

Data ONTAP™ 6.4.2

System Administration

Block Access Management Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-00570
September 2003

Copyright and trademark information

Copyright information

Copyright © 1994–2003 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER

IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use and purchase of this product do not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, and the European Union. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. FAServer is a registered trademark of Network Appliance, Inc. in the United States and the European Union. NetCache is a registered trademark of Network Appliance, Inc. in the United States, Canada, the European Union, and Japan. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. WAFL is a registered trademark of Network Appliance, Inc. in the United States, the European Union, and Canada. DataFabric, FilerView, SecureShare, SnapManager, SnapMirror and SnapRestore are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP is a trademark of Network Appliance, Inc. in the United States and Canada. Snapshot is a trademark of Network Appliance, Inc. in the United States and the European Union. NetApp—the Network Appliance Company is a registered trademark of Network Appliance, Inc. in the United States and other countries. ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, gFiler, MultiStore, NearStore, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Table of Contents

	Preface	vii
Chapter 1	NetApp SAN Overview	1
	Understanding the NetApp multiprotocol environment	2
	How NetApp implements SAN.	9
Chapter 2	Managing LUNs and Initiator Groups.	13
	Configuring disk space	14
	Methods of creating LUNs	26
	Creating LUNs with the lun setup program	29
	Creating LUNs and initiator groups with FilerView	35
	Creating LUNs and initiator groups using individual commands	40
Chapter 3	Managing LUNs	45
	Managing LUNs	46
	Modifying LUN attributes	50
	Displaying LUN information.	54
	Monitoring disk space	57
Chapter 4	Managing Initiator Groups.	67
	Understanding initiator groups	68
	Managing initiator groups	70
Chapter 5	Using Data Protection in the SAN Environment	73
	Data ONTAP data protection processes	74
	Using snapshots.	75
	Using LUN clones	77
	Deleting busy snapshots	79

Using SnapRestore 85

Using SyncMirror. 91

Using volume copy 92

Using SnapMirror. 93

Using SnapVault 96

Backing up data to tape. 101

Using NVFAIL 105

Using ndmpcopy 107

Chapter 6 **Managing the NetApp SAN. 109**

 Managing the FCP SAN 110

 Managing the iSCSI SAN 125

Index 143

Preface

Introduction

This guide describes how to use NetApp® storage appliances, such as FAS800 or FAS900 Series Enterprise Servers, or filers and in a SAN environment, how to create and manage LUNs and initiator groups (igroups), and how to monitor SAN protocols.

Audience

This guide is for system administrators who are familiar with operating systems such as UNIX® and Windows 2000®, that run on the filer's clients. It also assumes that you know how to configure the filer and you know how Network File System (NFS), Common Internet File System (CIFS), and Hypertext Transport Protocol (HTTP) are used for file sharing or transfers and how FCP and iSCSI are used for block sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology.

Terminology

When the term *filer* is used, the information applies to all NetApp filer models, including the FAS900 and R100 series of storage appliances.

This guide uses the following terms:

- ◆ *Enter* refers to pressing one or more keys and then pressing the Enter key.
- ◆ *LUN* refers to a logical unit number, which represents a logical unit of storage.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Type* refers to pressing one or more keys on the keyboard

Command conventions

You can enter filer commands either on the system console or from any client (host) computer that can access the filer through a Telnet session.

In examples that illustrate commands executed on a UNIX workstation, this guide uses the command syntax of SunOS 4.1.x. The command syntax and output might differ, depending on your version of UNIX.

**Keyboard
conventions**

When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, *Ctrl-D* means pressing the *Control* and *D* keys simultaneously. Also, this guide uses the term *Enter* to refer to the key that generates a carriage return, although the key is named *Return* on some keyboards.

**Typographic
conventions**

The following table describes typographic conventions used in this guide.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters <code>arp -d</code> followed by the actual name of the host. Book titles in cross-references.
Monospaced font	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note_____

A note contains important information that helps you install or operate the system efficiently.

Caution_____

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

About this chapter This chapter provides an overview of how Network Appliance implements SAN protocols in a NetApp SAN environment.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Understanding the NetApp multiprotocol environment”](#) on page 2
- ◆ [“How NetApp implements SAN”](#) on page 9

Understanding the NetApp multiprotocol environment

What NetApp storage appliances are

NetApp storage appliances are hardware and software-based storage systems, such as filers, that serve data using protocols for both SAN and NAS environments. Storage appliances include the main unit, or chassis, disk shelves and disk drives, and cables. Storage appliances are installed in telco racks or in NetApp system cabinets. For information about storage appliance product families, see <http://www.netapp.com/products/>.

Supported protocols

With Data ONTAP 6.4.2, you can configure a filer as a target device in a SAN environment using FCP or iSCSI to communicate with one or more hosts running operating systems such as Solaris or Windows. You can also configure LUNs for multiprotocol access. You configure the filer's storage as logical units of storage (LUNs) for block access or as files for file access, or both.

Data ONTAP 6.4.2 supports the following SAN protocols:

- ◆ Fibre Channel Protocol (FCP) on Linux
- ◆ Fibre Channel Protocol (FCP) on Solaris 8 and 9
- ◆ Fibre Channel Protocol (FCP) on Windows 2000
- ◆ iSCSI on Linux (iSCSI)
- ◆ iSCSI on Windows 2000 (iSCSI)

Data ONTAP 6.4.2 supports the following NAS protocols:

- ◆ CIFS
- ◆ DAFS
- ◆ FTP
- ◆ HTTP
- ◆ NFS
- ◆ VLD

How Data ONTAP works in a SAN environment

In a NetApp SAN, filers (targets) act on requests from hosts (initiators), or application servers, and process those requests by writing data to or retrieving data from disks in disk shelves. Data ONTAP provides a complete set of storage management tools through its command-line interface and through the FilerView® interface and DataFabric™ Manager interface.

Supported configurations for FCP

Data ONTAP 6.4.2 supports the following configurations for connecting targets (filers) to initiators (hosts) in the FCP SAN:

- ◆ Fabric-attached single-path connection, with one filer cluster, one or more Fibre Channel switches, and one or more hosts

A single-path host has only one initiator Host Bus Adapter (HBA).

- ◆ Fabric-attached multipath connection, with one filer cluster, two Fibre Channel switches, and one or more hosts

A multipath host has two initiator HBAs and requires additional multipathing software, such as the DMP (Dynamic Multipathing) module of VERITAS® volume Manager 3.2 for Solaris. This configuration does not support access of data on a LUN as an NFS file if it is part of a volume stored under VERITAS Volume Manager. Data on a LUN that has a VxFS file system on it is also not accessible through NFS.

Data on LUNs that are not part of a VERITAS volume and do not have a VERITAS file system on them can still be accessed through NFS if the partition table conforms to the rules defined in the *Host Bus Adapter Installation and Setup Guide for Fibre Channel Protocol on Solaris*.

- ◆ Direct-attached, with only one filer and only one host

Clustered filers have two target HBAs installed in the PCI slots of the chassis. Stand-alone filers in the direct-attached configuration have one target HBA. Slot assignments vary, depending on the filer model. For current slot configuration information for your specific filer, see the *System Configuration Guide* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

For more information about configurations, see one of the following guides, depending on your environment:

- ◆ The *Host Bus Adapter Installation and Setup Guide for Fibre Channel Protocol on Solaris*
- ◆ The *Host Bus Adapter Installation and Setup Guide for Fibre Channel Protocol on Windows*

For the latest information about all supported configurations, see the SAN Support Matrix at the NOW site at <http://now.netapp.com/NOW/knowledge/docs.shtml>.

Supported configurations for iSCSI

Data ONTAP 6.4.2 supports the following configurations for initiators and targets in the iSCSI SAN:

- ◆ Direct attached—A single filer or a cluster directly attached to one or more hosts.

- ◆ Switched—A single filer or a cluster connected to one or more hosts through a TCP/IP switch.

How network connectivity is achieved

For FCP, the NetApp SAN network provides basic Fibre Channel SAN operation; however, Network Appliance highly recommends you use a TCP/IP network for connectivity to perform the following tasks:

- ◆ Managing storage appliances with administration hosts through Telnet or Remote Shell (rsh) sessions
- ◆ Providing http service to use FilerView
- ◆ Managing Fibre Channel switches
- ◆ Providing support for the SnapDrive™ feature in a Windows environment

For iSCSI, use a TCP/IP network for all connectivity.

Components of a NetApp SAN network

In a NetApp SAN network using FCP, nodes are connected to each other with Network Appliance-supplied optical cables to communicate with each other. Nodes are one of the following Fibre Channel devices: filers, hosts, switches.

In a NetApp SAN network using iSCSI, nodes are connected to each other with Ethernet cable. iSCSI nodes are filers or hosts.

In either environment, every node must have a unique identifier.

How nodes are uniquely identified for FCP

World Wide Node Name: For FCP, every node has a unique World Wide Node Name (WWNN).

NetApp assigns a WWNN to a filer based on the serial number of its NVRAM card. The WWNN is stored on disk.

World Wide Port Name: Each Fibre Channel device has one or more ports that are used to connect it to the SAN network. Each port has a unique World Wide Port Name (WWPN).

Both the WWNN and the WWPN are a 64-bit address represented in the following format: *nn:nn:nn:nn:nn:nn:nn:nn*, where *n* represents a hexadecimal value.

For both filers and hosts, each port of an HBA has a unique WWPN.

How Data ONTAP displays WWPNS and WWNNs: For filers, Data ONTAP displays the WWNN as *FC nodename* or *nodename*, and displays the WWPNS as *FC portname* or *portname* when using commands such as `storage show adapter`, `fc show adapter`, `fc config`, or `sysconfig -v`.

How HBA ports are uniquely identified

Each target HBA has two ports, labeled Port 1 and Port 2. Data ONTAP displays information about Port 1 by referring to it as *a* of the slot number the HBA is installed in and information about Port 2 as slot number *b*. For example, if the target HBA is installed in slot 5, Port 1 is displayed as 5a and Port 2 is displayed as 5b.


For filers, when the port is active, the addresses of the WWNN and WWPNS have a pattern of 50:a9:80:nn:nn:nn:nn. The FC node name for both active ports (5a and 6a) are identical because they are the WWNN for the filer. The FC port names for the active ports are unique. The FC node name and port name for standby ports are based on the adapter's native node and port names so that when adapters log in with the fabric, they have a different node name and port name from the filer. This prevents initiators from seeing the standby adapter as a second path to the target's active node. After takeover occurs, the standby filer and its adapters become active, the node name and port name of its adapter's ports are changed to the node name and port name of the failed partner, and the standby filer exports its partner's LUNs.

When the port is in standby mode, the addresses of the WWNN and WWPNS have a pattern of 20:01:00:nn:nn:nn:nn.

For example, assume that Filer_A and Filer_B are clustered and their target HBAs are installed in slots 5 and 6 in the chassis of both filers. Also assume that the ports are configured as follows:

- ◆ Filer_A, port 5a is active, and Filer_B port 5b is its standby
- ◆ Filer_A, port 6a is active, and Filer_B port 6b is its standby
- ◆ Filer_B, port 5a is active, and Filer_A port 5b is its standby
- ◆ Filer_B, port 6a is active, and Filer_A port 6b is its standby

This configuration is illustrated in the following figure.



The following example shows the patterns of the port names (WWPNs) and node names (WWNNs) for the active and standby (passive) ports of Filer_A.

Filer_A Active ports:

5a: FC nodename 50:a9:80:00:02:00:bf:0e
 FC portname 50:a9:80:**01**:02:00:bf:0e
 6a: FC nodename 50:a9:80:00:02:00:bf:0e
 FC portname 50:a9:80:**03**:02:00:bf:0e

Filer_A Standby ports:

5b: FC nodename 20:01:00:e0:8b:**26:3a:b7**
 FC portname 20:01:00:e0:8b:**26:3a:b7**
 6b: FC nodename 20:01:00:e0:8b:**27:c0:6c**
 FC portname 20:01:00:e0:8b:**27:c0:6c**

When the filers are not in takeover mode, you can use the WWNN and WWPN of the standby ports to verify connectivity.

For hosts, Data ONTAP displays the WWPN as *portname* for commands such as `fc show initiator`. The host WWPN is derived by prepending 10:00 to its IEEE address and has a pattern of 10:00:00:00:nn:nn:nn:nn.

How WWNNs and WWPNS are used: When a Fibre Channel device is connected to the SAN network, it registers each of its ports with the switch's Fabric Name Server directory service using its WWNN and WWPNS.

The WWNN of the filer is used to create persistent bindings between Solaris hosts and the LUNs on a particular filer. See the *Host Bus Adapter Installation and Setup Guide 1.0 for Fibre Channel Protocol on Solaris* for information about the filer's WWNN and target IDs.

Note

The use of persistent bindings is not necessary on Windows hosts.

WWPNs of the host are used when creating initiator groups. You use initiator groups to map a LUN to a host, thereby making the LUN accessible to the host.

How WWNNs and WWPNS are assigned on switches: Fibre Channel switches have one WWNN for the device itself, and one WWPNS for each of its ports. For example, the following diagram shows how the WWPNS are assigned to each of the ports on a 16-port switch.

Brocade Fibre Channel switch
WWNN: 10:00:00:60:69:51:06:b4

Port numbers:															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

Port **0**, WWPNS 20:**00**:00:60:69:51:06:b4

Port **1**, WWPNS 20:**01**:00:60:69:51:06:b4

Port **2**, WWPNS 20:**02**:00:60:69:51:06:b4

Port **3**, WWPNS 20:**03**:00:60:69:51:06:b4

...

Port **10**, WWPNS 20:**0a**:00:60:69:51:06:b4

Port **11**, WWPNS 20:**0b**:00:60:69:51:06:b4

Port **12**, WWPNS 20:**0c**:00:60:69:51:06:b4

Port **13**, WWPNS 20:**0d**:00:60:69:51:06:b4

Port **14**, WWPNS 20:**0e**:00:60:69:51:06:b4

Port **15**, WWPNS 20:**0f**:00:60:69:51:06:b4

How nodes are uniquely identified for iSCSI

Each iSCSI entity has an iSCSI node name. It is a logical name based on the several components: the date when the naming authority acquired the domain name, the name of the naming authority, and a unique device name. It is not linked to an IP address. Only targets and initiators are iSCSI entities. Switches, routers, and ports are TCP/IP devices only and do not have iSCSI node names.

Format for iSCSI node names

An iSCSI node name has the following format:

iqn.yyyy-mm.backward_naming_authority:unique_device_name

yyyy-mm is the month and year in which the naming authority acquired the domain name.

backward_naming_authority is the reverse domain name of the entity responsible for naming this device. An example reverse domain name is com.netapp.

unique_device_name is a free-format unique name for this device assigned by the naming authority.

Example : An iSCSI node name for an initiator that is an application server:

iqn.1987-06.com.initvendor1:sn2345

Default iSCSI node name for the filer

Each filer has a default node name based on the Network Appliance reverse domain name and the serial number of the filer's NVRAM card.

Example : The default node name for a filer with the serial number 123456789:

iqn.1992-08.com.netapp:sn.123456789

How NetApp implements SAN

What initiators and targets are

In a NetApp SAN environment, hosts are initiators and filers are targets that have storage target devices, which are referred to as LUNs (Logical Unit Numbers). With Data ONTAP, you configure the filer's storage by creating LUNs that can be accessed by initiators.

What LUNs are

From the host (initiator), LUNs appear as local disks on the host that you can format and manage to store data.

From the filer (target), a LUN is a logical representation of a physical unit of storage. It is a collection of, or a part of, physical or virtual disks configured as a single disk. When you create a LUN, it is automatically striped across many physical disks.

Data ONTAP manages LUNs at the block level, so it cannot interpret the file system or the data in a LUN.

What LUN serial numbers are

A LUN serial number is a unique 12-byte filer-generated ASCII string. Many multipathing software packages use this serial number to identify redundant paths to the same LUN. For example, VERATIS DMP expects the serial number of a drive to reside in bytes 36 through 47. You display the LUN serial number with the `lun show -v` command.


Prerequisites for creating LUNs

Before you can create LUNs, you must configure the filer's storage by using the `vol create` command to create volumes and, optionally, the `qtree create` command to create qtrees if you prefer to organize and manage LUNs at the qtree level. A LUN must be created at the root level of a qtree or a volume.

LUN examples

The following diagram shows four LUNs that were created by using the storage from Filer_A. LUNs 0 and 1 were created at the root of the volume `/vol/vol1`, and LUNs 2 and 3 were created at the root of the qtree `/vol/vol1/qtree1`. No LUNs were created on `/vol/vol0` because that is the root volume Data ONTAP uses for system administration. To summarize, the four LUNs that were created on Filer_A are as follows:

- ◆ /vol/vol1/LUN0
- ◆ /vol/vol1/LUN1
- ◆ /vol/vol1/qtrees1/LUN2
- ◆ /vol/vol1/qtrees1/LUN3



What LUN IDs are

A LUN must have a unique LUN identification number (LUN ID) that the host requires to have access to it. You can assign a LUN ID or accept a default, depending on the method of creating the LUN. Typically, the default ID begins with 0 and increments by 1 for each additional LUN as it is created, but valid LUN ID numbers are host-dependent. The host associates the LUN ID with the location and path name of the LUN. For information about best practices or guidelines on the ranges of values to use when assigning LUN IDs for your host, see the appropriate HBA installation and setup guide for your host.

What an initiator group is

An initiator group is a collection of the unique identifiers, either iSCSI node names or WWPNs of initiators (hosts). When you create an initiator group, you give it a name so that it is easy to manage, and you specify which initiators are in the initiator group by adding their unique identifiers to it.

How to display unique identifiers

On the filer, you use the command line or FilerView to display the unique identifiers of the initiators. At the command line, use the `fcv show initiator` command to see the FCP unique identifiers that are registered with the Fabric Name Server directory. You use the `iscsi show initiator` command to see the iSCSI unique identifiers that are attached to the network or the filer. With FilerView, select LUNs > Initiator Groups > Add.

How to associate unique identifiers to initiators

To know which initiators are associated with the unique identifiers, see the appropriate HBA installation and setup guide for your host. These guides describe NetApp-supplied commands to run on the host that show the mapping between the host name and its unique identifier. For example, for Windows hosts, you use the `lputilnt` command, and for Solaris hosts, you use the `sanlun` command. For host specific commands, see the appropriate HBA installation and setup guide.

Example of creating an initiator group

Assume you are using FCP, you have a Solaris host whose name is Host1, and you installed the initiator HBA in that host. You obtained its WWPN by one of the following methods:

- ◆ Noting its IEEE address and prepending 10:00
- ◆ Using the `fcv show initiator` command on the filer
- ◆ Using the `sanlun fcv show adapter` command on the host

You can use the `igroup create` command or FilerView to create an initiator group called *solaris-group1* and add the WWPN of the initiator, Host1, to it.

What LUN mapping (or masking) is

To make the LUNs accessible to hosts, you use the `lun map` command or FilerView to map LUNs to an initiator group. By mapping LUNs to the hosts listed in the initiator group, you allow access to the LUNs to the specified hosts. If you do not map a LUN, the LUN is not accessible to any hosts. This effectively masks access to the LUN.

Cluster configurations

Prerequisites for clustered filers: Clustered filers in a NetApp SAN environment require that the following options are enabled to guarantee that takeover and giveback occur quickly enough to not interfere with host requests to the LUNs. These options are automatically enabled when the FCP or iSCSI service is turned on.

- ◆ volume option `create_ucose` to On

- ◆ coredump.timeout.enable to On
- ◆ coredump.timeout.seconds set to 60 or less
- ◆ cf.wafl.delay.enabled to On
- ◆ cf.failover.on_panic to On

Cluster configurations for FCP: If you have a filer cluster, there are two dual-port target HBAs in each filer chassis. Port 1 on each HBA is the primary port and normally active. Port 2 on each HBA is the secondary port, and is on standby for the partner filer. When failover occurs, port 2 becomes active.

Cluster configurations for iSCSI: The filer does not have iSCSI target HBAs. The iSCSI service on the filer is implemented by the iSCSI software target (iswt) driver. This is a software driver that enables the filer to function as an iSCSI target device.

There are two logical iswt drivers: iswta and iswtb. The iswta adapter controls all iSCSI traffic for the filer. The iswtb adapter is dedicated to iSCSI traffic for the partner filer. If a failover occurs and the filer is operating in takeover mode, the iswtb driver becomes active.

Setting up a NetApp SAN

For information about how to set up a NetApp SAN, see the *NetApp SAN Setup Overview for FCP* or the HBA.

Filer administration

You can issue filer commands at the filer's console or you can open a Telnet or rsh session from a host, depending on what licenses are enabled on the filer and whether there is a TCP/IP connection between the filer and the host. A network interface card (NIC) is preinstalled in the filer. For information about filer administration, see the *Data ONTAP 6.4 System Administration Storage Management Guide*.

About this chapter

This chapter describes how to configure the filer's storage in preparation of creating LUNs and the methods for creating LUNs and initiator groups.

This chapter assumes that the NetApp SAN is set up and configured, and that the FCP or iSCSI service is licensed and enabled. See [“Managing the FCP service”](#) on page 111 or [“Managing the iSCSI SAN”](#) on page 125 for information about these topics.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Configuring disk space”](#) on page 14
- ◆ [“Methods of creating LUNs”](#) on page 26
- ◆ [“Creating LUNs with the lun setup program”](#) on page 29
- ◆ [“Creating LUNs and initiator groups with FilerView”](#) on page 35
- ◆ [“Creating LUNs and initiator groups using individual commands”](#) on page 40

Configuring disk space

Filer storage entities

You use the following entities to configure and manage disk space on the filer:

- ◆ Volumes
- ◆ Qtrees
- ◆ Files
- ◆ Logical Unit Numbers (LUNs)

The volume is the largest container of storage. It can hold qtrees, files, and LUNs.

A qtree is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume to group LUNs. For more information about qtrees, see the *Data ONTAP 6.4 System Administration Storage Management Guide*.

You use files to store data in a NAS environment, and you use LUNs to store data in a SAN environment.

How LUNs appear to your host

LUNs appear as local disks attached to the initiator host. As a result, it is imperative that LUNs behave as attached disks, which means that when you create a LUN of a specified size, the application storing data in the LUN has guaranteed access to that amount of space.

Why space reservation is so important

Data ONTAP 6.4.2 uses the LUN space reservation option to reserve enough space to guarantee that all writes to the LUN are completed successfully. This ensures that writes to the LUN do not fail due to a lack of disk space.

Whether you create LUNs using SnapDrive on a Windows host or create LUNs using Data ONTAP on a UNIX host, space reservations for the LUN is enabled by default. For information about creating LUNs with SnapDrive, see the *SnapDrive Installation and Administration Guide*.

Note

Space reservation settings persist across reboots, takeovers, givebacks, and snap restores, but do not persist through reversions to an earlier software version. Single File Snap Restore restores the reserved state of a LUN to the reserved state at the time the snapshot was taken.

Using snapshots with LUNs

NetApp snapshots are required for many NetApp features, such as SnapMirror, SyncMirror, SnapVault, dump and restore, and ndmptcopy.

Snapshots cannot use reserved space, so snapshots can fail if there is not enough available space on the volume. When using snapshots, you need to calculate the amount of storage capacity required for the volume based on the information at [“Guidelines to calculate the size of a volume”](#) on page 21.

When you create a volume, NetApp snapshots are enabled by default, twenty percent of the space is reserved for snapshots (snap reserve, or snapshot reserve in FilerView), and snapshots are automatically scheduled. When using snapshots with LUNs, NetApp recommends that you set the percentage of snap reserve to zero per cent and that you turn off the automatic snap schedule.

In UNIX environments, NetApp recommends that you ensure the file system or application accessing the LUN is quiesced or synchronized before taking snapshots. In the Windows environment, SnapDrive and SnapManager ensure that applications accessing LUNs are quiesced or synchronized automatically before taking snapshots.

Guidelines for creating volumes

When you create a volume to store LUNs, you must specify its name and the number of disks to be used to create it.

The capacity and the number of disks you specify determine the size of the volume. For example, if you have 72GB disks, and you use three disks to create a volume, two disks are used for data, and one is used for parity, so the volume capacity is 144GB. Assuming that a 72GB disk results in 67.9 GB of physical capacity, two 72 GB disks provide 135.8 GB of storage capacity. The actual amount of physical storage capacity might vary, depending on the disk manufacturer.

You calculate the size of the volume by determining the type and the amount of data that will be stored in the LUNs on the volume. For more information, see [“Guidelines to calculate the size of a volume”](#) on page 21.

In addition to the number of disks, you can specify the size of the RAID group. For information about disk allocation or RAID groups and how to configure them, see the *Data ONTAP 6.4 System Administration Storage Management Guide*. For information about disk configurations for your filer model, see the *System Configuration Guide* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

NetApp strongly recommends that you configure volumes as follows:

- ◆ Do not create any LUNs on /vol/vol0.

When the filer is first set up and configured, Data ONTAP creates a /vol/vol0 root volume. This root volume is used by Data ONTAP to administer the filer.

- ◆ Create a volume large enough to accommodate the LUNs that will be stored on the volume, using the information at “[Guidelines to calculate the size of a volume](#)” on page 21.
- ◆ Ensure that the volume option `create_unicode` is enabled.

Data ONTAP requires that the path of a volume or qtree containing a LUN is in the Unicode format. This option is On by default when you create a volume, but it is important to verify that any existing volumes still have this option enabled before creating LUNs in them.
- ◆ Ensure that the snapshot functionality is modified as follows:
 - ❖ Set snap reserve to 0%.
 - ❖ Turn off the automatic snapshot schedule.
- ◆ Ensure that no other files or directories exist in a volume that contains a LUN.

If this is not possible, and you are storing LUNs and files in the same volume, use a separate qtree to contain the LUNs.
- ◆ If multiple hosts share the same volume, create a qtree on the volume to store all LUNs for the same host.

Verifying and modifying the volume option `create_unicode` using the command line

To verify the volume option to create Unicode format directories by default is enabled or to enable the option using the command line, complete the following steps.

Step	Action
1	To verify that the create_ucose option is enabled (On), enter the following command: <pre>vol status [volname] -v</pre> Example: vol status /vol1 -v Result: The following output is a sample of what is displayed. <pre>Volume State Status Options vol1 online normal nosnap=off, nosnapdir=off, minra=off, no_atime_update=off, raidsize=8, nvfail=off, snapmirrored=off, resyncsnaptime=60, create_ucose=on convert_ucose=off, maxdirsize=10240, fs_size_fixed=off, create_reserved=on raid_type=RAID4</pre> Plex /vol/vol1/plex0: online, normal, active RAID group /vol/vol1/plex0/rg0: normal Note_____ If you do not specify a volume, status on all volumes is displayed. _____
2	To enable the create_ucose option, enter the following command: <pre>vol options volname create_ucose on</pre> Example: vol options vol1 create_ucose on

Verifying and modifying the volume option create_ucose using FilerView

To verify the volume option to create Unicode format directories by default is enabled or to enable the option using FilerView, complete the following steps.

Step	Action
1	Open a browser on your host.

Step	Action
2	Enter the name of the filer followed by /na_admin/ as the location for the URL. Example: If you have a filer named toaster, enter the following URL in the browser: http://toaster/na_admin. Result: The Network Appliance Online administrative window appears. 
3	Click FilerView.
4	Click Volumes from the left panel.
5	Click Manage.
6	Locate the name of the volume you want to check, and click on the Modify icon for that volume.
7	Locate the Create New Directories in Unicode field and select On if it is Off. For more information about this option, click on the question mark that provides online help next to the On/Off field.
8	Click the Apply button.

Setting the percentage of snap reserve space using the command line

To set a percentage of snap reserve space on a volume and to verify what percentage is set using the command line, complete the following steps.

For volumes that contain LUNs and no snapshots, NetApp recommends that you set the percentage to zero.

Step	Action
1	To set the percentage, enter the following command: snap reserve <i>volname percent</i> Example: snap reserve vol1 0
2	To verify what percentage is set, enter the following command: snap reserve [<i>volname</i>] Example: snap reserve vol1 Result: The following output is a sample of what is displayed. Volume vol1: current snapshot reserve is 0% or 0 k-bytes.

Setting the percentage of snap reserve space using FilerView

To set a percentage of snap reserve space (or snapshot reserve space) on a volume and to verify what percentage is set using FilerView, complete the following steps.

For volumes that contain LUNs and no snapshots, NetApp recommends that you set the percentage to zero.

Step	Action
1	Open a browser on your host.
2	Enter the name of the filer followed by /na_admin/ as the location for the URL. Example: If you have a filer named toaster, enter the following URL in the browser: http://toaster/na_admin.
3	Click on FilerView.
4	Click Volumes.

Step	Action
5	Click Snapshots.
6	Click Configure.
7	Select the volume you want to configure.
8	At the Snapshot Reserve field, enter 0 as the percentage of space the volume reserves for snapshots. For more information about snap (or snapshot) reserve, click on the question mark that provides online help next to the percentage field.
9	Click the Apply button.

Turning off the automatic snapshot schedule using the command line

To turn off the automatic snapshot schedule on a volume and to verify that the schedule is off using the command line, complete the following steps.

Step	Action
1	To turn the automatic snapshot schedule off, enter the following command: snap sched volname 0 0 0 Example: snap sched vol1 0 0 0 Result: This effectively turns the snapshot schedule off because no times are specified for snapshots to be taken automatically. You can still take snapshots manually using the snap command.
2	To verify the automatic snapshot schedule is off, enter the following command: snap sched [volname] Example: snap sched vol1 Result: The following output is a sample of what is displayed. Volume vol1: 0 0 0

Turning off the automatic snapshot schedule using FilerView

To turn off the automatic snapshot schedule on a volume and to verify that the schedule is off using FilerView, complete the following steps.

Step	Action
1	Open a browser on your host.
2	Enter the name of the filer followed by /na_admin/ as the location for the URL. Example: If you have a filer named toaster, enter the following URL in the browser: http://toaster/na_admin.
3	Click on FilerView.
4	Click Volumes.
5	Click Snapshots.
6	Click Configure.
7	Select the volume you want to configure.
8	At the Hourly Snapshot Schedule, ensure that no time slots are selected. For example, if a check appears at 8:00 AM, click on it. For more information about the hourly snapshot schedule, click on the question mark that provides online help next to the Hourly Snapshot Schedule field.
9	Click the Apply button.

Guidelines to calculate the size of a volume

The size of the volume depends on the types and sizes of the entities you intend to create on the volume. The entities might be any of the following:

- ◆ LUNs
- ◆ Snapshots
- ◆ Writable snapshots of LUNs (or *backing store* LUNs)
- ◆ LUN clones

Use the following table to determine the size of the volume that will contain the entities you want to store in that volume.

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Do not use snapshots	Is at least 10% larger than the total size of the LUN (or the cumulative total of all the LUNs to be stored in the volume).	volume size = $x + 10\%x$ or volume size = $1.10x$ If LUN size = 10GB, then volume size = $1.10 \times 10\text{GB}$ or 11GB
Take one snapshot	Is at least two times the size of the LUN, plus delta Delta is the rate of change, given as a percentage.	volume size = $2x + \text{delta}$ If LUN size = 10GB, and the rate of change since the snapshot is 15%, volume size = $2 \times 10\text{GB} + (.15 \times 10\text{GB}) = 21.5\text{GB}$
Take more than one snapshot	Is at least two times the size of the LUN plus the delta between each snapshot	volume size = $2x + \text{delta}$ If you take 2 snapshots, and each snapshot differs from the previous version by y bytes, then $\text{delta} = 2y$ If LUN size = 10GB, and you take 2 snapshots, and the rate of change is 15%, volume size = $2 \times 10\text{GB} + 2(.15 \times 10\text{GB}) = 23\text{GB}$ Worst case: If LUN size = 10GB, and you take 3 snapshots, and the rate of change is 100%, volume size = $2 \times 10\text{GB} + 3 \times 10\text{GB} = 50\text{GB}$

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Create a writable snapshot of a LUN (a backing store LUN) using the <code>lun create -b</code> command	Has enough space to contain a backing store LUN, which is typically used for prototyping. This type of LUN only needs space for changed data, since the unchanged data it contains is stored in the snapshot from which it was created. You can create a backing store LUN with or without space reservation enabled, and you can create it as read-only, if desired.	With space reservation enabled (NOT using the <code>-o noreserve</code> option): volume size = original LUN size, plus snapshot requirements, plus space for this new LUN and its snapshot requirements With space reservation disabled (using the <code>-o noreserve</code> option): volume size = original LUN size, plus snapshot requirements
Do a single file snapshot restore	Has enough space to store the original LUN and the restored LUN.	volume size = $2x + 10\%x$ or volume size = $2.1x$
Create a clone of the writable snapshot of a LUN (a clone of a backing store LUN) using the <code>lun clone</code> command	Has enough space to store the original LUN and its snapshot, and the LUN clone and its snapshot. To clone a backing store LUN and make it readable-writable, enable space reservations on the LUN and treat it as you would any other LUN.	volume size = original LUN size, plus its snapshot requirements, plus the LUN clone and its snapshot requirements. Typically, this will be $4x + \text{delta}$.

Creating a volume using the command line

To create a volume using the command line, complete the following step.

Step	Action
1	On the filer's command line, enter the following command to create a volume: vol create volname disk-list <i>volname</i> is the name of the new volume (include the /vol/ prefix). <i>disk-list</i> is one of the following: ◆ <i>ndisks</i> [@disk-size] <i>ndisks</i> is the number of disks to use. It must be at least 2. <i>disk-size</i> is the disk size to use, in gigabytes. You must have at least two available disks of the size you specify. ◆ -d <i>disk_name1</i> <i>disk_name2</i>... <i>disk_nameN</i> <i>disk_name1</i>, <i>disk_name2</i>... <i>disk_nameN</i> are disk IDs of one or more available disks. Separate multiple disks with a space. Example: vol create vol1 8 For more information about the vol command, see the <i>Storage Management Guide</i>.

Creating a volume using FilerView

To create a volume using FilerView, complete the following steps.

Step	Action
1	Open a browser on your host.
2	Enter the name of the filer followed by /na_admin/ as the location for the URL. Example: If you have a filer named toaster, enter the following URL in the browser: http://toaster/na_admin.
3	Click on FilerView.
4	In the left panel, click on Volumes.
5	Click on Add. Result: The Add New Volume page appears.

Step	Action
6	Fill in the fields. For information about any of the fields, click on the question mark by the appropriate field.
7	Click the Add button.

Organizing LUNs by creating qtrees using the command line

You can use qtrees to organize related LUNs into groups. To create a qtree on the filer to host multiple LUNs, complete the following step.

Step	Action
1	On the filer's command line, enter the following command to create a qtree: <pre>qtree create /vol/volname/qtree-name</pre> Note_____ Create the qtree in any volume other than the root volume (/vol/vol0). Include /vol in the path because if it does not begin with a slash (/), the qtree is created in the root volume. _____ Example: qtree create /vol/vol1/qtree1 Result: qtree1 is a qtree in /vol/vol1.

Organizing LUNs by creating qtrees using FilerView

If you want to create a qtree in a volume using FilerView, launch Filerview, and click Volumes > Manage > Qtrees > Manage > Add Qtree button. Then fill in the fields, and click on the Add button.

Methods of creating LUNs

About creating LUNs

There are several ways to create a LUN, but whichever method you choose, you create a LUN by specifying the following attributes:

- ◆ The path name of the LUN
The path name must be at the root level of a qtree or a volume that specifies where the LUN is located. Example: /vol/vol1/lun1.

Note
You might find it useful to provide a meaningful path name, one that describes how the LUN is used, such as the name of an application or the kind of data that is stored in the LUN, or the user’s name who has exclusive or primary use of the data stored in the LUN.

- ◆ The type of operating system running on the host that accesses the LUN
 - ❖ The values are default, solaris, and windows.
default indicates that the contents of the LUN are not specified.
solaris indicates that the LUN stores a Solaris raw disk in a single-slice partition.
windows indicates that the LUN stores a raw disk device in a single-partition Windows disk.
linux indicates that the LUN stores a Linux raw disk with no partition table. The linux os type uses the “solaris” style geometry, and there is no prefix or suffix. So if you intend to create a LUN from a file or access the data on the LUN using NAS, you cannot put a partition table on a LUN with the “linux” os type. For more information about creating LUNs for Linux access, see the appropriate Host Attach Kit for Linux.
- ◆ The size of the LUN
When you create a LUN, you specify its size as raw disk space, depending on the filer and the host. The usable space in the LUN is dependent on host or application overhead requirements. For example, partition tables and host file system metadata reduce usable space for applications. The minimum and maximum sizes for LUNs are as follows:

Operating System	Minimum	Maximum
Solaris	4MB	1023GB

Operating System	Minimum	Maximum
Windows: Basic disk	31.5MB	2TB
Windows: Dynamic disk		1TB
Data ONTAP		2TB

You specify the size, in bytes (default), or by using the following multiplier suffixes.

Multiplier suffix	Size
c	bytes
w	words or double bytes
b	512-byte blocks
k	kilobytes
m	megabytes
g	gigabytes
t	terabytes

- ◆ A brief description of the LUN (optional)
- ◆ The name of the initiator group that is allowed access to the LUN

Note

You can create initiator groups that specify which initiators have access to the LUNs either before or after you create LUNs, but initiator groups must be created before you can map a LUN to an initiator group.

- ◆ The LUN ID number that the initiator uses to identify and access the LUN
- ◆ Whether space reservation is enabled

Methods of creating LUNs

You create LUNs using one of the following methods:

- ◆ Entering the `lun setup` command, which prompts you through the sequence of creating a LUN, creating an initiator group, and mapping the LUN to the initiator group
- ◆ Using FilerView
- ◆ Entering a series of individual commands (such as `lun create`, `igroup create`, and `lun map`)

Note

You can create as many LUNs as Data ONTAP supports using the `lun create` command before mapping them to LUN IDs seen by initiator groups.

- ◆ Using SnapDrive 2.0 or later for Windows
SnapDrive furnishes a user-friendly interface to manage LUNs on a filer. This product provides online storage configuration, LUN expansion, and streamlined storage management. Integrated with NetApp Snapshot functionality, it creates point-in-time images of data stored in LUNs that are connected to the filer, for backup, archiving, recovery, and asynchronous mirroring. For more information, see the appropriate *SnapDrive Installation and Administration Guide* for the license you purchased.

Note

If you choose to use SnapDrive, all LUN management functions must be performed using the SnapDrive interface.

Guidelines for creating LUNs

Use the following guidelines when creating LUNs:

- ◆ Do not put LUNs in the `/vol/vol0` root volume.
- ◆ Create LUNs at the root of a `qtree` in a volume or at the root of the volume. You can use a naming convention to help you identify where LUNs are located, such as prepending the name of the LUNs with *lun*. For example: `/vol/vol1/lun_host1`.
- ◆ Distribute LUNs across the filer cluster.

Creating LUNs with the lun setup program

What the lun setup program does

The lun setup program prompts you for information needed for creating a LUN and an initiator group, and for mapping the LUN to the initiator group. When a default is provided, you can press Enter to accept the default.

If you did not create volumes and qtrees to store LUNs in prior to running the lun setup program, you must terminate the program and create volumes and qtrees, if you want to use them.

If you did not create initiator groups, you can create them during the lun setup program.

Running the lun setup program

To run the lun setup program, complete the following steps. The answers given are an example of creating LUNs using FCP in a Solaris environment.

Step	Action
1	On the filer command line, enter the following command. If you press Enter, to select y, the lun setup program provides instructions and prompts you to enter your values, as shown in the remaining steps. If you press n, the program terminates. lun setup Result: The lun setup program displays the following instructions: This setup will take you through the steps needed to create LUNs and to make them accessible by initiators. You can type ^C (Control-C) at any time to abort the setup and no unconfirmed changes will be made to the system. Do you want to create a LUN? [y]:
2	Specify the operating system that will be accessing the LUN by responding to the next prompt: OS type of LUN (default/solaris/windows/linux) [default]: Example: solaris If you press Enter to select default, the contents of the LUN are not specified. Typing <i>solaris</i> indicates that the LUN will store a Solaris raw disk in a single-slice partition. Typing <i>windows</i> indicates that the LUN will store a raw disk device in a single-partition Windows disk. Typing <i>linux</i> indicates that the LUN will store a raw disk with no paratition table.

Step	Action
3	Specify where the LUN will be located by responding to the next prompt: A LUN path must be absolute. A LUN can only reside in a volume or qtree root. For example, to create a LUN with the name "lun0" in the qtree root /vol/vol/q0, specify the path as "/vol/vol1/q0/lun0". Enter LUN path: Example: If you previously created /vol/vol1/, you would enter /vol/vol1/lun0. Caution_____ Do not create LUNs in /vol/vol0. This is the root volume used for filer administration. Result: <i>lun0</i> is created in the root of /vol/vol1 if the configuration information is accepted later in this program.
4	Specify whether you want the LUN created with space reservations enabled. A LUN can be created with or without space reservations being enabled. Space reservation guarantees that data writes to that LUN will never fail. Do you want the LUN to be space reserved? [y]: Example: Press Enter to accept the default.
5	Specify the size of the LUN by responding to the next prompt: Size for a LUN is specified in bytes. You can use single-character multiplier suffixes: b(sectors), k(KB), m(MB), g(GB) or t(TB). Enter LUN size: Example: 5g Result: A LUN with 5 GB of raw disk space is created if the configuration information is accepted later in this program. The amount of usable disk space by the host varies depending on the operating system type and the application using the LUN.

Step	Action
6	Create a comment or a brief description about the LUN by responding to the next prompt: You can add a comment string to describe the contents of the LUN. Please type a string (without quotes), or hit ENTER if you don't want to supply a comment. Enter comment string: Example: 5 GB Solaris LUN for finance records If you do not want to provide a comment at this time, press Enter. You can add a comment later with the <code>lun comment</code> command or a description by using FilerView.
7	Create an initiator group by responding to the next prompt: The LUN will be accessible to an initiator group. You can use an existing group name, or supply a new name to create a new initiator group. Enter '?' to see existing initiator group names. Name of initiator group[]:
	Example 1: ? Result 1: If you did not create any initiator groups yet, the following message is displayed: Name of initiator group[]:? No existing initiator groups Name of initiator group[]: Result 2: If you already created one or more initiator groups, they are displayed. If you created two initiator groups, <code>solaris-igroup1</code> and <code>solaris-igroup2</code>, they appear as follows: Name of initiator group[]:? Existing initiator groups: solaris-igroup1 solaris-igroup2 Name of initiator group[solaris-igroup1]: Press Enter to accept the default for the last initiator group used; otherwise, enter the name of the initiator group that contains the initiator you want to allow access to the LUN you are creating. When you enter the name of an existing initiator group, it is displayed, along with the WWPNs it contains, as shown in the following message: solaris-igroup1 10:00:00:00:c9:2c:43:a0 10:00:00:00:c9:2c:46:60 The LUN will be accessible to all initiators in the initiator group. Enter '?' to display the LUNs already in use by one or more initiators in group "solaris-igroup1".

Step	Action								
	Example 2: solaris-igroup3 Result: The initiator group solaris-igroup3 is created. Note_____ The name of the initiator group that you create is independent of the name of the host that is used by the host operating system, host files, or DNS. _____								
8	Specify which protocol will be used by the host in the initiator group by responding to the next prompt. Press Enter to accept the default. Type of initiator group solaris-igroup3 (FCP/iSCSI)[FCP]: Result: The initiator unique identifiers (FCP WWPNs) in this initiator group use the FCP protocol.								
9	Add the unique identifiers (iSCSI node names or FCP WWPNs) of the hosts that will be in the initiator group by responding to the next prompt: A Fibre Channel Protocol (FCP) initiator group is a collection of initiator port names. Each port name (WWPN) is 16 hexadecimal digits, separated (only) by optional colon (:) characters. You can separate port names by commas. Enter '?' to display a list of connected initiators. Hit ENTER when you are done adding port names to this group. Enter comma separated portnames:								
	Example 1: ? Result: The following output is a sample of what is displayed. Initiators connected on adapter 4a <table> <tr> <th>Portname</th><th>Group</th></tr> <tr> <td>10:00:00:00:c9:2b:cc:51</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:dd:62</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:ee:5d</td><td></td></tr> </table> Adapter 4b is running on behalf of the partner. Initiators connected on adapter 5a: None connected. Enter comma separated portnames:	Portname	Group	10:00:00:00:c9:2b:cc:51		10:00:00:00:c9:2b:dd:62		10:00:00:00:c9:2b:ee:5d	
Portname	Group								
10:00:00:00:c9:2b:cc:51									
10:00:00:00:c9:2b:dd:62									
10:00:00:00:c9:2b:ee:5d									

Step	Action
	Example 2: 10:00:00:c9:2b:cc:51 Result: The initiator identified by this WWPN is added to the initiator group that was specified in Step 8. You are prompted for more portnames until you press Enter. Note_____ To determine which WWPN is associated with a host, use the NetApp-supplied commands you downloaded from the CD-ROM in the HBA kit, such as <code>sanlun</code> for Solaris or <code>lputilnt</code> for Windows hosts. See the HBA installation and setup guide for your host for instructions about how to use the commands. _____
10	Specify the operating system type that will be accessing the LUN by responding to the next prompt: OS type of initiator group "solaris-igroup3" (default/solaris/windows) [default]:
11	Specify the LUN ID that the host will map to the LUN by responding to the next prompt. Press Enter to accept the default. The LUN will be accessible to all the initiators in the initiator group. Enter '?' to display LUNs already in use by one or more initiators in group "solaris-igroup3". LUN ID at which initiator group "solaris-igroup3" sees "/vol/vol1/lun0" [0]: Result: If you press Enter to accept the default, Data ONTAP issues the lowest valid unallocated LUN ID to map it to the initiator, starting with zero. Alternatively, you can enter any valid number. See the HBA installation and setup guide for your host for information about valid LUN ID numbers. Note_____ Network Appliance recommends that you accept the default value for the LUN ID. _____ After you press Enter, the lun setup program displays the information you entered: <pre>LUN Path : /vol/vol1/q0/lun0 OS Type : default Size : 5g (5368709120) Comment : 5 GB Solaris LUN for finance records Initiator Group : solaris-igroup3 Initiator Group Type : FCP Initiator Group Members : 10:00:00:00:c9:2b:cc:51 Mapped to LUN-ID : 0</pre>

Step	Action
12	Commit the configuration information you entered by responding to the next prompt: Do you want to accept this configuration? [y] Result: If you type y, the LUNs are mapped to the specified initiator group. All changes are committed to the system and Ctrl-C cannot undo these changes. The LUN is created and mapped. If you want to modify the LUN, its mapping, or any of its attributes, you need to use individual commands or FilerView. See Chapter 3.
13	Either continue creating LUNs, or terminate the program by responding to the next prompt. Press Enter to terminate or type y to continue: Do you want to create another LUN? [n]

Creating LUNs and initiator groups with FilerView

What FilerView is

FilerView is a graphical user interface that you can use to create and manage LUNs and initiator groups and manage FCP and iSCSI service and traffic. After you launch FilerView, you can use the LUN Wizard, or you can use FilerView menus to create LUNs and initiator groups.

Launching FilerView


To launch FilerView, complete the following steps.

Step	Action
1	Open a browser on your host.
2	Enter the name of the filer followed by /na_admin/ as the location for the URL. Example: If you have a filer named tpubs-cf1, enter the following URL in the browser: http://tpubs-cf1/na_admin. Result: The Network Appliance Online administrative window appears.
3	Click FilerView. Result: The FilerView screen appears. The left panel provides a list of Data ONTAP functionality that you manage using FilerView.

Using the LUN Wizard

To use the LUN Wizard to create LUNs and initiator groups, complete the following steps.

Step	Action
1	Click LUNs in the left panel of the FilerView screen. Result : The management tasks you can perform on LUNs are displayed. 
2	Click Wizard. Result: The Wizard window appears and explains how to use it.

Step	Action
3	Follow the instructions in the Wizard and enter information in the appropriate fields.  Note— You can click the online help button (the question mark) to get more information.
4	At the Commit Changes window, review your input. If everything is correct, click Commit. Result: The LUN Wizard: Success! window appears and the LUN you created is mapped to the initiator groups you specified.

Using FilerView menus

To create a LUN using FilerView menu, complete the following steps.

Step	Action
1	Launch FilerView.
2	Click LUNs.
3	Click Add.

Step	Action
4	Fill in the fields. Note _____ Click the applicable question mark for online help for that field. _____
5	Click the Add button to commit changes.

Creating LUNs

To create a LUN using FilerView menus, complete the following steps.

Step	Action
1	Launch FilerView.
2	Go to LUNs > Add.
3	Fill in the fields.
4	Click the Add button to commit changes.

Creating initiator groups

To create an initiator group using FilerView menus, complete the following steps.

Step	Action
1	Launch FilerView.
2	Go to Initiator Groups > Add.
3	Fill in the fields.
4	Click the Add button to commit changes.

Mapping LUNs to initiator groups

To map LUNs to initiator groups using FilerView menus, complete the following steps.

Step	Action
1	Launch FilerView.
2	Go to LUNs > Manage.
3	Click “No Maps” in the Maps column that is on the same row as the path name of the LUN.
4	Click the Add Groups to Map link.
5	Select an initiator group name from the list on the right side of the window.
6	Click the Add button to commit the changes.

Creating LUNs and initiator groups using individual commands

When to use individual commands

There are several ways you can combine individual commands to create a LUN. The following table steps you through the basics.

To make a LUN accessible to the host, you must create an initiator group and map the LUN to the initiator group. For FCP, when the initiator HBA was installed in the host and the host was brought online on the SAN, its WWPN was registered with the switch’s Fabric Name Service. You can see what initiators are available to add to an initiator group by using the `fcpl show initiator`. For iSCSI, use the `iscsi show initiator` command.

The following table starts with this step.

To do this...	Use this command...
Display the unique identifiers of the initiators that are available to create initiator groups	<pre>fcpl show initiator</pre> Sample result: <pre>Initiators connected on adapter 4a: Portname Group 10:00:00:00c:2b:cc:92 10:00:00:00c:2b:cc:35 10:00:00:00c:2b:4d:f6 10:00:00:00c:2b:6b:38</pre> OR <pre>iscsi show initiator</pre> Sample result: <pre>Initiators connected on adapter iswta: Tgt_PG iSCSI Initiator Name / ISID 2 iqn.1986-03.com.intel:rtpqa-intel-2:hba-1 / 40:00:02:da:00:02</pre>
Determine which hosts are associated with the WWPNs	To determine which WWPN is associated with a host, use the NetApp-supplied commands you downloaded from the CD-ROM in the HBA kit, such as <code>sanlun</code> for Solaris hosts or <code>lputilnt</code> for Windows hosts. See the HBA installation and setup guide for your host for instructions about how to use the commands. Alternatively, if you recorded the IEEE address of the initiator (host) HBA, you can obtain its WWPN by prepending the address with “10:00”.

To do this...	Use this command...
Create an initiator group	<code>igroup create { -f -i } -t <i>ostype</i> <i>groupname</i> <i>unique_identifier</i></code> -f indicates that the initiator group contains Fibre Channel WWPNs, written as 16 hexadecimal characters. -i indicates that the initiator group contains iSCSI nodenames. -t indicates that the operating system type follows. <i>ostype</i> is the name of the operating system of the initiator: <code>default</code>, <code>solaris</code>, <code>windows</code>, or <code>linux</code>. If <code>default</code> is selected, the initiator group can contain initiators with different types of operating systems. <i>groupname</i> is the name you specify as the name of the initiator group. <i>unique_identifier</i> is either ◆ <i>WWPN</i>, which is the 64-bit address of the initiator's port name. ◆ <i>iSCSI nodename</i>, which is a string of characters defined as follows: <code>iqn.yyyy-mm.backward_naming_authority:unique_device_name</code> <i>yyyy-mm</i> is the year and month in which the naming authority acquired the domain name. <i>backward_naming_authority</i> is the reverse domain name of the entity responsible for naming this device. For example, <code>com.netapp</code>. <i>unique_device_name</i> is the unique name for this device assigned by the naming authority. Example: <pre>igroup create -f -t solaris solaris-igroup3 10:00:00:00:2b:cc:92</pre>

To do this...	Use this command...
Create a LUN	<pre>lun create -s size -t ostype lunname</pre> -s indicates the size of the LUN to be created, in bytes by default. You can use one of the following characters to multiply the size: c (bytes), w (words or double bytes), b (512), k (kilobytes), m (megabytes), g (gigabytes), t (terabytes). -t indicates that the operating system type follows. <i>ostype</i> values are <i>image</i> (default), which indicates that the contents of the LUN are not specified; <i>solaris</i>, which indicates that the LUN will store a Solaris raw disk in a single-slice partition, <i>windows</i>, which indicates that the LUN will store a raw disk device in a single-partition Windows disk, or <i>linux</i>, which indicates that the LUN will store a raw disk with no partition table. <i>lunname</i> is the LUN's path name that includes the volume and qtree. Example: <pre>lun create -s 4g -t solaris /vol/vol1/qtree1/lun3</pre> Result: A 4-GB LUN called /vol/vol1/qtree1/lun3 is accessible by a Solaris host. In general, when you format and partition LUNs as a disk on a host, the size of the disk is smaller due to the overhead required by the host.

To do this...	Use this command...
Map the LUN to an initiator group	<pre>lun map lunname initiator-group-name [LUN-ID]</pre> <i>lunname</i> is the path name of the LUN you created. <i>initiator-group-name</i> is the name of the initiator group you created. <i>LUN-ID</i> is the identification number that you specify the initiator should use when the LUN is mapped to it. The range of the number depends on the operating system of the initiator: ◆ For Solaris, the range is 0–4096. ◆ For Windows, the range is 0–254. ◆ For Linux, the range is 0– Example 1: <code>lun map /vol/vol1/mtree1/lun3 solaris-igroup3 0</code> Result: Data ONTAP maps /vol/vol1/mtree1/lun3 to the initiator group solaris-igroup3 at LUN ID 0. Example 2: <code>lun map /vol/vol1/lun0 solaris-igroup0</code> Result: Data ONTAP assigns the next lowest valid LUN ID to map the LUN to the initiator group. After entering the command in this example, Data ONTAP displays the following message: <pre>lun map: auto-assigned solaris-igroup0=0</pre>
Display the LUNs you created	<pre>lun show [-v -m] [all mapped offline online unmapped lun-path -g initiator-group]</pre> Example: <code>lun show -v</code> Sample result: <pre>/vol/vol1/mtree1/lun3 4g (4294967296) (r/w, online, mapped) Serial#: 0dCfh3bgaBTU Share: none Space Reservation: enabled Multiprotocol Type: solaris Maps: solaris-igroup0=0</pre>

To do this...	Use this command...																		
Display the LUN ID mapping	<pre>lun show -m</pre> Sample result: <table><tr><th>LUN path</th><th>Mapped to</th><th>LUN ID</th></tr><tr><td colspan="3">-----</td></tr><tr><td>/vol/vol1/lun0</td><td>solaris-igroup0</td><td>0</td></tr><tr><td>/vol/vol1/lun1</td><td>solaris-igroup1</td><td>0</td></tr><tr><td>/vol/vol1/qtrees1/lun2</td><td>solaris-igroup2</td><td>0</td></tr><tr><td>/vol/vol1/qtrees1/lun3</td><td>solaris-igroup3</td><td>0</td></tr></table> In this example, each LUN is mapped to a different initiator group, which was named to reflect the host whose WWPN was used when creating the initiator group. Note that the LUN ID mapping begins with 0 for each initiator group. If a second LUN is mapped to the same initiator group, the LUN ID automatically increments by one.	LUN path	Mapped to	LUN ID	-----			/vol/vol1/lun0	solaris-igroup0	0	/vol/vol1/lun1	solaris-igroup1	0	/vol/vol1/qtrees1/lun2	solaris-igroup2	0	/vol/vol1/qtrees1/lun3	solaris-igroup3	0
LUN path	Mapped to	LUN ID																	

/vol/vol1/lun0	solaris-igroup0	0																	
/vol/vol1/lun1	solaris-igroup1	0																	
/vol/vol1/qtrees1/lun2	solaris-igroup2	0																	
/vol/vol1/qtrees1/lun3	solaris-igroup3	0																	

About this chapter This chapter describes how to manage LUNs, change LUN attributes, and display LUN statistics.

Topics in this chapter This chapter discusses the following topics

- ◆ “[Managing LUNs](#)” on page 46
- ◆ “[Modifying LUN attributes](#)” on page 50
- ◆ “[Displaying LUN information](#)” on page 54
- ◆ “[Monitoring disk space](#)” on page 57

Managing LUNs

Management tasks

You can perform the following tasks to manage LUNs.

- ◆ Manage LUN maps
- ◆ Map a LUN to an initiator group to make it available to the initiators in the initiator group
- ◆ Unmap a LUN from an initiator group to make it unavailable to the initiators in the initiator group
- ◆ Bring a LUN online to make it available for block protocol access (or SAN)
- ◆ Take a LUN offline to make it available for file protocol access (or NAS)
- ◆ Remove (destroy) a LUN
- ◆ Access a LUN with NAS protocols

How to manage LUN maps

To support a large number of hosts and to enforce access control, Data ONTAP maintains a separate map for each initiator group. This map controls which LUN is accessible to a given initiator group at a given LUN ID.

After you create a LUN, you must map it to the specific initiator group for it to be accessible to the hosts in that initiator group. Use the following guidelines when mapping LUNs:

- ◆ You can map two different LUNs to the same LUN ID for two different initiator groups without having a conflict.
- ◆ You can map a LUN only once to an initiator group.
- ◆ You cannot map two LUNs to an initiator group at the same LUN ID.

Using the `lun map` command automatically brings a LUN online. The `lun online` and `lun offline` commands enable and disable block protocol accesses while preserving mappings.

Mapping a LUN to an initiator group

The LUN ID is mapped to the initiators (hosts) in the initiator group. You can specify a LUN ID or you can use the number automatically assigned by Data ONTAP, which begins with 0 and increments by 1 every time a new LUN ID needs to be assigned. Network Appliance recommends that you allow Data ONTAP to automatically assign LUN IDs.

To allow Data ONTAP to map a LUN to an initiator group at an automatically assigned LUN ID, complete the following step.

Step	Action
1	Enter the following command: lun map <i>lun-path</i> <i>igroup</i> LUN_ID Example: <code>lun map /vol/vol1/qtrees1/lun3 solhost2</code> Result: Data ONTAP maps lun3 to the initiator group solhost2 at the first available LUN ID, in this case LUN ID 0.

Unmapping a LUN from an initiator group

To remove the mapping of a LUN from an initiator group, complete the following steps.

Step	Action
1	Enter the following command: lun offline <i>lun-path</i> Example: <code>lun offline /vol/vol1/lun1</code>
2	Enter the following command: lun unmap <i>lun-path</i> <i>igroup</i> LUN_ID Example: <code>lun unmap /vol/vol1/lun1 solaris-igroup0 0</code>

Bringing a LUN online

To bring one or more LUNs online, complete the following step.

Step	Action
1	Enter the following command: lun online <i>lun-path</i> [<i>lun-path</i> ...] Example: <code>lun online /vol/vol1/lun0</code>

Taking a LUN offline Taking a LUN offline makes it unavailable for block protocol access; however, you can write to an offline LUN using NAS protocols. To take a LUN offline, complete the following step.

Step	Action
1	Enter the following command: <code>lun offline <i>lun-path</i> [<i>lun-path</i> ...]</code> Example: <code>lun offline /vol/vol1/lun0</code>

Removing a LUN To remove one or more LUNs, complete the following step.

Step	Action
1	Remove one or more LUNs by entering the following command: <code>lun destroy [-f] <i>lun-path</i> [<i>lun-path</i> ...]</code> -f forces the lun destroy command to execute even if the LUNs specified by one or more <i>lun-paths</i> are mapped or are online. Without the -f parameter, you must first take the LUN offline and unmap it, and then enter the <code>lun destroy</code> command.

Accessing a LUN with NAS protocols When you create a LUN, it can only be accessed with SAN protocols by default. However, you can make a LUN available to a host using NAS protocols if NAS protocols are licensed and enabled on the filer. The usefulness of accessing a LUN over NAS protocols is host application-dependent.

Note _____
A LUN cannot be extended or truncated using NFS or CIFS protocols.

If you want to write to a LUN over NAS protocols, you must take the LUN offline or unmap it to prevent overwriting data in the LUN using the SAN protocol. To make a LUN accessible to a host using a NAS protocol, complete the following step.

Step	Action
1	Determine whether you want to read or write to the LUN over the NAS protocol and take appropriate action: ◆ If you want read access, the LUN can remain online. ◆ If you want write access, ensure that the LUN is offline or unmapped.
2	Enter the following command: <pre>lun share <i>lun-path</i> {none read write all}</pre> Example: <code>lun share /vol/vol1/qtrees1/lun2 read</code> Result: The LUN is now readable over NAS.

Modifying LUN attributes

Attributes you can modify

You can modify the following attributes of a LUN:

- ◆ Name
- ◆ Space reservations setting
- ◆ Description (or comment)
- ◆ Size
- ◆ Serial number

Renaming a LUN

To rename a LUN, completing the following step.

Step	Action
1	Enter the following command: <code>lun move lun_path new_lun_path</code> Example: <code>lun move mylun mynewlun</code> Note_____ The existing <code>lun_path</code> and the new <code>lun_path</code> must be in the same qtree. _____

Enabling or disabling space reservations for LUNs

To enable or disable space reservations for a LUN, complete the following step.

Caution_____

If you disable space reservations, you might experience write failures due to insufficient disk space. This might lead to data corruption.

When write failures occur, Data ONTAP displays system messages (one message per file) to the console or sends these messages to log files and other remote systems, as specified by its `/etc/syslog.conf` configuration file.

Step	Action
1	Enter the following command: <pre>lun set reservation <i>lun_path</i> [enable disable]</pre> <i>lun_path</i> is the LUN in which space reservations are to be set. This must be an existing LUN. Note_____ Enabling space reservation on a LUN fails if there is not enough free space in the volume for the new reservation.

Modifying the LUN description

To modify the LUN description, complete the following step.

Step	Action
1	Enter the following command: <pre>lun comment <i>lun-path</i> [<i>comment</i>]</pre> Example: lun comment /vol/vol1/lun2 "10GB for payroll records" Note_____ If you use spaces in the comment, enclose the comment in quotation marks.

Resizing a LUN

You can increase or decrease the size of a LUN; however, the host operating system must be able to recognize changes to its disk partitions.

Caution

Before resizing a LUN, ensure that this feature is compatible with the host operating system.

Note the following:

- ◆ On Windows, resizing is only supported on basic disks. Resizing is not supported on dynamic disks.

- ◆ On Solaris, decreasing the LUN size is not supported. Only increasing the LUN size is supported.
- ◆ For VxVM, resizing LUNs is not supported.

To change the size a LUN, complete the following steps.

Step	Action
1	Take a LUN offline before resizing it by entering the following command: lun offline <i>lun-path</i> Example: lun offline /vol/vol1/qtrees/lun2
2	Change the size of the LUN by entering the following command: lun resize [-f] <i>lun-path new-size</i> -f overrides warnings when decreasing the size of the LUN. Example: (Assuming that lun2 is 5 GB and you are increasing it to 10 GB lun resize /vol/vol1/qtrees1/lun2 10g
3	From the host, rescan or rediscover the LUN so that the new size is recognized. Example: For Solaris hosts, use Solaris commands to erase the disk label and create a new one, then use the Solaris mkfs command to increase the file system on the disk slice to fill the newly enlarged slice.

Modifying a LUN's serial number

To modify a LUN's serial number, complete the following steps.

Step	Action
1	Take the LUN offline by entering the following command: lun offline <i>lun-path</i> Example: lun offline /vol/vol1/qtrees1/lun2

Step	Action
2	Modify a LUN's serial number by entering the following command: lun serial <i>lun-path new-serial-number</i> Example: lun serial /vol/vol1/qtree1/lun2 0cmePIdQxjcH

Displaying LUN information

Types of information you can display

You can display the following types of information about LUNs:

- ◆ Command-line help about LUN commands
- ◆ Read, write, and operations per second statistics
- ◆ Statistics and characteristics
- ◆ Space reservations setting

Displaying command-line help

To display command-line help, complete the following steps.

Step	Action																																		
1	On the filer's command line, enter the following command: lun help Result: A list of all LUN subcommands is displayed: <table><tbody><tr><td>lun help</td><td>- List LUN (logical unit of block storage) commands</td></tr><tr><td>lun clone</td><td>- Manage LUN cloning</td></tr><tr><td>lun comment</td><td>- Display/Change descriptive comment string</td></tr><tr><td>lun create</td><td>- Create a LUN</td></tr><tr><td>lun destroy</td><td>- Destroy a LUN</td></tr><tr><td>lun map</td><td>- Map a LUN to an initiator group</td></tr><tr><td>lun move</td><td>- Move (rename) LUN</td></tr><tr><td>lun offline</td><td>- Stop block protocol access to LUN</td></tr><tr><td>lun online</td><td>- Restart block protocol access to LUN</td></tr><tr><td>lun resize</td><td>- Resize LUN</td></tr><tr><td>lun serial</td><td>- Display/change LUN serial number</td></tr><tr><td>lun set</td><td>- Manage LUN properties</td></tr><tr><td>lun setup</td><td>- Initialize/Configure LUNs, mapping</td></tr><tr><td>lun share</td><td>- Configure NAS file-sharing properties</td></tr><tr><td>lun show</td><td>- Display LUNs</td></tr><tr><td>lun stats</td><td>- Displays or zeros read/write statistics for LUN</td></tr><tr><td>lun unmap</td><td>- Remove LUN mapping</td></tr></tbody></table>	lun help	- List LUN (logical unit of block storage) commands	lun clone	- Manage LUN cloning	lun comment	- Display/Change descriptive comment string	lun create	- Create a LUN	lun destroy	- Destroy a LUN	lun map	- Map a LUN to an initiator group	lun move	- Move (rename) LUN	lun offline	- Stop block protocol access to LUN	lun online	- Restart block protocol access to LUN	lun resize	- Resize LUN	lun serial	- Display/change LUN serial number	lun set	- Manage LUN properties	lun setup	- Initialize/Configure LUNs, mapping	lun share	- Configure NAS file-sharing properties	lun show	- Display LUNs	lun stats	- Displays or zeros read/write statistics for LUN	lun unmap	- Remove LUN mapping
lun help	- List LUN (logical unit of block storage) commands																																		
lun clone	- Manage LUN cloning																																		
lun comment	- Display/Change descriptive comment string																																		
lun create	- Create a LUN																																		
lun destroy	- Destroy a LUN																																		
lun map	- Map a LUN to an initiator group																																		
lun move	- Move (rename) LUN																																		
lun offline	- Stop block protocol access to LUN																																		
lun online	- Restart block protocol access to LUN																																		
lun resize	- Resize LUN																																		
lun serial	- Display/change LUN serial number																																		
lun set	- Manage LUN properties																																		
lun setup	- Initialize/Configure LUNs, mapping																																		
lun share	- Configure NAS file-sharing properties																																		
lun show	- Display LUNs																																		
lun stats	- Displays or zeros read/write statistics for LUN																																		
lun unmap	- Remove LUN mapping																																		

Step	Action
2	To display the syntax for any of the subcommands, enter the following command: <pre>lun help <i>subcommand</i></pre> Example: <code>lun help show</code>

Displaying statistics

To display the number of data reads, data writes, and operations per second for LUNs, complete the following step.

Step	Action
1	Enter the following command: <pre>lun stats {-z -k} [<i>lun-path</i> ...]</pre> -k displays the statistics in KBs. Note _____ The statistics start at zero at boot time. _____ -z resets the counter at any time. Sample Result: <pre>lun stats -k /vol/vol1/lun0 Read (kBytes) Write (kBytes) Ops 10241 26159 3290 /vol/vol1/lun1 Read (kBytes) Write (kBytes) Ops 10241 26279 3313</pre>

Displaying status and characteristics

To display status and characteristics about LUNs, complete the following steps.

Step	Action															
1	On the filer’s command line, enter the following command to display LUN status and characteristics: lun show -v Sample display: <pre> /vol/vol1/lun0 5g (5368709120) (r/w, online, mapped) Comment: "First LUN" Serial#: 0cmePidQwRb1 Share: none Space Reservations: enabled Multiprotocol Type: solaris Maps: solaris-igroup0=0 /vol/vol1/lun1 100m (104857600) (r/w, online, mapped) Comment: "Second LUN" Serial#: 0cmePidQuEtX Share: none Space Reservations: enabled Multiprotocol Type: solaris Maps: solaris-igroup1=0</pre>															
2	On the filer’s command line, enter the following command to display LUN mapping information: lun show -m Result: <table><thead><tr><th>LUN path</th><th>Mapped to</th><th>LUN ID</th></tr></thead><tbody><tr><td>/vol/vol1/lun0</td><td>solaris-igroup0</td><td>0</td></tr><tr><td>/vol/vol1/lun1</td><td>solaris-igroup1</td><td>0</td></tr><tr><td>/vol/vol1/qtrees1/lun2</td><td>solaris-igroup2</td><td>0</td></tr><tr><td>/vol/vol1/qtrees1/lun3</td><td>solaris-igroup2</td><td>1</td></tr></tbody></table>	LUN path	Mapped to	LUN ID	/vol/vol1/lun0	solaris-igroup0	0	/vol/vol1/lun1	solaris-igroup1	0	/vol/vol1/qtrees1/lun2	solaris-igroup2	0	/vol/vol1/qtrees1/lun3	solaris-igroup2	1
LUN path	Mapped to	LUN ID														
/vol/vol1/lun0	solaris-igroup0	0														
/vol/vol1/lun1	solaris-igroup1	0														
/vol/vol1/qtrees1/lun2	solaris-igroup2	0														
/vol/vol1/qtrees1/lun3	solaris-igroup2	1														

Displaying space reservation status for LUNs

To display the status of space reservations for LUNs in a volume, complete the following step.

Step	Action
1	Enter the following command: <pre>lun set reservation <i>lun-path</i></pre>

Monitoring disk space

Command to use

You use the `df [option] [pathname]` command to monitor the amount of free disk space that is available on one or all volumes on a filer. The amount of space is displayed in 1,024-byte blocks by default. You use the `-k`, `-m`, `-g`, or `-t` options to have space displayed in KB, MB, GB, or TB format, or you can use the `-h` option to let Data ONTAP scale the size to the appropriate format, based on the value. The `-r` option is explained below.

What the command output displays

The output of the `df` command displays four columns of statistics about the active file system in the volume and the snapshot directory for that volume. Blocks that are referenced by both the active file system and by one or more snapshots are counted only in the active file system (not in the snapshot). The statistics displayed are as follows:

- ◆ Amount of total space on the volume
- ◆ Amount of used space
- ◆ Amount of available space
- ◆ Percentage of the volume capacity being used

The `-r` option changes the last column to report on the amount of reserved space. The amount of reserved space, plus the amount of used space, plus the amount of available space is equal to the amount of total space in the volume.

When the snapshot reserve option is set to zero and snapshots are taken, the amount of space displayed in the total column is larger than what is actually present. This occurs because the space consumed by snapshots is reported as “used” by the active file system in addition to the snapshot.

Example of monitoring disk space

The following example illustrates how to use the `df -r` command to monitor disk space on a volume throughout a scenario of creating a volume, creating a LUN, writing data to the LUN, taking a snapshot, creating a writable snapshot of the LUN, and creating a clone of that LUN. Assuming that the amount of disk space required for the LUN is 3GB, then the information from [“Guidelines to calculate the size of a volume”](#) on page 21 recommends that the volume size be approximately four times 3GB plus the amount of change expected between snapshots. Assuming this rate is approximately 15%, the volume size is

calculated as $4 \times 3\text{GB} + (3\text{GB} \times .15 \times 2)$, or 12.9GB. If the size of the disks you plan to use is 72GB (which typically results in 67.9 GB of physical capacity), two disks provide more than enough space, one for data and one for parity.

This example does not include every step required to properly configure the filer or how to perform tasks on the host.

In the example, the filer is named *toaster* and the host is named *solaris_host*. In addition, there are two HBAs installed in the host, lpfc0 (WWPN: 10000000c92f6cd0) and lpfc1 (WWPN:10000000c92f6a24). This information is needed to create an igroup. Also, this example assumes that the original LUN requires 3 GB of disk space.

Step	Action																		
1	From the filer, create a new volume named <i>volspace</i> that has approximately 67 GB and observe the effect on disk space by entering the following commands: <pre>toaster> vol create volspace 2</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. There is a snap reserve of 20% on the volume even though the volume will be used for SAN or iSCSI LUNs. <table><tr><th>Filesystem</th><th>kbytes</th><th>used</th><th>avail</th><th>reserved</th><th>Mounted on</th></tr><tr><td>/vol/volspace</td><td>50119928</td><td>176</td><td>50119752</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>12529980</td><td>0</td><td>1252990</td><td>0</td><td>/vol/vol space/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace	50119928	176	50119752	0	/vol/volspace/	/vol/volspace/.snapshot	12529980	0	1252990	0	/vol/vol space/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace	50119928	176	50119752	0	/vol/volspace/														
/vol/volspace/.snapshot	12529980	0	1252990	0	/vol/vol space/.snapshot														
2	Set the percentage of snap reserve space to zero and observe the effect on disk space by entering the following commands: <pre>toaster> snap reserve volspace 0</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The amount of available snapshot space becomes zero, and the 20% of snapshot space is added to available space for /vol/volspace. <table><tr><th>Filesystem</th><th>kbytes</th><th>used</th><th>avail</th><th>reserved</th><th>Mounted on</th></tr><tr><td>/vol/volspace/</td><td>62649908</td><td>176</td><td>62649732</td><td>62649732</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td></td><td>0</td><td>0</td><td>0</td><td>0 /vol/vol space/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	176	62649732	62649732	/vol/volspace/	/vol/volspace/.snapshot		0	0	0	0 /vol/vol space/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	176	62649732	62649732	/vol/volspace/														
/vol/volspace/.snapshot		0	0	0	0 /vol/vol space/.snapshot														

Step	Action																		
3	Create a LUN (/vol/volspace/lun0) and observe the effect on disk space by entering the following commands: <pre>toaster> lun create -s 3g -t solaris /vol/volspace/lun0</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. Approximately 2GB of space is taken from available space and is displayed as reserved space for the LUN: <table><tr><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace/ volspace/</td><td>62649908</td><td>188</td><td>59500908</td><td>3148812</td><td>/vol/</td></tr><tr><td>/vol/volspace/.snapshot space/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>/vol/vol</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/ volspace/	62649908	188	59500908	3148812	/vol/	/vol/volspace/.snapshot space/.snapshot	0	0	0	0	/vol/vol
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/ volspace/	62649908	188	59500908	3148812	/vol/														
/vol/volspace/.snapshot space/.snapshot	0	0	0	0	/vol/vol														
4	Create an initiator group named <i>sol-igroup1</i> and map the LUN to the initiator group by entering the following commands (assuming your host has an HBA whose WWPN is 100000002bcc92). Depending on your host, you may need to create WWNN persistent bindings. These commands have no effect on disk space. <pre>toaster> igroup create -f -t solaris sol-igroup1 10:00:00:00:2b:cc:92</pre> <pre>toaster> lun map /vol/volspace/lun0 sol-igroup1 0</pre>																		
5	From the host, discover the LUNs, format them, and make the file system available to the host. For information on these procedures, see the <i>SAN Host Attach Kit Installation and Setup Guide</i> that came with your SAN Host Attach Kit.																		
6	From the filer, see the effect of creating the filesystem on the LUN by entering the following command: <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed: <table><tr><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace/ volspace/</td><td>62649908</td><td>183332</td><td>59500860</td><td>2965716</td><td>/vol/</td></tr><tr><td>/vol/volspace/.snapshot vol space/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>/vol/</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/ volspace/	62649908	183332	59500860	2965716	/vol/	/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/ volspace/	62649908	183332	59500860	2965716	/vol/														
/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/														
7	From the host, write data to the file system (the LUN on the filer).																		

Step	Action																		
8	From the filer, see the effect of writing to the LUN by entering the following command: toaster> df -r /vol/volspace Result: The following sample output is displayed: <table><tr><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace/ volspace/</td><td>62649908</td><td>654632</td><td>59500728</td><td>2494548</td><td>/vol/</td></tr><tr><td>/vol/volspace/.snapshot vol space/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>/vol/</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/ volspace/	62649908	654632	59500728	2494548	/vol/	/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/ volspace/	62649908	654632	59500728	2494548	/vol/														
/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/														
9	From the host, continue writing data to the file system until the LUN is full. As you continue writing data to the file system, you can observe the amount of reserved space and available space decrease and used space increase by executing the <code>df -r</code> command. Result: The following sample output is displayed: <table><tr><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace/ volspace/</td><td>62649908</td><td>2885456</td><td>59500004</td><td>264448</td><td>/vol/</td></tr><tr><td>/vol/volspace/.snapshot vol space/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>/vol/</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/ volspace/	62649908	2885456	59500004	264448	/vol/	/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/ volspace/	62649908	2885456	59500004	264448	/vol/														
/vol/volspace/.snapshot vol space/.snapshot	0	0	0	0	/vol/														

Step	Action																								
10	Take a snapshot named <i>snap1</i> of the active file system and observe the effect on disk space. Caution_____ From the host or the application, ensure that the active file system is in a quiesced or synchronized state prior to taking a snapshot. _____ Enter the following commands: <pre>toaster> snap create volspace snap1</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The amount of reserved space increases to approximately 3 GB to ensure that in addition to being able to write data to the LUN, data in the snapshot can be overwritten as well because when Data ONTAP writes to a block of snapshot data, both the old data and the new data must be stored. In this example, 722 KB of space is used for the snapshot. <table><tr><th>Filesystem</th><th>kbytes</th><th>used</th><th>avail</th><th>reserved</th><th>Mounted on</th></tr><tr><td>/vol/volspace/</td><td>62649908</td><td>2885652</td><td>56615440</td><td>3148816</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>772</td><td></td><td>0</td><td>0 /vol/</td></tr><tr><td>vol space/.snapshot</td><td></td><td></td><td></td><td></td><td></td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	2885652	56615440	3148816	/vol/volspace/	/vol/volspace/.snapshot	0	772		0	0 /vol/	vol space/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace/	62649908	2885652	56615440	3148816	/vol/volspace/																				
/vol/volspace/.snapshot	0	772		0	0 /vol/																				
vol space/.snapshot																									
11	Create a writable snapshot of LUN0 named <i>snaplun</i> and observe the effect on disk space by entering the following commands: <pre>toaster> lun create -b /vol/volspace/.snapshot/snap1/lun0 /vol/volspace/snaplun</pre> Note_____ After you create this LUN, you must go through the same procedure on the host to discover the new LUN, format it, and so on, except for creating a new file system on it. _____ <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The amount of reserved space increases to approximately 6 GB to ensure that both LUNs can be written to, and a small amount of space is increased for the snapshot. <table><tr><th>Filesystem</th><th>kbytes</th><th>used</th><th>avail</th><th>reserved</th><th>Mounted on</th></tr><tr><td>/vol/volspace/</td><td>62649908</td><td>2886408</td><td>53465868</td><td>6297632</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>792</td><td></td><td>0</td><td>0 /vol/</td></tr><tr><td>vol space/.snapshot</td><td></td><td></td><td></td><td></td><td></td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	2886408	53465868	6297632	/vol/volspace/	/vol/volspace/.snapshot	0	792		0	0 /vol/	vol space/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace/	62649908	2886408	53465868	6297632	/vol/volspace/																				
/vol/volspace/.snapshot	0	792		0	0 /vol/																				
vol space/.snapshot																									

Step	Action																								
12	Display the status of both LUNs by entering the following command: <pre>toaster> lun show -v</pre> Result: The following sample output is displayed: <pre>/vol/volSPACE/lun0 3g (3221225472) (r/w, online, mapped) Serial#: 0dDICn1A-Z0J Share: none Space Reservation: enabled Multiprotocol Type: solaris Maps: sol-igroup1=0 /vol/volSPACE/snaplun 3g (3221225472) (r/w, online) Serial#: 0dDICn1RRA0J Backed by: /vol/volSPACE/.snapshot/snap1/lun0 Share: none Space Reservation: enabled Multiprotocol Type: solaris Maps: sol-igroup1=0</pre>																								
13	Create a clone of the writable snapshot of the LUN you created in step 11 and observe the effect on disk space by entering the following commands: <pre>toaster> lun clone start /vol/volSPACE/snaplun toaster> df -r /vol/volSPACE</pre> Result: The following sample output is displayed. The amount of used space increases to 45 GB, available space decreases, and again, a small amount of space is used by snapshot. <table><tr><th>Filesystem</th><th>kbytes</th><th>used</th><th>avail</th><th>reserved</th><th>Mounted on</th></tr><tr><td>/vol/volSPACE/</td><td>62649908</td><td>4538176</td><td>51814100</td><td>6297632</td><td>/vol/volSPACE/</td></tr><tr><td>/vol/volSPACE/.snapshot</td><td>0</td><td>860</td><td></td><td>0</td><td>0 /vol/</td></tr><tr><td>vol space/.snapshot</td><td></td><td></td><td></td><td></td><td></td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volSPACE/	62649908	4538176	51814100	6297632	/vol/volSPACE/	/vol/volSPACE/.snapshot	0	860		0	0 /vol/	vol space/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volSPACE/	62649908	4538176	51814100	6297632	/vol/volSPACE/																				
/vol/volSPACE/.snapshot	0	860		0	0 /vol/																				
vol space/.snapshot																									

The effect of LUN and snapshot operations on disk space

The following table is another example of monitoring disk space. It illustrates the effect on disk space when you create a volume, create a LUN, write data to the LUN, take snapshots of the LUN, and expand the size of the volume. For this example, assume space reservation is enabled and snap reserve is set to 0%.

Action	Result	Comment
Create a 100-GB volume.	Used space = 0 GB Reserved space = 0 GB Available space = 100 GB Volume Total: 100 GB Snapshot creation is allowed.	
Create a 40-GB LUN on that volume.	Used space = 0 Reserved space = 40 Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	If the LUN size was limited to accommodate at least one snapshot when it was created, then it will always be less than one half of the volume size.
Write 40 GB of data to the LUN.	Used space = 40 GB Reserved space = 0 GB Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	When you write data to the LUN, it counts against the running Used total. The sum of Used, Reserved, and Available always equals Volume Total.
Create a snapshot of the LUN.	Used space = 40 GB Reserved space = 40 GB Available space = 20 GB Volume Total: 100 GB Snapshot succeeds.	The snapshot locks all the data on the LUN so that even if that data is later deleted, it remains in the snapshot until the snapshot is deleted. After a snapshot is created, the reserved space must now be large enough to ensure that any future writes to the LUN succeed.

Action	Result	Comment
Overwrite all 40 GB of data on the LUN with new data.	Used space = 60 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 100 GB Snapshot creation is blocked.	The amount of used space on the volume increases, because the original 40 GB of data belongs to the snapshot and therefore continues to count against the Used total. Reserved space must be equal to the size of the LUN (40 GB), and reserved and used space together cannot exceed the size of the volume. As a result, used space is displayed as 60 GB rather than the expected 80 GB. However, all data is preserved. You cannot take a snapshot now, because no space is available. That is, all space is used by data or held in reserve so that any and all changes to the content of the LUN can be written to the volume.
Expand the volume by 100 GB.	Used space = 80 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	After you expand the volume, free space becomes available again. Therefore, snapshot creation is no longer blocked. In addition, the Used and Available totals are adjusted to reflect the fact that reserved space is no longer being used to hold disk data.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 80 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	Because none of the overwritten data belongs to a snapshot, it disappears when the new data replaces it. As a result, the Used total remains unchanged.
Create a snapshot of the LUN.	Used space = 80 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	The snapshot locks all 40 GB of data currently on the LUN. A total of 80 GB of data now belongs to the two snapshots of the LUN. Reserved space remains the size of the LUN, or 40 GB.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 120 GB Reserved space = 40 GB Available space = 40 GB Volume Total: 200 GB Snapshot creation is allowed.	Because the data being replaced belongs to a snapshot, it remains on the volume.

Action	Result	Comment
Expand the LUN by 40 GB.	Used space = 120 GB Reserved space = 80 GB Available space = 0 GB Volume Total: 200 GB Snapshot creation is blocked.	The amount of reserved space increases to match the expanded size of the LUN. This guarantees that the entire contents of the LUN can be written to the volume. Because the available space has decreased to 0, snapshot creation is blocked.
Delete both snapshots of the volume.	Used space = 80 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 200 GB Snapshot creation is allowed.	The 80 GB of data locked by the two snapshots disappears from the Used total when the snapshots are deleted. Because there are no more snapshots of this LUN, the reserved space decreases to 40 GB, enough to guarantee any future write operations. Snapshot creation is once again allowed.
Delete the LUN.	Used space = 0 GB Reserved space = 0 GB Available space = 200 GB Volume Total: 200 GB	Because no snapshots exist for this volume, deletion of the LUN causes the Used space to decrease to 0 GB.

About this chapter This chapter explains how to create and manage initiator groups.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Understanding initiator groups](#)” on page 68
- ◆ “[Managing initiator groups](#)” on page 70

Understanding initiator groups

About initiator groups

Initiator groups are tables of unique identifiers (WWPNs or iSCSI nodenames) of initiators that you specify have access to LUNs. Initiator groups can have multiple initiators, and multiple initiator groups can have the same initiator.

Information needed to create an initiator group

When you create an initiator group, you need to specify the following information:

- ◆ A case-sensitive name that meets the following requirements:
 - v Contains 1 to 96 alphanumeric characters.
 - v Can contain any except the following special characters: &, #, -, ' , " , blank, or tab.

Note

You might find it useful to provide meaningful names for initiator groups, ones that describe the hosts that can access the LUNs mapped to them. However, the name you assign to an initiator group is independent of the name of the host. If you name an initiator group sun-1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

- ◆ The type of initiator group (FCP if it contains WWPNs, iSCSI if it does not)

Note

An initiator group can only contain one type of node, either FCP or iSCSI nodes.

- ◆ The unique identifier of the initiator (FCP WWPN or iSCSI nodename)
- ◆ The operating system of the initiator

Using initiator groups to control access to LUNs

You use initiator groups to control host access to LUNs. Typically, you only want one host to have access to a LUN. However, if you are using multipathing software, or have clustered hosts, you need to have redundant paths to the same LUN. The following table illustrates how four initiator groups are used to give access to the LUNs for four different hosts using the FCP service.

In the following table, note that the clustered hosts (Host3 and Host4) can use the same initiator group for common access to the LUNs mapped to solaris-group2. Also note that solaris-group3 contains the WWPNs of Host4 to store local information not intended to be seen by its partner.

Host with HBA WWPNs	Initiator groups	WWPNs added to initiator groups	LUNs mapped to initiator groups
Host1, single-path (one HBA) 10:00:00:00:c9:2b:7c:0f	solaris-group0	10:00:00:00:c9:2b:7c:0f	/vol/vol1/lun0
Host2, multipath (two HBAs) 10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	solaris-group1	10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	/vol/vol1/lun1
Host3, multipath, clustered (connected to Host4) 10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02	solaris-group2	10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol1/qtree1/lun2
Host4, multipath, clustered (connected to Host3) 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	solaris-group3	10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol1/qtree1/lun2 /vol/vol1/qtree1/lun3

Managing initiator groups

Tasks to manage initiator groups

You can use the command-line interface or FilerView to

- ◆ Create initiator groups
- ◆ Destroy initiator groups
- ◆ Add initiators (through their WWPNs) from initiator groups
- ◆ Remove initiators (through their WWPNs) from initiator groups
- ◆ Display all the initiators in an initiator group
- ◆ Set the operating system type (ostype) for an initiator group

Note

When using the command line, you can get command-line syntax help from the command line.

Creating an initiator group

To create an initiator group, complete the following step.

Step	Action
1	Enter the following command: <pre>igroup create -f -i [-t <i>ostype</i>] <i>igroup</i> [<i>unique identifier</i>...]</pre> <i>-f</i> indicates that it is an FCP initiator group. <i>-i</i> indicates that it is an iSCSI initiator. <i>ostype</i> indicates the operating system of the host (Windows or Solaris). <i>igroup</i> is the name of the initiator group you specify. <i>unique identifier</i> is either an iSCSI node name or an FCP WWPN. Example: <code>igroup create -f -t solaris solaris-group0 10:00:00:00:c9:2b:6b:3c</code> Result: You created an initiator group called solaris-group0 that contains the WWPN of the Solaris host associated with that WWPN.

Destroying an initiator group

To destroy one or more existing initiator groups, complete the following step.

Note

An initiator group cannot be destroyed if any LUNs are mapped to it, unless you use the `-f` parameter.

Step	Action	
1	If you want to...	Then enter this command...
	Remove LUNs mapped to an initiator group before deleting the initiator group	lun unmap <i>lun-path igroup</i> Example: lun unmap /vol/vol2/qtrees/LUN10 solaris-group5
	Delete one or more initiator groups	igroup destroy igroup [<i>igroup</i>, ...] Example: igroup destroy solaris-group5
	Remove all LUN maps for an initiator group and delete the initiator group with one command	igroup destroy -f igroup [<i>igroup</i> ...] Example: igroup destroy -f solaris-group5

Adding an initiator To add an initiator to an initiator group, complete the following step.

Step	Action
1	Enter the following command: igroup add <i>igroup WPN</i> Caution_____ When adding initiators to an initiator group, ensure that each initiator sees, at most, one LUN at a given LUN ID. _____ Example: igroup add solaris-group2 10:00:00:00:c9:2b:02:1f Result: You added the second port of Host2 to the initiator group solaris-group2.

Removing an initiator To remove an initiator from an initiator group, complete the following step.

Step	Action
1	Enter the following command: igroup remove <i>igroup</i> <i>WWPN</i> Example: <code>igroup remove solaris-group1 10:00:00:00:c9:2b:7c:0f</code>

Displaying initiators To display all the initiators in the specified initiator group, complete the following step.

Step	Action
1	Enter the following command: igroup show [<i>igroup</i>] Example: <code>igroup show solaris-group3</code>

Setting the ostype To set the operating system type (ostype) for an initiator group to Solaris or Windows, complete the following step.

Step	Action
1	Enter the following command. igroup set <i>igroup</i> <i>ostype</i> Example: <code>igroup set solaris-group3 solaris</code>

Getting command-line syntax help To get command-line syntax help, complete the following step.

Step	Action
1	Enter the following command: igroup help <i>subcommand</i>

About this chapter

This chapter provides information about how to use Data ONTAP data protection features in SAN environments.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Data ONTAP data protection processes”](#) on page 74
- ◆ [“Using snapshots”](#) on page 75
- ◆ [“Using LUN clones”](#) on page 77
- ◆ [“Deleting busy snapshots”](#) on page 79
- ◆ [“Using SnapRestore”](#) on page 85
- ◆ [“Using SyncMirror”](#) on page 91
- ◆ [“Using volume copy”](#) on page 92
- ◆ [“Using SnapMirror”](#) on page 93
- ◆ [“Using SnapVault”](#) on page 96
- ◆ [“Backing up data to tape”](#) on page 101
- ◆ [“Using NVFAIL”](#) on page 105
- ◆ [“Using ndmpcopy”](#) on page 107

Data ONTAP data protection processes

Processes supported in the SAN environment

The Data ONTAP data protection applications perform backup and restore functions on LUNs that you define as part of setting up your SAN networks. Data ONTAP offers the following data protection processes:

- ◆ Snapshot™ feature
- ◆ LUN cloning
- ◆ SnapRestore™ feature
- ◆ SyncMirror™ feature
- ◆ Volume copy
- ◆ SnapMirror™ feature
- ◆ SnapVault™ feature
- ◆ Backup to and restore from tape
- ◆ NVFAIL
- ◆ NDPM copy

The following sections explain the basics of what these applications do and how to use them in a SAN environment. You can find detailed explanations of each application in the *Data ONTAP 6.4 System Administration Data Protection Guide*.

Using snapshots

What snapshots are Snapshots are a standard feature of Data ONTAP. A snapshot is a frozen, read-only image of the entire Data ONTAP file system (or WAFL® volume) that reflects the state of the file system at the time the snapshot was created.

Snapshots are the first line of defense for backing up and restoring data.

Caution

You must save contents of host file system buffers to disk before taking snapshots or using other data protection applications, such as dump, SnapMirror, SyncMirror, SnapVault, and so on, on the filer.

Reasons to take snapshots Data ONTAP takes snapshots so that you can back up LUNs to tape or to a SAN host that is acting as a backup server.

If data corruption is found, a LUN or set of LUNs can be restored from a snapshot using the SnapRestore feature. Snapshots can also be replicated to remote sites, using the SnapMirror feature, in case they are needed for disaster recovery.

How snapshots are used in the SAN environment Snapshots are always consistent with regard to the underlying WAFL file system. Prior to creating a snapshot, you must quiesce the application or file system using the LUN (which flushes the host filer system buffers to disk), so that the snapshot is consistent with regard to the application.

In Data ONTAP 6.3 and later, volumes can be space reserved, which means that if free space within the volume is below a certain threshold, snapshots cannot be taken.

Note

On Windows hosts, you can use SnapDrive 2.0 or later to save the contents of the host file system buffers to disk and create snapshots. See the *SnapDrive Installation and Administration Guide*.

**How writable
snapshots are used
to copy LUNs for
data protection**

Only writable LUNs can be mapped to a host. Therefore, to map a LUN that exists in a snapshot, the snapshot must be made writable, that is, it must be made available in the active WAFL file system.

Network Appliance recommends that you do not use scheduled snapshots to replicate LUNs on a regular basis because multiple LUNs from different hosts can exist on the same WAFL volume.

To create, find, delete, or rename specific snapshots, see the detailed information about snapshots in the *Data ONTAP 6.4 Data Protection Guide*.

Using LUN clones

Creating a writable LUN of a snapshot

Before you can clone a LUN, you must create a writable snapshot of the LUN you want to clone. To create a writable snapshot, complete the following steps.

Step	Action
1	Create a LUN by entering the following command: lun create -s size lun_path Example: lun create -s 100g /vol/vol1/lun0
2	Create a snapshot of the volume containing the LUN to be cloned by entering the following command: snap create volume_name snapshot_name Example: snap create vol1 mysnap
3	Create a writable LUN by entering the following command: lun create -b snapshot_name lun_clone_path Example: lun create -b /vol/vol1/.snapshot/mysnap/lun0 /vol/vol1/snaplun0

What a LUN clone is A LUN clone is a copy of a writable snapshot of a LUN.

What LUN clones are used for

LUN clones are used to If you want to delete a snapshot of a LUN without taking the LUN offline or losing its contents, you can clone the contents of the writable snapshot of the LUN without any disruption to an application using the LUN. After the clone operation is complete, you can delete the snapshot.

Note
Cloning is not NVLOG protected, so if the filer panics during a clone operation, the operation is restarted from the beginning on a reboot or takeover.

With the lun clone command, you start and stop the operation, display all the LUNs being cloned, and display the progress of the operation.

Starting the cloning process

To clone a LUN, complete the following step.

Step	Action
1	Begin the clone operation by entering the following command: lun clone start <i>lun_clone_path</i>

Displaying the progress of the clone operation

To display the progress of the clone operation, complete the following step.

Step	Action
1	Enter the following command: lun clone status <i>lun_clone_path</i>

Displaying all the clones

To display all the clones, complete the following step.

Step	Action
1	Enter either of the following commands: lun clone show [<i>lun-path</i>] lun show -c

Stopping the cloning process

If you need to stop the clone process, complete the following step.

Step	Action
1	Enter the following the command: lun clone stop <i>lun-path</i>

Deleting busy snapshots

What a snapshot in a busy state means

A snapshot is in a busy state if there are any LUNs backed by data in that snapshot. In other words, the snapshot contains data that is used by the LUN. These LUNs can exist either in the active file system or in some other snapshot.

Process of a snapshot becoming busy

The following table describes a busy snapshot scenario: how a snapshot becomes busy and stays in a busy state through the various stages of LUN creation, snapshot activity, and LUN deletion.

Note

You see the status of the snapshot when you display it in stages 4, 6, 8, 12, and 14 in the following table.

Stage	Busy snapshot scenario
1	Create a LUN. Example: <code>lun create -s 10m /vol/vol2/lunA</code> Result: lunA exists in the active file system.
2	Display LUN status to see lunA. Example: <code>lun show -v</code> Result: The following message is displayed: /vol/vol2/lunA 10m (10485760) (r/w, online) Serial#: BnBt/YiART3H Share: none Space Reservation: enabled Multiprotocol Type: image
3	Take a snapshot of the volume containing lunA. Example: <code>snap create vol2 snap0</code> Result: snap0 of vol2 is created, but it is not busy.

Stage	Busy snapshot scenario
4	Display snapshot status to see that snap0 is not busy. Example: <code>snap list vol2</code> Result: The following message is displayed: <pre> Volume vol2 working ... %/used %/total date name ----- 0% (0%) 0% (0%) Mar 13 02:47 snap0 </pre>
5	Create another LUN, backed by lunA in snap0. Example: <pre> lun create -b /vol/vol2/.snapshot/snap0/lunA /vol/vol2/lunB </pre> Result: lunB is a writable snapshot of lunA. snap0 is now marked “(busy, LUNs).”
6	Display snapshot status to see that snap0 is marked busy. Example: <code>snap list vol2</code> Result: The following message is displayed: <pre> Volume vol2 working ... %/used %/total date name ----- 19% (19%) 0% (0%) Mar 13 02:47 snap0 (busy, LUNs) </pre>
7	Take another snapshot of vol2. Example: <code>snap create vol2 snap1</code> Result: snap1 now exists, but it is not busy. snap0 is still busy.

Stage	Busy snapshot scenario
8	Display snapshot status to see that snap1 is not busy and snap0 is still busy. Example: <code>snap list vol2</code> Result: The following message is displayed: <pre> Volume vol2 working ... %/used %/total date name ----- 0% (0%) 0% (0%) Mar 13 02:50 snap1 19% (19%) 0% (0%) Mar 13 02:47 snap0 (busy, LUNs) </pre>
9	Delete the LUN that was backed by lunA in snap0, lunB. Example: <code>lun destroy /vol/vol2/lunB</code> Result: lunB is deleted, but snap0 is still busy because snap1 contains a copy of lunB. If you attempt to delete snap0, the following message is displayed: <pre> vol2: Snapshot in use by dump, snap restore, SnapMirror, a CIFS share, RAID mirroring, LUNs or volume clone. </pre>
10	Identify the LUNs in existing snapshots that are backed by data in the specified snapshot. Example: <code>lun snap usage vol2 snap0</code> Result: The following message is displayed: <pre> snap1: LUN: /vol/vol2/.snapshot/snap1/lunB Backed By: /vol/vol2/.snapshot/snap0/lunA </pre>
11	Delete snap1. Example: <code>snap delete vol2 snap1</code> Result: snap1 is deleted, and snap0 is no longer marked busy.

Stage	Busy snapshot scenario
12	Display snapshot status to see that snap0 is no longer busy. Example: <code>snap list vol2</code> Result: The following message is displayed: <pre> Volume vol2 working ... %/used %/total date name ----- 19% (19%) 0% (0%) Mar 13 02:47 snap0 </pre>
13	Delete snap0. Example: <code>snap delete vol2 snap0</code> Result: snap0 is deleted.
14	Display snapshot status to see there are no more snapshots. Example: <code>snap list vol2</code> Result: The following message is displayed: <pre> Volume vol2 working ... No snapshots exist. </pre>

Command to use to find snapshots in a busy state

The `lun snap usage` command lists all the LUNs backed by data in the specified snapshot. It also lists the corresponding snapshots in which these LUNs exist. The `lun snap usage` command displays the following information:

- ◆ Writable snapshot LUNs that are holding a lock on the snapshot given as input to this command
- ◆ Snapshots in which these snapshot-backed LUNs exist

Deleting snapshots in a busy state

To delete a snapshot in a busy state, complete the following steps.

Step	Action																								
1	Identify all snapshots that are in a busy state, locked by LUNs, by entering the following command: snap list vol-name Example: snap list vol2 Result: The following message is displayed: Volume vol2 working... <table><thead><tr><th>%/used</th><th>%/total</th><th>date</th><th>name</th></tr></thead><tbody><tr><td>-----</td><td>-----</td><td>-----</td><td>-----</td></tr><tr><td>0% (0%)</td><td>0% (0%)</td><td>Jan 14 04:35</td><td>snap3</td></tr><tr><td>0% (0%)</td><td>0% (0%)</td><td>Jan 14 03:35</td><td>snap2</td></tr><tr><td>42% (42%)</td><td>22% (22%)</td><td>Dec 12 18:38</td><td>snap1</td></tr><tr><td>42% (0%)</td><td>22% (0%)</td><td>Dec 12 03:13</td><td>snap0 (busy, LUNs)</td></tr></tbody></table>	%/used	%/total	date	name	-----	-----	-----	-----	0% (0%)	0% (0%)	Jan 14 04:35	snap3	0% (0%)	0% (0%)	Jan 14 03:35	snap2	42% (42%)	22% (22%)	Dec 12 18:38	snap1	42% (0%)	22% (0%)	Dec 12 03:13	snap0 (busy, LUNs)
%/used	%/total	date	name																						
-----	-----	-----	-----																						
0% (0%)	0% (0%)	Jan 14 04:35	snap3																						
0% (0%)	0% (0%)	Jan 14 03:35	snap2																						
42% (42%)	22% (22%)	Dec 12 18:38	snap1																						
42% (0%)	22% (0%)	Dec 12 03:13	snap0 (busy, LUNs)																						
2	Identify the LUNs and the snapshots that contain them by entering the following command: lun snap usage vol_name snap_name Example: lun snap usage vol2 snap0 Result: The following message is displayed: active: LUN: /vol/vol2/lunC Backed By: /vol/vol2/.snapshot/snap0/lunA snap2: LUN: /vol/vol2/.snapshot/snap2/lunB Backed By: /vol/vol2/.snapshot/snap0/lunA snap1: LUN: /vol/vol1/.snapshot/snap1/lunB Backed By: /vol/vol2/.snapshot/snap0/lunA Note_____ The preceding LUNs are backed by lunA in the snap0 snapshot.																								

Step	Action
3	Delete all the LUNs in the active file system that are displayed by the <code>lun snap usage</code> command by entering the following command: <pre>lun destroy [-f] lun-path [lun-path ...]</pre> Example: <code>lun destroy /vol/vol2/lunC</code>
4	Delete all the snapshots that are displayed by the <code>lun snap usage</code> command in the order they appear, by entering the following command: <pre>snap delete vol-name snapshot-name</pre> Example: <code>snap delete vol2 snap2</code> <code>snap delete vol2 snap1</code> Result: All the snapshots containing lunB are now deleted and snap0 is no longer busy.
5	Delete the snapshot by entering the following command: <pre>snap delete vol-name snapshot-name</pre> Example: <code>snap delete vol2 snap0</code>

Using SnapRestore

What SnapRestore does

SnapRestore enables you use a snapshot to revert a an entity such as a volume, a file, or a LUN quickly to the state it was in when a particular snapshot was taken. In most cases, reverting an entity is much faster than restoring files from tape or copying files from a snapshot to the active file system.

After you select a snapshot for reversion, Data ONTAP reverts the specified entity to the data and timestamps that it contained when the selected snapshot was taken. Data that was written after the selected snapshot was taken is lost. If the entity you select for reversion is a root volume, Data ONTAP reboots the system.

Using SnapRestore

You can use SnapRestore with snapshots of LUNs. It is possible to use SnapRestore to restore a complete WAFL volume or a single LUN.

Caution

Using SnapRestore on a LUN, or on volume that contains LUNs, without stopping all host access to those LUNs, can cause data corruption and system errors.

You need to be aware of the following requirements:

- ◆ You must always unmount the LUN before you run the snaprestore command on a volume containing the LUN or a single file SnapRestore of the LUN. For a single file SnapRestore, you must also take the LUN offline.
- ◆ Check available space; SnapRestore does not revert the snapshot if sufficient space is unavailable.
- ◆ Make sure that the LUN is not the whole root volume (or it reboots the filer).

Restoring a snapshot of a LUN

To use SnapRestore to restore a snapshot of a LUN, complete the following steps.

Step	Action
1	From the host, stop all host access to the LUN.
2	From the host, if the LUN contains a host file system mounted on a host, unmount the LUN on that host.

Step	Action
3	From the filer, unmap the LUN by entering the following command: <pre>lun unmap <i>lun-path initiator-group</i></pre>
4	Enter the following command: <pre>snap restore [-f] [-t <i>vol</i>] <i>volume_name</i> [-s <i>snapshot_name</i>]</pre> -f suppresses the warning message and the prompt for confirmation. This option is useful for scripts. -t <i>vol volume_name</i> specifies the volume name to revert. <i>volume_name</i> is the name of the volume to be reverted. Enter the name only, not the complete path. You can enter only one volume name. -s <i>snapshot_name</i> specifies the name of the snapshot from which to revert the data. You can enter only one snapshot name. Example: <pre>filer> snap restore -s payroll_lun_backup.2 -t /vol/payroll_lun</pre> <pre>filer> WARNING! This will restore a volume from a snapshot into the active filesystem. If the volume already exists in the active filesystem, it will be overwritten with the contents from the snapshot. Are you sure you want to do this? y</pre> <pre>You have selected file /vol/payroll_lun, snapshot payroll_lun_backup.2 Proceed with restore? y</pre> Result: If you did not use the -f option, Data ONTAP displays a warning message and prompts you to confirm your decision to revert the volume.
5	Press y to confirm that you want to revert the volume. Result: Data ONTAP displays the name of the volume and the name of the snapshot for the reversion. If you did not use the -f option, Data ONTAP prompts you to decide whether to proceed with the reversion.

Step	Action	
6	If...	Then...
	You want to continue with the reversion	Press y . Result: The filer reverts the volume from the selected snapshot. If you are reverting the root volume, the filer reboots the system.
	You do not want to proceed with the reversion	Press n or Ctrl-C . Result: The volume is not reverted and you are returned to a filer prompt.
7	If the old maps aren't the ones you want unmap them (if any old maps exist). <code>lun unmap lun-path initiator-group</code>	
8	Remap the LUN. <code>lun map lun-path initiator-group</code>	
9	From the host, remount the LUN, if it was mounted on a host.	
10	From the host, restart access to the LUN.	
11	From the filer, bring the restored LUN online by entering the following command: <code>lun online lun-path</code>	

Note

After you have used SnapRestore to update a LUN from a snapshot, you also need to restart any database applications you closed down and remount the volume from the host side.

Restoring an online LUN from tape

If you try to restore a LUN from a NetApp NDMP/dump tape and the LUN being restored still exists and is exported or online, the restore fails with the following message:

RESTORE: Inode XXX: file creation failed.

Reverting a file

To revert a single file (rather than a volume), complete the following steps.

Note_____

You cannot use SnapRestore for single file reversion on files with NT streams or on directories.

Step	Action
1	Notify network users that you are going to revert a file so that they know that the current data in the file will be replaced by that of the selected snapshot. Note_____ NFS users who try to access a reverted file without first reopening it might get a “stale file handle” error message after the volume reversion. _____

Step	Action
2	Enter the following command: <pre>snap restore[-f] [-t file] [-s snapshot_name] [-r restore_as_path] path_and_LUN_name</pre> -f suppresses the warning message and the prompt for confirmation. -t <i>file</i> specifies that you are entering the name of a file to revert. -s <i>snapshot_name</i> specifies the name of the snapshot from which to revert the data. -r <i>restore_as_path</i> restores the file to a location in the volume different from the location in the snapshot. For example, if you specify /vol/vol0/vol3/myfile as the argument to -r, SnapRestore reverts the file called myfile to the location /vol/vol0/vol3 instead of to the file structure indicated by the path in <i>path_and_lun_name</i>. <i>path_and_LUN_name</i> is the complete path to the name of the file (LUN) to be reverted. You can enter only one path name. A file (LUN) can be restored only to the volume where it was originally. The directory structure to which a file is to be restored must be the same as specified in the path. If this directory structure no longer exists, you must re-create it before restoring the file. Unless you enter -r and a path name, only the file (LUN) at the end of the <i>path_and_lun_name</i> is reverted. Result: If you did not use the -f option, Data ONTAP displays a warning message and prompts you to confirm your decision to revert the file.
3	Press y to confirm that you want to revert the file. Result: Data ONTAP displays the name of the file and the name of the snapshot for the reversion. If you did not use the -f option, Data ONTAP prompts you to decide whether to proceed with the reversion.
4	Press y to continue with the reversion. Result: Data ONTAP reverts the LUN from the selected snapshot.

Example:

```
filer> snap restore -t file -s payroll_backup_friday  
/vol/vol1/payroll_luns
```

```
filer> WARNING! This will restore a file from a snapshot into the  
active filesystem. If the file already exists in the active  
filesystem, it will be overwritten with the contents from the  
snapshot.
```

```
Are you sure you want to do this? y
```

```
You have selected file /vol/vol1/payroll_luns, snapshot  
payroll_backup_friday  
Proceed with restore? y
```

Result: Data ONTAP restores the LUN called payroll_backup_friday to the existing volume and directory structure /vol/vol1/payroll_luns.

After a LUN is reverted with SnapRestore, all user-visible information (data and file attributes) for that LUN in the active file system is identical to that contained in the snapshot.

Using SyncMirror

What SyncMirror does

The SyncMirror product maintains a single mirrored WAFL volume as two physically separated and synchronously up-to-date copies (called plexes) on disks. The copy remains invisible until you split off the mirrored plex for test purposes, to use the space for other purposes, or if a multiple disk failure occurs within a single RAID group, when the filer automatically switches over access to the mirrored plex. SyncMirror operates identically in SAN and NAS environments.

For more information about SyncMirror, see the SyncMirror chapter in the *Data ONTAP 6.4 System Administration Data Protection Guide*.

Using volume copy

Using volume copy You can use the `volume copy` command to copy LUNs in a SAN environment.

The `vol copy start` command enables you to copy data from one WAFL volume to another, either within the same filer or to a different filer. The result of the `vol copy` command is a restricted volume containing the same data as is on the source filer at the time you initiated the copy operation.

For more information about using the `vol copy` command, see the chapter about volume copy in the *Data ONTAP 6.4 Data Protection Guide*.

To copy a WAFL volume containing a LUN to the same or different filer, complete the following step.

Caution—————
You must save contents of host file system buffers to disk before running `vol copy` commands on the filer.

Step	Action
1	Enter the following command: vol copy start -S source:source_volume dest:dest_volume Note————— If the copying takes place between two filers, you can enter the <code>vol copy start</code> command on either the source or destination filer. You cannot, however, enter the command on a third filer that does not contain the source or destination volume.

Example:
`vol copy start -S /vol/vol0 filerB:/vol/vol1`

Note—————
Because the `vol copy` command takes snapshots, remember that in Data ONTAP 6.3 and later, WAFL volumes can be space reserved, which means that if free space within the WAFL volume is below a certain threshold, snapshots cannot be taken.

Using SnapMirror

Using SnapMirror with LUNs

You can use SnapMirror to replicate WAFL volumes or qtrees containing LUNs. A qtree is a special subdirectory of the root directory of a volume.

Snap Mirror destination WAFL volumes and qtrees are read-only, which means that LUNs can be listed using the `lun show` command on the destination filer.

To set up SnapMirror to back up LUNs and then make the LUNs writable so that you can export them, complete the following steps.

Step	Action
1	Make sure that you followed the procedures in the <i>Data ONTAP 6.4 System Administration Data Protection Guide</i> for ◆ Preparing to run SnapMirror ◆ Specifying the destination filer on the source filer ◆ Creating or editing the <code>/etc/snapmirror.conf</code> file ◆ Starting SnapMirror Caution— You must save contents of host file system buffers to disk before running <code>snapmirror</code> commands on the filer.
2	From the destination filer, enter the following command. <pre>snapmirror initialize [options] [dest_filer:] {dest_volume qtree_path}</pre> Example: <code>snapmirror initialize host1:/vol/vol_name/qtree_1 filerA:/vol/vol_name/qtree_1</code> Note— If you need to update the snapshot at a time not specified in the <code>snapmirror.conf</code> file, use the <code>snapmirror update</code> command. See complete information in the SnapMirror chapter of the <i>Data ONTAP 6.4 Data Protection Guide</i>.

Recovering LUNs from the SnapMirror destination

You can see the LUN residing in SnapMirror destination volumes or qtrees using the `lun show` command. You cannot map or export these LUNs to other hosts, however, until you use the `mirror break` command to make the qtree or volume writable.

For qtree-based SnapMirror, you can create a writable LUN outside of the read-only destination qtree instead of having to break the mirror, which enables you to take backups while continuing to mirror the source qtree.

Note

In Data ONTAP 6.3 and later, WAFL volumes can be “space reserved,” which means that if free space within the WAFL volume is below a certain threshold, snapshots cannot be taken and SnapMirror transfers can fail.

To recover data from a LUN that resides in a SnapMirror destination, complete the following steps.

Step	Action
1	Enter the following command to see the LUN backup that you want to recover: lun show Result: Data ONTAP shows all the LUNs on the filer on which the command is run. Example: <pre>backup_server> lun show /vol/vol0/qtree1/lun1 2g (2147483648) /vol/vol0/qtree1/ payroll_backup_lun 1g (1073741824)</pre>
2	Make the WAFL volume or qtree writable so that it can be accessed or exported by entering the following command: snapmirror break {volume_name /vol/volume_name/qtree_name} Example: <pre>snapmirror break /vol/vol_name/qtree_1/payroll_backup_lun</pre>

Step	Action
3	Export the LUN to the administrative host (if that is where you want it) by entering the following command: <pre>lun map /vol/vol_name/qtree_name</pre> Example: <pre>lun map /vol/vol1/qtree_1/payroll_backup_lun server 3</pre>
4	Create a new LUN in the active file system from the existing snapshot so that it is writable by entering the following command: <pre>lun create -b <snapshot lun-path> <lun-path></pre>

Note

SnapDrive is tightly integrated with SnapMirror and creates application-consistent snapshots that can then be replicated by SnapMirror. See the *SnapDrive 2.0 Installation and Administration Guide*.

Using SnapVault

What SnapVault does

SnapVault protects data resident in individual qtrees on SnapVault clients by transferring data from SnapVault client snapshots to a separate SnapVault server.

Using SnapVault with LUNs

To use SnapVault to back up LUNs, complete the following steps.

Caution—————
You must save contents of host file system buffers to disk before running `snapvault` commands on the filer.

Step	Action
1	Make sure that you followed the procedures in the <i>Data ONTAP 6.4 Data Protection Guide</i> for ◆ Preparing to run SnapVault ◆ Specifying the host names of SnapVault clients that are allowed to restore data from the server or creating application-consistent snapshots from the host ◆ Specifying the host name of the server on each SnapVault client, using the <code>snapvault.access</code> option ◆ Using the <code>snapvault snap sched</code> command to configure the SnapVault client snapshot schedule, if not using application-consistent snapshots created by the host ◆ Using the <code>snapvault snap sched -v</code> command to configure the SnapVault server snapshot schedule ◆ Starting SnapVault

Step	Action
2	From the SnapVault server, enter the following command: <pre>snapvault start [-S [source_filer:]src_qtree_path] [dest_filer:]/vol/dest_volume/dest_qtree_with_luns</pre> <i>-S [source_filer:]src_qtree_path</i> specifies the source (SnapVault client) qtree to be backed up. This option must be set the first time the command is run for each SnapVault client qtree you want to copy. <i>dest_filer</i> is the name of the SnapVault server to which the snapshots from the SnapVault client are transferred. If no SnapVault server is specified, the local host's name is used. <i>dest_volume</i> is the name of the volume on the SnapVault server to which the data from the SnapVault client is transferred. <i>dest_qtree_path</i> is the path to and the name of the qtree on the SnapVault server. The qtree on the SnapVault server must have the same name as the qtree on the SnapVault client server. Note If the server qtree name exists before the command runs, the command fails. Example: filerA> snapvault start -S filerB:/vol/vol2/payroll_qtree /vol/vol1/payroll_qtree Result: SnapVault creates a snapshot of payroll_qtree (which contains LUNs) in vol2 of the SnapVault client (filerB), copies the data in it to the SnapVault server (filerA) at /vol/vol1/payroll_qtree, and configures the qtree for future updates.

Note

If you need to update the snapshot at a time not specified in the SnapVault schedule, use the `snapvault update` command. See complete information in the SnapMirror chapter of the *Data ONTAP 6.3 Data Protection Guide*.

Reason to the SnapVault client

You might need to restore data from the SnapVault server to the SnapVault client if the data on the SnapVault client is damaged or corrupted or to migrate data from one filer to another.

After you restore data to the SnapVault client you must do one of the following:

- ◆ Restart the backup process from the restored data (because you want to continue SnapVault protection of the data).
- ◆ Cancel any further backups and release the resources on the SnapVault server that were used in the SnapVault relationship (because you no longer want to back up that data).

Restoring data to the SnapVault client

To restore data from the SnapVault server to the SnapVault client, complete the following step.

Step	Action	
1	Enter the following command from the SnapVault client: <pre>snapvault restore [-s] [-k]] -S [source_filer:]source_qtree_path [dest_filer:]dest_qtree_path</pre> -S [source_filer:]source_qtree_path specifies the name of the SnapVault server from which you want to restore the data. dest_filer is the name of the SnapVault client that you want to restore to. If specified, this name must match the name of the SnapVault client. dest_qtree_path is the name of the SnapVault client qtree that you want to restore to. For more information about SnapVault options, see the <i>Data ONTAP 6.3 Data Protection Guide</i>.	
2	If you...	Then...
	Want to resume SnapVault backups to the newly restored qtree on the SnapVault client	Cancel backup of the newly restored qtree on the SnapVault client Go to Step 3.
	Don't want to resume SnapVault backups to the newly restored qtree on the SnapVault client	Go to Step 6.

Step	Action
3	From the SnapVault server, enter the following command on a single line: <pre>snapvault start -r [-s] [-k] [-S][dest_filer:]dest_qtree_path</pre> -r is required to restart backups from a restored SnapVault client. -S <i>[source_filer:]src_qtree_path</i> sets the new SnapVault client for the qtree. For more information about SnapVault options, see the <i>Data ONTAP 6.4 Data Protection Guide</i>.
4	Export the LUN to the administrative host by entering the following command: <pre>lun map lun-path igroup LUN_ID</pre> Result: You are done. Example: This example shows the SnapVault client (filerB) requesting data to be restored from the SnapVault server (filerA) and then the server restarting the SnapVault backup relationship: <pre>filerB> snapvault restore -S filerA:/vol/sv_vol/payroll_qtree /vol/vol1/payroll_qtree_with_luns</pre> <pre>filerA> snapvault start -r -S filerB:/vol/vol1/payroll_qtree_with_luns</pre> <pre>lun map /vol/vol_name/qtree_1/payroll_backup_lun server 3</pre>
5	Enter the following command on the server: <pre>snapvault release source_qtree_path [dest_filer:]dest_qtree_path</pre> <i>source_qtree_path</i> is the name of the SnapVault client snapshot that you want to release. <i>[dest_filer:]dest_qtree_path</i> is the name of the SnapVault server snapshot that you want to release.

Step	Action
6	Export the LUN to the administrative host. <pre>lun map lun-path igroup LUN_ID</pre> Result: You are done. Example 2: This example shows the SnapVault client (filerB) requesting data to be restored from the SnapVault server (filerA) and then the server canceling the SnapVault backup relationship on both filers to release the resources used: <pre>filerB> snapvault restore -S filerA:/vol/sv_vol/payroll_qtree_with_luns /vol/vol1/payroll_qtree</pre> <pre>filerA> snapvault release filerB:/vol/vol1/payroll_qtree_with_luns filerA:/vol/sv_vol/payroll_qtree_with_luns</pre> <pre>lun map /vol/vol_name/qtree_1/payroll_backup_lun server 1</pre>

Note


In Data ONTAP 6.3 and later, WAFL volumes can be space reserved, which means that if free space within the WAFL volume is below a certain threshold, snapshots cannot be taken and SnapVault might fail.

Backing up data to tape

Structure of SAN backups

In most cases, backup of SAN systems to tape takes place through a separate host, the backup host, to avoid performance degradation on the application host.

From the point of view of the SAN host, LUNs can be confined to a single WAFL volume or qtree or spread across multiple WAFL volumes, qtrees, or filers, as shown in the following diagram, which shows a SAN setup that uses two applications hosts and a clustered pair of filers.



Volumes on the SAN host can consist of a single LUN mapped from the filer or multiple LUNs using a volume manager such as VxVM on Solaris.

Prior to creating a snapshot, it is required to quiesce the application or file system using the LUN (which flushes the host filer system buffers to disk), so that the snapshot is consistent with regard to the application.

Backing up a single LUN to tape

To map a LUN within a snapshot to the backup server for backup, the snapshot must be made writable.

To map a LUN within a snapshot for backup, complete the following steps.

Step	Action
1	Enter the following command to create an initiator group (igroup): <pre>igroup create {-f -i} [-t <ostype>] <group> [<node> ...]</pre> Example: <code>igroup create -f -t solaris payroll_server 20:00:00:00:01:02:03:04</code> Result: Data ONTAP creates an initiator group with the FCP option, the operating system Solaris, the name payroll, and the WWPN 20:00:00:00:01:02:03:04.
2	Enter the following command to create a LUN: <pre>lun create -s <size> [-t <type>] <lun-path></pre> Example: <code>lun create -s 48g -t solaris /vol/vol1/qtree_1/payroll_lun</code> Result: Data ONTAP creates a LUN with a size of 48 GB, of the type Solaris, and with the name and path /vol/vol1/qtree_1/payroll_lun.
3	Enter the following command to map the LUN to the initiator group with a SCSI ID: <pre>lun map <lun-path> <initiator-group> <scsi-lun></pre> Example: <code>lun map /vol/vol1/qtree_1/payroll_lun payroll_server 1</code> Result: Data ONTAP maps the /vol/vol_name/qtree_1/payroll_lun LUN to the payroll_server initiator group with a LUN ID of 1. Note _____ Steps 4, 5, and 6 can be part of your SAN backup application's pre-processing script.
4	Save the contents of host file system buffers to disk using SnapDrive 2.0 for Windows systems or the <code>sync file</code> command for Solaris systems.

Step	Action
5	Create a snapshot by using SnapDrive 2.0 for Windows systems or enter the following command for Solaris systems: snap create volume_name snapshot_name Example: snap create vol1 payroll_backup
6	Enter the following command to create a new LUN in the active file system so that it can be mapped to the backup server from the existing snapshot: lun create -b <snapshot lun-path> <lun-path> Example: lun create -b vol/vol1/.snapshot/payroll_backup/qtrees_1/payroll_lun /vol/vol1/qtrees_1/payroll_backup_lun
7	Enter the following command to map the new LUN you created to the backup host: (<i>backup_server</i> is an initiator group that already exists for the backup host.) lun map <lun-path> <initiator-group> <scsi-lun> Example: lun map /vol/vol_name/qtrees_1/payroll_lun payroll_server 1 Result: Data ONTAP maps the LUN /vol/vol_name/qtrees_1/payroll_lun to the initiator group payroll_server with a SCSI ID of 1.
8	Back up the data in the new LUN from the backup host to tape using your SAN backup application. Note _____ Steps 9 and 10 can be part of your SAN backup application's post-processing script. _____
9	Take the backup LUN offline by entering the following command: lun offline /vol/vol_name/qtrees_name/lun_name Example: lun offline /vol/vol1/qtrees1/payroll_backup_lun

Step	Action
10	Remove the snapshot by entering the following command: snap delete <i>volume_name</i> <i>lun_name</i> Example: snap delete vol1 payroll_backup

Note

In Data ONTAP 6.3, WAFL volumes can be space reserved, which means that if free space within the WAFL volume is below a certain threshold, snapshots cannot be taken.

Keep NAS and SAN data separate

It is strongly recommended that you keep SAN and NAS data separated for backup purposes. Configure volumes as SAN-only or NAS-only and configure qtrees within a single volume as SAN-only or NAS-only.

Using NVFAIL

How the NVFAIL option provides database protection

Enabling the `vol options nvfail` option on a volume provides database file protection by detecting NVRAM failure at boot time.

If this option is enabled and problems are found, it warns database administrators of NVRAM problems by

- ◆ Offlining LUNs in the specified volumes
- ◆ Sending error messages to the console advising you that Data ONTAP has taken LUNs offline or NFS file handles are stale

How you can provide additional protection for databases

You can provide additional protection to specific database LUNs by adding them to an optional file you create called `/etc/nvfail_rename`. When you enable `nvfail` and Data ONTAP detects NVRAM errors at bootup, Data ONTAP renames any database LUNs specified in the `nvfail_rename` file by appending `.nvfail` to the original name. When Data ONTAP renames a database LUN, the database cannot restart automatically. This gives you the opportunity to examine it for inconsistencies before you remove the `.nvfail` extension and make it accessible again.

How NVFAIL works in the SAN environment on the filer

LUNs accessed over CIFS, NFS or FCP can be entered in `/etc/nvfail_rename`. This causes a LUN rename upon boot after an NVRAM failure. This process is described in more detail in the NVRAM chapter of the *Data ONTAP 6.4 System Administration Data Protection Guide*.

Caution

NVRAM failure can lead to possible data inconsistencies. Upon a reboot after an NVRAM failure, all LUNs are taken offline and are no longer exported over FCP. You need to bring the LUNs online (and reexport) each LUN manually to the initiator after data recovery.

Enabling the NVFAIL option

To enable the NVFAIL option on WAFL volumes, complete the following step.

Step	Action
1	Enter the following command: vol options <i>volume-name</i> nvfail on

Creating the nvfail_rename file

To create the nvfail_rename file, complete the following steps.

Step	Action
1	Use an editor to create or modify the nvfail_rename file in the filer's /etc directory.
2	List the full path and file name, one file per line, within the nvfail_rename file. Example: /vol/vol1/home/dbs/oracle-WG73.dbf
3	Save the file.

Using ndmpcopy

When to use the ndmpcopy command

With Data ONTAP 6.4.2, you can use the `ndmpcopy` command to copy a directory, `qtree`, or volume that contains a LUN. However, tape backup and recovery operations of LUNs should generally only be performed on the filer for disaster recovery scenarios, applications with transaction logging, or when combined with other filer-based protection elements, such as SnapMirror and SnapVault. All tape operations local to the filer operate on the entire LUN and cannot interpret the data or file system within the LUN. Thus, you can only recover LUNs to a specific point-in-time unless transaction logs exist to roll forward. When finer granularity is required, use host-based backup and recovery methods.

If you do not specify an existing snapshot when performing a native or NDMP backup operation, the filer creates one before proceeding. This snapshot is deleted when the backup is completed. When a file system contains FCP data, Network Appliance recommends you specify a snapshot that was created at a point in time when the data was consistent by quiescing an application or placing it in “Hot Backup” mode before creating the snapshot. After the snapshot is created, normal application operation can resume and tape backup of the snapshot can occur at any convenient time.

For information on how to use the `ndmpcopy` command, see the *Data ONTAP 6.4 System Administration Data Protection Guide*.

About this chapter

This chapter provides an overview of how to manage filers, as a target; FCP HBAs, FCP traffic; iSCSI service, initiators; and initiator groups in a SAN environment.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Managing the FCP SAN](#)” on page 110
- ◆ “[Managing the iSCSI SAN](#)” on page 125

Managing the FCP SAN

About this section

This section describes how to manage the FCP service and display information about adapters in a Fibre Channel SAN.

For detailed information

For details about the following topics, see:

- ◆ [“Managing the FCP service”](#) on page 111
- ◆ [“Displaying information about FCP adapters”](#) on page 114
- ◆ [“Managing FCP adapters”](#) on page 123

Managing the FCP service

Commands to use

You use the `fc` commands for most of the tasks involved in managing the FCP service and the target and initiator HBAs. For a quick look at all the `fc` commands, enter the `fc help` command at the filer prompt.

You can use FilerView and go to

- ◆ LUNs > FCP to manage FCP adapters and view FCP statistics
- ◆ LUNs > iSCSI to view iSCSI statistics
- ◆ Filer > Manage Licenses to manage the FCP or iSCSI license

Verifying that FCP service is running

If FCP service is not running, target HBAs are automatically taken offline. They cannot be brought online until the FCP service is started.

To verify that the FCP service is running, complete the following step.

Step	Action
1	Enter the following command: fc status Result: A message is displayed indicating whether FCP service is running. Note _____ If the FCP service is not running, verify that the FCP license is enabled and start the FCP service. _____

Verifying that the FCP service is licensed

To verify whether the FCP service is licensed, complete the following step.

Step	Action
1	Enter the following command: license Result: A list of all available services appears, and those services that are enabled show the license code; those that are not enabled are indicated as “not licensed.”

Enabling the FCP service

To enable the FCP service, complete the following step.

Step	Action
1	Enter the following command: license add <i>license_code</i> <i>license_code</i> is the license code you received from NetApp when you purchased the FCP license.

Starting and stopping FCP service

To start and stop the FCP service, complete the following step.

Step	Action
1	Enter the following command: fcp {start stop} Example: fcp start Result: The FCP service begins running. If you enter fcp stop, the FCP stops running.

Taking HBA adapters offline and bringing them online

To take a target HBA adapter offline or bring it online, complete the following step.

Step	Action
1	Enter the following command: <pre>fcp config <i>adapter</i> [up down]</pre> Example: <pre>fcp config 4a down</pre> Result: The target HBA 4a is offline. If you enter <code>fcp config 4a up</code>, the target HBA is brought online.

Disabling the FCP license

To disable the FCP license, complete the following step.

Step	Action
1	Enter the following command: <pre>license delete <i>service</i></pre> <i>service</i> is any service you can license. Example: <pre>license delete fcp</pre>

Displaying information about FCP adapters

How to display HBA adapter information

You can display information and manage HBA adapters in a variety of ways, depending on what you are trying to accomplish. For some commands, you can display additional information by using the `-v` flag, or you can specify which adapter to display information about by entering the adapter number (either the number of the slot it is installed in, or the slot and port number).

If you want to display...	Use this command...
The WWNN (node name) of the target HBA	<code>fcv nodename</code>
Configuration and status information about all adapters (including HBAs, NICs, and switch port information)	<code>sysconfig [-v] [<i>adapter</i>]</code> <i>adapter</i> is a numerical value only, for example, 5. <code>-v</code> displays additional information about all adapters.
Configuration information about disks, disk loops, and options that affect core dumps and takeover	<code>sysconfig -c</code>
Configuration information about target HBAs	<code>fcv config</code>
The node name, port name, and link state of target HBAs	<code>fcv show adapter [-p] [-v][<i>adapter&portnumber</i>]</code> <code>-p</code> displays information about adapters running on behalf of the partner node (filer). <code>-v</code> displays additional information about target adapters. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
Adapter information, including firmware level, PCI bus width and clock speed, node name, cacheline size, FC packet size, link data rate, SRAM parity, and various states	<code>storage show adapter</code>

If you want to display...	Use this command...
The port address, port name, node name, and initiator group name of initiator HBAs connected to target HBAs	<pre>fcv show initiator [-v] [adapter&portnumber]</pre> -v displays the Fibre Channel host address of the initiator. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
Statistics for target HBAs	<pre>fcv stats [-z] [adapter&portnumber]</pre> -z zeros the statistics. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
FCP traffic information	<pre>sysstat -f</pre>
How long FCP has been running	<pre>uptime</pre>
Service statistics	<pre>availtime</pre>

Displaying the HBA's WWNN

To display the WWNN of a target HBA, complete the following step.

Step	Action
1	Enter the following command: <pre>fcv nodename</pre> Result: Fibre Channel nodename: 50:a9:80:00:02:00:8d:b2 (50a9800002008db2)

Displaying information about all adapters

To display information about all adapters installed in the filer, complete the following step.

Step	Action
1	At the filer, enter the following command to see information about all adapters. <pre>sysconfig -v</pre> Result: System configuration information and adapter information for each slot that is used is displayed on the screen. Look for “Fibre Channel Target Host Adapter” to get information about target HBAs. Example: A partial display of information about a target HBA installed in slot 4 appears as follows: <pre>slot 4: Fibre Channel Target Host Adapter 4a (Dual-channel, QLogic 2312 (2342) rev. 2, 64-bit, <ONLINE>) Firmware rev: 3.1.15 Host Port Addr: 060200 Cacheline size: 8 SRAM parity: Yes Link Data Rate: 1 Gbit FC Nodename: 50:a9:80:00:02:00:b6:18 (50a980000200b618) FC Portname: 50:a9:80:01:02:00:b6:18 (50a980010200b618) Connection: PTP, Fabric slot 4: Fibre Channel Target Host Adapter 4b (Dual-channel, QLogic 2312 (2342) rev. 2, 64-bit, <ONLINE>) Firmware rev: 3.1.15 Host Port Addr: 060300 Cacheline size: 8 SRAM parity: Yes Link Data Rate: 1 Gbit FC Nodename: 20:01:00:e0:8b:26:39:b7 (200100e08b2639b7) FC Portname: 20:01:00:e0:8b:26:39:b7 (200100e08b2639b7) Connection: PTP, Fabric, Standby</pre> Note— In the display, the information about the Dual-channel QLogic HBA, the value 2312, does not specify the model number of the HBA; it refers to the device ID set by QLogic.

Step	Action
	If the adapter is attached to a switch, switch port information is displayed immediately following SRAM parity, as follows: <pre>Switch Port: brcd:1 brcd:5.6 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.5 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.4 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.3 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.2 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.1 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.0 : NETAPP X225_ST336704FC NA08 34.5GB (71687368 512B/sect) brcd:5.shelf0 : VEM Kernel Version : 2.5 App. Version : 3.0</pre>

Displaying brief target HBA information

To display configuration information about target HBAs, and to quickly detect whether they are active and online, complete the following step.

Step	Action
1	At the filer, enter the following command. <pre>fcp config</pre> Result: <pre>4a: ONLINE <ADAPTER UP> PTP Fabric host address 060200 portname 50:a9:80:01:02:00:b6:18 nodename 50:a9:80:00:02:00:b6:18 mediatype ptp partner adapter None 4b: ONLINE <ADAPTER UP> PTP Fabric Standby host address 060300 portname 20:01:00:e0:8b:26:39:b7 nodename 20:01:00:e0:8b:26:39:b7 mediatype auto partner adapter None</pre>

Displaying detailed target HBA information

To display the nodename, portname, and link state of all target HBAs, complete the following step. Notice that the portname and nodename are displayed with the separating colons, and without. For Solaris hosts, you use the WWPN without separating colons when you map adapter portnames (or these target WWPNs) to the host.

Step	Action
1	At the filer, enter the following command: fcp show adapter Result: The following information is displayed for all target HBAs. This example displays information for the HBA in slot 4: slot: 4a Description Fibre Channel Target Adapter 4a (Dual-channel, QLogic 2312 rev.2 Status: ONLINE FC Nodename: 50:a9:80:00:02:c7:15 (50a9800002c715) FC Portname: 50:a9:80:01:02:c7:15 (50a9800102c715) Standby No slot: 4b Description Fibre Channel Target Adapter 4a (Dual-channel, Qlogic 2312 rev.2 Status: LINK NOT CONNECTED FC Nodename: 20:01:00:e0:8b:27:b2:a8 (200100e08b27b2a8) FC Portname: 20:01:00:e0:8b:27:b2:a8 (200100e08b27b2a8) Standby YES

Displaying initiator HBA information To display the portname and group name (of the initiator group) of initiator HBAs (host) connected to target HBAs, complete the following step.

Step	Action																				
1	At the filer, enter the following command: <pre>fcp show initiator</pre> Result : The following output is displayed: Initiators connected on adapter 4a: <table> <tr> <th>Portname</th><th>Group</th></tr> <tr> <td>10:00:00:00:c9:2b:fc:bf</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:fe:fe</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:cc:ca</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:fd:2d</td><td></td></tr> </table> Adapter 4b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode. Initiators connected on adapter 5a: <table> <tr> <th>Portname</th><th>Group</th></tr> <tr> <td>10:00:00:00:c9:2b:8e:29</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:83:41</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:8d:fc</td><td></td></tr> <tr> <td>10:00:00:00:c9:2b:9f:15</td><td></td></tr> </table> Adapter 5b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode.	Portname	Group	10:00:00:00:c9:2b:fc:bf		10:00:00:00:c9:2b:fe:fe		10:00:00:00:c9:2b:cc:ca		10:00:00:00:c9:2b:fd:2d		Portname	Group	10:00:00:00:c9:2b:8e:29		10:00:00:00:c9:2b:83:41		10:00:00:00:c9:2b:8d:fc		10:00:00:00:c9:2b:9f:15	
Portname	Group																				
10:00:00:00:c9:2b:fc:bf																					
10:00:00:00:c9:2b:fe:fe																					
10:00:00:00:c9:2b:cc:ca																					
10:00:00:00:c9:2b:fd:2d																					
Portname	Group																				
10:00:00:00:c9:2b:8e:29																					
10:00:00:00:c9:2b:83:41																					
10:00:00:00:c9:2b:8d:fc																					
10:00:00:00:c9:2b:9f:15																					

Displaying statistics

To display information about the activity on target HBAs, complete the following step.

Step	Action																																												
1	Enter the following command: <pre>fcps stats [adapter]</pre> <i>adapter</i> is the slot and port number where the target HBA is installed. Example: Statistics for adapter 4a <table><tr><td>Adapter Resets:</td><td>0</td></tr><tr><td>Frame Overruns:</td><td>0</td></tr><tr><td>Frame Underruns:</td><td>0</td></tr><tr><td>Initiators Connected:</td><td>4</td></tr><tr><td>Link Breaks:</td><td>0</td></tr><tr><td>LIP Resets:</td><td>0</td></tr><tr><td>Scsi Requests Dropped:</td><td>0</td></tr><tr><td>Spurious Interrupts:</td><td>65</td></tr><tr><td>Total Logins:</td><td>40</td></tr><tr><td>Total Logouts:</td><td>36</td></tr><tr><td>CRC Errors:</td><td>0</td></tr></table> Adapter 4b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode. Statistics for adapter 5a <table><tr><td>Adapter Resets:</td><td>0</td></tr><tr><td>Frame Overruns:</td><td>0</td></tr><tr><td>Frame Underruns:</td><td>0</td></tr><tr><td>Initiators Connected:</td><td>0</td></tr><tr><td>Link Breaks:</td><td>0</td></tr><tr><td>LIP Resets:</td><td>0</td></tr><tr><td>Scsi Requests Dropped:</td><td>0</td></tr><tr><td>Spurious Interrupts:</td><td>0</td></tr><tr><td>Total Logins:</td><td>0</td></tr><tr><td>Total Logouts:</td><td>0</td></tr><tr><td>CRC Errors:</td><td>0</td></tr></table> Adapter 5b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode.	Adapter Resets:	0	Frame Overruns:	0	Frame Underruns:	0	Initiators Connected:	4	Link Breaks:	0	LIP Resets:	0	Scsi Requests Dropped:	0	Spurious Interrupts:	65	Total Logins:	40	Total Logouts:	36	CRC Errors:	0	Adapter Resets:	0	Frame Overruns:	0	Frame Underruns:	0	Initiators Connected:	0	Link Breaks:	0	LIP Resets:	0	Scsi Requests Dropped:	0	Spurious Interrupts:	0	Total Logins:	0	Total Logouts:	0	CRC Errors:	0
Adapter Resets:	0																																												
Frame Overruns:	0																																												
Frame Underruns:	0																																												
Initiators Connected:	4																																												
Link Breaks:	0																																												
LIP Resets:	0																																												
Scsi Requests Dropped:	0																																												
Spurious Interrupts:	65																																												
Total Logins:	40																																												
Total Logouts:	36																																												
CRC Errors:	0																																												
Adapter Resets:	0																																												
Frame Overruns:	0																																												
Frame Underruns:	0																																												
Initiators Connected:	0																																												
Link Breaks:	0																																												
LIP Resets:	0																																												
Scsi Requests Dropped:	0																																												
Spurious Interrupts:	0																																												
Total Logins:	0																																												
Total Logouts:	0																																												
CRC Errors:	0																																												

Displaying FCP traffic information

To display FCP traffic information (FCP ops/s, KB/s), complete the following step.

Step	Action										
1	Enter the following command:										
	sysstat -f										
	Result: The following output is displayed:										
	CPU	NFS	CIFS	FCP	Net	kB/s	diskB/s	FCP	kB/s	Cache	
					in	out	read	write	in	out	age
	1%	0	0	0	2	0	8	20	0	0	>60
1%	0	0	0	1	0	19	40	0	0	>60	
1%	0	0	0	2	0	8	27	0	0	>60	

Displaying how long FCP has been running

To display information about how long FCP has been running, complete the following step.

Step	Action
1	Enter the following command: <pre>uptime</pre> Result: The following output is displayed: <pre>12:46am up 2 days, 8:59 102 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 DAFS ops, 2 FCP ops, 0 iSCSI ops</pre>

Displaying FCP service statistics

To display FCP service statistics, complete the following step.

Step	Action
1	Enter the following command: <pre>availtime</pre> Result: The following output is displayed: <pre>Service statistics as of Mon Jul 1 00:28:37 GMT 2002 System (UP). First recorded (3894833) on Thu May 16 22:34:44 GMT 2002 P 28, 230257, 170104, Mon Jun 10 08:31:39 GMT 2002 U 24, 131888, 121180, Fri Jun 7 17:39:36 GMT 2002 NFS (UP). First recorded (3894828) on Thu May 16 22:34:49 GMT 2002 P 40, 231054, 170169, Mon June 10 08:32:44 GMT 2002 U 36, 130363, 121261, Fri Jun 7 17:40:57 GMT 2002 FCP P 19, 1417091, 1222127, Tue Jun 4 14:48:59 GMT 2002 U 6, 139051, 121246, Fri Jun 7 17:40:42 GMT 2002</pre>

Managing the FCP SAN

Managing FCP adapters

How to configure media types

Use the following table to decide how to configure the media type of an HBA.

When you...	Use this media type
Want Data ONTAP to detect what media type to use for the either node, the initiator (host) or the target (filer).	autosense
Are connecting the host directly to the filer. Note _____ NetApp recommends that you use ptp whenever the host supports it.	ptp
Are connect the host and filer to a Fibre Channel switch.	loop

Changing the filer's WWNN

The WWNN of a filer is generated by a serial number in its NVRAM, but it is stored on disk. If you ever replace a filer chassis and reuse it in the same NetApp SAN, it is possible, although extremely rare, that the WWNN of the replaced filer is duplicated. In this unlikely event, you can change the WWNN of the filer by completing the following step.

Step	Action
1	Enter the following command: fcp nodename <i>nodename</i> <i>nodename</i> is a 64-bit WWNN address Example: fcp nodename 50:a9:80:00:02:00:8d:ff

Specifying a partner adapter

To specify the name of the adapter in the partner filer that the adapter in the local filer should take over in the event the partner filer fails, complete the following step.

Step	Action
1	Enter the following command: <code>fcp config partner <i>adapter</i></code>

Managing the iSCSI SAN

About this section

This section describes how to manage the iSCSI service and the filer as a target in the iSCSI SAN environment.

Topics in this section

This section discusses the following topics

- ◆ [“Understanding the iSCSI service on the filer”](#) on page 126
- ◆ [“Registering the filer with an iSNS server”](#) on page 127
- ◆ [“Managing the iSCSI service”](#) on page 130
- ◆ [“Managing security for iSCSI initiators”](#) on page 133
- ◆ [“Displaying statistics for iSCSI sessions”](#) on page 137
- ◆ [“Managing the iSCSI service on filer interfaces”](#) on page 140

Understanding the iSCSI service on the filer

How the filer implements iSCSI

The filer does not use HBAs to implement the iSCSI service. An iSCSI software adapter driver implements the iSCSI service. Communication between an initiator and the filer occurs over an iSCSI session. The initiator and the filer conduct a session over one or more TCP/IP connections, and iSCSI messages are sent over the filer's standard Ethernet interfaces.

For a quick look at all of the `iscsi` commands, enter the `iscsi help` command at the filer prompt.

What portal groups are

A portal group is a set of network portals within an iSCSI node over which an iSCSI session is conducted. In a target, a network portal is identified by its IP address and listening TCP port. The iSCSI protocol allows only one session between a specific iSCSI initiator port and a single portal group on the target.

The iSCSI software adapter maps each Ethernet interface on the filer to a single portal group. In this way, an initiator port can conduct simultaneous iSCSI sessions on the filer, with one session for each portal group.

Registering the filer with an iSNS server

What an iSNS server does

The Internet Storage Name Service (iSNS) is a protocol that enables automated discovery and management of iSCSI devices on an IP storage network. An iSNS server maintains information about active iSCSI devices on the network, including their IP addresses, iSCSI node names, and portal groups. An iSCSI initiator can query the iSNS server to discover iSCSI target devices.

How the filer interacts with an iSNS server

The filer automatically registers its IP address, node name, and portal groups with the iSNS server when the iSCSI service is started and iSNS is enabled. You use the `iscsi isns update` command to notify the iSNS server when the filer's iSCSI configuration information changes.

Command to register the filer

You use the `iscsi isns` command to configure the filer to register with an iSNS server. This command specifies the information the filer sends to the iSNS server.

Note

You cannot configure or manage the iSNS server from the filer console. To manage the iSNS server, use the iSNS server administration tools.

Configuring the filer to register with an iSNS server

To configure the filer to register with the iSNS server, complete the following steps.

Step	Action
1	On the filer console, enter the following command to identify the iSNS server that the filer registers with: <pre>iscsi isns config -e <i>entity_id_string</i> -i <i>ip_addr</i></pre> <i>entity_id_string</i> is the filer's network entity on the iSNS server. A network entity consists of storage node objects and portal objects. A storage node object is the filer's iSCSI node name. A portal object is the IP address and port through which the filer is accessed. <i>ip_addr</i> is the IP address of the iSNS server.
2	Enter the following command: <pre>iscsi isns start</pre> Result: The iSNS service is started and the filer registers with the iSNS server. Note_____ iSNS registration is persistent across reboots if the iSCSI service is running and iSNS is started. _____

Reregistering the filer with the iSNS server

You must reregister the filer with the iSNS server when you change the filer's iSCSI configuration information, such as its node name or portal information.

To reregister the filer, complete the following step.

Step	Action
1	On the filer console, enter the following command: <pre>iscsi isns update</pre>

Disabling iSNS

When you stop the iSNS service, the filer stops registering its iSCSI information with the iSNS server.

To stop the iSNS service, complete the following step.

Step	Action
1	Enter the following command: iscsi isns stop

Managing the iSCSI service

Verifying that the iSCSI service is running

To verify that the iSCSI service is running, complete the following step.

Step	Action
1	Enter the following command: iscsi status Result: A message is displayed indicating whether iSCSI service is running. Note_____ If the iSCSI service is not running, verify that the iSCSI license is enabled and start the service. _____

Verifying that iSCSI is licensed

To verify that the iSCSI service is licensed, complete the following step.

Step	Action
1	Enter the following command: license Result: A list of all available licenses appears. An enabled license shows the license code.

Enabling the iSCSI license

To enable the iSCSI license, complete the following step.

Step	Action
1	Enter the following command: license add <i>License_code</i> <i>license_code</i> is the license code you received from NetApp when you purchased the iSCSI license.

Starting and stopping the iSCSI service

To start the iSCSI service, complete the following step.

Step	Action
1	Enter the following command: iscsi {start stop}

Displaying the filer's node name

To display the filer's node name, complete the following step.

Step	Action
1	Enter the following command: iscsi nodename Example: iscsi nodename iSCSI target nodename: iqn.1992-08.com.netapp:sn.123456789

Changing the filer's node name

To change the filer's node name, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi nodename iqn.1992-08.com.netapp:unique_device_name</pre> <i>unique_device_name</i> is the unique name for the filer. Example: <pre>iscsi nodename iqn.1992-08.com.netapp:filer_hq</pre>

Displaying initiators connected to the filer

To display a list of initiators connected to the filer, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi show initiator</pre> Result: The filer console displays the initiators currently connected to the filer and the portal groups to which they are connected. Example: <pre>iscsi show initiator Initiators connected on adapter iswta: Tgt_PG iSCSI Initiator Name/ISID 2 iqn.1999-07.com.mycorp:svr1 / 40:01:01:dc:00:05 2 iqn.1999-07.com.mycorp:svr2 / 40:01:01:dc:00:08</pre>

Managing security for iSCSI initiators

How iSCSI authentication works

During the initial stage of an iSCSI session, the initiator sends a login request to the filer to begin an iSCSI session. The filer permits or denies the login request according to one of the following authentication methods:

- ◆ Challenge Handshake Authentication Protocol (CHAP)—The initiator logs in using a CHAP user name and password. You can specify a CHAP password or generate a random password.

There are two types of CHAP user names and passwords:

- ❖ Inbound—The filer authenticates the initiator. Inbound settings are required if you are using CHAP authentication.
- ❖ Outbound—This is an optional setting to enable the initiator to authenticate the filer. You can use outbound settings only if you defined an inbound user name and password on the filer.
- ◆ deny—The initiator is denied access to the filer.
- ◆ none—The filer does not require authentication for the initiator.

You can define a list of initiators and their authentication methods. You can also define a default authentication method for initiators that are not on this list. If you do not specify a list of initiators and authentication methods, the default method is none—any initiator can access the filer without authentication.

Guidelines for using CHAP authentication

The following guidelines apply to CHAP authentication:

- ◆ If you define an inbound user name and password on the filer, you must use the same user name and password for outbound CHAP settings on the initiator.
- ◆ If you also define an outbound user name and password on the filer to enable bidirectional authentication, you must use the same user name and password for inbound CHAP settings on the initiator.
- ◆ You cannot use the same user name and password for inbound and outbound settings on the filer.

Displaying initiator authentication methods

To view a list of initiators and their authentication methods, complete the following step.

Step	Action
1	Enter the following command: iscsi security show

Defining an authentication method for an initiator

To define an authentication method for an initiator, complete the following steps.

Step	Action	
1	If you want to...	Then...
	Use CHAP authentication and generate a random password	1. Enter the following command: iscsi security generate Result: The filer generates a 128-bit random password. 2. Proceed to Step 2 .
	◆ Use CHAP authentication and specify a password ◆ Use another security method	Proceed to Step 2 .

Step	Action
2	For each initiator, enter the following command: <pre>iscsi security add -i initiator -s method -p inpassword -n inname [-o outpassword -m outname]</pre> <i>initiator</i> is the initiator name in the iSCSI nodename format. <i>method</i> is one of the following: ◆ <i>CHAP</i>—Authenticate using a CHAP user name and password. ◆ <i>none</i>—The initiator can access the filer without authentication. ◆ <i>deny</i>—The initiator cannot access the filer. <i>inpassword</i> is the inbound password for CHAP authentication. The filer uses the inbound password to authenticate the initiator. <i>inname</i> is a user name for inbound CHAP authentication. The filer uses the inbound user name to authenticate the initiator. <i>outpassword</i> is a password for outbound CHAP authentication. The filer uses this password for authentication by the initiator. <i>outname</i> is a user name for outbound CHAP authentication. The filer uses this user name for authentication by the initiator. Note— If you generated a random password in Step 1, you can use this string for either <i>inpassword</i> or <i>outpassword</i>. If you enter a string, the filer interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

Defining a default authentication method

To define a default authentication method for initiators that are not on the authentication list, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi security default -s <i>method</i> -p <i>inpassword</i> -n <i>inname</i> [-o <i>outpassword</i> -m <i>outname</i>]</pre> <i>method</i> is one of the following: ◆ <i>CHAP</i>—Authenticate using a CHAP user name and password. ◆ <i>none</i>—Initiators that are not on the list do not require authentication to access the filer. ◆ <i>deny</i>—Initiators that are not on the list are denied access to the filer. <i>inpassword</i> is the inbound password for CHAP authentication. The filer uses the inbound password to authenticate the initiator. <i>inname</i> is a user name for inbound CHAP authentication. The filer uses the inbound user name to authenticate the initiator. <i>outpassword</i> is a password for outbound CHAP authentication. The filer uses this password for authentication by the initiator. <i>outname</i> is a user name for outbound CHAP authentication. The filer uses this user name for authentication by the initiator.

Removing an initiator from the authentication list

To remove an initiator from the authentication list, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi security delete -i <i>initiator</i></pre> -i <i>initiator</i> is the initiator name in the iSCSI node name format. Result: The initiator is removed from the authentication list and logs in to the filer using the default authentication method.

Managing the iSCSI SAN

Displaying statistics for iSCSI sessions

Displaying iSCSI statistics

To display iSCSI statistics, complete the following step.

Step	Action
1	Enter the following command: iscsi stats Result: The following information is displayed: Statistics for adapter iswta iSCSI PDUs Received SCSI-Cmd: 0 Nop-Out: 0 SCSI TaskMgtCmd: 0 LoginReq: 0 LogoutReq: 0 Text Req: 0 DataOut: 0 SNACK: 0 Unknown: 0 Total: 0 iSCSI PDUs Transmitted SCSI-Rsp: 0 Nop-In: 0 SCSI TaskMgtRsp: 0 LoginRsp: 0 LogoutRsp: 0 TextRsp: 0 Data_In: 0 R2T: 0 Asyncmsg: 0 Reject: 0 Total: 0 iSCSI CDBs DataIn Blocks: 0 DataOut Blocks: 0 Error Status: 0 Success Status: 0 Total CDBs: 0 iSCSI ERRORS Failed Logins: 0 Failed TaskMgt: 0 Failed Logouts: 0 Failed TextCmd: 0 Protocol: 0 Digest: 0 Unexpected session disconnects: 0 PDU discards (outside CmdSN window): 0 PDU discards (invalid header): 0

Interpreting iSCSI statistics

The `iscsi stats` command displays the following statistics:

iSCSI PDUs Received: This section lists the iSCSI Protocol Data Units (PDUs) sent by the initiator. It includes the following statistics:

- ◆ SCSI-CMD—SCSI-level command descriptor blocks.
- ◆ LoginReq—Login request PDUs sent by initiators during session setup.
- ◆ DataOut—PDUs containing write operation data that did not fit within the PDU of the SCSI command. The PDU maximum size is set by the filer during the operation negotiation phase of the iSCSI login sequence.
- ◆ Nop-Out—A message sent by initiators to check whether the target is still responding.
- ◆ Logout-Req—A request sent by initiators to terminate active iSCSI sessions.
- ◆ SNACK—Serial Number Acknowledge requests sent by initiators to request retransmission of PDUs previously sent by the target.

Note

This release of Data ONTAP does not support SNACK messages.

- ◆ SCSI TaskMgtCmd—SCSI-level task management messages, such as ABORT_TASK and RESET_LUN.
- ◆ Tex-Req—Text request PDUs that initiators send to request target information and renegotiate session parameters.

iSCSI PDUs transmitted: This section lists the iSCSI PDUs sent by the filer and includes the following statistics:

- ◆ SCSI-Rsp—SCSI response messages.
- ◆ LoginRsp—Responses to login requests during session setup.
- ◆ DataIn—Messages containing data requested by SCSI read operations.
- ◆ Nop-In—Responses to initiator Nop-Out messages.
- ◆ Logout-Rsp—Responses to Logout-Req messages.
- ◆ R2T—Read to transfer messages indicating that the target is ready to receive data during a SCSI write operation.
- ◆ SCSI TaskMgtRsp—Responses to task management requests.
- ◆ TexRsp—Responses to Tex-Req messages.
- ◆ Asyncmsg—Messages the target sends to asynchronously notify the initiator of an event, such as the termination of a session.
- ◆ Reject—Messages the target sends to report an error condition to the initiator, for example:
 - ❖ Data Digest Error (checksum failed)

- ❖ Target does not support command sent by the initiator
- ❖ Initiator sent a command PDU with an invalid PDU field

iSCSI CDBs: This section lists statistics associated with the handling of iSCSI Command Descriptor Blocks, including the number of blocks of data transferred, and the number of SCSI-level errors and successful completions.

iSCSI Errors: This section lists login failures, session disconnects, and other SCSI protocol errors.

Managing the iSCSI service on filer interfaces

Command to use You use the `iswt` command to manage the iSCSI service on the filer’s Ethernet interfaces. The `iswt` command enables you to control the network interfaces used for iSCSI communication. For example, you can enable iSCSI communication over specific Gigabit Ethernet (GbE) interfaces.

Displaying iSCSI status on filer interfaces To display the status of the iSCSI service on filer interfaces, complete the following step.

Step	Action
1	Enter the following command: <pre>iswt interface show [-a <i>interface</i>]</pre> <code>-a</code> specifies all interfaces. <code><i>interface</i></code> is a specific Ethernet interface. Example: The following example shows the iSCSI service enabled on two filer Ethernet interfaces: <pre>iswt interface -a Interface e0 enabled Interface e5 enabled</pre>

**Enabling iSCSI on a
filer interface**

To enable the iSCSI service on an interface, complete the following step.

Step	Action
1	Enter the following command: <pre>iswt interface enable {-a <i>interface</i> ... }</pre> -a specifies all interfaces. <i>interface</i> is a specific Ethernet interface. Example: The following example enables the iSCSI service on interface e5: <pre>iswt interface enable e5</pre>

**Disabling iSCSI on
a filer interface**

To disable the iSCSI service on an interface, complete the following step.

Step	Action
1	Enter the following command: <pre>iswt interface disable [-f] {-a <i>interface</i> ... }</pre> -f forces the termination of any outstanding iSCSI sessions without prompting you for confirmation. If you do not use this option, the <code>iswt</code> command displays a message notifying you that active sessions are in progress on the interface and requests confirmation before terminating these sessions and disabling the interface. -a specifies all interfaces. <i>interface</i> is a specific Ethernet interface.

**Enabling iSCSI on
vifs**

To enable the iSCSI service for a vif, you must ensure that all the interfaces in the vif are enabled for the iSCSI service. If the iSCSI service is disabled on any of the interfaces in a vif, then you cannot transmit iSCSI traffic on that vif.

Index

Symbols

/etc/nvfail_rename file, for database protection 105
/etc/syslog.conf, space reservations in 50

A

adapters
 bringing online 113
 displaying information about 114, 116
 managing 114
 partner filer, specifying 124
 taking offline 113
applications, backup 74

B

backup applications 74
backup to tape, of LUNs 101
backups, keeping NAS and SAN separate in 104

C

clusters, prerequisites for SAN 12
configurations, supported in NetApp SAN environments 3

D

data protection applications 74
database file protection, renaming files 105
database protection 105
direct-attached connection 3

F

Fabric Name Service, registering WWPNs to 40
fabric-attached
 multipath connection 3
 single-path connection 3
FCP
 commands 111
 disabling 113
 enabling 112

 license 112
FCP commands
 fcpl show initiator 40
 fcpl start 112
 fcpl stop 112
FCP service
 how long running 121
 statistics about 122
 status of 111
 traffic information 121
FCP traffic, displaying information about 122
filers (targets) role in NetApp SAN environments 3
FilerView
 creating initiator groups with 38
 mapping LUNs to initiator groups with 38
files
 nvfail_rename 106
 protecting with nvfail option 105

G

guidelines for creating LUNs 28

H

HBAs. *See* host bus adapters
host bus adapters
 displaying information about 121, 122
 finding world wide node name (WWNN) 115
 initiator, displaying information about 119
 media type, configuring 123
 target, displaying information about 117
hosts (initiators), role in NetApp SAN environments 3

I

igroup commands
 igroup add 71
 igroup create 41, 70
 igroup destroy 71
 igroup help 72
 igroup remove 72

- igroup set 72
- igroup show 72
- igroups. *See* initiator groups
- initializing SnapMirror 93
- initiator groups
 - adding 71
 - creating 70
 - creating with FilerView 38
 - defined 10
 - destroying 71
 - displaying command line help 72
 - displaying contents of 72
 - information needed to create 68
 - mapping LUNs to, with FilerView 38
 - mapping tables of 46
 - naming 10
 - removing 72
 - setting the operating system type 72
 - unmapping LUNs from 47
- initiator host bus adapters, displaying information about 119
- initiators (hosts), role in NetApp SAN environments 3
- iSCSI node name 8

L

Logical Unit Numbers. *See* LUNs

lputilnt command 11

lun commands

- lun create 42
- lun destroy 48
- lun map 43
- lun offline 48
- lun online 47
- lun resize 52
- lun serial 53
- lun setup 29
- lun show 43, 56
- lun stats 55
- lun unmap 47

LUN IDs, defined 10

LUN masking, defined 11

LUNs

- access controlled by mapping tables 46

- accessing with NAS protocols 48
- adding comments to 51
- backing up to tape 102
- backing up with SnapMirror 93
- backing up with SnapVault 96
- backup applications for 74
- bringing online 47
- changing comments to 51
- changing serial number 53
- creating
 - guidelines for 28
 - methods of 27
- creating volumes for 24
- creating with individual commands 40
- creating with lun setup command 29
- defined 9
- displaying information about 54, 56
- displaying reads, writes, and operations for 55
- information required for creating 26
- mapping to initiator groups using FilerView 38
- masking 11
- modifying attributes 50
- multiple initiator groups for 68
- naming 26
- operating statistics 55
- qtrees for 25
- recovering from SnapMirror destinations 94
- removing 48
- resizing 52
- restoring snapshots of with SnapRestore 85
- size information 26
- snapshots of 75
- taking offline 48
- unmapping from initiator group 47
- using Single File SnapRestore with 88
- using with SyncMirror 91
- using with the vol copy command 92
- where to create 9

M

- mapping tables, for accessing initiator groups 46
- media type, of host bus adapter, configuring 123
- minimum and maximum LUN size
 - Windows 27

N

- network interface card (NIC), for TCP/IP connection 12
- nodenames, of initiator host bus adapters, displaying 119
- nvfail option, of vol options command
 - using with LUNs 105
- nvfail_rename file 106

P

- port names of initiator adapters, displaying 119

Q

- qtrees
 - creating 15
 - creating for LUNs 25

S

- serial numbers, changing LUN 53
- Single File SnapRestore, using with LUNs 88
- snap restore -t vol command 86
- SnapMirror
 - recovering LUNs from destinations 94
 - using with LUNs 93
- snapmirror initialize command 93
- SnapRestore
 - how it works 85
 - using with LUNs 85
- snapshots
 - creating with SnapDrive 2.0 75
 - reason to use 75
 - using with LUNs 75
 - using with SnapRestore 85
- SnapVault
 - initializing backup relationships with 97
 - releasing snapshots with 99
 - restoring data from server to client 98
 - restoring LUNs to the client 98
 - using with LUNs 96

- snapvault release command 99
- snapvault restore command 98
- snapvault start command 97
- space reservations
 - enabling 19
 - requirements for in LUN volumes 94
- SyncMirror, using with LUNs 91
- sysconfig -v command 116

T

- tape backup, of LUNs 101
- target host bus adapters, displaying information about 117
- targets (filers), role of in NetApp SAN environments 3
- TCP/IP network, recommendation for 4

V

- vol copy command, copying LUNs with 92
- vol option nvfail, using with LUNs 105
- volumes
 - creating 15
 - creating for LUNs 24
 - reverting from snapshots with SnapRestore 86

W

- Windows, minimum and maximum LUN size 27
- WNNS, defined 4
- WWNNs
 - displaying 5
 - how assigned 7
 - how to change 123
 - how used 7
- WWPNs
 - defined 4
 - displaying 5
 - how assigned 7
 - how used 7

