

Data ONTAP™ 6.5

System Administration

Block Access Management Guide for FCP

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 210-00477
December 2003

Copyright and trademark information

Copyright information

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

Portions © 1998–2001 The OpenSSL Project. All rights reserved.

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.

Portions of the software were created by Netscape Communications Corp.

The contents of those portions are subject to the Netscape Public License Version 1.0 (the “License”); you may not use those portions except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>.

Software distributed under the License is distributed on an “AS IS” basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corp. Portions created by Netscape are Copyright © 1998 Netscape Communications Corp. All rights reserved.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from <http://www.xmlsoft.org/>.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.
2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [Sdate-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de

Recherche en Informatique et en Automatique, Keio University). All Rights Reserved.
<http://www.w3.org/Consortium/Legal/>"

3. Notice of any changes or modifications to the W3C files, including the date changes were made.
(We recommend you provide URLs to the location from which the code is derived.)

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

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 and other countries. 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, NearStore, 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, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	How NetApp Implements an FC Network.	1
	Understanding NetApp storage appliances	2
	Understanding how NetApp implements an FC SAN network	5
	Related documents	10
Chapter 2	Creating LUNs, igroups, and LUN maps	11
	Configuring disk space	12
	Guidelines for creating volumes	15
	Guidelines for calculating the size of a volume	21
	Methods of creating LUNs, igroups, and LUN maps	26
	Creating LUNs with the lun setup program	32
	Creating LUNs and igroups with FilerView	37
	Creating LUNs and igroups using individual commands	40
Chapter 3	Managing LUNs	45
	Managing LUNs and LUN maps	46
	Displaying LUN information	52
	Monitoring disk space	56
Chapter 4	Managing Initiator Groups	71
	Managing igroups.	72
	Managing initiator requests.	75
Chapter 5	Using Data Protection with FCP	81
	Data ONTAP protection methods	82
	Using snapshots.	84
	Using LUN clones	86

Deleting busy snapshots 88

Using SnapRestore 94

Backing up data to tape 100

Using NDMP 104

Using volume copy 105

Using NVFAIL 106

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

Managing the FCP service 110

Displaying information about HBAs 113

Managing FCP adapters 122

Glossary 125

Index 129

Preface

About this guide

This guide describes how to use a NetApp® storage appliance as a Fibre Channel Protocol (FCP) target in a SCSI storage network. Specifically, this guide describes how to calculate the size of volumes containing LUNs, how to create and manage LUNs and initiator groups (igroups), and how to monitor FCP traffic. This guide assumes you have completed the following tasks to install, set up, and configure your storage appliance:

- ◆ Referred to the *FCP/iSCSI Configuration Guide* to ensure that your configuration is supported.
- ◆ Installed your storage appliance according to the instructions in the *Site Requirements Guide*, other installation documentation, such as the *System Cabinet Guide*, and the hardware and service guide for your specific storage appliance.
- ◆ Configured your filers according to the instructions in the following documents:
 - ❖ *SAN Setup Overview for FCP*
 - ❖ *Data ONTAP Software Setup Guide*
 - ❖ *SAN Host Attach Kit for Fibre Channel Protocol* for your specific host
 - ❖ Any SAN switch documentation for your specific switch, which you can find at http://now.corp.netapp.com/NOW/knowledge/docs/client_filer_index.shtml

Audience

This guide is for system and storage administrators who are familiar with operating systems, such as Windows® 2000 and UNIX®, that run on the hosts that access storage managed by NetApp storage appliances. It also assumes that you know how block access protocols 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

This guide uses the following terms:

- ◆ *Enter* refers to pressing one or more keys on the keyboard and then pressing the Enter key.
- ◆ *Filer* refers to any NetApp storage appliance.
- ◆ *Type* refers to pressing one or more keys on the keyboard.

Command conventions

In examples that illustrate commands executed on a UNIX workstation 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. 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 <i>hostname</i></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 a brief introduction to NetApp storage appliances, how they are administered, and an overview of how NetApp implements the Fibre Channel Protocol (FCP) in a NetApp FC network.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Understanding NetApp storage appliances”](#) on page 2
- ◆ [“Understanding how NetApp implements an FC SAN network”](#) on page 5
- ◆ [“Related documents”](#) on page 10

Understanding NetApp storage appliances

What NetApp storage appliances are

NetApp storage appliances are storage systems, referred to as *filers*, that serve and protect data using protocols for both SAN and NAS networks. For information about filer product families, see <http://www.netapp.com/products/>.

In an FC SAN network, 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 hosts, which are the initiators.

What Data ONTAP is

Data ONTAP is the operating system for all NetApp filers. It provides a complete set of storage management tools through its command-line interface and through the FilerView® interface and DataFabric™ Manager interface.

Data ONTAP supports a multiprotocol environment. You can configure a filer as a target device in a SAN network using the SCSI protocol over FCP (using the fcp service) or in an iSCSI network using the SCSI protocol over TCP/IP (using the iscsi service) to communicate with one or more hosts. You can also configure a filer as a storage device in a NAS network using NFS, CIFS, DAFS, HTTP, or FTP.

Ways to administer a filer

You can administer a filer by using the following methods:

- ◆ Command line
- ◆ FilerView
- ◆ DataFabric Manager. You must purchase the DataFabric Manager license to use this product. For more information about DataFabric Manager, see the DataFabric Manager Information Library at http://now.corp.netapp.com/NOW/knowledge/docs/DFM_win/dfm_index.shtml.

Command-line administration: You can issue Data ONTAP commands at the filer's console or you can open a Telnet or Remote Shell (rsh) session from a host. An Ethernet network interface card (NIC) is preinstalled in the filer, and NetApp highly recommends you use this to connect to a TCP/IP network to perform the following tasks:

- ◆ Manage storage appliances with administration hosts through Telnet or rsh sessions
- ◆ Access FilerView
- ◆ Manage Fibre Channel switches
- ◆ Provide support for the SnapDrive™ feature in a Windows environment

There are two levels of administration at the command level, admin and advanced. The admin level is identified by the *filer_name>* prompt, and the advanced level is identified by the *filer_name*>* prompt. You set the level with the `priv set` command.


Online command-line information aids: When using the command line, you can get command-line syntax help from the command line by entering the name of the command followed by `help` or `?`. You can also access online manual (man) pages by entering the `man na_command_name` command. For example, if you want to read the man page about the `lun` command, you would enter the following command: `man na_lun`.

For more information about filer administration, see the *Data ONTAP 6.5 System Administration Storage Management Guide*.

FilerView administration: As an alternative to entering commands at the command line or using scripts or configuration files, you can use FilerView to perform many common tasks. FilerView is the graphical management interface for managing a filer from a Web browser or for viewing information about the filer, its storage units (such as volumes), LUNs, and adapters, and statistics about the storage units and FCP or iSCSI and network traffic. FilerView is easy to use, and you can access online Help by clicking the `?` button at the topic and field levels. Help information explains Data ONTAP features and how to use them.

Launching FilerView: To launch 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. Result: ◆ If the filer is password protected, you are prompted for a user name and password. ◆ Otherwise, FilerView is launched, and a screen appears with a list of topics in the left panel and the system status in the main panel.
4	Click any of the topics in the left panel to expand navigational links.

Understanding how NetApp implements an FC SAN network

What FCP is

FCP is a licensed service on the filer that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over a Fibre Channel fabric. For information about enabling the fcp license, see [“Managing the FCP service”](#) on [page 110](#).

What nodes are

In an FCP network, nodes include targets, initiators, and switches. Targets are filers and initiators are hosts. Filers have storage devices, which are referred to as LUNs. Nodes register with the Fabric Name Server directory when they are connected to a Fibre Channel (FC) switch.

What LUNs are

From the filer, 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.

From the host, LUNs appear as local disks on the host that you can format and manage to store data, using either the FCP or iSCSI protocols.

What a LUN serial number is

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, VERITAS® 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.

How nodes are connected

Filers and hosts have Host Bus Adapters (HBAs) so they can be connected directly to each other or to FC switches with optical cable. In addition, they can be connected to each other or to TCP/IP switches with Ethernet cable for filer and FC switch administration.

When a node is connected to the FC SAN network, it registers each of its ports with the switch’s Fabric Name Server directory service using a unique identifier.

How nodes are uniquely identified

Every FCP node has a unique identifier. Every manufacturer of FCP nodes assigns their equipment with a World Wide Node Name (WWNN) or a World Wide Port Name (WWPN).

Host: 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*.

To know which WWPN is associated with a specific host, see the SAN HBA attach kit documentation for your host. These documents describe NetApp-supplied or initiator-vendor supplied commands or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, you use the `lputilnt` command, and for Solaris hosts, you use the `sanlun` command.

You can use the `fcv show initiator` command or FilerView (click LUNs > Initiator Groups > Manage) to see all of the WWPNs of the FCP initiators that have logged on to the filer. Data ONTAP displays the WWPN as *Portname*.

Filer: When the FCP service is first initialized, it assigns a WWNN to a filer based on the serial number of its NVRAM adapter. The WWNN is stored on disk. Each target port on the HBAs installed in the filer has a unique 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.

You can use commands, such as `fcv show adapter`, `fcv config`, or `sysconfig -v` or FilerView (click LUNs > FCP > Report) to see the filer's WWNN as *FC nodename* or *nodename* or the filer's WWPN as *FC portname* or *portname*.

Switch: Fibre Channel switches have one WWNN for the device itself, and one WWPN 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 Brocade switch. For details about how the ports are numbered for a particular switch, see the vendor-supplied documentation for that 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**, WWPN 20:**00**:00:60:69:51:06:b4

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

...

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

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

How WWNNs and WWPNs are used

WWNNs and WWPNs are used for the following purposes.

Creating an initiator group: The host's WWPN is used to create an initiator group (igroup). An igroup is a collection of WWPNs of initiators in an FCP network that you specify when you create the igroup. When you map a LUN on a filer to the igroup, you grant all the initiators in that group access to that LUN. If a host's WWPN is not in an igroup that is mapped to a LUN, that host does not have access to the LUN. This effectively masks access to the LUN.

Uniquely identifying a filer's HBA target ports: The combination of the filer's WWNN and WWPN is used to uniquely identify each HBA target port on a filer. Some operating systems require persistent binding to ensure that the LUN appears at the same target ID on the host.

How Data ONTAP supports FCP with non-clustered filers

If your filer is not in a cluster, Port A handles access to LUNs. Port B can provide a second path to the LUNs for configurations with one or more hosts using multipathing software.

How Data ONTAP supports FCP with clustered filers

If your filers are in cluster, Data ONTAP provides multiple modes of operation required to support homogeneous and heterogeneous host operating systems.

FCP is implemented over the target's and initiator's HBA ports. Target HBAs can have one or two ports, and are labeled Port 1 and Port 2 (if there is a second port).

Data ONTAP uses the FCP cfmode setting to control how these target ports:

- ◆ Log into the fabric
- ◆ Handle local and partner traffic for a cluster, in normal operation and in failover

These modes can be viewed with the `fcpl show cfmode` command.

Caution

Changing the FCP cfmode setting on your filer might prevent hosts from being able to access data on mapped LUNs. Please contact Network Appliance Customer Support to modify the FCP cfmode setting.

FCP cfmode settings for clustered configurations: Clustered filers in an FC network require that the FCP cfmode settings are the same. You can use the NetApp Cluster Configuration Checker to detect inconsistent FCP cfmode settings on clusters licensed for FCP.

To download and run the Cluster Configuration Checker, log in to the NOW Web site at <http://now.netapp.com> and go to Software Downloads > Tools & Utilities.

Enabled options for cluster configurations: Clustered filers in an FC network 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 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

How FCP cfmode settings affect target ports

The following settings for FCP cfmode determine how the FCP target ports provide access to LUNs:

- ◆ Standby mode—If you upgrade your filer to Data ONTAP 6.5, the FCP cfmode is set to standby mode. Port A operates as the active port and Port B operates as a standby port. When the cluster is in normal operation, Port A provides access to local LUNs, and Port B is not used. When one filer fails, Port B on the partner filer becomes active and provides access to the LUNs on the failed filer.
- ◆ Partner mode—If you have a filer cluster of F800 series or FAS900 series filers with Data ONTAP 6.5 newly installed, the FCP cfmode setting is set to partner mode. In this mode, Port A and Port B are both active. Port A provides access to local LUNs, and Port B provides access to LUNs on the partner filer.
- ◆ Mixed mode—Each FCP target port supports three virtual ports:
 - ❖ Virtual local port, which provides access to LUNs on the local filer
 - ❖ Virtual standby port, which provides access to LUNs on the failed filer when a takeover occurs. The standby virtual port assumes the WWPN and WWNN of the corresponding B port on the failed partner.
 - ❖ Virtual partner port, which provides access to LUNs on the partner filer. This port enables hosts to bind the physical switch port address to the target device, and allows hosts to use active/passive multipathing software.

Note

Mixed mode requires switches that support public loop and cannot be used if you are accessing a FAS270 cluster.

- ◆ **Dual_fabric**—This is the only supported mode of operation for FAS270 clusters. The FAS270 cluster consists of two storage appliances integrated into a DiskShelf14mk2 FC disk shelf. Each storage appliance has two Fibre Channel ports. The port labeled Fibre Channel 1 operates as a Fibre Channel target port after you license the FCP service and reboot the storage appliance. The port labeled Fibre Channel 2 connects to the internal disks, enabling you to connect additional disk shelves to an FAS270 cluster. Port 1 of each FAS270 appliance in the cluster supports three virtual ports:
 - ❖ Virtual local port, which provides access to LUNs on the local FAS270
 - ❖ Virtual standby port, which is not used
 - ❖ Virtual partner port, which provides access to LUNs on the partner node

How to display FCP target ports

Data ONTAP displays information about the ports by using the slot number where the HBA is installed in the filer. The display also depends on the FCP cfmode setting. For example, if the FCP cfmode setting is standby or partner and the target HBA is installed in slot 5 of a filer, the FCP commands display Port 1 as 5a and Port 2 as 5b. In addition, the local port name and node name addresses of the WWNN and WWPNN have a pattern of 50:a9:80:nn:nn:nn:nn. The standby port name and node name addresses of the WWNN and WWPNN have a pattern of 20:01:00:nn:nn:nn:nn.

If the FCP cfmode setting is mixed, and the target HBA is installed in slot 6, FCP commands display the HBA as 6a, the virtual local port as 6a_0, the virtual standby port as 6a_1, and the virtual partner port as 6a_2.

For FAS270c appliances only: For FAS270c appliances, the FCP cfmode setting is dual_fabric, and each port is configured as a virtual port. The ports are displayed as 0c_0, 0c_1, and 0c_2.

How to manage target port resources

Each target port has a fixed number of resources, or command blocks, for incoming initiator requests. When all command blocks are used, an initiator receives a QFull message on subsequent requests. Data ONTAP enables you to monitor these requests and manage the number of command blocks available for specified initiators. You can limit the number of command blocks used by the initiator in an igroup. This is known as igroup throttling. For information about igroup throttling, see [“Managing initiator requests”](#) on page 75.

Related documents

Where to go for more information

The following table lists documents on NetApp's NOW Web site at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>, unless specified otherwise, with the most current information about host initiator and filer system requirements and additional documentation.

If you want...	Go to...
The most current system requirements for your host	<i>FCP/iSCSI Configuration Guide</i>
The latest information about the supported filer models for Data ONTAP licensed with FCP	◆ <i>FCP/iSCSI Configuration Guide</i> ◆ <i>System Configuration Guide</i>
Information about how to install and configure SAN HBAs	◆ <i>NetApp SAN Setup Overview for FCP</i> ◆ <i>SAN Host Attach Kit Installation and Setup Guide</i> for your specific host that was supplied with the adapter and at http://now.netapp.com/NOW/knowledge/docs/client_filer_index.shtml
The latest information about how to configure the FCP service on a filer	<i>Data ONTAP 6.5 Release Notes</i> (if available)

About this chapter

This chapter describes the preparation process for creating LUNs and the methods for creating LUNs, igroups, and LUN maps.

This chapter assumes that your NetApp SAN is set up and configured, and that the FCP service is licensed and enabled. If that is not the case, see [“Managing the FCP service”](#) on page 110 for information about these topics.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Configuring disk space”](#) on page 12
- ◆ [“Guidelines for creating volumes”](#) on page 15
- ◆ [“Guidelines for calculating the size of a volume”](#) on page 21
- ◆ [“Methods of creating LUNs, igroups, and LUN maps”](#) on page 26

Configuring disk space

Filer storage units for managing disk space

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


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

The volume is the largest unit 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. You use LUNs to store data in a SAN environment, and you use files to store data in a NAS environment.

Where LUNs reside: You create LUNs at the root of a volume or a qtree, with the exception of /vol/vol0. Do not create LUNs at /vol/vol0 because this is used by Data ONTAP for system administration.

The following diagram shows four LUNs that were created by using the storage from Filer_A. The four LUNs are:

- ◆ /vol/vol1/LUN0—created at volume /vol/vol1
- ◆ /vol/vol2/LUN1—created at /vol/vol2
- ◆ /vol/vol2/qtree1/LUN2—created at the qtree of volume /vol/vol2, at /vol/vol2/qtree1
- ◆ /vol/vol2/qtree1/LUN3—also created at the qtree of volume /vol/vol2, at /vol/vol2/qtree1



Why the space reservation option is important

LUNs appear as local disks attached to the 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 must have guaranteed access to that amount of space. Data ONTAP achieves this with the space reservation option.

When you create a LUN of a specified size, Data ONTAP turns the space reservation on by default and reserves the number of blocks (not the physical blocks) in the WAFL® volume that are required to complete writes to the LUN. This ensures that writes to the LUN do not fail due to a lack of disk space.

How space reservation persists

Space reservation settings persist across reboots, takeovers, givebacks, and snap restores. The space reservation setting persists if you revert from Data ONTAP 6.5 to Data ONTAP 6.4, but not if you revert from Data ONTAP 6.4 to Data ONTAP 6.3. Single File SnapRestore restores the reserved state of a LUN to the reserved state at the time the snapshot was taken.

Why you need to change snapshot defaults

NetApp snapshots are required for many NetApp features, such as the SnapMirror® feature, SyncMirror™ feature, SnapVault™ feature, dump and restore, and ndmpcopy.

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 for calculating the size of a volume”](#) on page 21.

When you create a volume, Data ONTAP automatically

- ◆ Enables snapshots enabled by default
- ◆ Reserves 20% of the space for snapshots (snap reserve, or *snapshot reserve* in FilerView)
- ◆ Schedules snapshots

Because the internal scheduling mechanism for taking snapshots within Data ONTAP has no means of ensuring that the data within a LUN is in a consistent state, NetApp recommends that you change these snapshot settings by performing the following tasks:

- ◆ Set the percentage of snap reserve to 0%
- ◆ Turn off the automatic snap schedule

With Windows hosts, you use SnapDrive and SnapManager to ensure that applications accessing LUNs are quiesced or synchronized automatically before taking snapshots. With hosts that are not supported with SnapDrive, NetApp recommends that you ensure that the file system or application accessing the LUN is quiesced or synchronized before taking snapshots. For information about how to use snapshots, see [“Using snapshots”](#) on page 84.

Guidelines for creating volumes

Required information to create a volume

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 72-GB disks and you use seven disks to create a volume, six disks are used for data, and one is used for parity; the volume capacity is 432 GB. Assuming that a 72-GB disk results in 67.9 GB of physical capacity, six 72-GB disks provide 407.4 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 for calculating the size of a volume”](#) on page 21.

Optional information to create a volume

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 6.5 System Administration Storage Management Guide*. For information about disk configurations for your filer model, see the *System Configuration Guide*.

Guidelines to use when creating volumes

NetApp strongly recommends that you create volumes using the following guidelines:

- ◆ Do not create any LUNs on /vol/vol0.
When you first set up and configure the filer, 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 for calculating the size of a volume”](#) on page 21. At a minimum, create volumes using seven disks.
- ◆ 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.
- ◆ 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.

Setting the percentage of snap reserve space using the command line

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

Step	Action
1	To set the percentage, enter the following command: <pre>snap reserve <i>volname percent</i></pre> Note_____ For volumes that contain LUNs and no snapshots, NetApp recommends that you set the percentage to zero. _____ Example: <code>snap reserve vol1 0</code>
2	To verify what percentage is set, enter the following command: <pre>snap reserve [<i>volname</i>]</pre> Example: <code>snap reserve vol1</code> 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 use FilerView to set a percentage of snap reserve space (or *snapshot* reserve space, as it appears in FilerView) on a volume and to verify what percentage is set, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Click Volumes > Snapshots > Configure.
3	Select the volume you want to configure.
4	At the Snapshot Reserve field, enter 0 as the percentage of space the volume reserves for snapshots. Note _____For volumes that contain LUNs and no snapshots, NetApp recommends that you set the percentage to zero._____
5	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 off the automatic snapshot schedule, enter the following command: snap sched volname 0 0 0 Example: snap sched vol1 0 0 0 Result: This effectively turns off the snapshot schedule because no times are specified for snapshots to be taken automatically. You can still take snapshots manually using the snap command.
2	To verify that 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	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Click Volumes > Snapshots > Configure.
3	Select the volume you want to configure.
4	At the Hourly Snapshot Schedule, ensure that no time slots are selected. For example, if a check appears at 8:00 AM, click it to deselect it.
5	Click the Apply button.

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

To use the command line to verify that the Create Unicode Format Directories By Default volume option is enabled, or to enable the option, complete the following steps.

Step	Action
1	To verify that the <code>create_ucose</code> option is enabled (On), enter the following command: <pre>vol status [volname] -v</pre> Example: <code>vol status /vol1 -v</code> 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 Plex /vol/vol1/plex0: online, normal, active RAID group /vol/vol1/plex0/rg0: normal</pre> Note _____ If you do not specify a volume, status on all volumes is displayed.
2	To enable the <code>create_ucose</code> option, enter the following command: <pre>vol options volname create_ucose on</pre> Example: <code>vol options vol1 create_ucose on</code>

**Verifying and
modifying the
volume option
create_unicode using
FilerView**

To use FilerView to verify that the Create Unicode Format Directories By Default volume option is enabled, or to enable the option, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Click Volumes from the left panel.
3	Click Manage.
4	Locate the name of the volume you want to check, and click the Modify icon for that volume.
5	Locate the Create New Directories in Unicode field and select On if it is Off.
6	Click the Apply button.

Guidelines for calculating the size of a volume

What the volume size depends on

NetApp recommends that you use a minimum of seven disks to create a volume, so the size of your smallest volume should be six times your disk capacity (because one disk is for parity). In the example at “[Required information to create a volume](#)” on page 15, the disk size was 72 GB. A volume created with six of these disks provided a volume size of 407.4 GB.

If your volume requires more than 407 GB, you need to calculate its size based on the types and sizes of the storage units you intend to create on the volume. The storage units might be any of the following:

- ◆ LUNs
- ◆ Snapshots
- ◆ Writable snapshots of LUNs (or *backing store* LUNs)
A backing store LUN means some or all of the unchanged data in the LUN is stored in the snapshot that was used to create the LUN.
- ◆ LUN clones

Use the following table to determine the size of the volume that will contain the storage units 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 = 10 GB, volume size = 1.10×10 GB or 11 GB

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Take one or more snapshots	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 delta. If LUN size = 10 GB, and the rate of change since the snapshot is 15%, volume size = $2 \times 10 \text{ GB} + (.15 \times 10 \text{ GB}) = 21.5 \text{ GB}$. The first snapshot requires enough space to overwrite every block of data in the active file system; however, subsequent snapshots only require enough space to store what is overwritten in snapshot data. They do not require additional reserved space.
Create 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, because 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 snap restore	Has enough space to store the original LUN and the restored LUN.	Volume size = $2x + 10\%x$ or volume size = $2.1x$.

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Create a clone of the 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 is $4x + \text{delta}$.

Creating a volume using the command line

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

Step	Action
1	On the filer's command line, enter the following command: <pre>vol create volname disk-list</pre> <i>volname</i> is the name of the new volume (include the /vol/ prefix). <i>disk-list</i> is one of the following: ◆ <code>ndisks [@disk-size]</code> <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. ◆ <code>-d disk_name1 disk_name2... disk_nameN</code> <i>disk_name1</i>, <i>disk_name2</i>... and <i>disk_nameN</i> are disk IDs of one or more available disks. Separate multiple disks with a space. Example: <code>vol create vol1 8</code> For more information about the <code>vol</code> command, see the <i>Storage Management Guide</i>.

Creating a volume using FilerView

To use FilerView to create a volume, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	In the left panel, click Volumes.
3	Click Add. Result: The Volume Wizard appears. If you do not have enough disks, you see a message at the top of the Wizard. For example, if you only have one spare disk, you see the following message: The number of spares is 1. At least two spares must exist to create a new volume. Add enough disks to the filer to create a new volume.
4	Fill in the fields.
5	Click the Add button.

Organizing LUNs by creating qtrees using the command line

You can use qtrees to organize related LUNs into groups. To use the command line 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: <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

To use FilerView to create a qtree in a volume, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	In the left panel, click Volumes > Manage > Qtrees > Manage > Add Qtree button.
3	Fill in the fields.
4	Click the Add button.

Methods of creating LUNs, igroups, and LUN maps

About the methods

You create LUNs and igroups using one of the following methods:

- ◆ Entering the lun setup command
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup. For information about this method, see [“Creating LUNs with the lun setup program”](#) on page 32.
- ◆ Using FilerView
This method provides a LUN Wizard that steps you through the process of creating and mapping new LUNs. For information about this method, see [“Creating LUNs and igroups with FilerView”](#) on page 37.
- ◆ Entering a series of individual commands (such as lun create, igroup create, and lun map)
This method allows you to create one or more LUNs and igroups in any order. For information about this method, see [“Creating LUNs and igroups using individual commands”](#) on page 40.
- ◆ Using SnapDrive 2.0 or later for Windows
SnapDrive furnishes a Windows-like interface to manage LUNs on a filer. This product provides online storage configuration, LUN expansion, and streamlined storage management. Integrated with NetApp Snapshot™ functionality, SnapDrive 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 *SnapDrive Installation and Administration Guide*.

Note

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

What is required to create a LUN

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. Do not create LUNs at /vol/vol0.

For clustered filer configurations, NetApp recommends that you distribute LUNs across the filer cluster.

Note

You might find it useful to provide a meaningful path name for the LUN. For example, you might choose a name that describes how the LUN is used, such as the name of the application the kind of data that it is, or the user using the data. Examples are `/vol/database/lun0`, `/vol/finance/lun1`, or `/vol/bill/lun2`.

The host operating system type: The host operating system type (`ostype`) indicates the type of operating system running on the host that accesses the LUN, which also determines the geometry used to access data on the LUN.

- ◆ When you use the `lun setup` command, FilerView, or the `lun create` command, the LUN `ostype` values are *image*, *solaris*, *windows*, *hpux*, *aix*, and *linux*.
- ◆ Specifying the operating system determines the geometry used to store data on the LUN. For example, selecting *solaris* specifies that the LUN will store a Solaris raw disk in a single-partition. If you select *image*, the contents of the LUN are not specified.
- ◆ If you know that the operating system is supported, specify that `ostype`; otherwise, specify *image* to prevent Data ONTAP from making any assumptions about the geometry used for the LUN.

Not all versions of Data ONTAP support all operating system types. Neither protocol (FCP and iSCSI) supports all operating systems. For information about what is supported, see the *FCP/iSCSI Configuration Guide* or the release notes that came with your version of Data ONTAP.

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. 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

Multiplier suffix	Size
g	gigabytes
t	terabytes

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. 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.

The disk geometry used by the operating system determines the minimum and maximum values of LUNs. Windows geometry is characterized by 63 sectors/track and 255 heads. Solaris geometry is characterized by 16 heads and a cylinder size which is a multiple of 1MB. The minimum LUN size is 4 cylinders. For Solaris, that dictates a minimum LUN size of 4MB. For Windows, that dictates a minimum of 31.5MB. For information about the maximum sizes for LUNs, see the SAN HBA host attach kit documentation for your host.

A brief description of the LUN (optional): You use this attribute to store alphanumeric information about the LUN. You can edit this description at the command line or with FilerView.

A LUN identification number (LUN ID). A LUN must have a unique LUN ID so the host can identify and access the LUN. This is used to create the map between the LUN and the host.

Space reservation setting: When you create a LUN, the space reservation is turned on. NetApp recommends that you keep this setting on.

About igroups

Initiator groups are tables of WWPNs of hosts that you use to control access to LUNs. Typically, you 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.

You can create igroups that specify which initiators have access to the LUNs either before or after you create LUNs, but you must create igroups before you can map a LUN to an igroup.

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. The following table illustrates how four igroups give access to the LUNs for four different hosts using the FCP service. The clustered hosts (Host3

and Host4) use the same igroup for common access to the LUNs mapped to solaris-group2. The igroup named solaris-group3 contains the WWPNs of Host4 to store local information not intended to be seen by its partner.

Host with HBA WWPNs	igroups	WWPNs added to igroups	LUNs mapped to igroups
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/qtrees1/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/qtrees1/lun2 /vol/vol1/qtrees1/lun3

What is required to create an igroup

Whichever method you choose, you create an igroup by specifying the following attributes:

The name of the igroup: This is a case-sensitive name that meets the following requirements:

- ◆ Contains 1 to 96 alphanumeric characters.
- ◆ Can contain any except the following special characters:
&, #, -, ' , " , blank, or tab.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or DNS. If you name an igroup `sun1`, for example, it is not mapped to the actual IP host name (DNS name) of the host.

Note

You might find it useful to provide meaningful names for igroups, ones that describe the hosts that can access the LUNs mapped to them.

The type of igroup: The type is `-f` for FCP if it contains WWPNs or `-i` for iSCSI if it contains iSCSI node names. An igroup can contain only one type of unique identifier of nodes, either FCP or iSCSI.

The ostype of the initiators: The ostype indicates the type of operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype.

- ◆ When you use the `lun setup` command, FilerView, or the `igroup create` command, the igroup ostype values are *default*, *solaris*, *windows*, *hpux*, *aix*, and *linux*.
- ◆ If you know that the operating system is supported, specify that ostype; otherwise, specify *default*.

Not all versions of Data ONTAP support all ostyles. Neither protocol (FCP and iSCSI) supports all operating systems. For information about what is supported, see the *FCP/iSCSI Configuration Guide* or the release notes that came with your version of Data ONTAP.

WWPNs of the initiators: You can specify the WWPNs of the initiators when you create an igroup. You can also add them or remove them at a later time.

What is required to map a LUN to an igroup

Whichever method you choose, you map a LUN to an igroup by specifying the following attributes:

LUN name: You must specify the path name of the LUN to be mapped.

Initiator group: You must specify the name of the igroup that contains the host you want to give LUN access to.

LUN ID: You assign a number for the LUN ID or accept the default LUN ID. Typically, the default LUN ID begins with 0 and increments by 1 for each additional LUN as it is created. The range of valid LUN ID numbers are host-dependent. For example,

- ◆ For Solaris, the range is 0–4096

- ◆ For Windows, the range is 0-254

For information about the ranges of values to use when assigning LUN IDs for your host, see the appropriate *SAN Host Attach Kit Installation and Setup Guide* for your host.

The host associates the LUN ID with the location and path name of the LUN.

When you map the LUN ID to the igroup, you grant the initiators in the igroup access to the LUN. If you do not map a LUN, the LUN is not accessible to any hosts. This effectively masks access to the LUN. Data ONTAP maintains a separate LUN map for each igroup to support a large number of hosts and to enforce access control.

Guidelines for mapping LUNs

Use the following guidelines when mapping LUNs:

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

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 igroup, and for mapping the LUN to the igroup. When a default is provided in brackets in the prompt, you can press Enter to accept it.

Prerequisites for running the lun setup program

If you did not create volumes to store LUNs in prior to running the lun setup program, you must terminate the program and create volumes. If you want to use qtrees, you must also create them before running 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. <pre>lun setup</pre> Result: The lun setup program displays the following instructions. Press Enter to continue or n to terminate the program. <pre> 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]:</pre>
2	Specify the operating system that will be accessing the LUN by responding to the next prompt: OS type of LUN (image/solaris/windows/hpux/linux) [image]: Example: solaris For information about specifying the ostype of the LUN, see “The host operating system type” on page 27.

Step	Action
3	Specify the name of the LUN where it 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/ and want to create a LUN called lun0, 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 you accept the configuration information later in this program.
4	Specify whether you want the LUN created with space reservations enabled by responding to the prompt: 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: Enter Caution_____ If you choose n, space reservation is disabled. This might cause writes to the filer to fail, which can cause data corruption. NetApp strongly recommends that you enable space reservations. _____
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 you accept the configuration information 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: GB Solaris LUN for finance records If you choose not to provide a comment at this time, you can add a comment later with the lun comment command or fill in the description field using FilerView.
7	Create or use an igroup 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[]: Result: If you have already created one or more igroups, you can enter ? to list them. The last igroup you used appears as the default. If you press Enter, that igroup will be used. If you have not created any igroups, enter a name of the igroup you want to create now. For information about naming an igroup, see “The name of the igroup” on page 29.
8	Specify which protocol will be used by the hosts in the igroup by responding to the next prompt: Type of initiator group solaris-igroup3 (FCP/iSCSI)[FCP]: Result: The initiators in this igroup use the FCP protocol.
9	Add the WWPNs of the hosts that will be in the igroup 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:

Step	Action
	Example 1a: ? Result: The following output is a sample of what is displayed: <pre> Initiators connected on adapter 4a 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 Adapter 4b is running on behalf of the partner. Initiators connected on adapter 5a: None connected. Enter comma separated portnames: </pre> Example 1b: Enter a WWPN, for example, 10:00:00:c9:2b:cc:51. Result: The initiator identified by this WWPN is added to the igroup that you specified in Step 7. You are prompted for more port names until you press Enter. For information about how to determine which WWPN is associated with a host, see “Host” on page 6.
10	Specify the operating system type initiators in the igroup use to access LUNs by responding to the next prompt: The initiator group has an associated OS type. The following are currently supported: solaris, windows, hpux, aix, linux, or default. OS type of initiator group “solaris-igroup3” (default/solaris/windows) [default]: For information about specifying the ostype of an igroup, see “About igroups” on page 28.

Step	Action
11	Specify the LUN ID that the host will map to the LUN by responding to the next prompt: 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 : solaris 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>
12	Commit the configuration information you entered by responding to the next prompt: Do you want to accept this configuration? [y] Result: If you press Enter the default, the LUNs are mapped to the specified igroup. 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.
13	Either continue creating LUNs or terminate the program by responding to the next prompt: Do you want to create another LUN? [n]

Creating LUNs and igroups with FilerView

Methods of creating LUNs

You can use FilerView to create LUNs and igroups with the following methods:

- ◆ LUN Wizard
- ◆ Menu
 - ❖ Create LUN
 - ❖ Create igroup
 - ❖ Map LUN

Creating LUNs with the LUN Wizard

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

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Click LUNs in the left panel of the FilerView screen. Result: The management tasks you can perform on LUNs are displayed.
3	Click Wizard. Result: The LUN Wizard window appears and explains how to use it.

Step	Action
4	Click the Next button to continue. Result: The first window of fields in the LUN Wizard appears. 
5	Follow the instructions in the LUN Wizard and enter information in the appropriate fields.
6	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 igroups you specified.

Creating LUNs and igroups with FilerView menus

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

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Go to LUNs > Add.
3	Fill in the fields.

Step	Action
4	Click the Add button to commit changes.

Creating igroups: To create an igroup using FilerView menus, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Go to Initiator Groups > Add.
3	Fill in the fields.
4	Click the Add button to commit changes.

Mapping LUNs to igroups: To map LUNs to igroups using FilerView menus, complete the following steps.

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Go to LUNs > Manage.
3	In the first column, find the LUN you want to map an igroup to. ◆ If the LUN is mapped, <i>yes</i> or the name of the igroup and the LUN ID appears in the last column. ◆ If the LUN is not mapped, <i>no</i> or <i>No Maps</i> appears in the last column.
4	Click the Add Groups to Map link.
5	Click the link under the Maps column for the LUN you want to map an igroup to.
6	Select an igroup name from the list on the right side of the window.
7	Click the Add button to commit the changes.

Creating LUNs and igroups using individual commands

How to use individual commands

The commands in the following table occur in a logical sequence for creating LUNs and igroups for the first time. However, you can use the commands in any order, or you can skip a command if you already have the information that a particular command displays.

For more information about all of the options for these commands, see the online man pages. For information about how to view man pages, see “[Command-line administration](#)” on page 2.

To do this...	Use this command...
Display the WWPNs of the initiators that are available to create igroups	<pre>fcv 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>
Determine which hosts are associated with the WWPNs	For information about how to determine which WWPN is associated with a host, see “ Host ” on page 6.

To do this...	Use this command...
Create an igroup	<code>igroup create -f -t <i>ostype</i> <i>initiator_group</i> [<i>node</i>]</code> -f indicates that the igroup contains Fibre Channel WWPNs -t <i>ostype</i> indicates the operating system type of the initiator. The values are: default, solaris, windows, hpux, aix, or linux. For information about specifying the ostype of an igroup, see “About igroups” on page 28. <i>initiator_group</i> is the name you specify as the name of the igroup. <i>node</i> is a WWPN, which is the 64-bit address of the initiator’s port name Example: <pre>igroup create -f -t solaris solaris-igroup3 10:00:00:00c:2b:cc:92</pre>
Create a LUN	<code>lun create -s <i>size</i> -t <i>ostype</i> <i>lun_path</i></code> -s indicates the size of the LUN to be created, in bytes by default. For information about LUN size, see “The size of the LUN” on page 27. -t <i>ostype</i> indicates the operating system type that determines the geometry used to store data on the LUN. For information about specifying the ostype of the LUN, see “The host operating system type” on page 27. <i>lun_path</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.

To do this...	Use this command...
Map the LUN to an igroup	<code>lun map <i>lun_path</i> <i>initiator_group</i> [<i>lun_id</i>]</code> <i>lun_path</i> is the path name of the LUN you created. <i>initiator_group</i> is the name of the igroup you created. <i>lun_id</i> is the identification number that you specify the initiator should use when the LUN is mapped to it. If you do not enter a number, Data ONAP generates the next available LUN ID number. Example 1: <code>lun map /vol/vol1/mtree1/lun3 solaris-igroup3 0</code> Result: Data ONTAP maps /vol/vol1/mtree1/lun3 to the igroup 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 igroup. After entering the command in this example, Data ONTAP displays the following message: <code>lun map: auto-assigned solaris-igroup0=0</code>
Display the LUNs you created	<code>lun show -v</code> -v provides additional information, such as the comment string, serial number, and LUN mapping. 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	lun show -m -m provides mapping information in a tabular format. Sample result: <table><tr><th>LUN path</th><th>Mapped to</th><th>LUN ID</th></tr><tr><td>-----</td><td>-----</td><td>-----</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 igroup, which was named to reflect the host whose WWPN was used when creating the igroup. Note that the LUN ID mapping begins with 0 for each igroup. If a second LUN is mapped to the same igroup, 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 and LUN maps](#)” on page 46
- ◆ “[Displaying LUN information](#)” on page 52
- ◆ “[Monitoring disk space](#)” on page 56

Managing LUNs and LUN maps

Tasks to manage LUNs and LUN maps

You can use the command-line interface or FilerView to

- ◆ Control block protocol access
- ◆ Unmap a LUN from an igroup
- ◆ Rename a LUN
- ◆ Modify the LUN description
- ◆ Modify a LUN's serial number
- ◆ Enable or disable space reservations
- ◆ Remove a LUN
- ◆ Access a LUN with NAS protocols

Controlling block protocol access

The `lun online` and `lun offline` commands enable and disable block protocol accesses while preserving mappings.

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: lun offline <i>lun-path</i> [<i>lun-path</i> ...] Example: <code>lun offline /vol/vol1/lun0</code>

Unmapping a LUN from an igroup

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

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

Renaming a LUN

To rename a LUN, completing the following step.

Step	Action
1	Enter the following command: lun move <i>lun_path new_lun_path</i> Example: lun move /vol/vol1/mylun /vol/vol1/mynewlun Note _____The existing <i>lun_path</i> and the new <i>lun_path</i> must be in the same qtree. _____

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.

Restrictions on resizing a LUN: The following restrictions apply:

- ◆ 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.
- ◆ On Solaris, if you are running VxVM, resizing LUNs is not supported.

If you want to increase the size of the LUN, the SCSI disk geometry imposes an upper limit for how large a LUN can be resized, and Data ONTAP imposes a maximum increase of 2 TB, in 1 MB increments.

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

Caution

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

Step	Action
1	Take a LUN offline before resizing it by entering the following command: lun offline <i>lun-path</i> Example: <code>lun offline /vol/vol1/qtree/lun2</code>
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 <code>lun resize /vol/vol1/qtree1/lun2 10g</code>
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, and then use the Solaris <code>mkfs</code> command to increase the file system on the disk slice to fill the newly enlarged slice.

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. _____

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: <pre>lun offline <i>lun-path</i></pre> Example: lun offline /vol/vol1/qtree1/lun2
2	Modify the LUN’s serial number by entering the following command: <pre>lun serial <i>lun-path</i> <i>new-serial-number</i></pre> Example: lun serial /vol/vol1/qtree1/lun2 0cmePIdQxjch

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. When this occurs, Data ONTAP displays system messages (one message per file) on 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.

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: <pre>lun destroy [-f] <i>lun-path</i> [<i>lun-path</i> ...]</pre> -f forces the <code>lun destroy</code> 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 steps.

Step	Action
1	Determine whether you want to read, write, or do both to the LUN over the NAS protocol and take the 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.

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
- ◆ Status 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><tr><td>lun help</td><td>- List LUN (logical unit of block storage) commands</td></tr><tr><td>lun config-check</td><td>- Check all lun/igroup/fcp settings for correctness</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 snap</td><td>- Manage LUN and snapshot interactions</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></table>	lun help	- List LUN (logical unit of block storage) commands	lun config-check	- Check all lun/igroup/fcp settings for correctness	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 snap	- Manage LUN and snapshot interactions	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 config-check	- Check all lun/igroup/fcp settings for correctness																																						
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 snap	- Manage LUN and snapshot interactions																																						
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> -z resets the counter at any time. -k displays the statistics in KBs. Note_____ The statistics start at zero at boot time. _____ Sample Result: <pre>lun stats -k /vol/vol1/lun0 (1 day, 23 hours, 9 minutes, 27 seconds) Read (kBytes) Write (kBytes) Read Ops Write Ops 10241 26159 3290 4995 /vol/vol1/lun1 Read (kBytes) Write (kBytes) Read Ops Write Ops 10241 26279 3313 7195</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: lun set reservation <i>lun-path</i>

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, respectively. You can use the `-h` option to let Data ONTAP scale the size to the appropriate format, based on the value. You can also use the `-r` option, which is explained in the next section.

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. The statistics displayed are as follows:

- ◆ Amount of total space on the volume, in the size of bytes you specify
Total space = used space + available space
- ◆ Amount of used space
Used space = space storing data + space storing snapshots + space storing reserved space
- ◆ Amount of available space
Available space = space that is not used or reserved; it is free space
- ◆ Percentage of the volume capacity being used

The `-r` option changes the last column to report on the amount of reserved space; that is, how much of the used space is reserved for writing data to existing LUNs.

In the statistics displayed for the snapshot directory, the sum of used space and available space can be larger than the total space for that volume. This is because the additional space used by snapshots is also counted in the used space of the active file system.

Effect of LUN and snapshot operations on disk space The following table 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 that 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.	N/A
Create a 40-GB LUN on that volume.	Used space = 40 GB Reserved space = 0 GB Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	Used space is 40 GB for the LUN. If the LUN size was limited to accommodate at least one snapshot when it was created, 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.	The amount of used space does not change because with space reservations set to On, the same amount of space is used when you write to the LUN as when you created the LUN.
Create a snapshot of the LUN.	Used space = 80 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. As soon as a snapshot is created, the reserved space must be large enough to ensure that any future writes to the LUN succeed. Reserved space is now 40 GB, the same size of the LUN. Data ONTAP always displays the amount of reserved space required for successful writes to LUNs. Because reserved space is also counted as used space, used space is 80 GB.

Action	Result	Comment
Overwrite all 40 GB of data on the LUN with new data.	Used space = 100 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 100 GB Snapshot creation is blocked.	Data ONTAP manages the space so that the overwrite increases used space to 100 GB and decreases available space to 0. The 40 GB for reserved space is still displayed. You cannot take another snapshot 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 = 120 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	After you expand the volume, the amount of used space displays the amount needed for the 40 GB LUN, the 40 GB snapshot, and 40 GB of reserved space. Free space becomes available again, so snapshot creation is no longer blocked.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 120 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 = 160 GB Reserved space = 40 GB Available space = 40 GB Volume Total: 200 GB Snapshot creation is allowed.	The snapshot locks all 40 GB of data currently on the LUN. The used space is the sum of 40 GB for the LUN, 40 GB for each snapshot, and 40 GB for reserved space.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 160 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 = 200 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 200 GB Snapshot creation is blocked.	The amount of used space increases by the amount of LUN expansion. The amount of reserved space remains at 40 GB. Because the available space has decreased to 0, snapshot creation is blocked.
Delete both snapshots of the volume.	Used space = 80 GB Reserved space = 0 GB Available space = 120 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 0 GB. 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.

Examples of disk space monitoring

The following examples illustrate how to monitor disk space when you create LUNs in various scenarios. They do not include every step required to configure the filer or to perform tasks on the host.

- ◆ Without using snapshots
- ◆ Using snapshots
- ◆ Using backing store LUNs and LUN clones

In the examples, assume that the filer is named *toaster* and the host is named *solaris_host*, and that two HBAs are installed in the host, lpfc0 (WWPN: 10000000c92f6cd0) and lpfc1 (WWPN:10000000c92f6a24).

Monitoring disk space without using snapshots: The following example illustrates how to monitor disk space on a volume when you create a LUN without using snapshots. For the purpose of this example, assume that you require less than the minimum amount of volume capacity based on the NetApp recommendation of creating a seven-disk volume.

For simplicity, assume the LUN only requires 3 GB of disk space. The information from “[Guidelines for calculating the size of a volume](#)” on page 21 recommends that the volume size be approximately 3 GB plus 10%. If you plan

to use 72 GB disks (which typically results in 67.9 GB of physical capacity, depending on the manufacturer), two disks provide more than enough space, one for data and one for parity.

To work through the example, complete the following steps.

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 LUNs because snap reserve is set to 20% by default. <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>1440</td><td>50118488</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>12529980</td><td>0</td><td>12529980</td><td>0</td><td>/vol/volsp ace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace	50119928	1440	50118488	0	/vol/volspace/	/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/volsp ace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/														
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/volsp ace/.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>1440</td><td>62648468</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0 /vol/vols pace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/	/vol/volspace/.snapshot	0	0	0	0	0 /vol/vols pace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/														
/vol/volspace/.snapshot	0	0	0	0	0 /vol/vols pace/.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. 3 GB of space is used because this is the amount of space specified for the LUN, and space reservation is enabled by default. <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/</td><td>62649908</td><td>3150268</td><td>59499640</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0 /vol/volsp</td></tr><tr><td>ace/.snapshot</td><td></td><td></td><td></td><td></td><td></td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	3150268	59499640	0	/vol/volspace/	/vol/volspace/.snapshot	0	0	0	0	0 /vol/volsp	ace/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace/	62649908	3150268	59499640	0	/vol/volspace/																				
/vol/volspace/.snapshot	0	0	0	0	0 /vol/volsp																				
ace/.snapshot																									
4	Create an igroup named <i>sol-igroup1</i> and map the LUN to it by entering the following commands (assuming that your host has an HBA whose WWPN is 100000002bcc92). Depending on your host, you might 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 LUN, format it, make the file system available to the host, and write data to the file system. For information about these procedures, see the <i>SAN Host Attach Kit Installation and Setup Guide</i> that came with your SAN Host Attach Kit. These commands have no effect on disk space.																								
6	From the filer, ensure that creating the file system on the LUN and writing data to it has no effect on space on the filer by entering the following command: <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. From the filer, the amount of space used by the LUN remains 3 GB. <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/</td><td>62649908</td><td>3150268</td><td>59499640</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0 /vol/vol</td></tr><tr><td>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	3150268	59499640	0	/vol/volspace/	/vol/volspace/.snapshot	0	0	0	0	0 /vol/vol	space/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace/	62649908	3150268	59499640	0	/vol/volspace/																				
/vol/volspace/.snapshot	0	0	0	0	0 /vol/vol																				
space/.snapshot																									

Step	Action																		
7	Turn off space reservations and see the effect on space by entering the following commands: <pre>toaster> lun set reservation /vol/volspace/lun0 disable</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The 3 GB of space for the LUN is no longer reserved, so it is not counted as used space; it is now available space. Any other requests to write data to the volume can occupy all the available space, including the 3 GB that the LUN expects to have. If the available space is used before the LUN is written to, writes to the LUN fail. To restore the reserved space for the LUN, turn space reservations on. <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/</td><td>62649908</td><td>144</td><td>62649584</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td></td><td>0</td><td>0 /vol/volspace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	144	62649584	0	/vol/volspace/	/vol/volspace/.snapshot	0	0		0	0 /vol/volspace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	144	62649584	0	/vol/volspace/														
/vol/volspace/.snapshot	0	0		0	0 /vol/volspace/.snapshot														

Monitoring disk space using snapshots: The following example illustrates how to monitor disk space on a volume when taking snapshots. Assume that you start with a new volume and the LUN requires 3 GB of disk space. The information from “[Guidelines for calculating the size of a volume](#)” on page 21 recommends that the volume size be approximately 2*3 GB plus delta. Assuming the amount of change is small, delta would be minimal, so using two 72 GB disks should still provide more than enough space.

To work through the example, complete the following steps.

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><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace</td><td>50119928</td><td>1440</td><td>50118488</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>12529980</td><td>0</td><td>12529980</td><td>0</td><td>/vol/vol s pace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace	50119928	1440	50118488	0	/vol/volspace/	/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/vol s pace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/														
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/vol s pace/.snapshot														
2	Set the percentage of snap reserve space to zero by entering the following command: <pre>toaster> snap reserve volspace 0</pre>																		
3	Create a LUN (/vol/volspace/lun0) by entering the following commands: <pre>toaster> lun create -s 6g -t solaris /vol/volspace/lun0</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. Approximately 6 GB of space is taken from available space and is displayed as used 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/</td><td>62649908</td><td>6300536</td><td>56169372</td><td>0</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/volsp ace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	6300536	56169372	0	/vol/volspace/	/vol/volspace/.snapshot		0	0	0	0 /vol/volsp ace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	6300536	56169372	0	/vol/volspace/														
/vol/volspace/.snapshot		0	0	0	0 /vol/volsp ace/.snapshot														
4	Create an igroup named <i>sol-igroup1</i> and map the LUN to the igroup by entering the following commands. 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 about these procedures, see the <i>SAN Host Attach Kit Installation and Setup Guide</i> that came with your SAN Host Attach Kit. These commands have no effect on disk space.																		

Step	Action																		
6	From the host, write data to the file system (the LUN on the filer). This has no effect on disk space.																		
7	Take a snapshot named <i>snap1</i> of the active file system, write 1 GB of data to it, 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 first snapshot reserves enough space to overwrite every block of data in the active file system, so you see 12 GB of used space, the 6-GB LUN (which has 1 GB of data written to it), and one snapshot. Notice that 6 GB appears in the reserved column to ensure writes to the LUN do not fail. If you disable space reservation, this space is returned to available space. <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>12601072</td><td>49808836</td><td>6300536</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>180</td><td></td><td>0</td><td>0 /vol/volspace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	12601072	49808836	6300536	/vol/volspace/	/vol/volspace/.snapshot	0	180		0	0 /vol/volspace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	12601072	49808836	6300536	/vol/volspace/														
/vol/volspace/.snapshot	0	180		0	0 /vol/volspace/.snapshot														

Step	Action																		
8	From the host, write another 1 GB of data to the LUN. Then, from the filer, observe the effect on disk space by entering the following commands: <pre>toaster> df -r /vol/volSPACE</pre> Result: The following sample output is displayed. The amount of data stored in the active file system does not change. You just overwrote 1 GB of old data with 1 GB of new data. However, the snapshot requires the old data to be retained. Before the write, there was only 1 GB of data, after the write, there is 1 GB of new data and 1 GB of data in a snapshot. Notice that the used space goes up for the snapshot by 1 GB and the available space for the volume goes down by 1 GB. <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/</td><td>62649908</td><td>12601072</td><td>47758748</td><td>0</td><td>/vol/volSPACE/</td></tr><tr><td>/vol/volSPACE/.snapshot</td><td>0</td><td>1050088</td><td>0</td><td>0</td><td>/vol/volSPACE/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volSPACE/	62649908	12601072	47758748	0	/vol/volSPACE/	/vol/volSPACE/.snapshot	0	1050088	0	0	/vol/volSPACE/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volSPACE/	62649908	12601072	47758748	0	/vol/volSPACE/														
/vol/volSPACE/.snapshot	0	1050088	0	0	/vol/volSPACE/.snapshot														
9	Take a snapshot named <i>snap2</i> of the active file system and observe the effect on disk space by entering the following command: 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. _____ <pre>toaster> snap create volSPACE snap2</pre> Result: The following sample output is displayed. Because the first snapshot reserved enough space to overwrite every block, only 44 blocks are used to account for the second snapshot. <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/</td><td>62649908</td><td>12601072</td><td>47758748</td><td>6300536</td><td>/vol/volSPACE/</td></tr><tr><td>/vol/volSPACE/.snapshot</td><td>0</td><td>1050136</td><td>0</td><td>0</td><td>/vol/volSPACE/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volSPACE/	62649908	12601072	47758748	6300536	/vol/volSPACE/	/vol/volSPACE/.snapshot	0	1050136	0	0	/vol/volSPACE/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volSPACE/	62649908	12601072	47758748	6300536	/vol/volSPACE/														
/vol/volSPACE/.snapshot	0	1050136	0	0	/vol/volSPACE/.snapshot														

Step	Action																		
10	From the host, write 2 GB of data to the LUN and observe the effect on disk space by entering the following command: <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The second write requires the amount of space actually used if it overwrites data in a snapshot. <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/</td><td>62649908</td><td>12601072</td><td>4608427</td><td>6300536</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>3150371</td><td>0</td><td>0</td><td>/vol/volspace/.s napshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	12601072	4608427	6300536	/vol/volspace/	/vol/volspace/.snapshot	0	3150371	0	0	/vol/volspace/.s napshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	12601072	4608427	6300536	/vol/volspace/														
/vol/volspace/.snapshot	0	3150371	0	0	/vol/volspace/.s napshot														

Monitoring disk space using backing store LUNs and LUN clones

The following example illustrates how to monitor disk space on a volume when taking snapshots, using backing store LUNs and LUN clones. Assume that you start with a new volume and the LUN requires 3 GB of disk space. The information from [“Guidelines for calculating the size of a volume”](#) on page 21 recommends that the volume size be approximately 4*3 GB plus delta. Assuming that the amount of change is high, delta would be 20%, so using two 72 GB disks should still provide more than enough space.

To work through the example, complete the following steps.

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>1440</td><td>50118488</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>12529980</td><td>0</td><td>12529980</td><td>0</td><td>/vol/vol space</td></tr><tr><td>/.snapshot</td><td></td><td></td><td></td><td></td><td></td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace	50119928	1440	50118488	0	/vol/volspace/	/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/vol space	/.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/																				
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/vol space																				
/.snapshot																									

Step	Action																		
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><td>Filesystem</td><td>kbytes</td><td>used</td><td>avail</td><td>reserved</td><td>Mounted on</td></tr><tr><td>/vol/volspace/</td><td>62649908</td><td>1440</td><td>62648468</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0 /vol/volspace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/	/vol/volspace/.snapshot	0	0	0	0	0 /vol/volspace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/														
/vol/volspace/.snapshot	0	0	0	0	0 /vol/volspace/.snapshot														
3	Create a LUN (/vol/volspace/lun0) and observe the effect on disk space by entering the following commands: <pre>toaster> lun create -s 12g -t solaris /vol/volspace/lun0</pre> <pre>toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. Approximately 12 GB 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/</td><td>62649908</td><td>0</td><td>0</td><td>0</td><td>/vol/volspace/</td></tr><tr><td>/vol/volspace/.snapshot</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0 /vol/volspace/.snapshot</td></tr></table>	Filesystem	kbytes	used	avail	reserved	Mounted on	/vol/volspace/	62649908	0	0	0	/vol/volspace/	/vol/volspace/.snapshot	0	0	0	0	0 /vol/volspace/.snapshot
Filesystem	kbytes	used	avail	reserved	Mounted on														
/vol/volspace/	62649908	0	0	0	/vol/volspace/														
/vol/volspace/.snapshot	0	0	0	0	0 /vol/volspace/.snapshot														
4	Create an igroup named <i>sol-igroup1</i> and map the LUN to the igroup by entering the following commands. 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 about on these procedures, see the <i>SAN Host Attach Kit Installation and Setup Guide</i> that came with your SAN Host Attach Kit. These commands have no effect on disk space.																		
6	From the host, write data to the file system (the LUN on the filer). This has no effect on disk space.																		

Step	Action
7	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>
8	Create a backing store 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>
9	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>

Step	Action
10	Create a clone of the backing store LUN you created in Step 8 and observe the effect on disk space by entering the following commands: <pre>toaster> lun clone start /vol/volspace/snaplun</pre> <pre>toaster> df -r /vol/volspace</pre>

About this chapter This chapter explains how to manage igroups and initiator requests.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Managing igroups](#)” on page 72
- ◆ “[Managing initiator requests](#)” on page 75

Managing igroups

Tasks to manage igroups	You can use the command-line interface or FilerView to ◆ Create igroups ◆ Destroy igroups ◆ Add initiators (through their WWPNs) to igroups ◆ Remove initiators (through their WWPNs) from igroups ◆ Display all the initiators in an igroup ◆ Set the operating system type (ostype) for an igroup
-------------------------	--

Creating an igroup To create an igroup, complete the following step.

Step	Action
1	Enter the following command: <pre>igroup create -f -i [-t <i>ostype</i>] <i>initiator_group</i> [<i>node_name</i>...]</pre> -f indicates that it is an FCP igroup. -i indicates that it is an iSCSI initiator. -t <i>ostype</i> indicates the operating system of the host. The values are <i>default</i>, <i>solaris</i>, <i>windows</i>, <i>hpux</i>, <i>aix</i>, or <i>linux</i>. Use <i>default</i> if you are using another operating system. <i>initiator_group</i> is the name of the igroup you specify. <i>node_name</i> is 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 igroup called solaris-group0 that contains the WWPN of the Solaris host associated with that WWPN.

Destroying an igroup

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

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

Step	Action
1	Enter the following command: igroup add <i>igroup WPN</i> Caution_____ When adding initiators to an igroup, 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 igroup solaris-group2.

Removing an initiator

To remove an initiator from an igroup, 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 igroup, 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 igroup, 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>

Managing initiator requests

Why you need to manage initiator requests

Each physical port on the target HBA in the filer has a fixed number of command blocks for incoming initiator requests. When initiators send large numbers of requests, they can monopolize the command blocks and prevent other initiators from accessing the command blocks at that port.

With an igroup throttle, you can perform the following tasks:

- ◆ Limit the number of concurrent I/O requests an initiator can send to the filer
- ◆ Prevent initiators from flooding a port and preventing other initiators from accessing a LUN
- ◆ Ensure that specific initiators have guaranteed access to the queue resources.

How Data ONTAP implements initiator request management

When you use igroup throttles, Data ONTAP calculates the total amount of command blocks available and allocates the appropriate number to reserve for an igroup based on the percentage you specify when you create a throttle for that igroup. Data ONTAP does not allow you to reserve more than 99% of all the resources. The remaining command blocks are unreserved and are available for use by igroups without throttles.

How to manage initiator requests

You use igroup throttles to specify what percentage of the queue resources they can reserve for their use. For example, if you set an igroup's throttle to be 20%, 20% of the queue resources available at the filer's ports are reserved for the initiators in that igroup. The remaining 79% of the queue resources is unreserved. In another example, if you have four hosts and they are in separate igroups, you might set the igroup throttle of the most critical host at 30%, the least critical at 10%, and the remaining two at 20%, leaving 19% of the resources unreserved.

How to use igroup throttles

When you create igroup throttles, you can use them to ensure that critical initiators are guaranteed access to the queue resources and that less critical initiators are not flooding the queue resources. You can perform the following tasks.

- ◆ Create one igroup throttle per igroup (if desired; it is not required)

Note

Igroups without a throttle share all the unreserved queue resources.

- ◆ Assign a specific percentage of the queue resources on each physical port to the igroup
- ◆ Reserve a minimum percentage of queue resources for a specific igroup
- ◆ Restrict an igroup to a maximum percentage of use
- ◆ Allow an igroup throttle to exceed its limit by borrowing resources from
 - ❖ The pool of unreserved resources to handle unexpected I/O requests
 - ❖ The pool of unused reserved resources, if those resources belong to an igroup throttle whose `igroup throttle throttle_borrow` setting is set to y

Creating up an igroup throttle

To create an igroup throttle, complete the following step.

Step	Action
1	Enter the following command: igroup set <i>igroup_name</i> throttle_reserve <i>percentage</i> Example: <code>igroup set solaris-igroup1 throttle_reserve 20</code> Result: The igroup throttle is created for solaris-igroup1, and it will persist through reboots.

Destroying an igroup throttle

To destroy an igroup throttle, complete the following step.

Step	Action
1	Enter the following command: igroup set <i>igroup_name</i> throttle_reserve 0

Allowing or disallowing an igroup to borrow resources

To allow or disallow an igroup to borrow queue resources from the unreserved pool, complete the following step with the appropriate option (yes or no). The default when you create an igroup throttle is no.

Step	Action
1	Enter the following command: <pre>igroup set <i>igroup_name</i> throttle_borrow {yes no }</pre> Example: <code>igroup set solaris-igroup1 throttle_borrow y</code> Result: When you set the <code>throttle_borrow</code> setting to y, the percentage of queue resources used by the initiators in the igroup might be exceeded if resources are available.

Displaying throttle information

To display information about the throttles assigned to igroups, complete the following step. The *exceeds* column displays the number of times the initiator sends more requests than the throttle allows. The *borrow*s column displays the number of times the throttle is exceeded and the filer uses queue resources from the unreserved pool. In the *borrow*s column, “N/A” indicates that the igroup `throttle_borrow` option is set to no.

Step	Action												
1	Enter the following command: igroup show -t Result: <table><tr><td>name</td><td>reserved</td><td>exceeds</td><td>borrow</td></tr><tr><td>solaris-igroup1</td><td>20%</td><td>0</td><td>N/A</td></tr><tr><td>solaris-igroup2</td><td>10%</td><td>0</td><td>0</td></tr></table>	name	reserved	exceeds	borrow	solaris-igroup1	20%	0	N/A	solaris-igroup2	10%	0	0
name	reserved	exceeds	borrow										
solaris-igroup1	20%	0	N/A										
solaris-igroup2	10%	0	0										

Displaying igroup throttle usage

To display real-time information about how many command blocks the initiator in the igroup is using (the first number under the port name) and the number of command blocks reserved for the igroup on the specified port (the second number under the port name), complete the following step.

Step	Action																								
1	Enter the following command: <pre>igroup show -t -i <i>number_of_seconds</i> [-a]</pre> -a displays all throttles, including ones that are not using any command blocks. Example: <code>igroup show -t -i 1</code> Result: The following is a sample display: <table><tr><td>name</td><td>reserved</td><td>4a</td><td>4b</td><td>5a</td><td>5b</td></tr><tr><td>igroup1</td><td>20%</td><td>45/98</td><td>0/98</td><td>0/98</td><td>0/98</td></tr><tr><td>igroup2</td><td>10%</td><td>0/49</td><td>0/49</td><td>17/49</td><td>0/49</td></tr><tr><td>unreserved</td><td></td><td>87/344</td><td>0/344</td><td>112/344</td><td>0/344</td></tr></table> In this example, the display indicates that igroup1 is using 45 of the 98 reserved command blocks on adapter 4a, and igroup2 is using 17 of the 49 reserved command blocks on adapter 5a. Igroups without throttles are counted as unreserved.	name	reserved	4a	4b	5a	5b	igroup1	20%	45/98	0/98	0/98	0/98	igroup2	10%	0/49	0/49	17/49	0/49	unreserved		87/344	0/344	112/344	0/344
name	reserved	4a	4b	5a	5b																				
igroup1	20%	45/98	0/98	0/98	0/98																				
igroup2	10%	0/49	0/49	17/49	0/49																				
unreserved		87/344	0/344	112/344	0/344																				

Displaying LUN statistics on exceeding throttles

To display statistics about I/O requests for LUNs that exceed the igroup throttle, complete the following steps.

Step	Action
1	Display the QFULL column with LUN statistics, which indicates the number of initiator requests that exceeded the number allowed by the igroup throttle, and, as a result, received the QFULL message, by entering the following command: <pre>lun stats -o -i <i>time_in_seconds</i></pre> Example: <code>lun stats -o -i 1</code>

Step	Action
2	Display the total count of QFULLS sent for each LUN by entering the following command: lun stats -o

How a cluster failover affects igroup throttles

Igroup throttles manage physical ports, so during a cluster failover, their behavior varies according to the FCP cfmodes that are in effect, as shown in the following table.

FCP cfmodes	How igroup throttles behave when failover occurs
Standby	Throttles apply to the A ports: ◆ A ports have local throttles ◆ B ports have partner throttles
Partner	Throttles apply to the appropriate ports: ◆ A ports have local throttles ◆ B ports have partner throttles
Mixed	Throttles apply to all ports, and are divided by two

Displaying igroup throttle usage after takeover

To display information about how many command blocks the initiator in the igroup is using and the number of command blocks reserved for the igroup on the specified port after a takeover occurs, complete the following step.

Step	Action												
1	Enter the following command: igroup show -t Example: igroup show -t Result: The following is a sample display: <table><thead><tr><th>name</th><th>reserved</th><th>exceeds</th><th>borrows</th></tr></thead><tbody><tr><td>solaris-igroup1</td><td>20%</td><td>0</td><td>N/A (Reduced by takeover to 10%)</td></tr><tr><td>solaris-igroup2</td><td>10%</td><td>0</td><td>0 (Reduced by takeover to 5%)</td></tr></tbody></table>	name	reserved	exceeds	borrows	solaris-igroup1	20%	0	N/A (Reduced by takeover to 10%)	solaris-igroup2	10%	0	0 (Reduced by takeover to 5%)
name	reserved	exceeds	borrows										
solaris-igroup1	20%	0	N/A (Reduced by takeover to 10%)										
solaris-igroup2	10%	0	0 (Reduced by takeover to 5%)										

About this chapter

This chapter provides information about how to use Data ONTAP data protection features using the SCSI protocol in an FCP Network.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Data ONTAP protection methods”](#) on page 82
- ◆ [“Using snapshots”](#) on page 84
- ◆ [“Using LUN clones”](#) on page 86
- ◆ [“Deleting busy snapshots”](#) on page 88
- ◆ [“Using SnapRestore”](#) on page 94
- ◆ [“Backing up data to tape”](#) on page 100
- ◆ [“Using NDMP”](#) on page 104
- ◆ [“Using volume copy”](#) on page 105
- ◆ [“Using NVFAIL”](#) on page 106

Data ONTAP protection methods

Data protection methods

Data ONTAP provides a variety of methods for protecting data using FCP. These methods, described in the following table, are based on NetApp's Snapshot™ technology, which enables you to maintain multiple read-only versions of LUNs online per filer volume.

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 LUN or the file system at the time the snapshot was created. All the other data protection methods rely on snapshots or create, use, and destroy snapshots, as required.

For information about NetApp data protection products and solutions, see the *Network Appliance Data Protection Portal* at http://www.netapp.com/solutions/data_protection.html.

Method	Used to...
Snapshot	◆ Take point-in-time copies of a volume.
SnapRestore®	◆ Restore a LUN or file system to an earlier preserved state in less than a minute without rebooting the filer, regardless of the size of the LUN or volume being restored. ◆ Recover from a corrupted database or a damaged application, a file system, a LUN, or a volume by using an existing snapshot.
SnapMirror®	◆ Replicate data or asynchronously mirror data from one filer to another over local or wide area networks (LANs or WANs). ◆ Transfer snapshots taken at specific points in time to other filers or NetApp NearStore® systems. These replication targets can be in the same data center through a LAN or distributed across the globe connected through metropolitan area networks (MANs) or WANs. Because SnapMirror operates at the changed block level instead of transferring entire files or file systems, it generally reduces bandwidth and transfer time requirements for replication.
SnapVault™	◆ Back up data by using snapshots on the filer and transferring them on a scheduled basis to a destination filer or NearStore system. ◆ Store these snapshots on the destination filer for weeks or months, allowing recovery operations to occur nearly instantaneously from the destination filer to the original filer.

Method	Used to...
SnapDrive™	◆ Manage a filer's LUNs that serve as virtual storage devices for application data in Windows 2000 Server and Windows 2003 Server environments in an integrated environment with the Windows Volume Manager. ◆ Perform online storage configuration, LUN expansion, and streamlined management. Note_____ For more information about SnapDrive, see the <i>SnapDrive Installation and Administration Guide</i>.
Native tape backup and recovery	◆ Store and retrieve data on tape. Note_____ Data ONTAP supports native tape backup and recovery from local, Gigabit Ethernet, and Fibre Channel SAN-attached tape devices. Support for most existing tape drives is included, as well as a method for tape vendors to dynamically add support for new devices. In addition, Data ONTAP supports the Remote Magnetic Tape (RMT) protocol, allowing backup and recovery to any capable system. Backup images are written using a derivative of the BSD dump stream format, allowing full file-system backups as well as nine levels of differential backups.
NDMP	◆ Control native backup and recovery facilities in NetApp filers and other file servers. Backup application vendors provide a common interface between backup applications and file servers. Note_____ The Network Data Management Protocol (NDMP) is an open standard for centralized control of enterprise-wide data management. For more information about how NDMP-based topologies can be used by filers to protect data, see the <i>Data Protection Solutions Overview</i>, Technical Report TR3131 at http://www.netapp.com/tech_library/3131.html.

Using snapshots

How Data ONTAP snapshots work in an FC network

Data ONTAP cannot take snapshots of some applications that are implemented on top of file systems, such as NTFS. If you want to take a snapshot of these types of applications, you must first ensure that the files are closed and cannot be modified, and that the application is quiesced, or taken offline, so that the file system caches are committed before the snapshot is taken. The snapshot takes less than one second to complete, at which time the application can resume normal operation.

If the application requires a lot of time to quiesce, this can result in some amount of unavailable time. To avoid this scenario, some applications have a built-in *hot backup* mode. This allows a snapshot or a backup to occur while the application operates in a degraded mode, with limited performance.

Data ONTAP cannot take snapshots of applications that have the ability to work with raw device partitions. NetApp recommends that you use specialized modules from a backup software vendor tailored for such applications.

If you want to back up raw partitions, it is best to use the hot backup mode for the duration of the backup operation. For more information about backup and recovery of databases using NetApp SAN configurations, see the appropriate Technical Report for the database at http://www.netapp.com/tech_library.

How snapshots are used in the SAN environment

Data ONTAP cannot ensure that the data within a LUN is in a consistent state with regard to the application accessing the data inside the LUN. Therefore, 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. For example, you can use batch files and scripts on a host that has administrative access to the filer. You use these scripts to perform the following tasks:

- ◆ Make the data within the LUN consistent with the application, possibly by quiescing a database, placing the application in hot backup mode, or taking the application offline
- ◆ Use the `rsh` or `ssh` command to create the snapshot on the filer (this only takes a few seconds, regardless of volume size or use)
- ◆ Return the application to normal operation

Note

On Windows hosts, you can use the Windows Task Scheduler service to execute this script at specified intervals.
In addition, 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*.

The relationship between a LUN and a snapshot

When you take a snapshot of a LUN, it is initially backed by data in the snapshot. After the snapshot is taken, writes to the LUN reside in the active file system. The backing snapshot cannot be deleted without destroying this LUN.

What you can do with a snapshot

After you have a snapshot, you can use it to create a LUN for temporary use as a prototype for testing data or scripts in applications or databases by creating a writable snapshot of a LUN (a backing store LUN) with the `lun create -b` command. Because the LUN is backed by the snapshot, you cannot delete the snapshot until the backing store LUN is deleted.

If you want to retain the LUN on a permanent basis, or want to delete the snapshot that backs the backing store LUN, you can clone the backing store LUN and then delete the snapshot. LUN clones require the same amount of space as the original LUN, whereas backing store LUNs require the amount of space of the original LUN plus the delta in the snapshot.

If you want to restore the LUN from a snapshot, you can use SnapRestore, but it will not have any updates to the data since the snapshot was taken.

What snapshots require

In Data ONTAP 6.5, LUNs are space reserved by default. This means that enough space is reserved so that writes to the LUNs are guaranteed. The more space that is reserved, the less free space is available. If free space within the volume is below a certain threshold, snapshots cannot be taken. For information about how to manage available space, see “[Monitoring disk space](#)” on page 56.

Using LUN clones

What a LUN clone is A LUN clone is a copy of a writable snapshot of a LUN. After the clone process is completed, the LUN resides entirely in the active file system and does not depend on any snapshots.

What LUN clones are used for You can use LUN clones to create a permanent copy of the LUN at the time the snapshot of the original LUN was taken. You can use it when you need 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.

Creating a writable snapshot of a LUN Before you can clone a LUN, you must create a writable snapshot of a 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

Step	Action
3	Create a writable snapshot of a LUN by entering the following command: lun create -b <i>snapshot_lunpathname lun_clone_path</i> Example: <code>lun create -b /vol/vol1/.snapshot/mysnap/lun0 /vol/vol1/snaplun0</code>

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 show <i>lun_clone_path</i>

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>lunpath</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. 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.

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 (a backing store LUN). 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 that 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 (backing store 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 <i>vol-name</i> Example: snap list vol2 Result: The following message is displayed: Volume vol2 working... <table><tr><th>%/used</th><th>%/total</th><th>date</th><th>name</th></tr><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></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 <i>vol_name snap_name</i> 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 uses a snapshot to revert an entire volume or a LUN to the state it was in when the snapshot was taken, within seconds, regardless of size or utilization. You can use SnapRestore to restore an entire volume or a LUN, or you can perform a single file SnapRestore on a LUN.

Requirements for using SnapRestore

Before using SnapRestore, you must perform the following tasks:

- ◆ 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 (otherwise, it reboots the filer).

Caution

When restoring a single LUN, it must be taken offline or be unmapped prior to recovery if the LUN has not been destroyed.

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.

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.
3	From the filer, unmap the LUN by entering the following command: <code>lun unmap lun-path initiator-group</code>

Step	Action
4	Enter the following command: <pre>snap restore [-f] [-t vol] volume_name [-s snapshot_name]</pre> -f suppresses the warning message and the prompt for confirmation. This option is useful for scripts. -t vol <i>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 press 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) by entering the following command: <code>lun unmap lun-path initiator-group</code>	
8	Remap the LUN by entering the following command: <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 use 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.

Backing up data to tape

Structure of SAN backups


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

Note

NetApp strongly recommends 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.

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.

The following diagram shows a SAN setup that uses two applications hosts and a clustered pair of filers.



Volumes on the FCP host can consist of a single LUN mapped from the filer or multiple LUNs using a volume manager, such as VxVM on HP-UX.

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.

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.

Note

Steps 4, 5, and 6 can be part of your SAN backup application's pre-processing script.

Steps 9 and 10 can be part of your SAN backup application's post-processing script.

Step	Action
1	Enter the following command to create an igroup: <pre>igroup create {-f -i} [-t ostype] group [node ...]</pre> Example: <code>igroup create -f -t windows payroll_server iqn.1991-05.com.microsoft:eng1</code> Result: Data ONTAP creates an igroup with the FCP option, the operating system Windows, the name payroll, and the node name 10:00:00:00:c3:4a:0e:e1.
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 windows /vol/vol1/qtree_1/payroll_lun</code> Result: Data ONTAP creates a LUN with a size of 48 GB, of the type Windows, and with the name and path /vol/vol1/qtree_1/payroll_lun.

Step	Action
3	Enter the following command to map the LUN to the igroup with a SCSI ID: <pre>lun map lun-path initiator-group scsi-lun</pre> Example: <code>lun map /vol/vol1/qtrees_1/payroll_lun payroll_server 1</code> Result: Data ONTAP maps the /vol/vol_name/qtrees_1/payroll_lun LUN to the payroll_server igroup with a LUN ID of 1.
4	Save the contents of host file system buffers to disk using the command provided by your host operating system, or by using SnapDrive 2.0 or later for Windows systems.
5	Create a snapshot by using SnapDrive 2.0 for Windows systems, or enter the following command for non-Windows systems: <pre>snap create volume_name snapshot_name</pre> Example: <code>snap create vol1 payroll_backup</code>
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: <pre>lun create -b snapshot lun-path lun-path</pre> Example: <code>lun create -b vol/vol1/.snapshot/payroll_backup/qtrees_1/payroll_lun /vol/vol1/qtrees_1/payroll_backup_lun</code>
7	Enter the following command to map the new LUN you created to the backup host: (<i>backup_server</i> is an igroup that already exists for the backup host.) <pre>lun map lun-path initiator-group scsi-lun</pre> Example: <code>lun map /vol/vol_name/qtrees_1/payroll_lun payroll_server 1</code> Result: Data ONTAP maps the LUN /vol/vol_name/qtrees_1/payroll_lun to the igroup <i>payroll_server</i> with a SCSI ID of 1.

Step	Action
8	Back up the data in the new LUN from the backup host to tape using your SAN backup application.
9	Take the backup LUN offline by entering the following command: lun offline /vol/vol_name/qtree_name/lun_name Example: lun offline /vol/vol1/qtree1/payroll_backup_lun
10	Remove the snapshot by entering the following command: snap delete volume_name lun_name Example: snap delete vol1 payroll_backup

Using NDMP

When to use native or NDMP backup

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. For information about these features, see the *Data ONTAP 6.5 Data Protection Online Backup and Recovery* guide.

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.

When to use the ndmpcopy command

You can use the `ndmpcopy` command to copy a directory, qtree, or volume that contains a LUN. For information about how to use the `ndmpcopy` command, see the *Data ONTAP 6.5 Data Protection Online Backup and Recovery* guide.

Using volume copy

Command to use You can use the `volume copy` command to copy LUNs; however, this requires that applications accessing the LUNs are quiesced and offline prior to the copy operation.

The `vol copy 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.

Copying a volume To copy a 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: <pre>vol copy start -S source:source_volume dest:dest_volume</pre> 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: <code>vol copy start -S /vol/vol0 filerB:/vol/vol1</code>

Using NVFAIL

How NVFAIL works with LUNs

If an NVRAM failure occurs on a volume, Data ONTAP detects the failure at boot up time. If you enabled the `vol options nvfail` option for the volume and it contains the LUNs, Data ONTAP performs the following actions:

- ◆ Offlines the LUNs in the volumes that had the NVRAM failure
- ◆ Stops exporting LUNs over FCP
- ◆ Sends error messages to the console stating that Data ONTAP took the LUNs offline or that NFS file handles are stale (this is also useful if the LUN is accessed over NAS protocols.)

Caution

NVRAM failure can lead to possible data inconsistencies.

How you can provide additional protection for databases

In addition, you can protect specific LUNs, such as database LUNs, by creating a file called `/etc/nvfail_rename` and adding their names to the file. In this case, if NVRAM failures occur, Data ONTAP renames the LUNs specified in `/etc/nvfail_rename` file by appending the extension `.nvfail` to the name of the LUNs. When Data ONTAP renames a LUN, the database cannot start automatically. As a result, you must perform the following actions:

- ◆ Examine the LUNs for any data inconsistencies and resolve them
- ◆ Remove the `.nvfail` extension with the `lun move` command (for information about this command, see [“Renaming a LUN”](#) on page 47)

How you make the LUNs accessible to the host after an NVRAM failure

To make the LUNs accessible to the host or the application after an NVRAM failure, you must perform the following actions:

- ◆ Ensure that the LUNs data is consistent
- ◆ Bring the LUNs online
- ◆ Export each LUN manually to the initiator

For information about NVRAM, see the *Data ONTAP 6.5 Data Protection Online Backup and Recovery* guide.

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.

About this chapter This chapter provides an overview of how to manage adapters, initiators, igroups, and traffic in an NetApp FC environment.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Managing the FCP service](#)” on page 110
- ◆ “[Displaying information about HBAs](#)” on page 113
- ◆ “[Managing FCP adapters](#)” on page 122

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 also use FilerView and go to

- ◆ LUNs > FCP to manage FCP adapters and view FCP statistics
- ◆ Filer > Manage Licenses to manage the FCP 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.

For FAS270 appliances: After you license the FCP service on an FAS270 appliance, you must reboot. When the appliance boots up, the port labeled Fibre Channel 2 is in SAN target mode. When you enter Data ONTAP commands that display adapter statistics, this port is slot 0, so the virtual ports are shown as 0a_0, 0a_1, and 0a_2.

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 service 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 HBAs

How to display HBA information

You can display information and manage HBAs 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).

The display will vary, depending on the FCP `cfmode` setting, and the filer model. The examples of output generated by commands in this chapter assume you are using an FAS960 in standby mode.

If you want to display...	Use this command...
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>
All adapters (including HBAs, NICs, and switch port) configuration and status 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.
Disks, disk loops, and options configuration information that affect coredumps and takeover	<code>sysconfig -c</code>
FCP <code>cfmode</code> setting	<code>fcv show cfmode</code>
FCP Traffic information	<code>sysstat -f</code>
How long FCP has been running	<code>uptime</code>
Initiator HBA port address, port name, node name, and igroup name connected to target HBAs	<code>fcv show initiator [-v] [<i>adapter&portnumber</i>]</code> <code>-v</code> 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.
Service statistics	<code>availtime</code>

If you want to display...	Use this command...
Target HBAs configuration information	<code>fcv config</code>
Target HBAs node name, port name, and link state	<code>fcv show adapter [-p] [-v][<i>adapter&portnumber</i>]</code> -p displays information about adapters running on behalf of the partner node (filer). -v displays additional information about target adapters. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
Target HBAs statistics	<code>fcv stats [-z] [<i>adapter&portnumber</i>]</code> -z zeros the statistics. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
Virtual target interconnect (vtic) statistics	<code>fcv stats vtic</code>
WWNN (node name) of the target HBA	<code>fcv nodename</code>

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, in 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>fcg 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 node name, port name, and link state of all target HBAs, complete the following step. Notice that the port name and node name are displayed with the separating colons, and without. For Solaris hosts, you use the WWPN without

separating colons when you map adapter port names (or these target WWPNs) to the host.

Step	Action
1	At the filer, enter the following command: <pre>fcp show adapter</pre> Result: The following information is displayed for all target HBAs. This example displays information for the HBA in slot 4: <pre>slot: 4a Description Fibre Channel Target Adapter 4a (Dual-channel, QLogic 2312 (2342) rev.2) Adapter Type: Local 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 (2342) rev.2) Adapter Type: Local 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</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. _____

Displaying initiator HBA information

To display the port name and group name (of the igroup) of initiator HBAs (host) connected to target HBAs, complete the following step.

Step	Action
1	At the filer, enter the following command: <pre>fcpx show initiator</pre> Result: The following output is displayed: <pre>Initiators connected on adapter 4a: 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</pre> Adapter 4b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode. <pre>Initiators connected on adapter 5a: 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</pre> Adapter 5b is running on behalf of the partner. Use the partner command to view this adapter when in takeover mode.

Displaying statistics

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

Step	Action																																												
1	Enter the following command: <pre>fcg 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>

Displaying vtic statistics in partner mode

When the FCP cfmode setting is in partner mode, a virtual target interconnect (vtic) adapter handles the initiator requests. You use the `fcv stats vtic` command to display vtic adapter statistics. You can use the `igroup show` command to verify that vtic appears in the output, which indicates that the initiator is logged in to the partner port.

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: <pre>Fibre Channel nodename: 50:a9:80:00:02:00:8d:b2 (50a9800002008db2)</pre>

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 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 connecting the host and filer to a Fibre Channel switch Note _____Some switches support both ptp and loop. For more specific information about which media type to use, see the vendor-supplied documentation for your 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 if the partner filer fails, complete the following step.

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

Glossary

client	A computer that shares files on a filer. See also host.
FCP	Fibre Channel Protocol. A licensed service on the filer that enables you to export LUNs to hosts using the SCSI protocol over a Fibre Channel fabric.
filer	A NetApp storage appliance, including the R100, FAS200, FAS800, and FAS900 series models.
HBA	Host bus adapter. An I/O adapter that connects a host I/O bus to a computer's memory system in SCSI environments. The HBA might be an FCP adapter or an iSCSI adapter.
host	Any computer system that accesses data on a filer as blocks using the iSCSI or FCP protocols, or is used to administer a filer.
igroup	Initiator group. A collection of unique identifiers, either iSCSI node names of initiators (hosts) in an IP network, or FCP WWPNs in a SCSI network, that are given access to LUNs when they are mapped to those LUNs.
initiator	The system component that originates an I/O command over an I/O bus or network; a host that has iSCSI initiator software installed on it, or a host that has a host bus adapter (HBA) installed in it, which is connected to the iSCSI or FCP network with the appropriate license enabled.
initiator group	See igroup.
iSCSI	A licensed service on the filer that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

iSCSI node name	A logical name to identify an iSCSI node, with the format <code>iqn.yyyy-mm.backward_naming_authority:sn.unique_device_name</code>. <code>yyyy-mm</code> is the month and year in which the naming authority acquired the domain name. <code>backward_naming_authority</code> is the reverse domain name of the entity responsible for naming this device. An example reverse domain name is <code>com.netapp</code>. <code>unique_device_name</code> is a free-format unique name for this device assigned by the naming authority, preceded by <code>sn</code>. Typically, the <code>unique_device_name</code> is a serial number.									
iswt	The iSCSI Software Target driver that provides the ability for filers to connect to initiators using the iSCSI protocol over TCP/IP. In a filer cluster, <code>iswta</code> refers to the local filer and <code>iswtb</code> refers to the partner. If the filer is not clustered, <code>iswta</code> refers to the single filer.									
LUN	A logical unit of storage.									
LUN clone	A complete copy of a LUN, which was initially created to be backed by a LUN or a file in a snapshot. The clone creates a complete copy of the LUN and frees the snapshot, which you can then delete.									
LUN ID	The numerical identifier that the filer exports for a given LUN. The LUN ID is mapped to an <code>igroup</code> to enable host access.									
LUN path	The path to a LUN on the filer. The following example shows a LUN path: <table><tr><td>LUN path</td><td>Mapped to</td><td>LUN ID</td></tr><tr><td>-----</td><td>-----</td><td>-----</td></tr><tr><td><code>/vol/vol01/fcpdb.lun</code></td><td><code>igroup_1</code></td><td><code>6</code></td></tr></table>	LUN path	Mapped to	LUN ID	-----	-----	-----	<code>/vol/vol01/fcpdb.lun</code>	<code>igroup_1</code>	<code>6</code>
LUN path	Mapped to	LUN ID								
-----	-----	-----								
<code>/vol/vol01/fcpdb.lun</code>	<code>igroup_1</code>	<code>6</code>								
LUN serial number	The unique serial number for a LUN, as defined by the filer.									

map	To create an association between a LUN and an igroup. A LUN mapped to an igroup is exported to the nodes in the igroup (iqn or eui) when the LUN is online. LUN maps are used to secure access relationships between LUNs and the host.
online	Signifies that a LUN is exported to its mapped igroups. A LUN can be online only if it is enabled for read/write access.
offline	Disables the export of the LUN to its mapped igroups. The LUN is not available to hosts.
qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes. Qtrees can be used to group LUNs.
SAN	Storage Area Network. A storage network composed of one or more filers connected to one or more hosts in either a direct-attached or network-attached configuration using the iSCSI protocol over TCP/IP or the SCSI protocol over FCP.
share	An entity that allows the LUN's data to be accessible through multiple file protocols such as NFS and iSCSI. You can share a LUN for read or write access, or all permissions.
space reservations	An option that determines whether disk space is reserved for a specified LUN or file remains available for writes to any LUNs, files, or snapshots. Required for guaranteed space availability for a given LUN with or without snapshots.
storage appliance	Hardware and software-based storage systems, such as filers, that serve and protect data using protocols for both SAN and NAS networks.
target	The system component that receives a SCSI I/O command. A filer with the iSCSI or FCP license enabled and serving the data requested by the initiator.

volume	A file system. Volume refers to a functional unit of filer storage, based on one or more RAID groups, that is made available to the host. LUNs are stored in volumes.
WWN	World Wide Number. A unique 48- or 64-bit number assigned by a recognized naming authority (often through block assignment to a manufacturer) that identifies a connection for an FCP node to the storage network. A WWN is assigned for the life of a connection (device).
WWNN	World Wide Node Name. A unique 64-bit address represented in the following format: nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value. NetApp assigns a WWNN to a filer based on the serial number of its NVRAM. The WWNN is stored on disk. Data ONTAP refers to this number as a <i>Fibre Channel Nodename</i> , or simply, a node name.
WWPN	World Wide Port Name. A unique 64-bit address represented in the following format: nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value. Each Fibre Channel device has one or more ports that are used to connect to a SCSI network. Each port has a unique WWPN, which Data ONTAP refers to as an <i>FC Portname</i> , or simply, a port name.

Index

Symbols

/etc/nvfail_rename, database protection 106

A

adapters
 displaying information about 113

B

backing store LUN 21
backup
 data to tape 100
 single LUNs to tape 101
 tape, when to use 104

C

clustered configuration
 cfmode settings 7
clustered configurations
 options required 8

D

database 106
database protection
 using /etc/nvfail_rename 106
 using vol options nvfail 106
dual_fabric mode 9

F

FAS270
 ports, how displayed 9
FAS270, dual-fabric mode 9
FAS270, switch requirement 9
FCP
 cfmode setting 7
 licensed service 5
 nodes defined 5
 nodes, filer 6

nodes, host 6
nodes, how connected 5
nodes, how identified 6
nodes, switch 6

FCP cfmode settings
 affect on target ports 8
FCP commands
 fcv show initiator 40
FCP service
 displaying how long running 120
 displaying traffic information about 120
filer administration 2
 using FilerView 3
 using the command line 2
filer, defined as target 2

H

host bus adapters
 displaying information about 120
 initiator, displaying information about 118

I

igroup commands
 igroup add 73
 igroup create 41, 72
 igroup destroy 73
 igroup remove 74
 igroup set 74
 igroup show 74
initiator group, defined 7
initiator groups
 adding 73
 creating 72
 destroying 73
 displaying contents of 74
 removing 74
 setting the operating system type 74
 unmapping LUNs from 47
initiator host bus adapters, displaying information about 118

L

lun commands

- lun online 46

- lun serial 49

- lun unmap 47

LUNs

- accessing with NAS protocols 50

- backing store 21

- bringing online 46

- changing serial number 49

- creating volumes for 23

- defined 5

- displaying reads, writes, and operations for 53

- resizing restrictions 47

- serial number 5

- unmapping from initiator group 47

M

man page command 3

mixed mode 8

N

nodenames, of initiator host bus adapters,

displaying 118

nvfail option, of vol options command 106

P

partner mode 8

port resources, managing 9

portnames of initiator adapters, displaying 118

ports

- used in clustered configurations 7

ports, used in non-clustered configurations 7

Q

qtrees, creating for LUNs 24

R

restoring snapshots of LUNs 94

S

serial number, of LUNs, changing 49

Single File SnapRestore, using with LUNs 97

snap reserve, setting the percentage 16

snapshot schedule, turning off at the command line 17

snapshots, using with SnapRestore 94

standby mode 8

V

vol option nvfail, using with LUNs 106

volumes

- creating for LUNs 23

W

WWPN

- creating igroups with 7

- identifying filer ports with 7

WWPNs

- how assigned 6

/etc/nvfail_rename file, for database protection 113