

Data ONTAP™ 7.0.1

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 210-01113-A0
April 2005

Copyright and trademark information

Copyright information

Copyright © 1994–2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document 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.

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:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

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 © [\$date-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/>.

Notice of any changes or modifications to the W3C files, including the date changes were made.

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 or purchase of this product does 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, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. Data ONTAP, gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, vFiler, VFM, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, and SpinServer are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, and SpinStor are trademarks of Spinnaker Networks, LLC in the United States and/or other countries.

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	vii
Chapter 1	How NetApp Implements an FCP Network	1
	Understanding NetApp storage appliances	2
	Understanding how NetApp implements an FC SAN network	5
	How Data ONTAP supports FCP with clustered filers.	8
	Related documents	13
Chapter 2	Configuring Filer Storage.	15
	Understanding filer storage units	16
	Understanding space reservation for volumes and LUNs	18
	Understanding how fractional reserve affects available space	21
	How 100 percent fractional reserve affects available space	22
	How reducing fractional reserve affects available space.	28
	How guarantees on flexible volumes affect fractional reserve	32
	Calculating the size of a volume	34
	Guidelines for creating volumes that contain LUNs	39
	Creating LUNs, igroups, and LUN maps.	45
	Creating LUNs with the lun setup program	52
	Creating LUNs and igroups with FilerView	57
	Creating LUNs and igroups by using individual commands.	61
Chapter 3	Managing LUNs	65
	Managing LUNs and LUN maps	66
	Displaying LUN information.	72
	Reallocating LUN and volume layout	77
	Monitoring disk space	87
Chapter 4	Managing Initiator Groups.	99
	Managing igroups.	100

	Managing initiator requests.	105
Chapter 5	Using Data Protection with FCP	111
	Data ONTAP protection methods	112
	Using snapshots.	114
	Using LUN clones	116
	Deleting busy snapshots	119
	Using SnapRestore	122
	Backing up data to tape	127
	Using NDMP	131
	Using volume copy	132
	Cloning flexible volumes.	133
	Using NVFAIL	138
	Using SnapValidator	140
Chapter 6	Managing the NetApp SAN.	151
	Managing the FCP service	152
	Managing the FCP service on systems with onboard ports	156
	Displaying information about HBAs	160
	Glossary	171
	Index	175

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 Logical Unit Numbers (LUNs), how to create and manage LUNs and initiator groups (igroups), and how to monitor FCP traffic. This guide assumes that you have completed the following tasks to install, set up, and configure your storage appliance:

- ◆ Ensured that your configuration is supported by referring to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.
- ◆ 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 <i>arp -d</i> 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 upper case.

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 FCP 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
- ◆ [“How Data ONTAP supports FCP with clustered filers”](#) on page 8
- ◆ [“Related documents”](#) on page 13

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 Logical Unit Numbers (LUNs). 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 SnapDrive™ software in a Windows environment.

Online command-line information aids: When using the command line, you can get command-line syntax help 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 enter the following command: `man na_lun`.

For more information about filer administration, see the *Data ONTAP 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 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 <code>/na_admin/</code> as the location for the URL. Example: If you have a filer named toaster, enter the following URL in the browser: <code>http://toaster/na_admin</code>. 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 152](#).

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 when they are connected to a Fibre Channel 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.

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. 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 service, using a unique identifier.

How nodes are uniquely identified

Each FCP node is identified by a World Wide Node Name (WWNN) or a World Wide Port Name (WWPN).

How WWPNs are used: WWPNs identify each port on an HBA. WWPNs are used for the following purposes.

- ◆ Creating an initiator group

The WWPNs of the host's HBAs are used to create an initiator group (igroup). An igroup is used to control host access to specific LUNs. You create an igroup by specifying a collection of WWPNs of initiators in an FCP network. When you map a LUN on a filer to an 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 means that the LUNs do not appear as disks on that host. For detailed information about mapping LUNs to igroups, see [“What is required to map a LUN to an igroup”](#) on page 50.

- ◆ Uniquely identifying a filer's HBA target ports

The filer's WWPNs uniquely identify each target port on a filer. The host operating system uses the combination of the WWNN and WWPN to identify filer HBAs and host target IDs. Some operating systems require persistent binding to ensure that the LUN appears at the same target ID on the host.

How filers are identified: 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 `fcg show adapter`, `fcg config`, or `sysconfig -v`, `fcg nodename`, 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.

Note

The target WWPNs might change if you add or remove HBAs on the filer.

Filer system serial numbers: The filer also has a unique system serial number that you can view by using the `sysconfig` command. The system serial number is a unique 7-digit identifier that is assigned by NetApp manufacturing. You cannot modify this serial number. Some multipathing software products use the system serial number together with the LUN serial number to identify a LUN.

How hosts are identified: To know which WWPNs are associated with a specific host, see the SAN Host Attach Kit documentation for your host. These documents describe commands supplied by NetApp or the vendor of the initiator or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, you use the `lputilnt` utility, and for UNIX 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.

How switches are identified: 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 target ports are labeled

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

How to manage target port resources

Each target port has a fixed number of resources, or command blocks, for incoming initiator requests. When all the 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 command blocks used by the initiators in an igroup, or you can reserve a pool of command blocks for the exclusive use of initiators in an igroup. This is known as *igroup throttling*. For information about igroup throttling, see [“Managing initiator requests”](#) on page 105.

How Data ONTAP supports FCP with clustered filers

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 so that they do 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
- ◆ `cf.wafl.delay.enable` to On
- ◆ `cf.takeover.on_panic` to On

About the FCP cfmode setting

If your filers are in a cluster, Data ONTAP provides multiple modes of operation required to support homogeneous and heterogeneous host operating systems. Each target HBA has two ports: Port A and Port B. The FCP cfmode setting controls how the target ports:

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

The FCP cfmode settings must be set to the same value for both nodes in a cluster. You view how these modes are set for your filer by using `fc 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 your Network Appliance Professional Services representative to modify the FCP cfmode setting.

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 a filer cluster of F800 series or FAS900 series filers to Data ONTAP 6.5 or later, the FCP cfmode is set to standby mode. Port A on each target HBA 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 available to the initiator. When one filer fails, Port B

on the partner filer becomes active and provides access to the LUNs on the failed filer. The Port B assumes the WWPN of the Port A on the failed partner.

Standby mode behavior is the only available behavior for Data ONTAP versions 6.3.x and 6.4.x. When you upgrade your filer to Data ONTAP version 6.5 software from these versions, standby is the default cfmode setting. This enables you to upgrade existing SAN configurations with Windows or Solaris hosts and enable these hosts to continue to access the filer after the upgrade without any configuration changes.

Some operating systems, such as HP-UX and AIX, do not support standby mode. For detailed information, see the documentation for your Host Attach Kit.

◆ Partner mode

If you have a filer cluster of F800 series or FAS900 series filers with Data ONTAP 6.5 or later newly installed, the FCP cfmode setting is set to partner mode. In this mode, Port A and Port B are both active. Port A on each HBA provides access to local LUNs, and Port B provides access to LUNs on the partner filer.

Partner mode requires that multipathing software be installed on the host. For information about the multipathing software supported for your host, see the documentation for your SAN Host Attach Kit.

In partner mode, the target ports connect to the fabric in point-to-point mode.

◆ 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 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.

In mixed mode, the target ports connect to the fabric in loop mode. this means that you cannot use mixed mode with switches that do not support public loop.

Mixed mode also requires that multipathing software be installed on the host. For information about the multipathing software supported for your host, see the documentation for your SAN Host Attach Kit.

◆ **Dual_fabric**

This is the only supported mode of operation for FAS270 clusters. You cannot change the `cfmode` from `dual_fabric` to a different setting for the FAS270. The `dual_fabric` mode is not supported for other filer models.

The FAS270 cluster consists of two storage appliances integrated into a DiskShelf14mk2 FC disk shelf. Each storage appliance has two Fibre Channel ports. The orange port labeled Fibre Channel C operates as a Fibre Channel target port after you license the FCP service and reboot the storage appliance. The blue port labeled Fibre Channel B connects to the internal disks, enabling you to connect additional disk shelves to an FAS270 cluster. The Fibre Channel target port 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

Note

For switched configurations, `dual_fabric` mode requires switches that support public loop.

How Data ONTAP displays information about 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. You use the `fcv config` or `fcv show adapter` commands to display information about the target ports.

Standby mode: When the FCP `cfmode` setting is `standby`, the local WWNN and WWPNN have a pattern of `50:a9:80:nn:nn:nn:nn:nn` or `50:0a:09:nn:nn:nn:nn:nn`. Each port has a unique WWPNN. The standby WWNN and WWPNN have a pattern of `20:01:00:nn:nn:nn:nn:nn`.

The following `fcv config` output shows target port information for a filer in `standby` mode. The target HBAs are installed in slots 9 and 11. Port 1 in slot 9 is displayed as 9a. Port 2 in slot 9 is displayed as 9b.

```

filer> fcp config
9a:  ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021b00
      portname 50:a9:80:01:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
      mediatype ptp  partner adapter None

9b:  ONLINE <ADAPTER UP>  PTP  Fabric  Standby
      host address 021a00
      portname 20:01:00:e0:8b:28:71:54  nodename 20:01:00:e0:8b:28:71:54
      mediatype ptp  partner adapter 9a

11a: ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021500
      portname 50:a9:80:03:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
      mediatype ptp  partner adapter None

11b: ONLINE <ADAPTER UP>  PTP  Fabric  Standby
      host address 021600
      portname 20:01:00:e0:8b:28:70:54  nodename 20:01:00:e0:8b:28:70:54
      mediatype ptp  partner adapter 11a

```

Partner mode: : When the FCP cfmode setting is partner, the local and partner addresses of the WWNN and WWPN have a pattern of 50:a9:80:nn:nn:nn:nn. The WWPN and WWNN of the B ports are based on the WWNN of the partner filer in the cluster. For example, port B on the local filer represents the WWNN of its partner. The following `fcp config` command output shows how Data ONTAP displays WWNN and WWPN when the filer's cfmode is set to partner and the cluster is in normal operation.

```

filer> fcp config
9a:  ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021b00
      portname 50:a9:80:01:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
      mediatype ptp  partner adapter 9a

9b:  ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021a00
      portname 50:a9:80:0a:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
      mediatype ptp  partner adapter 9b

11a: ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021500
      portname 50:a9:80:03:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
      mediatype ptp  partner adapter 11a

11b: ONLINE <ADAPTER UP>  PTP  Fabric
      host address 021600
      portname 50:a9:80:0c:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
      mediatype ptp  partner adapter 11b

```

Mixed mode: When the `cfmode` setting is mixed, FCP commands display three virtual ports for each physical port. For example, if a target HBA is installed in slot 9, the `fcp config` command shows the physical ports as 9a and 9b. The virtual ports associated with 9a are 9a_0 (local), 9a_1 (standby), and 9a_2 (partner), and the partner virtual port as 9a_2. The virtual ports associated with 9b are 9b_0 (local), 9b_1 (standby), and 9b_2 (partner).

dual_fabric (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.

Related documents

Where to go for more information

The following table lists documents on the NetApp on the Web™ (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 and the supported filer models for Data ONTAP licensed with FCP	<i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Target HBA slot assignments	<i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/
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, which is supplied with the adapter and also available 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 Release Notes</i> (if available)

About this chapter

This chapter describes how Data ONTAP reserves space for storing data in LUNs and provides guidelines for estimating the amount of space you for your LUNs. It also describes 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 152 for information about these topics.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Understanding filer storage units](#)” on page 16
- ◆ “[Understanding space reservation for volumes and LUNs](#)” on page 18
- ◆ “[Understanding how fractional reserve affects available space](#)” on page 21
- ◆ “[How guarantees on flexible volumes affect fractional reserve](#)” on page 32
- ◆ “[Calculating the size of a volume](#)” on page 34
- ◆ “[Guidelines for creating volumes that contain LUNs](#)” on page 39
- ◆ “[Creating LUNs, igroups, and LUN maps](#)” on page 45

Understanding filer storage units

Filer storage units for managing disk space

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

- ◆ Aggregates
- ◆ Traditional or flexible volumes
- ◆ Qtrees
- ◆ Files
- ◆ LUNs

The aggregate is the physical layer of storage that consists of the disks within the Redundant Array of Independent Disks (RAID) groups and the plexes that contain the RAID groups. Aggregates provide the underlying physical storage for traditional and flexible volumes.

A traditional volume is directly tied to the underlying aggregate and its properties. When you create a traditional volume, Data ONTAP creates the underlying aggregate based on the properties you assign with the `vol create` command, such as the disks assigned to the RAID group and RAID-level protection.

A flexible volume is loosely tied to the underlying aggregate. You create an aggregate by specifying its physical properties, such as its size and number of disks. Within each aggregate you can create one or more flexible volumes—the logical file systems that share the physical storage resources, RAID configuration, and plex structure of that common containing aggregate. This means that the flexible volume is not tied directly to the physical storage.

You use either traditional or flexible volumes to organize and manage system and user data. A volume can hold qtrees and LUNs. A qtree is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume in order to group LUNs.

For detailed information

For detailed information about filer storage units, including aggregates, and traditional and flexible volumes, see the *Data ONTAP System Administration Storage Management Guide*.

Where LUNs reside You create LUNs in the root of a volume (traditional or flexible) or in a qtree, with the exception of the root volume. Do not create LUNs in the root volume because it is used by Data ONTAP for system administration. The default root volume is /vol/vol0.

Understanding space reservation for volumes and LUNs

What space reservation is

Data ONTAP uses space reservation to guarantee that space is available for completing writes to a LUN or for overwriting data in a LUN. When you create a LUN, Data ONTAP reserves enough space in the traditional or flexible volume so that write operations to those LUNs do not fail because of a lack of disk space on the filer. Other operations, such as taking a Snapshot™ copy or the creation of new LUNs, can occur only if there is enough available unreserved space; other operations are restricted from using reserved space.

What fractional reserve is

Fractional reserve controls the amount of space Data ONTAP reserves in a traditional or flexible volume to enable overwrites to space-reserved LUNs. When you create a space-reserved LUN, fractional reserve is by default set to 100 percent. This means that Data ONTAP automatically reserves 100 percent of the total LUN size for overwrites. For example, if you create a 500-GB space-reserved LUN, Data ONTAP by default ensures that the host-side application storing data in the LUN always has access to 500 GB of space.

You can reduce the amount of space reserved for overwrites to less than 100 percent when you create LUNs in the following types of volumes:

- ◆ Traditional volumes
- ◆ Flexible volumes that have the `guarantee` option set to `volume`.

If the `guarantee` option for a flexible volume is set to `file`, then fractional reserve is set to 100 percent and is not adjustable.

For detailed information about how guarantees affect fractional reserve, see [“How guarantees on flexible volumes affect fractional reserve”](#) on page 32.

How the total LUN size affects reserved space

The amount of space reserved for overwrites is based on the total size of all space-reserved LUNs in a volume. For example, if there are two 200 GB LUNs in a volume, and the `fractional_reserve` option is set to 50 percent, then Data ONTAP guarantees that the volume has 200 GB available for overwrites to those LUNs.

Note

Fractional overwrite is set at the volume level. It does not control how the total amount of space reserved for overwrites in a volume is applied to individual LUNs in that volume.

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, write operations to a LUN might fail due to insufficient disk space and the host application or operating system might crash. The LUN goes offline when the volume is full.

When write operations fail, 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 lun_path [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.

Command for setting fractional reserve

You use the following command to set fractional reserve:

```
vol options vol-name fractional_reserve pct
```

pct is the percentage of the LUN you want to reserve for overwrites. The default setting is 100. For traditional volumes and flexible volumes with the volume guarantee, you can set *pct* to any value from 0 to 100. For flexible volumes with the file guarantee, *pct* is set to 100 by default and is not adjustable.

Example: The following command sets the fractional reserve space on a volume named testvol to 50 percent:

```
vol options testvol fractional_reserve 50
```

How space reservation settings persist

Space reservation settings persist across reboots, takeovers, givebacks, and snap restores. A single file SnapRestore® action restores the reserved state of a LUN to the reserved state at the time the snapshot was taken. For example, if you restore a LUN or volume from a snapshot, the space reservation setting on the LUN is restored and the fractional reserve setting for that volume is restored.

If you revert from Data ONTAP 7.0 to Data ONTAP 6.5, or from Data ONTAP 6.5 to 6.4, the space reservation option remains on. If you revert from Data ONTAP 6.4 to 6.3, the space reservation option is set to `off`.

How revert operations affect fractional reserve

Fractional reserve is available in Data ONTAP 6.5.1 or later. Data ONTAP 6.4.x does not support setting the amount of reserve space to less than 100 percent of the total LUN size. If you want to revert from Data ONTAP 6.5.1 to Data ONTAP 6.4.x, and are using fractional reserve, make sure you have enough available space for 100 percent overwrite reserve. If you do not have enough space when you revert, Data ONTAP displays the following prompt:

```
You have an over committed volume. You are required to set the
fractional_reserve to 100. This can be done by either disabling
space reservations on all objects in the volume or making more
space available for full reservations or deleting all the snapshots
in the volume.
```

Understanding how fractional reserve affects available space

What fractional reserve provides

Fractional reserve enables you to tune the amount of space reserved for overwrites based on application requirements and the rate of change of your data. You define fractional reserve settings per volume. For example, you can group LUNs with a high rate of change in one volume and leave the fractional reserve setting of the volume at the default setting of 100 percent. You can group LUNs with a low rate of change in a separate volume with a lower fractional reserve setting and therefore make better use of available volume space.

Risk of using fractional reserve

Fractional reserve requires to you actively monitor space consumption and the rate of change of data in the volume to ensure you do not run out of space reserved for overwrites. If you run out of overwrite reserve space, writes to the active file system fail and the host application or operating system might crash. This section includes an example of how a volume might run out of free space when you are using fractional reserve. For details, see [“How a volume with fractional overwrite reserve runs out of free space”](#) on page 30.

Data ONTAP provides tools for monitoring available space in your volumes. After you calculate the initial size of your volume and the amount of overwrite reserve space you need, you can monitor space consumption by using these tools. For details, see [“Monitoring disk space”](#) on page 87.

For detailed information

For detailed information, see the following sections:


- ◆ [“How 100 percent fractional reserve affects available space”](#) on page 22
- ◆ [“How reducing fractional reserve affects available space”](#) on page 28


Understanding how fractional reserve affects available space

How 100 percent fractional reserve affects available space

What happens when the fractional overwrite option is set to 100 percent

When you create a space-reserved LUN, fractional reserve is by default set to 100 percent. The following example shows how this setting affects available space in a 1-TB volume with a 500-GB LUN.


Stage	Status
1	The following illustration shows a 1-TB volume with a 500-GB LUN after 200 GB of data are written to the LUN. The volume has 200 GB of space intended for overwrite reserve. This space is actually reserved only when you take a snapshot by using either the <code>snap</code> command or snapshot-methods, such as the SnapMirror® method. For example, if you take a snapshot in the volume shown in the illustration, the original 200 GB of data in the LUN are locked in the snapshot. The reserve space guarantees that you can write over the original 200 GB of data inside the LUN even after you take the snapshot. It guarantees that an application storing data in the LUN always has 500 GB of space available for writes. 


Stage	Status
2	The following illustration shows that the volume still has enough space for the following: ◆ 500-GB LUN (containing 200 GB of data) ◆ 200 GB intended reserve space for overwrites ◆ An additional 200 GB of other data At this point, there is enough space for one snapshot. 

How the volume runs out of free space

The following two examples show how the volume might run out of free space when the fractional overwrite option is set to 100 percent.

Example 1:


Stage	Status
1	The following illustration shows the 1-TB volume with a 500-GB LUN that contains 200 GB of data. There are 200 GB intended for overwrite reserve. At this point, you have not taken a snapshot, and the volume has 500 GB of available space.  The diagram illustrates a 1 TB Volume as a cylinder. It is divided into three horizontal sections. The bottom section, representing 200 GB, is filled with diagonal hatching and labeled '200 GB Data writes into the LUN'. The middle section, representing 200 GB, is solid gray and labeled '200 GB intended for overwrite reserve'. The top section, representing 100 GB, is white. A bracket on the right side of the cylinder spans the height of the hatched and gray sections, labeled '500 GB LUN'. A larger bracket on the right side spans the entire height of the cylinder, labeled '1 TB Volume'.

Stage	Status
2	The following illustration shows the volume after you write 400 GB of other data. Data ONTAP reports that the volume is full when you try to take a snapshot. This is because the 400 GB of other data does not leave enough space for the intended overwrite reserve. The snapshot requires Data ONTAP to reserve 200 GB of space, but you have only 100 GB of available space.  The diagram illustrates a 1 TB volume represented as a cylinder. It is divided into three main sections: a top section with diagonal hatching labeled '400 GB Other data', a middle solid black section labeled '200 GB intended for overwrite reserve', and a bottom section with diagonal hatching labeled '200 GB Data writes into the LUN'. A bracket on the right side of the cylinder indicates the total '1 TB Volume'. Another bracket on the right side, spanning the bottom section and the reserve section, is labeled '500 GB LUN'.

Example 2:

Stage	Status
1	A 1-TB volume has a 500-GB LUN that contains 200 GB of data. There are 200 GB of intended reserve space in the free area of the volume.

Stage	Status
2	The following illustration shows the volume with a snapshot. The volume has 200 GB reserved for overwrites to the original data and 300 GB of free space remaining for other data. 300 GB free for other data 200 GB reserved for overwrites after the first snapshot 200 GB reserved for overwrites after the first snapshot 200 GB Data writes into the LUN 500 GB LUN 1 TB Volume
3	The following illustration shows the volume after you write 300 GB of other data to the volume. 300 GB Other data 200 GB reserved for overwrites after the first snapshot 200 GB reserved for overwrites after the first snapshot 200 GB Data writes into the LUN 500 GB LUN 1 TB Volume

Stage	Status
4	The following illustration shows the volume after you write another 100 GB of data to the LUN. At this point, the volume does not have enough space for another snapshot. The second snapshot requires 300 GB of reserve space because the total size of the data in the LUN is 300 GB.  The diagram illustrates a 1 TB storage volume divided into five horizontal layers. From top to bottom, the layers are: 300 GB of other data (diagonal hatching), 200 GB reserved for overwrites after the first snapshot (solid black), 100 GB of new data written to the LUN (solid light gray), 100 GB reserved for overwrites after the second snapshot (solid black), and 200 GB of data written into the LUN (diagonal hatching). A bracket on the right indicates the total volume is 1 TB. A bracket on the left indicates the 500 GB LUN consists of the 100 GB new data and the 400 GB reserved space (200 GB + 200 GB).

Understanding how fractional reserve affects available space


How reducing fractional reserve affects available space


**When you can
reduce fractional
reserve**

You can reduce fractional reserve to less than 100 percent for traditional volumes or for volumes that have the guarantee option set to volume.

**What happens
when the fractional
reserve option is
set to 50 percent**


The following example shows how a fractional reserve setting of 50 percent affects available space in the same 1-TB volume with a 500-GB LUN.


Stage	Status
1	The following illustration shows a 1-TB volume with a 500-GB LUN after 200 GB of data are written to the LUN. The volume has 100 GB intended for overwrite reserve because the fractional reserve for this volume is set to 50 percent. 

Stage	Status
2	The following illustration shows the volume with an additional 300 GB of other data. The volume still has 100 GB of free space, which means there is space for one of the following: ◆ Writing up to 200 GB of new data to the LUN and maintaining the ability to take a snapshot ◆ Writing up to 100 GB of other data and maintaining the ability to take a snapshot Compare this example with the volume shown in “Example 2” on page 25, in which the same volume has an overwrite reserve of 100 percent, but the volume has run out of free space.  The diagram illustrates a 1 TB volume represented as a cylinder. It is divided into four horizontal sections. From top to bottom: a hatched section labeled '300 GB Other data'; a solid black section labeled '100 GB intended overwrite reserve'; a light gray section labeled '500 GB LUN'; and a hatched section at the bottom labeled '200 GB Data writes into the LUN'. A bracket on the right side of the cylinder spans the entire height and is labeled '1 TB Volume'. Another bracket on the right side spans the light gray LUN section and the bottom hatched section, labeled '500 GB LUN'.

How a volume with fractional overwrite reserve runs out of free space

The following example shows how the volume might run out of space when the fractional reserve option is set to 50 percent.

Stage	Status
1	The following illustration shows a 1-TB volume with a 500-GB LUN after you write 500 GB to the LUN and then take a snapshot. The volume has 250 GB reserved for overwrites to the LUN and 250 GB available for other data.  The diagram illustrates a 1TB volume represented as a cylinder. The cylinder is divided into three horizontal sections. The bottom section is light gray and is labeled '500GB LUN' with a bracket. Above this is a dark gray section labeled '250 GB overwrite reserve' with a bracket. The top section is white and labeled '250 GB free for other data' with a bracket. A large bracket on the right side of the cylinder encompasses all three sections and is labeled '1TB Volume'.

Stage	Status
2	The following illustration shows that you have 50 GB of free space after you write 200 GB of other data (for example, files) to the volume. You try to write more than 300 GB of data to the LUN, and the write fails. The volume has 50 GB of free space plus 250 GB of space reserved for overwrites to the LUN. The volume has enough space for you to write no more than 300 GB of data to the LUN. 

How guarantees on flexible volumes affect fractional reserve

What guarantees are

Guarantees on a flexible volume ensure that write operations to a specified flexible volume or write operations to LUNs with space reservation on that file do not fail because of lack of available space in the containing aggregate. Guarantees determine how the aggregate preallocates space to the flexible volume. Guarantees are set at the volume level. There are three types of guarantees:

- ◆ `volume`

A guarantee of `volume` ensures that the amount of space required by the flexible volume is always available from its aggregate. This is the default setting for flexible volumes. Fractional reserve is an adjustable value. For example, if you set the fractional reserve to 50 percent in a 200-GB flexible volume, you have 100 GB of intended reserve space in the volume. By default, guarantees for flexible volumes are set to `volume`.

- ◆ `file`

The aggregate guarantees that space is always available for overwrites to space-reserved LUNs. Fractional reserve is set to 100 percent and is not adjustable.

- ◆ `none`

A flexible volume with a guarantee of `none` reserves no space, regardless of the space reservation settings for LUNs in that volume. Write operations to space-reserved LUNs in that volume might fail if its containing aggregate does not have enough available space.

Command for setting guarantees

You use the following command to set volume guarantees:

```
vol options f_vol_name guarantee guarantee_value
```

f_vol_name is the name of the flexible volume whose space guarantee you want to change.

guarantee_value is the space guarantee you want to assign to this volume. The possible values are `volume`, `file`, and `none`.

For detailed information about setting guarantees, see the *Data ONTAP Storage Management Guide*.


Overcommitting an aggregate

You might want to overcommit an aggregate to enable flexible provisioning. For example, you might need to assign large volumes to specific users, but you know they will not use all their available space initially. When your users require additional space, you can increase the size of the aggregate on demand by assigning additional disks to the aggregate.

To overcommit an aggregate, you create flexible volumes with a guarantee of `none` or `file`, so that the volume size is not limited by the aggregate size. The total size of the flexible volumes you create might be larger than the containing aggregate.

The following example shows a 1-TB aggregate with two flexible volumes. The guarantee is set to `file` for each flexible volume. Each flexible volume contains a 200-GB LUN. The `file` guarantee ensures that there are 200 GB of intended reserve space in each flexible volume so that write operations to the space-reserved LUNs do not fail, regardless of the size of the flexible volumes that contain the LUNs.

Each flexible volume has space for other data. For example, you can create non-space-reserved LUNs in a flexible volume, but write operations to these LUNs or LUNs might fail when the aggregate runs out of free space.



For detailed information

For detailed information about using guarantees, see the *Data ONTAP Storage Management Guide*.

Calculating the size of a volume

What the volume size depends on

Before you create the volumes that contain qtrees and LUNs, calculate the size of the volume and the amount of reserve space required by determining the type and the amount of data that you want to store in the LUNs on the volume.

The size of the volume depends on the following:

- ◆ Total size of all the LUNs in the volume
- ◆ Whether you want to maintain snapshots
- ◆ If you want to maintain snapshots, the number of snapshots you want to maintain and the amount of time you want to retain them (retention period).
- ◆ Rate at which data in the volume changes
- ◆ Amount of space you need for overwrites to LUNs (fractional reserve).

The amount of fractional reserve depends on the rate at which your data changes and how quickly you can adjust your system when you know that available space in the volume is scarce.


Decision process for estimating the size of a volume

Use the decision process in the flowchart shown on the following page to estimate the size of the volume. For detailed information about each step in the decision process, see the following sections:

- ◆ [“Calculating the total LUN size”](#) on page 35
- ◆ [“Determining the volume size when you don’t need snapshots”](#) on page 36
- ◆ [“Calculating the amount of space for snapshots”](#) on page 36
- ◆ [“Calculating the fractional reserve”](#) on page 37

Example: Your database needs two 20 GB disks. You must create two 20 GB LUNs.

Note: Some filer data protection mechanisms, such as Snapmirror rely on snapshots.



Calculating the total LUN size

The total LUN size is the sum of the LUNs you want to store in the volume. The size of each LUN depends on the amount of data you want to store in the LUNs. For example, if you know your database needs two 20-GB disks, you must create two 20-GB LUNs. The total LUN size in this example is 40 GB.

Determining the volume size when you don't need snapshots

If you are not using snapshots, the size of your volume depends on the size of the LUNs and whether you are using traditional or flexible volumes.

- ◆ Traditional volumes

Traditional volumes are tied directly to the physical storage. When you create traditional volumes, you specify the number of disks used to create it. The capacity and number of disks you specify determines the size of the volume. For example, a 72-GB disk provides approximately 69.7 GB of usable space. If 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 actual amount of usable space for six 72-GB disks is 407.4 GB.

If you are using traditional volumes, create a volume that has enough disks to accommodate the size of your LUNs. For example, if you need two 200-GB LUNs, create a volume with enough disks to provide 400 GB of storage capacity.

- ◆ Flexible volumes

If you are using flexible volumes, the size of the flexible volume is the total size of all the LUNs in the volume.

ONTAP data protection methods and snapshots: Before you determine that you do not need snapshots, verify the method for protecting data in your configuration. Most data protection methods, such as the SnapRestore, SnapMirror, SyncMirror®, dump and restore, and the ndmccopy methods rely on snapshots. If you are using these methods, calculate the amount of space required for these snapshots.

Note

Host based backup methods do not require additional space.

Calculating the amount of space for snapshots

The amount of space you need for snapshots depends on the following:

- ◆ Estimated Rate of Change (ROC) of your data per day.

The ROC is required to determine the amount of space you need for snapshots and fractional overwrite reserve. The ROC depends on how often you overwrite data.

- ◆ Number of days that you want to keep old data in snapshots. For example, if you take one snapshot per day and want to save old data for two weeks, you need enough space for 14 snapshots.

You can use the following guideline to calculate the amount of space you need for snapshots:

Space for snapshots = ROC in bytes per day * number of snapshots

Example: You need a 20-GB LUN, and you estimate that your data changes at a rate of about 10 percent, or 2 GB each day. You want to take one snapshot each day and want to keep three weeks' worth of snapshots, for a total of 21 snapshots. The amount of space you need for snapshots is $21 * 2$ GB, or 42 GB.

Calculating the fractional reserve

The fractional reserve setting depends on the following:

- ◆ Amount of time you need to enlarge your volume by either adding disks or deleting old snapshots when free space is scarce.
- ◆ ROC of your data
- ◆ Size of all LUNs that will be stored in the volume

Example: You have a 20-GB LUN and your data changes at a rate of 2 GB each day. You want to keep 21 snapshots. You want to ensure that write operations to the LUNs do not fail for three days after you take the last snapshot. You need $2 \text{ GB} * 3$, or 6 GB of space reserved for overwrites to the LUNs. Thirty percent of the total LUN size is 6 GB, so you must set your fractional reserve to 30 percent.

Calculating the size of a sample volume

The following example shows how to calculate the size of a volume based on the following information:

- ◆ You need to create two 50-GB LUNs.
The total LUN size is 100 GB.
- ◆ Your data changes at a rate of 10 percent of the total LUN size each day.
Your ROC is 10 GB per day (10 percent of 100 GB).
- ◆ You take one snapshot each day and you want to keep the snapshots for 10 days.
You need 100 GB of space for snapshots (10 GB ROC * 10 snapshots).
- ◆ You want to ensure that you can continue to write to the LUNs through the weekend, even after you take the last snapshot and you have no more free space.
You need 20 GB of space reserved for overwrites (10 GB per day ROC * 2 days). This means you must set fractional reserve to 20 percent ($20 \text{ GB} = 20$ percent of 100 GB).

Calculate the size of your volume as follows:

Volume size = Total LUN size + Amount of space for snapshots + Space for overwrite reserve

The size of the volume in this example is 220 GB (100 GB + 100 GB + 20 GB).

How fractional reserve settings affect the total volume size: When you set the fractional reserve to less than 100 percent, writes to LUNs are not unequivocally guaranteed. In this example, writes to LUNs will not fail for about two days after you take your last snapshot. You must monitor available space and take corrective action by increasing the size of your volume or aggregate or deleting snapshots to ensure you can continue to write to the LUNs.

If you leave the fractional reserve at the default setting of 100 percent in this example, Data ONTAP sets aside 100 GB as intended reserve space. The volume size must be 300 GB, which breaks down as follows:

- ◆ 100 GB for 100 percent fractional reserve.
- ◆ 100 GB for the total LUN size (50 GB plus 50 GB).
- ◆ 100 GB for snapshots.

This means you initially need an extra 80 GB for your volume.

Calculating the size of the volume with LUN clones

If you want to create a readable-writable clone of a LUN, ensure that space reservation is enabled for the LUN and consider the clone a LUN that is the same size as the parent. When you calculate the size of the volume, make sure you have enough space for:

- ◆ The parent LUNs and their snapshots
- ◆ The LUN clones and their snapshots

Guidelines for creating volumes that contain LUNs

Guidelines to use when creating volumes

NetApp strongly recommends that you use the following guidelines to create traditional or flexible volumes that contain LUNs:

- ◆ Do not create any LUNs in the filer's root volume. Data ONTAP uses this volume to administer the filer. The default root volume is /vol/vol0.
- ◆ Ensure that the snapshot functionality is modified as follows:
 - ❖ Set the snap reserve to zero.
 - ❖ Turn off the automatic snapshot schedule.

For detailed procedures, see [“Changing snapshot defaults”](#) on page 39.

- ◆ 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.

For detailed procedures, see [“Verifying and modifying the volume option `create\_unicode`”](#) on page 43.

- ◆ Use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

For information about creating aggregates, volumes and qtrees

For detailed procedures that describe how to create and configure aggregates, volumes, and qtrees, see the *Data ONTAP Storage Management Guide*.

Changing snapshot defaults

Why you need to change snapshot defaults: NetApp snapshots are required for many NetApp features, such as the SnapMirror feature, SyncMirror feature, dump and restore, and ndmcopy.

When you create a volume, Data ONTAP automatically does the following:

- ◆ Reserves 20 percent 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 zero.
- ◆ Turn off the automatic snap schedule.

For Windows systems and some UNIX hosts, you use SnapDrive™ for Windows or SnapDrive™ for UNIX to ensure that applications accessing LUNs are quiesced or synchronized automatically before taking snapshots. With UNIX 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 whether your UNIX host is supported by SnapDrive for UNIX, see the *NetApp FCP SAN Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml. Click the link for your host operating system (OS). The compatibility matrix for your host lists the version of SnapDrive supported in a row called “Snapshot Integration.”

For information about how to use snapshots, see “[Using snapshots](#)” on page 114.

Setting the percentage of snap reserve space by 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 volname percent</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>

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

Setting the percentage of snap reserve space by using FilerView: To use FilerView 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	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Click Volumes > Snapshots > Configure.
3	Select the volume you want to configure.
4	In 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 0. _____
5	Click Apply.

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

Step	Action
1	To turn off the automatic snapshot schedule, enter the following command: <pre>snap sched volname 0 0 0</pre> Example: <code>snap sched vol1 0 0 0</code> Result: This command turns off the snapshot schedule because there are no weekly, nightly, or hourly snapshots scheduled. You can still take snapshots manually by using the <code>snap</code> command.
2	To verify that the automatic snapshot schedule is off, enter the following command: <pre>snap sched [volname]</pre> Example: <code>snap sched vol1</code> Result: The following output is a sample of what is displayed: <pre>Volume vol1: 0 0 0</pre>

Turning off the automatic snapshot schedule by using FilerView: To turn off the automatic snapshot schedule on a volume and to verify that the schedule is off, 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	In the Hourly Snapshot Schedule field, ensure that no time slots are selected. For example, if a check appears at 8:00 AM, click it to deselect it.
5	Click Apply.

Verifying and
modifying the
volume option
create_ucose

Modifying the create_ucose option using the command line: To use the command line to verify that the create_ucose volume option is enabled, or to enable the option, complete the following steps.

Step	Action
1	To verify that the create_ucose option is enabled (On), enter the following command: <pre>vol status [volname] -v</pre> Example: vol status vol1 -v Result: The following output example shows that the create_ucose option is on: <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, the status of all volumes is displayed.
2	To enable the create_ucose option, enter the following command: <pre>vol options volname create_ucose on</pre> Example: vol options vol1 create_ucose on

Modifying the create_unicode option by using FilerView: The create_unicode option is displayed as the “Create Unicode Format Directories By Default” volume option field. To verify that this 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	From the left panel, click Volumes.
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.
6	Click Apply.

Creating LUNs, igroups, and LUN maps

Methods for creating LUNs, igroups, and LUN maps

You use one of the following methods to create LUNs and igroups:

- ◆ 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 52.
- ◆ 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 57.
- ◆ Entering a series of individual commands (such as `lun create`, `igroup create`, and `lun map`)
Use this method to create one or more LUNs and igroups in any order. For information about this method, see “[Creating LUNs and igroups by using individual commands](#)” on page 61.

Caution about using SnapDrive

For Windows hosts, you can use SnapDrive™ for Windows to create and manage LUNs. If you use SnapDrive to create LUNs, you must use it for all LUN management functions. Do not use the Data ONTAP command line interface or FilerView to manage LUNs.

For information about the version of SnapDrive supported for your host, see the *NetApp FCP SAN Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

Click the link for your host operating system. The compatibility matrix for your host lists the version of SnapDrive supported in a row called “Snapshot Integration”.

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 in which the LUN is located. Do not create LUNs in the root volume. The default root volume is `/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 type of data that it stores, or the user accessing 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 following:

- ◆ Geometry used to access data on the LUN
- ◆ Minimum LUN sizes
- ◆ Layout of data for multiprotocol access

The LUN ostype values are *solaris*, *windows*, *hpux*, *aix*, *linux*, and *image*. When you create a LUN, specify the ostype that corresponds to your host. If your host OS is not one of these values but it is listed as a supported OS in the *NetApp FCP SAN Compatibility Matrix*, specify *image*.

For information about supported hosts, see the *NetApp FCP SAN Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

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
g	gigabytes
t	terabytes

The usable space in the LUN depends on host or application requirements for overhead. For example, partition tables and metadata on the host file system reduce the usable space for applications. In general, when you format and partition LUNs as a disk on a host, the actual usable space on the disk depends on the overhead required by the host.

The disk geometry used by the operating system determines the minimum and maximum size values of LUNs. For information about the maximum sizes for LUNs and disk geometry, see the vendor documentation for your host OS. If you are using third-party volume management software on your host, consult the vendor's documentation for more information about how disk geometry affects LUN size.

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 it. This is used to create the map between the LUN and the host. When you map a LUN to an igroup, you can specify a LUN ID. If you do not specify a LUN ID, Data ONTAP automatically assigns one.

Space reservation setting: When you create a LUN by using the `lun setup` command or FilerView, you specify whether you want to enable space reservation. When you create a LUN using the `lun create` command, space reservation is automatically turned on.

Note

NetApp recommends that you keep this setting on.

About igroups

Initiator groups (igroups) are tables of WWPNs of hosts and are used to control access to LUNs. Typically, you want all host bus adapters (HBAs) to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA of each clustered host needs 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. However, you cannot map a LUN to multiple igroups that have the same initiator.

Note

An initiator cannot be a member of igroups of differing otypes.

The following table illustrates how four igroups give access to the LUNs for four different hosts accessing the filer. The clustered hosts (Host3 and Host4) are both members of the same igroup (solaris-group2) and can access the LUNs mapped to this igroup. 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/qtree1/lun2
Host4, multipath, clustered (connected to Host3) 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	solaris-group3	10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol1/qtree1/lun2 /vol/vol1/qtree1/lun3

Required information for creating 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 character, 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 Domain Name Service (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 igroup type is FCP in a Fibre Channel SAN.

The ostype of the initiators: The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostyles of initiators are *solaris*, *windows*, *hpux*, *aix*, and *linux*. If your host OS is not one of these values but it is listed as a supported OS in the *NetApp FCP SAN Compatibility Matrix*, specify *default*.

For information about supported hosts, see the *NetApp FCP SAN Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

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.

To know which WWPNs are associated with a specific host, see the SAN Host Attach Kit documentation for your host. These documents describe commands supplied by NetApp or the vendor of the initiator or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, you use the `lputilnt` utility, and for UNIX hosts, you use the `sanlun` command. For information about using the `sanlun` command on UNIX hosts, see “[Creating an igroup using the sanlun command \(UNIX hosts\)](#)” on page 100.

What is required to map a LUN to an igroup

When you map the LUN 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. Data ONTAP maintains a separate LUN map for each igroup to support a large number of hosts and to enforce access control.

You map a LUN to an igroup by specifying the following attributes:

LUN name: Specify the path name of the LUN to be mapped.

Initiator group: Specify the name of the igroup that contains the hosts that will access the LUN.

LUN ID: 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 host associates the LUN ID with the location and path name of the LUN. The range of valid LUN ID numbers depend on the host. For detailed information, see the documentation provided with your SAN Host Attach Kit.

Guidelines for mapping LUNs to igroups: Use the following guidelines when mapping LUNs to igroups:

- ◆ You can map two different LUNs with the same LUN ID to two different igroups without having a conflict, provided that the igroups do not share any initiators or only one of the LUNs is online at a given time.
- ◆ You can map a LUN only once to an igroup or a specific initiator.
- ◆ You can add a single initiator to multiple igroups, but the initiator can be mapped to a LUN only once. You cannot map a LUN to multiple igroups that contain the same initiator.
- ◆ You cannot use the same LUN ID for two LUNs mapped to the same igroup.

Guidelines for LUN layout and space requirements

When you create LUNs, use the following guidelines for layout and space requirements:

- ◆ Group LUNs according to their rate of change.

If you plan to take snapshots, do not create LUNs with high rate of change in the same volumes as LUNs with a low rate of change. When you calculate the size of your volume, the rate of change of data enables you determine the amount of space you need for snapshots. Data ONTAP takes snapshots at the volume level, and the rate of change of data in all LUNs affects the amount of space needed for snapshots. If you calculate your volume size based on a low rate of change, and you then create LUNs with a high rate of change in that volume, you might not have enough space for snapshots.
- ◆ Keep backup LUNs in separate volumes.

Network Appliance recommends that you keep backup LUNs in separate volumes because the data in a backup LUN changes 100 percent for each backup period. For example, you might copy all the data in a LUN to a backup LUN and then move the backup LUN to tape each day. The data in the backup LUN changes 100 percent each day. If you want to keep backup LUNs in the same volume, calculate the size of the volume based on a high rate of change in your data.

- ◆ Quotas are another method you can use to allocate space. For example, you might want to assign volume space to various database administrators and allow them to create and manage their own LUNs. You can organize the volume into qtrees with quotas and enable the individual database administrators to manage the space they have been allocated.

If you organize your LUNs in qtrees with quotas, make sure the quota limit can accommodate the sizes of the LUNs you want to create. Data ONTAP does not allow you to create a LUN in a qtree with a quota if the LUN size exceeds the quota.

Host-side procedures required

The host detects LUNs as disk devices. When you create a new LUN and map it to an igroup, you must configure the host to detect the new LUNs. The procedure you use depends on your host operating system. On HP-UX hosts, for example, you use the `ioscan` command. For detailed procedures, see the documentation for your SAN Host Attach Kit.

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 for storing LUNs before running the `lun setup` program, terminate the program and create volumes. If you want to use qtrees, 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 <code>lun setup</code> program displays the following instructions. Press Enter to continue or <code>n</code> 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: <pre>OS type of LUN (image/solaris/windows/hpux/linux) [image]:</pre> Example: <code>solaris</code> For information about specifying the <code>ostype</code> of the LUN, see “The host operating system type” on page 46.

Step	Action
3	Specify the name of the LUN and 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/finance/ and want to create a LUN called records, you enter, /vol/finance/records. Note_____ Do not create LUNs in the root volume because it is used for filer administration. _____ Result: A LUN called records is created in the root of /vol/finance 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]: Caution_____ If you choose n, space reservation is disabled. This might cause write operations 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 disk space usable 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 <code>lun comment</code> command or fill in the description field by 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 is used. If you have not created any igroups, enter a name of the group you want to create now. For information about naming an igroup, see “The name of the igroup” on page 49.
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 an example 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, 10:00:00:00: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 “How hosts are identified” on page 7.
10	Specify the operating system type that the 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]: For information about specifying the ostype of an igroup, see “About igroups” on page 47.

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/finance/records 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, which is 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 and igroups 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	In the left panel of the FilerView screen, click LUNs, Result: The management tasks you can perform on LUNs are displayed.
3	Click Wizard. Result: The LUN Wizard window appears.

Step	Action
4	Click the Next button to continue. Result: The first window of fields in the LUN Wizard appears.  LUN Wizard: Specify LUN Parameters Path: The full path to the LUN, for example <code>/vol/luns/lunOne</code>. The LUN must be created in the root directory of a volume or a qtree. Size: The size of the LUN. Multiprotocol Type: Select the multiprotocol type for the LUN. Space-reserved: If checked, indicates that the LUN should be space-reserved. Description: An optional description of the LUN. < Back Cancel Next >
5	Enter LUN information in the appropriate fields and click Next
6	Specify the following information in the next windows: ◆ Whether you want to add an igroup. ◆ Whether to you want to use an existing igroup or create a new one. ◆ WWPNs of the initiators in the igroup ◆ LUN mapping
7	In 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 use FilerView menus to create LUNs, 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.
4	Click Add to commit changes.

Creating igroups: To use FilerView menus to create an igroup, 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 Add to commit changes.

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

Step	Action
1	Launch FilerView, as described in “ Launching FilerView ” on page 3.
2	Go to LUNs > Manage.
3	If the maps are not displayed, click the Hide Maps link.
4	In the first column, find the LUN to which you want to map an igroup. ◆ If the LUN is mapped, yes or the name of the igroup and the LUN ID appears in the last column. Click yes to add igroups to the LUN mapping. ◆ If the LUN is not mapped, no or No Maps appears in the last column. Click no to map the LUN to an igroup.

Step	Action
5	Click Add Groups to Map.
6	Select an igroup name from the list on the right side of the window.
7	To commit your changes click Add.

Creating LUNs, igroups, and LUN maps

Creating LUNs and igroups by 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 connected to the filer.	<pre>fcpx show initiator</pre> Sample result: <pre>Initiators connected on adapter 7a: Portname Group 10:00:00:00:c9:39:4d:82 50:06:0b:00:00:11:35:62 10:00:00:00:c9:34:05:0c 10:00:00:00:c9:2f:89:41 10:00:00:00:c9:2d:56:5f Initiators connected on adapter 7b: Portname Group 10:00:00:00:c9:2f:89:41 10:00:00:00:c9:2d:56:5f 10:00:00:00:c9:39:4d:82 50:06:0b:00:00:11:35:62 10:00:00:00:c9:34:05:0c</pre>
Determine which hosts are associated with the WWPNs	For information about how to determine which WWPN is associated with a host, see “ How hosts are identified ” on page 7.

To do this...	Use this command...
Create an igroup	igroup create -f -t <i>ostype</i> <i>initiator_group</i> [<i>node</i>] -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 47. <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 space-reserved LUN	lun create -s <i>size</i> -t <i>ostype</i> <i>lun_path</i> -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 46. -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 46. <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. Space reservation is enabled for the LUN.

To do this...	Use this command...
Map the LUN to an igroup	<pre>lun map lun_path initiator_group [lun_id]</pre> <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 the initiator uses when the LUN is mapped to it. If you do not enter a number, Data ONTAP generates the next available LUN ID number. Example 1: <code>lun map /vol/vol1/qtree1/lun3 solaris-igroup3 0</code> Result: Data ONTAP maps /vol/vol1/qtree1/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 the command in this example is entered, Data ONTAP displays the following message: <pre>lun map: auto-assigned solaris-igroup0=0</pre>
Display the LUNs you created	<pre>lun show -v</pre> <code>-v</code> provides additional information, such as the comment string, serial number, and LUN mapping. Example: <code>lun show -v</code> Sample result: <pre>/vol/vol1/qtree1/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><th>Protocol</th></tr><tr><td colspan="4">-----</td></tr><tr><td>/vol/tpcc_disks/ctrl_0</td><td>solaris_cluster</td><td>0</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/ctrl_1</td><td>solaris_cluster</td><td>1</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/crash1</td><td>solaris_cluster</td><td>2</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/crash2</td><td>solaris_cluster</td><td>3</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/cust_0</td><td>solaris_cluster</td><td>4</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/cust_1</td><td>solaris_cluster</td><td>5</td><td>FCP</td></tr><tr><td>/vol/tpcc_disks/cust_2</td><td>solaris_cluster</td><td>6</td><td>FCP</td></tr></table>	LUN path	Mapped to	LUN ID	Protocol	-----				/vol/tpcc_disks/ctrl_0	solaris_cluster	0	FCP	/vol/tpcc_disks/ctrl_1	solaris_cluster	1	FCP	/vol/tpcc_disks/crash1	solaris_cluster	2	FCP	/vol/tpcc_disks/crash2	solaris_cluster	3	FCP	/vol/tpcc_disks/cust_0	solaris_cluster	4	FCP	/vol/tpcc_disks/cust_1	solaris_cluster	5	FCP	/vol/tpcc_disks/cust_2	solaris_cluster	6	FCP
LUN path	Mapped to	LUN ID	Protocol																																		

/vol/tpcc_disks/ctrl_0	solaris_cluster	0	FCP																																		
/vol/tpcc_disks/ctrl_1	solaris_cluster	1	FCP																																		
/vol/tpcc_disks/crash1	solaris_cluster	2	FCP																																		
/vol/tpcc_disks/crash2	solaris_cluster	3	FCP																																		
/vol/tpcc_disks/cust_0	solaris_cluster	4	FCP																																		
/vol/tpcc_disks/cust_1	solaris_cluster	5	FCP																																		
/vol/tpcc_disks/cust_2	solaris_cluster	6	FCP																																		
Determine the maximum possible size of a LUN in a volume or qtree	lun maxsize vol-path vol-path is the path to the volume or qtree in which you want to create the LUN. Result: The lun maxsize command displays the maximum possible size of a LUN in the volume or qtree, depending on the LUN type and geometry. It also shows the maximum size possible for each LUN type with or without snapshots. Sample result: <pre>lun maxsize /vol/lunvol Space available for a LUN of type: solaris, aix, hpux, linux, or image Without snapshot reserve: 184.9g (198508019712) With snapshot reserve: 89.5g (96051658752) Space available for a LUN of type: windows Without snapshot reserve: 184.9g (198525358080) With snapshot reserve: 89.5g (96054819840)</pre>																																				

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 66
- ◆ “[Displaying LUN information](#)” on page 72
- ◆ “[Reallocating LUN and volume layout](#)” on page 77
- ◆ “[Monitoring disk space](#)” on page 87

Managing LUNs and LUN maps

Tasks to manage LUNs and LUN maps

You can use the command-line interface or FilerView to

- ◆ Control LUN availability
- ◆ Unmap a LUN from an igroup
- ◆ Rename a LUN
- ◆ Resize a LUN
- ◆ Modify the LUN description
- ◆ Enable or disable space reservations
- ◆ Remove a LUN
- ◆ Access a LUN with NAS protocols

Actions that require host-side procedures

The host detects LUNs as disk devices. The following actions make LUNs unavailable to the host and require host-side procedures so that the host detects the new configuration.

- ◆ Taking a LUN offline
- ◆ Bringing a LUN online
- ◆ Unmapping a LUN from an igroup
- ◆ Removing a LUN
- ◆ Resizing a LUN
- ◆ Renaming a LUN

The procedure depends on your host operating system. For example, on HP-UX hosts, you use the `ioscan` command. For detailed procedures, see the documentation for your SAN Host Attach Kit.

Controlling LUN availability

The `lun online` and `lun offline` commands enable and control the availability of LUNs while preserving mappings.

Before you bring a LUN online or take it offline, make sure that you quiesce or synchronize any host application accessing the LUN.

Bringing a LUN online: To bring one or more LUNs online, complete the following step.

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

Taking a LUN offline: Taking a LUN offline makes it unavailable for block protocol access. To take a LUN offline, complete the following step.

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

Renaming a LUN

To rename a LUN, completing the following step.

Step	Action
1	Enter the following command: <pre>lun move <i>lun_path</i> <i>new_lun_path</i></pre> Example: <code>lun move /vol/vol1/mylun /vol/vol1/mynewlun</code> Note_____ If you are organizing LUNs in qtrees, the existing path (<i>lun_path</i>) and the new path (<i>new_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 systems, resizing is supported only on basic disks. Resizing is not supported on dynamic disks.
- ◆ If you are running VxVM version 3.5 or lower, resizing LUNs is not supported.
- ◆ If you want to increase the size of the LUN, the SCSI disk geometry imposes an upper limit for the size of the LUN, and Data ONTAP imposes a maximum increase to 2 TB.

For additional restrictions on resizing a LUN, see the following documents:

- ◆ *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
- ◆ Documentation for your SAN Host Attach Kit.
- ◆ Vendor documentation for your operating system.

To change the size of 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: <pre>lun offline lun_path</pre> Example: <code>lun offline /vol/vol1/qtrees/lun2</code>
2	Change the size of the LUN by entering the following command: <pre>lun resize [-f] lun_path new_size</pre> -f overrides warnings when you are decreasing the size of the LUN. Example: (Assuming that lun2 is 5 GB and you are increasing it to 10 GB) <pre>lun resize /vol/vol1/qtrees1/lun2 10g</pre>
3	From the host, rescan or rediscover the LUN so that the new size is recognized. For detailed procedures see the documentation for your SAN Host Attach Kit.

Modifying the LUN description

To modify the LUN description, complete the following step.

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

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, write operations to a LUN might fail due to insufficient disk space and the host application or operating system might crash. The LUN goes offline when the volume is full.

When write operations fail, 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 lun_path [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] lun_path [lun_path ...]</pre> <code>-f</code> 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 <code>-f</code> 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 be accessed only with SAN protocols by default. However, you can use NAS protocols to make a LUN available to a host if the NAS protocols are licensed and enabled on the filer. The usefulness of accessing a LUN over NAS protocols depends on the host application.

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 a FCP SAN host from overwriting data in the LUN. To make a LUN accessible to a host that uses 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: <code>lun share lun_path {none read write all}</code> 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
 - ◆ Statistics about read operations, write operations, and the number of operations per second
 - ◆ LUN mapping
 - ◆ Settings for space reservation
 - ◆ Additional information, such as serial number or ostype.

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: <pre>lun help</pre> 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 read and write operations and the number of operations per second for LUNs, complete the following step.

Step	Action																																																																																																			
1	Enter the following command: <pre>lun stats -z -k -i interval -c count -o [-a lun_path]</pre> -z zeros statistics Note_____ The statistics start at zero at boot time. _____ -k displays the statistics in KBs. -i interval is the interval, in seconds, at which the statistics are displayed. -c <i>count</i> is the number of intervals. For example, the <code>lun stats -i 10 -c 5</code> command displays statistics in ten-second intervals, for five intervals. -o <i>displays</i> additional statistics, including the number of QFULL messages the filer sends when its SCSI command queue is full and the amount of traffic received from the partner filer. -a shows statistics for all LUNs <i>lun_path</i> displays statistics for a specific LUN Example: <pre>lun stats -o -i 1</pre> <table><tr><th>Read Ops</th><th>Write Ops</th><th>Other Ops</th><th>QFull</th><th>Read kB</th><th>Write kB</th><th>Average Latency</th><th>Queue Length</th><th>Partner Ops</th><th>Partner kB</th><th>Lun</th></tr><tr><td>0</td><td>351</td><td>0</td><td>0</td><td>0</td><td>44992</td><td>11.35</td><td>3.00</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr><tr><td>0</td><td>233</td><td>0</td><td>0</td><td>0</td><td>29888</td><td>14.85</td><td>2.05</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr><tr><td>0</td><td>411</td><td>0</td><td>0</td><td>0</td><td>52672</td><td>8.93</td><td>2.08</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr><tr><td>2</td><td>1</td><td>0</td><td>0</td><td>16</td><td>8</td><td>1.00</td><td>1.00</td><td>0</td><td>0</td><td>/vol/tpcc/ctrl_0</td></tr><tr><td>1</td><td>1</td><td>0</td><td>0</td><td>8</td><td>8</td><td>1.50</td><td>1.00</td><td>0</td><td>0</td><td>/vol/tpcc/ctrl_1</td></tr><tr><td>0</td><td>326</td><td>0</td><td>0</td><td>0</td><td>41600</td><td>11.93</td><td>3.00</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr><tr><td>0</td><td>353</td><td>0</td><td>0</td><td>0</td><td>45056</td><td>10.57</td><td>2.09</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr><tr><td>0</td><td>282</td><td>0</td><td>0</td><td>0</td><td>36160</td><td>12.81</td><td>2.07</td><td>0</td><td>0</td><td>/vol/tpcc/log_22</td></tr></table>	Read Ops	Write Ops	Other Ops	QFull	Read kB	Write kB	Average Latency	Queue Length	Partner Ops	Partner kB	Lun	0	351	0	0	0	44992	11.35	3.00	0	0	/vol/tpcc/log_22	0	233	0	0	0	29888	14.85	2.05	0	0	/vol/tpcc/log_22	0	411	0	0	0	52672	8.93	2.08	0	0	/vol/tpcc/log_22	2	1	0	0	16	8	1.00	1.00	0	0	/vol/tpcc/ctrl_0	1	1	0	0	8	8	1.50	1.00	0	0	/vol/tpcc/ctrl_1	0	326	0	0	0	41600	11.93	3.00	0	0	/vol/tpcc/log_22	0	353	0	0	0	45056	10.57	2.09	0	0	/vol/tpcc/log_22	0	282	0	0	0	36160	12.81	2.07	0	0	/vol/tpcc/log_22
Read Ops	Write Ops	Other Ops	QFull	Read kB	Write kB	Average Latency	Queue Length	Partner Ops	Partner kB	Lun																																																																																										
0	351	0	0	0	44992	11.35	3.00	0	0	/vol/tpcc/log_22																																																																																										
0	233	0	0	0	29888	14.85	2.05	0	0	/vol/tpcc/log_22																																																																																										
0	411	0	0	0	52672	8.93	2.08	0	0	/vol/tpcc/log_22																																																																																										
2	1	0	0	16	8	1.00	1.00	0	0	/vol/tpcc/ctrl_0																																																																																										
1	1	0	0	8	8	1.50	1.00	0	0	/vol/tpcc/ctrl_1																																																																																										
0	326	0	0	0	41600	11.93	3.00	0	0	/vol/tpcc/log_22																																																																																										
0	353	0	0	0	45056	10.57	2.09	0	0	/vol/tpcc/log_22																																																																																										
0	282	0	0	0	36160	12.81	2.07	0	0	/vol/tpcc/log_22																																																																																										

Displaying LUN mapping information

To display LUN mapping information, complete the following step.

Step	Action																																				
1	On the filer's command line, enter the following command: <pre>lun show -m</pre> Result: <table><tr><th>LLUN path</th><th>Mapped to</th><th>LUN ID</th><th>Protocol</th></tr><tr><td colspan="4">-----</td></tr><tr><td>/vol/tpcc/ctrl_0</td><td>solaris_cluster</td><td>0</td><td>FCP</td></tr><tr><td>/vol/tpcc/ctrl_1</td><td>solaris_cluster</td><td>1</td><td>FCP</td></tr><tr><td>/vol/tpcc/crash1</td><td>solaris_cluster</td><td>2</td><td>FCP</td></tr><tr><td>/vol/tpcc/crash2</td><td>solaris_cluster</td><td>3</td><td>FCP</td></tr><tr><td>/vol/tpcc/cust_0</td><td>solaris_cluster</td><td>4</td><td>FCP</td></tr><tr><td>/vol/tpcc/cust_1</td><td>solaris_cluster</td><td>5</td><td>FCP</td></tr><tr><td>/vol/tpcc/cust_2</td><td>solaris_cluster</td><td>6</td><td>FCP</td></tr></table>	LLUN path	Mapped to	LUN ID	Protocol	-----				/vol/tpcc/ctrl_0	solaris_cluster	0	FCP	/vol/tpcc/ctrl_1	solaris_cluster	1	FCP	/vol/tpcc/crash1	solaris_cluster	2	FCP	/vol/tpcc/crash2	solaris_cluster	3	FCP	/vol/tpcc/cust_0	solaris_cluster	4	FCP	/vol/tpcc/cust_1	solaris_cluster	5	FCP	/vol/tpcc/cust_2	solaris_cluster	6	FCP
LLUN path	Mapped to	LUN ID	Protocol																																		

/vol/tpcc/ctrl_0	solaris_cluster	0	FCP																																		
/vol/tpcc/ctrl_1	solaris_cluster	1	FCP																																		
/vol/tpcc/crash1	solaris_cluster	2	FCP																																		
/vol/tpcc/crash2	solaris_cluster	3	FCP																																		
/vol/tpcc/cust_0	solaris_cluster	4	FCP																																		
/vol/tpcc/cust_1	solaris_cluster	5	FCP																																		
/vol/tpcc/cust_2	solaris_cluster	6	FCP																																		

Displaying status of space reservations

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 lun_path Example: lun set reservation /vol/lunvol/hpux/lun0 Space Reservation for LUN /vol/lunvol/hpux/lun0 (inode 3903199): enabled

Displaying additional LUN information

To display additional information about LUNs, such as the serial number and ostype (displayed as Multiprotocol Type), and maps, complete the following step.

Step	Action
1	On the filer's command line, enter the following command to display LUN status and characteristics: <pre>lun show -v</pre> Example: <pre>/vol/tpcc_disks/cust_0_1 382m (400556032) (r/w, online, mapped) Serial#: VqmOVYoe3BUf Share: none Space Reservation: enabled Multiprotocol Type: solaris SnapValidator Offset: 1m (1048576) Maps: sun_hosts=0 /vol/tpcc_disks/cust_0_2 382m (400556032) (r/w, online, mapped) Serial#: VqmOVYoe3BV6 Share: none Space Reservation: enabled Multiprotocol Type: solaris SnapValidator Offset: 1m (1048576) Maps: sun_hosts=1</pre>

Reallocating LUN and volume layout

Reasons to use reallocation scans

You use reallocation scans to ensure that blocks in a LUN, large file, or volume are laid out sequentially. If a LUN, large file, or volume is not laid out in sequential blocks, sequential read commands take longer to complete because each command might require an additional disk seek operation. Sequential block layout improves the read/write performance of host applications that access data on the filer.

What a reallocation scan is

A reallocation scan evaluates how the blocks are laid out in a LUN, file, or volume. Data ONTAP performs the scan as a background task, so applications can rewrite blocks in the LUN or volume during the scan. Repeated layout checks during a scan ensure that the sequential block layout is maintained during the current scan.

A reallocation scan does not necessarily rewrite every block in the LUN. Rather, it rewrites whatever is required to optimize the layout of the LUN.

How a reallocation scan works

Data ONTAP performs a reallocation scan in the following steps:

1. Scans the current block layout of the LUN.
2. Determines the level of optimization of the current layout on a scale of 3 (moderately optimal) to 10 (not optimal).
3. Performs one of the following tasks, depending on the optimization level of the current block layout:
 - If the layout is optimal, the scan stops.
 - If the layout is not optimal, blocks are reallocated sequentially.
4. Scans the new block layout.
5. Repeats steps 2 and 3 until the layout is optimal.

Reallocation scans and LUN availability

You can perform reallocation scans on LUNs when they are online. You do not have to take them offline. You also do not have to perform any host-side procedures when you perform reallocation scans.

**How you manage
reallocation scans**

You manage reallocation scans by performing the following tasks:

- ◆ First, enable reallocation scans.
- ◆ Then, either define a reallocation scan to run at specified intervals (such as every 24 hours), or define a reallocation scan to run on a specified schedule that you create (such as every Thursday at 3:00 p.m.).

You can define only one reallocation scan for a single LUN.

You can also initiate scans at any time, force Data ONTAP to reallocate blocks sequentially regardless of the optimization level of the LUN layout, and monitor and control the progress of scans.

If you delete a LUN, you do not delete the reallocation scan defined for it. If you take the LUN offline, delete it, and then reconstruct it, you still have the reallocation scan in place. However, if you delete a LUN that has a reallocation scan defined and you do not restore the LUN, the filer console displays an error message the next time the scan is scheduled to run.

**Enabling
reallocation scans**

Reallocation scans are disabled by default. You must enable reallocation scans globally on the filer before you run a scan or schedule regular scans.

To enable reallocation scans, complete the following step.

Step	Action
1	On the filer’s command line, enter the following command: <code>reallocate on</code>

**Defining a
reallocation scan**

To define a reallocation scan for a LUN, complete the following step:

Step	Action
1	On the filer's command line, enter the following command: <pre>reallocate start [-t <i>threshold</i>] [-n] [-i <i>interval</i>] <i>lun_path</i></pre> -t <i>threshold</i> is a number between 3 (layout is moderately optimal) and 10 (layout is not optimal). The default is 4. A scan checks the block layout of a LUN before reallocating blocks. If the current layout is below the threshold, the scan does not reallocate blocks in the LUN. If the current layout is equal to or above the threshold, the scan reallocates blocks in the LUN. -n reallocates blocks in the LUN without checking its layout. -i <i>interval</i> is the interval, in hours, minutes, or days, at which the scan is performed. The default interval is 24 hours. Specify the interval as follows: [m h d] For example, 30m is a 30-minute interval. The countdown to the next scan begins only after the first scan is complete. For example, if the interval is 24 hours and a scan starts at midnight and lasts for an hour, the next scan begins at 1:00 a.m. the next day—24 hours after the first scan is completed. Examples: The following example creates a new LUN and a normal reallocation scan that runs every 24 hours. <pre>lun create -s 100g /vol/vol2/lun0 reallocate start /vol/vol2/lun0</pre>

Step	Action	
2	If...	Then...
	You want to run the reallocation scan according to a schedule	Proceed to “ Creating a reallocation scan schedule ” on page 81.
	You do not want to define a schedule	Proceed to “ Tasks for managing reallocation scans ” on page 82.

**Creating a
reallocation scan
schedule**

You can run reallocation scans according to a schedule. The schedule you create replaces any interval you specified when you entered the `reallocate start` command.

To create a reallocation scan schedule, complete the following step.

Step	Action
1	Enter the following command: <pre>reallocate schedule [-s <i>schedule</i>] <i>lun_path</i></pre> <code>-s <i>schedule</i></code> is a string with the following fields: <pre>"<i>minute hour day_of_month day_of_week</i>"</pre> ❖ <i>minute</i> is a value from 0 to 59. ❖ <i>hour</i> is a value from 0 (midnight) to 23 (11:00 p.m.). ❖ <i>day_of_month</i> is a value from 1 to 31. ❖ <i>day_of_week</i> is a value from 0 (Sunday) to 6 (Saturday). A wildcard character (*) indicates every value for that field. For example, a * in the <i>day_of_month</i> field means every day of the month. You cannot use the wildcard character in the <i>minute</i> field. You can enter a number, a range, or a comma-separated list of values for a field. For example, entering "0,1" in the <i>day_of_week</i> field means Sundays and Mondays. You can also define a range of values. For example, "0-3" in the <i>day_of_week</i> field means Sunday through Wednesday. Examples: The following example schedules a reallocation scan for every Saturday at 11:00 PM. <pre>reallocate schedule -s "0 23 * 6" /vol/myvol/lun1</pre>

**Deleting a
reallocation scan
schedule**

You can delete an existing reallocation scan schedule that is defined for a LUN. If you delete a schedule, the scan runs according to the interval that you specified when you initially defined the scan using the `reallocate start` command.

A reallocation scan is not automatically deleted if you delete its corresponding LUN. However, if you destroy a volume, all reallocation scans defined for LUNs in that volume are deleted.

To delete a reallocation scan schedule, complete the following step:

Step	Action
1	Enter the following command: <code>reallocate schedule -d lun_path</code> Example: <code>reallocate schedule -d /vol/myvol/lun1</code>

Tasks for managing reallocation scans

You perform the following tasks to manage reallocation scans:

- ◆ Start a one-time reallocation scan.
- ◆ Start a scan that reallocates every block in a LUN or volume, regardless of layout.
- ◆ Display the status of a reallocation scan.
- ◆ Stop a reallocation scan.
- ◆ Quiesce a reallocation scan.
- ◆ Restart a reallocation scan.
- ◆ Disable reallocation.

Starting a one-time reallocation scan

You can perform a one-time reallocation scan on a LUN. This type of scan is useful if you do not want to schedule regular scans for a particular LUN.

To start a one-time reallocation scan, complete the following step:

Step	Action
1	Enter the following command: <code>reallocate start -o -n lun_path</code> -o performs the scan only once. -n performs the scan without checking the LUN's layout.

Performing a full reallocation scan of a LUN or volume

You can perform a scan that reallocates every block in a LUN or a volume regardless of the current layout by using the `-f` option of the `reallocate start` command. A full reallocation optimizes layout more aggressively than a normal reallocation scan. A normal reallocation scan moves blocks only if the move improves LUN layout. A full reallocation scan always moves blocks, unless the move makes the LUN layout even worse.

Using the `-f` option of the `reallocate start` command implies the `-o` and `-n` options. This means that the full reallocation scan is performed only once, without checking the LUN's layout first.

You might want to perform this type of scan if you add a new RAID group to a volume and you want to ensure that blocks are laid out sequentially throughout the volume or LUN.

If the volume contains snapshots, then use the full reallocation with caution. In this case, a full reallocation might result in using significantly more space in the volume, because the old, unoptimized blocks are still present in the snapshot after the scan.

To perform a full reallocation scan, complete the following step:

Step	Action
1	Enter the following command: <code>reallocate start -f lun_path volume-path</code>

Quiescing a reallocation scan

You can quiesce a reallocation scan that is in progress and restart it later. The scan restarts from the beginning of the reallocation process. For example, if you want to back up a LUN, but a scan is already in progress, you can quiesce the scan.

To quiesce a reallocation scan, complete the following step.

Step	Action
1	Enter the following command: <code>reallocate quiesce lun_path</code>

Restarting a reallocation scan

You might restart a scan for the following reasons:

- ◆ You quiesced the scan by using the `reallocate quiesce` command, and you want to restart it.
- ◆ You have a scheduled scan that is idle (it is not yet time for it to run again), and you want to run it immediately.

To restart a scan, complete the following step:

Step	Action
1	Enter the following command: <pre>reallocate restart lun_path</pre> Result: The command restarts a quiesced scan. If there is a scheduled scan that is idle, the <code>reallocate restart</code> command runs the scan.

Viewing the status of a scan

To view the status of a scan, complete the following step:

Step	Action
1	Enter the following command: <pre>reallocate status [-v] lun_path</pre> <code>-v</code> provides verbose output <i>lun_path</i> is the path to the LUN for which you want to see reallocation scan status. If you do not specify a value for <i>lun_path</i>, then the status for all scans is displayed. Result: The <code>reallocate status</code> command displays the following information: ◆ State—whether the scan is in progress or idle. ◆ Schedule—schedule information about the scan. If there is no schedule, then the <code>reallocate status</code> command displays n/a. ◆ Interval—intervals at which the scan runs, if there is no schedule defined. ◆ Optimization—information about the LUN layout.

Deleting a reallocation scan

You use the `reallocate stop` command to permanently delete a scan you defined for a LUN. The `reallocate stop` command also stops any scan that is in progress on the LUN.

To delete a scan, complete the following step:

Step	Action
1	Enter the following command: <pre>reallocate stop lun_path</pre> Result: The <code>reallocate stop</code> command stops and deletes any scan on the LUN, including a scan in progress, a scheduled scan that is not running, or a scan that is quiesced.

Disabling reallocation scans

You use the `reallocate off` command to disable reallocation on the filer. When you disable reallocation scans, you cannot start or restart any new scans. Any scans that are in progress are stopped. If you want to re-enable reallocation scans at a later date, use the `reallocate on` command.

To disable reallocation scans, complete the following step:

Step	Action
1	On the filer's command line, enter the following command: <pre>reallocate off</pre>

Best practice recommendations

NetApp recommends the following best practices for using reallocation scans:

- ◆ Define a reallocation scan when you first create the LUN. This ensures that the LUN layout remains optimized as a result of regular reallocation scans.
- ◆ Define regular reallocation scans by using either intervals or schedules. This ensures that the LUN layout remains optimized. If you wait until most of the blocks in the LUN layout are not sequential, a reallocation scan will take more time.

- ◆ Define intervals according to the type of read/write activity associated with the LUN:
 - ❖ Long intervals—Define long reallocation scan intervals for LUNs in which the data changes slowly, for example, LUNs in which data changes as a result of infrequent large write operations.
 - ❖ Short intervals—Define short reallocation scan intervals for LUNs that are characterized by workloads with many small random write and many sequential read operations. These types of LUNs might become heavily fragmented over a shorter period of time.
- ◆ If a LUN has an access pattern of random write operations followed by periodic large sequential read operations (for example, it is accessed by a database or a mail backup application), you can schedule reallocation scans to take place before you back up the LUN. This ensures that the LUN is optimized before the backup.

Monitoring disk space

Commands for monitoring disk space

You use the following commands to monitor disk space:

- ◆ `snap delta`—Estimates the rate of change of data between snapshots in a volume. For detailed information, see [“Estimating the rate of change of data between snapshots”](#) below.
- ◆ `snap reclaimable`—Estimates the amount of space freed if you delete the specified snapshots. If space in your volume is scarce, you can reclaim free space by deleting a set of snapshots. For detailed information, see [“Estimating the amount of space freed by snapshots”](#) on page 89.
- ◆ `df`—Displays the statistics about the active file system and the snapshot directory in a volume or aggregate. For detailed information, see [“Displaying statistics about free space”](#) on page 89.

Estimating the rate of change of data between snapshots

When you initially set up volumes and LUNs, you estimate the rate of change of your data to calculate the volume size. After you create the volumes and LUNs, you use the `snap delta` command to monitor the actual rate of change of data. You can adjust the fractional overwrite reserve or increase the size of your aggregates or volumes based on the actual rate of change.

Displaying the rate of change: To display the rate of change of data between snapshots, complete the following steps:

Step	Action																																																							
1	Enter the following command: <pre>snap delta [-A] vol_name snapshot snapshot</pre> -A displays the rate of change of data between snapshots for all aggregates in the system. <i>vol_name</i> is the name of the volume. <i>snapshot</i> is the name of the snapshot. If you do not specify an argument, the <code>snap delta</code> command displays the rate of change of data between snapshots for all volumes in the system. Example: The following example displays the rate of change of data between all snapshots in <code>vol0</code>. <pre>filer_1> snap delta vol0 Volume vol0 working...</pre> <table><tr><th>From Snapshot</th><th>To</th><th>KB changed</th><th>Time</th><th>Rate (KB/hour)</th></tr><tr><td>hourly.0</td><td>Active File System</td><td>1460</td><td>0d 02:16</td><td>639.961</td></tr><tr><td>nightly.0</td><td>hourly.0</td><td>1492</td><td>0d 07:59</td><td>186.506</td></tr><tr><td>hourly.1</td><td>nightly.0</td><td>368</td><td>0d 04:00</td><td>91.993</td></tr><tr><td>hourly.2</td><td>hourly.1</td><td>1420</td><td>0d 04:00</td><td>355.000</td></tr><tr><td>hourly.3</td><td>hourly.2</td><td>1960</td><td>0d 03:59</td><td>490.034</td></tr><tr><td>hourly.4</td><td>hourly.3</td><td>516</td><td>0d 04:00</td><td>129.000</td></tr><tr><td>nightly.1</td><td>hourly.4</td><td>1456</td><td>0d 08:00</td><td>182.000</td></tr><tr><td>hourly.5</td><td>nightly.1</td><td>364</td><td>0d 04:00</td><td>91.000</td></tr></table> <pre>Summary...</pre> <table><tr><th>From Snapshot</th><th>To</th><th>KB changed</th><th>Time</th><th>Rate (KB/hour)</th></tr><tr><td>hourly.5</td><td>Active File System</td><td>9036</td><td>1d 14:16</td><td>236.043</td></tr></table>	From Snapshot	To	KB changed	Time	Rate (KB/hour)	hourly.0	Active File System	1460	0d 02:16	639.961	nightly.0	hourly.0	1492	0d 07:59	186.506	hourly.1	nightly.0	368	0d 04:00	91.993	hourly.2	hourly.1	1420	0d 04:00	355.000	hourly.3	hourly.2	1960	0d 03:59	490.034	hourly.4	hourly.3	516	0d 04:00	129.000	nightly.1	hourly.4	1456	0d 08:00	182.000	hourly.5	nightly.1	364	0d 04:00	91.000	From Snapshot	To	KB changed	Time	Rate (KB/hour)	hourly.5	Active File System	9036	1d 14:16	236.043
From Snapshot	To	KB changed	Time	Rate (KB/hour)																																																				
hourly.0	Active File System	1460	0d 02:16	639.961																																																				
nightly.0	hourly.0	1492	0d 07:59	186.506																																																				
hourly.1	nightly.0	368	0d 04:00	91.993																																																				
hourly.2	hourly.1	1420	0d 04:00	355.000																																																				
hourly.3	hourly.2	1960	0d 03:59	490.034																																																				
hourly.4	hourly.3	516	0d 04:00	129.000																																																				
nightly.1	hourly.4	1456	0d 08:00	182.000																																																				
hourly.5	nightly.1	364	0d 04:00	91.000																																																				
From Snapshot	To	KB changed	Time	Rate (KB/hour)																																																				
hourly.5	Active File System	9036	1d 14:16	236.043																																																				

Interpreting snap delta output: The first row of the `snap delta` output displays the rate of change between the most recent snapshot and the active file system. The following rows provide the rate of change between successive snapshots. Each row displays the names of the two snapshots that are compared, the amount of data that has changed between them, the time elapsed between the two snapshots, and how fast the data changed between the two snapshots.

If you do not specify any snapshots when you enter the `snap delta` command, the output also displays a table that summarizes the rate of change for the volume between the oldest snapshot and the active file system.

Estimating the amount of space freed by snapshots

To estimate the amount of space freed by deleting a set of snapshots, complete the following step.

Step	Action
1	Enter the following command: <pre>snap reclaimable vol_name snapshot snapshot...</pre> <i>vol_name</i> is the name of the volume. <i>snapshot</i> is the name of the snapshot. You can specify more than one snapshot. Example: The following example shows the approximate amount of space that would be freed by deleting two snapshots. <pre>filer_1> snap reclaimable vol0 hourly.1 hourly.5 Processing (Press Ctrl-C to exit) ... snap reclaimable: Approximately 1860 Kbytes would be freed.</pre>

Displaying statistics about free space

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 display space in KB, MB, GB, or TB format, respectively.

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 overwrites to existing LUNs.

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 following statistics are displayed:

- ◆ Amount of total space on the volume, in the byte format you specify.
Total space = used space + available space.
- ◆ Amount of used space.

Used space = space storing data + space storing snapshots + space reserved for overwrites.

- ◆ Amount of available space.

Available space = space that is not used or reserved; it is free space.

- ◆ Percentage of the volume capacity being used.

This information is displayed if you do not use the `-r` option.

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.

How LUN and snapshot operations affect disk space

The following table illustrates the effect on disk space when you create a sample 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, fractional overwrite reserve is set to 100 percent, and snap reserve is set to 0 percent.

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, the LUN 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.

Action	Result	Comment
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 write operations 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 write operations to LUNs. Because reserved space is also counted as used space, used space is 80 GB.
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 total amount of used space remains unchanged.

Action	Result	Comment
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.
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 using the df command

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*.

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 this example, assume that you require less than the minimum capacity based on the NetApp recommendation of creating a seven-disk volume.

For simplicity, assume the LUN requires only 3 GB of disk space. For a traditional volume, the volume size must be approximately 3 GB plus 10 percent. If you plan to use 72-GB disks (which typically provide 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 traditional 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 toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. There is a snap reserve of 20 percent on the volume even though the volume will be used for LUNs because snap reserve is set to 20 percent 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 toaster> df -r /vol/volspace</pre> Result: The following sample output is displayed. The amount of available snapshot space becomes zero, and the 20 percent 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 aix /vol/volspace/lun0 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>aix_cluster</i> and map the LUN to it by entering the following commands (assuming that your host has an HBA whose WWPN is 10:00:00:00:c9:2f:98:44). 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 aix aix_host 10:00:00:00:c9:2f:98:44 toaster> lun map /vol/volspace/lun0 aix_host 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 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, write operations 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, and fractional overwrite reserve is set to 100 percent. The recommended volume size is approximately 2*3 GB plus the rate of change of data. Assuming the amount of change is small, the rate of change is minimal, so using two 72-GB disks still provides 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 percent on the volume even though the volume will be used for 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 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														

Step	Action																								
2	Set the percentage of snap reserve space to zero by entering the following command: toaster> snap reserve volspace 0																								
3	Create a LUN (/vol/volspace/lun0) by entering the following commands: toaster> lun create -s 6g -t aix /vol/volspace/lun0 toaster> df -r /vol/volspace 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>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	6300536	56169372	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	6300536	56169372	0	/vol/volspace/																				
/vol/volspace/.snapshot	0	0	0	0	0 /vol/volsp																				
ace/.snapshot																									
4	Create an igroup named <i>aix_host</i> and map the LUN to the igroup by entering the following commands. These commands have no effect on disk space. toaster> igroup create -f -t aix aix_host 10:00:00:00:c9:2f:98:44 toaster> lun map /vol/volspace/lun0 aix_host 0																								
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.																								
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, 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 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 write operations 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														
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 operation, there was only 1 GB of data, after the write operation, there were 1 GB of new data and 1 GB of data in a snapshot. Notice that the used space increases for the snapshot by 1 GB, and the available space for the volume decreases by 1 GB. <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>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														

Step	Action																								
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. _____ toaster> snap create volspace snap2 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/</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/	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																									
10	From the host, write 2 GB of data to the LUN and observe the effect on disk space by entering the following command: toaster> df -r /vol/volspace Result: The following sample output is displayed. The second write operation 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/</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/	62649908	12601072	4608427	6300536	/vol/volspace/	/vol/volspace/.snapshot	0	3150371	0	0	/vol/volspace/	.snapshot					
Filesystem	kbytes	used	avail	reserved	Mounted on																				
/vol/volspace/	62649908	12601072	4608427	6300536	/vol/volspace/																				
/vol/volspace/.snapshot	0	3150371	0	0	/vol/volspace/																				
.snapshot																									

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 100
- ◆ “[Managing initiator requests](#)” on page 105

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 using the filer command line

To create an igroup, complete the following step.

Step	Action
1	Enter the following command: <pre>igroup create -f [-t ostype] initiator_group [node_name...]</pre> -f indicates that it is an FCP igroup. -t ostype indicates the operating system of the host. The values are <i>solaris</i>, <i>windows</i>, <i>hpux</i>, <i>aix</i>, or <i>linux</i>. initiator_group is the name of the igroup you specify. node_name is an FCP WWPN. You can specify more than one WWPN. Example: <code>igroup create -f -t hpux hpux 50:06:0b:00:00:10:a7:00 50:06:0b:00:00:10:a6:06</code>

Creating an igroup using the sanlun command (UNIX hosts)

If you have a UNIX host, you can run the `sanlun` command on the host to create an igroup. The command obtains the host's WWPNs and prints out the `igroup create` command with the correct arguments. You can then copy and paste this command into the filer's command line.

To create igroup by using the `sanlun` command, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the <code>/opt/NetApp/santools/bin</code> directory.
3	Enter the following command to print a command to be run on the filer that creates an igroup containing all the HBAs on your host: <pre>./sanlun fcp show adapter -c</pre> <code>-c</code> option prints the full <code>igroup create</code> command on the screen. Result: An <code>igroup create</code> command with the host's WWPNs appears on the screen. The igroup's name matches the name of the host. Example: Enter this filer command to create an initiator group for this system: <pre>igroup create -f -t solaris "hostA" 10000000AA11BB22 10000000AA11EE33</pre> In this example, the name of the host is "hostA," so the name of the igroup with the two WWPNs is "hostA."
4	On the host in a different session, use the <code>telnet</code> command to access the filer.
5	Copy the <code>igroup create</code> command from Step 3, paste the command on the filer's command line, and press Enter to run the <code>igroup</code> command on the filer. Result: An igroup is created on the filer.

Step	Action
6	On the filer's command line, enter the following command to verify the newly created igroup: <pre>igroup show</pre> Result: The newly created igroup with the host's WWPNs is displayed. Example: <pre>filerX> igroup show hostA (FCP) (ostype: solaris): 10:00:00:00:AA:11:BB:22 10:00:00:00:AA:11:EE:33</pre>

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	<pre>lun unmap lun_path igroup</pre> Example: lun unmap /vol/vol2/qtrees/LUN10 solaris-group5
	Delete one or more igroups	<pre>igroup destroy igroup [igroup,...]</pre> Example: igroup destroy solaris-group5
	Remove all LUN maps for an igroup and delete the igroup with one command	<pre>igroup destroy -f igroup [igroup ...]</pre> 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: <pre>igroup add igroup WWPN</pre> Caution_____ When adding initiators to an igroup, ensure that each initiator sees only one LUN at a given LUN ID. _____ Example: <code>igroup add solaris-group2 10:00:00:00:c9:2b:02:1f</code> 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: <pre>igroup remove igroup WWPN</pre> 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: <pre>igroup show [igroup]</pre> 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. <pre>igroup set igroup ostype value</pre> <i>igroup</i> is the name of the igroup. <i>value</i> is the ostype of the igroup. The ostyles of initiators are <i>solaris</i>, <i>windows</i>, <i>hpux</i>, <i>aix</i>, and <i>linux</i>. If your host OS is not one of these values but it is listed as a supported OS in the NetApp FCP SAN Compatibility Matrix, specify <i>default</i>. For information about supported hosts and ostyles, see the <i>NetApp FCP SAN Compatibility Matrix</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml. Example: <code>igroup set solaris-group3 ostyle 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 manages initiator requests

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 percent of all the resources. The remaining command blocks are always 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 percent, 20 percent of the queue resources available at the filer's ports are reserved for the initiators in that igroup. The remaining 80 percent of the queue resources are 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 percent, the least critical at 10 percent, and the remaining two at 20 percent, leaving 20 percent 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 from these resources.
 - ❖ The pool of unreserved resources to handle unexpected I/O requests
 - ❖ The pool of unused reserved resources, if those resources are available

Creating up an igroup throttle

To create an igroup throttle, complete the following step.

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

Destroying an igroup throttle

To destroy an igroup throttle, complete the following step.

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

Defining whether an igroup can borrow resources

To define whether an igroup can 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 igroup_name 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 <code>y</code>, 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.

Step	Action												
1	Enter the following command: <pre>igroup show -t</pre> Sample output: <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> Explanation of output: The exceeds column displays the number of times the initiator sends more requests than the throttle allows. The borrow column displays the number of times the throttle is exceeded and the filer uses queue resources from the unreserved pool. In the borrow column, “N/A” indicates that the igroup throttle_borrow option is set to no.	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 and the number of command blocks reserved for the igroup on the specified port, complete the following step.

Step	Action																								
1	Enter the following command: <pre>igroup show -t -i interval -c count [igroup -a]</pre> -t displays information on igroup throttles. -i <i>interval</i> displays statistics for the throttles over an interval in seconds. -c <i>count</i> determines how many intervals are shown. <i>igroup</i> is the name of a specific group for which you want to show statistics. -a displays statistics for all igroups, including idle igroups. Example: <code>igroup show -t -i 1</code> Result: The following is a sample display: <table><tr><th>name</th><th>reserved</th><th>4a</th><th>4b</th><th>5a</th><th>5b</th></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> The first number under the port name indicates the number of command blocks the initiator is using. The second number under the port name indicates the number of command blocks reserved for the igroup on that port. 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	Enter the following command: <pre>lun stats -o -i <i>time_in_seconds</i></pre> <i>-i time_in_seconds</i> is the interval over which performance statistics are reported. For example, <i>-i 1</i> reports statistics each second. <i>-o</i> displays additional statistics, including the number of QFULL messages. Example: <code>lun stats -o -i 1 /vol/vol1/lun2</code> Result: The output displays performance statistics, including the QFULL column. This column indicates the number of initiator requests that exceeded the number allowed by the igroup throttle, and, as a result, received the SCSI Queue Full response.
2	Display the total count of QFULLS sent for each LUN by entering the following command: <pre>lun stats -o <i>lun_path</i></pre>

How a cluster failover affects igroup throttles

Throttles manage physical ports, so during a cluster takeover, their behavior varies according to the FCP cmode that is in effect, as shown in the following table.

FCP cmode	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 or dual_fabric	Throttles apply to all ports and are divided by two when the cluster is in takeover.

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: <code>igroup show -t</code> Example: <code>igroup show -t</code> Result: The following is a sample display: <table><tr><th>name</th><th>reserved</th><th>exceeds</th><th>borrows</th></tr><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></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 112
- ◆ [“Using snapshots”](#) on page 114
- ◆ [“Using LUN clones”](#) on page 116
- ◆ [“Deleting busy snapshots”](#) on page 119
- ◆ [“Using SnapRestore”](#) on page 122
- ◆ [“Backing up data to tape”](#) on page 127
- ◆ [“Using NDMP”](#) on page 131
- ◆ [“Using volume copy”](#) on page 132
- ◆ [“Cloning flexible volumes”](#) on page 133
- ◆ [“Using NVFAIL”](#) on page 138
- ◆ [“Using SnapValidator”](#) on page 140

Data ONTAP protection methods

Data protection methods

Data ONTAP provides a variety of methods for protecting data in a Fibre Channel SAN. 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 is created. The other data protection methods listed in the table below 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™ for Windows or UNIX	◆ 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. For some UNIX environments, you can use SnapDrive for UNIX to create snapshots. To see if your UNIX host is supported by SnapDrive, see the <i>NetApp FCP SAN Compatibility Matrix</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml. Click the link for your host operating system (OS). The compatibility matrix for your host lists the version of SnapDrive supported in a row called "Snapshot Integration". ◆ 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_____ 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 FCP network

Snapshots of applications running on a file system may result in the snapshot containing inconsistent data unless measures are taken (such as quiescing the application prior to the snapshot) to ensure the data on disk is logically consistent before you take the snapshot. 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, it might be unavailable for some amount of 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. This action flushes the host file system buffers to disk. Quiescing ensures 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 takes only 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 to 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, data written to the LUN is in the active file system.

After you have a snapshot, you can use it to create a LUN clone for temporary use as a prototype for testing data or scripts in applications or databases. Because the LUN clone is backed by the snapshot, you cannot delete the snapshot until you split the clone from it.

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 and later, space reservation is enabled when you create the LUN. This means that enough space is reserved so that write operations 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 87.

Using LUN clones

What a LUN clone is A LUN clone is a point-in-time, writable copy of a LUN in a snapshot. Changes made to the parent LUN after the clone is created are not reflected in the clone.

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

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.

Reasons for cloning LUNs You can use LUN clones to create multiple read/write copies of a LUN. You might want to do this for the following reasons:

- ◆ You need to create a temporary copy of a LUN for testing purposes.
- ◆ You need to make a copy of your data available to additional users without giving them access to the production data.
- ◆ You want to create a clone of a database for manipulation and projection operations, while preserving the original data in unaltered form.

Creating a snapshot of a LUN Before you can clone a LUN, you must create a snapshot (the backing snapshot) of a LUN you want to clone. To create a snapshot, complete the following steps.

Step	Action
1	Create a LUN by entering the following command: <code>lun create -s size lun_path</code> Example: <code>lun create -s 100g /vol/vol1/lun0</code>

Step	Action
2	Create a snapshot of the volume containing the LUN to be cloned by entering the following command: <pre>snap create volume_name snapshot_name</pre> Example: snap create vol1 mysnap

Creating a clone of the LUN

After you create the snapshot of the LUN, you create the LUN clone. To create the LUN clone, complete the following step.

Step	Action
1	Enter the following command: <pre>lun clone create clone_lun_path -b parent_lun_path parent_snap</pre> <i>clone_lun_path</i> is the path to the clone you are creating, for example, /vol/vol1/lun0clone. <i>parent_lun_path</i> is the path to the original LUN. <i>parent_snap</i> is the name of the snapshot of the original LUN. Example: lun clone create /vol/vol1/lun0clone -b vol/vol1/lun0 mysnap

Splitting the clone from the backing snapshot

You can split the LUN clone from the backing snapshot and then delete the snapshot without taking the LUN offline or losing its contents. To begin the process of splitting the clone from the backing snapshot, complete the following step.

Step	Action
1	Begin the clone operation by entering the following command: <pre>lun clone split start lun_path</pre> <i>lun_path</i> is the path to the parent LUN. Result: The clone does not share data blocks with the snapshot of the original LUN. This means you can delete the snapshot.

Displaying or stopping the progress of a clone splitting operation

Because clone splitting is a copy operation and might take considerable time to complete, you can stop or check the status of a clone splitting operation.

Displaying the progress of a clone-splitting operation: To display the progress of the clone-splitting operation, complete the following step.

Step	Action
1	Enter the following command: <pre>lun clone split status lun_path</pre> <i>lun_path</i> is the path to the parent LUN.

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

Step	Action
1	Enter the following the command: <pre>lun clone split stop lun_path</pre> <i>lun_path</i> is the path to the parent LUN.

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.

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: <pre>snap list vol-name</pre> Example: <pre>snap list vol2</pre> Result: The following message is displayed: <pre>Volume vol2 working...</pre> <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)																						

Step	Action
2	Identify the LUNs and the snapshots that contain them by entering the following command: <pre>lun snap usage vol_name snap_name</pre> Example: <pre>lun snap usage vol2 snap0</pre> Result: The following message is displayed: <pre>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</pre> Note_____ The LUNs are backed by lunA in the snap0 snapshot.
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: <pre>lun destroy /vol/vol2/lunC</pre>
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: <pre>snap delete vol2 snap2 snap delete vol2 snap1</pre> Result: All the snapshots containing lunB are now deleted and snap0 is no longer busy.

Step	Action
5	Delete the snapshot by entering the following command: <code>snap delete vol-name snapshot-name</code> 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 its state when the snapshot was taken. You can use SnapRestore to restore an entire volume, 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 `snap restore` command on a volume containing the LUN or before you run 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.

Caution—————
When a single LUN is restored, it must be taken offline or be unmapped prior to recovery. Using SnapRestore on a LUN, or on a 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 restore. <i>volume_name</i> is the name of the volume to be restored. 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 restore the data. You can enter only one snapshot name. Example: <pre> filer> snap restore -s payroll_lun_backup.2 -t /vol/payroll_lun 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 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 restore the volume.
5	Press y to confirm that you want to restore 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.
	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	Enter the following command to unmap the existing old maps that you don't want to keep. <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.
```

Restore a single LUN

To restore a single LUN (rather than a volume), complete the following steps.

Note—
You cannot use SnapRestore to restore LUNs with NT streams or on directories.

Step	Action
1	Notify network users that you are going to restore a LUN so that they know that the current data in the LUN will be replaced by that of the selected snapshot.
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 file specifies that you are entering the name of a file to revert. -s snapshot_name specifies the name of the snapshot from which to restore the data. -r restore_as_path restores the file to a location in the volume different from the location in the snapshot. For example, if you specify /vol/vol0/vol3/mylun as the argument to -r, SnapRestore restores the file called mylun to the location /vol/vol0/vol3 instead of to the path structure indicated by the path in path_and_lun_name. path_and_LUN_name is the complete path to the name of the LUN to be restored. You can enter only one path name. A LUN can be restored only to the volume where it was originally. The directory structure to which a LUN 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 LUN at the end of the path_and_lun_name 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 restore the LUN.

Step	Action
3	Press y to confirm that you want to restore the file. Result: Data ONTAP displays the name of the LUN and the name of the snapshot for the restore operation. If you did not use the <code>-f</code> option, Data ONTAP prompts you to decide whether to proceed with the restore operation.
4	Press y to continue with the restore operation. Result: Data ONTAP restores 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 restored 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 systems.

Backing up a single LUN to tape

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 for the production application server: <pre>igroup create -f [-t ostype] group [node ...]</pre> Example: <code>igroup create -f -t windows payroll_server 10:00:00:00:c3:4a:0e:e1</code> Result: Data ONTAP creates an igroup called payroll_server, which includes the WWPN (10:00:00:00:c3:4a:0e:e1) of the Windows application server used in the production environment.
2	Enter the following command to create the production 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.
3	Enter the following command to map the production LUN to the igroup that includes the WWPN of the application server. <pre>lun map lun_path initiator-group LUN_ID</pre> Example: <code>lun map /vol/vol1/qtree_1/payroll_lun payroll_server 1</code> Result: Data ONTAP maps the production LUN (/vol/vol_name/qtree_1/payroll_lun) to the payroll_server igroup with a LUN ID of 1.

Step	Action
4	From the host, discover the new LUN, format it, 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.
5	When you are ready to do backup (usually after your application has been running for some time in your production environment), save the contents of host file system buffers to disk using the command provided by your host operating system, or by using SnapDrive for Windows or UNIX systems.
6	Create a snapshot by entering the following command: <pre>snap create volume_name snapshot_name</pre> Example: <code>snap create vol1 payroll_backup</code>
7	Enter the following command to create a clone of the production LUN: <pre>lun clone create clone_lunpath -b parent_lunpath parent_snap</pre> Example: <code>lun clone create /vol/vol1/qtrees_1/payroll_lun_clone -b /vol/vol1/qtrees_1/payroll_lun payroll_backup</code>
8	Create an igroup that includes the WWPN of the backup server. <pre>igroup create -f [-t ostype] group [node ...]</pre> Example: <code>igroup create -f -t windows backup_server 10:00:00:00:d3:6d:0f:e1</code> Result: Data ONTAP creates an igroup that includes the WWPN (10:00:00:00:d3:6d:0f:e1) of the Windows backup server.

Step	Action
9	Enter the following command to map the LUN clone you created in Step 7 to the backup host: <pre>lun map lun_path initiator-group LUN_ID</pre> Example: lun map /vol/vol1/qtrees_1/payroll_lun_clone backup_server 1 Result: Data ONTAP maps the LUN clone (/vol/vol1/qtrees_1/payroll_lun_clone) to the igroup called backup_server with a SCSI ID of 1.
10	From the host, discover the new LUN, format it, 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.
11	Back up the data in the LUN clone from the backup host to tape by using your SAN backup application.
12	Take the LUN clone offline by entering the following command: <pre>lun offline /vol/vol_name/qtrees_name/lun_name</pre> Example: lun offline /vol/vol1/qtrees_1/payroll_lun_clone
13	Remove the LUN clone by entering the following command: <pre>lun destroy lun_path</pre> Example: lun destroy /vol/vol1/qtrees_1/payroll_lun_clone
14	Remove the snapshot by entering the following command: <pre>snap delete volume_name lun_name</pre> 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 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 that 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 Data Protection Online Backup and Recovery Guide*.

Using volume copy

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

The `vol copy` command enables you to copy data from one WAFL volume to another, either within the same filer or to a different filer. The result of the `vol copy` command is a restricted volume containing the same data that was on the source filer at the time you initiate 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> -S copies all snapshots in the source volume to the destination volume. If the source volume has snapshot-backed LUNs, you must use the -S option to ensure that the snapshots are copied to the destination volume. Note If the copying takes place between two filers, you can enter the <code>vol copy start</code> command on either the source or destination filer. You cannot, however, enter the command on a third filer that does not contain the source or destination volume. Example: <code>vol copy start -S /vol/vol0 filerB:/vol/vol1</code>

Cloning flexible volumes

What clones volumes are

A clone volume is a writable, point-in-time copy of a parent flexible volume. Clone volumes reside in the same aggregate as their parent volume. Changes made to the parent volume after the clone is created are not inherited by the clone.

Because clone volumes and parent volumes share the same disk space for any data common to both, creating a clone is instantaneous and requires no additional disk space. You can split the clone from its parent if you do not want the clone and parent to share disk space.

Clone volumes are fully functional volumes; you manage them using the `vol` command, just as you do the parent volume. Clone volumes themselves can be cloned.

Reasons to clone flexible volumes

You can clone flexible volumes when you want a writable, point-in-time copy of a flexible volume. For example, you might want to clone flexible volumes in the following scenarios:

- ◆ You need to create a temporary copy of a volume for testing or staging purposes.
- ◆ You want to create multiple copies of data for additional users without giving them access to production data.
- ◆ You want to copy a database for manipulation or projection operations without altering the original data.

How clone volumes affect LUNs

When you create a clone volume, LUNs in the parent volume are present in the clone but they are not mapped and they are offline. To bring the LUNs in the clone online, you must map them to igroups. When the LUNs in the parent volume are backed by snapshots, the clone also inherits the snapshots.

You can also clone individual LUNs. If the parent volume has LUN clones, the clone volume inherits the LUN clones. A LUN clone has a base snapshot, which is also inherited by the volume clone. The LUN clone's base snapshot in the parent volume shares blocks with the LUN clone's base snapshot in the volume clone. You cannot delete the LUN clone's base snapshot in the parent volume until you delete the base snapshot in the volume clone.

How volume cloning affects space reservation

Volume-level guarantees: Clone volumes inherit the same volume-level space guarantee setting as the parent volume, but the space guarantee is disabled for the clone. This means that the containing aggregate does not ensure that space is always available for write operations to the clone volume, regardless of the clone's guarantee setting.

The following example shows guarantee settings for two volumes: a parent volume called `testvol` and its clone, `testvol_c`. For `testvol` the guarantee option is set to `volume`. For `testvol_c`, the guarantee option is set to `volume`, but the guarantee is disabled.

```
filer_1> vol options testvol
nosnap=off, nosnapdir=off, minra=off, no_atime_update=off,
nvfail=off, snapmirrored=off, create_ucose=off, convert_ucose=off,
maxdirsize=5242, fs_size_fixed=off, guarantee=volume,
svo_enable=off, svo_checksum=off, svo_allow_rman=off,
svo_reject_errors=off, fractional_reserve=100

filer_1> vol status testvol_c
Volume State      Status      Options
cl      online    raid_dp, flex  maxdirsize=5242,
                                     guarantee=volume (disabled)
Clone, backed by volume 'testvol', snapshot 'hourly.0'
Containing aggregate: 'a1'
```

Volume-level space guarantees are enabled on the clone volume only after you split the clone volume from its parent. After the clone-splitting process, space guarantees are enabled for the clone volume, but the guarantees are enforced only if there is enough space in the containing aggregate.

Space reservation and fractional overwrite reserve: LUNs in clone volumes inherit the space reservation setting from the LUNs in the parent volume. This means if space reservation is enabled for a LUN in the parent volume, it is also enabled for the LUN in the clone volume. Clone volumes inherit fractional overwrite reserve settings from the parent volume. For example, if fractional overwrite is set to 50 percent on the parent volume, it is also set to 50 percent on the clone volume. Space reservation and fractional overwrite reserve settings are enabled, but they are enforced only if there is enough space in the containing aggregate.

Commands for cloning flexible volumes

You use the following commands to clone flexible volumes:

- ◆ `vol clone create`—creates a volume clone and a base snapshot of the parent volume.

- ◆ `vol clone split`—splits the volume clone from the parent so that they no longer share data blocks.

Cloning a flexible volume

To clone a flexible volume, complete the following steps.

Step	Action
1	Enter the following command to clone the volume: <pre>vol clone create cl_vol_name [-s {volume file none}] -b f_p_vol_name [parent_snap]</pre> <i>cl_vol_name</i> is the name of the clone volume that you want to create. <code>-s {volume file none}</code> specifies the space guarantee for the new volume clone. If no value is specified, the clone is given the same space guarantee setting as its parent. For more information, see “How volume cloning affects space reservation” on page 134. Note For Data ONTAP 7.0, space guarantees are disabled for clone volumes until they are split from the parent volume. <i>f_p_vol_name</i> is the name of the flexible parent volume that you intend to clone. <i>[parent_snap]</i> is the name of the base snapshot of the parent volume. If no name is specified, Data ONTAP creates a base snapshot with the name <code>clone_cl_name_prefix.id</code>, where <i>cl_name_prefix</i> is the name of the new clone volume (up to 16 characters) and <i>id</i> is a unique digit identifier (for example 1,2, etc.). The base snapshot cannot be deleted as long as the parent volume or any of its clones exists. Example snapshot name: To create a clone <i>newclone</i> of the volume named <i>flexvol1</i>, the following command is entered: <pre>vol clone create newclone -b flexvol1</pre> The snapshot created by Data ONTAP is named <code>clone_newclone.1</code>.

Step	Action
2	Verify the success of the clone creation by entering the following command: <pre>vol status -v <i>cl_vol_name</i></pre>

Splitting a cloned volume

You might want to split your cloned volume into two independent volumes that occupy their own disk space.

Note

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

If you take the clone offline while the splitting operation is in progress, the operation is suspended; when you bring the clone back online, the splitting operation resumes.

To split a clone from its parent volume, complete the following steps.

Step	Action
1	Verify that enough additional disk space exists in the containing aggregate to support the clone and its parent volume unsharing their shared disk space by entering the following command: <pre>df -A <i>aggr_name</i></pre> <i>aggr_name</i> is the name of the containing aggregate of the flexible volume clone that you want to split. The <code>avail</code> column tells you how much available space you have in your aggregate. When a volume clone is split from its parent, the resulting two flexible volumes occupy completely different blocks within the same aggregate.

Step	Action
2	Enter the following command to split the volume: <pre>vol clone split start cl_vol_name</pre> <i>cl_vol_name</i> is the name of the clone that you want to split from its parent. The original volume and its clone begin to split apart, unsharing the blocks that they formerly shared.
3	If you want to check the status of a clone-splitting operation, enter the following command: <pre>vol clone status cl_vol_name</pre>
4	If you want to stop the progress of an ongoing clone-splitting operation, enter the following command: <pre>vol clone stop cl_vol_name</pre> The clone-splitting operation halts; the original and clone volumes will remain clone partners, but the disk space that was duplicated up to that point will remain duplicated.
5	Display status for the newly split volume to verify the success of the clone-splitting operation by entering the following command: <pre>vol status -v cl_vol_name</pre>

For detailed information

For detailed information about volume cloning, including limitations of volume cloning, see the *Data ONTAP Storage Management Guide*.

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 68).

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 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: <code>vol options volume-name nvfail on</code>

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: <code>/vol/vol1/home/dbs/oracle-WG73.dbf</code>
3	Save the file.

Using SnapValidator

What the SnapValidator feature does

Oracle Hardware Assistant Resilient Data (H.A.R.D.) is a system of checks embedded in Oracle data blocks that enable a storage system to validate write operations to an Oracle database. The SnapValidator™ feature implements Oracle H.A.R.D. checks to detect and reject invalid Oracle data before it is written to the filer.

Note

SnapValidator is not based on SnapShot technology.

H.A.R.D. checks that SnapValidator implements

SnapValidator implements the following Oracle H.A.R.D. validations:

- ◆ Checks for writes of corrupted datafile blocks. This includes the checksum value and validation of selected fields in the block.
- ◆ Checks for writes of corrupted redo log blocks. This includes the checksum value and validation of selected fields in the block.
- ◆ Checks for writes of corrupted controlfile blocks. This includes the checksum value and validation of selected fields in the block.
- ◆ Verifies that writes of Oracle data are multiples of a valid Oracle blocksize for the target device.

When to use SnapValidator

You use SnapValidator if you have existing Oracle database files or LUNs on a filer or if you want to store a new Oracle database on the filer.

Supported protocols

SnapValidator checks are supported for the following protocols:

- ◆ LUNs accessed by FCP or iSCSI protocols
- ◆ Files accessed by NFS

Guidelines for preparing a database for SnapValidator

You prepare database files or LUNs for SnapValidator checks by using the following guidelines.

1. Make sure you are working in your test environment, not your production environment.
2. Make sure the Oracle data files or LUNs are in single volume.
3. Do not put the following types of files in the same volume as the Oracle data:
 - ❖ Oracle configuration files
 - ❖ Files or LUNs that are not Oracle-owned (for example, scripts or text files)

For an existing database, you might have to move configuration files and other non-Oracle data to another virtual volume.

4. If you are using new LUNs for Oracle data, and the LUN is accessed by non-Windows hosts, set the LUN Operating System type (ostype) to image. If the LUNs are accessed by Windows hosts, the ostype must be windows. LUNs in an existing database can be used, regardless of their ostype. For more information about LUN Operating System types, see [“Creating LUNs, igroups, and LUN maps”](#) on page 45.
5. Make sure Oracle H.A.R.D. checks are enabled on the host running the Oracle application server. You enable H.A.R.D. checks by setting the `db_block_checksum` value in the `init.ora` file to `true`.

Example: `db_block_checksum=true`

6. License SnapValidator. For more information, see [“Licensing SnapValidator”](#) on page 142.
7. Enable SnapValidator checks on your volumes. For more information, see [“Enabling SnapValidator checks on volumes”](#) on page 143.

Make sure you set SnapValidator to return an error log to the host and filer consoles for all invalid operations by entering the following command:

```
vol options volume-name svo_reject_errors off
```

8. Test your environment by writing data to the filer.
 9. Set SnapValidator to reject invalid operations and return an error log to the host and filer consoles for all invalid operations by entering the following command:
- ```
vol options volume-name svo_reject_errors on
```
10. Put your database into production.

## Tasks for implementing SnapValidator checks

After you prepare the database, you implement SnapValidator checks by completing the following tasks on the filer:

- ◆ License SnapValidator  
For detailed information, see “[Licensing SnapValidator](#)” on page 142.
- ◆ Enable SnapValidator checks on the volume that contains the Oracle data.  
For detailed information, see “[Enabling SnapValidator checks on volumes](#)” on page 143.
- ◆ If you are using LUNs for Oracle data, configure the disk offset for each LUN in the volume to enable SnapValidator checks on those LUNs.  
For detailed information, see “[Enabling SnapValidator checks on LUNs](#)” on page 144.

## Licensing SnapValidator

To license SnapValidator complete the following steps:

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                           |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Verify whether SnapValidator is licensed by entering the following command:<br><br><b>license</b><br><br><b>Result:</b> A list of all available services appears. Services that are enabled show the license code. Services that are not enabled are indicated as “not licensed.” For example, the following line indicates that SnapValidator is not licensed.<br><br>snapvalidator not licensed |                                                                                                                                                                                           |
| 2    | <b>If SnapValidator is...</b>                                                                                                                                                                                                                                                                                                                                                                     | <b>Then...</b>                                                                                                                                                                            |
|      | Licensed                                                                                                                                                                                                                                                                                                                                                                                          | Proceed to “ <a href="#">Enabling SnapValidator checks on volumes</a> ” on page 143.                                                                                                      |
|      | Not licensed                                                                                                                                                                                                                                                                                                                                                                                      | Enter the following command:<br><br><b>license add license_code</b><br><br><i>license_code</i> is the license code you received from NetApp when you purchased the SnapValidator license. |

## Enabling SnapValidator checks on volumes

You enable SnapValidator checks at the volume level. To enable SnapValidator checks on a volume, complete the following steps:

### Note

You cannot enable SnapValidator on the root volume.

| Step | Action                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the filer command line, enable SnapValidator by entering the following command:</p> <pre>vol options volume-name svo_enable on</pre> <p><b>Result:</b> All SnapValidator checks are enabled on the volume, with the exception of checksums.</p> |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 2    | <b>If you want to...</b>                                                                                                                                                                                                                              | <b>Then enter the following command:</b>                                                                                                                                                                                                                                                                                                                                                                                                          |
|      | Enable data checksumming on the volume                                                                                                                                                                                                                | <pre>vol options volume-name svo_checksum on</pre>                                                                                                                                                                                                                                                                                                                                                                                                |
|      | Disable block number checks because the volume contains Oracle Recovery Manager (RMAN) backup data.                                                                                                                                                   | <pre>vol options volume-name svo_allow_rman on</pre>                                                                                                                                                                                                                                                                                                                                                                                              |
|      | Set SnapValidator to return an error log to the host and filer consoles for all invalid operations. You might want to do this when you are testing SnapValidator before you put your database into production.                                        | <pre>vol options volume-name svo_reject_errors off</pre> <p>When you set this option to off, SnapValidator only logs errors but does not reject invalid operations.</p>                                                                                                                                                                                                                                                                           |
|      | Set SnapValidator to reject all invalid operations and return an error log to the host and filer consoles.                                                                                                                                            | <pre>vol options volume-name svo_reject_errors on</pre> <p>If this option is not set to on, then SnapValidator detects invalid operations but only logs them as errors. The following shows a SnapValidator error example displayed on the filer console:</p> <pre>Thu May 20 08:57:08 GMT [filer_1: wafl.svo.checkFailed:error]: SnapValidator: Validation error Bad Block Number:: v:9r2 vol:flextest inode:98 length:512 Offset: 1298432</pre> |

| Step | Action                                                                                                               |
|------|----------------------------------------------------------------------------------------------------------------------|
| 3    | If the volume contains LUNs, proceed to <a href="#">“Enabling SnapValidator checks on LUNs”</a> in the next section. |

## Enabling SnapValidator checks on LUNs

If you enable SnapValidator on volumes that contain database LUNs, you must also enable SnapValidator checks on the LUNs by defining the offset to the Oracle data on each LUN. The offset separates the Oracle data portion of the LUN from the host volume manager’s disk label or partition information. The value for the offset depends on the Operating System (OS) of the host accessing the data on the LUN. By defining the offset for each LUN, you ensure that SnapValidator does not check write operations to the disk label or partition areas as if they were Oracle write operations.

To define the offset, you must first identify the offset on your host and then define that offset to the filer. The method you use to identify the offset depends on your host. For details see:

- ◆ [“Identifying the disk offset for Solaris hosts”](#) on page 145
- ◆ [“Identifying the disk offset for other hosts”](#)
- ◆ [“Defining the disk offset on the filer”](#)

**Identifying the disk offset for Solaris hosts:** To identify the disk offset for Solaris hosts, complete the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the host, enter the following command:</p> <pre>prtvtoc /dev/rdisk/device_name</pre> <p><b>Result:</b> The host console displays a partition map for the disk.</p> <p><b>Example:</b> The following output example shows the partition map for disk c3t9d1s2:</p> <pre>prtvtoc /dev/rdisk/c3t9d1s2 * /dev/rdisk/c3t9d1s2 partition map * * Dimensions: *   512 bytes/sector *   384 sectors/track *   16 tracks/cylinder *   6144 sectors/cylinder *   5462 cylinders *   5460 accessible cylinders * * Flags: *   1: unmountable *   10: read-only * * * Partition  Tag  Flags      First      Sector      Last * Partition  Tag  Flags      Sector      Count      Sector  Mount Directory *   0         0    00           0        6144        6143 *   2         5    01           0  33546240  33546239 *   6         0    00        6144  33540096  33546239</pre> |
| 2    | <p>Obtain the offset value by multiplying the value of the first sector of partition 6 by the bytes/sector value listed under Dimensions. In the example shown in <a href="#">Step 1</a>, the disk offset is <math>6144 * 512 = 3145728</math>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Identifying the disk offset for other hosts:** To identify the disk offset for non-Solaris hosts, complete the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                              |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the host console, enter the following command:</p> <pre>dd if=/dev/zero of=/dev/rdisk/device_name bs=4096 count=1 conv=notrunc</pre> <p><i>device_name</i> is the name of the device—for example c0t0d3s6. Use slice 6 of the device.</p> <p><b>Result:</b> The host writes an Oracle 4K block of zeros to the filer.</p>                                                                                     |
| 2    | <p>Check the SnapValidator error message displayed on the filer console. The error message displays the offset.</p> <p><b>Example:</b> The following error message example shows that the disk offset is 1048576 bytes.</p> <pre>filerA&gt; Thu Mar 10 16:26:01 EST [filerA:waf1.svo.checkFailed:error]: SnapValidator: Validation error Zero Data:: v:9r2 vol:test inode:3184174 length:4096 Offset: 1048576</pre> |

**Defining the disk offset on the filer:** To define the disk offset on the filer, complete the following step.

| Step | Action                                                                                                                                                                                                                                            |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Use the volume manager tools for your host OS to obtain the value of the offset. For detailed information about obtaining the offset, see the vendor-supplied documentation for your volume manager.</p>                                       |
| 2    | <p>On the filer command line, enter the following command:</p> <pre>lun set svo_offset lun_path offset</pre> <p><i>offset</i> is specified in bytes, with an optional multiplier suffix: c(1), w(2), b(512), k(1024), m(k*k), g(k*m), t(m*m).</p> |

## Disabling SnapValidator on a volume

To disable SnapValidator, complete the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the filer command line, enter the following command:</p> <pre>vol options volume-name svo_enable off</pre> <p><b>Result:</b> SnapValidator does not check Oracle write operations to files or LUNs. The settings for each type of check (for example, checksumming) are not disabled. If you re-enable SnapValidator, the settings for each type of check are saved.</p>                                                                                                                                                                                                                |
| 2    | <p>To disable a specific SnapValidator option, enter the following command:</p> <pre>vol options volume-name option off</pre> <p><i>option</i> is one of the following:</p> <ul style="list-style-type: none"><li>◆ <i>svo_checksum</i>—disables data checksumming on the volume.</li><li>◆ <i>svo_allow_rman</i>—allows block number checks on the volume. You disable this option (set it to off) if the volume does not contain RMAN data.</li><li>◆ <i>svo_reject_errors</i>—detects invalid operations but does not reject them. Invalid operations are only logged as errors.</li></ul> |

## Disabling SnapValidator checks on a LUN

To disable SnapValidator checks on a LUN, complete the following step:

| Step | Action                                                                                                        |
|------|---------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the filer command line, enter the following command:</p> <pre>lun set lun_path svo_offset disable</pre> |

## How SnapValidator checks are set for upgrades and reverts

When you upgrade to Data ONTAP 7.0 from a previous release, all SnapValidator options on all volumes are disabled. The offset attribute (the *svo\_offset* option) for LUNs is also disabled.

When you revert to a previous version of Data ONTAP, all SnapValidator options on all volumes are disabled. The value for the LUN offset is retained, but the earlier version of Data ONTAP does not apply it.

## SnapValidator error messages

**When write operations to LUNs fail:** SnapValidator displays two messages similar to the following when write operations to a LUN fail:

- ◆ The first message is generated by SnapValidator and indicates that the filer detected invalid data. The error message does not show the full path to the LUN. The following is an example error message:  

```
Thu May 20 08:57:08 GMT [fas940: waf1.svo.checkFailed:error]:
SnapValidator: Validation error Bad Block Number:: v:9r2
vol:dbtest inode:98 length:512 Offset: 1298432
```
- ◆ The second error message is a scsitarget.write error, which shows the full path to the LUN. The following is an example error message that indicates a write to a specific LUN failed:

```
Thu May 20 14:19:00 GMT [fas940:
scsitarget.write.failure:error]: Write to LUN
/vol/dbtest/oracle_lun1 failed (5)
```

If you receive a message indicating that a write operation to a LUN failed, verify that you set the correct disk offset on the LUN. Identify the disk offset and reset the offset defined for the LUN by using the procedures described in [“Enabling SnapValidator checks on LUNs”](#) on page 144.

**Other invalid data error messages:** The following messages indicate that SnapValidator detected invalid data:

- ◆ Checksum Error
- ◆ Bad Block Number
- ◆ Bad Magic Number
- ◆ No Valid Block Size
- ◆ Invalid Length for Log Write
- ◆ Zero Data
- ◆ Ones Data
- ◆ Write length is not aligned to a valid block size
- ◆ Write offset is not aligned to a valid block size

If you receive a message indicating that SnapValidator detected or rejected invalid data, verify the following:

1. You enabled the SnapValidator checks on the volumes that contain your data files. For more information, see [“Enabling SnapValidator checks on volumes”](#) on page 143.
2. You set the SnapValidator checks correctly. For example, if you set the `svo_allow_rman` volume option to `on`, then make sure that the volume contains Oracle Recovery Manager (RMAN) backup data. If you store

RMAN data in a volume that does not have this option set, then you might receive an error message indicating that SnapValidator detected invalid data.

If the SnapValidator checks are enabled and the options on the filer are correctly set but you still receive the above errors, you might have the following problems:

- ◆ Your host is writing invalid data to the filer. Consult your database administrator to check Oracle configuration on the host.
- ◆ You might have a problem with network connectivity or configuration. Consult your system administrator to check the network path between your host and filer.



**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 152
- ◆ [“Managing the FCP service on systems with onboard ports”](#) on page 156
- ◆ [“Displaying information about HBAs”](#) on page 160

# 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    | <p>Enter the following command:</p> <p><code>fc status</code></p> <p><b>Result:</b> A message is displayed indicating whether FCP service is running.</p> <p><b>Note</b>_____</p> <p>If the FCP service is not running, verify that the FCP license is enabled, and start the FCP service.</p> _____ |

## Verifying that the FCP service is licensed

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

| Step | Action                                                                                                                                                                                                                                        |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>license</pre> <p><b>Result:</b> 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.”</p> |

## Enabling the FCP service

To enable the FCP service, complete the following step.

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

**For FAS270 appliances:** After you license the FCP service on an FAS270 appliance, you must reboot. When the appliance boots up, the orange port labeled Fibre Channel C 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. For detailed information, see “[Managing the FCP service on systems with onboard ports](#)” on page 156.

## Starting and stopping FCP service

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

| Step | Action                                                                                                                                                                                                                                  |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>fcpx {start stop}</pre> <p><b>Example:</b></p> <pre>fcpx start</pre> <p><b>Result:</b> The FCP service begins running. If you enter <code>fcpx stop</code>, the FCP service stops running.</p> |

## 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    | <p>Enter the following command:</p> <pre>fcpx config adapter [up down]</pre> <p><b>Example:</b></p> <pre>fcpx config 4a down</pre> <p><b>Result:</b> The target HBA 4a is offline. If you enter <code>fcpx config 4a up</code>, the target HBA is brought online.</p> |

## Disabling the FCP license

To disable the FCP license, complete the following step.

| Step | Action                                                                                                                                                                            |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>license delete service</pre> <p><i>service</i> is any service you can license.</p> <p><b>Example:</b></p> <pre>license delete fcpx</pre> |

## 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    | <p>Enter the following command:</p> <pre><b>fcg nodename <i>nodename</i></b></pre> <p><i>nodename</i> is a 64-bit WWNN address</p> <p><b>Example:</b> <code>fcg nodename 50:a9:80:00:02:00:8d:ff</code></p> |

# Managing the FCP service on systems with onboard ports

---

## FAS270 storage appliances

FAS270 unit provides two independent Fibre Channel ports identified as Fibre Channel B (with a blue label) and Fibre Channel C (with an orange label):


- ◆ You use the Fibre Channel B port to communicate to internal and external disks.
- ◆ You can configure the Fibre Channel C port in one of two modes:
  - ❖ You use initiator mode to communicate with tape backup devices such as in a TapeSAN backup configuration.
  - ❖ You use target mode to communicate with SAN hosts or a front end SAN switch.

The Fibre Channel C port does not support mixed initiator/target mode. The default mode for this port is initiator mode. If you want to license the FCP service and connect the FAS270 to a SAN, you have to configure this port to operate in SAN target mode.

## FAS270 cluster configuration example

FAS270 cluster configurations must be cabled to switches that support public loop topology. To connect a FAS270 cluster to a fabric topology that includes switches that only support point-to-point topology, such as McDATA Director class switches, you must connect the cluster to an edge switch and use this switch as a bridge to the fabric.

The following figure shows an example configuration in which a multi-attached host accesses a FAS270 cluster. For information about specific switch models supported and fabric configuration guidelines, see the online *NetApp Fibre Channel Configuration Guide* at the following URL:  
[http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/QuickRef/FCPCConfigurationGuide.pdf](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPCConfigurationGuide.pdf).



## Configuring the Fibre Channel port for target mode

After you cable your configuration and enable the cluster, configure port Fibre Channel C for target mode by completing the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>If the FCP protocol is not licensed, install the license by entering the following command:</p> <pre>license add <i>FCP_code</i></pre> <p><i>FCP_code</i> is the FCP service license code provided to you by NetApp.</p> <p><b>Example:</b></p> <pre>fas270a&gt; license add XXXXXXXX A fcp site license has been installed. cf.takeover.on_panic is changed to on Run 'fcp start' to start the FCP service. Also run 'lun setup' if necessary to configure LUNs. A reboot is required for FCP service to become available.       FCP enabled. fas270a&gt; Fri Dec  5 14:54:24 EST [fas270a: rc:notice]: fcp licensed</pre> |
| 2    | <p>Reboot the FAS270 by entering the following command:</p> <pre>reboot</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Verify that the Fibre Channel C port is in target mode by entering the following command:</p> <pre>sysconfig</pre> <p><b>Example:</b></p> <pre>fas270a&gt; sysconfig NetApp Release R6.5xN_031130_2230: Mon Dec  1 00:07:33 PST 2003 System ID: 0084166059 (fas270a) System Serial Number: 123456 (fas270a) slot 0: System Board     Processors:          2     Processor revision: B2     Processor type:      1250     Memory Size:         1022 MB slot 0: FC Host Adapter 0b     14 Disks:             952.0GB     1 shelf with EFH slot 0: <b>Fibre Channel Target Host Adapter 0c</b> slot 0: SB1250-Gigabit Dual Ethernet Controller     e0a MAC Address:     00:a0:98:01:29:cd (100tx-fd-up)     e0b MAC Address:     00:a0:98:01:29:ce (auto-unknown-cfg_down) slot 0: NetApp ATA/IDE Adapter 0a (0x000000000000001f0)     0a.0                  245MB</pre> <p><b>Note</b>_____</p> <p>The Fibre Channel C port is identified as Fibre Channel Target Host Adapter 0c.</p> |
| 4    | <p>Start the FCP service by entering the following command:</p> <pre>fcp start</pre> <p><b>Example:</b></p> <pre>fas270a&gt; fcp start FCP service is running. Wed Sep 17 15:17:04 GMT [fas270a: fcp.service.startup:info]: FCP service startup</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## Displaying information about HBAs

---

### How to display HBA information

The following table lists the commands available for displaying information about HBAs. The output varies depending on the FCP cfmodesetting and the filer model.

| If you want to display...                                                                                                                                                                       | Use this command...                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Information for all adapters in the system, 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] [ adapter ]</code><br><i>adapter</i> is a numerical value only, for example, 5.<br>-v displays additional information about all adapters.                                                                   |
| Disks, disk loops, and options configuration information that affects coredumps and takeover                                                                                                    | <code>sysconfig -c</code>                                                                                                                                                                                                        |
| FCP cfmodesetting                                                                                                                                                                               | <code>fcv show cfmodesetting</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] [adapter&amp;portnumber]</code><br>-v displays the Fibre Channel host address of the initiator.<br><i>adapter&amp;portnumber</i> is the slot number with the port number, a or b; for example, 5a. |
| Service statistics                                                                                                                                                                              | <code>availtime</code>                                                                                                                                                                                                           |
| Target HBAs configuration information                                                                                                                                                           | <code>fcv config</code>                                                                                                                                                                                                          |

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

## Displaying information about all adapters

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>At the filer, enter the following command to see information about all adapters.</p> <pre>sysconfig -v</pre> <p><b>Result:</b> 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.</p> <p><b>Note</b>_____</p> <p>In the output, 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.</p> <p>_____</p> <p><b>Note</b>_____</p> <p>The output varies according to filer model. For example, if you have a FAS270, the target port is displayed as slot 0: Fibre Channel Target Host Adapter 0c.</p> <p>_____</p> <p><b>Example:</b> A partial display of information about a target HBA installed in slot 7 appears as follows:</p> <pre>slot 7: Fibre Channel Target Host Adapter 7a       (Dual-channel, QLogic 2312 (2352) rev. 2, 64-bit, &lt;ONLINE&gt;)       Firmware rev:   3.2.18       Host Port Addr: 170900       Cacheline size: 8       SRAM parity:    Yes       FC Nodename:    50:0a:09:80:86:87:a5:09 (500a09808687a509)       FC Portname:    50:0a:09:83:86:87:a5:09 (500a09838687a509)       Connection:     PTP, Fabric  slot 7: Fibre Channel Target Host Adapter 7b       (Dual-channel, QLogic 2312 (2352) rev. 2, 64-bit, &lt;ONLINE&gt;)       Firmware rev:   3.2.18       Host Port Addr: 171800       Cacheline size: 8       SRAM parity:    Yes       FC Nodename:    50:0a:09:80:86:57:11:22 (500a098086571122)       FC Portname:    50:0a:09:8c:86:57:11:22 (500a098c86571122)       Connection:     PTP, Fabric</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.

The output of the `fcv config` command also depends on the filer's `cfmode` setting. For examples, see [“How Data ONTAP displays information about target ports”](#) on page 10.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>At the filer, enter the following command.</p> <pre>fcv config</pre> <p><b>Sample output:</b></p> <pre>7a:  ONLINE &lt;ADAPTER UP&gt;  PTP  Fabric       host address 170900       portname 50:0a:09:83:86:87:a5:09  nodename 50:0a:09:80:86:87:a5:09       mediatype ptp  partner adapter 7a  7b:  ONLINE &lt;ADAPTER UP&gt;  PTP  Fabric       host address 171800       portname 50:0a:09:8c:86:57:11:22  nodename 50:0a:09:80:86:57:11:22       mediatype ptp  partner adapter 7b</pre> <p><b>Sample output for FAS270:</b> For the FAS270, the <code>fcv config</code> command displays the target virtual local, standby, and partner ports.</p> <pre>0c:  ONLINE &lt;ADAPTER UP&gt;  Loop  Fabric       host address 0100da       portname 50:0a:09:81:85:c4:45:88  nodename 50:0a:09:80:85:c4:45:88       mediatype loop  partner adapter 0c  0c_0: ONLINE Local       portname 50:0a:09:81:85:c4:45:88  nodename 50:0a:09:80:85:c4:45:88       loopid 0x7 portid 0x0100da  0c_1: OFFLINED BY USER/SYSTEM Standby       portname 50:0a:09:81:85:c4:45:91  nodename 50:0a:09:80:85:c4:45:91       loopid 0x0 portid 0x000000  0c_2: ONLINE Partner       portname 50:0a:09:89:85:c4:45:91  nodename 50:0a:09:80:85:c4:45:91       loopid 0x9 portid 0x0100d6</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 and without the separating colons. 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    | <div>At the filer, enter the following command:<br/><br/><code>fcv show adapter</code><br/><br/><b>Sample output for F8xx or FAS9xx series filers:</b> The following sample output displays information for the HBA in slot 7:<br/><br/>Slot: 7a<br/>Description: Fibre Channel Target Adapter 7a (Dual-channel, QLogic 2312 (2352) rev. 2)<br/>Adapter Type: Local<br/>Status: ONLINE<br/>FC Nodename: 50:0a:09:80:86:87:a5:09 (500a09808687a509)<br/>FC Portname: 50:0a:09:83:86:87:a5:09 (500a09838687a509)<br/>Standby: No<br/><br/>Slot: 7b<br/>Description: Fibre Channel Target Adapter 7b (Dual-channel, QLogic 2312 (2352) rev. 2)<br/>Adapter Type: Partner<br/>Status: ONLINE<br/>FC Nodename: 50:0a:09:80:86:57:11:22 (500a098086571122)<br/>FC Portname: 50:0a:09:8c:86:57:11:22 (500a098c86571122)<br/>Standby: No<br/><br/><b>Note</b><br/>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.<br/><br/><b>Note</b><br/>For the FAS270, the <code>fcv show adapter</code> command displays the target virtual local (0c_0), standby (0c_1), and partner (0c_2) ports.</div> |

**Displaying initiator HBA information**      To display the port name and igroup name of initiator HBAs connected to target HBAs, complete the following step.

| Step                    | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-------|-------------------------|-----------|-------------------------|--------|-------------------------|-----------|-------------------------|---------|----------|-------|-------------------------|---------|-------------------------|-----------|-------------------------|--------|-------------------------|-----------|
| 1                       | <p>At the filer, enter the following command:</p> <pre><b>fcg show initiator</b></pre> <p><b>Result:</b> The following output is displayed:</p> <p>Initiators connected on adapter 7a:</p> <table> <tr> <td>Portname</td><td>Group</td></tr> <tr> <td>10:00:00:00:c9:39:4d:82</td><td>sunhost_1</td></tr> <tr> <td>50:06:0b:00:00:11:35:62</td><td>hphost</td></tr> <tr> <td>10:00:00:00:c9:34:05:0c</td><td>sunhost_2</td></tr> <tr> <td>10:00:00:00:c9:2f:89:41</td><td>aixhost</td></tr> </table> <p>Initiators connected on adapter 7b:</p> <table> <tr> <td>Portname</td><td>Group</td></tr> <tr> <td>10:00:00:00:c9:2f:89:41</td><td>aixhost</td></tr> <tr> <td>10:00:00:00:c9:39:4d:82</td><td>sunhost_1</td></tr> <tr> <td>50:06:0b:00:00:11:35:62</td><td>hphost</td></tr> <tr> <td>10:00:00:00:c9:34:05:0c</td><td>sunhost_2</td></tr> </table> | Portname | Group | 10:00:00:00:c9:39:4d:82 | sunhost_1 | 50:06:0b:00:00:11:35:62 | hphost | 10:00:00:00:c9:34:05:0c | sunhost_2 | 10:00:00:00:c9:2f:89:41 | aixhost | Portname | Group | 10:00:00:00:c9:2f:89:41 | aixhost | 10:00:00:00:c9:39:4d:82 | sunhost_1 | 50:06:0b:00:00:11:35:62 | hphost | 10:00:00:00:c9:34:05:0c | sunhost_2 |
| Portname                | Group                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:39:4d:82 | sunhost_1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 50:06:0b:00:00:11:35:62 | hphost                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:34:05:0c | sunhost_2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:2f:89:41 | aixhost                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| Portname                | Group                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:2f:89:41 | aixhost                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:39:4d:82 | sunhost_1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 50:06:0b:00:00:11:35:62 | hphost                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |
| 10:00:00:00:c9:34:05:0c | sunhost_2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |       |                         |           |                         |        |                         |           |                         |         |          |       |                         |         |                         |           |                         |        |                         |           |

Displaying statistics

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <div>Enter the following command:</div> <div><code>fcpx stats -i interval [ -c count ] [ -a   adapter ]</code></div> <div>-i interval is the interval, in seconds, at which the statistics are displayed.</div> <div>-c count is the number of intervals. For example, the <code>lun stats -i 10 -c 5</code> command displays statistics in ten-second intervals, for five intervals.</div> <div>-a shows statistics for all adapters.</div> <div>adapter is the slot and port number of a specific target HBA.</div> <div><b>Example output:</b></div> <div><pre>fcpx stats -i 1 r/s   w/s   o/s   ki/s   ko/s   asvc_t   qlen hba   0    0    0     0     0    0.00    0.00 7a 110  113    0  7104  12120   9.64    1.05 7a 146   68    0  6240  13488  10.28    1.05 7a 106   92    0  5856  10716  12.26    1.06 7a 136  102    0  7696  13964   8.65    1.05 7a</pre></div> <div><b>Explanation of output:</b> Each column displays the following information:</div> <div>r/s—The number of SCSI read operations per second.</div> <div>w/s—The number of SCSI write operations per second.</div> <div>o/s—The number of other SCSI operations per second.</div> <div>ki/s— Kilobytes per second of received traffic</div> <div>ko/s—Kilobytes per second send traffic.</div> <div>asvc_t—Average time in milliseconds to process a request</div> <div>qlen—The average number of outstanding requests pending.</div> <div>hba—The HBA slot and port number.</div> |

## 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:<br><br><b>sysstat -f</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | <b>Result:</b> The following output is displayed:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | <table><tr><th>CPU</th><th>NFS</th><th>CIFS</th><th>FCP</th><th>Net<br/>in</th><th>kB/s<br/>out</th><th>Disk<br/>read</th><th>kB/s<br/>write</th><th>FCP<br/>in</th><th>kB/s<br/>out</th><th>Cache<br/>age</th></tr><tr><td>81%</td><td>0</td><td>0</td><td>6600</td><td>0</td><td>0</td><td>105874</td><td>56233</td><td>40148</td><td>232749</td><td>1</td></tr><tr><td>78%</td><td>0</td><td>0</td><td>5750</td><td>0</td><td>0</td><td>110831</td><td>37875</td><td>36519</td><td>237349</td><td>1</td></tr><tr><td>78%</td><td>0</td><td>0</td><td>5755</td><td>0</td><td>0</td><td>111789</td><td>37830</td><td>36152</td><td>236970</td><td>1</td></tr><tr><td>80%</td><td>0</td><td>0</td><td>5732</td><td>0</td><td>0</td><td>111222</td><td>44512</td><td>35908</td><td>235412</td><td>1</td></tr><tr><td>81%</td><td>0</td><td>0</td><td>7061</td><td>0</td><td>0</td><td>107742</td><td>49539</td><td>42651</td><td>232778</td><td>1</td></tr><tr><td>78%</td><td>0</td><td>0</td><td>5770</td><td>0</td><td>0</td><td>110739</td><td>37901</td><td>35933</td><td>237980</td><td>1</td></tr><tr><td>79%</td><td>0</td><td>0</td><td>5693</td><td>0</td><td>0</td><td>108322</td><td>47070</td><td>36231</td><td>234670</td><td>1</td></tr><tr><td>79%</td><td>0</td><td>0</td><td>5725</td><td>0</td><td>0</td><td>108482</td><td>47161</td><td>36266</td><td>237828</td><td>1</td></tr><tr><td>79%</td><td>0</td><td>0</td><td>6991</td><td>0</td><td>0</td><td>107032</td><td>39465</td><td>41792</td><td>233754</td><td>1</td></tr><tr><td>80%</td><td>0</td><td>0</td><td>5945</td><td>0</td><td>0</td><td>110555</td><td>48778</td><td>36994</td><td>235568</td><td>1</td></tr><tr><td>78%</td><td>0</td><td>0</td><td>5914</td><td>0</td><td>0</td><td>107562</td><td>43830</td><td>37396</td><td>235538</td><td>1</td></tr></table> | CPU | NFS  | CIFS | FCP       | Net<br>in   | kB/s<br>out  | Disk<br>read  | kB/s<br>write | FCP<br>in   | kB/s<br>out  | Cache<br>age | 81% | 0 | 0 | 6600 | 0 | 0 | 105874 | 56233 | 40148 | 232749 | 1 | 78% | 0 | 0 | 5750 | 0 | 0 | 110831 | 37875 | 36519 | 237349 | 1 | 78% | 0 | 0 | 5755 | 0 | 0 | 111789 | 37830 | 36152 | 236970 | 1 | 80% | 0 | 0 | 5732 | 0 | 0 | 111222 | 44512 | 35908 | 235412 | 1 | 81% | 0 | 0 | 7061 | 0 | 0 | 107742 | 49539 | 42651 | 232778 | 1 | 78% | 0 | 0 | 5770 | 0 | 0 | 110739 | 37901 | 35933 | 237980 | 1 | 79% | 0 | 0 | 5693 | 0 | 0 | 108322 | 47070 | 36231 | 234670 | 1 | 79% | 0 | 0 | 5725 | 0 | 0 | 108482 | 47161 | 36266 | 237828 | 1 | 79% | 0 | 0 | 6991 | 0 | 0 | 107032 | 39465 | 41792 | 233754 | 1 | 80% | 0 | 0 | 5945 | 0 | 0 | 110555 | 48778 | 36994 | 235568 | 1 | 78% | 0 | 0 | 5914 | 0 | 0 | 107562 | 43830 | 37396 | 235538 | 1 |
|                                                                                                       | CPU                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | NFS | CIFS | FCP  | Net<br>in | kB/s<br>out | Disk<br>read | kB/s<br>write | FCP<br>in     | kB/s<br>out | Cache<br>age |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 81%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 6600 | 0         | 0           | 105874       | 56233         | 40148         | 232749      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 78%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5750 | 0         | 0           | 110831       | 37875         | 36519         | 237349      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 78%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5755 | 0         | 0           | 111789       | 37830         | 36152         | 236970      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 80%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5732 | 0         | 0           | 111222       | 44512         | 35908         | 235412      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 81%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 7061 | 0         | 0           | 107742       | 49539         | 42651         | 232778      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 78%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5770 | 0         | 0           | 110739       | 37901         | 35933         | 237980      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 79%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5693 | 0         | 0           | 108322       | 47070         | 36231         | 234670      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
|                                                                                                       | 79%                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 0   | 0    | 5725 | 0         | 0           | 108482       | 47161         | 36266         | 237828      | 1            |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| 79%                                                                                                   | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 0   | 6991 | 0    | 0         | 107032      | 39465        | 41792         | 233754        | 1           |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| 80%                                                                                                   | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 0   | 5945 | 0    | 0         | 110555      | 48778        | 36994         | 235568        | 1           |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| 78%                                                                                                   | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 0   | 5914 | 0    | 0         | 107562      | 43830        | 37396         | 235538        | 1           |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| <b>Explanation of FCP statistics:</b> The following columns provide information about FCP statistics. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| CPU—The percentage of the time that one or more CPUs were busy.                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| FCP—The number of FCP operations per second.                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |
| FCP kB/s—The number of kilobytes per second of incoming and outgoing FCP traffic.                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |     |      |      |           |             |              |               |               |             |              |              |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |     |   |   |      |   |   |        |       |       |        |   |

**Displaying information about traffic from the partner**

If you have a cluster and your filer’s cfmode setting is partner, mixed, or dual\_fabric, you might want to obtain information about the amount of traffic coming to the filer from its partner.

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                     |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>sysstat -b</pre> <p><b>Result:</b> The following columns display information about partner traffic:</p> <ul style="list-style-type: none"><li>◆ Partner—The number of partner operations per second</li><li>◆ Partner kB/s—The number of kilobytes per second of incoming and outgoing partner traffic.</li></ul> |

**Displaying how long FCP has been running**

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

| Step | Action                                                                                                                                                                                                                             |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>uptime</pre> <p><b>Result:</b> The following output is displayed:</p> <pre>12:46am up 2 days, 8:59 102 NFS ops, 2609 CIFS ops, 0 HTTP ops, 0 DAFS ops, 1933084 FCP ops, 0 iSCSI ops</pre> |

## Displaying FCP service statistics

To display FCP service statistics, complete the following step.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Enter the following command:</p> <pre>availtime</pre> <p><b>Result:</b> The following output is displayed:</p> <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 the HBA's WWNN

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

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



# Glossary

---

|                        |                                                                                                                                                                                                                       |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>client</b>          | A computer that shares files on a filer. See also host.                                                                                                                                                               |
| <b>FCP</b>             | 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.                                                                 |
| <b>filer</b>           | A NetApp storage appliance, including the R100, FAS200, FAS800, and FAS900 series models.                                                                                                                             |
| <b>HBA</b>             | Host bus adapter. An I/O adapter that connects a host I/O bus to a computer's memory system in SCSI environments.                                                                                                     |
| <b>host</b>            | Any computer system that accesses data on a filer as blocks using the FCP protocol, or is used to administer a filer.                                                                                                 |
| <b>igroup</b>          | Initiator group. A collection of unique identifiers, either FCP WWPNs in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs. |
| <b>initiator</b>       | The system component that originates an I/O command over an I/O bus or network.                                                                                                                                       |
| <b>initiator group</b> | See igroup.                                                                                                                                                                                                           |
| <b>LUN</b>             | A logical unit of storage.                                                                                                                                                                                            |

|                          |                                                                                                                                                                                                                                                               |          |           |        |       |  |  |                      |          |   |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-----------|--------|-------|--|--|----------------------|----------|---|
| <b>LUN clone</b>         | A complete copy of a LUN, which was initially created to be backed by a LUN in a snapshot. The clone creates a complete copy of the LUN and frees the snapshot, which you can then delete.                                                                    |          |           |        |       |  |  |                      |          |   |
| <b>LUN ID</b>            | The numerical identifier that the filer exports for a given LUN. The LUN ID is mapped to an igroup to enable host access.                                                                                                                                     |          |           |        |       |  |  |                      |          |   |
| <b>LUN path</b>          | <p>The path to a LUN on the filer. The following example shows a LUN path:</p> <table><tr><td>LUN path</td><td>Mapped to</td><td>LUN ID</td></tr><tr><td colspan="3">-----</td></tr><tr><td>/vol/vol01/fcpdb.lun</td><td>igroup_1</td><td>6</td></tr></table> | LUN path | Mapped to | LUN ID | ----- |  |  | /vol/vol01/fcpdb.lun | igroup_1 | 6 |
| LUN path                 | Mapped to                                                                                                                                                                                                                                                     | LUN ID   |           |        |       |  |  |                      |          |   |
| -----                    |                                                                                                                                                                                                                                                               |          |           |        |       |  |  |                      |          |   |
| /vol/vol01/fcpdb.lun     | igroup_1                                                                                                                                                                                                                                                      | 6        |           |        |       |  |  |                      |          |   |
| <b>LUN serial number</b> | The unique serial number for a LUN, as defined by the filer.                                                                                                                                                                                                  |          |           |        |       |  |  |                      |          |   |
| <b>map</b>               | 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.                       |          |           |        |       |  |  |                      |          |   |
| <b>online</b>            | Signifies that a LUN is exported to its mapped igroups. A LUN can be online only if it is enabled for read/write access.                                                                                                                                      |          |           |        |       |  |  |                      |          |   |
| <b>offline</b>           | Disables the export of the LUN to its mapped igroups. The LUN is not available to hosts.                                                                                                                                                                      |          |           |        |       |  |  |                      |          |   |
| <b>qtree</b>             | 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.                                                                                                                    |          |           |        |       |  |  |                      |          |   |
| <b>SAN</b>               | 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.                       |          |           |        |       |  |  |                      |          |   |

|                           |                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>share</b>              | 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.                                                                                                                                                                             |
| <b>space reservations</b> | 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.                                                                                                                  |
| <b>storage appliance</b>  | Hardware and software-based storage systems, such as filers, that serve and protect data using protocols for both SAN and NAS networks.                                                                                                                                                                                                                    |
| <b>target</b>             | 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.                                                                                                                                                                                                      |
| <b>volume</b>             | 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.                                                                                                                                                                                      |
| <b>WWN</b>                | 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).                                                                                |
| <b>WWNN</b>               | 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.    |
| <b>WWPN</b>               | 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 138

## A

adapters  
    displaying information about 160

## B

backup  
    data to tape 127  
    single LUNs to tape 128  
    tape, when to use 131

## C

clustered configurations  
    options required 8

## D

database 138  
database protection  
    using /etc/nvfail\_rename 138  
    using vol options nvfail 138  
dual\_fabric mode 10

## F

FAS270  
    ports, how displayed 12  
FAS270, dual-fabric mode 10  
FAS270, switch requirement 10  
FCP  
    cfmode setting 8  
    licensed service 5  
    nodes defined 5  
    nodes, filer 6  
    nodes, host 7  
    nodes, how connected 5  
    nodes, how identified 6

    nodes, switch 7  
FCP cfmode settings  
    affect on target ports 8  
FCP commands  
    fcv show initiator 61  
FCP service  
    displaying how long running 168  
    displaying traffic information about 167  
filer administration 2  
    using FilerView 3  
    using the command line 2  
filer, defined as target 2

## H

host bus adapters  
    displaying information about 167, 168  
    initiator, displaying information about 165

## I

igroup commands  
    igroup add 103  
    igroup create 62, 100  
    igroup destroy 102  
    igroup remove 103  
    igroup set 104  
    igroup show 103  
igroup show 102  
initiator group, creating 101  
initiator group, defined 6  
initiator groups  
    adding 103  
    creating 100  
    destroying 102  
    displaying contents of 103  
    removing 103  
    requirements for creation 49  
    setting the operating system type 104  
    unmapping LUNs from 67  
initiator host bus adapters, displaying information about 165

## L

lun commands

lun online 67

lun unmap 67

LUNs

accessing with NAS protocols 70

bringing online 67

defined 5

displaying reads, writes, and operations for 74

resizing restrictions 68

serial number 5

unmapping from initiator group 67

## M

man page command 3

mixed mode 9

## N

nodenames, of initiator host bus adapters,  
displaying 165

nvfail option, of vol options command 138

## P

partner mode 9

port resources, managing 7

portnames of initiator adapters, displaying 165

ports

used in clustered configurations 8

## R

restoring snapshots of LUNs 122

## S

sanlun fcp show adapter 101

Single File SnapRestore, using with LUNs 124

snap reserve, setting the percentage 40

snapshot schedule, turning off at the command line  
42

snapshots, using with SnapRestore 122

standby mode 8

## V

vol option nvfail, using with LUNs 138

volume commands

vol destroy (destroys an off-line volume) 135,  
136

volumes

destroying (vol destroy) 135, 136

## W

WWPN

creating igroups with 6

identifying filer ports with 6

WWPNs

how assigned 7