258.

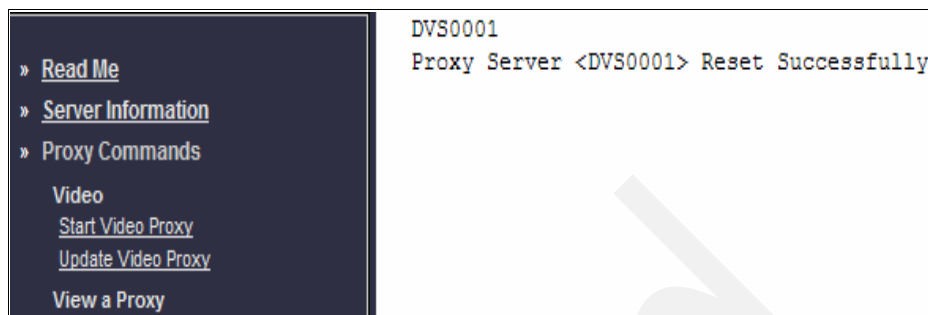


Figure 8-9 Reset successfully

### 8.1.3 Stopping a proxy

The **stop proxy** command stops a proxy and any archives that are running against it on the CMS host.

**Note:** When you stop a proxy, all of the archives that are connected to this proxy will stop also.

To stop a proxy:

1. To access the stop proxy page, on the left navigation menu click, **Stop Proxy**, as shown in Figure 8-10.



Figure 8-10 Stop Proxy

2. From the name field, Figure 8-11, choose the name of the proxy that you want to stop, and then click the **execute command** button to stop the selected proxy.



Figure 8-11 Stop Proxy

A message is displayed stating that the proxy server stopped successfully, as shown in Figure 8-12.



Figure 8-12 Stopped successfully

### 8.1.4 Viewing a proxy

If you already have a proxy running or if you create a new proxy, you view this proxy using the View a proxy option.

To view a proxy:

1. On the left navigation menu, click AX Client <OBJECT>, and the View a Video Proxy with AX Client page is displayed. The proxy view window is black until you select a proxy from the Running Proxies pull-down menu, as shown in Figure 8-13.

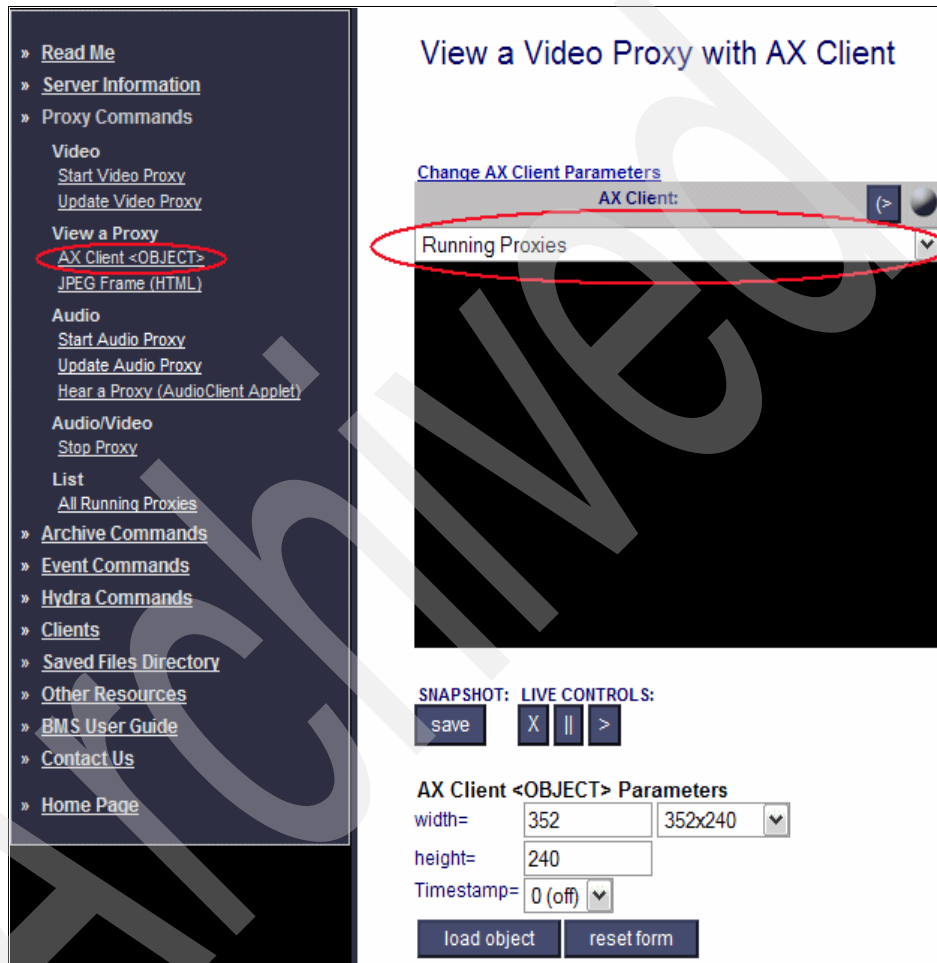


Figure 8-13 View a Video Proxy

After you select the proxy that you want to view, the proxy will start to play, as shown in Figure 8-14 on page 261. You have the live controls and the option to stop, play, and to take a picture of the proxy using a snapshot.



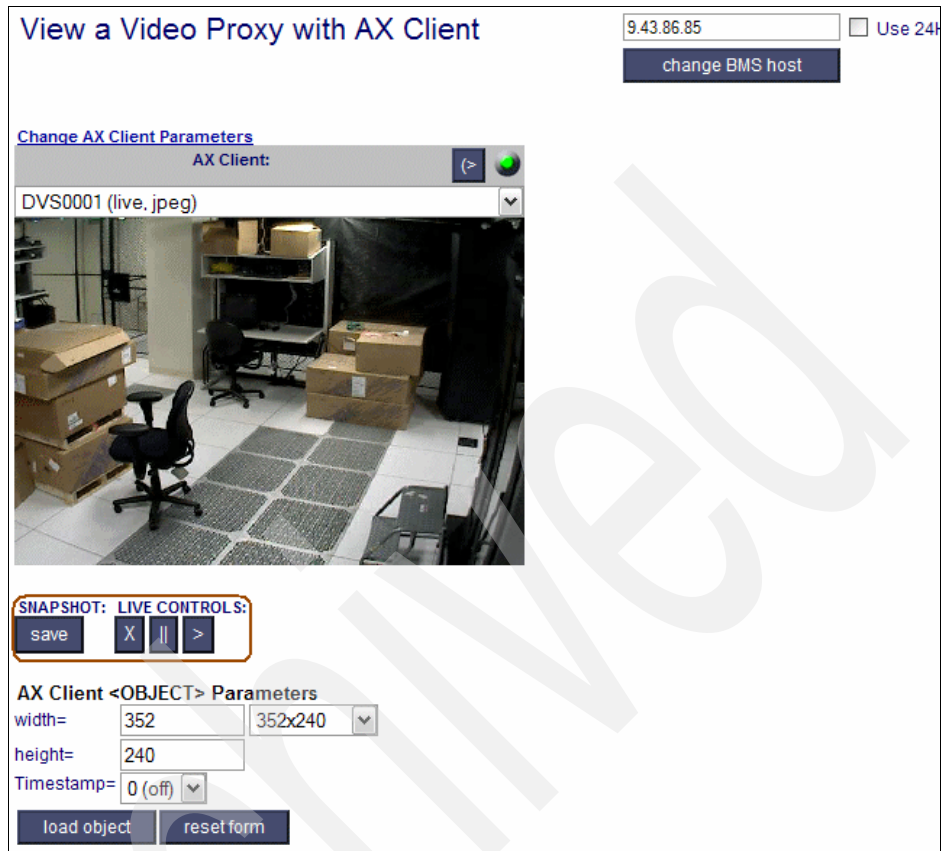


Figure 8-14 View a Video Proxy

2. You can also use the AX Client <OBJECT> Parameters in Figure 8-14 to resize the proxy view windows size, as shown in Figure 8-15 on page 262. To resize the proxy view windows, populate the fields under the AX Client<OBJECT> field, and click the **load object** button.

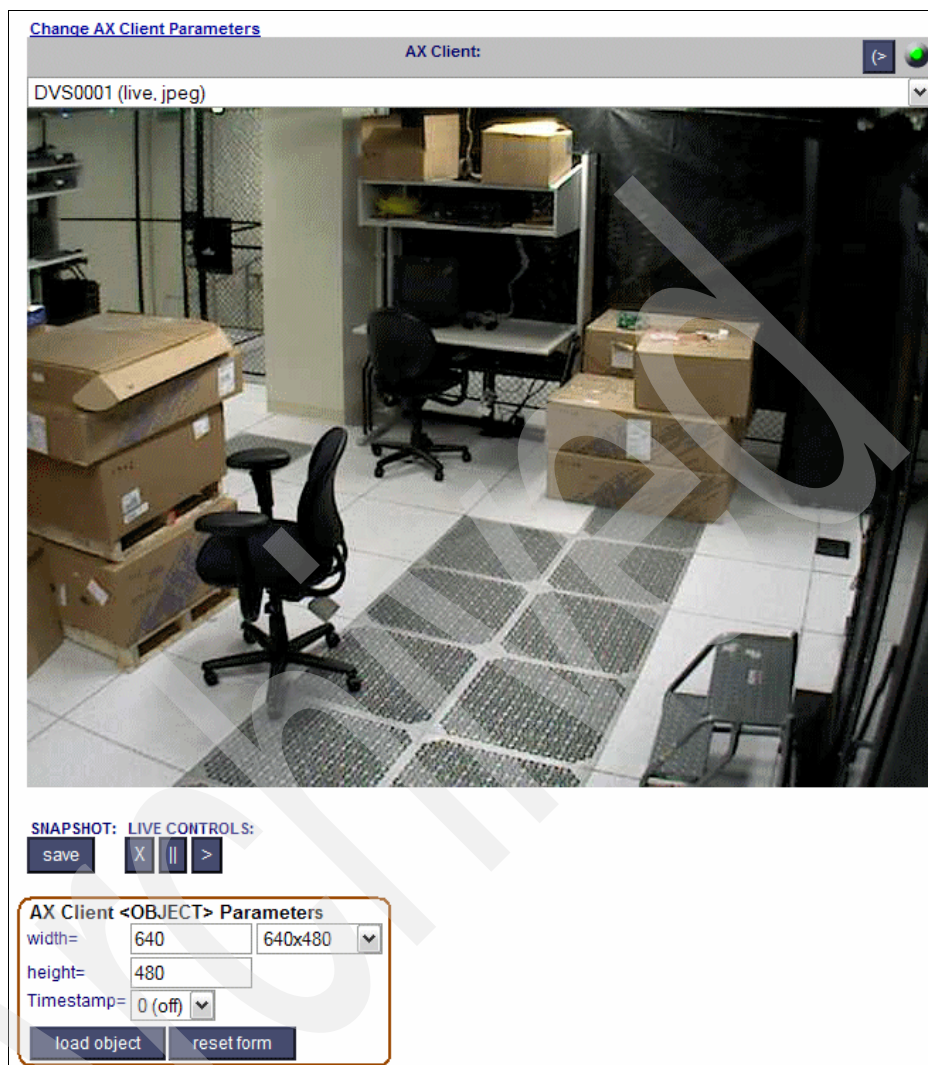


Figure 8-15 View a Video proxy options

## 8.1.5 Listing all proxies

To get a list of all of the running proxies on the CMS host, in the left navigation men, click **all running proxies**, as shown in Figure 8-16 on page 263.

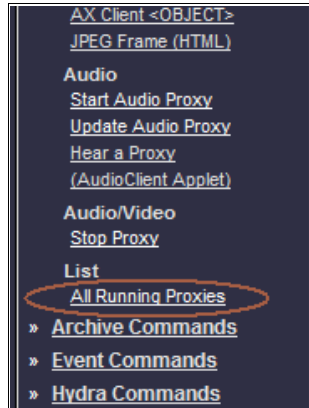


Figure 8-16 All Running Proxies

The information about all running proxies is displayed, as shown in Figure 8-17.

Proxy Server Info - Windows Internet Explorer  
 http://9.43.86.85/info.bwt?type=proxy&display=html

Proxy Server Information										
name	status	type	exec	source	mediatype	f/b-rate	quality	width	height	model
DVS0001	Running	axis207	proxy	1@9.43.86.199:80	jpeg	5.00000	50	320	240	57
DVS0002	Running	axis207	proxy	1@9.43.86.199:80	mpeg4-v	512	50	640	480	57
dvs001	Suspended	axis207	proxy	1@9.43.86.199:80	jpeg	5.00000	50	160	120	57
dvsJpeg	Suspended	axis207	proxy	1@9.43.86.199:80	jpeg	10.00000	50	640	480	57
Total: 4 Proxies										

Figure 8-17 All Running Proxy details

## 8.2 Retaining and archiving data

In this section, we discuss the archive commands. The **archive** commands permit you to record, store, and manage resources for audio and video archives. All commands that control archiving activities are URL-based. You can manually issue commands or program them after an HTTP connection is established.

## 8.2.1 The archive commands

There are two types of archives, and those archives are stored on the N series. Previously those archives needed to be sized on N Series. To obtain more information about sizing, refer to 6.2, “Sizing N series storage to DVS” on page 129:

The two types of archives are:

- ▶ Regular archives, where the archive terminates after the specified duration lapses.
- ▶ Loop archives, where the archive keeps recording until the archive is stopped. Loop archives reuse the space (first-in/first-out) that is allocated after every completion of the specified loop time.

The archive commands page lists all archive-related functions and their shortcuts:

<b>Start Archive</b>	Begins to archive the existing video/audio stream.
<b>Update Recording Frame</b>	Rate updates the frame rate of a running archive.
<b>Stop Archive</b>	Stops an existing archive process.
<b>Remove Archive</b>	Removes an existing archive.
<b>Set “Day-to-Live”</b>	Sets the number of days (from the date that the archive stops) that the archive is stored before system removal.
<b>Save Video Clip Command URL</b>	Extracts a clip from an existing archive with a starting time and an ending time.

To access in the Archive page, in the left navigation menu, click **Archive Commands**, as shown in Figure 8-16 on page 263, and all of the **archive** command options are displayed, as shown in Figure 8-18 on page 265.

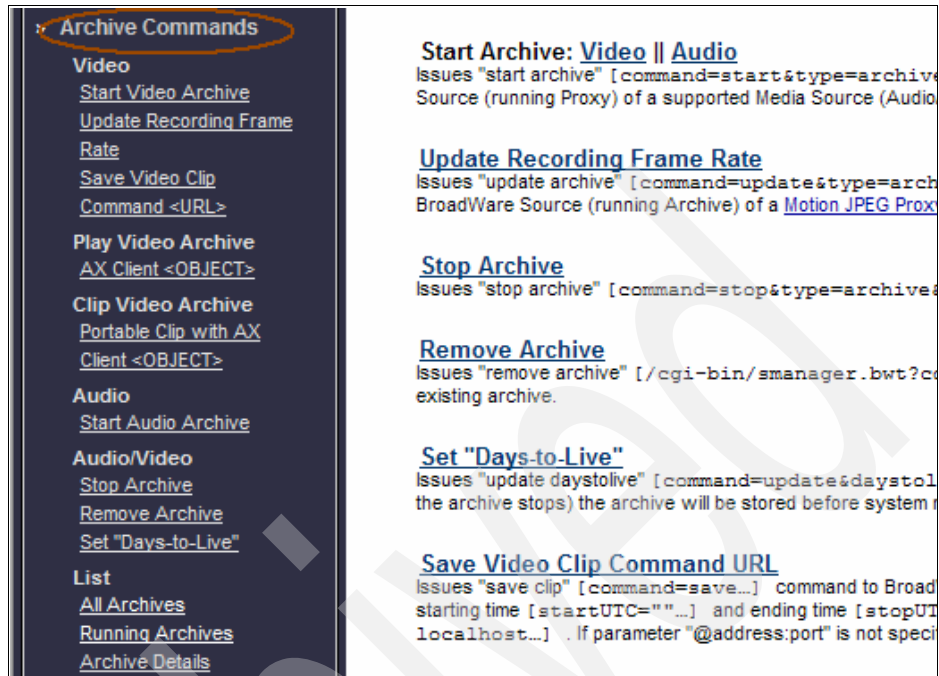


Figure 8-18 Archive Commands

## Start archive command

The **start archive** command initiates an archive against an existing proxy.

To start an archive:

1. On the left navigation menu, click **Archive commands** → **Start Video Archive** or **Start Archive video**. The Start Archive page is displayed, as shown in Figure 8-19 on page 266.

Start Video Archive

9.43.86.85 ☐ Use 24H

change BMS host

name= REQUIRED

source= Running Proxies-REQUIRED

loop= optional 0 (Off)

duration= optional seconds 

get sec »

framerate= optional 5

killproxy= optional 0 (Off)

daystolive= 0 days (Ø never expires)

repos= optional

count= optional integer (must have PS/AS configured on BMS host)

desc= optional

execute command

show command

reset form

Second Calculator

0 days

0 hours

0 minutes

0 seconds

0 total

sum

clear

Figure 8-19 Start Video Archive

- Complete the REQUIRED fields to start video archive. Table 8-3 contains descriptions for all of the required fields on the Start Video Archive page.

Table 8-3 Required fields on the Start Video Archive page

Field	Description
host	Format: [host name.domain   IP address] This is the Web address of the host where CMS is running.  CMS runs on port 80 by default. Specify additional ports in this field (for example, cms_host.broadware.com:8080).
command	Reserved value: [start]
type	Reserved value: [archive]
name	Character class: [0-9   A-Z   a-z   _   -]; Reserved: (-1); Length: (1-127) This is the name of the archive that is being recorded.  Each archive must have a unique name on the CMS host.
source	Format: [name] name (required) This is the name of the source proxy being archived on the CMS host.

Table 8-4 on page 267 contains descriptions for all of the optional fields on the Start Video Archive page.

**Fields that affect the storage size of N series:** The following fields affect the storage size of N Series: duration, desc. To obtain more information about sizing, go to 6.2, “Sizing N series storage to DVS” on page 129.

Table 8-4 Optional fields on the Start Video Archive page

Field	Description
duration	<p>Format: [integer in seconds](default=900) The total archive time in seconds.</p> <p>Note: Duration of 0 starts an archive of 900 seconds. To start an archive of maximum length, first calculate the available storage, then use "get average frame size" for an archive from a similar environment. <math>AVERAGE\ FRAME\ SIZE \times FPS \times SECONDS = AVERAGE\ REQUIRED\ STORAGE</math></p>
desc	<p>Character class: [0-9   A-Z   a-z   _   &lt;space&gt;   -](proxy source); Reserved: (-1); Length: (0-20) Brief description of the archive.</p> <p>If desc is left blank, CMS will store the name of the proxy source in the description field.</p>
daystolive	<p>Format: [integer in days](0) The number of days, from the date that the archive stops, until the archive is stored before system removal. For permanent storage set daystolive=0.</p>
killproxy	<p>Boolean values: [0   1](0) Stops the source proxy when the archive is stopped [1] or leaves the source proxy running [0].</p>
count	<p>Format: [integer] Total number of archives to start.</p>
repos	<p>Format: [repository_mount] The location where the new archive will be saved.</p> <p>Saves clips directly to the repository mount location. Only one mount is recognized. If no mount is specified, the clip repository must be specified using the "repos" field in the save clip XML API request. This repository also serves as a workspace area for remote clip generation.</p>
framerate	<p>Range: [0.001-30](proxy framerate) Maximum number of MJPEG frames per second requested from source proxy.</p> <p>The framerate cannot be higher than source proxy.</p>
loop	<p>Boolean values: [0   1](0) Record a loop [1] or regular archive [0]. A loop archive continuously records over its beginning after it reaches the end of its duration.</p>

Figure 8-20 is a sample creation of a video archive.

**Start Video Archive**

9.43.86.85 ☐ Use 24H  
 change BMS host

name= DVS\_Archive  
 source= DVS0001 (live, jpeg)   
 loop= 1 (On)   
 Choose from possible optional values.  
 duration= 1800 seconds   
 get sec »  
 framerate= optional 5   
 Choose from typical framerates for JPEG feeds.  
 killproxy= optional 0 (Off)   
 Choose from possible optional values.  
 daystolive= 7 days (Ø never expires)  
 repos= optional  
 count= optional integer (must have PS/AS configured on BMS host)  
 desc= optional

Second Calculator  
 0 days  
 0 hours  
 30 minutes  
 0 seconds  
 1800 total  
 sum clear

execute command show command reset form

Figure 8-20 Sample of Start Video Archive

- After you populate the required and necessary optional fields, click the **execute command** button to finish.

» [Read Me](#)  
 » [Server Information](#)  
 » [Proxy Commands](#)  
 » [Archive Commands](#)  
 Video  
[Start Video Archive](#)  
[Update Recording Frame Rate](#)  
[Save Video Clip Command <URL>](#)  
 Play Video Archive  
[AX Client <OBJECT>](#)  
 Clip Video Archive  
[Portable Clip with AX Client <OBJECT>](#)  
 Audio  
[Start Audio Archive](#)  
 Audio/Video  
[Stop Archive](#)  
[Remove Archive](#)  
[Set "Days-to-Live"](#)

```
DVS_Archive
1 (tonga)
DVS_Archive started successfully!
```

Figure 8-21 Video Archive Started successful



## Stop archive command

The **stop archive** command stops a running archive process. Use this command to stop a regular archive before its specified time elapses or to stop a loop archive. If you set killproxy to [1], stopping an archive also stops the source proxy.

To stop a proxy:

1. On the Stop Archive page, Figure 8-22, select an existing archive from the name= pull-down menu.



Figure 8-22 Stop Archive

2. Click the **execute command** button stop an archive, as shown in Figure 8-23 on page 270.



Figure 8-23 Select an Archive to be stopped

A message is displayed stating that the stop was successful, as shown in Figure 8-24.

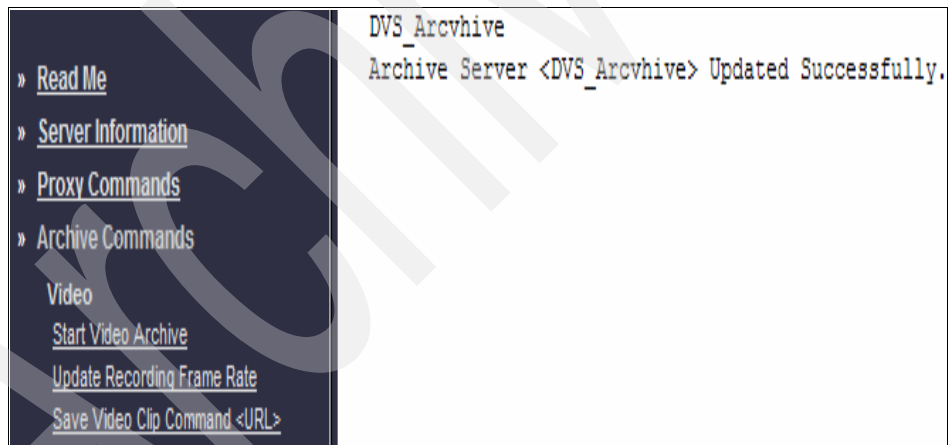


Figure 8-24 Stop an Archive successfully

## Update recording frame rate command

The **update MJPEG archive framerate** command updates the archive recording frame rate for a running MJPEG archive.

**Framerate:** Changing the framerate for an archive affects how the allocated storage gets used up. Allocated storage is based on the framerate that is specified in the framerate parameter at the time that the archive is started and allocated at that framerate for its duration. Increasing the recorded framerate might shorten the duration of the archive because the allocated storage gets used up sooner as more frames are written onto storage, so use a thin provisioning from N series to avoid this problem. For loop archives, each loop can have more storage allocated, but unused storage is not de-allocated if it has already been written to before; therefore, space might be wasted if the frequency of framerate updates is significantly less than the previous number.

To open the update Recording Frame Rate page:

1. From the left navigation menu, click **Update Recording Frame Rate**, as shown Figure 8-25.

» [Read Me](#)

» [Server Information](#)

» [Proxy Commands](#)

» [Archive Commands](#)

Video

[Start Video Archive](#)

[Update Recording Frame Rate](#)

[Save Video Clip Command <URL>](#)

Play Video Archive

[AX Client <OBJECT>](#)

Clip Video Archive

[Portable Clip with AX Client <OBJECT>](#)

## Update Recording Frame R

name=

framerate=

NOTE: New frame rate cannot exceed m

Figure 8-25 Update Recording Frame Rate

2. On the Update Recording Frame Ratepage, locate the name field pull-down menu, and choose the archive whose frame rate you want to change. In the framerate field, type the new framerate, as shown in Figure 8-26 on page 272.

Figure 8-26 Update Recording Frame Rate

3. Click the **execute command** button, and the new value is updated. A confirmation message is displayed, similar to Figure 8-27.

Figure 8-27 Update Successfully

## Remove archive command

The **remove archive** command removes a shelved archive from storage allocation. An archive cannot be deleted while it is still running. Stop the archive prior to removing it. To stop the archive, follow the procedure in “Stop archive command” on page 269.

To access the Remove Archive Page:

1. In the left navigation menu, click **Remove Archive**, as shown in Figure 8-28 on page 273.

Figure 8-28 Remove Archive

2. On the Remove Archive page, from the name field pull-down menu, choose the archive that you want to remove, and then click the **execute command** button. The selected archive is removed (Figure 8-29).

Figure 8-29 Remove Archive

## Set archive days-to-live command

The set archive days-to-live command sets the number of days for the archive to be stored after it has stopped recording, which directly affects your N series storage capacity. Shelved archives whose days-to-live date has past are removed by the system nightly at midnight. You can modify this nightly system cron job time by editing the user root's crontab file. Refer to Linux documentation for information about how to modify the root crontab file. You can also set archives for permanent storage, `daystolive=0`.

To access the Set Days-to-Live page:

1. In the left navigation menu, click **Set "Days-to-Live"**, as shown in Figure 8-30 on page 274.



Figure 8-30 Set Archive "Days-to-Live"

- From the name field pull-down menu, choose the target archive. In the daystolive field, type the number of days that you want to keep the archive when the archive is stopped, as shown in Figure 8-31. If you want to keep the archive permanently, input the value 0. Click the **execute command** button, for the setting to take effect.

Figure 8-31 Days to Live setting

## Save video clip command URL

The **save clip** command extracts a clip from an existing archive. If you do not specify a clip name, CMS generates the name. A generated clip name is comprised of the archive name from where the clip was taken and the date (in GMT) that the clip was generated, for example, an archive clip name of archive1\_05\_23\_2003\_20\_33\_15\_05 means that the source name is called archive1 and the clip was saved on May 23, 2006 at 20:33:15.05 (hours, minutes, and seconds).

To open the Save Clip page:

1. in the left navigation, click **Save Video Clip Command <URL>**, as shown in Figure 8-32.



Figure 8-32 Save Clip

2. On the Save Clip page, you must choose the source of video or audio archive from the pull-down menu, determine if you want local or remote to record this clip data, and specify the clip's start time and ending time. The time must be in UTC format. Just select the actual time on the top of the video-screen, and click the **get utc>>** button, and you will get the UTC time (refer to Figure 8-33 on page 277 and Figure 8-34 on page 278, we want to get a five minute clip in that case) You can also draw the slide button under the video-screen to export UTC time. Also, input the format of the clip. Table 8-5 contains detailed descriptions about the required fields on the Save Clip page.

Table 8-5 Required fields on the Save Clip page

Field	Descriptions
host	Format: [host name.domain   IP address] Web address of host where CMS is running.  Note: CMS runs on port 80 by default. If another port is configured, then specify it here (for example, cms_host.broadware.com:8080).
command	Reserved value: [save]
source	Format: [source_id] source_id (required): The source archive name. This is the parent archive to create a clip from.

Field	Descriptions
startutc	Format: [UTC milliseconds] Start date of the child clip in UTC milliseconds. Make sure the parent archive contains data for this date.
stoputc	Format: [UTC milliseconds] Stop date of the child clip in UTC milliseconds. Make sure the parent archive contains data for this date.
savemode	Format: [local   remote   localandremote] Indicates where to save the archive clip. local: Save clip to this CMS host. remote: If name parameter (in form of target_id@address:port) is specified, then CMS saves the clip to the specified host. localandremote: Save clip to this CMS host and a remote host.

Table 8-6 provides descriptions about the fields that you are not required to complete on the Save Clip page.

*Table 8-6 Optional fields on the Save Clip page*

Field	Description
name	Format: [target_id@address] Character class: [0-9   A-Z   a-z   _   -]; Reserved: (-1); Length: (1-127) target_id: The new archive clip name. If no name is specified, CMS generates the name. <source ID>_<month MM>_<day DD>_<year YYYY>_<hours hh-24>_<minutes mm>_<seconds ss>_<hundredths .00> (e.g. lobby_07_02_2003_15_25_52_01) address (optional): IP address/host name.domain.extension of target host to where clip is saved.
desp	Character class: [0-9   A-Z   a-z   _   <space>   -]; Reserved: (-1); Length: (0-20) Brief description for the new archive clip.  For an archive clip to have type "clip" in the archive listing (/info.bwt?type=archive), DO NOT provide a description. The description defaults to "clip".  For a BWM file clip to have type "bwm" in the archive listing (/info.bwt?type=archive), DO NOT provide a description. The description defaults to "BWM".
repos	Format: [repository_mount] Location where the new archive clip will be saved. Saves clips directly to the repository mount location. Only one mount is recognized. If no mount is specified, then the clip repository must be specified using the "repos" field in the save clip XML API request. This repository also serves as a workspace area for remote clip generation.



Field	Description
saveformat	Format: [ regular   bwm   bwx ](default=regular) Type of archive clip to generate: regular archive, bwm archive, or signed bwm archive.
key	Format: Character class: [0-9   A-Z   a-z   _   <space>   -]; Reserved: (-1); Length: (6-64) The key used to sign the bmx archive. The key is not stored in the archive.
daystolive	Format: [integer in days](0) Number of days from the date archive clip ends the archive is stored before system removal.  For permanent storage set daystolive=Ø.
notifyurl	Format: [http://<host>/handler_path] URL to send upon the successful completion or failure of a clip.

3. Scroll down to complete the rest of the options on the page, as shown in Figure 8-33.

**Save Clip**

source=

name=

savemode=  requires configuration for remote clips

startUTC=

stopUTC=

desc=

daystolive=  days (Ø never expires)

repos=

saveformat=

key=  (for bwx clips)

notifyURL=

Figure 8-33 Save Clip settings

Figure 8-34 shows an example of the full screen with the clip shown in the right side.

Figure 8-34 Save Clip settings

4. To get the UTC time for the stopUTC, click the **getutc>>** button, which is located on the right side of the field, as shown in Figure 8-35 on page 279.

Save Clip

9.43.86.85

☐ Use 24H

change BMS host

source= All Archives-REQUIRED

name= optional

savemode= REQUIRED

requires configuration for remote clips

startUTC= 1190845659000 

get utc >

stopUTC= 1190845960000 

get utc >

desc= optional

daystolive= 0 days (0 never expires)

repos= optional

saveformat= REQUIRED

key= optional (for bwx clips)

notifyURL= optional

execute command

show command

reset form

UTC Calculator

9

26

2007

03

32

40

PM

Use IMC time to populate UTC Calculator:

>

Wednesday, September 26, 2007 3:32:40 PM

Figure 8-35 Save Clip Settings

When you finish setting up the parameters, click the **execute command** button to save a clip. The CMS will display a message that the clip save is in progress, as shown in Figure 8-36. You can check it in the 'Portable Clip with AX Client <OBJECT>' page.

» [Read Me](#)

» [Server Information](#)

0

FOTO\_09\_26\_2007\_00\_50\_16\_00 save clip in progress

Figure 8-36 Save Clip in progress

## 8.2.2 List all archives command

The **list all archives** command displays information for all archives (running or shelved) on a CMS host.

Figure 8-37 on page 280 shows what to select on the Archive Commands page to access the List of All Archives page.

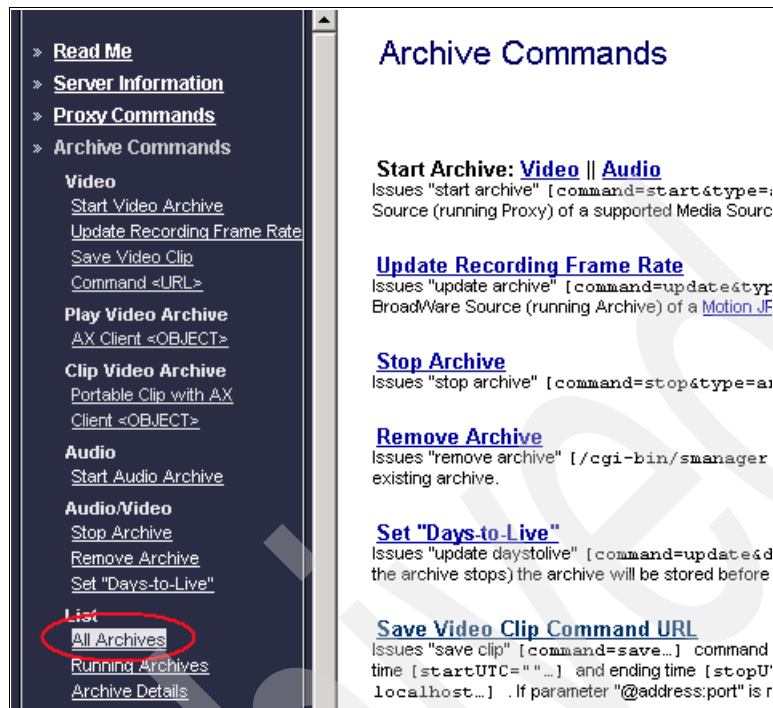


Figure 8-37 Accessing the List All Archives page

Figure 8-38 shows some listed examples. In the List All Archive page, you can see the name of the archive, its path, type, size, status, the begin time, the end time of the archive, and its expire time.

BWT_STORAGE_MANAGER: archive listing on tonga.					
ID	Directory Path	Type	Requested Size(K)	Actual Size(K)	Status
<a href="#">DVS000</a>	/Clips/DVS000	loop	72975.000000	0.000000	RUNNING
<a href="#">DVS_Archive</a>	/Archives/DVS_Archive	loop	356634.000000	353325.537000	RUNNING
<a href="#">DVS_Archive0005</a>	/Clips/DVS_Archive0005	regular	670187.000000	429718.257000	SHELVED

Figure 8-38 Example of list all archives

### 8.2.3 List all running archives command

The **list running archive** command displays information for all currently running archives. When specifying an archive name, only that archive's information is displayed.

Archive Server Information									
exec	media type	name	source	frame/bit rate	quality	width	height	loop	duration
archiver	jpeg	DVS_Archive	DVS0001	5.00000	50	640	480	1	1920
archiver	jpeg	DVS_Archive0005	DVSjpeg	10.00000	50	640	480	0	1800

Figure 8-39 List Running Archives

### 8.2.4 Archive details command

The **get archive file information** command displays the media master file contents for an archive.

Figure 8-39 shows what to select on the Archive Commands page to access the Archive Details page.

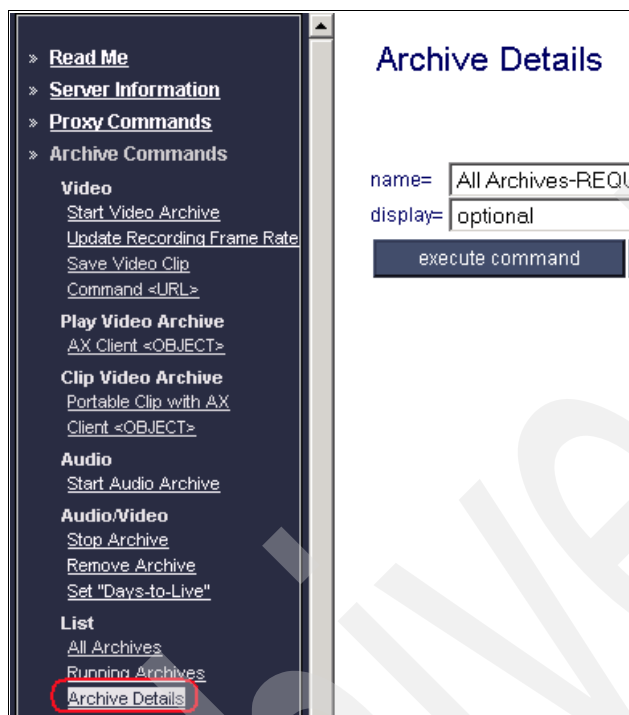


Figure 8-40 List Archive Details page

On the Archive Details page, Figure 8-41 on page 283, you can get detailed information about an archive:

1. From the name pull-down list, choose the archive of which you want to see details, and choose the display format.
2. Click the **execute command** button. You will get the detail information about the archive, such as video resolution information, framerate information, and so on.

Figure 8-41 on page 283 shows an archive's detail information. You can also access the archive's detail information from the List All Archives page too by clicking the list name of the archive.

BWT_STORAGE_MANAGER: meta data of archive /Archives/FOTO@tonga.	
Attribute Name	Attribute Value
MediaMaster ID	/Archives/FOTO/media_master
Version	4.8
Description	DVS0001
Type	1
Media Type	jpeg
Video Width	640
Video Height	480
Video Quality	50
Framerate	5.000000
Bit rate	600
Duration	900
First Offset	0
Last Offset	197
Loopsize	0
Avg Frame Size	35939.000000
# of Loops	0
Days To Live	0
Start Time	Tue Sep 25 17:50:15 2007
Stop Time	Tue Sep 25 17:50:55 2007
<a href="#">ShowArchive</a>	Wed Sep 26 11:36:46 2007

Figure 8-41 Detail information of an archive

## 8.3 Clip management

The clip function lets you extract a part from an existing archive, and then you can add the clip to the storage repository. You can save the clip to a local host or a remote host. With the clip, you can extract the information you want from the video archives, and you can also allow client's to remotely access your clips.

### 8.3.1 Creating a clip

We previously explained to you how to save a clip from an existing archive in "Save video clip command URL" on page 274. In this section, we tell you step-by-step how to create a clip and how to use it:

1. From the Save Clip page, from the source pull-down menu, choose the original archive, as shown in Figure 8-42 on page 284.

Figure 8-42 Choose the archive that you want to do clip action

2. Name the archive; otherwise, the system automatically gives it a name. In Figure 8-43, we give it a name - OptionalName\_DVS001. Choose the savemode: local, remote, or local and remote. In Figure 8-43, we choose local.

Figure 8-43 Input clip name and choose the savemode of a clip

3. Choose which part of the archive you want to clip. You can select it based on time, such as from 09/24/2007 9:00.00 AM to 09/24/2007 9:05.00 AM. Input



the actual time in the UTC, and click the **get utc** button, and the system exports the UTC value automatically for you. The same applies for the stopUTC setting shown in Figure 8-44.

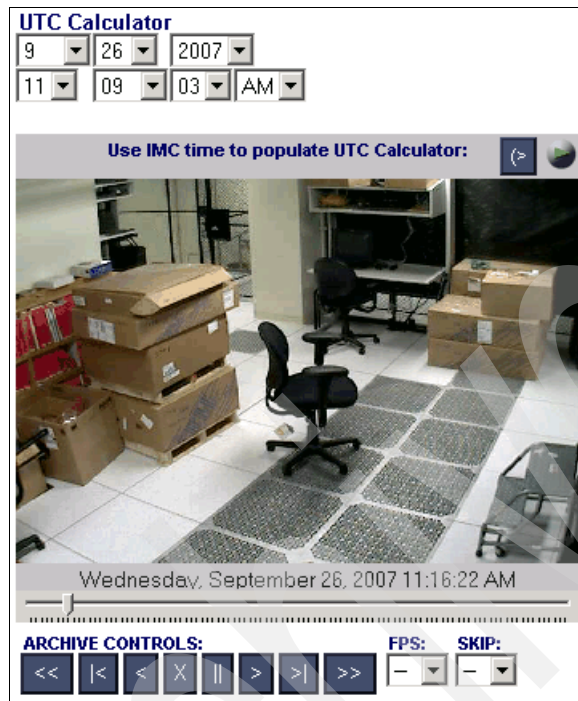


Figure 8-44 Import the start UTC time

4. You can also take a glance of the archive video from the little screen on the right-bottom side. Move the slide button anywhere, and then export it as the begin time or end time of the clip, as shown in Figure 8-45 on page 286.

### Save Clip

source= DVS\_Archive (archive, jpeg) ▼

name= optional

savemode= REQUIRED ▼ requires configuration for remote clips

startUTC= 1190830143000

stopUTC= 1190831070000

desc= optional

daystolive= 0 days (∅ never expires)

Figure 8-45 Setting times of clip

5. Select the format of the clip: regular, bwm, or bwx. In Figure 8-46, we choose regular.

### Save Clip

source= DVS\_Archive (archive, jpeg) ▼

name= optional

savemode= REQUIRED ▼ requires configuration for remote clips

startUTC= 1190830143000

stopUTC= 1190831070000

desc= optional

daystolive= 0 days (∅ never expires)

repos= optional

saveformat= REQUIRED ▼

key= REQUIRED (for bwx clips)

notifyURL= regular  
bwm  
bwx

Figure 8-46 Select the saveformat of the clip

6. Click the **execute command** button to generate a clip. If you successfully create a clip, it will return a code of '0', as shown in Figure 8-47 on page 287.



Figure 8-47 A clip is created

### 8.3.2 Extracting a clip with the AX Client

AX Client can display video and control cameras (but in our environment, our camera does not support to be controlled). The AX client can play live and archived video and audio sources using a single control. A single AX Client renders video in a single-video panel. We can view the clip on the Extract A Portable Clip with AX Client page:

1. When you get into this page, from the pull-down menu, choose an archive. You will extract a clip from that archive. In Figure 8-48, we use archive 'OptionalName\_DVS001'.

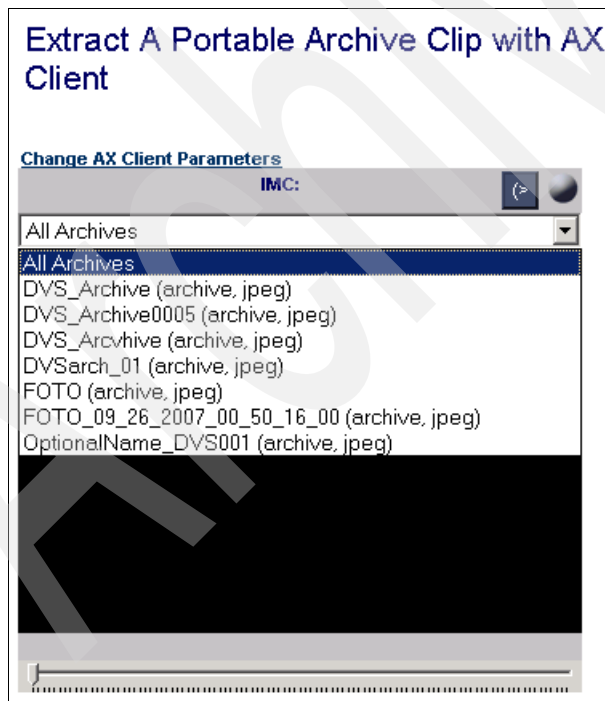


Figure 8-48 Extract a clip with AX Client page

After your selection, you will get a video-screen of the archive, as shown in Figure 8-49.




Figure 8-49 Video on the Extract Clip page

2. In this page, you can control your video show using the buttons in the ARCHIVE CONTROLS section. You can also input a specific time that you want to monitor in the JUMP-TO-DATE field, as shown in Figure 8-50 on page 289.




Figure 8-50 Video screen control in 'extract clip' page

- Now you can configure to save a specific clip. You can select the source archive from the pull-down menu, and the system automatically gives you the path of it. Next, specify the time point in the START CLIP and END CLIP fields. Determine if you want to save the clip to your local disk, as shown in Figure 8-51.

**SAVE CLIP:**  
**SOURCE URL (optional):**  
 bwims://9.43.86.85/FOTO (select from "All Archives" menu below)  
 FOTO  
**START CLIP:**  
 Tuesday, September 25, 2007 5:50:27 PM   
**END CLIP:**  
 Tuesday, September 25, 2007 5:50:32 PM   
**DESTINATION PATH (optional):**  
 FOTO001.wav  
**PROFILE (optional):**  
 WMV CBR 200Kbits

Figure 8-51 Extract a clip

In this Example, we save the clip to our local disk, and the clip name is FOTO001.wav, as shown in Figure 8-52. You can use the media player to play this video.




Figure 8-52 Video clip data on local disk

As a remote access user, you can also access the video data remotely from the source URL in Figure 8-53.



Figure 8-53 Access a Clip by explorer remotely

The Explorer calls related applications, such as Media Player, to play the video that you specified.

4. Finally, if you want to change the display resolution of the video-screen on the Extract Clip' page, click the **Change AX Client Parameters** link, which is located above the video screen, as shown in Figure 8-54 on page 291.

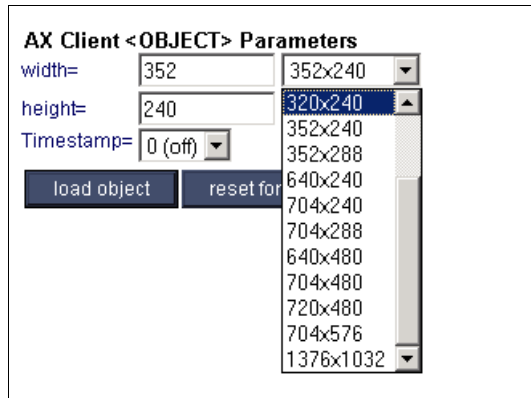


Figure 8-54 Reset the display resolution of the video on the Extract Clip page

5. Click the **load object** button to get the ideal screen size that you want.

## 8.4 Migrating data

If you are migrating data or have a need to migrate using N series storage, the N series provides features, such as SnapMirror and FlexClone, to assist you with Migration. Additionally, the mixed storage capabilities of the N series with FC disks and SATA disks behind the same storage system will assist those video surveillance applications, which can move data from one file system or volume to another. Get more information about SnapMirror in 6.5, “SnapMirror” on page 142 and about FlexClone in 6.5, “SnapMirror” on page 142.

## 8.5 Expiring data

Some type of expiration of the video capture data must occur so that storage can be reutilized and to keep storage costs down. The N series does not have an expiration feature at this time so the video surveillance application or assist addon must perform this function. Expiring data depends on how long your archives and clips needs to remain stored. Generally this time can be anywhere from 30 days to three years or more. This information is directly connected with your different security requirements. Probably, if you already have a video surveillance system and you are migrating to N series storage solution, you already have some policies.

If you are just implementing the video surveillance technology for the first time, you can use the following factors to know how long this data needs to be retained on storage:

- ▶ Number of cameras
- ▶ How many hours per day each camera records
- ▶ The type of camera resolution
- ▶ The frame rate
- ▶ Type of video compression
- ▶ How much storage you have available

Information about all of these factors are in 6.2, “Sizing N series storage to DVS” on page 129.

On CMS, the expiring date of each archive is provided using the field Days-to-live, and you can get more information about this field in “Set archive days-to-live command” on page 273.



## Exploiting N series features and DVS

In this chapter, we introduce the implementation of the following N series features:

- ▶ Flexvol
- ▶ MultiStore
- ▶ SnapMirror
- ▶ SnapLock

# 9.1 FlexVol

In this section, we explain how to resize a volume. We worked on the volume that was dedicated to archives and resized it at 25 GB.

## 9.1.1 Adding space to the aggregate (optional)

In the initial configuration, as we could see under **Aggregates** → **Manage**, we had 1.5 GB available in the aggregate aggr1\_node2. We could have used this space, but we preferred to add a new disk.

To add space to the aggregate:

1. Click the name of the aggregate, as shown in Figure 9-1.

**Important:** It is not possible to extend a volume by taking space from another aggregate that is not the initial aggregate that contains this volume. We work on aggregate aggr1\_node2 because VolArchCam2, the volume we want to extend, was located in this aggregate.

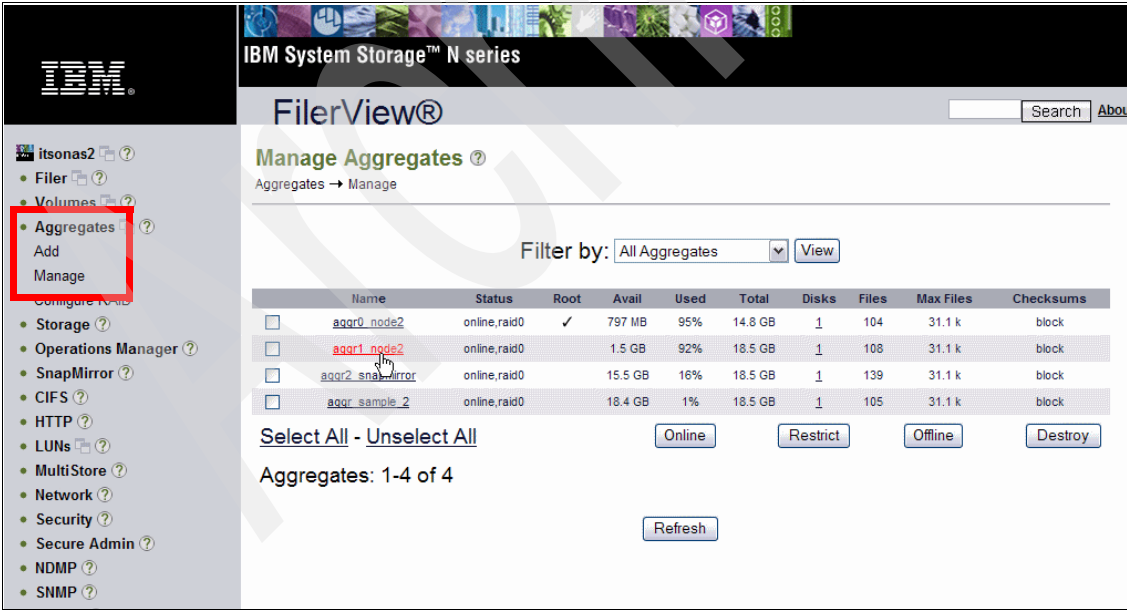


Figure 9-1 Aggregates / Manage

- A description of the aggregate is displayed.
2. As shown in Figure 9-2, click **Add Disks**.

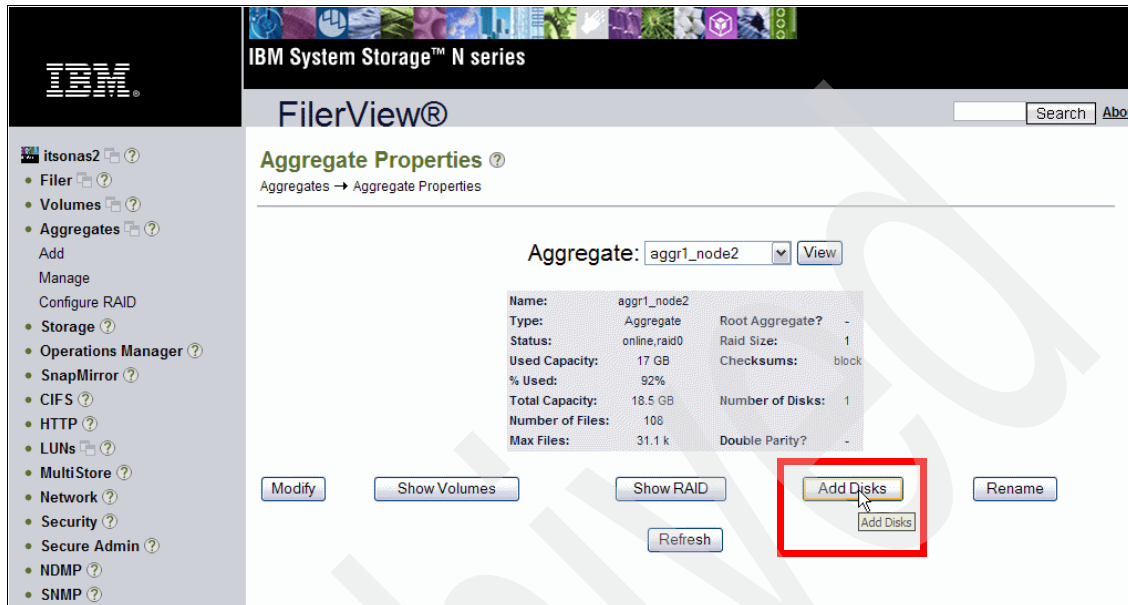


Figure 9-2 Adding disks to an aggregate.

The Aggregate Wizard is displayed, as shown in Figure 9-3.

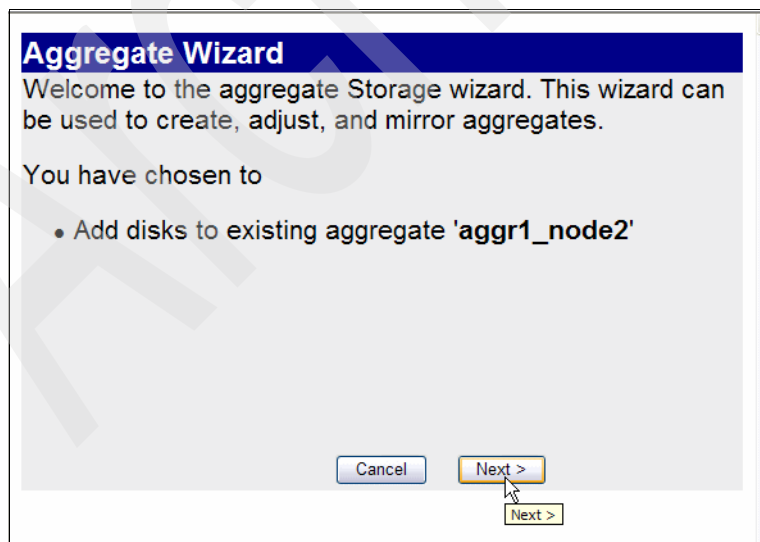


Figure 9-3 Aggregate wizard

- Next, indicate if you want disks automatically added or if you want to manually select. We chose **Manual Selection**. Click **Next**, as shown in Figure 9-4.



Figure 9-4 Disk Selection Method

As shown in Figure 9-5, we select one disk to add. Click **Next**.

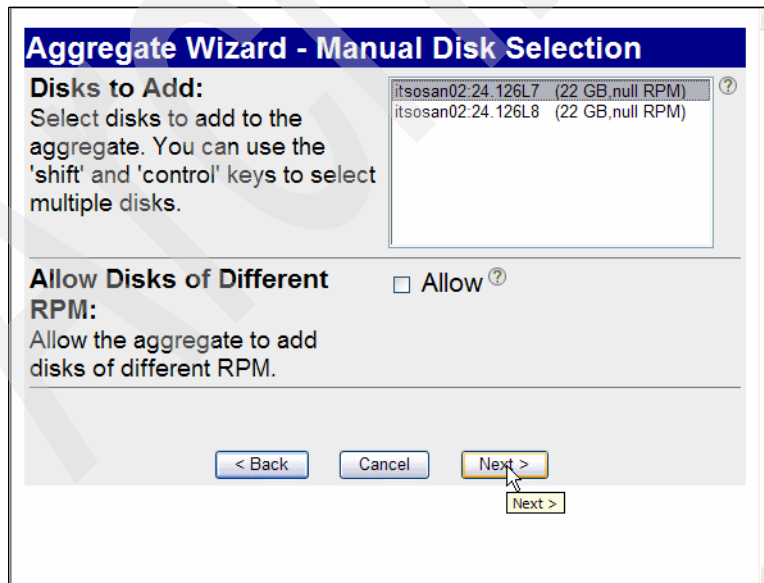


Figure 9-5 Disk selection

4. A summary of the modifications is displayed for confirmation, as shown in Figure 9-6. Click **Commit**.



Figure 9-6 Summary of the changes

A message is displayed, as shown in Figure 9-7, to indicate that the aggregate is updated.



Figure 9-7 Aggregate Updated

5. Select **Aggregates** → **Manage**, verify that the new disk was added to the aggregate and that the size changed, as shown in Figure 9-8.

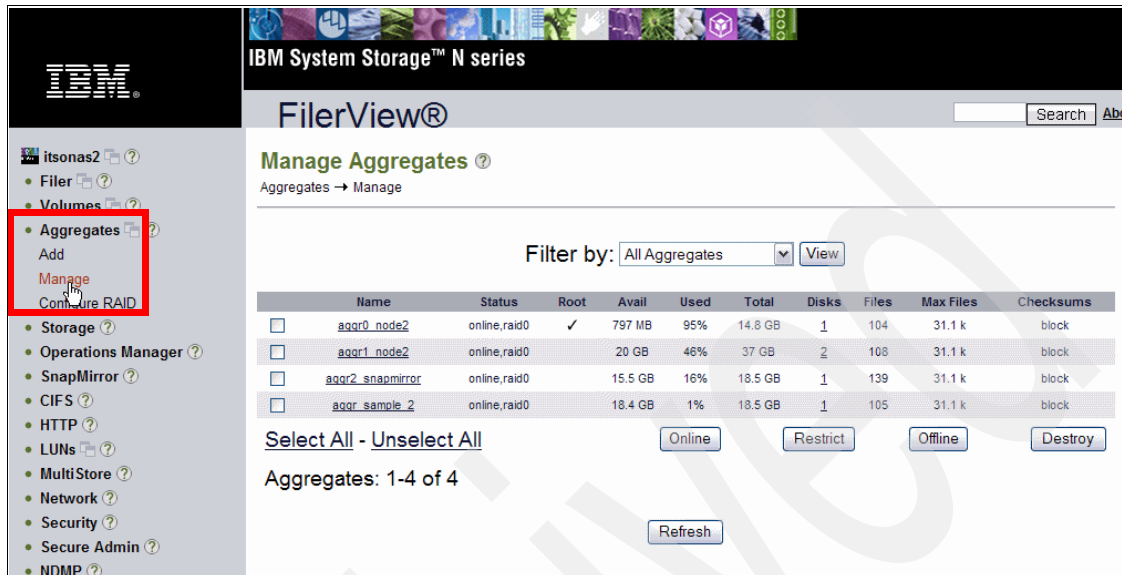


Figure 9-8 Aggregates / Manage: new configuration

## 9.1.2 Changing the size of a volume

To change the size of a volume:

1. In FilerView, as shown in Figure 9-9 on page 299, select **Volumes** → **Manage**, you can see the original size of the volume you want to modify. Click the name of this volume.

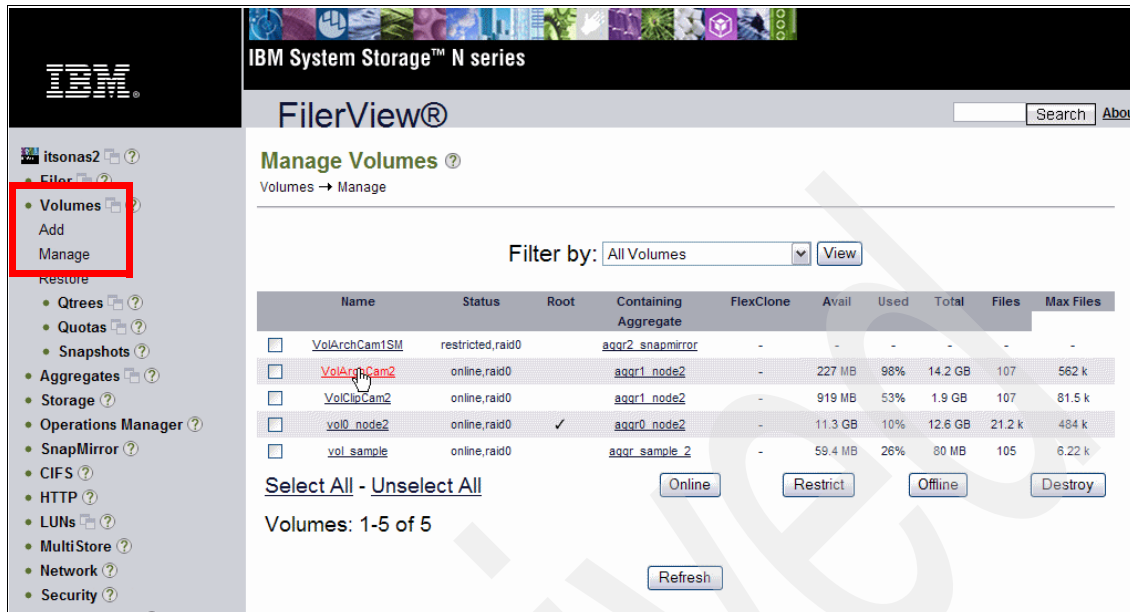


Figure 9-9 Resizing a volume

2. A window with information about the volume is displayed. Click **Resize Storage**, as shown in Figure 9-10 on page 300.

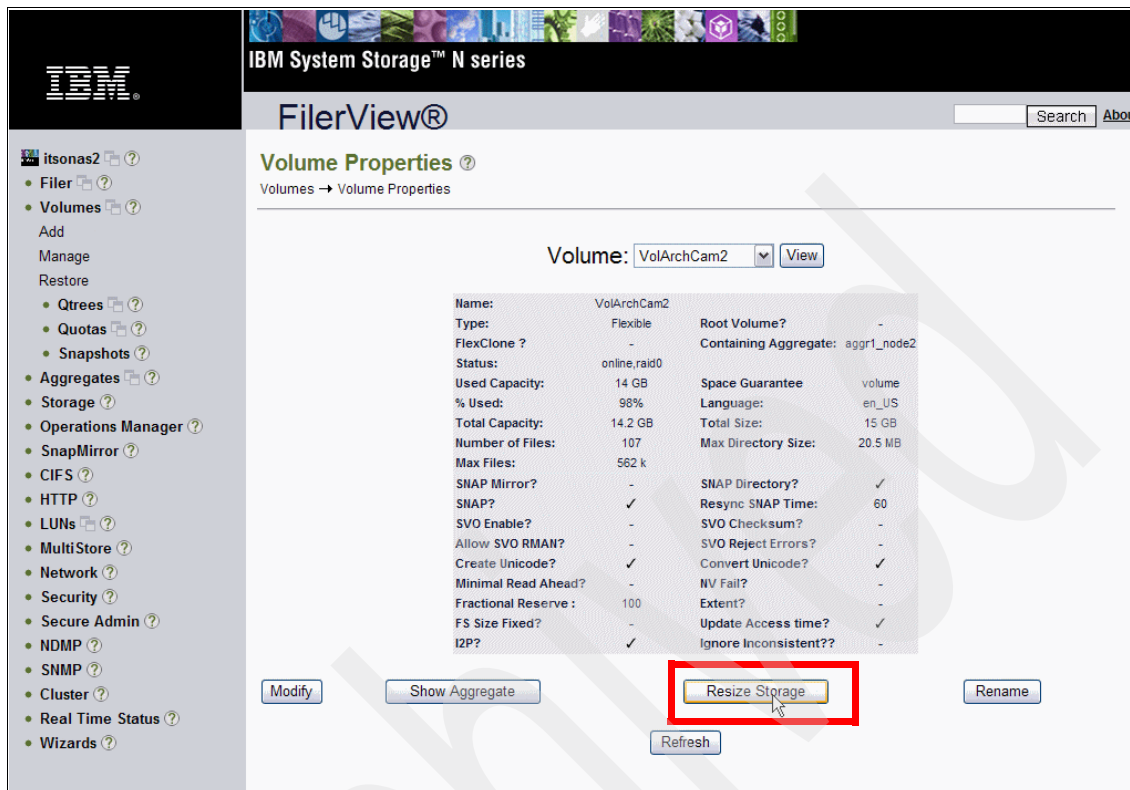


Figure 9-10 Volume Properties

- When the Volume wizard is displayed, click **Next**, as shown in Figure 9-11 on page 301.



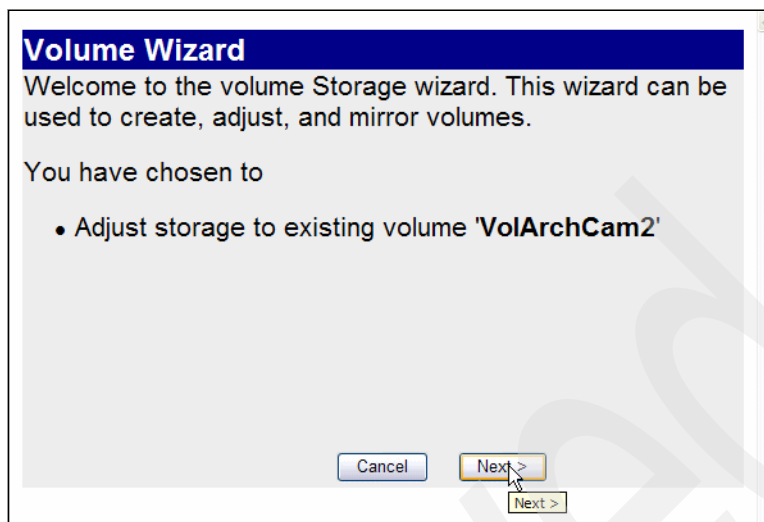


Figure 9-11 Volume Wizard

4. As shown in Figure 9-12, we chose Space Guarantee Type. For more information about this feature, refer to 7.1.3, “Creating the volumes” on page 211. Accept the default choice, volume, and click **Next**.

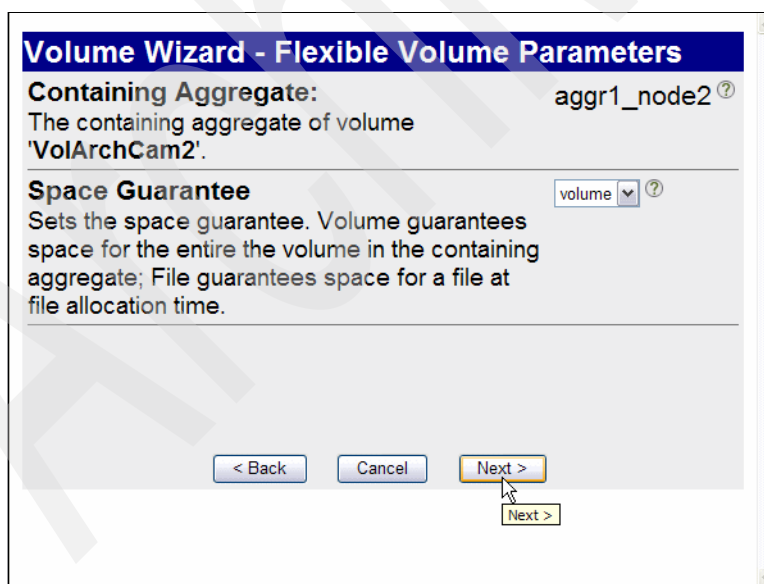
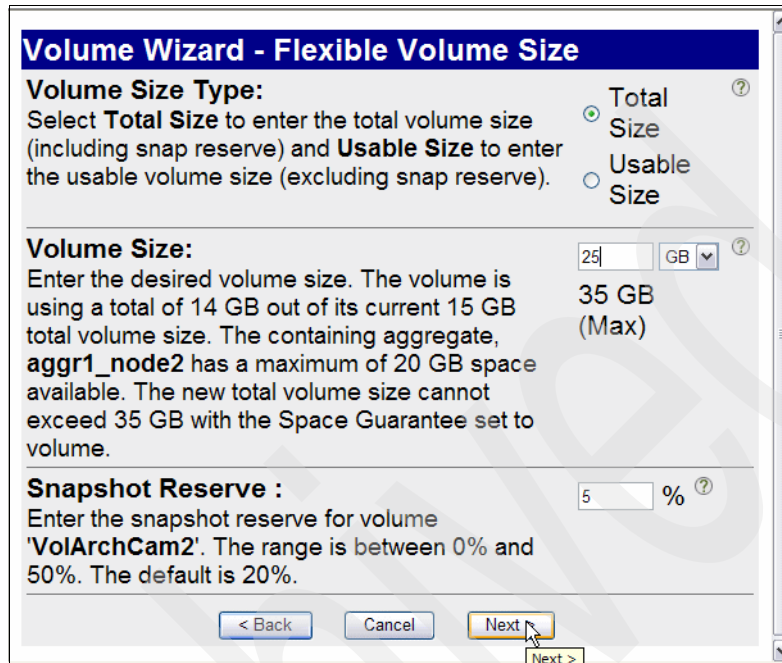


Figure 9-12 Flexible Volume Parameters

5. Choose the new size of the volume, as shown in Figure 9-13. We defined 25 GB.



The image shows a 'Volume Wizard - Flexible Volume Size' dialog box. It has three main sections: 'Volume Size Type', 'Volume Size', and 'Snapshot Reserve'. In the 'Volume Size Type' section, 'Total Size' is selected with a radio button. The 'Volume Size' section has a text input field with '25' and a unit dropdown menu set to 'GB'. To the right of the input field, it says '35 GB (Max)'. The 'Snapshot Reserve' section has a text input field with '5' and a percentage sign. At the bottom, there are three buttons: '< Back', 'Cancel', and 'Next >'. A mouse cursor is pointing at the 'Next >' button.

**Volume Wizard - Flexible Volume Size**

**Volume Size Type:**  
Select **Total Size** to enter the total volume size (including snap reserve) and **Usable Size** to enter the usable volume size (excluding snap reserve).

☒ Total Size  
☐ Usable Size

**Volume Size:**  
Enter the desired volume size. The volume is using a total of 14 GB out of its current 15 GB total volume size. The containing aggregate, **aggr1\_node2** has a maximum of 20 GB space available. The new total volume size cannot exceed 35 GB with the Space Guarantee set to volume.

25 GB 35 GB (Max)

**Snapshot Reserve :**  
Enter the snapshot reserve for volume 'VolArchCam2'. The range is between 0% and 50%. The default is 20%.

5 %

< Back Cancel Next >

Figure 9-13 Flexible Volume Size

6. As shown in Figure 9-14 on page 303, a window is displayed that summarizes the configuration. Click **Commit**.




Figure 9-14 Summary of the changes

A message is displayed that indicates that the changes were applied, as shown in Figure 9-15.




Figure 9-15 Configuration applied

7. **Volumes / Manage**, and verify that the size of the volume was modified, as shown in Figure 9-16 on page 304.




Figure 9-16 Volumes / manage: new configuration

### 9.1.3 Having more space at the server level

Now that we have more space available in the volume, we have two possibilities to use this new space at the server level:

- ▶ Create a new LUN, and map the LUN to the server. We do not detail these steps again because we provide them in Chapter 7, “Setting up the Digital Video Surveillance Solution” on page 201.
- ▶ Resize the current LUN, and make the new space available at the server level. It is developed in the following parts: for the iSCSI server first and then for our FC environment. The two procedures differ because the second one is a multipath configuration.

#### Resizing LUN in our no multipath configuration (iSCSI server)

To resize the LUN in our multipath configuration:

1. Resize the LUN at the N series level.

As shown in Example 9-6 on page 307, we:

- Put the LUN offline with the command: `lun offline lun_path`
- Resized the LUN with the command: `lun resize lun_path new_size`
- Put the LUN online again with the command: `lun online lun_path`

#### Example 9-1 Resizing LUN at the N series level

---

```
itsonas2> lun offline /vol/VolArchCam2/LUNArchCam2
Fri Sep 28 00:31:10 GMT [itsonas1: lun.offline:warning]: LUN
/vol/VolArchCam2/LUNArchCam2 has been taken offline

itsonas2> lun resize /vol/VolArchCam2/LUNArchCam2 20g
lun resize: resized to: 20.0g (21475885056)

itsonas2> lun online /vol/VolArchCam2/LUNArchCam2
```

---

2. Scan the LUN again at the server level for the new size to be discovered, as shown in Example 9-2.

#### Example 9-2 Discovering the new size of the disk

---

```
lochnese:~ # /etc/init.d/iscsi restart
Stopping iSCSI: sync
Message from syslogd@lochnese at Wed Sep 26 17:37:15 2007 ...
lochnese kernel: Disabling IRQ #2
  umount sync iscsid
done
Starting iSCSI: iscsi iscsid fsck/mount
done
```

---

3. Resize the partition using fdisk, as shown in Example 9-3, where we can see that the disk is resized and is now 20 GB but not the partition, which is still 14 GB. Delete the previous partition, and recreate it to cover the entire disk, as shown in Example 9-4 on page 306.

#### Example 9-3 Disk had been resized but not the partition

---

```
lochnese:/mnt # fdisk -l /dev/sda

Disk /dev/sda: 21.4 GB, 21475885056 bytes
64 heads, 32 sectors/track, 20481 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1		1	14337	<b>14681072</b>	83	Linux

---

#### Example 9-4 Recreate the partition disk

---

```
lochnese:/mnt # fdisk /dev/sda
```

The number of cylinders for this disk is set to 20481.

There is nothing wrong with that, but this is larger than 1024, and could in certain setups cause problems with:

- 1) software that runs at boot time (for example, old versions of LILO)
- 2) booting and partitioning software from other OSs (for example, DOS FDISK, OS/2 FDISK)

```
Command (m for help): d
```

```
Selected partition 1
```

```
Command (m for help): n
```

```
Command action
```

```
  e   extended
```

```
  p   primary partition (1-4)
```

```
p
```

```
Partition number (1-4): 1
```

```
First cylinder (1-20481, default 1):
```

```
Using default value 1
```

```
Last cylinder or +size or +sizeM or +sizeK (1-20481, default 20481):
```

```
Using default value 20481
```

```
Command (m for help): w
```

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
```

```
Syncing disks.
```

---

4. Increase the size of the file system, as shown in Example 9-5.

#### Example 9-5 Resizing the file system

---

```
lochnese:/mnt # xfs_growfs /dev/sda1
```

```
meta-data=/Archives          isize=256    agcount=16, agsize=229391  
blks
```

```
        =                      sectsz=512
```

```
data      =                      bsize=4096   blocks=3670256,
```

```
imaxpct=25
```

```
        =                      sunit=0      swidth=0 blks,
```

```
unwritten=1
```

```
naming    =version 2            bsize=4096
```

```
log       =internal            bsize=4096   blocks=2560, version=1
```

```
        =                      sectsz=512    sunit=0 blks
```

```

realtime =none                      extsz=65536  blocks=0, rtextents=0
data blocks changed from 3670256 to 5243132
lochnese:/mnt # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda3              43306144    2617412  40688732   7% /
tmpfs                  1876552         12   1876540   1% /dev/shm
/dev/sdb1              1043552        144   1043408   1% /Clips
/dev/sda1             20962288        384  20961904   1% /Archives

```

---

## Resizing LUN in our multipath configuration (FC server)

To resize the LUN in a multipath configuration (FC server):

1. Resize the LUN at the N series level, as shown in Example 9-6:
  - We put the LUN offline with the command: **lun offline lun\_path**
  - We resized the LUN with the command: **lun resize lun\_path new\_size**
  - We put the LUN online again with the command: **lun online lun\_path**

*Example 9-6 Resizing LUN at the N series level*

```

itsonas1> lun offline /vol/VolArchCam1/LUNArchCam1
Fri Sep 28 00:31:10 GMT [itsonas1: lun.offline:warning]: LUN
/vol/VolArchCam1/LUNArchCam1 has been taken offline

itsonas1> lun resize /vol/VolArchCam1/LUNArchCam1 20g
lun resize: resized to:  20.0g (21475885056)

itsonas1> lun online /vol/VolArchCam1/LUNArchCam1

```

---

2. Scan the LUN again at the server level by unmounting and remounting the qla2300 driver:

```

modprobe -r qla2300
modprobe qla2300

```

Example 9-7 shows that the new size of the disk of 20 GB was discovered, but /Archives were still 14 GB because the file system was not resized.

*Example 9-7 Disk had been resized*

```

tonga:~ # fdisk -l /dev/sda

Disk /dev/sda: 21.4 GB, 21475885056 bytes
64 heads, 32 sectors/track, 20481 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes

```

Disk /dev/sda doesn't contain a valid partition table

```
tonga:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda2              73911032    3611308   70299724   5% /
tmpfs                 1876552         12   1876540   1% /dev/shm
/dev/dm-0             14670848    692120   13978728   5% /Archives
/dev/dm-1             1038336    915372    122964  89% /Clips
/dev/dm-2             1038336      148   1038188   1% /test
```

---

3. Resize the file system, as shown in Example 9-8 with the command:

`xfs_growfs file_system_name`

Using the output of the command **df**, we can verify that our repository is now 20 GB.

*Example 9-8 Resizing the file system*

---

```
tonga:~ # xfs_growfs /dev/dm-0
meta-data=/Archives      isize=256    agcount=16, agsize=229392
blks
        =                sectsz=512
data      =                bsize=4096   blocks=3670272,
imaxpct=25
        =                sunit=0       swidth=0 blks,
unwritten=1
naming    =version 2      bsize=4096
log       =internal       bsize=4096   blocks=2560, version=1
        =                sectsz=512   sunit=0 blks
realtime  =none          extsz=65536  blocks=0, rtextents=0
data blocks changed from 3670272 to 5243136
tonga:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda2              73911032    3611352   70299680   5% /
tmpfs                 1876552         12   1876540   1% /dev/shm
/dev/dm-0             20962304    692232   20270072   4% /Archives
/dev/dm-1             1038336    915372    122964  89% /Clips
/dev/dm-2             1038336      148   1038188   1% /test
```

---



## 9.2 Multistore

In this section, we explain how to set up Multistore.

### 9.2.1 Adding a Multistore license

We put Multistore licenses in both nodes of the cluster, as shown in Figure 9-17 on page 310.

#### Clustering Considerations:

- ▶ Each member of a cluster must have a MultiStore license to take over its partner with a MultiStore license.
- ▶ The Vfiler™ units hosted by the storage systems of the cluster are created and configured independently. That is, each storage system can host a different number of Vfiler units, and the Vfiler unit configurations on the storage systems can be different from each other.
- ▶ In takeover mode, the functioning storage system takes over all Vfiler units that are created on the failed storage system. These Vfiler units include the Vfiler units you create and the unit called Vfiler0. Therefore, for Vfiler units on the failed storage system to work correctly after the takeover, each network interface that a Vfiler unit uses in a cluster must have a partner interface.

To add a Multistore license:

1. Select **Filer** → **Manage Licenses**, as shown in Figure 9-17 on page 310.
2. At the bottom of the page, click **Apply**, and a message is displayed that indicates that the license updated successfully.




Figure 9-17 Adding Multistore License

## 9.2.2 Preparing the node for Vfiler creation

You must decide which resources to allocate to the filer. These resources are:

- ▶ IP addresses
- ▶ Volumes

### Internet Protocol addresses

We had two Ethernet interfaces in our node e0a and e0b. Because we wanted to assign e0b to the Vfiler that we are about to create, we put it down.

To prepare the node for Vfiler creation:

Click **Network** → **Manage Interfaces**. The Manage Network Interfaces page is displayed.

Under the Operations section, click **down** for **e0b**, as shown in Figure 9-18.

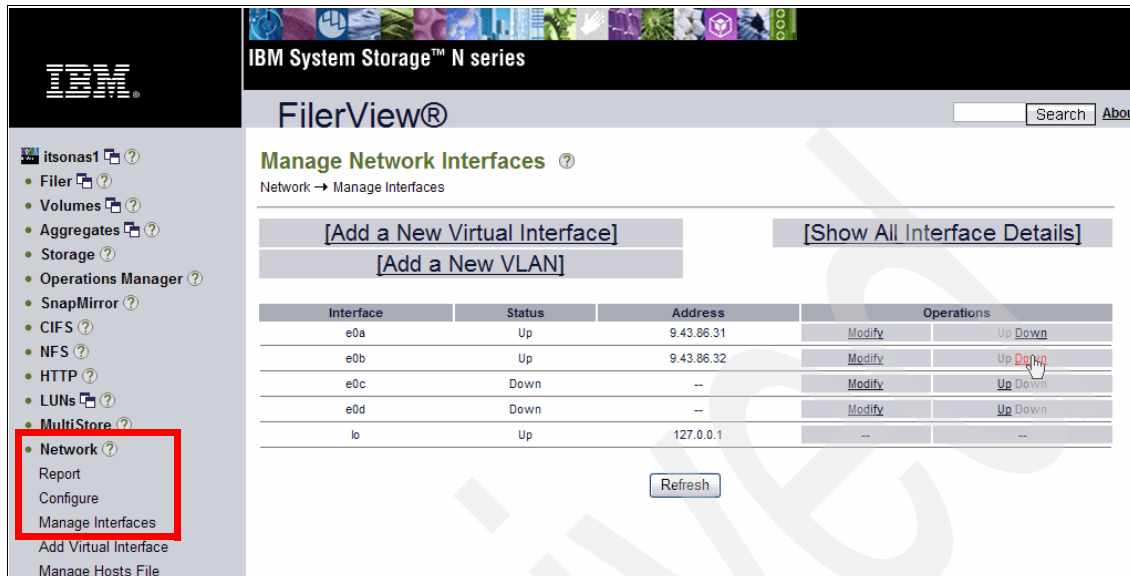


Figure 9-18 Putting down the interface to be attached to the Vfiler

Notice that we use e0a to access the node; therefore, we can put down e0b without losing connection.

3. A confirmation message is displayed, as shown in Figure 9-19. Click **OK**.

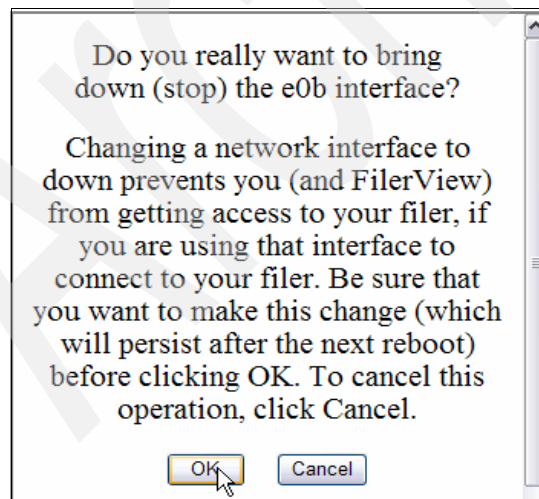


Figure 9-19 Confirmation Message

The e0b appears as down, as shown in Figure 9-20.

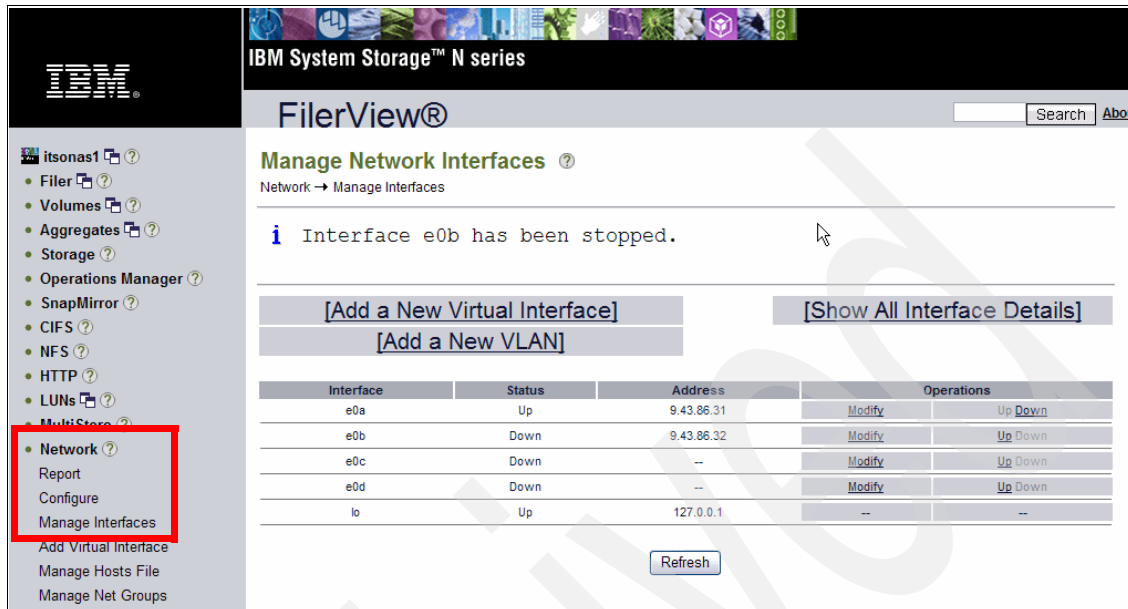


Figure 9-20 Interface down

## Volumes

We need to define a volume that the Vfiler can use as a root volume. Create a new aggregate and a volume in this aggregate. This volume must not be a read-only volume.

Because we previously discussed the process of creating an aggregate and a volume, we do not discuss it here. Review 7.1.2, "Creating the aggregate" on page 206 and 7.1.3, "Creating the volumes" on page 211 for detailed information about creating an aggregate and a volume.

Create one aggregate, aggr\_vfiler, with the following options:

- ▶ name: aggr\_vfiler
- ▶ RAID group size: 1
- ▶ Disk Selection: Manual
- ▶ Disk used: itsosan02:24.126L3

Create one volume, VfilerB, with the following options:

- ▶ Volume Type: Flexible
- ▶ Volume Name: Vfiler
- ▶ Language: English US
- ▶ Containing Aggregate: aggr\_vfiler

- ▶ Space Guarantee: volume
- ▶ Volume Size Type: Total Size
- ▶ Volume Size: 15 GB
- ▶ Snapshot Reserve: 5%

### 9.2.3 Creating a Vfiler

To create a Vfiler:

1. Click **MultiStore** → **Manage Vfilers**, as shown in Figure 9-21. The initial configuration includes a unique Vfiler, `vfiler0`, which is the default filer we are using.




Figure 9-21 MultiStore - Initial Configuration

2. Click **MultiStore** → **Vfiler Wizard**, as shown in Figure 9-22 on page 314.

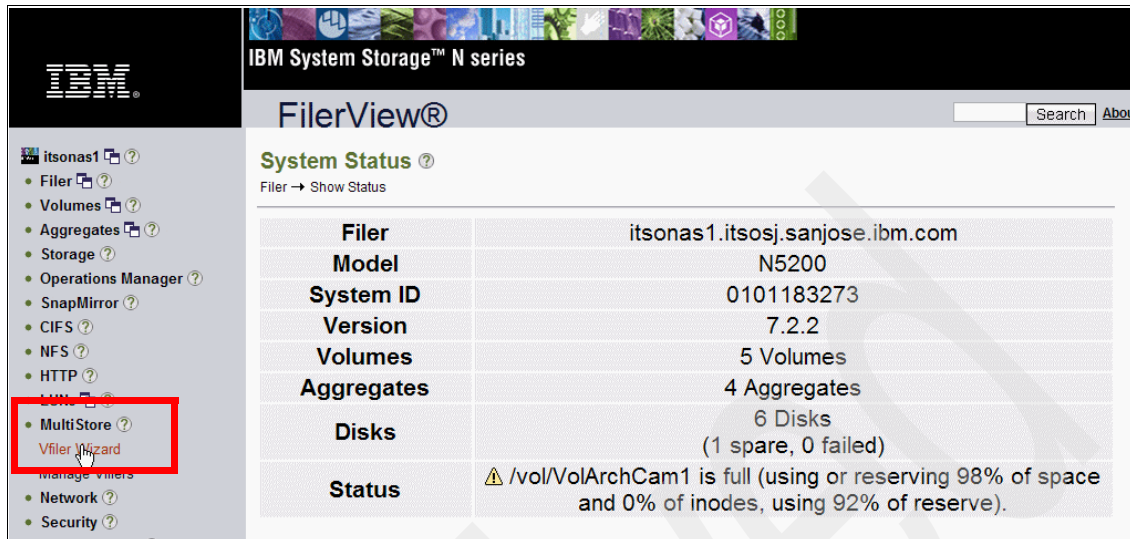


Figure 9-22 Configuring MultiStore

- The Vfiler Wizard is displayed with options to **Create a new vfiler** and **Configure the new filer**, which are selected by default, as shown in Figure 9-23. Click **Next**.

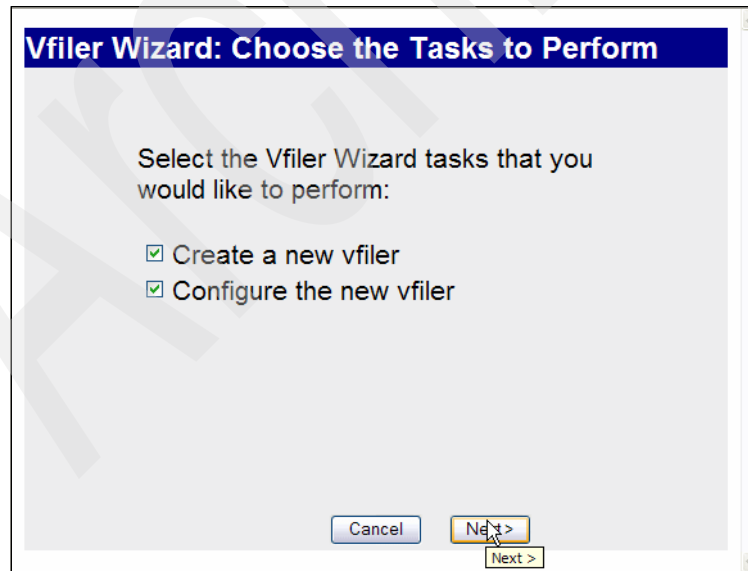
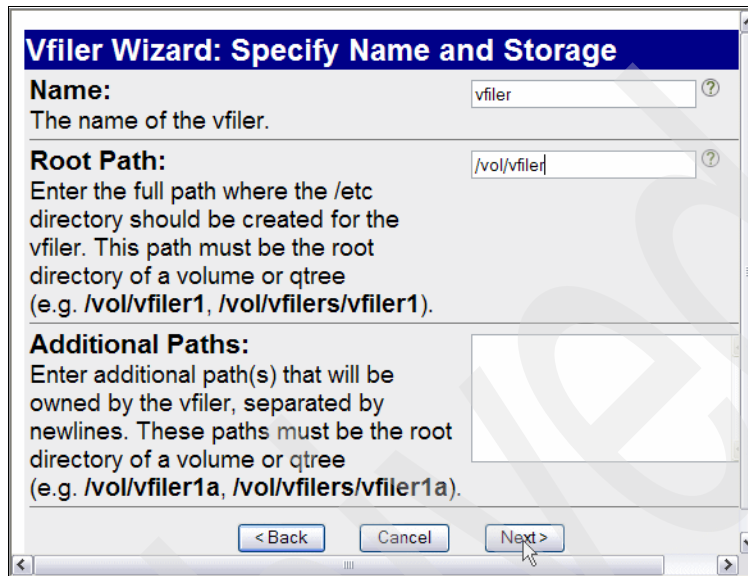


Figure 9-23 Vfiler Wizard

4. Define `vfiler` as the name of the vfiler and specify the root volume to use. Use the volume, `/vol/vfiler`, which we created in 9.2.2, “Preparing the node for Vfiler creation” on page 310, as shown in Figure 9-24.



The image shows a dialog box titled "Vfiler Wizard: Specify Name and Storage". It has three sections: "Name:", "Root Path:", and "Additional Paths:". The "Name:" section has a text box containing "vfiler". The "Root Path:" section has a text box containing "/vol/vfiler". The "Additional Paths:" section has an empty text box. At the bottom, there are three buttons: "< Back", "Cancel", and "Next >". A mouse cursor is pointing at the "Next >" button.

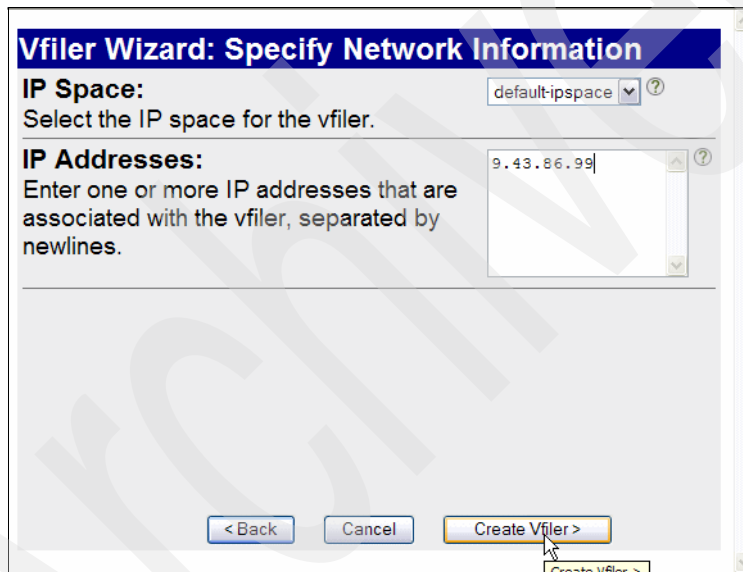
Figure 9-24 Specifying name and root volume of the Vfiler

5. Define the IP address to be used for the Vfiler, as shown in Figure 9-25 on page 316.

In our configuration, we use an IP address in the same range of addresses as the network of our primary filer, and we do not configure any IPspace because we do not need to isolate the filer and the Vfiler. If we were using a different range of addresses for the primary and virtual filer, they would be isolated by default. If we were using addresses in the same range and wanted to isolate them, we would use the IPspaces.

### About IPspaces:

An IPspace defines a distinct IP address space in which vFiler units can participate. IP addresses that are defined for an IPspace are meaningful only within that IPspace. A distinct routing table is maintained for each IPspace. No cross-IPspace traffic is routed, for example, if you have two distinct enterprises that access the same filer with two virtual filers defined and one different Ethernet card that belongs to each Vfiler but using the same range of IP addresses. You do not want computers from one network to communicate to computers of the other network through the filer. So, for each Vfiler, you define one IPspace to isolate the two networks.



The image shows a screenshot of a software window titled "Vfiler Wizard: Specify Network Information". The window has a blue header bar with the title. Below the header, there are two main sections. The first section is labeled "IP Space:" and contains a dropdown menu with "default-ipspace" selected and a help icon. Below this is the instruction "Select the IP space for the vfiler." The second section is labeled "IP Addresses:" and contains a text input field with "9.43.86.99" entered and a help icon. Below this is the instruction "Enter one or more IP addresses that are associated with the vfiler, separated by newlines." At the bottom of the window, there are three buttons: "< Back", "Cancel", and "Create Vfiler >". A mouse cursor is pointing at the "Create Vfiler >" button.

Figure 9-25 IP address definition

6. Create and configure a Vfiler, as shown in Figure 9-26 on page 317.



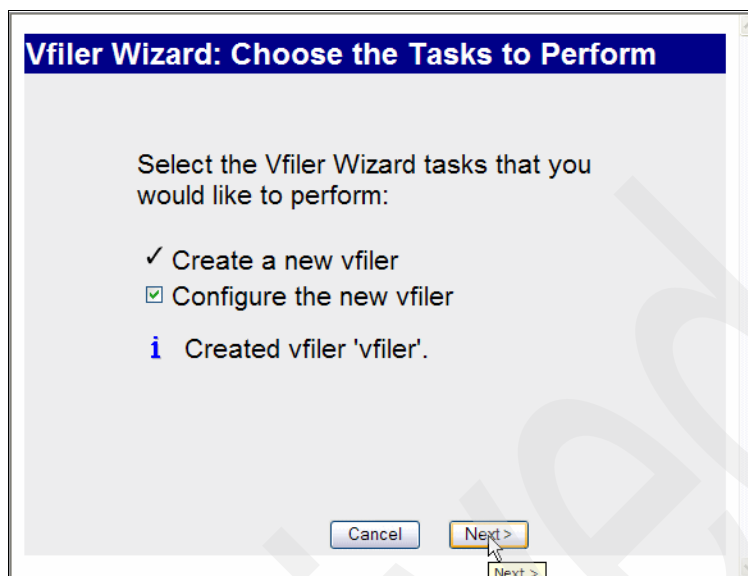
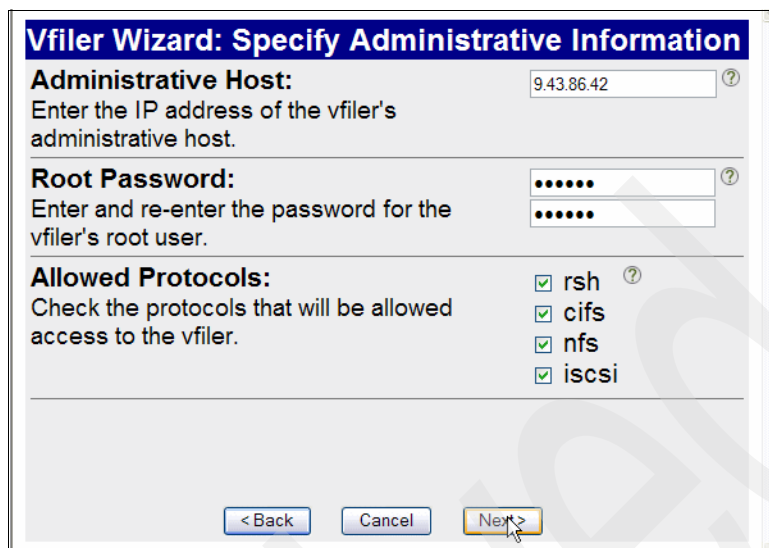


Figure 9-26 Vfiler Created

7. Define the administration parameters, as shown in Figure 9-27 on page 318:

- Administration host
- Password

Click **Next**.



**Vfiler Wizard: Specify Administrative Information**

**Administrative Host:** 9.43.86.42 ?  
Enter the IP address of the vfiler's administrative host.

**Root Password:** ..... ?  
Enter and re-enter the password for the vfiler's root user.

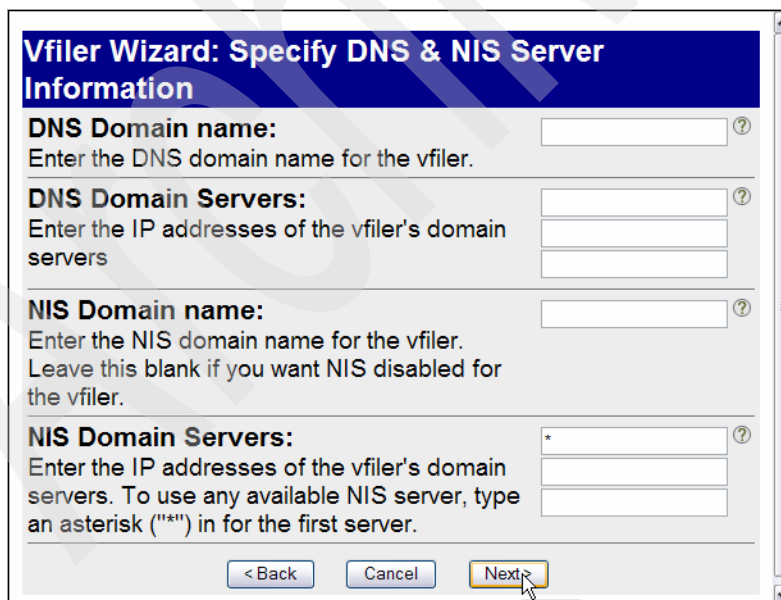
**Allowed Protocols:** ?  
Check the protocols that will be allowed access to the vfiler.

- ☒ rsh
- ☒ cifs
- ☒ nfs
- ☒ iscsi

< Back    Cancel    Next >

Figure 9-27 Administration parameters

8. Do not define a DNS or NIS server right now. Click **Next**, as shown in Figure 9-28.



**Vfiler Wizard: Specify DNS & NIS Server Information**

**DNS Domain name:** ?  
Enter the DNS domain name for the vfiler.

**DNS Domain Servers:** ?  
Enter the IP addresses of the vfiler's domain servers

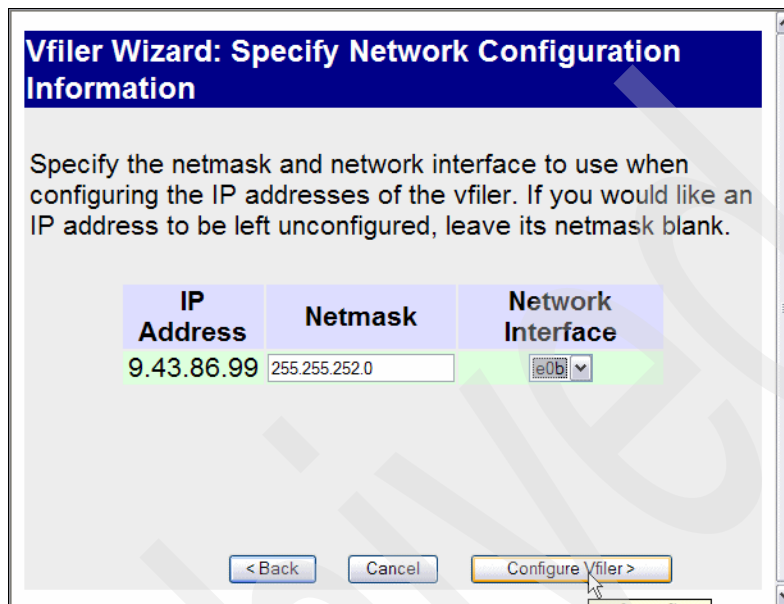
**NIS Domain name:** ?  
Enter the NIS domain name for the vfiler.  
Leave this blank if you want NIS disabled for the vfiler.

**NIS Domain Servers:** ?  
Enter the IP addresses of the vfiler's domain servers. To use any available NIS server, type an asterisk ("\*") in for the first server.

< Back    Cancel    Next >

Figure 9-28 DNS and NIS server

9. Configure the network, as shown in Figure 9-29, and specify the netmask to use and the physical Ethernet card to be associated to the Vfiler and to the IP address that is already denied. Click **Configure Vfiler>**.



The dialog box titled "Vfiler Wizard: Specify Network Configuration Information" contains the following text: "Specify the netmask and network interface to use when configuring the IP addresses of the vfiler. If you would like an IP address to be left unconfigured, leave its netmask blank." Below this text is a table with three columns: "IP Address", "Netmask", and "Network Interface". The "IP Address" column contains the value "9.43.86.99". The "Netmask" column contains the value "255.255.252.0". The "Network Interface" column contains a dropdown menu with the value "e0b" selected. At the bottom of the dialog box are three buttons: "< Back", "Cancel", and "Configure Vfiler >". A mouse cursor is pointing at the "Configure Vfiler >" button.

IP Address	Netmask	Network Interface
9.43.86.99	255.255.252.0	e0b

< Back   Cancel   Configure Vfiler >

Figure 9-29 Network Configuration

10. Vfiler is configured. Click **Close Window**, as shown in Figure 9-30 on page 320.

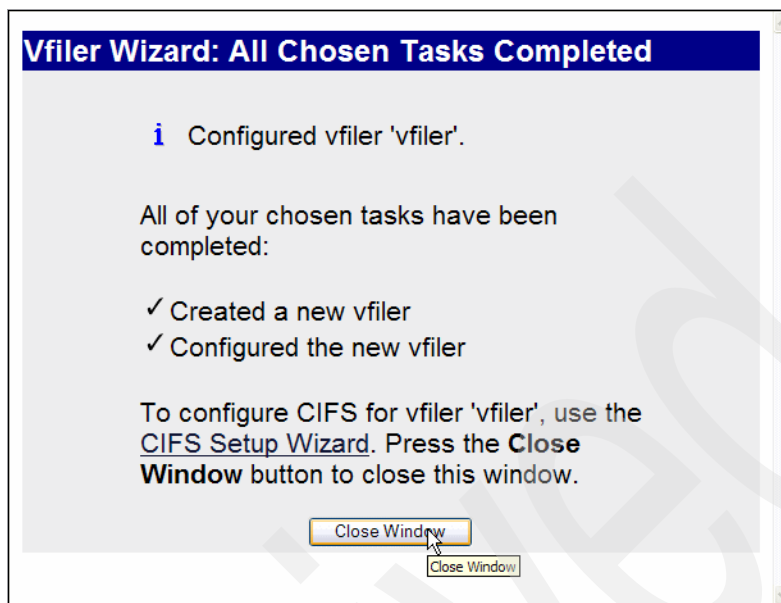


Figure 9-30 Vfiler configured

## 9.2.4 Using Vfiler

By default, Vfiler:

- Is accessible through rsh from the server that we indicated as the administration host, as shown in Example 9-9.

*Example 9-9 Accessing Vfiler through rsh*

---

```
tonga:~ # rsh 9.43.86.99 hostname
vfiler
```

---

- Exported as NFS to share the root volume with permissions that authorize the administration host to access it. We can see in Example 9-10 that files are present in the root volume of the Vfiler.

*Example 9-10 Root Volume of the Vfiler*

---

```
tonga:~ # mount 9.43.86.99:/ /mnt
tonga:~ # cd /mnt
tonga:/mnt # ls -l
total 16
drwxr-xr-x  4 root root 4096 Sep 20 14:10 .
```

---

```

drwxr-xr-x  4 root root 4096 Sep 20 14:10 ..
drwxrwxrwx  2 root root 4096 Sep 19 17:52 .snapshot
drwxr-xr-x  7 root root 4096 Sep 20 16:16 etc
tonga:/mnt # cd etc
tonga:/mnt/etc # ls -l
total 132
drwxr-xr-x  7 root root 4096 Sep 20 16:16 .
drwxr-xr-x  4 root root 4096 Sep 20 14:10 ..
-rw-r--r--  1 root root  600 Sep 20 14:10 .a00100
-rw-r--r--  1 root root  105 Sep 20 15:01 exports
-rw-r--r--  1 root root  105 Sep 20 14:56 exports.bak
drwxr-xr-x  3 root root 4096 Sep 20 14:14 exports_arc
-r--r----- 1 root root   28 Sep 20 14:10 filersid.cfg
-rw-r--r--  1 root root   74 Sep 20 14:26 hosts
-rw-r--r--  1 root root   65 Sep 20 15:01 hosts.equiv
-rw-r--r--  1 root root   65 Sep 20 14:56 hosts.equiv.bak
drwxr-xr-x  6 root root 4096 Sep 20 14:10 keymgr
-rwxrwx---  1 root root  397 Sep 20 14:26 lclgroups.cfg
drwxr-xr-x  2 root root 4096 Sep 20 14:10 log
-rw-r--r--  1 root root  202 Sep 20 14:26 nsswitch.conf
-rw-r--r--  1 root root   54 Sep 20 14:26 quotas
-rw-r--r--  2 root root 11834 Sep 20 16:12 registry
-rw-r--r--  1 root root 11834 Sep 20 16:12 registry.0
-rw-r--r--  2 root root 11834 Sep 20 16:12 registry.1
-rw-r--r--  1 root root 11834 Sep 20 16:12 registry.bck
-rw-r--r--  1 root root  2567 Sep 20 14:10 registry.default
-rw-r--r--  1 root root  2640 Sep 20 14:10 registry.lastgood
-rw-r--r--  2 root root  1876 Sep 20 16:12 registry.local
-rw-r--r--  2 root root  1876 Sep 20 16:12 registry.local.0
-rw-r--r--  1 root root  1876 Sep 20 16:12 registry.local.1
-rw-r--r--  1 root root  1876 Sep 20 16:12 registry.local.bck
-rwxr-xr-x  1 root root   16 Sep 20 15:37 rmtab
drwxr-xr-x  2 root root 4096 Sep 20 14:10 sm
drwxr-xr-x  2 root root 4096 Sep 20 15:03 sshd
-r-----  1 root root   12 Sep 20 14:10 vfiler-storage

```

---

**The vfiler command:** In this book, we do not go into details about administering Vfiler. However, notice that one of the best ways to administer the Vfiler is to use from Vfiler0 (the primary filer) the command **vfiler**, as shown in Example 9-11 and Example 9-12 on page 323.

With the **vfiler** command, you can, for example, define resources that are attributed to the Vfiler and can also send commands directly to the Vfiler with the following syntax:

```
vfiler run vfiler_hostname command
```

Refer to the official documentation for more details.

Example 9-11 shows the user for the **vfiler** command.

*Example 9-11 Vfiler command*

---

```
itsonas1> vfiler
The following commands are available; for more information
type "vfiler help <command>"
add                disallow                migrate                run
allow              dr                      move                  start
context            help                    remove                status
create             limit                   rename                stop
destroy
```

```
vfiler help - Help for vfiler command.
vfiler context - Set the vfiler context of the CLI.
vfiler create - Create a new vfiler.
vfiler rename - Rename an existing vfiler.
vfiler destroy - Release vfiler resources.
vfiler dr - Configure a vfiler for disaster recovery.
vfiler add - Add resources to a vfiler.
vfiler remove - Remove resources from a vfiler.
vfiler migrate - Migrate a vfiler from a remote filer.
vfiler move - Move resources between vfilers.
vfiler start - Restart a stopped vfiler.
vfiler stop - Stop a running vfiler.
vfiler status - Provide status on vfiler configuration.
vfiler run - Run a command on a vfiler.
vfiler allow - Allow use of a protocol on a vfiler.
vfiler disallow - Disallow use of a protocol on a vfiler.
vfiler limit - Limit the number of vfilers that can be created.
```

---

Example 9-12 shows the **run** command in a Virtual Filer through the **vfiler** command.

*Example 9-12 Run command in Virtual Filer through vfiler command*

```
itsonas1> vfiler run vfiler help

===== vfiler
?                hostname          nfsstat
snapmirror
arp              igroup          nis
snapvault
cifs             ipsec           options
traceroute
config           iscsi           passwd
useradmin
df               keymgr          ping            vol
dns              lock            qtree           vscan
echo             lun             quota           wcc
exportfs         nbtstat         route           ypcat
filestats        ndmpcopy        secureadmin     ypgroup
fpolicy          ndmpd           setup           ypmatch
fsecurity        netstat         snap            ypwhich
help             nfs
```

### 9.3 SnapMirror

In this section, we describe the implementation of Synchronous SnapMirror. We implement SnapMirror between the two nodes of our cluster.

**Disable clustering:** Disable clustering because it is not possible to implement SnapMirror between the two nodes of a cluster.

**Synchronous SnapMirror:** SnapMirror can be established between qtrees and between volumes, but it is not possible to do Synchronous SnapMirror between qtrees, only between volumes. In our configuration, we implement Synchronous SnapMirror between volumes.

To implement Synchronous SnapMirror:

1. Click **Filer** → **Manage** to access FilerView.
2. Add licenses for SnapMirror and SnapMirror Synchronous for both nodes, as shown in Figure 9-31. In a production environment, SnapMirror occurs from one filer in one site to another filer in another site. Because we only have a cluster for our test, we implement SnapMirror from the node1 to the node2.




Figure 9-31 Adding Licenses

3. To configure SnapMirror, we need access to the configuration files of both N series. Mount for the source filer, /etc partition of both N series, as shown in Example 9-13.

*Example 9-13 Accessing the configuration files*

```
tonga:~ # mount 9.43.86.31:/etc /mnt
tonga:~ # cd /mnt
```



```
tonga:/mnt # ls
.          hosts.equiv.bak      registry
..         http         registry.0
.avail     httpd.mimetypes.sample registry.1
.snapshot inst_perm.bat         registry.bck
.zapi      install.bat   registry.default
apc_ups     java         registry.lastgood
asup_content.conf keymgr    registry.local
asuptriggers.sample lang      registry.local.0
boot       lclgroups.cfg    registry.local.1
cifs_homedir.cfg log       registry.local.bck
cifs_nbalias.cfg man        resolv.conf.bak
cifsconfig_setup.bak messages  rlm_config_to_filer
cifsconfig_setup.cfg messages.0 rlm_stats
cifsconfig_share.bak messages.1 rmtab
cifsconfig_share.cfg messages.2 serialnum
cifssec.bak  messages.3 services
crash       messages.4 shelf_fw
dgateways  messages.5 shm_fcsl_stat.txt
dgateways.bak mib       shm_med_err.log
disk_fw    mlnx      shm_stat.bin
exports    modules    shm_stat.txt
exports.bak netapp_filer.dtd sm
exports.old nsswitch.conf sshd
exports_arc nsswitch.conf.bak stats
filersid.cfg passwd   sysconfigtab
firmware   qual_devices syslog.conf.sample
group      qual_devices_v2 tape_config
hosts      quotas    templates
hosts.bak  rc          usermap.cfg
hosts.bak.20070919 rc.back    vfiler
hosts.equiv rc.bak     zoneinfo
```

---

4. Configure `/etc/hosts` to have the definition of the filers, as shown in Example 9-14 and Example 9-15 on page 326, and decide which Ethernet card to use. It is not possible to do this in our configuration; however, it is possible to aggregate two Ethernet cards.

*Example 9-14 Set up of `/etc/hosts` for the source filer*

---

```
#Generated by setup Wed Sep 19 23:57:14 GMT 2007
#Auto-generated by setup Tue Jun 26 22:44:41 GMT 2007
127.0.0.1 localhost
9.43.86.31 itsonas1 itsonas1-e0a
9.43.86.32 itsonas1-e0b
```

```
# 0.0.0.0 itsonas1-e0c
# 0.0.0.0 itsonas1-e0d
9.43.86.33 itsonas2 itsonas2-e0a
```

---

*Example 9-15 Set up of /etc/hosts for the destination filer*

---

```
#Auto-generated by setup Wed Sep 12 18:32:07 GMT 2007
127.0.0.1 localhost
9.43.86.33 itsonas2 itsonas2-e0a
9.43.86.34 itsonas2-e0b
# 0.0.0.0 itsonas2-e0c
# 0.0.0.0 itsonas2-e0d
9.43.86.31 itsonas1 itsonas1-e0a
```

---

5. Create a volume on the destination filer. The size of this volume must be equal to or superior to the size of the source volume.

Because we previously explained the process of creating the volumes (7.1.3, “Creating the volumes” on page 211), we do not provide those instructions here. The parameters that are specified for the volume are:

- Volume Name: VolArchCam1SM
- Volume Type: Flexible
- Language: English US
- Space Guarantee: volume
- Volume Size Type: Total Size
- Volume Size: 18 GB
- Snapshot Reserve: 5%

6. Specify the destination system on the source filer to authorize the destination filer to access data. We used the `snapmirror.access` option and entered the command on the source filer, as shown in Example 9-16.

*Example 9-16 Defining accesses*

---

```
itsonas1> options snapmirror.access host=itsonas2
```

---

7. Create `/etc/snapmirror.conf` on the destination filer. If this file already existed, we could use a text editor to modify it. In this file, add the lines to specify which volumes to put in the mirroring configuration. We added the line as shown in Example 9-17 on page 327.

```
itsonas1:VolArchCam1 itsonas2:VolArchCam1SM - sync
```

---

Use the following format to modify the lines to add in this file:

```
source_system:{source_volume | /vol/volume_name/qtree_name}  
dest_system:{dest_volume | /vol/volume_name/qtree_name}  
arguments schedule
```

- source system: This is the name of the storage system from which you are copying data. This name is followed by the name of the volume or qtree that you are copying. Use the volume name alone for volumes. Use the full path name for qtrees.
- dest\_system: This is the host name of the system to which the data is copied, which is followed by the name of the destination volume or qtree to which you are copying data. Use the volume name alone for volumes. Use the full path name for qtrees.
- arguments: These are optional arguments that you can define, such as the transfer speed, for example (kbs), or the TCP window size (wsize). We specified “-” to use default values.
- schedule: Determines the schedule the destination uses to update data. The schedule is mandatory. We used “sync” to indicate that we wanted Synchronous SnapMirror. Read the following note for more information about schedule.

### Schedule option:

The schedule consists of four space-separated fields in the following order:

minute hour dayofmonth dayofweek

- ▶ Minute can be a value from 0 to 59.
- ▶ Hour can be a value from 0 (midnight) to 23 (11 p.m.).
- ▶ Dayofmonth can be a value from 1 to 31.
- ▶ Dayofweek can be a value from 0 (Sunday) to 6 (Saturday).

You can enter multiple values, separated by commas, in any field.

You can apply all possible values for a field with an asterisk (\*). If you specify an asterisk in each field of the schedule, SnapMirror updates the destination every minute.

A single dash (-) in any field means *never* and prevents this schedule entry from executing, which is useful if you want the server to appear in the `/etc/snapmirror.conf` file so that the SnapMirror update can find it, but you do not want the SnapMirror scheduler to run automatically.

Indicate a range of values for any field with a low value and a high value that are separated by a dash, for example, you can indicate that you want an update every hour from 8:00 a.m. to 5:00 p.m by entering this value in the hour field: 8-17.

A range of values followed by a slash and a number indicates the frequency of the update, for example, you can indicate that you want an update every five minutes by entering this value in the minutes field: 0-59/5.

Typing `sync` instead of the four space-separated fields specifies synchronous replication

### Note: Changing the configuration

You can make changes to the schedule, for example, if SnapMirror is enabled, changes to the `/etc/snapmirror.conf` file take effect within two minutes. If SnapMirror is not enabled, changes to the `/etc/snapmirror.conf` file take effect immediately after you type `snapmirror` on the command line to enable SnapMirror.

8. Turn on SnapMirror for both filers using the command in Example 9-18 for the source filer.

*Example 9-18 Turning on SnapMirror*

---

```
itsonas1> options snapmirror.enable on
```

---

9. Restrict the destination volume for SnapMirror to be initialized. We used the command in Example 9-19.

*Example 9-19 Restricting destination volume*

---

```
itsonas2> vol restrict VolArchCam1SM
Volume 'VolArchCam1SM' is now restricted.
```

---

10. Initialize relationship. We used the **snapmirror initialize** command, as shown in Example 9-20. The basic syntax is:
  - - S source system: source volume or qtree (specifies the source system and volume or qtree to copy. This parameter is optional because we already specified the source system in /etc/snapmirror.conf)
  - destination system: destination volume or qtree

**Note:** To specify volume, you just indicate name. For qtrees, you must give the full path name (/vol/...).

*Example 9-20 Initializing the SnapMirror relationship*

---

```
itsonas2> snapmirror initialize -S itsonas1:VolClipCam1
itsonas2:VolArchCam1SM
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.
```

---

While the SnapMirror relationship is initializing, on the destination filer, use the command **snapmirror status** to monitor the system and the first phase of the SnapMirror relationship, transfer of the data, from the source system to the destination system to create a replica. Example 9-21 shows that data is transferred with the quantity already regularly increasing.

*Example 9-21 Monitoring the initialization*

---

```
itsonas2> snapmirror status
SnapMirror is on.
Source           Destination           State           Lag
Status
```

```
itsonas1:VolArchCam1  itsonas2:VolArchCam1SM  Uninitialized  00:09:13
Transferring  (85 MB done)
```

```
itsonas2> snapmirror status
SnapMirror is on.
Source           Destination           State           Lag
Status
itsonas1:VolArchCam1  itsonas2:VolArchCam1SM  Uninitialized  00:09:19
Transferring  (150 MB done)
```

```
itsonas2> snapmirror status
SnapMirror is on.
Source           Destination           State           Lag
Status
itsonas1:VolArchCam1  itsonas2:VolArchCam1SM  Uninitialized  00:09:55
Transferring  (548 MB done)
```

Another interesting command is **ifstat**, as shown in Example 9-22. As shown in Example 9-22, we monitor the transfer rate through the interface that we used for SnapMirror. Examples of interesting information is the transfer rate in Bytes / second.

**Note:** We could verify that the transfer rate of 12MB/sec was normal for our link at 100MB/sec.

*Example 9-22 Monitoring transfer rate*

```
itsonas2> ifstat e0a

-- interface  e0a  (9 days, 5 hours, 17 minutes, 31 seconds) --

RECEIVE
Frames/second:  8118 | Bytes/second:  12321k | Errors/minute:
0
Discards/minute:  0 | Total frames:  5438k | Total bytes:
2433m
Total errors:      0 | Total discards:  0 | Multi/broadcast:
3873k
No buffers:        0 | Non-primary u/c:  0 | Tag drop:
0
Vlan tag drop:     0 | Vlan untag drop:  0 | CRC errors:
0
```

```

Runt frames:      0 | Fragment:      0 | Long frames:
0
Jabber:           0 | Alignment errors: 0 | Bus overruns:
0
Queue overflows:  0 | Xon:           0 | Xoff:
0
Jumbo:           0 | Reset:         0 | Reset1:
0
Reset2:           0
TRANSMIT
Frames/second:    4036 | Bytes/second:    266k | Errors/minute:
0
Discards/minute:  0 | Total frames:    1382k | Total bytes:
161m
Total errors:      0 | Total discards:   0 | Multi/broadcast:
2659
Queue overflows:  0 | No buffers:      0 | Max collisions:
0
Single collision:  0 | Multi collisions: 0 | Late collisions:
0
Timeout:          0 | Xon:           0 | Xoff:
0
Jumbo:           0
LINK_INFO
Current state:     up | Up to downs:     0 | Auto:
on
Speed:            100m | Duplex:          full | Flowcontrol:
full

```

---

After the initial transfer is complete, the two volumes are announced as SnapMirrored. There is no more lag. From this moment, modifications to the source volume are automatically and immediately applied to the destination volume. Example 9-23 shows the SnapMirrored volumes.

#### Example 9-23 SnapMirrored volumes

```

itsonas2> snapmirror status
SnapMirror is on.
Source           Destination           State           Lag
Status
itsonas1:VolArchCam1 itsonas2:VolArchCam1SM Snapmirrored    -
In-sync

```

---

## 9.4 SnapLock

We want the administrator to be able to delete data; therefore, we use **SnapLock Enterprise**. A user with administration privileges on the N series can still delete SnapLock volumes. In strictly regulated environments that require information to be retained for specified lengths of time, such as those governed by SEC Rule 17a-4, SnapLock Compliance is more suited.

During the major part of this book, we used a Gateway filer. Because the Snaplock feature is directly related to disk management, it is not available in a Gateway filer; therefore, we used a non Gateway filer to implement SnapLock.

### 9.4.1 Adding the license

You can add the license through FilerView, which we previously explained in 9.2.1, “Adding a Multistore license” on page 309. You can also add it through the command line interface (CLI), as shown in Example 9-24.

*Example 9-24 Adding License*

---

```
itsotuc4*> license add XXXXXXXX
A snaplock_enterprise 90 day site license has been installed.
SnapLock(tm) Enterprise enabled.
```

---

### 9.4.2 Creating SnapLock volumes

You can create SnapLock traditional volumes or SnapLock FlexVol volumes. We used only FlexVol volumes in this book, therefore, we do not discuss creating Snaplock traditional volumes.

**Important:** To create a SnapLock FlexVol volume, you must create an aggregate with SnapLock as an attribute of that aggregate. Afterwards, FlexVol volumes that are created in this aggregate will be SnapLock volumes.

#### Specifying SnapLock attributes during aggregate creation

The process of creating a SnapLock aggregate is similar to the process of creating a normal aggregate, which we discussed in 7.1.2, “Creating the aggregate” on page 206.



To specify SnapLock attributes while you create aggregates:

1. In FilerView, select **Aggregates** → **Add**. A wizard similar to Figure 9-32 is displayed. Click **Next**.



Figure 9-32 Aggregate Wizard: Add an Aggregate

2. In the next wizard, as shown in Figure 9-33 on page 334, specify the name of the aggregate.

Because we are using a non Gateway filer, we can use the *double parity* option; however, we disabled it because we want to create a little aggregate with only 2 disks.

Select the **snaplock** option.



Figure 9-33 Aggregate Wizard: Aggregate Name and SnapLock option

3. Choose the SnapLock type you want to use to use. In our case, we had only Snaplock Enterprise license installed, as shown in Figure 9-34.

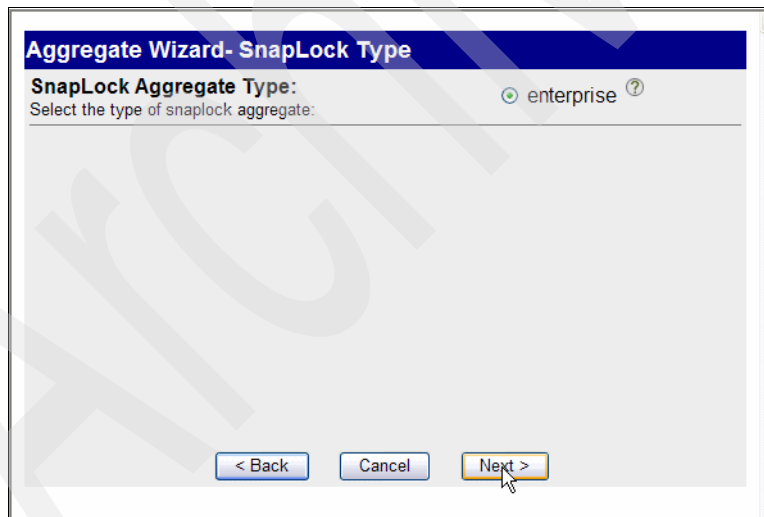


Figure 9-34 SnapLock Type

4. In the next wizard, shown in Figure 9-35 on page 335, define the number of disks in the RAID groups. If you want to use RAID-DP, this is where you decide to use RAID-DP groups of 14 disks, for example, with 12 disks for data

and two disks for parity. In our case, we want to have two disks in the aggregate with RAID4, so we chose **2** disks per RAID group.



Figure 9-35 Aggregate Wizard: RAID Group Size

5. As shown in Figure 9-36, choose automatic or manual disk selection. Automatic disk selection means the disks are selected for you automatically. Manual disk selection means that you can choose the disks yourself. We selected **Manual** disk selection so that we can choose the disks.

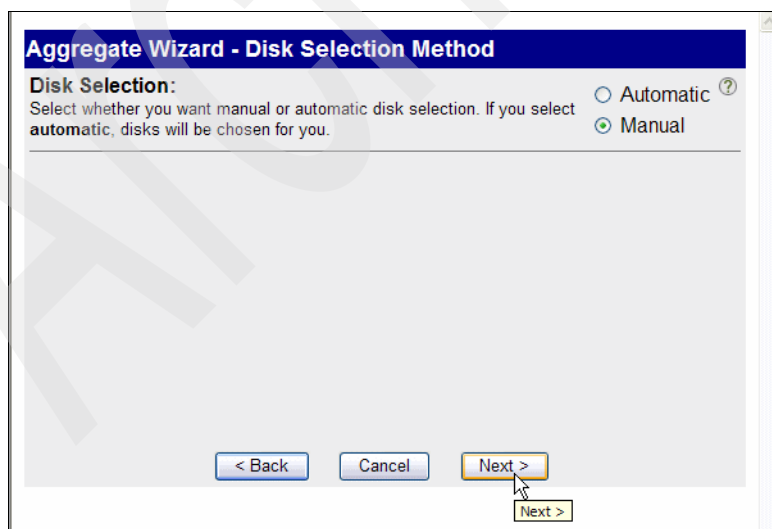


Figure 9-36 Aggregate Wizard: Disk Selection

6. In Figure 9-37, choose the disk that you want to add to the aggregate. In our case, we checked the ID and disk size and then selected the two first disks. Click **Next**. The disks that you select from the list will form the aggregate.



Figure 9-37 Aggregate Wizard: Disks to add

A message is displayed that shows a summary of the new aggregate, as shown in Figure 9-38. Verify the settings, and click **Commit**.

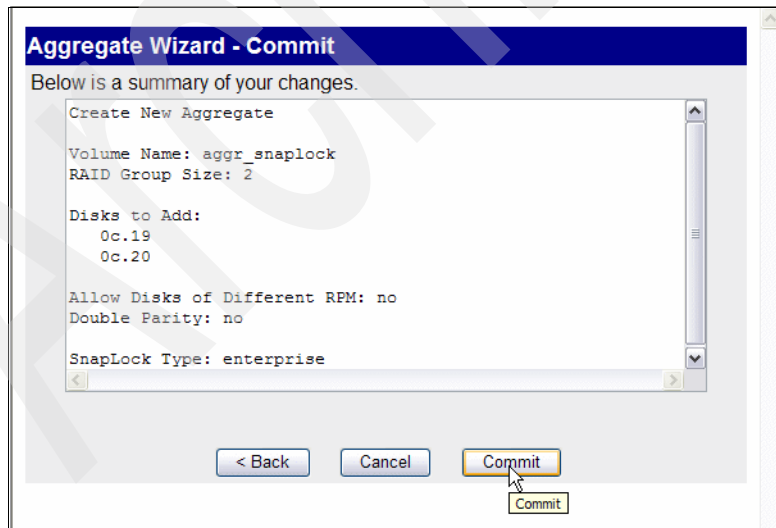


Figure 9-38 Aggregate Wizard: Summary

A status message about the creation of the new aggregate is displayed.

7. Check the status of the aggregate that was just created at the CLI using the **aggr status** command.

As shown in Example 9-25, the aggregate has the option `snaplock_enterprise` and is in the creating state with an initializing status.

Example 9-25 Aggregate status

---

itsotuc4> aggr status			
Aggr State		Status	Options
aggr_snaplock <b>creating</b>		raid4, aggr	raidsize=2,
<b>snaplock_enterprise,</b>		<b>initializing</b>	
snapshot_autodelete=off,			lost_write_protect=off

---

After a few hours, you can use the aggregate, as shown in Example 9-26 with the output of the **aggr status** command.

Example 9-26 Aggregate status 2

---

itsotuc4> aggr status			
Aggr State		Status	Options
aggr_snaplock online		raid4, aggr	raidsize=2,
snaplock_enterprise			

---

## Specifying the SnapLock attribute during volume creation

The process of creating a SnapLock volume is similar to the process of creating a normal volume, which we discussed in 7.1.3, “Creating the volumes” on page 211.

To specify the SnapLock attribute while creating a volume:

1. Under FilerView, click **Volumes** → **Add**. The Volume Wizard is displayed, as shown in Figure 9-39. Click **Next**.



Figure 9-39 Volume Wizard: Add a new volume

2. Choose either a traditional or a flexible volume. We could create a traditional SnapLock volume, and select disks to make it in a future wizard window; however, we want to create a flexible volume in the aggregate that we previously created. We chose **Flexible**, as shown in Figure 9-40 on page 339



Figure 9-40 Volume Wizard: Volume Type Selection

3. In Figure 9-41 on page 340, specify the volume parameters. In the Volume Name field, you must specify a SnapLock Volume; otherwise, if you specify something else, you cannot choose a SnapLock aggregate in which to build the aggregate later on. In the Language field, choose the appropriate language. We did not choose unicode.

**Important:** If you do not specify that you want to build a SnapLock volume, you cannot choose a SnapLock aggregate to build the volume in afterwards. Possible aggregates are in this case only normal aggregates.



The image shows a 'Volume Wizard - Volume Parameters' dialog box. It has a blue title bar. Below the title bar, there are four sections: 'Volume Name:' with a text field containing 'vol\_snaplock' and a help icon; 'Language:' with a dropdown menu set to 'English (US)' and a help icon; 'UTF-8:' with an unchecked checkbox and a help icon; and 'SnapLock Volume:' with a checked checkbox and a help icon. At the bottom, there are three buttons: '< Back', 'Cancel', and 'Next >'. A mouse cursor is pointing at the 'Next >' button, and a tooltip with 'Next >' is visible below it.

**Volume Wizard - Volume Parameters**

**Volume Name:**  
Enter a name for the new volume.  ?

**Language:**  
Select the language to use on this volume.  ?

**UTF-8:**  
Select to make language of this volume UTF-8 encoded. ☐ UTF-8 ?

**SnapLock Volume:**  
Select to create a snaplock volume. ☒ snaplock ?

< Back Cancel Next > Next >

Figure 9-41 Volume Wizard: Volume Parameters

4. Choose the type of SnapLock you want to implement, as shown in Figure 9-42 on page 341.

**Note:**

- ▶ The only choice we had here was SnapLock Enterprise because this was the only license that we had installed.
- ▶ Depending on the type of SnapLock that you choose, you have corresponding choices of aggregates afterwards in which to build your volume.





Figure 9-42 SnapLock Volume Type

5. Choose the Containing Aggregate that your volume will reside in and the space guarantee, as shown in Figure 9-43 on page 342.

There are three types of space guarantee:

- A space guarantee of *volume* preallocates space in the aggregate for the volume. The preallocated space cannot be allocated to any other volume in that aggregate. The space management for a FlexVol volume that has a space guarantee of volume is equivalent to a traditional volume.
- A space guarantee of *file* preallocates space in the aggregate so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are pinned for a Snapshot copy.
- A FlexVol volume that has a space guarantee of *none* reserves no extra space, which means that writes to LUNs or files that are contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

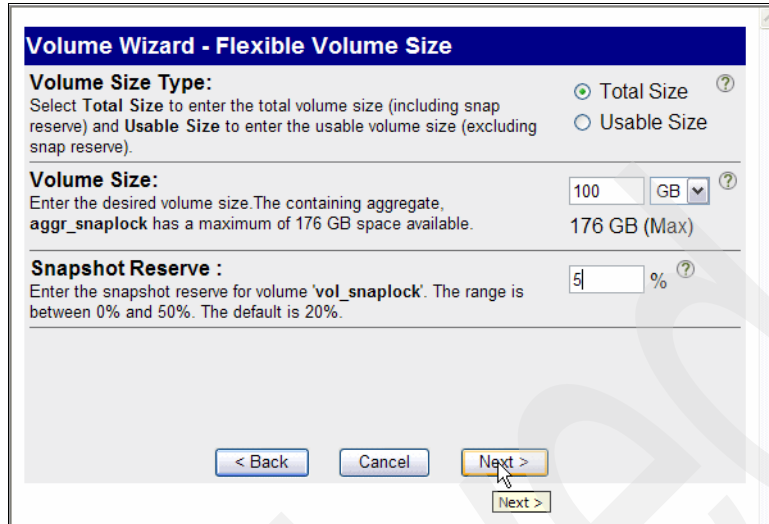
We chose **Volume** as our space guarantee type for our volumes.

**Note:** Because we chose to create a Snaplock Enterprise Flexible Volume, we had only Snaplock Enterprise aggregates to choose between.



Figure 9-43 Volume Wizard: Volume Parameters

6. Choose the size of the volume, as shown in Figure 9-44 on page 343. Here you define how much space you want for the volume, and how much space in percentage will be allocated for Snapshot, which includes the size you enter (Total size option) or added to the size which will be really usable (Usable size option).
7. Choose the Snapshot reserve. We chose only 5% (default is 20%) because DVS data is not normally modified and users do not need often need to return to a previous state. In DVS context, recovery is needed more in disaster recover situations, and to prevent this type of failure from impacting the infrastructure, a strategy like SnapMirror is more suited than Snapshot.



**Volume Wizard - Flexible Volume Size**

**Volume Size Type:**  
 Select **Total Size** to enter the total volume size (including snap reserve) and **Usable Size** to enter the usable volume size (excluding snap reserve).  
☒ Total Size ?  
☐ Usable Size

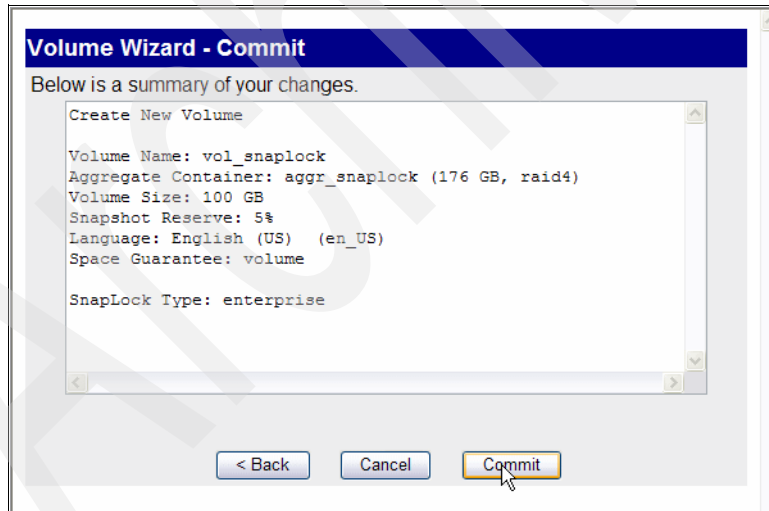
**Volume Size:**  
 Enter the desired volume size. The containing aggregate, **aggr\_snaplock** has a maximum of 176 GB space available.  
 100 GB ?  
 176 GB (Max)

**Snapshot Reserve :**  
 Enter the snapshot reserve for volume 'vol\_snaplock'. The range is between 0% and 50%. The default is 20%.  
 5% ?

< Back    Cancel    **Next >**  
 Next >

Figure 9-44 Volume Wizard: Flexible Volume Size

8. FilerView displays a summary similar to Figure 9-45 about the volume that will be created. Select **Commit**.



**Volume Wizard - Commit**

Below is a summary of your changes.

```

Create New Volume

Volume Name: vol_snaplock
Aggregate Container: aggr_snaplock (176 GB, raid4)
Volume Size: 100 GB
Snapshot Reserve: 5%
Language: English (US) (en_US)
Space Guarantee: volume

SnapLock Type: enterprise
  
```

< Back    Cancel    **Commit**

Figure 9-45 Volume Wizard: Summary

9. When you receive a message stating that the Volume updated successfully, as shown in Figure 9-46 on page 344, click **Close** to exit the wizard.

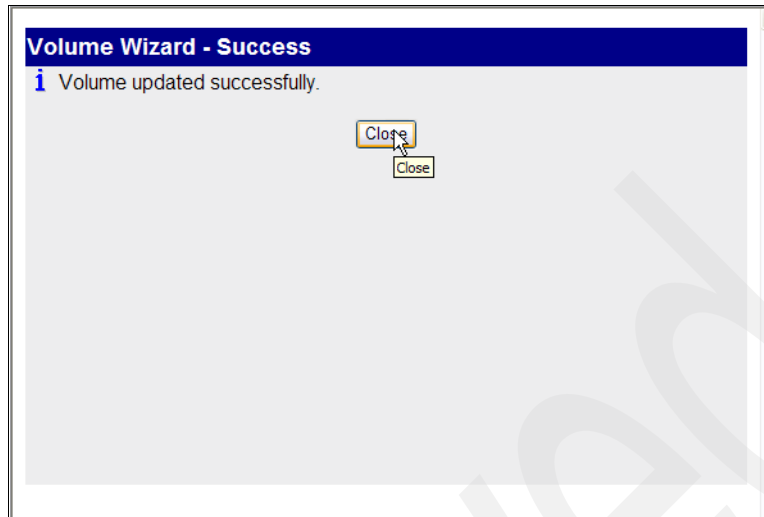


Figure 9-46 Confirmation of the volume creation

### 9.4.3 Managing WORM files

After you place a file into a SnapLock volume, you must explicitly commit it to a WORM state before it becomes WORM data, by transitioning the file from a writable state to a read-only, while in the SnapLock volume.

Data can be transitioned to WORM state **interactively** or **automatically**.

**Important:** Be careful to synchronize the date and time of the server and of the N series. In case of a WORM volume, this always the N series time which will be considered.

#### Interactively changing data to WORM state

Let us use the case where we have our WORM volume mounted under /mnt.

To interactively change data to WORM state:

1. Create a file in this volume, as shown in Example 9-27. We created a file called archive.

*Example 9-27 We created a file in the WORM volume*

---

```
[root@itsohelvio mnt]# touch archive
```

---

Our file, archive, is currently in read write mode. In Example 9-28, there are three important things:

- Last accessed date is 2007-09-27 15:50:03
- The file is writable, which is shown by the result of the command **echo**
- Date/time considered is not the time of the server but the time of the N series (this is a point we verified on the filer).

*Example 9-28 Status of the file we created*

---

```
[root@itsohelvio mnt]# stat archive
  File: `archive'
  Size: 0                Blocks: 0          IO Block: 32768   Regular
File
Device: ah/10d  Inode: 2548897      Links: 1
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/
root)
Access: 2007-09-27 15:50:03.000000000 -0700
Modify: 2007-09-27 15:50:03.000000000 -0700
Change: 2007-09-27 15:50:03.000000000 -0700

[root@itsohelvio mnt]# echo "test" > archive

[root@itsohelvio mnt]# date
Thu Sep 27 15:42:53 MST 2007
```

---

In Example 9-28, the last accessed time stamp of the file at the time it is committed to WORM state becomes its retention date.

2. Because we are in a Linux environment, use the two following commands to put the file in a WORM state, with a retention date of 15:30.

```
touch -a -t 200709271552 archive
chmod -w archive
```

As shown in Example 9-29, our file is now announced as a read-only file, with the last access time stamp becoming the retention date.

*Example 9-29 Putting the file in WORM state*

---

```
[root@itsohelvio mnt]# touch -a -t 200709271552 archive
[root@itsohelvio mnt]# chmod -w archive

[root@itsohelvio mnt]# stat archive
  File: `archive'
```

```
Size: 5          Blocks: 0          IO Block: 32768  Regular
File
Device: ah/10d  Inode: 2548897      Links: 1
Access: (0444/-r--r--r--)  Uid: (   0/   root)   Gid: (   0/
root)
Access: 2007-09-27 15:52:00.000000000 -0700
Modify: 2007-09-27 15:50:38.000000000 -0700
Change: 2007-09-27 15:50:56.000000000 -0700
```

---

3. Use the syntax in Example 9-30 to verify that the file is in WORM test.

*Example 9-30 Verification of the WORM state*

---

```
[root@itsohelvio mnt]# chmod 777 archive
chmod: changing permissions of `archive': Read-only file system

[root@itsohelvio mnt]# echo "test" > archive
-bash: archive: Permission denied

[root@itsohelvio mnt]# rm -f archive
rm: cannot remove `archive': Read-only file system
```

---

## Automatically changing data to WORM state

The feature *Auto-commit time delay* automatically converts files on the volume to WORM status without involving the application that created the files, which permits existing applications to be used with SnapLock without any application changes. You can set the global option, `snaplock.autocommit_period`, to specify the delay. If the file does not change during the delay period, the file is committed to WORM at the end of the delay period. The minimum delay that can be specified is two hours. Auto-commit does not take place instantly when the delay period ends. Auto-commits are performed using a scanner and can take some time.

To automatically change data to WORM state using Auto-commit:

1. Enter this command to set the auto-commit time delay:  
`options snaplock.autocommit_period none|(count|h|d|m|y)`

The `autocommit_period` is the time delay in hours, days, months, or years. We defined it as shown in Example 9-31 on page 347.

### Example 9-31 Setting auto-commit time delay

---

```
itsotuc4> options snaplock.autocommit_period 1m
```

You are changing option `snaplock.autocommit_period` which applies to both members of

the cluster in takeover mode.

This value must be the same in both cluster members prior to any takeover

or giveback, or that next takeover/giveback may not work correctly.

Thu Sep 27 16:14:42 MST [itsotuc4: waf1.scan.start:info]: Starting autocommit volume scan on volume `vol_snaplock`.

Thu Sep 27 16:14:42 MST [itsotuc4: reg.options.cf.change:warning]: Option `snaplock.autocommit_period` changed on one cluster node.

---

2. Create a file in the WORM volume. We created the file `autocommit`. This file is, by default, read-writable, as shown in Example 9-32.

### Example 9-32 Creating a file in the WORM volume

---

```
[root@itsohelvio mnt]# touch autocommit
[root@itsohelvio mnt]# date
Thu Sep 27 16:08:00 MST 2007
[root@itsohelvio mnt]# stat autocommit
  File: `autocommit'
  Size: 0             Blocks: 0          IO Block: 32768   Regular
File
Device: ah/10d  Inode: 2548898      Links: 1
Access: (0644/-rw-r--r--)  Uid: (   0/   root)   Gid: (   0/
root)
Access: 2007-09-27 16:15:15.000000000 -0700
Modify: 2007-09-27 16:15:15.000000000 -0700
Change: 2007-09-27 16:15:15.000000000 -0700
```

---

## About volume retention periods

There are three volume retention periods that you can specify:

- ▶ **Minimum retention period:** this is the shortest amount of time a WORM file can be set in a SnapLock volume. You set the minimum retention period to ensure that applications or users do not assign retention periods that do not conform with customer minimum retention period values.
- ▶ **Maximum retention period:** this is the longest amount of time that the WORM file can be set in a SnapLock volume. You set the maximum retention period to ensure that applications or users do not assign excessive retention periods that do not conform with customer maximum retention period values, or you can set the maximum retention period to infinity.

- ▶ Default retention period: specifies the retention period that is assigned to any WORM file on the SnapLock volume that was not explicitly assigned a retention period.

To specify retention periods, use the following commands:

- ▶ `vol options vol_name snaplock_minimum_period [period | infinite]`
- ▶ `vol options vol_name snaplock_maximum_period [period | infinite]`
- ▶ `vol options vol_name snaplock_default_period [period | min | max | infinite]`

In the commands:

- *vol\_name* is the SnapLock volume name.
- *period* is the retention period, specified by a numeral, followed by days (d), months (m), or years (y). Alternatively, you can specify infinite, meaning that all files on the volume have infinite retention. See the `na_vol(1)` man page for details.



# Administering DVS data with N series

In this chapter, we provide information about how to manage DVS data on N series. After we take you through the initial configuration, you can access the N series storage system in one of the three ways, as seen in Figure 10-1 on page 350:

- ▶ FilerView
- ▶ Operations Manager
- ▶ Command Line Interface:
  - Secure Shell (SSH)
  - Telnet
  - Serial Console Access

We introduce these methods of managing DVS data.

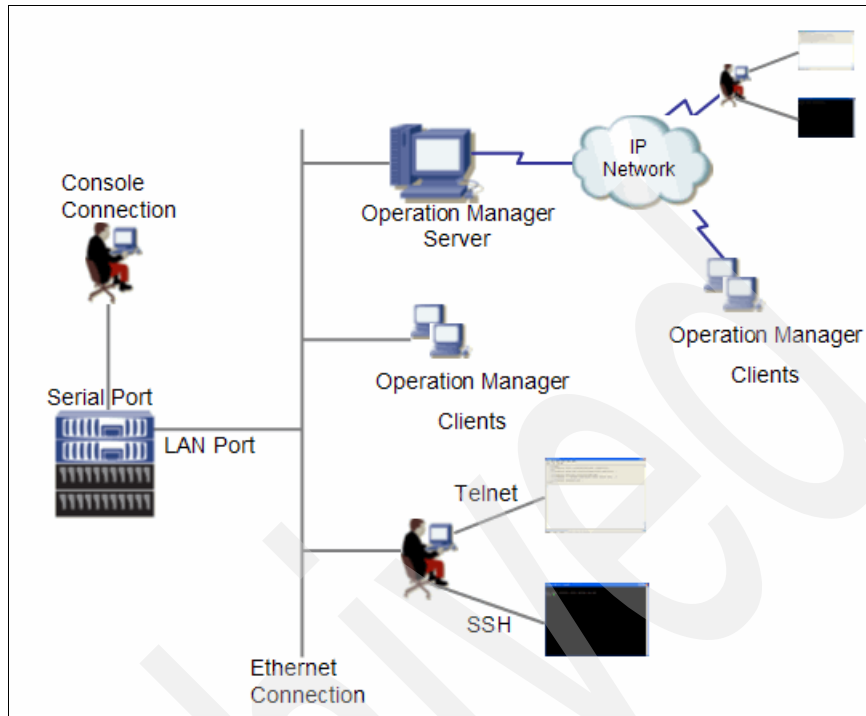


Figure 10-1 N series access options

## 10.1 FilerView

FilerView is a graphical user interface (GUI) management interface that enables you to manage N series storage system from a Web browser rather than from the Command Line Interface (CLI). Using the FilerView tool, you can view information about the storage system and its physical storage units, such as adapters, disks, and RAID groups. You can also view information about the storage system's logical storage units, such as aggregates, volumes, and LUNs. Additionally, you can view statistics about network traffic and the performance of the filer. The administrator can access FilerView remotely by Web browser to manage the N series storage system.

### 10.1.1 Before you use FilerView

Before you can access the FilerView, you must complete the follow tasks:

- ▶ Make sure that your browser is supported: Microsoft Internet Explorer 6.x or later or Mozilla Firefox 1.5.x or later.<sup>1</sup>
- ▶ Make sure that your browser has Java and JavaScript enabled. If your system does not include Java support, you must download a Java run-time environment separately to ensure that FilerView functions properly.
- ▶ Make sure that your platform is in the FilerView support list:
  - Windows Server® 2003
  - Windows XP
  - Solaris 9
  - Solaris 10
  - Linux AS V3
  - Linux AS V4
  - Linux ES V4
  - SUSE Linux 9.0
- ▶ The following options control access to FilerView:
  - options httpd.admin.access<sup>2</sup>
  - options httpd.admin.enable<sup>3</sup>
  - options admin.ssl.enable<sup>4</sup>

### 10.1.2 Accessing storage system using FilerView

To access a storage system from a client by using FilerView:

1. Start your Web browser.
2. Access FilerView by URL: **http://filename\_OR\_filerIP/na\_admin.**
3. When the authentication window is displayed, type your username and password, as shown in Figure 10-1 on page 350.


---

<sup>1</sup> Other browsers that support Java and JavaScript™ might also be compatible with FilerView.

<sup>2</sup> Restricts HTTP access to FilerView, the administration area of the filer, through a private IBM URL: any URL beginning with /na\_admin. If this value is set, trusted.hosts is ignored for FilerView access. Default value is legacy.

<sup>3</sup> Enables HTTP access to FilerView, the administration area of the filer, using a private IBM URL: any URL beginning with /na\_admin is mapped to the directory /etc/http. Thus, you can access a man page on the filer toaster with the file name /etc/http/man/name with the URL http://toaster/na\_admin/man/name. Default value is on.

<sup>4</sup> Enables HTTPS access to FilerView.



*Figure 10-2 Authentication window for FilerView*

When you successfully authenticate, you can access the FilerView, as shown in Figure 10-3 on page 353.

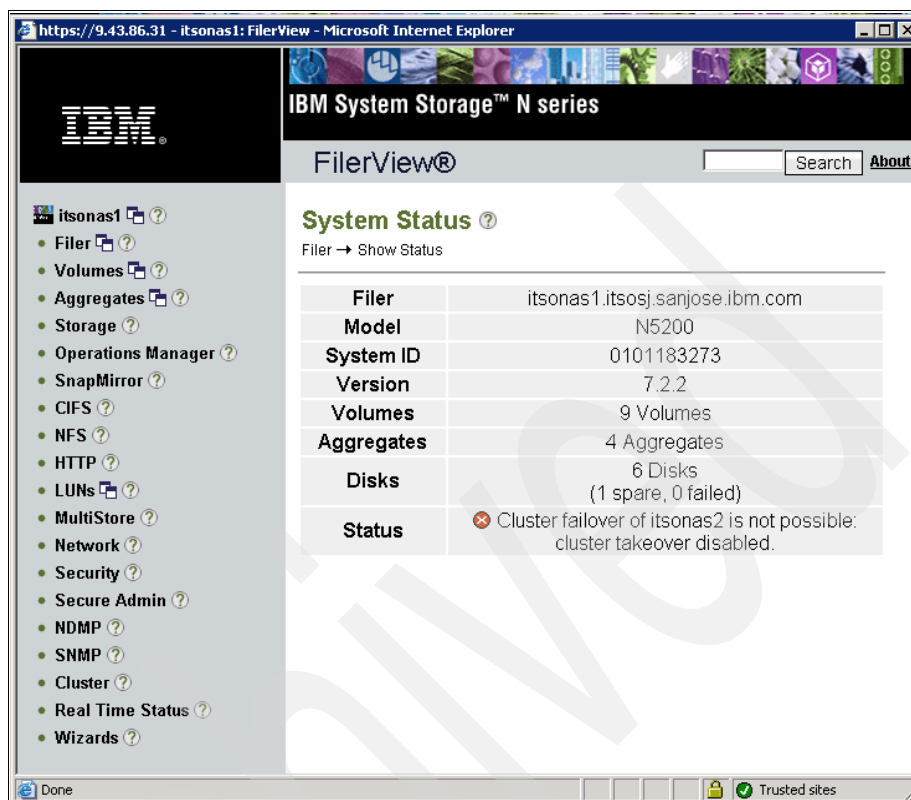


Figure 10-3 FilerView

### 10.1.3 Managing storage system with FilerView

You can do most maintenance work for N series Storage System in FilerView. The FilerView interface consists of three main frames, as shown in Figure 10-4:

- ▶ Left frame: This is the Topic frame, which contains manageable items.
- ▶ Right frame: This is the Result frame that displays a result based on the selection you make in the left frame.
- ▶ Title frame: Contains the name and functions that you select from the left frame.

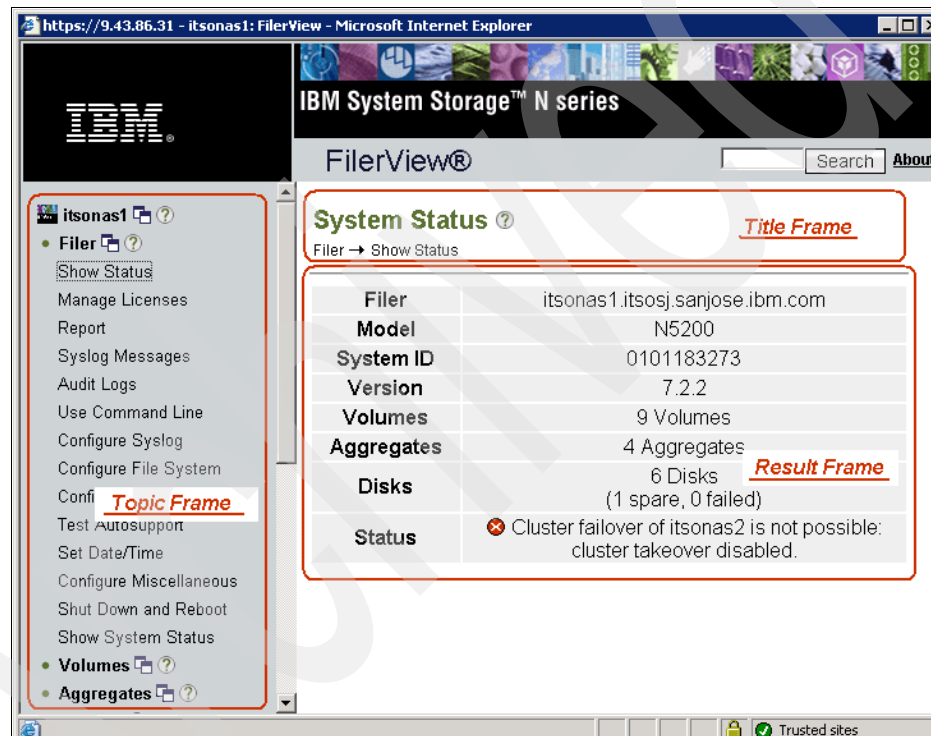


Figure 10-4 FilerView Function model

On the Topic frame of the FilerView interface, the following items are mostly used:

- ▶ Filer: performs some maintenance work on the storage system, such as add licenses, view storage system's status, view logging information, set time and date on storage system, use CLI function in FilerView, and reboot or shutdown storage system. Figure 10-5 on page 355 give you a glance of the 'Filer' item.

- ▶ **Volumes:** manages the logical storage resource independently of the physical layer of storage. This item performs volume and qtree's management and their quota's management.
- ▶ **Aggregates:** containers that capture all of the physical aspects of storage: disks, RAID groups, and plexes. This item performs aggregate management.
- ▶ **Storage:** configures storage disks and view report of disks, adapters, fabric connection, and hub connection.
- ▶ **CIFS:** Windows clients use the CIFS protocol to access data on a storage system. Using CIFS management pages, makes it easy to configure, set up, and manage CIFS on a storage system.
- ▶ **LUNs:** a unit of storage on a storage system that clients access. LUN management includes: LUN setup, FCP/iSCSI related function setup, initiator group management, LUN management, and FCP management.
- ▶ **Cluster:** a cluster consists of a pair of storage systems that are connected through an interconnect adapter and cables, and are configured so that both storage systems share access to a set of Fibre Channel disks, subnets, and tape drives to provide fault tolerance. We can use this item to enable/disable a takeover or to initiate a takeover/giveback.
- ▶ **Real Time Status:** offers the health monitor, performance meter show, and traffic statistics show on the storage system.
- ▶ **Wizards:** is a page-based program that takes you through the necessary steps to accomplish a task.

Figure 10-5 is an example of the Filer in FilerView.

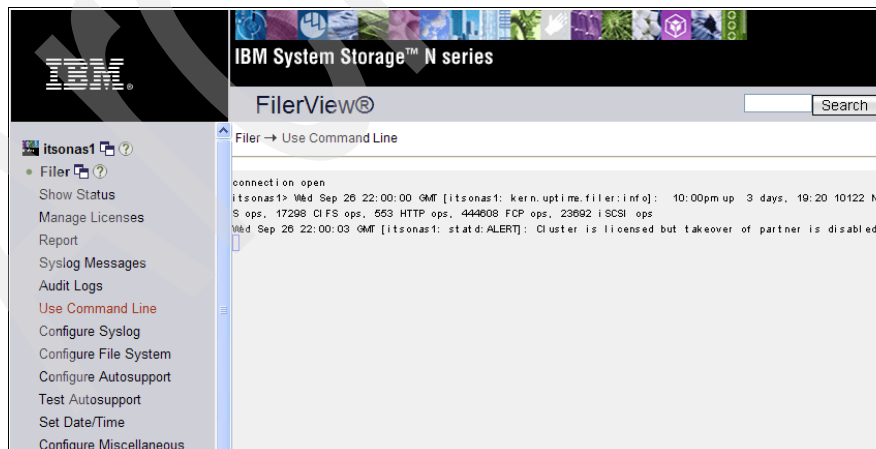


Figure 10-5 Filer item in FilerView

Under each major item, there are some sub-items. Three of them are more widely used, as shown in Figure 10-6:

- ▶ Add: an 'add object' task in FilerView with interactive method.
- ▶ Manage: modify related objects with interactive method.
- ▶ Report: view the status of selected objectives.

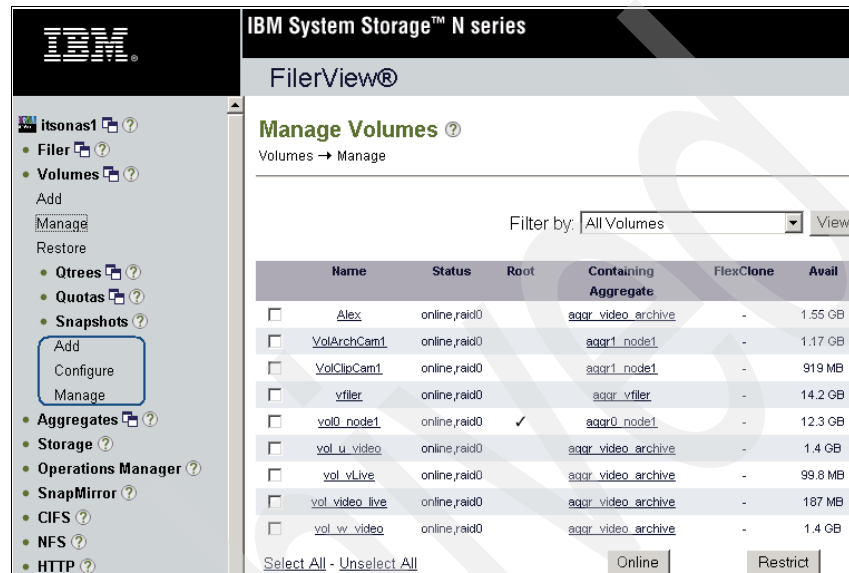


Figure 10-6 More often used sub-items

## 10.1.4 Commonly used items in FilerView

Here are the most used items in the FilerView in the DVS management work:

- ▶ **Filer → Show Status:** Displays information about the storage system and its status.
- ▶ **Filer → Licenses:** The storage system requires a software license to enable most services. This page enables you to manage license codes for many of the services on your storage system.
- ▶ **Filer → Report:** Displays the current storage system configuration information.
- ▶ **Filer → Syslog Messages:** Displays information and error messages that the storage system displays on the console and logs in the /etc/messages file.
- ▶ **Filer → Use Command Line:** FilerView cannot accomplish every task on storage system. Sometimes you have to achieve your goal using the command line method. Use this page to enter storage system commands.



- ▶ **Volume → Add:** Use to create both flexible volumes and traditional volumes.
- ▶ **Volume → Manage:** Use this page to modify volume-related parameters and statuses.
- ▶ **Volume → Qtrees:** A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within either a traditional volume or a flexible volume. This page enables you to add or manage a new qtree to a volume on your storage system.
- ▶ **Volume → Quotas:** A disk quota is a set of rules that restrict disk space and the number of files that a user or a group uses. Use this page to buildup quotas' rule and add or manage a quotas setting.
- ▶ **Volume → Snapshots:** A Snapshot is a read-only image of a unit of data. This page enables you to add or manage Snapshots for a specified volume.
- ▶ **Aggregates → Add:** Aggregates are containers that capture all of the physical aspects of storage: disks, RAID groups, and plexes. Use this page to add aggregates in an interactive method.
- ▶ **Aggregates → Manage:** Displays information about aggregates on the storage system.
- ▶ **Storage → Disk → Manage:** Displays information about the storage system disks, including type (parity, data, or hot spare), size, and aggregate assignment (if applicable). You can also remove a disk and fail a disk here.
- ▶ **Storage → Adapter:** Displays the current data about all adapters that are installed in the storage system.
- ▶ **CIFS → Configure → Setup Wizard:** Helps you set up and configure CIFS in an interactive way. It is the same as **Wizard → CIFS Setup Wizard**.
- ▶ **CIFS → Shares → Add:** As the Administrator of the storage system, you can create directories on the storage system. However, these directories do not automatically become accessible to users. You must create shares that correspond to these directories so that users can share them. This page enables you to add new CIFS shares on the storage system.
- ▶ **CIFS → Shares → Manage:** Enables you to manage the shares that are currently on the storage system.
- ▶ **LUNs → Wizard:** Helps you to create and map new LUNs and create new initiator groups. It is the same as **Wizard → LUN Wizard**. Figure 10-7 on page 358 is a LUN Wizard Welcome page in FilerView.
- ▶ **LUNs → Manage:** Enables you to perform LUN management tasks, such as adding LUN, modifying LUN, viewing LUN map, mapping/unmapping a LUN, modifying LUN ID, and deleting a LUN.

- ▶ **LUNs → Initiator Groups:** An initiator group (igroup) specifies which initiators can have access to a LUN. This page enables you to add (create) or manage initiator groups on your storage system.
- ▶ **LUNs → FCP:** Fibre Channel is a common, efficient transport system that supports multiple protocols or raw data using native Fibre Channel guaranteed delivery services. Use this page to view information, such as adapter name, status, and statistics about the FCP adapters on your storage system. You can also modify the status and media type of these adapters.
- ▶ **LUNs → iSCSI:** iSCSI is a transport protocol that allows standard SCSI block access over a TCP/IP network. Use this page to manage iSCSI-related parameters and settings.
- ▶ **Wizards → Setup Wizard:** A wizard is a page-based program that takes you through the necessary steps to accomplish a task. The Setup Wizard helps you get your storage system up and running. You can also use this wizard to change or view the current setup of the storage system.
- ▶ **Wizards → CIFS Setup Wizard:** Helps you set up and configure CIFS in an interactive way. It is the same as **CIFS → Configure → CIFS Setup Wizard**.
- ▶ **Wizards → LUN Wizard:** Helps you to create and map new LUNs and create new initiator groups, as shown in Figure 10-7. It is the same as **LUNs → Wizard**.

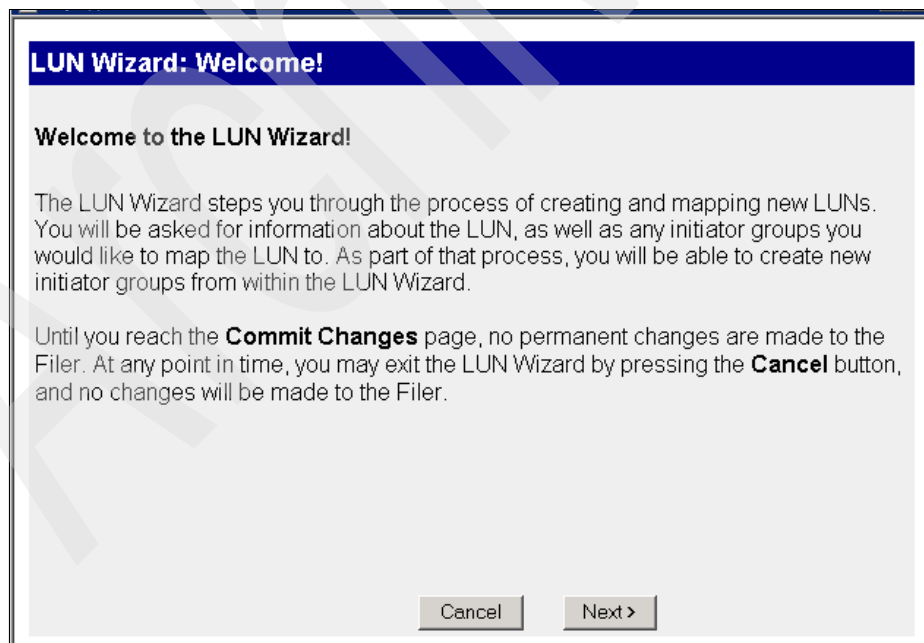


Figure 10-7 LUN Setup Wizard

## 10.2 Operations Manager

The Operations Manager is an integrated management system that provides infrastructure services, such as discovery, monitoring, role-based access control, auditing, and logging for products in the Storage and Data Suites. The Operations Manager is a Web-based user interface of DataFabric® Manager, which we show in Figure 10-8. You can use it to manage multiple storage systems.

You can do the management from a single location by using different management tools:

- ▶ Operations Manager: A Web-based user interface from which you can monitor and manage multiple storage systems and filer clusters.
- ▶ FilerView: A GUI interface to manage individual storage systems.
- ▶ Performance Advisor: Runs on the IBM N series Management Console and provides a single location to view filer's performance information.
- ▶ Protection Manager: Runs on the IBM N series Management Console and helps to backup and do disaster-recovery operations automatically.
- ▶ Host Agent: An independent software application that resides on a host with which DataFabric Manager interacts.
- ▶ IBM N series Management Console: The client host of the Operations Manager.

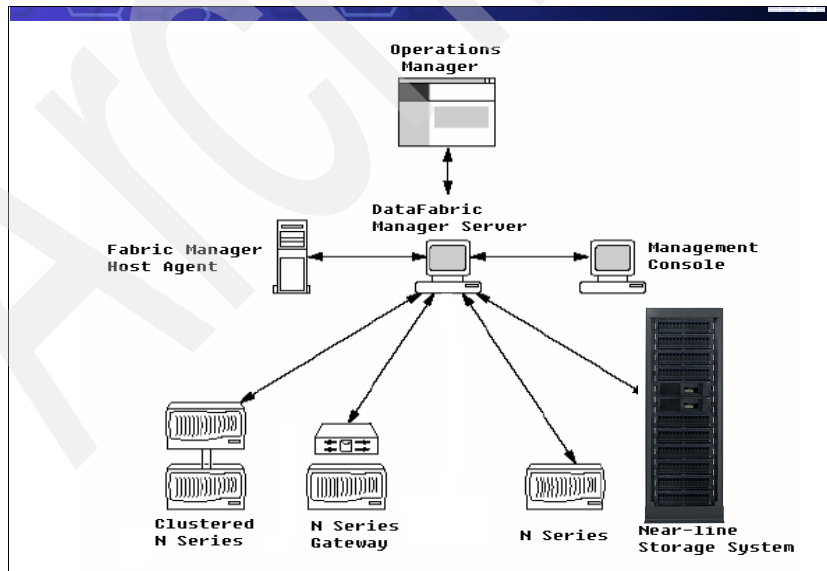


Figure 10-8 Operations Manager in DFM

## 10.2.1 Before you use the Operations Manager

The Operations Manager runs on a separate workstation or server. It does not run on the N series storage system. Prior to installing or upgrading the Operations Manager software, you must ensure that you met requirements in the following areas:

- ▶ Hardware and software requirements that are specified in Table 10-1
- ▶ License requirements
- ▶ Data ONTAP requirement

Table 10-1 Hardware and software requirements

Windows 2000 and Windows 2003 server	
Hardware requirements	Software requirements
<ul style="list-style-type: none"><li>▶ Intel®-based PC with single 2 GHz CPU (Xeon or Pentium® 4)</li><li>▶ 4 GB of free disk space minimum, 8 GB recommended</li><li>▶ 512 MB of memory minimum</li></ul>	<ul style="list-style-type: none"><li>▶ Windows 2000 server (Service Pack 2 or later)</li><li>▶ Windows 2003 server</li></ul>
Solaris 9 and Solaris 10 servers	
Hardware requirements	Software requirements
<ul style="list-style-type: none"><li>▶ Single UltraSPARC III processor at 1.34 GHz (such as Sun™ Fire™ V120 or V240)</li><li>▶ 4 GB of free disk space minimum, 8 GB recommended</li><li>▶ 1 GB of memory minimum</li></ul>	<ul style="list-style-type: none"><li>▶ Solaris 9 SPARC</li><li>▶ Solaris 10 SPARC</li></ul>
Linux workstation or server	
Hardware requirements	Software requirements
<ul style="list-style-type: none"><li>▶ Intel-based PC with single 2 GHz CPU (Xeon or Pentium 4™)</li><li>▶ 4 GB of free disk space minimum, 8 GB recommended</li><li>▶ 512 MB of memory minimum</li></ul>	<ul style="list-style-type: none"><li>▶ Red Hat Enterprise Linux ES version 3 (Update 3 or later) for x86, 32-bit</li><li>▶ Red Hat Enterprise Linux AS version 3 (Update 3 or later) for x86, 32-bit</li><li>▶ Red Hat Enterprise Linux AS Version 4 for x86, 32-bit and 64-bit</li></ul>
Browser	
Hardware requirements	Software requirements
None	<ul style="list-style-type: none"><li>▶ Microsoft Internet Explore 6.0 or later</li><li>▶ Mozilla FireFox 1.0 or later</li><li>▶ Mozilla 1.7 or later</li></ul>

**Note:** Operations Manager 3.5.1 is not supported on Windows NT® 4.0, Windows XP, Solaris 8, or distributions of Linux that are not listed in Table 10-1 on page 360.

### License requirements

You must have a valid Operations Manager license key to complete the Operations Manager installation. After the installation, you can enter additional license keys on the Options page.

**Note:** The Operations Manager core license defines the DataFabric Manager system serial number. All optional licenses must have the same serial number that the Operations Manager license has.

### Data ONTAP requirements

You must be running Data ONTAP version 7.1 with DataFabric Manager 3.3.1 or later.

**Note:** For configuring multiple storage systems across multiple Data ONTAP software versions, an Operations Manager plug-in is required for the version of Data ONTAP that you are running. DataFabric Manager 3.5.1 includes the plug-ins for Data ONTAP 7.1, 7.1.1, and 7.2. You do not need to download a plug-in unless you are using a different version of Data ONTAP.

## 10.2.2 Accessing storage system using the Operations Manager

After being installed, the Operations Manager starts discovering, monitoring, collecting, and saving information about objects in its database. Objects are entities, such as, storage systems the vFiler units, disks, aggregates, volumes, and qtrees that are on these storage systems, LUNs, and user quotas.

You can access the Operations Manager by using the following URL:

`http://server_IP_address_OR_server_dnsname:8080`

After you input the URL, a Welcome page, similar to Figure 10-9 on page 362, is displayed.

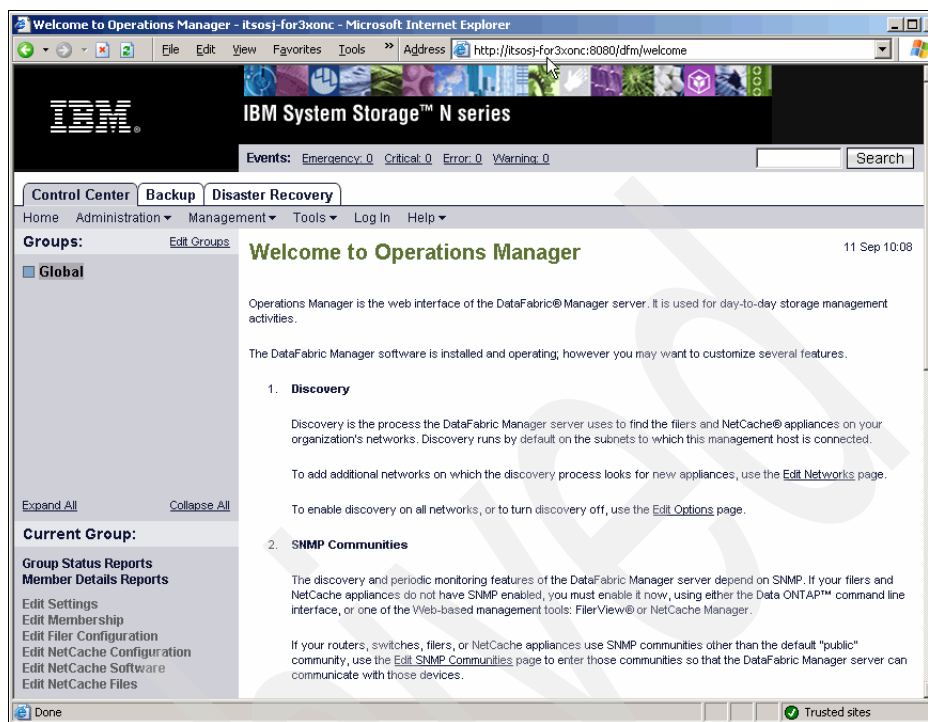


Figure 10-9 Welcome page of the Operations Manager

## 10.2.3 Setting up the Operations Manager

You must do some basic setup on the Operations Manager:

- ▶ IP addresses or names of the storage systems to be discovered by DataFabric
- ▶ Manager (in addition to those it automatically discovers)
- ▶ Set up SNMP communities
- ▶ Set up administrator accounts
- ▶ Configure administrator access control
- ▶ Create groups
- ▶ Configure alarms

See the *N series DataFabric Manager 3.5.1 with Operations Manager Administration* for details about the initial setup. Also, view the following related Web page:

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

## 10.2.4 Managing the storage system with Operations Manager

In this section, we introduce how you can manage the Operations Manager. For more detailed information, read *N series DataFabric Manager 3.5.1 with Operations Manager Administration*. Also, view the following related Web page:

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

We introduce:

- ▶ The page structure of the Operations Manager
- ▶ Logging in to DataFabric Manager
- ▶ Groups and Objects
- ▶ Data collection and event generation
- ▶ Operations Manager usage

### The Operations Manager's page structure

Figure 10-10 shows the Group Summary page of the Operations Manager, which has tabs that support different administrative tasks.

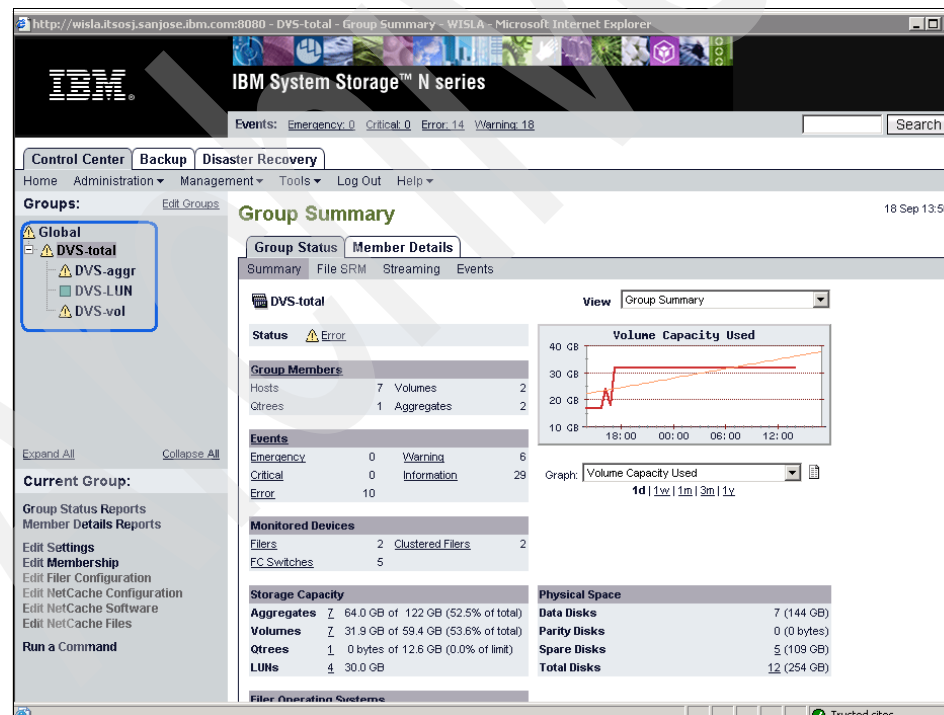


Figure 10-10 Operations Manager's main page

An out-of-space condition with your video surveillance software could lead to a failure. The Operations Manager's reporting and alerts can help you to prevent this condition:

- ▶ **Events Tab:** enables you to view a listing of all current SNMP traps or to sort them by severity. Each view provides information about each SNMP trap, for example, the name of the trap, the severity, and the condition that led to the error.
- ▶ **Main Tabs:** includes the following three tabs:
  - **Control Center:** Allows you to configure and view results for discovery, monitoring, reporting, and alerting for storage systems.
  - **Backup:** Allows users to monitor and manage SnapVault and Open Systems SnapVault disk-to-disk backups.
  - **Disaster Recovery:** Allows users to monitor and manage SnapMirror disk-to-disk mirroring.
- ▶ **Groups Tabs:** located in the left-pane area. We organize objects to a logical group for management reasons, and by default, a group called Global exists in the DataFabric Manager database. All discovered objects belong to this group. We can create a new group and assign objects to this group. The current group is updated with the new group.
- ▶ **Group Status, on the right of the Group Tab:** Used to view groups' status, such as group members, events summary, storage capacity, monitored devices.
- ▶ **Member Details:** Used to view group member's status in detail.

## Logging into the Operations Manager

The Operations Manager uses role-based access control (RBAC) for administrative user login. To log into the Operations Manager, go to the Control Center, shown in Figure 10-11 on page 365, and select **Log In**.



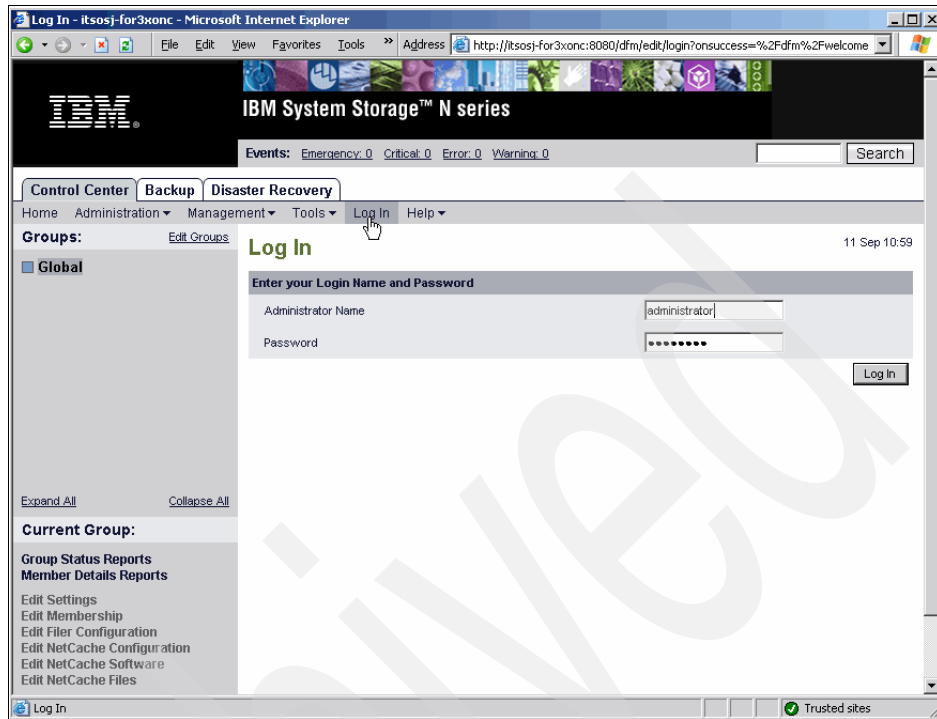


Figure 10-11 Logging into the DataFabric Manager

## Groups and objects

A group is a collection of objects that are related in some way. You can group objects based on characteristics, such as storage system operating system, all storage systems at a location, or all file systems that belong to a specific project or group in your organization. Storage system elements that the Operations Manager monitors, such as storage systems, aggregates, file systems (volumes and qtrees), and logical unit numbers (LUNs), are referred to as objects.

We can create the following type of groups:

- ▶ Appliance Resource group: Contains storage systems, vFiler units, and host agents.
- ▶ Aggregate Resource group: Contains aggregates only.
- ▶ File System Resource group: Contains volumes, qtrees, or both LUN Resource groups, and contains LUNs only.
- ▶ Configuration Resource group: Contains storage systems that are associated with one or more configuration files.

- ▶ Data set: The data stored in a collection of primary storage containers, which includes all of the copies of the data in those containers.
- ▶ Resource pool: A collection of storage objects from which other storage containers are allocated.
- ▶ Storage Resource Manager (SRM) path group: Contains SRM paths only. SRM paths define the location in the file system that is to be indexed for data.

Figure 10-12 shows the groups that we created: DVS-group, DVS-aggr, DVS-LUN, DVS-total, and DVS-vol.

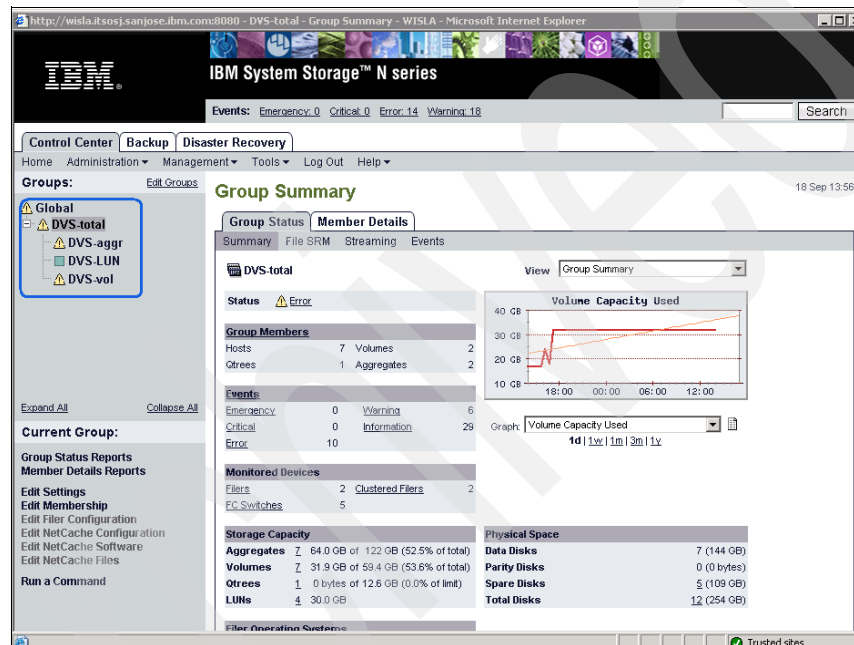


Figure 10-12 User-defined groups

## Collecting data and generating events

The Operations Manager automatically discovers the storage systems on your network. It also periodically monitors data that it collects, such as filer's CPU usage, interface statistics, free disk space, qtree usage, and chassis environmental. The Operations Manager generates events when it discovers abnormalities or when a predefined threshold is crossed. Thresholds can determine at what point you want the Operations Manager to generate an event. In other words, the threshold is a predefined situation or status, such as Volume Full, and if configured, the Operations Manager can send a notification to a recipient when an event triggers an alarm. The flow chart in Figure 10-13 on page 367 illustrates the Operations Manager's monitoring progress.

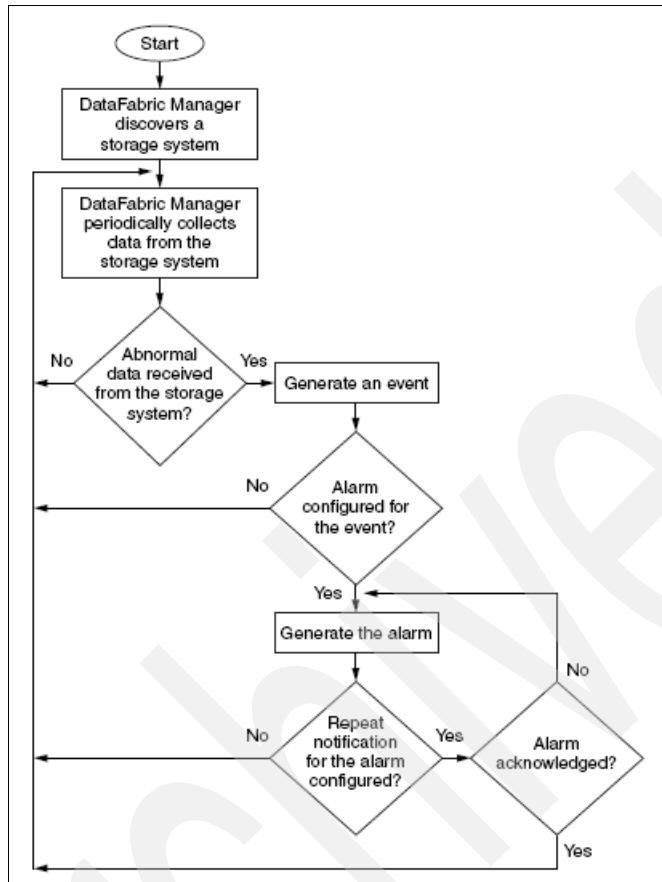


Figure 10-13 Flow chart of DFM monitor process

To configure a threshold:

1. Select **Content Center** → **Tools** → **Options**. The Options page is displayed.
2. Choose **Default Thresholds**. Figure 10-14 on page 368 shows the default threshold on the DataFabric Manager.




Figure 10-14 Default thresholds on DFM

The configurable thresholds include: Host CPU Too Busy, HBA port Too Busy, Aggregate Full, Volume Full, Qtree Full and User Quota, and so on. You can change the default value as required.

All events are associated with one of the following severity levels:

- ▶ **Emergency:** indicates that an object stopped performing unexpectedly and experienced unrecoverable data loss. You must take corrective action immediately to avoid extended downtime.
- ▶ **Critical:** indicates that a problem occurred that might lead to service disruption, if corrective action is not taken immediately.
- ▶ **Error:** indicates that an object is still performing, but corrective action is required to avoid service disruption.
- ▶ **Warning:** indicates that the object experienced an occurrence that you should be aware of. Such events will not cause service disruption, and corrective action might not be required.
- ▶ **Information:** a normal occurrence just like the Normal severity level event, and does not require you to take any action.

- ▶ Normal: indicates that an object is in normal status and is operating within the desired thresholds.

The Operations Manager automatically logs and reports the events; however, you have to manually check the events log. So, if you want a specific severity event that can auto-trigger an alarm to inform the storage administrator about the situation, configure alarms. You can configure the alarm notification to be sent to one or more specified recipients: an email address, a pager number, an SNMP trap host, or a script that you write.

The Operations administrator needs to set up which events trigger alarms, whether the alarm repeats until it is acknowledged and how many recipients an alarm has. When you create alarms, consider the following guidelines:

- ▶ Alarms must be created by group, either an individual group or the Global group.
- ▶ Alarms that you create for a specific event are triggered when that event occurs.
- ▶ Alarms you create for a type of event are triggered when any event of that severity level occurs.
- ▶ Alarms can be for events of severity information or higher.

To configure alarms:

1. Select **Control Center** → **Administration** → **Alarms** as Figure 10-15 on page 370. The Alarms page is displayed, as shown in Figure 10-15 on page 370.




Figure 10-15 Configure alarms in DFM

2. After you configure alarms for a group, specify what triggers the alarm: an event or the severity of event.
3. Specify the recipients of the alarm notification on the Recipients Administrative Users page, which you can get to by clicking **Control Center** → **Administration** → **Administrative Users**, as shown in Figure 10-16 on page 371. On this page, you can configure the recipients name, role, and contact information, such as e-mail address or pager address. In Figure 10-16 on page 371, we can see that two other administrators already exist.




Figure 10-16 Configure recipients in DFM

## Using the Operations Manager

To use the Operations Manager:

1. Define a group that includes objects, such as storage systems, aggregates, file systems (volume and qtree), and LUNs, as shown in Figure 10-12 on page 366.
2. Check the newly defined group into the 'Group Status area' at the left-side of the page. The Operations Manager should automatically discover the device; however, if the Operations Manager does not automatically discover the device, make sure that the Operation Manager's autodiscover is setup by selecting **Control Center** → **Tools** → **Options** → **Discover**.

If the Operations Manager still cannot discover the device, try to discover the device manually. You can add appliances by running the CLI command **dfm host add <host\_IP>**. Normally you only need to add one appliance from each network.

3. Click the group that you defined in step 1, and a Group Summary report is displayed, as shown in, Figure 10-17 on page 372.

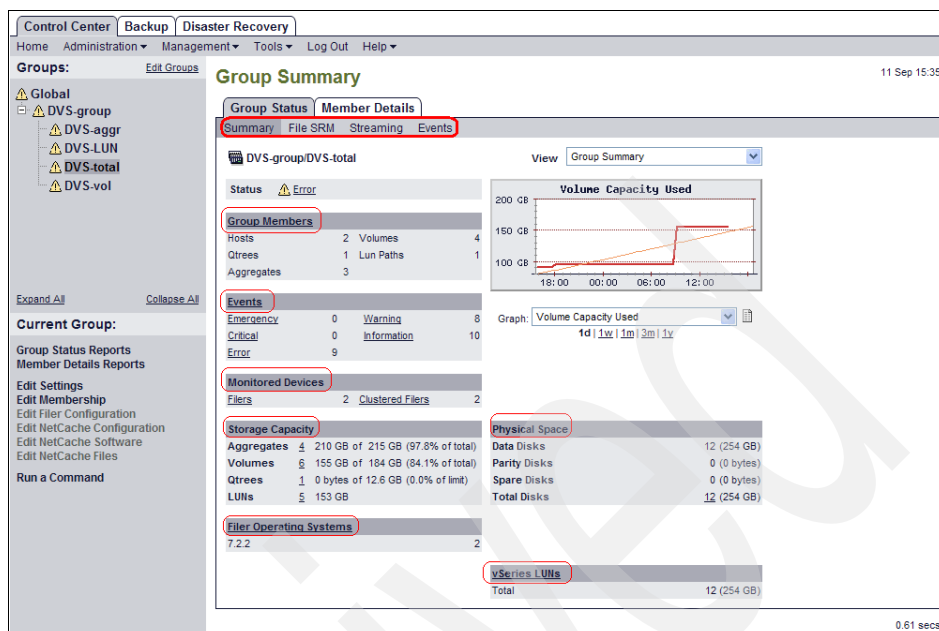


Figure 10-17 Group Summary Report

In the Group Summary Report, shown in Figure 10-17, we can see the Group Members. In this example, we have two hosts, four volumes, one qtree, one LUN Path, and three aggregates. You can also check the Group Members to obtain detail member information.

You can also see the Event report. In this example, there are nine errors, eight warnings, and 10 informations. You can click the type that you want to access to view the event report with more detail event information. Also in the report, the discovered device and its operation system are listed, and most important, the report lists the group's storage information, which includes totally physical storage space, aggregate capacity, which in this example is approaching maximum capacity, volume/qtree/LUN's capacity, and the utilization of this space.

- To see each member's detail report, click the **Member Detail** tab, which is on the top the report, as shown in Figure 10-18 on page 373. Click the object (in Figure 10-18 on page 373, it is Appliance).



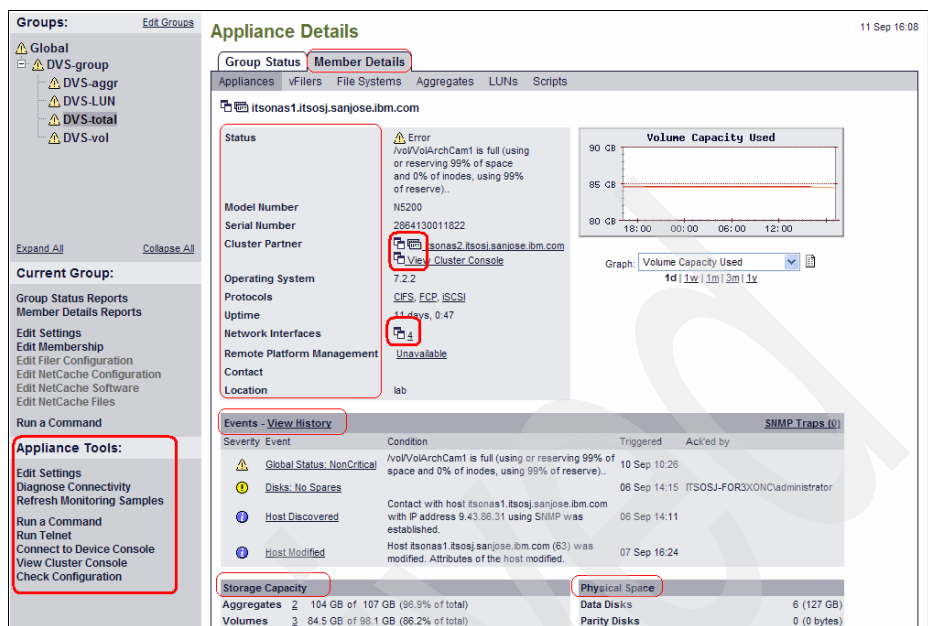


Figure 10-18 Member Detail report

In the Member Detail report, Figure 10-18, you can see the detail information for a specific object (in this example, it is Appliance: itsonas1 in domain itsosj.sanjose.ibm.com).

- The object's information, such as status, model number, serial number, cluster partner, operation system version, supported protocol, uptime, network interface, and so on. You can click the tiny double windows' button (associated with cluster partner and network interface in this example), which calls FilerView to manage related object. You can also click **Protocols** (which is underlined in this example), which supported in this appliance, calls the Related Set up wizard in FilerView, where you can see all events that happened on this object and their severity. Click the specific severity or event to see detailed information. As you can see there is a event generated for the volume full condition. Though non-critical to the N series this could be a catastrophic error to the surveillance application and would be an event that is set up to trigger a alert. Notice the storage information about this object. If you want to see more detailed information, click the number with an underline, and then you can see the name of it (such as, aggr1, vol1). You can go ahead to see detailed information of the new pop-up object (such as aggr1, vol1).

At the left-bottom pane, there is a tools for the object, and in this example it is appliance tools. Here, you can set up Operations Manager for appliance and do some diagnostic testing. You can remotely run a command just like RSH

mode. If the object is an appliance, you can telnet directly to the object or connect to the Console Terminal Server of the appliance. In the 'check configuration' item, you can get an analysis of your appliance configuration.

We discussed the common usage of the Operations Manager, using 'appliance' as a sample object. For other objects, the method should be similar, but might have different reports or settings.

There is one more configurable page that is important, the editing options for the Operations Manager. Normally, the Operations Manager has a default value for these parameters, but under some conditions, you might need to modify some of them, such as SNMP community string. As Figure 10-19 shows, we are modifying an option value: **Control\_Center>Tools>Options>edited-target**.

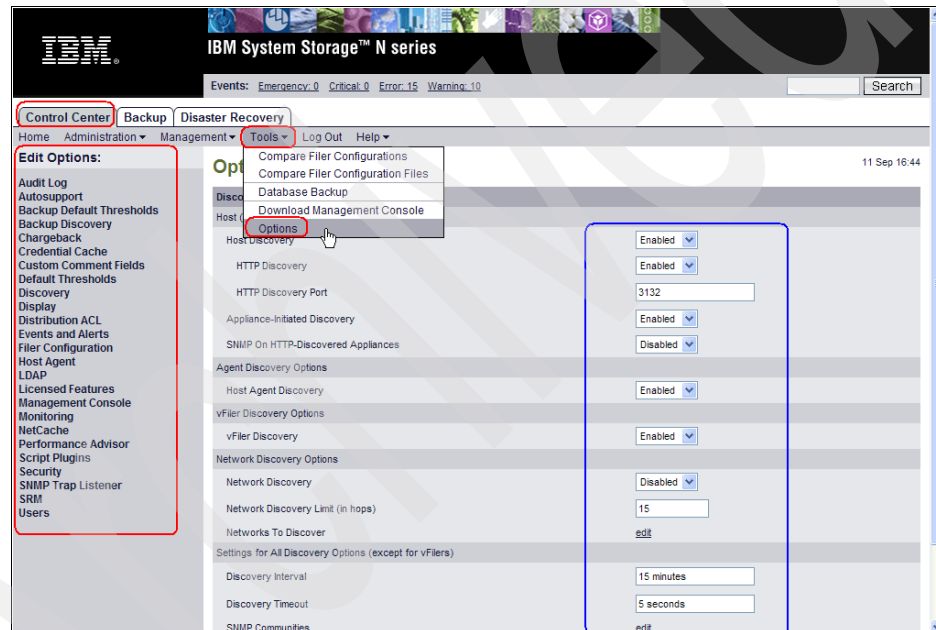


Figure 10-19 Edit options of the Operations Manager

## 10.3 Secure Shell

Secure Shell is widely known as SSH. The traditional network applications, such as FTP and telnet, do not have a security mechanism; therefore, it is not safe to transmit data across the network. SSH was developed to achieve this goal— to make secure data that is being transmitted safer. Even if someone does capture the data, they cannot easily extract the information.

SSH improves storage system security by providing a means to authenticate the client and by generating a session key that encrypts data that is sent between the client and the storage system.

### 10.3.1 Before you use Secure Shell

SecureAdmin uses the following options to enable secure sessions using SSH:

- ▶ options `ssh.passwd_auth.enable`: controls password-based authentication.
- ▶ options `ssh.pubkey_auth.enable`: controls public key authentication.
- ▶ options `ssh.access`: controls access to a storage system.
- ▶ options `ssh.port`: assigns the port number to a storage system.

#### Default ssh values:

- ▶ The default value for `ssh.passwd_auth.enable` and `ssh.pubkey_auth.enable` is On.
- ▶ The default value for `ssh.access` allows everyone to access the storage system.
- ▶ The default value for `ssh.port` is 22.

Next, you need to enable the SSH function on the storage system. We enable it using the FilerView method.

To enable the SSH function on the storage system:

1. In the Topic area of FilerView, as shown in Figure 10-20 on page 376, go to the Secure Admin page, and click the **Generate Keys** button to generate a SSH key.

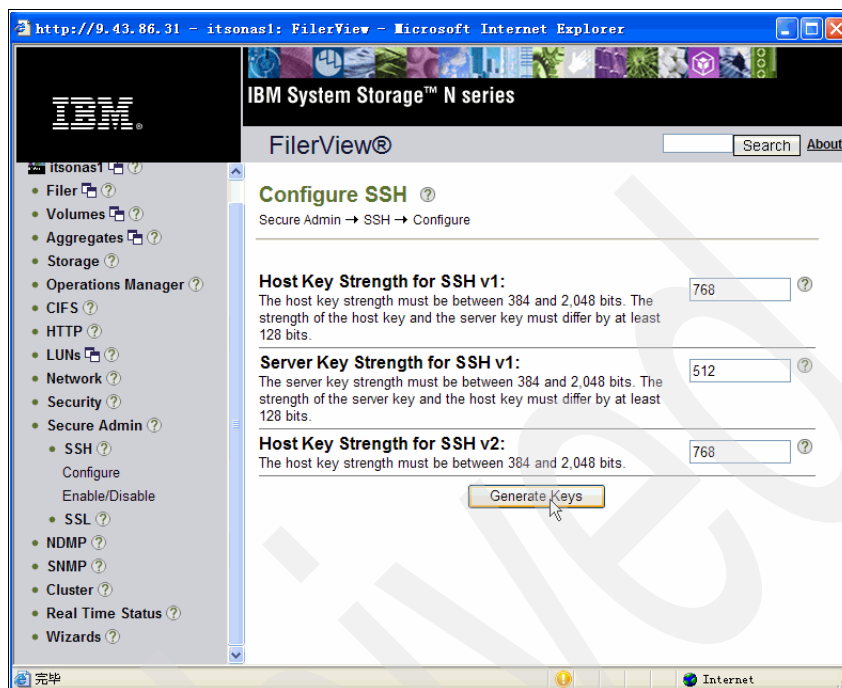


Figure 10-20 Generate a SSH key

2. After you successfully generate a key, go to the Enable/Disable page to enable the SSH function, as shown in Figure 10-21 on page 377.




Figure 10-21 Enable SSH

### 10.3.2 Using Secure Shell

Now we can use SSH protocol to secure communication between the storage system and the management host. We use a popular SSH tool, PuTTY, as an example, which we show in Figure 10-22 on page 378:

1. In the Destination area, enter the IP address (or host name) of the storage system, make sure that the port number is 22, and click **Open**.

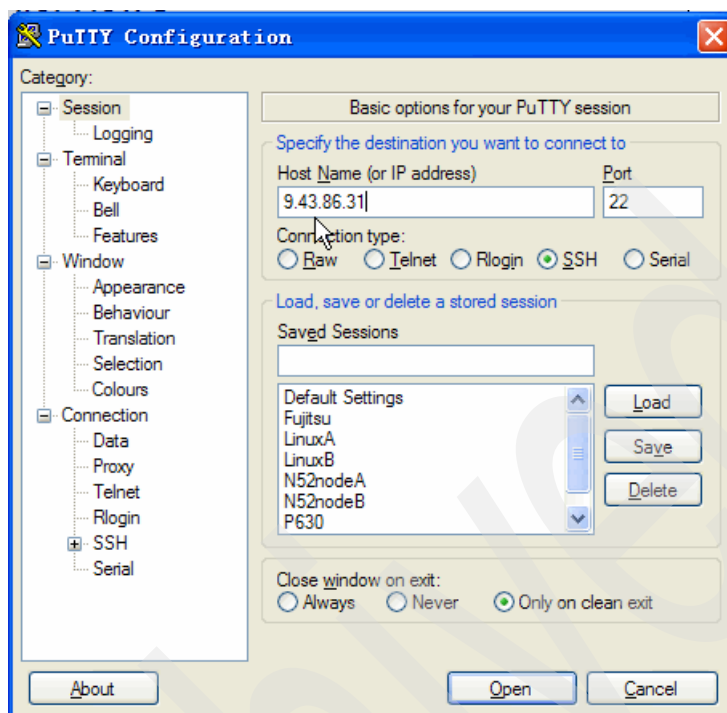


Figure 10-22 Set up the PuTTY

2. A PuTTY security alert window is displayed to query you to accept the new key. Click **Yes** to accept it. Now you have a SSH session, as shown in Figure 10-23 on page 379, connected to the storage system, just like a telnet session, but more secure.



*Figure 10-23 An SSH session*

Figure 10-24 on page 380 is an example of a DVS volume related implementation. We assume that the volume (vol\_video\_live) is getting full (from DFM's event log), and the storage system administrator gets a notification. You can remotely access the storage system using PuTTY and increase the volume size.

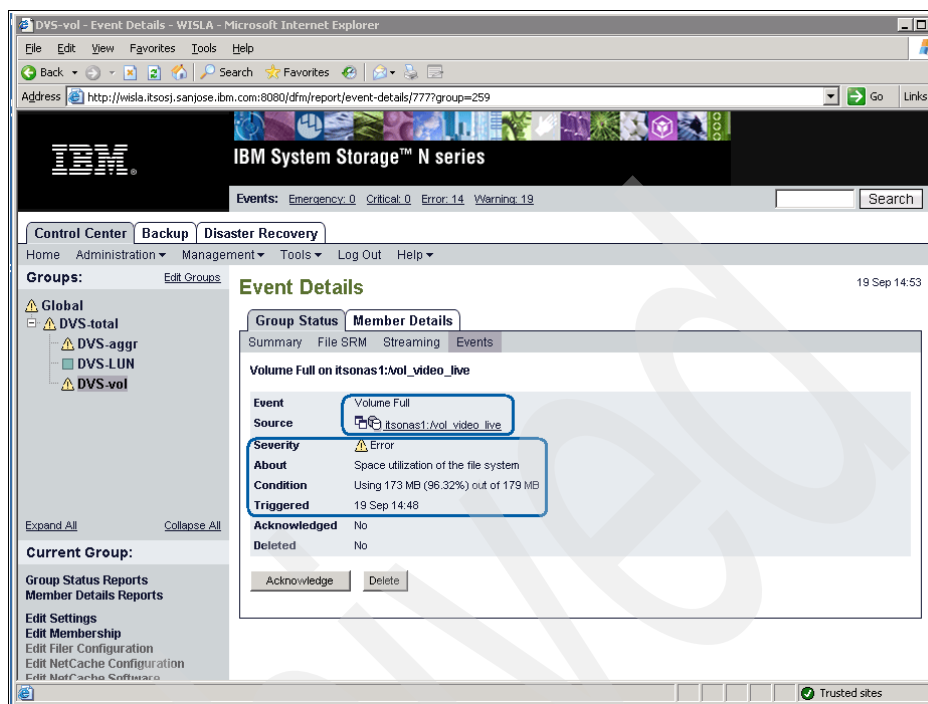


Figure 10-24 An event on DFC that indicates the volume is Full

3. Check the volume utilization, which the event indicated, as shown in Example 10-1.

#### Example 10-1 Volume increase activity

```
itsonas1> df -Vm vol_video_live
Filesystem          total      used      avail capacity  Mounted on
/vol/vol_video_live/ 179MB    172MB        6MB      96% /vol/vol_video_live/
/vol/vol_video_live/.snapshot 44MB        0MB      44MB        0%
/vol/vol_video_live/.snapshot
itsonas1>
```

4. Increase the size of the volume by 100%, as shown in Example 10-2.

#### Example 10-2 Increasing volume size

```
itsonas1> vol size vol_video_live +225m
vol size: Flexible volume 'vol_video_live' size set to 449m.
itsonas1> df -Vm vol_video_live
Filesystem          total      used      avail capacity  Mounted on
/vol/vol_video_live/ 359MB    172MB    186MB      48% /vol/vol_video_live/
/vol/vol_video_live/.snapshot 89MB        0MB      89MB        0%
/vol/vol_video_live/.snapshot
```



```
itsonas1>
```

---

5. You can also increase the size of the aggregate on which the volume resides, as shown in Example 10-3.

*Example 10-3 Volume location*

---

```
itsonas1> df -Ag aggr_video_archive
Aggregate      total      used    avail capacity
aggr_video_archive      18GB       0GB     18GB      2%
aggr_video_archive/.snapshot      0GB     0GB     0GB      0%
itsonas1> aggr add aggr_video_archive 1
Addition of 1 disk to the aggregate has been initiated. The disk needs
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.
itsonas1> Wed Sep 19 22:30:58 GMT [itsonas1: raid.vol.disk.add.done:notice]: Addition
of Disk /aggr_video_archive/plex0/rg0/itsosan02:24.126L2 Shelf - Bay - [IBM
1742-900 0000] S/N [600A0B80001742330000006846C49D7F] to aggregate
aggr_video_archive has completed successfully
itsonas1> df -Ag aggr_video_archive
Aggregate      total      used    avail capacity
aggr_video_archive      36GB       0GB     36GB      1%
aggr_video_archive/.snapshot      1GB     0GB     1GB      0%
itsonas1>
```

---

You can now manage the storage system using the Command Line method.

In the Command Line Mode, use the **help** command to show every command that you can execute in the storage system, as shown in Figure 10-25 on page 382.

```

9.43.86.31 - PuTTY
itsonas1>
itsonas1> help
?
aggr      ftp      nis      sis
arp       halt     options  snap
backup    help     orouted  snapmirror
bmc       hostname outb      snapvault
cf        httpstat partner  snmp
charmap   ifconfig passwd   software
cifs      ifstat   ping     source
config    igroup   ping6    stats
date      ipsec    pktt     storage
df        ipspace  portset  sysconfig
disk      iscsi    priority sysstat
disk_fw_update license  priv     timezone
dns       lock     qtree    traceroute
download  logger   rdate    traceroute6
dump      logout   reallocate
echo      lun      reboot   uptime
ems       man      restore  useradmin
environment maxfiles rlm      version
exportfs  mt       rmc      vfiler
fcdadmin  nbtstat route    vif
fcp       ndmcopy  routed   vlan
fcstat    ndmpd    sasadmin vol
file      ndp      sasstat  vscan
filestats netdiag  sascore  wcc
flexcache netstat  savecore ypcat
fpolicy   nfs      secureadmin
fsecurity nfsstat  setup    ypgroup
itsonas1>

```

Figure 10-25 Command Line help function (1)

If you want to execute one command, type the command at the prompt, and press Enter. If the command cannot be executed without a parameter, the system gives you command syntax for help. If there is a second-level parameter of the command, the system lists them, as shown in Figure 10-26 on page 383.

```
9.43.86.31 - PuTTY
itsonas1> date
Mon Sep 17 17:29:34 GMT 2007
itsonas1>
itsonas1> config
Usage:
    config clone <filer> <remote_user>
    config diff [-o <output_file>] <config_file1> [ <config_file2> ]
    config dump [-f] [-v] <config_file>
    config restore [-v] <config_file>
itsonas1>
itsonas1> cifs
The following commands are available; for more information
type "cifs help <command>"
access                gpresult                prefdc                sidcache
audit                 gpupdate                resetdc               stat
broadcast             help                    restart               terminate
changefilerpwd        homedir                 sessions              testdc
comment               lookup                  setup                 top
domaininfo            nbalias                 shares
itsonas1>
itsonas1> cifs access

Usage:
    cifs access <share> [-g] <user|group> <rights>
    cifs access <share> -m
    cifs access -delete <share> [-g] <user|group>
    cifs access -delete <share> -m
        rights can be Unix-style combinations of r w x -
        or NT-style "No Access", "Read", "Change", and "Full Control"

itsonas1>
```

Figure 10-26 Command Line help function (2)

For the command reference, there is an online manual page in FilerView, as shown in Figure 10-27 on page 384, where you can get an explanation of every command.

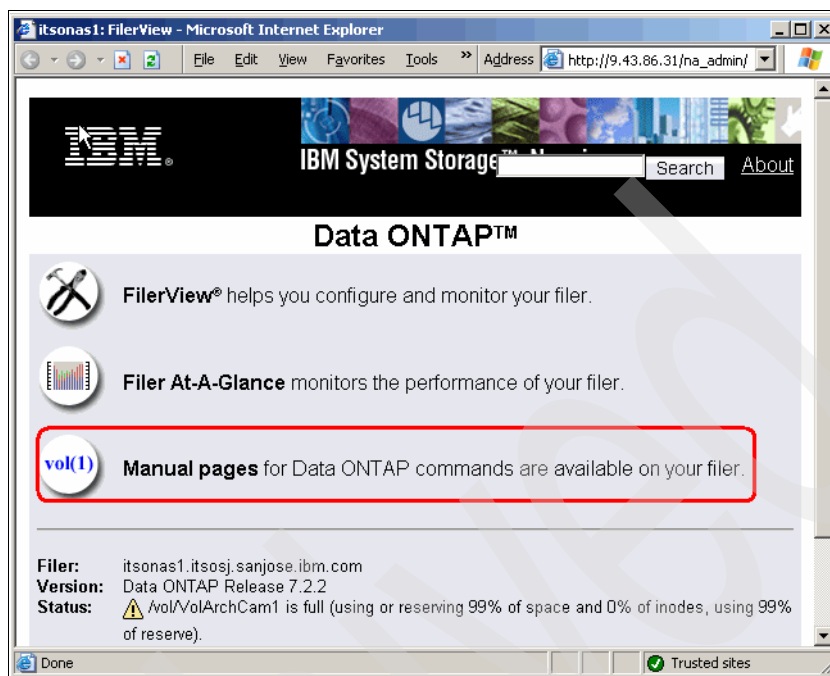


Figure 10-27 Online Manual in FilerView

## 10.4 Telnet

Telnet is a very popular remote access tool that almost every operating system supports.

### 10.4.1 Before you use Telnet

The storage system requires the following configurations before you connect to it using a Telnet session:

- ▶ The telnet.enable option must be set to On, which is the default setting.
- ▶ The telnet.access option must be set so that the protocol access control that is defined for the storage system allows Telnet access.

**Telnet sessions:** Only one Telnet session can be active at a time. You can, however, open a console session at the same time that a Telnet session is open, but they share one output.

## 10.4.2 Using Telnet

To use Telnet:

From your operating system, open a terminal window, and then initiate a Telnet session, as shown in Figure 10-28.




Figure 10-28 Initiate a Telnet session

After your telnet session connects to the storage system successfully, everything is the same as SSH.

Now we will give an example about the SnapVault usage on DVS, which configured using the CLI:

1. Enable SnapVault, as shown in Example 10-4.

---

### Example 10-4 SnapVault Example

```
itsonas1> options snapvault.enable on
itsonas1>
itsonas2> options snapvault.enable on
itsonas2>
itsonas1> options snapvault.access host=itsonas2
itsonas1>
```

```
itsonas2> options snapvault.access host=itsonas1
itsonas2>
```

---

## 2. Turn off normal Snapshot schedules, as shown in Example 10-5

### *Example 10-5 Turning off Snapshot schedules*

```
itsonas1> snap reserve vol_CamSource 0
itsonas1>
itsonas2> snap reserve vol_CamDest 0
itsonas2>
```

---

## 3. Set up SnapVault Snapshot schedules on the primary storage system, as shown in Example 10-6.

### *Example 10-6 Setting up Snapshot schedules*

```
itsonas1> snapvault snap sched vol_CamSource sv_hourly 4@9-18
itsonas2> snapvault snap sched vol_CamDest sv_hourly 4@9-18
```

---

## 4. Initialize the baseline transfer, as shown in Example 10-7.

### *Example 10-7 Baseline transfer*

```
itsonas2> snapvault start -S itsonas1:/vol/vol_CamSource/- /vol/vol_CamDest
Snapvault configuration for the qtree has been set.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
itsonas2> Mon Sep 17 22:23:19 GMT [itsonas2: waf1.inode.fill.enable:info]: fill
reservation enabled for inode 115102 (vol vol_CamDest).
Mon Sep 17 22:23:19 GMT [itsonas2: waf1.inode.overwrite.enable:info]: overwrite
reservation enabled for inode 115102 (vol vol_CamDest).
itsonas2>
```

---

## 5. Check the status of the SnapVault, as shown in Example 10-8.

### *Example 10-8 SnapVault status*

```
itsonas2> snapvault status
Snapvault secondary is ON.
Source          Destination          State          Lag
Status
itsonas1:/vol/vol_CamSource/- itsonas2:/vol/vol_CamDest/bkp Snapvaulted 00:00:27
Idle
itsonas2>
```

---

## 10.5 Serial Console Access

Serial Console Access is another connection mode for the N series storage system administrator. The administrator uses the serial console to do initial set up on storage system. It is only for local administrative task. The default setting for the Serial Console on hyperterminal tools is: 9600-8-N-1. Figure 10-29 shows how to access the storage system by hyperterminal.




Figure 10-29 Serial Console Access

After connecting to the storage system by console, you can execute the command, just like with SSH and Telnet.

Next, we give an example about the SnapMirror usage on DVS, which you can configure using the CLI:

1. Create a volume on the source storage system and destination storage system, as shown in Example 10-9.

### Example 10-9 SnapMirror implementation using the CLI

```
itsonas1> vol create vol_CamSource aggr_video 1g
itsonas1>
```

```
itsonas2> vol create vol_CamDest aggr_video 1g
itsonas2>
```

---

2. On the destination storage system, restrict the volume, as shown in Example 10-10.

*Example 10-10 Restricting the volume*

```
itsonas2> vol restrict vol_CamDest
Volume 'vol_CamDest' is now restricted.
itsonas2>
```

---

3. On the N series source storage system, specify the target host, as shown in Example 10-11.

*Example 10-11 Specifying the target host*

```
itsonas1> options snapmirror.access host=itsonas2
```

---

4. On the destination storage system, configure the `/etc/snapmirror.conf`, as shown in Example 10-12.

*Example 10-12 Configuring the snapmirror.conf file*

```
itsonas2> wrfile /etc/snapmirror.conf
itsonas1:vol_CamSource itsonas2:vol_CamDest restart=always sync
itsonas2>
itsonas2> rdfile /etc/snapmirror.conf
itsonas1:vol_CamSource itsonas2:vol_CamDest restart=always sync
itsonas2>
```

---

5. Enable SnapMirror, as shown in Example 10-13.

*Example 10-13 Enable SnapMirror*

```
snapmirror on
snapmirror on
```

---

6. On the destination N series storage system, initialize SnapMirror, as shown in Example 10-14.

*Example 10-14 SnapMirror initialization*

```
itsonas2> snapmirror initialize -S itsonas1:vol_CamSource itsonas2:vol_CamDest
Transfer started.
Monitor progress with 'snapmirror status' or the snapmirror log.
itsonas2>
```

---



7. Check the status of SnapMirror with the **snapmirror status** command, as shown in Example 10-15.

*Example 10-15 SnapMirror status*

---

```
itsonas2> snapmirror status
```

Snapmirror is on.

Source	Destination	State	Lag	Status
itsonas1:vol_CamSource	itsonas2:vol_CamDest	Snapmirrored	00:00:16	Idle

```
itsonas2>
```

---



# SnapDrive for UNIX on a Linux host

In this appendix, we explain how to install and use SnapDrive for UNIX on a Linux host.

## SnapDrive overview

SnapDrive for UNIX is a tool that simplifies the backup of data so that you can recover it if it is accidentally deleted or modified. SnapDrive for UNIX uses IBM N series Snapshot technology to create an image (that is, a Snapshot) of the data on a storage system that is attached to a UNIX host at a specific point-in-time. If the need arises later, you can restore the data to the storage system. When you restore a Snapshot, it replaces the current data on the storage system with the image of the data in the Snapshot. Figure A-1 on page 392 is an example of the SnapDrive command on the Linux host.

```
lochnese:~ # snapdrive
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
lochnese:~ #
```

Figure A-1 *SnapDrive commands*

In addition, SnapDrive for UNIX lets you provision storage on the storage system. It provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file through the volume manager to the storage system-side LUNs that provides the actual repository.

With SnapDrive for UNIX installed, you can:

- ▶ Create a Snapshot of one or more volume groups on a storage system. Then you can rename the Snapshot, restore it, or delete it. You can also connect a Snapshot to a different location on the host, on a different host, or disconnect it. After you connect the Snapshot, you can view and modify the content of the Snapshot copy, or you can disconnect the Snapshot copy.
- ▶ Create storage on a storage system in the form of LUNs, file systems, logical volumes, or disk groups. You can increase the storage or delete it. You can also connect the storage to a host or disconnect it. In addition, SnapDrive for UNIX lets you display information about the storage that you created with it.

**Note:** SnapDrive for UNIX works only with the Snapshots that it creates. It cannot restore Snapshots that it did not create.

The SnapDrive for UNIX software has commands that you can use to create, restore, and manage Snapshot copies of storage entities. SnapDrive software commands can also work with the logical volume manager (LVM) to create the LVM objects and file systems that use the N series storage system. The LVM combines LUN from the storage system into disk or volume groups and then divides it into logical volumes. SnapDrive for UNIX together with LVM determine which LUNs make up each disk group, host volume, and file system that requested Snapshot copy. SnapDrive for UNIX can perform the storage operation without using the host LVM. You can also use SnapDrive to create, delete, connect, and disconnect LUNs, and file systems that they contain without activating the LVM.

**Note:** SnapDrive for UNIX only supports Ext3 file systems on the Linux host at this time.

## Installing SnapDrive for UNIX

Before you install SnapDrive for UNIX, you need to verify that your N series storage systems are ready:

- ▶ The storage systems are online
- ▶ Data ONTAP 7.1 or later is recommended
- ▶ Make sure HBAs and NICs are ready to work
- ▶ IP network connection is available
- ▶ Related license are ready

Before you install SnapDrive for UNIX, you also need to ensure that the hosts are configured properly:

- ▶ If your configuration requires an FCP Host Utilities, you must install it and get it working. You can get the FCP Host Utilities from:

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

Refer to the documentation that comes with the FCP Host Utilities. It contains information about volume managers, multipathing, and other features that you need to set up before you install SnapDrive.

- ▶ If your configuration uses a iSCSI Host Utilities, refer to the iSCSI Host Utilities documentation to ensure that the system is being set up properly.
- ▶ Make sure the communication between the host and the storage system is working correctly.

## Downloading SnapDrive for UNIX on a Linux host

After you meet the prerequisites, start the installation procedure. In this appendix, we use a Linux host with a iSCSI as an sample. Here is the requirement for using SnapDrive on a Linux host in a iSCSI environment:

- ▶ iSCSI Linux Host Utilities: To make sure that you have the correct version of the utility, go to the IBM NAS product Web page at:

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

Set up the host and storage system according to the instructions in the installation and set up guide for the iSCSI Linux Host Utilities. You must do this before you install SnapDrive for UNIX.

- ▶ Additional disk space: SnapDrive for UNIX maintains the audit, recovery, and trace log files. While SnapDrive for UNIX rotates the files when they reach a maximum size, make sure that you have sufficient disk space for them. Based on the default settings for the audit and trace log files, you need at least 1.1 MB of space. There is no default size for the recovery log because it rotates only after an operation completes and not when it reaches a specific size.

## Moving the downloaded file to a local directory

If you downloaded the file and did not place it on your Linux host, you must move it to that host. To move the downloaded file to a local directory:

Copy the downloaded file to your Linux host. You can place it in any directory on the host, for example, you can use commands similar to the following ones to move the file you download to your Linux machine:

```
# mkdir /tmp/snapdrive
# cd /tmp/snapdrive
# cp /mnt/tmp/ontap.snapdrive.linux_3_0.rpm
```

Make sure you include the period (.) at the end of the copy command line.

**Note:** Ensure that all of the supported service packs are installed on the host before you install SnapDrive.

## Installing SnapDrive for UNIX on a Linux host

To install SnapDrive for UNIX on a Linux host:

1. Change to the directory on your Linux host where you put the software you download, and use the **rpm** command to install the software, as shown in Example A-1:

```
# rpm -U -v ontap.snapdrive.linux_3_0.rpm
```

2. After the installation, run the following commands to verify the SnapDrive installation, as shown in Example A-2:

```
# rpm -qa ontap.snapdrive
# ls -R /opt/Ontap
# ls -l /usr/sbin/snapdrive
```

3. Complete your setup by configuring *snapdrive.conf*. Most of this information is set by default; however, you need to specify the login information for the storage system. See “Specifying the current login information for storage systems” on page 396.

### Example: A-1 Install SnapDrive

---

```
lochnese:/tmp/snapdrive # rpm -U -v ontap.snapdrive.linux_3_0.rpm
Preparing packages for installation...
ontap.snapdrive-3.0-1
lochnese:/tmp/snapdrive #
```

---

### Example: A-2 Verify the installation of SnapDrive

---

```
lochnese:/tmp/snapdrive # rpm -qa ontap.snapdrive
ontap.snapdrive-3.0-1
lochnese:/tmp/snapdrive # ls -R /opt/Ontap
/opt/Ontap:
. .. snapdrive

/opt/Ontap/snapdrive:
. .. .pwfile bin diag docs snapdrive.conf

/opt/Ontap/snapdrive/bin:
. .. snapdrive

/opt/Ontap/snapdrive/diag:
. .. SHsupport.pm Telnet.pm filer_info linux_info snapdrive.dc

/opt/Ontap/snapdrive/docs:
. .. man1 snapdrive.1.html
```

```
/opt/On tap/snapdrive/docs/man1:  
. .. filer_info.1 linux_info.1 snapdrive.1 snapdrive.dc.1  
lochnese:/tmp/snapdrive # ls -l /usr/sbin/snapdrive  
lrwxrwxrwx 1 root root 39 Sep 17 15:43 /usr/sbin/snapdrive ->  
../../opt/On tap/snapdrive/bin/snapdrive  
lochnese:/tmp/snapdrive #
```

---

## Specifying the current login information for storage systems

When SnapDrive for UNIX accesses each storage system, it is authenticated with a user name and password. On the Linux host, in addition to being logged in as root, the person running SnapDrive for UNIX must supply the correct user name and password when prompted for it. If a login is compromised, you can delete it, and set a new user login.

You create the user login for each storage system when you set the storage system up. For SnapDrive for UNIX to work with the storage system, you must give SnapDrive the login information for the storage system. Depending on what you specified when you set up the storage systems, each storage system could use either the same login or a unique login.

SnapDrive for UNIX stores the storage system logins and passwords in encrypted form on each host. You can specify that SnapDrive encrypts this information when you send it across the wire by setting the `snapdrive.conf` variable `use-https-to-filer=on`.

## Specifying login information

To specify the user login information for a storage system:

**Note:** Depending on what you specified when you set up the N series storage system, each N series storage system can use either the same user name and password or unique ones. If all of the storage systems use the same user name and password information, you only need to perform the following steps one time. If not, repeat these steps for each storage system.

1. Enter the following command:

```
snapdrive config set user_name filename [filename ...]
```

*user\_name* is the user name that you specified for that storage system when you first set it up.

*filename* is the name of the storage system. You can enter multiple storage system names on one command line, if they all have the same user name and password. You must enter the name of at least one storage system.



2. At the prompt, enter the password, as shown in Example A-3. If no password was set, press Enter (the null value) when prompted for a password.
3. If you have another storage system with a different user name and password, repeat these steps.

*Example: A-3 Setting login information*

---

```
lochnese:/tmp # snapdrive config set root itsonas2
Password for root:
Retype password:
lochnese:/tmp #
```

---

### **Verifying storage system user names associated with SnapDrive**

You can verify which user name SnapDrive for UNIX has associated with a storage system by executing the **snapdrive config list** command.

**Note:** This command does not query the storage system to determine whether additional user names were configured for it. Nor does it display the password that is associated with a storage system.

To verify the storage system user names that are associated with SnapDrive, enter the following command, as shown in Example A-4:

```
snapdrive config list
```

This command displays the user name/storage system pairs for all systems that have users specified within SnapDrive for UNIX. It does not display the passwords for the storage systems.

*Example: A-4 Verifying SnapDrive login information setting*

---

```
lochnese:/tmp # snapdrive config list
user name      filer name
-----
root           itsonas2
lochnese:/tmp #
```

---

## **Using SnapDrive**

After you install SnapDrive, consider the `snapdrive.conf` file, which contains configurable variables of SnapDrive. This file includes a name/value pair for each configurable variable and is located under the SnapDrive installation directory. Most of the content of the `snapdrive.conf` file is set to default values, which you

can use in most cases. SnapDrive software provides some commands that enable you to modify these default values. In the former topic, we introduced the most important value for SnapDrive, the login information. After that, SnapDrive can communicate with the storage system, and now you can create a LUN.

## Creating LUNs using SnapDrive

In this section, we explain how to create LUNs with SnapDrive. When you use SnapDrive for UNIX to create storage, use the **snapdrive storage create** command. You can use SnapDrive for UNIX to create:

- ▶ LUNs
- ▶ A file system that is created directly on a LUN
- ▶ Disk groups, host volumes, and file systems that are created on LUNs

For the later two situations, SnapDrive automatically handles all of the needed tasks to set up LUNs, which includes preparing the host, performing discovery mapping, and connecting to each LUN that you create.

In our case, we show you the first situation to let you understand the creation procedure better. We create LUNs without a file system on it, do a partition manually, and make a file system on it manually. After you create the LUN, you can use the **snapdrive storage show** command to display information about the LUNs, disk groups, host volumes, file systems, or NFS directory trees that you create.

At the beginning, we do not have any LUNs, as shown in Example A-5, only the internal disk.

### Example: A-5 LUN Verification

---

```
lochnese:~ # ls -l /dev/disk/by-id/
total 0
drwxr-xr-x 3 root root 24 Sep 13 13:35 .
drwxr-xr-x 4 root root 32 Aug 29 10:22 ..
drwxr-xr-x 2 root root 25 Sep 13 13:31 HL-DT-ST_RW
lochnese:~ #
```

---

1. Create the LUNs that you need. The LUN is used to record live video data. The live video path is: /vol/vol\_vLive/LUNlive. So, name the LUN LUNlive, and create it in volume vol\_vLive, as shown in Example A-6.

### Example: A-6 LUN Creation

---

```
lochnese:/ # snapdrive storage create -lun
9.43.86.33:/vol/vol_vLive/LUNlive -lunsize 50m -noreserve
```

```

LUN itsonas2:/vol/vol_vLive/LUNlive ... created

mapping new lun(s) ... done
discovering new lun(s) ... done

LUN to device file mappings:
- itsonas2:/vol/vol_vLive/LUNlive => /dev/sda

lochnese:/ #

lochnese:/ # ls -l /dev/disk/by-id/*
lrwxrwxrwx 1 root root 9 Sep 19 11:04
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-0 -> ../../sda
lrwxrwxrwx 1 root root 9 Sep 19 11:04
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-0-generic ->
../../sg0

/dev/disk/by-id/HL-DT-ST_RW:
total 4
drwxr-xr-x 2 root root 25 Sep 19 09:50 .
drwxr-xr-x 3 root root 4096 Sep 19 11:04 ..
lrwxrwxrwx 1 root root 12 Sep 19 09:50 DVD_GCC-4480 -> ../../../hdc

```

---

You can see the device that you created, which is now mounted as /dev/sda.

2. Create the partition for /dev/sda, as shown in Example A-7.

*Example: A-7 SnapDrive Partition Creation*

---

```

lochnese:~ # fdisk
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-0
Device contains neither a valid DOS partition table, nor Sun, SGI or
OSF disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)

Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 1

```

First cylinder (1-1024, default 1):  
Using default value 1  
Last cylinder or +size or +sizeM or +sizeK (1-1024, default 1024):  
Using default value 1024

Command (m for help): **p**

Disk /dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-2: 52 MB,  
52428800 bytes  
2 heads, 50 sectors/track, 1024 cylinders  
Units = cylinders of 100 \* 512 = 51200 bytes

Start	End	Blocks	Id	System	Device	Boot
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101184428-2p1						
1	1024	51175	83	Linux		

Command (m for help): **w**  
The partition table has been altered!

Calling ioctl() to re-read partition table.  
Syncing disks.  
lochnese:/ #

3. Verify the result, as shown in Example A-8.

*Example: A-8 Verifying the result*

---

```
lochnese:/ # ls -l /dev/disk/by-id
total 4
drwxr-xr-x 3 root root 4096 Sep 19 11:26 .
drwxr-xr-x 4 root root 32 Aug 29 10:22 ..
drwxr-xr-x 2 root root 25 Sep 19 09:50 HL-DT-ST_RW
lrwxrwxrwx 1 root root 9 Sep 19 09:50
iscsi-iqn.1986-03.com.ibm:sn.101184428-0 -> ../../sda
lrwxrwxrwx 1 root root 9 Sep 19 09:50
iscsi-iqn.1986-03.com.ibm:sn.101184428-0-generic -> ../../sg0
lrwxrwxrwx 1 root root 10 Sep 19 09:50
iscsi-iqn.1986-03.com.ibm:sn.101184428-0p1 -> ../../sda1
lochnese:/ #
```

---

4. Now let us make the file systems. Create an EXT3 file system, in this example, because EXT3 is the supported file system to SnapDrive on Linux. (In our DVS environment, the Cisco VSM requires the XFS file system, which

SnapDrive does not support at this time). Example A-9 shows how to make the file system Ext3 on /dev/sda1.

*Example: A-9 File system creation*

---

```
lochnese:/ # mkfs -V -t ext3
/dev/disk/by-id/iscsi-qn.1986-03.com.ibm:sn.101184428-0p1
mkfs version 2.12 (Nov 17 2005)
mkfs.ext3 /dev/disk/by-id/iscsi-qn.1986-03.com.ibm:sn.101184428-0p1
mke2fs 1.38 (30-Jun-2005)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
12824 inodes, 51172 blocks
2558 blocks (5.00%) reserved for the super user
First data block=1
7 block groups
8192 blocks per group, 8192 fragments per group
1832 inodes per group
Superblock backups stored on blocks:
    8193, 24577, 40961

Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 33 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
lochnese:/ #
```

---

**Note:** Mount the partition in /etc/fstab by adding the following lines:

```
/dev/sda1 /LiveVideo ext3          defaults          0 0
```

## Using SnapDrive to make Snapshot copies

Now you can use the **snapdrive snap create** command to create Snapshot copies, which are point-in-time, read-only images of data that is on the storage system volumes. The snap create operation ensures that you backed up your LUNs or NFS files and directory trees. You can use the Snapshot copy that you create to restore your data if you encounter data corruption or other problems.

In our environment, we use Cisco VSM on an XFS file system; however, SnapDrive has not supported the XFS file system, until now. So, in the following example, we show a simple snapshot demo by using SnapDrive on the Ext3 file

system. We assume that we have some video surveillance data recorded on the video server's file system and that the video surveillance system is running on an Ext3 file system. First, we do a snapshot of it to back up the data on it, and then we delete some data on the file system. Last, we restore the snapshot and find that the deleted data is back.

To use SnapDrive to create snapshot copies:

1. Check the mountpoint of the LUN, as shown in Example A-10.

*Example: A-10 Check the mountpoint of the LUN*

---

```
lochnese:/ # mount
/dev/hda3 on / type xfs (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
tmpfs on /dev/shm type tmpfs (rw)
devpts on /dev/pts type devpts (rw,mode=0620,gid=5)
/dev/hdc on /media/cdrecorder type subfs
(ro,nosuid,nodev,fs=cdfss,procuid,icharset=utf8)
/dev/fd0 on /media/floppy type subfs
(rw,nosuid,nodev,sync,fs=floppyfss,procuid)
usbfs on /proc/bus/usb type usbfs (rw)
/dev/sda1 on /Video type ext3 (rw)
lochnese:/ #
```

---

2. Check the disk's utilization and its content, as shown in Example A-11.

*Example: A-11 Check disk utilization and files in the directory*

---

```
lochnese:/ # df -m /Video
Filesystem          1M-blocks      Used Available Use% Mounted on
/dev/sda1             49          19         28  41% /Video
lochnese:/ #
lochnese:/ # cd Video
lochnese:/Video # ls
.
..
Demo_Loitering_Event2_220061222143808393.wmv
Demo_Loitering_Event3_220061222143707784.wmv
Demo_DirectionalMotion_HC3-1_200612291595231_Comp2_Quad1.mlv
Demo_Loitering_Event4_220061222143500471.wmv
Demo_Finding DHL Trucks_Event10_320061218095124406.wmv
Demo_Loitering_Event5_220061222143135237.wmv
Demo_Finding DHL Trucks_Event6_320061222092641393.wmv
ibm_linux_host_utils_3_0.tar.gz
Demo_Finding DHL Trucks_Event7_320061221151015940.wmv
```

```

iscsi-initiator-utils-4.0.3.0-2.i386.rpm
Demo_Finding DHL Trucks_Event8_320061219151014906.wmv
Demo_Finding DHL Trucks_Event9_320061218151159875.wmv
lost+found
lochnese:/Video #

```

---

3. Create a snapshot for this LUN, as shown in Example A-12.

*Example: A-12 Create a snapshot*

```

lochnese:/ # snapdrive snap create -fs /Video -snapname snapVideo
Starting snap create /Video
WARNING: DO NOT CONTROL-C!
        If snap create is interrupted, incomplete snapdrive
        generated data may remain on the filer volume(s)
        which may interfere with other snap operations.
Successfully created snapshot snapVideo on itsonas2:/vol/vol_vLive

        snapshot snapVideo contains:
        file system: /Video
lochnese:/ #

```

---

4. Verify the snapshot that you created, as shown in Example A-13.

*Example: A-13 Verify the snapshot*

```

lochnese:/ # snapdrive snap show -filer itsonas2
snap name          host          date
snapped
-----
itsonas2:/vol/vol_vLive:snapVideo  lochnese      Sep 19 15:36
/Video

lochnese:/ # snapdrive snap list -fs /Video -v
snap name          host          date
snapped
-----
itsonas2:/vol/vol_vLive:snapVideo  lochnese      Sep 19 15:36
/Video
host OS: Linux 2.6.5-7.244-bigsmg #1 SMP Mon Dec 12 18:32:25 UTC 2005
snapshot name: snapVideo
file system:      type: ext3          mountpoint: /Video
lun path          dev paths
-----

```

```
itsonas2:/vol/vol_vLive/LUNlive      /dev/sda  
lochnese:/ #
```

---



5. Delete some files on this LUN, as shown in Example A-14.

*Example: A-14 Delete some files from the original directory*

---

```
lochnese:/Video # ls
.
..
Demo_Loitering_Event2_220061222143808393.wmv
Demo_Loitering_Event3_220061222143707784.wmv
Demo_DirectionalMotion_HC3-1_200612291595231_Comp2_Quad1.m1v
Demo_Loitering_Event4_220061222143500471.wmv
Demo_Finding DHL Trucks_Event10_320061218095124406.wmv
Demo_Loitering_Event5_220061222143135237.wmv
Demo_Finding DHL Trucks_Event6_320061222092641393.wmv
ibm_linux_host_utils_3_0.tar.gz
Demo_Finding DHL Trucks_Event7_320061221151015940.wmv
iscsi-initiator-utils-4.0.3.0-2.i386.rpm
Demo_Finding DHL Trucks_Event8_320061219151014906.wmv
Demo_Finding DHL Trucks_Event9_320061218151159875.wmv
lost+found
lochnese:/Video # rm D*
lochnese:/Video # ls
.
..
ibm_linux_host_utils_3_0.tar.gz
iscsi-initiator-utils-4.0.3.0-2.i386.rpm
lost+found
lochnese:/Video #
```

---

6. Do a snapdrive restore, as shown in Example A-15.

*Example: A-15 Do a snapshot restore*

---

```
lochnese:/ # snapdrive snap restore -fs /Video -snapname  
itsonas2:/vol/vol_vLive:snapVideo  
Starting to restore /Video  
  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  
Successfully restored snapshot snapVideo on itsonas2:/vol/vol_vLive  
file system: /Video
```

```
lochnese:/ #
```

---

7. Check the result of the restoration, as shown in Example A-16.

*Example: A-16 Check the result of the restoration*

---

```
lochnese:/ # cd Video
lochnese:/Video # ls
.
..
Demo_Loitering_Event2_220061222143808393.wmv
Demo_Loitering_Event3_220061222143707784.wmv
Demo_DirectionalMotion_HC3-1_200612291595231_Comp2_Quad1.mlv
Demo_Loitering_Event4_220061222143500471.wmv
Demo_Finding DHL Trucks_Event10_320061218095124406.wmv
Demo_Loitering_Event5_220061222143135237.wmv
Demo_Finding DHL Trucks_Event6_320061222092641393.wmv
ibm_linux_host_utils_3_0.tar.gz
Demo_Finding DHL Trucks_Event7_320061221151015940.wmv
iscsi-initiator-utils-4.0.3.0-2.i386.rpm
Demo_Finding DHL Trucks_Event8_320061219151014906.wmv
Demo_Finding DHL Trucks_Event9_320061218151159875.wmv
lost+found
lochnese:/Video #
lochnese:/Video # df -m
Filesystem            1M-blocks      Used Available Use% Mounted on
/dev/sda1              49             19         28  41% /Video
lochnese:/Video #
```

---

## Resizing the LUN using SnapDrive

SnapDrive for UNIX lets you increase (but *not* decrease) the size of the storage system volume group or disk group. You can 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. You can use the LVM commands to resize host volumes and file systems after you resize the underlying disk group.

## Resizing an Ext3 File System

SnapDrive requires that the host file system on Linux is Ext3. You can use SnapDrive to:

- ▶ Create an Ext3 file system directly.
- ▶ Resize the volume group or disk group and then extend the host volume manually.
- ▶ Take a snapshot of the file system and do a restoration when data is corrupted, which we showed in Example A-10 on page 402 to Example A-16 on page 406, in the previous section.

Example A-17 through Example A-20 on page 408, we explain how to increase the size of the storage with SnapDrive for UNIX on an Ext3 file system.

1. Create an LVM, and set up the LUN automatically, as shown in Example A-17.

*Example: A-17 Create an LVM and setting up the LUN automatically*

---

```
lochnese:/ # snapdrive storage create -fs /vArch -fstype ext3 -lun  
itsonas1:/vol/vol_u_video/lunA -lunsize 100m
```

```
LUN itsonas1:/vol/vol_u_video/lunA ... created
```

```
mapping new lun(s) ... done  
discovering new lun(s) ... done
```

```
LUN to device file mappings:  
- itsonas1:/vol/vol_u_video/lunA => /dev/sde
```

```
disk group vArch_SdDg created  
host volume vArch_SdHv created  
file system /vArch created
```

```
lochnese:/ #  
lochnese:/ # df -h /vArch  
Filesystem          Size  Used Avail Use% Mounted on  
/dev/mapper/vArch_SdDg-vArch_SdHv  
                    93M   4.1M   85M   5% /vArch  
lochnese:/ #
```

---

2. Resize the disk group, as shown in Example A-18 on page 407.

*Example: A-18 Resize the disk group with the SnapDrive*

---

```
lochnese:/ # snapdrive storage resize -dg vArch_SdDg -growby 100m  
-addlun
```

```

discovering filer LUNs in disk group vArch_SdDg...done
LUN itsonas1:/vol/vol_u_video/vArch_SdLun ... created

mapping new lun(s) ... done
discovering new lun(s) ... done.
initializing LUN(s) and adding to disk group vArch_SdDg...done
Disk group vArch_SdDg has been resized
Desired resize of host volumes or file systems
contained in disk group must be done manually
lochnese:/ #

```

---

3. Resize the disk group. Extend the logical volume manually, as shown in Example A-19.

*Example: A-19 Extend the logical volume*

```

lochnese:/ # lvextend -L +96m /dev/vArch_SdDg/vArch_SdHv
Extending logical volume vArch_SdHv to 196.00 MB
Logical volume vArch_SdHv successfully resized

```

---

4. After changing the size of the logical volume, manually resize the file system, as shown in Example A-20.

*Example: A-20 Resize the file system*

```

lochnese:/ # umount /vArch
lochnese:/ # resize2fs -f /dev/mapper/vArch_SdDg-vArch_SdHv
resize2fs 1.38 (30-Jun-2005)
Resizing the filesystem on /dev/mapper/vArch_SdDg-vArch_SdHv to 200704
(1k) blocks.
The filesystem on /dev/mapper/vArch_SdDg-vArch_SdHv is now 200704
blocks long.
lochnese:/ # mount /dev/mapper/vArch_SdDg-vArch_SdHv /vArch
lochnese:/ # df -h /vArch

```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/vArch_SdDg-vArch_SdHv	190M	4.1M	178M	3%	/vArch

```

lochnese:/ #

```

---

### ***Resizing an XFS file system***

We explained the procedure of resizing an Ext3 file system in “Resizing the LUN using SnapDrive” on page 406. However, in our environment, we use Cisco MediaServer to archive video data to the XFS file system, which uses storage on N series. The Cisco MediaServer, currently, is only supported on an XFS file system, and the SnapDrive is only supported on an Ext3 file system. So in our

environment, we cannot use SnapDrive to create a file system that Cisco MediaServer will use.

Using Example A-21 through Example A-26 on page 411, we explain how to create a LUN with SnapDrive and let LVM do the resizing work manually:

1. Use SnapDrive to create a LUN without a file system on it, as shown in Example A-21.

*Example: A-21 Create a LUN without file system on it*

---

```
lochnese:/ # snapdrive storage create -lun
itsonas1:/vol/vol_u_video/LUN1 -lunsize 100m

LUN itsonas1:/vol/vol_u_video/LUN1 ... created

mapping new lun(s) ... done
discovering new lun(s) ... done

LUN to device file mappings:
- itsonas1:/vol/vol_u_video/LUN1 => /dev/sdh

lochnese:/ #
lochnese:/ # ls -l /dev/disk/by-id/*
.....
lrwxrwxrwx 1 root root 9 Sep 24 13:48
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101183273-4 -> ../../sdh
lrwxrwxrwx 1 root root 9 Sep 24 13:48
/dev/disk/by-id/iscsi-iqn.1986-03.com.ibm:sn.101183273-4-generic ->
../../sg7
.....

/dev/disk/by-id/HL-DT-ST_RW:
total 4
drwxr-xr-x 2 root root 25 Sep 19 14:57 .
drwxr-xr-x 3 root root 4096 Sep 24 13:48 ..
lrwxrwxrwx 1 root root 12 Sep 19 14:57 DVD_GCC-4480 -> ../../../hdc
lochnese:/ #
```

---

2. Create a physical volume, volume group, and a logical volume, as shown in Example A-22.

*Example: A-22 Create the volumes*

---

```
lochnese:/ # pvcreate /dev/sdh
Physical volume "/dev/sdh" successfully created
lochnese:/ #

lochnese:/ # vgcreate clipvg /dev/sdh
Volume group "clipvg" successfully created
lochnese:/ #

lochnese:/ # vgchange -a y clipvg
0 logical volume(s) in volume group "clipvg" now active
lochnese:/ #

lochnese:/ # lvcreate -L 96 -ncliplv clipvg
Logical volume "cliplv" created
lochnese:/ #
```

---

3. Create an XFS file system, and mount it on, as shown in Example A-23.

*Example: A-23 Create an XFS file system and mount*

---

```
lochnese:/ # mkfs -t xfs /dev/clipvg/cliplv
meta-data=/dev/clipvg/cliplv      isize=256    agcount=6, agsize=4096
blks
        =                               sectsz=512
data      =                               bsize=4096    blocks=24576, imaxpct=25
        =                               sunit=0       swidth=0 blks,
unwritten=1
naming     =version 2                bsize=4096
log        =internal log             bsize=4096    blocks=1200, version=1
        =                               sectsz=512    sunit=0 blks
realtime   =none                     extsz=65536   blocks=0, rtextents=0
lochnese:/ #
lochnese:/ # mkdir Clip_lvm
lochnese:/ # mount /dev/clipvg/cliplv /Clip_lvm

lochnese:/ # df -h Clip_lvm
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/clipvg-cliplv
                92M  112K   92M   1% /Clip_lvm
lochnese:/ #
```

---

Now we have an XFS file system, which is mounted on /Clip\_lvm. The size of it is about 92 MB. Next, we increase its size.

4. Use SnapDrive to create another new LUN without the file system on it, as shown in Example A-24.

*Example: A-24 Create a new LUN without file system on it*

---

```
lochnese:/ # snapdrive storage create -lun
itsonas1:/vol/vol_u_video/LUN2 -lunsize 100m

LUN itsonas1:/vol/vol_u_video/LUN2 ... created

mapping new lun(s) ... done
discovering new lun(s) ... done

LUN to device file mappings:
- itsonas1:/vol/vol_u_video/LUN2 => /dev/sdi

lochnese:/ #
```

---

5. Create the new physical volume base on this new LUN, and ask LVM to extend the related volume group and logical volume, as shown in Example A-25.

*Example: A-25 LVM operations*

---

```
lochnese:/ # pvcreate /dev/sdi
Physical volume "/dev/sdi" successfully created
lochnese:/ #
lochnese:/ # vgextend clipvg /dev/sdi
Volume group "clipvg" successfully extended
lochnese:/ #
lochnese:/ # lvextend -L +96m /dev/clipvg/cliplv
Extending logical volume cliplv to 192.00 MB
Logical volume cliplv successfully resized
```

---

6. Use the **xfs\_growfs** command to extend the XFS file system, as shown in Example A-26.

*Example: A-26 Extend the XFS file system*

---

```
lochnese:/ # xfs_growfs /Clip_lvm
meta-data=/Clip_lvm          isize=256    agcount=6, agsize=4096
blks
      =                      sectsz=512
data      =                  bsize=4096   blocks=24576, imaxpct=25
```

```

=                                sunit=0      swidth=0 blks,
unwritten=1
naming =version 2                bsize=4096
log    =internal                 bsize=4096   blocks=1200, version=1
=                                sectsz=512    sunit=0 blks
realtime =none                   extsz=65536   blocks=0, rtextents=0
data blocks changed from 24576 to 49152
lochnese:/ #
lochnese:/ # df -h /Clip_lvm
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/clipvg-cliplv
                  188M  208K  188M   1% /Clip_lvm
lochnese:/ #

```

---

## Disconnecting, connecting, and deleting LUNs and storage entities to the host

We can use the **snapdrive storage connect** command to connect storage entities to the host, disconnect the storage entities, and delete them when required.

When you enter the **snapdrive storage connect** command to connect LUNs to the host, SnapDrive for UNIX performs the necessary discovery and mapping; however, it does not modify LUN contents.

The **storage disconnect** operation removes the LUNs, or the LUNs and storage entities, that were mapped to the host using the **storage create** or **storage connect** command. This action, which marks the disk and file system as exported, is the only change that disconnecting the mappings has on the contents of the LUNs. When you use the **storage disconnect** command on some operating systems, you lose information, such as the host volume names, the file system mount point, the storage system volume names, and the names of the LUNs. Without this information, reconnecting the storage at a later point-in-time is difficult.

The **snapdrive storage delete** command removes the storage entities on the host and all underlying host-side entities and storage system LUNs that are backing them.

**Important:** The **snapdrive storage delete** command deletes data. Use caution in running it.

Using Example A-27 on page 413 through Example A-30 on page 414, we explain how to use the **storage disconnect**, **storage connect**, and **snapdrive**



**storage delete** operations. We disconnect the LUNS first, connect it, and then we delete it.

1. Check the original file system, as shown in Example A-27.

*Example: A-27 Check the file system*

---

```
lochnese:/ # snapdrive storage list -all
```

Connected LUNs and devices:

device filename	adapter	path	size	proto	state	clone
lun path		backing snapshot				
-----	-----	----	----	-----	-----	-----
-----		-----				
/dev/sdh	-	P	1g	iscsi	online	No
itsonas1:/vol/vol_u_video/LUN1		-				

Host devices and file systems:

```
dg: vArch_SdDg          dgtype lvm
hostvol: /dev/mapper/vArch_SdDg-vArch_SdHv    state: AVAIL
fs: /dev/mapper/vArch_SdDg-vArch_SdHv    mount point: /vArch
(persistent) fstype ext3
```

device filename	adapter	path	size	proto	state	clone
lun path		backing snapshot				
-----	-----	----	----	-----	-----	-----
-----		-----				
/dev/sde	-	P	100m	iscsi	online	No
itsonas1:/vol/vol_u_video/lunA		-				
/dev/sdf	-	P	104m	iscsi	online	No
itsonas1:/vol/vol_u_video/vArch_SdLun		-				
/dev/sdg	-	P	104m	iscsi	online	No
itsonas1:/vol/vol_u_video/vArch-1_SdLun		-				

```
raw device: /dev/sda1    mount point: /Archives (persistent) fstype xfs
```

---

2. Disconnect a LUN, as shown in Example A-28.

*Example: A-28 Disconnect a LUN*

---

```
lochnese:/ # snapdrive storage disconnect -lun
itsonas1:/vol/vol_u_video/LUN1
```

```
    - LUN itsonas1:/vol/vol_u_video/LUN1 ... disconnected
0001-669 Warning:
```

Please save information provided by this command.  
You will need it to re-connect disconnected filespecs.

---

3. Reconnect the LUN to the host, as shown in Example A-29.

*Example: A-29 Connect a LUN to the host*

---

```
lochnese:/ # snapdrive storage connect -lun  
itsonas1:/vol/vol_u_video/LUN1  
  
mapping lun(s) ... done  
discovering lun(s) ... done  
  
LUN itsonas1:/vol/vol_u_video/LUN1 connected  
- device filename(s): /dev/sdh  
lochnese:/ #
```

---

4. Now, delete the LUN, as shown in Example A-30.

**Important:** The **snapdrive storage delete** command deletes data. Use caution in running it.

*Example: A-30 Delete a LUN*

---

```
lochnese:/ # snapdrive storage delete -lun  
itsonas1:/vol/vol_u_video/LUN1  
- LUN itsonas1:/vol/vol_u_video/LUN1 ... deleted  
lochnese:/ #
```

---

## Summary

In this appendix, we explain how to use SnapDrive for UNIX to manage your data on a Linux host. SnapDrive for UNIX is also supported on an AIX, HP-UX, and Solaris host. Visit the following Web site for documentation with detailed information:

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

Although in our installation we could not use SnapDrive because of the XFS limitation, your DVS, Linux, and N series might benefit from the management features that SnapDrive brings.

# Related publications

The publications that we list in this section are considered particularly suitable for a more detailed discussion of the topics that we covered in this book.

## IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 416. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *Using the IBM System Storage N Series with IBM Tivoli Storage Manager*, SG24-7243-00
- ▶ *Using an IBM System Storage N series with VMware to Facilitate Storage and Server Consolidation*, REDP-4211-00
- ▶ *IBM N Series Storage Systems in a Microsoft Windows Environment*, REDP-4083-00
- ▶ *IBM System Storage N series A-SIS Deduplication Deployment and Implementation Guide*, REDP-4320-01

## Other publications

These publications are also relevant as further information sources:

- ▶ *BM System Storage N series Introduction and Planning Guide*, GA32-0543-11
- ▶ *IBM System Storage N series Data ONTAP 7.3 System Administration Guide*, GC52-1279-00
- ▶ *IBM System Storage N series Data ONTAP 7.3 Network Management Guide*, GC52-1280-00

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM System Storage N series Data ONTAP 7.3RC1 Filer Publication Matrix  
<http://www-1.ibm.com/support/docview.wss?rs=1147&uid=ssg1S7002181>
- ▶ Support for Data ONTAP  
<https://www-304.ibm.com/systems/support/myview/supportsite.wss/supportresources?taskind=7&brandind=5000029&familyind=5329797&typeind=0&modelind=0&osind=0>
- ▶ Support for Network attached storage (NAS) & iSCSI  
<https://www-304.ibm.com/systems/support/supportsite.wss/allproducts?&taskind=1&brandind=5000029&familyind=0&typeind=0&modelind=0&osind=0>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

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

## Help from IBM

IBM Support and downloads

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

IBM Global Services

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

# Index

## Numerics

2864-A10 28  
2864-A20 28  
2865-A10 28  
2865-A20 29  
2866-A10 34–35  
2866-A20 34–35  
2867-A10 35  
2867-A20 35  
2868 - A10 29  
2868 -A20 29

## A

A10 18, 25  
Abandoned object 79  
access control 77  
aggregates 4  
alarm conditions 80  
all-IP 77  
analog cameras 75  
analog input 75  
analog into digital 82  
analog video surveillance 75  
ancillary data 70  
application environments 103  
applications, 103  
archive libraries 73  
archived data 106  
ASCII 23  
assets 82  
asynchronous 107  
asynchronous mode 144  
automated storage systems 82  
availability 2

## B

backup 110  
backup/recovery 2  
best-in-class video surveillance systems 114  
bezel 38  
biometric 73  
biometric data 82

biometric databases 81  
biometric system 81  
biometrics 81  
block level 2  
block options 103  
BM System Storage N series Gateway 52

## C

calculations 105  
cameras 69, 111  
capacity for growth 109  
captured data 78  
cassette recorder solution 83  
CCTV 69, 77  
centric digital video 102  
CIFs 102  
Cisco 114  
Cisco Intelligent Converged Environment 114  
Cisco video surveillance 114  
Cisco video surveillance products 114  
Cisco's Video Surveillance Media 113  
clip 75  
Cluster Failover 24, 109  
Clustered systems 109  
command center 73  
communications technology 77  
Compact flash 23  
compression 105  
configuration file 107  
console port 42  
constituents 76  
Converged Environment architecture 114  
converged network 76–78  
converged network infrastructure 77  
convergence 77  
copper 23  
cost considerations 82  
cost effectiveness 77  
CPU 20  
CPU tray 23

## D

data access system 76

- data availability 4
- data blocks 107
- data integrity 4
- data management 4
- data network 78
- Data ONTAP® 4
- data protection 52
- data streams 77
- data volume 82
- decentralization 78
- decision-support information 70
- detention 81
- Diagnostic 23
- different applications 103
- digital 82
- Digital data collection 72
- digital environment 73
- digital frame 72
- digital hardware 76
- digital images 72
- Digital surveillance 73
- digital surveillance solution 75
- digital video 102
- digital video capture 72
- digital video surveillance 72, 78–79
- digital video surveillance systems 73
- Directional motion 79
- disaster 69, 107
- disk drives 104
- disparate systems 78
- drive flexibility 14
- dual-node 37
- DVS 101–103, 107, 109–110

## E

- educational institution 69
- employee 111
- employee operations 111
- environmental data 79
- equipment 82
- ESH2 23
- Ethernet 23
- Ethernet ports 43
- evidentiary documentation 73
- EXN1000 106
- expiration date 82

## F

- Fault tolerance 104, 109
- FCP 3
- features 82
- feedback 81
- Fibre Channel 23
- Fibre drives 106
- file servers 103
- flexible storage 1
- flexible volumes 110
- FlexVol 110
- footage 81
- fragmentation 106
- frames per sec 105
- FTP 110
- future-ready strategy 73

## G

- Gateways 52
- grow on-demand 103
- Growth spike 110

## H

- halted node 109
- Hardware 24
- hiccups 102
- high maintenance 82
- high-bandwidth 76
- Higher availability 104
- higher resolution 102
- hot swappable 40
- human-caused 107

## I

- IBM 39, 102, 113
- IBM business partners 113
- IBM DVS 113
- IBM N series 13, 102
  - Gateway versus IBM N series storage systems 3
  - hardware 2
  - hardware quick reference
    - A models 4
    - A & G models 5
    - standard software features 7
    - storage systems A models 12
- IBM N series Gateway 4, 48, 51

- IBM N series storage systems 12–13
- IBM redbooks 113
- IBM Storage System N series 1
- IBM System Storage 17, 101
  - introduction to N3700 16, 28, 34
- IBM System Storage N series 1, 14
- IBM System Storage N series cluster 108
- IBM System Storage N series Gateways 48
- identifiers 81
- illegal 79
- illegal activity 72
- incident prevention 70
- information technology 71
- infrastructure 28, 77, 82
- initial surveillance efforts 76
- installations 103
- integrated digital solution 73
- Integration 81
- Internet 77
- intruders 83
- IP network infrastructure 114
- IP Next-Generation Network 76
- IP SAN 102
- IT 83
- IT-based network 83
- IT department 83
- IT infrastructure 71
- IT staff 71

## **L**

- LAN 76
- law enforcement 73
- lawsuits 72
- LEDs 46
- legal proceedings 75
- License-plate recognition 79
- litigious society 70
- local storage 82
- LockVault Compliance 4
- LUN 4

## **M**

- maintenance requirements 77
- Metro environment 83
- minimum video data retention period 82
- mirror 107
- Model A20 108
- MTTDL 104

- Multi-disk 33
- MultiStore 110
- multivendor environments 102

## **N**

- N series 105
- N series storage subsystems 110
- N series storage systems 107
- N3700 2, 17, 19, 21, 23
- N5000 4, 28–29, 31, 35
- N5000 series 2
- N5200 29, 33
- N5500 28, 32
- N7000 34, 36, 44
- N7000 series 2
- N7600 34–35
- N7800 35, 45
- NAS 3
  - IBM TotalStorage NAS 13
- natural disasters 70
- Near-line Feature 4
- near-line storage 15
- network bandwidth utilization 107
- network connections 110
- network resource 76
- network storage platform 102
- network-centric 101
- Network-centric DVS 101
- networked environment 73
- NFS 49
- node 108, 109
- node filesystem data 109
- Non disruptive storage system 109
- non-disruptively 1
- NVRAM 36
- NVRAM6 44

## **O**

- Object recognition 81
- Object removal 79
- Online 109
- open architecture 83
- optical 23
- optional software 8

## **P**

- pain point 83

- P-based network infrastructure 77
- PCI 44
- PCI slots 41
- PCI-Express 44
- People count 79
- personal digital assistants 77
- petrochemical facilities 83
- physical security 69, 82, 102
- physical storage 110
- point-in-time 110
- police chiefs 76
- policy 82, 110
- power 101
- prevention 101
- proactive management 70
- proprietary 76

## Q

- QVGA 105

## R

- RAID 4
- RAID Double Parity 104
- RAID-DP 104
- read-only access 107
- record 69
- recovery purposes 107
- Redbooks Web site 418
  - Contact us xv
- Redundant 23
- Redundant Cooling 23
- regulatory concerns 82
- remote users 102
- requirements 102
- resource management 77
- retention period 74
- robust analytics 101

## S

- SAN 3, 48
- SATA 15, 29, 35
- SATA drives, 106
- SCSI 3
- security 83
- security application environments 102
- security breaches 73
- security program 69–70

- security requirements 69
- security sensors 77
- serviceability 4
- short-term 82
- short-term storage 82
- simultaneous camera feeds 101
- size/shape/mass parameters 81
- SnapMirror 107
- Snapshots 110
- software quick reference 10
- solution design 82
- source system 107
- storage administrator 107
- storage capacity requirements 105
- storage expansion 23
- storage footprints 105
- storage management 52
- storage options 82
- storage provisioning 52
- storage system 4, 23, 52, 101–102, 106
- storage utilization 52
- stored footage 102
- surveillance 80, 101
- surveillance customer 82
- surveillance footage 74, 102
- surveillance infrastructure 75
- surveillance software 77
- surveillance system 81
- SWAT 70
- synchronizes 81
- synchronizing 107
- Synchronous mode 107
- system initialization 4
- System Storage N series 1
- System Storage N series solutions 102

## T

- takeover 108
- Tape 17
- TCO 102
- telephony service 76
- terabytes 83
- thin server 13
- transmit 82

## U

- utilization 49



## **V**

video analytics 83  
video and data transmission 77  
video cameras 69, 74  
video capture 70–71  
video cassette recorder 82  
video content 77  
video data 83, 110  
video evidence 72  
video footage 70, 72  
video frame 70  
video image 71  
video information 77  
video input 74  
video surveillance 69, 71–72, 76, 78, 83, 114  
video surveillance deployments 82  
video surveillance footage 79  
video surveillance service 83  
video surveillance system 79, 82  
video surveillance's new capability 80  
video surveillance design 82  
video transmission 114  
Video's 101  
videotape 72, 74  
videotape archive 75, 77  
videotape library 75  
virtualization 48  
voiceprints 81

## **W**

WAFL 3  
watermarking 73  
wireless capability 73  
wireless technology 73  
wireless transmission 73  
wireless wide area networks 77  
worldwide enterprises 102



## IBM System Storage N series and Digital Video Surveillance

(0.5" spine)  
0.475" <-> 0.875"  
250 <-> 459 pages







# IBM System Storage N series and Digital Video Surveillance

**Meeting the Digital Video Surveillance Storage Capacity demands with N series**

**Protecting Digital Video Data with N series**

**Implementing Digital Video Surveillance with N series**

In this IBM® Redbooks publication, we introduce Digital Video Surveillance (DVS) and the role that IBM plays in this industry. We provide the benefits of and define the role of the IBM System Storage™ N series in the DVS solution. We also include an example DVS application software installation and how it operates with the N series storage system.

## INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

### BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)