Data ONTAP® 7.1
Block Access Management Guide for iSCSI and 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 210-01094_B0
December 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 Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

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, Data ONTAP, 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. 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, 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 v

Chapter 1 **Introducing Block Access** 1

Understanding storage systems 2

Understanding how hosts connect to storage systems 5

Understanding how SnapDrive connects to storage 7

Related documents 9

Chapter 2 **How Data ONTAP Implements an iSCSI Network** 11

Changes for Data ONTAP 7.1 12

Understanding how Data ONTAP implements an iSCSI network 13

Setup Procedure Overview 19

Chapter 3 **How Data ONTAP Implements an FCP Network** 21

Understanding how Data ONTAP implements a Fibre Channel SAN 22

Chapter 4 **Configuring Storage** 27

Understanding storage units 28

Understanding space reservation for volumes and LUNs 30

Understanding how fractional reserve affects available space 33

- How 100 percent fractional reserve affects available space 34
- How reducing fractional reserve affects available space 40
- Reasons to set fractional reserve to zero 44

How guarantees on FlexVol volumes affect fractional reserve 45

Calculating the size of a volume 48

Guidelines for creating volumes that contain LUNs 53

Creating LUNs, igroups, and LUN maps 57

- Creating LUNs with the lun setup program 65
- Creating LUNs and igroups with FilerView 70
- Creating LUNs and igroups by using individual commands 73

Creating iSCSI LUNs on vFiler units for MultiStore 78

Chapter 5 **Managing LUNs** 81

Managing LUNs and LUN maps 82

Displaying LUN information 88

Chapter 6 **Managing iSCSI igroups** 93

Managing igroups. 94

Using igroups on vFiler units 97

Chapter 7 **Managing FCP Initiator Groups** 99

Managing igroups. 100

Managing Fibre Channel initiator requests 105

Chapter 8 **Managing FCP in a clustered environment** 111

How FCP cfmode settings work 112
- Overview of partner mode 114
- Overview of single_image mode. 118
- Overview of standby mode. 122
- Overview of dual_fabric mode. 125
- Overview of mixed mode 129

Changing the cluster's cfmode setting 131

Making LUNs available on specific FCP target ports 133

Chapter 9 **Managing Disk Space** 139

Monitoring disk space 140

Defining a space management policy. 152

Chapter 10 **Using Data Protection with iSCSI and FCP** 159

Data ONTAP protection methods 160

Using Snapshot copies 163

Using LUN clones 165

Deleting busy Snapshot copies 168
Using SnapRestore 171
Backing up data to tape 176
Using NDMP 180
Using volume copy 181
Cloning FlexVol volumes 182
Using NVFAIL 187
Using SnapValidator 189

Chapter 11 **Improving Read/Write Performance** 199
Reallocating LUN and volume layout 200
Improving Microsoft Exchange read performance 210

Chapter 12 **Managing the iSCSI Network** 211
Management changes for iSCSI in Data ONTAP 7.1 212
Managing the iSCSI service 216
Registering the storage system with an iSNS server 222
Displaying initiators connected to the storage system 228
Managing security for iSCSI initiators 229
Managing target portal groups 236
Displaying statistics for iSCSI sessions 243
Displaying information for iSCSI sessions and connections 247
Managing the iSCSI service on storage system interfaces 252
Using iSCSI on clustered storage systems 256
Troubleshooting common iSCSI problems 259

Chapter 13 **Managing the Fibre Channel SAN** 263
Managing the FCP service 264
Managing the FCP service on systems with onboard ports 268
Displaying information about HBAs 280

Glossary 291

Index 295

Preface

About this guide

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

- Ensured that your configuration is supported by referring to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, available on the NetApp on the Web™ (NOW) site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.
- Installed your storage system according to the instructions in the documentation for your specific storage system.
- Configured your storage system according to the instructions in the following documents:
  - *Data ONTAP™ Software Setup Guide*
  - *iSCSI Support Kit* for your specific host
  - *SAN Setup Overview for FCP*
  - *SAN Host Attach Kit for Fibre Channel Protocol* for your specific host
  - Any SAN switch documentation for your specific FCP switch, which you can find at http://now.netapp.com/NOW/knowledge/docs/san/

Audience

This guide is for system and storage administrators who are familiar with operating systems, such as Microsoft Windows® 2003 and UNIX®, that run on the hosts that access storage on N series storage systems. It also assumes that you know how block access protocols are used for block sharing or transfers. This guide does not cover basic system or network administration topics, such as IP addressing, routing, and network topology.

Terminology

This guide uses the following terms:

- NetApp storage products (filers, FAS appliances, and NearStore systems) are all *storage systems*—also sometimes called filers or storage appliances.
- *Enter* refers to pressing one or more keys on the keyboard and then pressing the Enter key.
- *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
Italic font	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the `arp -d` *`hostname`* command, you enter the characters *arp -d* 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

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

Introducing Block Access 1

About this chapter

This chapter provides a brief introduction to storage systems and how they are administered.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding storage systems" on page 2
- "Understanding how hosts connect to storage systems" on page 5
- "Understanding how SnapDrive connects to storage" on page 7
- "Related documents" on page 9

Understanding storage systems

What storage systems are

NetApp storage systems (filers, FAS appliances, and NearStore systems) are all *storage systems*—also sometimes called filers or storage appliances—that serve and protect data using protocols for storage area network (SAN), Internet small computer system interface (iSCSI), and network attached storage (NAS) networks. For information about storage system product families, see http://www.netapp.com/products/.

In iSCSI and FCP networks, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). Using the Data ONTAP® operating system, you configure the storage by creating LUNs. The LUNs are accessed by hosts, which are initiators in the storage network.

What Data ONTAP is

Data ONTAP is the operating system for the storage system. 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 storage system as a target device in an iSCSI network using the SCSI protocol over TCP/IP (using the iSCSI service) and in a SAN network using the SCSI protocol over FCP (using the FCP service) to communicate with one or more hosts. You can also configure a storage system as a storage device in a NAS network using Network File System (NFS), CIFS, HTTP, and File Transfer Protocol (FTP). You can configure a single storage system to serve data over all these protocols.

Ways to administer a storage system

You can administer a storage system by using the following methods:

- Command line
- FilerView
- DataFabric Manager. You must purchase the DataFabric Manager license to use this product.

Command-line administration: You can issue Data ONTAP commands at the storage system'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 storage system.

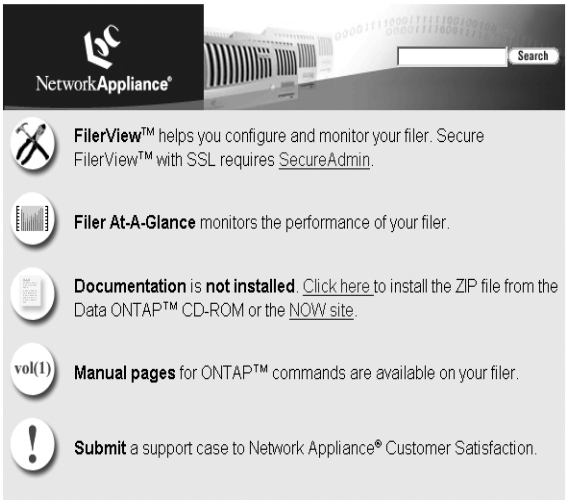
When using the command line, you can get command-line syntax help from the command line by entering the name of the command followed by `help` or `?`. You can also access the online manual (man) pages by entering `man` *command_name*. For example, if you want to read the man page about the `lun` command,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

For more information about storage system 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 storage system from a Web browser or for viewing information about the storage system, its storage units (such as volumes), LUNs, and adapters, and statistics about the storage units, iSCSI, and network traffic. FilerView is easy to use, and it includes Help that 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 storage system, followed by /na_admin/ as the location for the URL. **Example:** If you have a storage system named toaster, enter the following URL in the browser: http://toaster/na_admin. **Result:** The Network Appliance™ online administrative window is displayed. 
3	Click FilerView. **Result:** ◆ If the storage system 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 hosts connect to storage systems

Understanding connection options

Hosts can connect to block storage using Internet small computer systems interface (iSCSI) or Fibre Channel protocol (FCP) networks. To connect to iSCSI networks, hosts can use standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, or dedicated iSCSI HBAs. To connect to FCP networks, hosts require Fibre Channel host bus adapters (HBAs).

What iSCSI Host Support Kits are

A Host Support Kit includes support software and documentation for connecting a supported host to an iSCSI network. The support software includes programs that display information about storage, and programs to collect information needed by customer support to diagnose problems. Depending on the host operating system, you may need to download iSCSI initiator software. You may choose to use an iSCSI HBA instead of a software initiator.

Separate kits are offered for each supported host operating system. In some cases, different versions of the support kit are available for different versions of the host operating system.

The documentation included with the Host Support Kits describes how to set up an iSCSI connection to your storage system. It includes the commands and procedures for the particular host operating system. You should use the support kit documentation along with this guide to set up and manage your iSCSI network.

What FCP host attach kits are

An FCP host attach kit includes support software and documentation for connecting a supported host to an FCP network. The support software includes programs that display information about storage, and programs to collect information needed by NetApp to diagnose problems. The attach kit may include a host bus adapter (HBA) and drivers, or you may obtain an HBA separately. Attach kits are offered for each host operating system. In some cases, different versions of the attach kit are available for different versions of the host operating system.

The documentation included with the host attach kits describes how to set up an FCP connection to your storage system. It includes the commands and procedures for the particular host operating system. You should use the attach kit documentation along with this guide to set up and manage your FCP network.

Downloading files and documentation

You can download iSCSI and FCP documentation from the NOW™ (NetApp on the Web) site at http://now.netapp.com/NOW/knowledge/docs/san/.

You can download support and attach kit software from the NOW site at http://now.netapp.com/NOW/cgi-bin/software.

Be sure to check the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* to verify that your host operating system version, and HBA model and firmware, are qualified to work with your Data ONTAP version and storage system platform.

Understanding how SnapDrive connects to storage

What SnapDrive is

SnapDrive™ software is an optional management package for Microsoft Windows and some UNIX hosts. SnapDrive can simplify some of the management and data protection tasks associated with iSCSI and FCP storage.

About SnapDrive for Windows

SnapDrive for Windows software integrates with the Windows Volume Manager so that storage systems can serve as storage devices for application data in Windows 2000 Server and Windows Server 2003 environments.

SnapDrive manages LUNs on a storage system, making this storage available as local disks on Windows hosts. This allows Windows hosts to interact with the LUNs just as if they belonged to a directly attached disk array.

SnapDrive for Windows provides the following additional features:

- It enables online storage configuration, LUN expansion, and streamlined management.
- It integrates Snapshot™ technology, which creates point-in-time images of data stored on LUNs.
- It works in conjunction with SnapMirror® software to facilitate disaster recovery from asynchronously mirrored destination volumes.

About SnapDrive for UNIX

SnapDrive for UNIX is a tool that simplifies data backup management so that you can recover should data be accidentally deleted or modified. SnapDrive for UNIX uses Snapshot technology to create an image of the data stored on a storage system attached to a UNIX host. You can then restore that data at a later time.

In addition, SnapDrive for UNIX lets you provision storage on the storage system. SnapDrive for UNIX provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file down through the volume manager to the storage-system-side LUNs providing the actual repository.

With SnapDrive for UNIX installed, you can perform the following tasks:

- Create and restore consistent Snapshot copies of one or more volume groups on a storage system. Host volume groups can span multiple storage system volumes and even multiple storage systems.

- Rename a Snapshot copy of one or more host volume groups.
- Restore or delete a Snapshot copy.
- Display information about Snapshot copies that SnapDrive for UNIX created.
- Display information about which Data ONTAP LUNs are used for a specific host volume group, host volume, or file system.
- Connect objects captured by a Snapshot copy at a new location on a host.
- Disconnect objects captured by a Snapshot copy from the host.
- Create storage on a storage system. This storage can be in the form of LUNs, file systems, logical volumes, or disk groups.
- Resize or delete storage.
- Connect storage to and disconnect storage from the host.

SnapDrive limitations

In general, SnapDrive software works only with the storage it provisions. If you use SnapDrive, do not create LUNs manually.

Be sure that you have a supported version of SnapDrive for your version of Data ONTAP, your host environment, and your iSCSI support kit. See the support and interoperability matrices at: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Related documents

Where to go for more information

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

If you want...	Go to...
The most current system requirements for your host and the supported storage system models for Data ONTAP licensed with iSCSI	*Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Configuration limits for iSCSI and FCP environments	*iSCSI Configuration Guide* and *Fibre Channel Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Information about TCP/IP network features supported by Data ONTAP	*Data ONTAP Network Guide*
Information about how to install and configure iSCSI and FCP initiator hardware and software	iSCSI Host Support Kit and FCP Host Attach Kit documentation for your specific host at http://now.netapp.com/NOW/knowledge/docs/san/
Slot assignments for network adapters in the storage system	*System Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/
The latest information about Data ONTAP updates, new features, and limitations.	*Data ONTAP Release Notes*
Information about installing and using SnapDrive	The SnapDrive section of the SAN/IPSAN Information Library page at: http://now.netapp.com/NOW/knowledge/docs/san/

How Data ONTAP Implements an iSCSI Network 2

About this chapter

This chapter provides an overview of how Data ONTAP implements the iSCSI protocol in an iSCSI network.

Topics in this chapter

This chapter discusses the following topics:

- "Changes for Data ONTAP 7.1" on page 12
- "Understanding how Data ONTAP implements an iSCSI network" on page 13
- "Setup Procedure Overview" on page 19

Changes for Data ONTAP 7.1

Administration model changes

In Data ONTAP 7.1, the administrative model for managing iSCSI networks has changed. If you have used iSCSI with earlier versions of Data ONTAP, you should pay special attention to this section and to "Management changes for iSCSI in Data ONTAP 7.1" on page 212.

In earlier releases of Data ONTAP, you managed the iSCSI software target driver using the `iswt` command. In Data ONTAP 7.1, you do not need to manage the iswt driver. Instead, you manage the iSCSI service using FilerView or the `iscsi` command. You manage the underlying networking interfaces using the standard networking commands or FilerView pages.

Target portal group changes

In earlier releases, each interface (Ethernet port or vif) was in its own target portal group and there was no way to change this. In Data ONTAP 7.1, a portal group can contain multiple interfaces. This change is required to support multi-connection sessions. Each iSCSI session between an initiator and target can have more than one underlying TCP connection.

Multi-connection session changes

Data ONTAP 7.1 supports multi-connection sessions. An iSCSI session between an initiator and the storage system can use as many as 16 TCP/IP connections. By default, this feature is turned off and only one TCP/IP connection is allowed for each session. See "Enabling multi-connection sessions" on page 213.

Error recovery level changes

The iSCSI specification (RFC 3720) defines three error recovery levels: 0, 1, and 2. The specification requires initiators and targets to support level 0; the other two levels are optional. With ErrorRecoveryLevel=0, whenever a problem is detected in an iSCSI session, the session is typically dropped and then reestablished. Levels 1 and 2 enable more sophisticated error detection and recovery.

Data ONTAP 7.1 can support all three error recovery levels; earlier versions supported only level 0.

By default, the storage system allows only error recovery level 0. See "Enabling error recovery levels 1 and 2" on page 214.

Understanding how Data ONTAP implements an iSCSI network

What iSCSI is

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3720 (http://www.ietf.org/).

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The storage system does not have a hardware iSCSI HBA. The iSCSI protocol is implemented over the storage system's standard gigabit Ethernet interfaces using a software driver.

The connection between the initiator and target uses a standard TCP/IP network. No special network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your regular public network. The storage system listens for iSCSI connections on TCP port 3260.

What LUNs are

From the storage system, 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 nodes are

In an iSCSI network, there are two types of nodes: targets and initiators. *Targets* are storage systems, and *initiators* are hosts. Switches, routers, and ports are TCP/IP devices only and are not iSCSI nodes.

How nodes are connected

Supported configurations: Storage systems and hosts can be direct-attached or they can be connected through Ethernet switches. Both direct-attached and switched configurations use Ethernet cable and a TCP/IP network for connectivity.

How iSCSI is implemented on the host: iSCSI can be implemented on the host in one or more of the following ways:

- Initiator software that uses the host's standard Ethernet interfaces.
- An iSCSI host bus adapter (HBA). An iSCSI HBA appears to the host operating system as a SCSI disk adapter with local disks.
- TCP Offload Engine (TOE) adapter that offloads TCP/IP processing. The iSCSI protocol processing is still performed by host software.

For information about the types of initiators supported, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

How target nodes are connected to the network: The storage system does not use a hardware iSCSI HBA to implement the iSCSI protocol. The iSCSI protocol on the storage system is implemented over the storage system's standard Ethernet interfaces using software that is integrated into Data ONTAP. iSCSI can be implemented over multiple storage system Ethernet interfaces. An interface used for iSCSI can also transmit traffic for other protocols, such as CIFS or NFS.

Note

For F800 series and FAS900 series models, the e0 interface is a 10/100 interface. Although you can use this interface for iSCSI traffic, NetApp strongly recommends using only gigabit Ethernet (GbE) interfaces for iSCSI traffic.

How nodes are uniquely identified

Every iSCSI node must have a node name. The two formats, or type designators, for iSCSI node names are *iqn* and *eui*. The storage system always uses the iqn-type designator. The initiator can use either the iqn-type or eui-type designator.

iqn-type designator: This is a logical name. It is not linked to an IP address; rather, it is based on the following components:

- The type designator itself, iqn, followed by a period (.)
- The date when the naming authority acquired the domain name, followed by a period
- The name of the naming authority, optionally followed by a colon (:)
- A unique device name

Note

Some initiators might provide variations on the above format. For detailed information about the default initiator-supplied node name, see the documentation provided with your iSCSI Host Attach Kit or Support Kit.

The format is:
iqn.*yyyy-mm.backward_naming_authority*:*unique_device_name*

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

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

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

The following example shows the iSCSI node name for an initiator that is an application server:

iqn.1987-06.com.initvendor1:123abc

Storage system node name: Each storage system has a default node name based on a reverse domain name and the serial number of the storage system's non-volatile RAM (NVRAM) card in the following format:

iqn.1992-08.com.netapp:sn.*serial_number*

The following example shows the default node name for a storage system with the serial number 12345678:

iqn.1992-08.com.netapp:sn.12345678

eui type designator: The format is based on the following components:

- The type designator itself, eui, followed by a period (.)
- Sixteen hexadecimal digits

The format is:
eui.*nnnnnnnnnnnnnnnn*

How the storage system checks initiator node names

The storage system checks the format of the initiator node name at session login time. If the initiator node name does not comply with storage system node name requirements, the storage system rejects the session.

How node names are used

The host's node name is used to create initiator groups (igroups). When you create an igroup, you specify a collection of node names of iSCSI initiators. You map a LUN on a storage system to the igroup to grant all the initiators in that group access to that LUN. If a host's node name is not in an igroup that is mapped to a LUN, that host does not have access to the LUN and the LUNs do not appear as local disks on that host.

Only one host should be allowed to access any given LUN. Clustered hosts can access the same LUNs as long as the clustering software allows only one host to write to any given LUN at a time.

Default port for iSCSI

The iSCSI protocol is configured in Data ONTAP to use TCP port number 3260. Data ONTAP does not support changing the port number for iSCSI. Port number 3260 is registered as part of the iSCSI specification and cannot be used by any other application or service.

What target portal groups are

A target portal group is a set of network portals within an iSCSI node over which an iSCSI session is conducted. In a target, a network portal is identified by its IP address and listening TCP port. For storage systems, each network interface can have one or more IP addresses and therefore one or more network portals. A network interface can be an Ethernet port, virtual local area network (VLAN), or virtual interface (vif).

The assignment of target portals to portal groups is important for two reasons:

- The iSCSI protocol allows only one session between a specific iSCSI initiator port and a single portal group on the target.
- All connections within an iSCSI session must use target portals that belong to the same portal group.

By default, Data ONTAP maps each Ethernet interface on the storage system to its own default portal group. You can create new portal groups that contain multiple interfaces.

You can have only one session between an initiator and target using a given portal group. To support some multipath I/O (MPIO) solutions, you need to have separate portal groups for each path. Other initiators, including the Microsoft iSCSI initiator version 2.0, support MPIO to a single target portal group by using different initiator session IDs (ISIDs) with a single initiator node name.

Understanding iSNS

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

You obtain an iSNS server from a third-party vendor. If you have an iSNS server on your network, and it is configured and enabled for use by both the initiator and the storage system, the storage system automatically registers its IP address, node name, and portal groups with the iSNS server when the iSNS service is started. The iSCSI initiator can query the iSNS server to discover the storage system as a target device.

If you do not have an iSNS server on your network, you must manually configure each target to be visible to the host. For information on how to do this, see the appropriate iSCSI host initiator Support Kit or the iSCSI Host Bus Adapter Attach Kit documentation for your specific host.

Currently available iSNS servers support different versions of the iSNS specification. Depending on which iSNS server you are using, you may have to set a configuration parameter in the storage system.

Understanding CHAP authentication

The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators and targets. When you use CHAP authentication, you define CHAP user names and passwords on both the initiator and the storage system.

During the initial stage of an iSCSI session, the initiator sends a login request to the storage system to begin the session. The login request includes the initiator's CHAP user name and CHAP algorithm. The storage system responds with a CHAP challenge. The initiator provides a CHAP response. The storage system verifies the response and authenticates the initiator. The CHAP password is used to compute the response.

Communication sessions

During an iSCSI session, the initiator and the target communicate over their standard Ethernet interfaces, unless the host has an iSCSI HBA. The storage system appears as a single iSCSI target node with one iSCSI node name. For storage systems with a MultiStore™ license enabled, each vFiler™ unit is a target with a different node name.

On the storage system, the interface can be an Ethernet port, virtual network interface (vif), or a virtual LAN (VLAN) interface.

Each interface on the target belongs to its own portal group by default. This enables an initiator port to conduct simultaneous iSCSI sessions on the target, with one session for each portal group. The storage system supports up to 1,024 simultaneous sessions, depending on its memory capacity. To determine whether your host's initiator software or HBA can have multiple sessions with one storage system, see your host OS or initiator documentation.

You can change the assignment of target portals to portal groups as needed to support multi-connection sessions, multiple sessions, and multipath I/O.

Each session has an Initiator Session ID (ISID), a number that is determined by the initiator.

Options that are automatically enabled

The following options are automatically enabled when the iSCSI service is turned on. Do not change these options.

- `volume option create_ucode` to On
- `cf.takeover.on_panic` to On

How vFiler units are used

If you purchased a MultiStore® license and created vFiler™ virtual storage systems, you can enable the iSCSI license for each vFiler to manage LUNs and igroups on a per vFiler basis. For information about vFiler units, see "Creating iSCSI LUNs on vFiler units for MultiStore" on page 78 and the sections on iSCSI service on vFiler units or LUNs on vFiler units in the *MultiStore Management Guide*.

Using iSCSI with clustered storage systems

Clustered storage systems provide high availability because one system in the cluster can take over if its partner ever fails. During cluster failover (CFO), the working storage system assumes the IP addresses of the failed partner and can continue to support iSCSI LUNs.

The two systems in the cluster should have identical networking hardware with equivalent network configurations. The target portal group tags associated with each networking interface must be the same on both systems in the cluster. This ensures that the hosts see the same IP addresses and target portal group tags whether connected to the original storage system or connected to the partner during CFO.

Setup Procedure Overview

Setup procedure

The procedure for setting up the iSCSI protocol on a host and storage system follows the same basic sequence for all host types:

Step	Action
1	Install the initiator HBA and driver or software initiator on the host and record or change the host's iSCSI node name. It is recommended that you use the host name as part of the initiator node name to make it easier to associate the node name with the host.
2	Configure the storage system, including: ◆ Licensing and starting the iSCSI service ◆ Optionally configuring CHAP ◆ Creating LUNs, creating an igroup that contains the host's iSCSI node name, and mapping the LUNs to that igroup **Note** If you are using SnapDrive, do not configure LUNs manually. Configure them using SnapDrive after it is installed.
3	Configure the initiator on the host, including: ◆ Setting initiator parameters, including the IP address of the target on the storage system ◆ Optionally configuring CHAP ◆ Starting the iSCSI service
4	Access the LUNs from the host, including: ◆ Creating file systems on the LUNs and mounting them, or configuring the LUNs as raw devices ◆ Creating persistent mappings of LUNs to file systems

Note

You must alternate between setting up the host and the storage system in the order shown above.

How Data ONTAP Implements an FCP Network 3

About this chapter

This chapter provides an overview of how Data ONTAP implements the Fibre Channel Protocol (FCP) in an FCP network.

Topics in this chapter

This chapter discusses the following topic:

- "Understanding how Data ONTAP implements a Fibre Channel SAN" on page 22

Understanding how Data ONTAP implements a Fibre Channel SAN

What FCP is

FCP is a licensed service on the storage system 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 264.

What nodes are

In an FCP network, nodes include targets, initiators, and switches. *Targets* are storage systems, and *initiators* are hosts. Storage systems 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 storage system, 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, ASCII string generated by the storage system. Many multipathing software packages use this serial number to identify redundant paths to the same LUN. You can display the LUN serial number with the `lun show -v` command.

How nodes are connected

Storage systems and hosts have Host Bus Adapters (HBAs) so they can be connected directly to each other or to Fibre Channel switches with optical cable. For switch or storage system management, they might be connected to each other or to TCP/IP switches with Ethernet cable.

When a node is connected to the Fibre Channel storage area network (FC SAN), 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 worldwide node name (WWNN) and a worldwide 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 storage system 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 62.

  You can also create portsets to make a LUN visible only on specific target ports. A portset consists of a group of FCP target ports. You bind a portset to an igroup. Any host in the igroup can access the LUNs only by connecting to the target ports in the portset. For detailed information about portsets, see "Making LUNs available on specific FCP target ports" on page 133.

- Uniquely identifying a storage system's HBA target ports

  The storage system's WWPNs uniquely identify each target port on the system. The host operating system uses the combination of the WWNN and WWPN to identify storage system 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 storage systems are identified: When the FCP service is first initialized, it assigns a WWNN to a storage system based on the serial number of its NVRAM adapter. The WWNN is stored on disk. Each target port on the HBAs installed in the storage system 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 `fcp show adapter, fcp config`, or `sysconfig -v, fcp nodename,` or FilerView (click LUNs > FCP > Report) to see the system's WWNN as `FC Nodename` or `nodename` or the system's WWPN as `FC portname` or `portname`.

Note

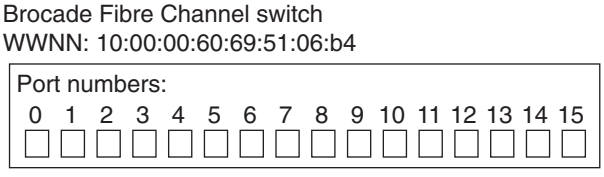
The target WWPNs might change if you add or remove HBAs on the storage system.

System serial numbers: The storage system 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 when the storage system is manufactured. 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 FCP Host Attach Kit documentation for your host. These documents describe commands supplied by the Attach Kit 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 `fcp 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 storage system. 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.




```
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
```

About FCP ports

The FCP service is implemented over the target's and initiator's FCP ports. Initiator HBAs can have one or two ports. Depending on the specific model, a storage system can have two types of target ports:

- Host Bus Adapter (HBA) ports—The storage system has a target FCP HBA with two ports that are labeled Port A and Port B (if there is a second port). F800 series and FAS900 series systems use target HBAs.

- Onboard ports—The following systems have onboard FCP adapters, or ports, that you can configure to connect to disk shelves or to operate in SAN target mode:
  - FAS270 models—These storage systems have a port labeled Fibre Channel C (with an orange label). You can configure the Fibre Channel C port in initiator or target mode. You use initiator mode to connect to 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.
  - FAS3000 models—These storage systems have four onboard Fibre Channel ports that have orange labels and are numbered 0a, 0b, 0c, 0d. You use the `fcadmin` command to configure the ports to operate in SAN target mode or initiator mode. In SAN target mode, the onboard ports connect to Fibre Channel switches or fabric. In initiator mode, they connect to disk shelves.

  For detailed information about systems with integrated target ports, see "Managing the FCP service on systems with onboard ports" on page 268.

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 Fibre Channel initiator requests" on page 105.

How Data ONTAP supports FCP with clustered systems

Enabled options for cluster configurations: Clustered storage systems in a Fibre Channel SAN 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 (value set to `on`) when the FCP service is turned on. Do not change them.

- `vol options create_ucode`
- `cf.takeover.on_panic`

cfmode settings: If your storage systems are in a cluster, Data ONTAP provides multiple modes of operation required to support homogeneous and heterogeneous host operating systems. 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
- Provide access to local and partner LUNs in a cluster

For detailed information, see Chapter 8, “Managing FCP in a clustered environment,” on page 111.

Configuring Storage 4

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 need to store your LUNs. It also describes the methods for creating LUNs, igroups, and LUN maps.

This chapter assumes that your SAN is set up and configured, and that the iSCSI service or FCP service is licensed and enabled. If that is not the case, see "Managing the iSCSI service" on page 216 or "Managing the FCP service" on page 264 for information about these topics.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding storage units" on page 28
- "Understanding space reservation for volumes and LUNs" on page 30
- "Understanding how fractional reserve affects available space" on page 33
- "How guarantees on FlexVol volumes affect fractional reserve" on page 45
- "Calculating the size of a volume" on page 48
- "Guidelines for creating volumes that contain LUNs" on page 53
- "Creating LUNs, igroups, and LUN maps" on page 57
- "Creating iSCSI LUNs on vFiler units for MultiStore" on page 78

Understanding storage units

Storage units for managing disk space

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

- 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.

A *plex* is a collection of one or more RAID groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system volumes. Data ONTAP uses plexes as the unit of RAID-level mirroring when the SyncMirror® software is enabled.

An aggregate is a collection of one or two plexes, depending on whether you want to take advantage of RAID-level mirroring. If the aggregate is unmirrored, it contains a single plex. Aggregates provide the underlying physical storage for traditional and FlexVol 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.

Once you set up the underlying aggregate, you can create, clone, or resize FlexVol volumes without regard to the underlying physical storage.You do not have to manipulate the aggregate frequently.

You use either traditional or FlexVol 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 storage units, including aggregates, and traditional and FlexVol volumes, see the *Data ONTAP Storage Management Guide*.

Where LUNs reside

You create LUNs in the root of a volume (traditional or flexible) or in the root of 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 FlexVol volume so that write operations to those LUNs do not fail because of a lack of disk space on the storage system. Other operations, such as taking a Snapshot copy or the creation of new LUNs, can occur only if there is enough available unreserved space; these operations are restricted from using reserved space.

What happens when space reservation is disabled

You can create LUNs with space reservation enabled or disabled. If you disable space reservation, write operations to a LUN might fail due to insufficient disk space and the host application or operating system might crash. 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.

See "How to use individual commands" on page 73 for information about creating LUNs with space reservation enabled or disabled.

What fractional reserve is

Fractional reserve controls the amount of space Data ONTAP reserves in a traditional or FlexVol 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
- FlexVol volumes that have the `guarantee` option set to `volume`

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

For detailed information about how guarantees affect fractional reserve, see "How guarantees on FlexVol volumes affect fractional reserve" on page 45.

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. LUNs that do not have space reservation enabled are not included in the total LUN size. For example, if there are two 200-GB LUNs in a volume (400 GB total), 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 (400 GB total * 50% = 200 GB).

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.

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 FlexVol volumes with the volume guarantee, you can set *pct* to any value from 0 to 100. For FlexVol 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 on a volume maintains the fractional reserve setting of the volume and the space reservation settings of the LUNs in that volume. A single file SnapRestore of a LUN restores the space reservation setting of that LUN, provided there is enough space in the volume.

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* versions do 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 Snapshot
copies 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 42.

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 Chapter 9, "Managing Disk Space".

For detailed information

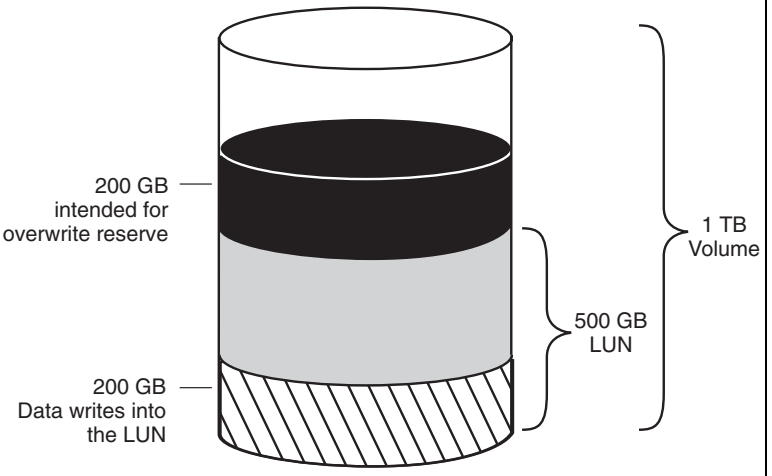
For detailed information, see the following sections:

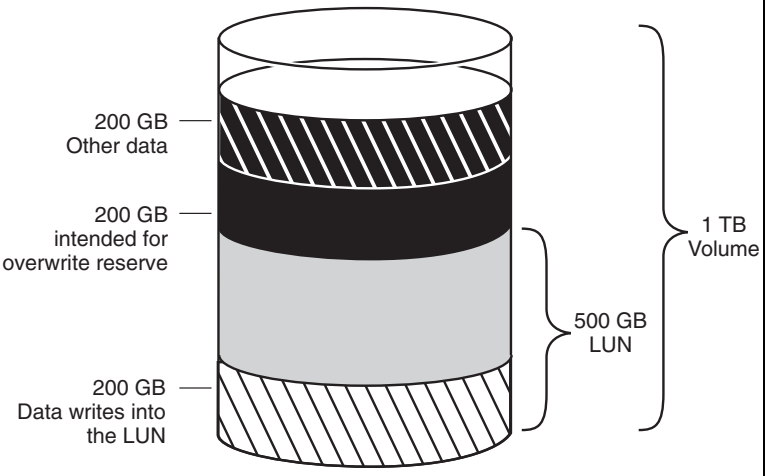
- "How 100 percent fractional reserve affects available space" on page 34
- "How reducing fractional reserve affects available space" on page 40
- "Reasons to set fractional reserve to zero" on page 44

How 100 percent fractional reserve affects available space

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

The following example shows how the default fractional reserve setting of 100 affects available space in a 1-TB volume with a 500-GB LUN.

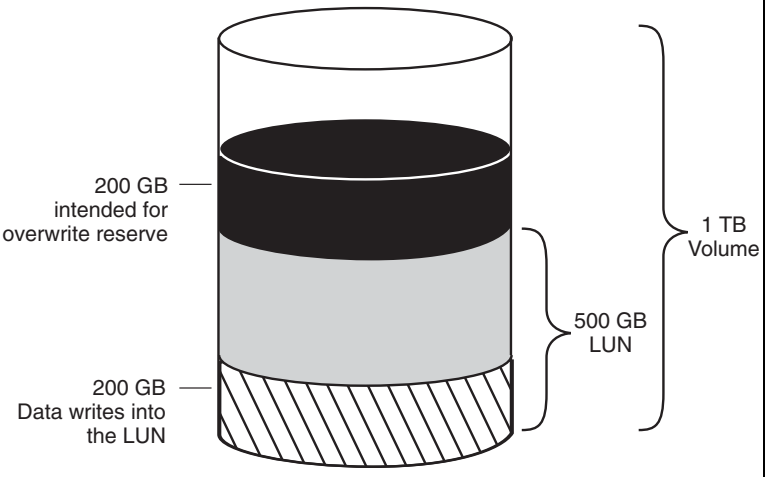
Stage	Status
1	The following illustration shows a 1-TB volume with a 500-GB LUN after 200 GB of data are written to the LUN. The volume has 200 GB of space intended for overwrite reserve. This space is actually reserved only when you take a Snapshot copy by using either the `snap` command or snapshot-methods,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

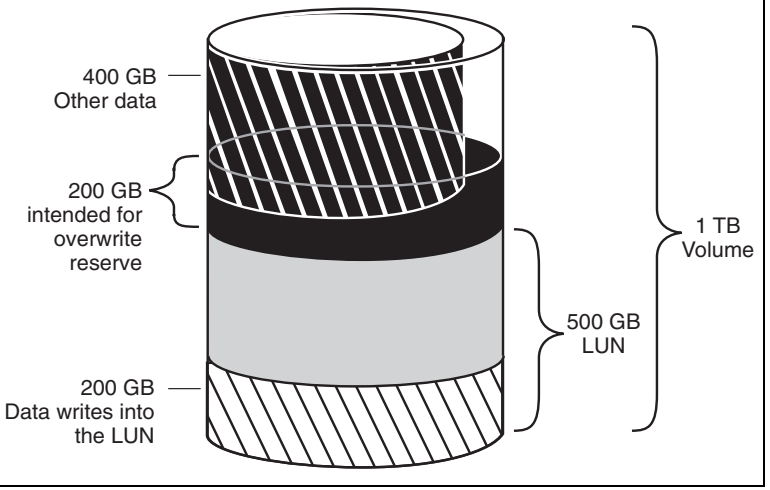
Stage	Status
2	The following illustration shows that the volume still has enough space for the following: ◆ 500-GB LUN (containing 200 GB of data) ◆ 200 GB intended reserve space for overwrites ◆ An additional 200 GB of other data At this point, there is enough space for one Snapshot copy. 

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.

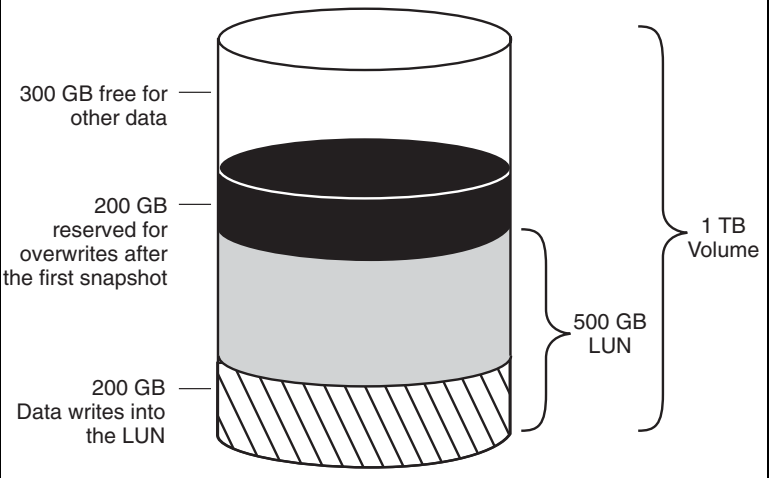
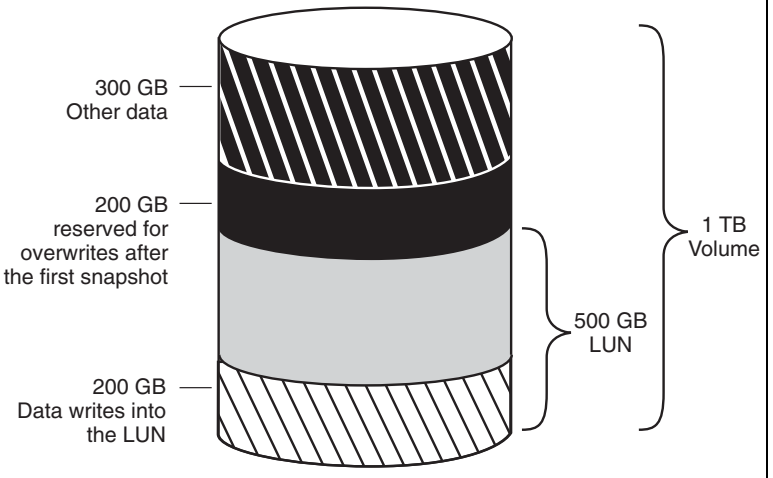
Example 1:

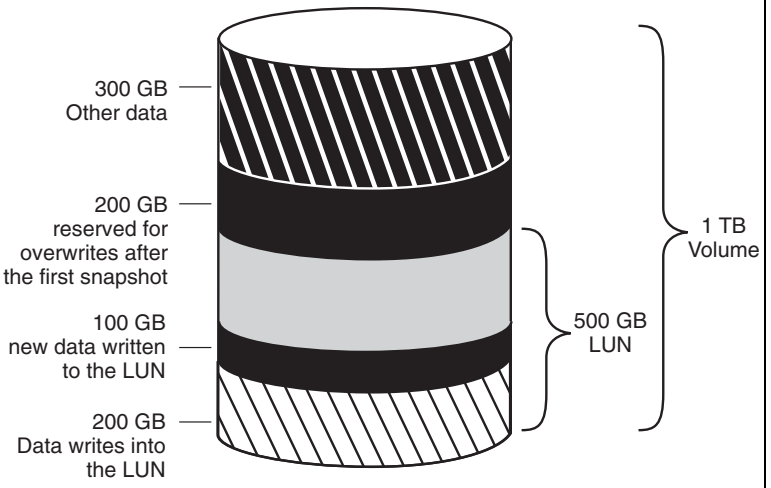
Stage	Status
1	The following illustration shows the 1-TB volume with a 500-GB LUN that contains 200 GB of data. There are 200 GB intended for overwrite reserve. At this point, you have not taken a Snapshot copy, and the volume has 500 GB of available space. 

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 copy. This is because the 400 GB of other data does not leave enough space for the intended overwrite reserve. The Snapshot copy requires Data ONTAP to reserve 200 GB of space, but you have only 100 GB of available space. 

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 copy. The volume has 200 GB reserved for overwrites to the original data and 300 GB of free space remaining for other data. 
3	The following illustration shows the volume after you write 300 GB of other data (not in the LUN) to the volume. The volume reports that it is full because you have used all free space, but you can write data to the LUN indefinitely. 

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 copy. The second Snapshot copy requires 300 GB of reserve space because the total size of the data in the LUN is 300 GB. 

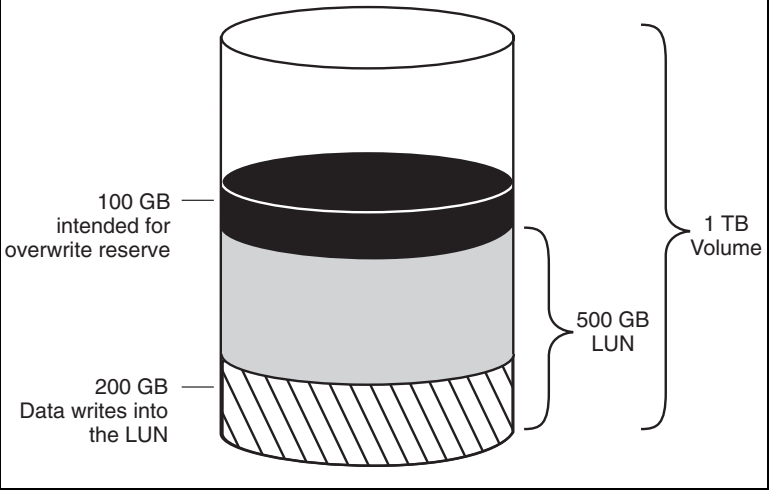
How reducing fractional reserve affects available space

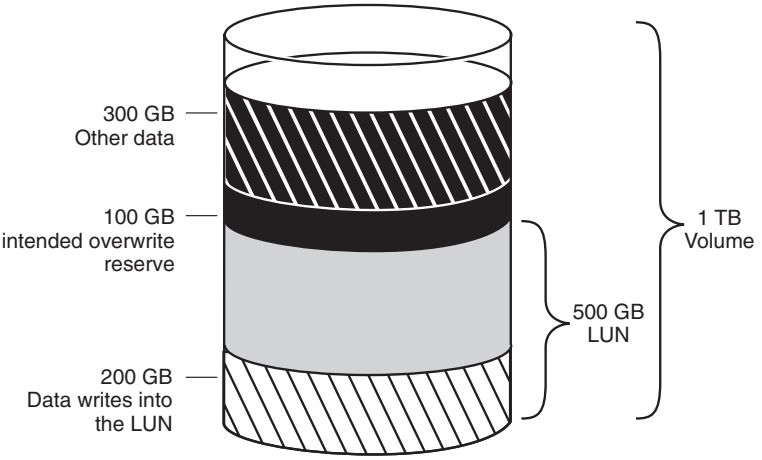
When you can reduce fractional reserve

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

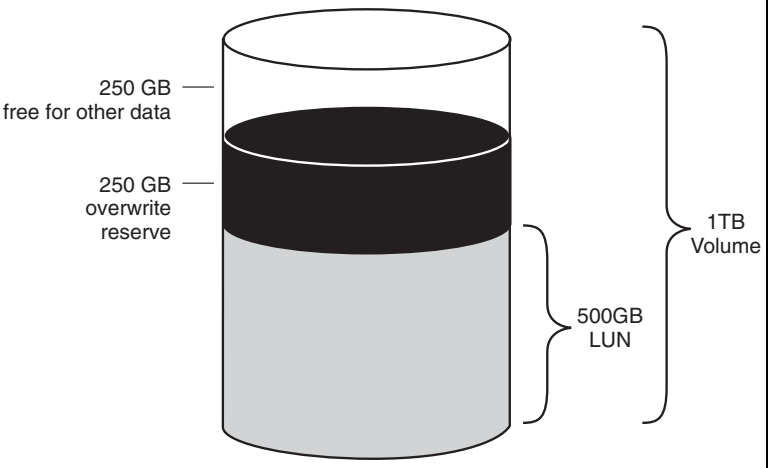
The following example shows how a fractional reserve setting of 50 percent affects available space in the same 1-TB volume with a 500-GB LUN.

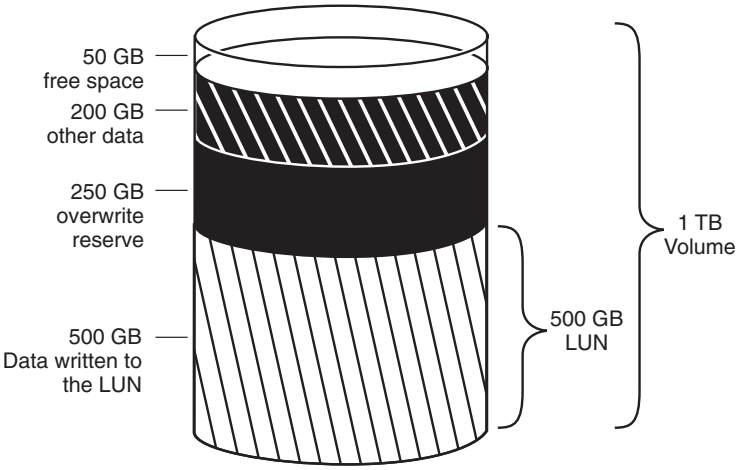
Stage	Status
1	The following illustration shows a 1-TB volume with a 500-GB LUN after 200 GB of data are written to the LUN. The volume has 100 GB intended for overwrite reserve because the fractional reserve for this volume is set to 50 percent. 

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 copy ◆ Writing up to 100 GB of other data and maintaining the ability to take a Snapshot copy Compare this example with the volume shown in "Example 2" on page 37, in which the same volume has an overwrite reserve of 100 percent, but the volume has run out of free space. 

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 copy. The volume has 250 GB reserved for overwrites to the LUN and 250 GB available for other data. 

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. 

Reasons to set fractional reserve to zero

Use case scenarios

You might want to set the fractional reserve to 0 on a volume that is a dedicated target for SnapMirror or SnapVault Snapshot copies of LUNs. For example, the volume might be on a NearStore system for long-term retention. You mount the LUN Snapshot copy only for data recovery purposes. This means the rate of change in this volume is low, and you can set the fractional reserve to 0 to provide more usable space. You set the fractional reserve option to 0 in this case if maintaining usable space is your priority because if you do a large number of write operations during the recovery process and the volume runs out of space, the LUN will go offline.

How guarantees on FlexVol volumes affect fractional reserve

What guarantees are

Guarantees on a FlexVol volume ensure that write operations to a specified FlexVol 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 FlexVol 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 FlexVol volume is always available from its aggregate. This is the default setting for FlexVol volumes. Fractional reserve is adjustable from the default of 100 percent only when a FlexVol volume has guarantees.

- `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 FlexVol 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 can specify guarantees when you create a FlexVol volume by using the `-s` option of the `vol create` command:

vol create *f_vol_name* [-l *language_code*] [-s {volume|file|none}] *aggr_name size*{k|m|g|t}

You can change the guarantee setting of the volume by using the `vol options` command:

vol options *f_vol_name* guarantee *guarantee_value*

f_vol_name is the name of the FlexVol 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*.

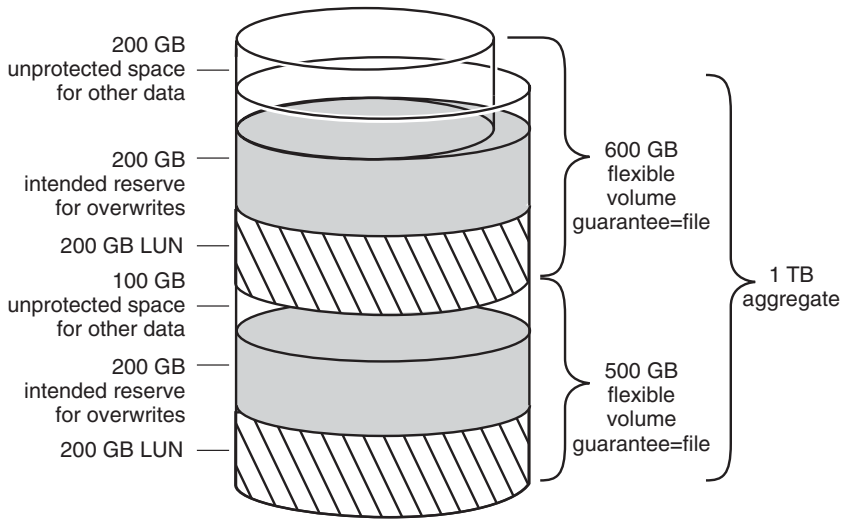
Overcommitting an aggregate

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 FlexVol 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 FlexVol volumes you create might be larger than the containing aggregate.

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

Each FlexVol volume has space for other data. For example, you can create non-space-reserved LUNs in a FlexVol 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:

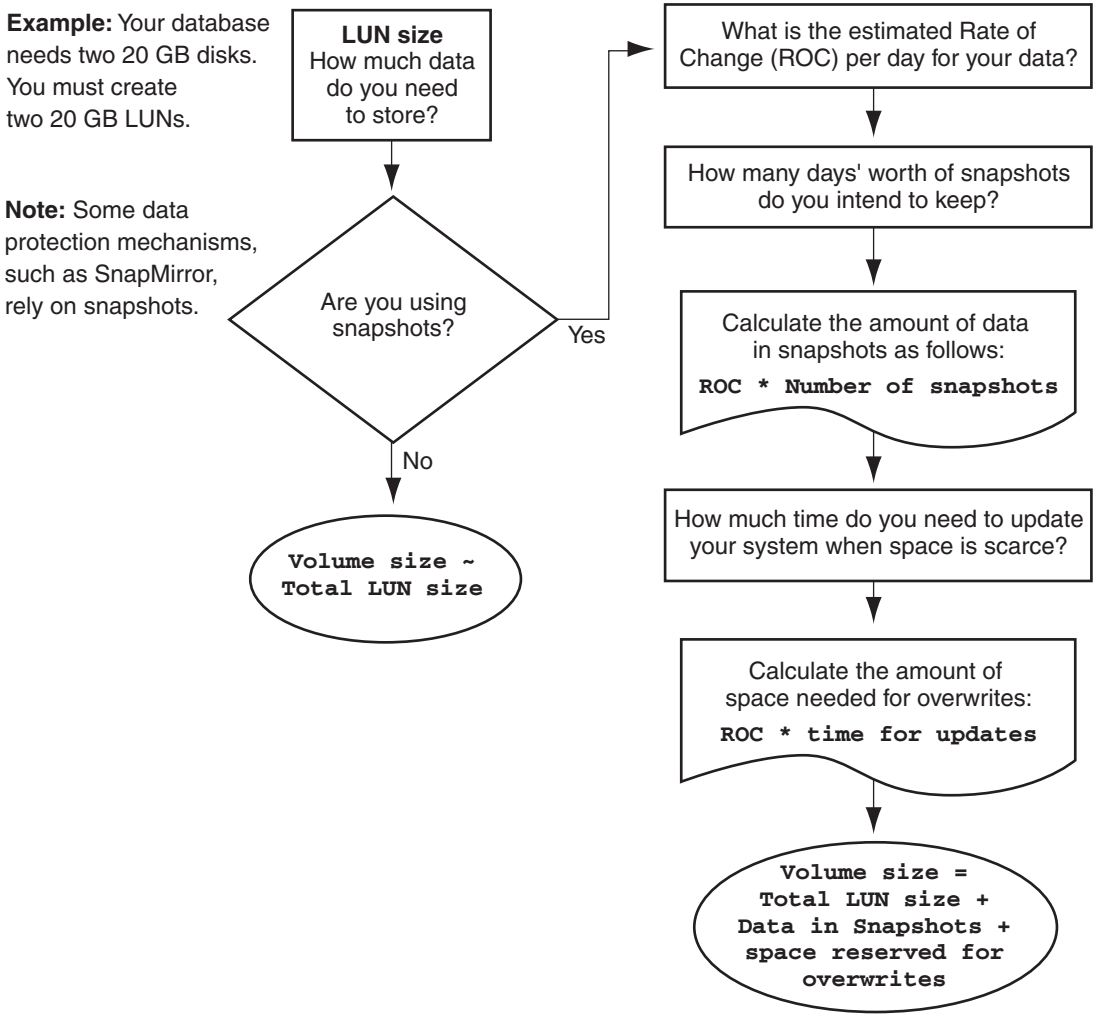
- Total size of all the LUNs in the volume.
- Whether you want to maintain Snapshot copies.
- If you want to maintain Snapshot copies, the number of Snapshot copies you want to maintain and the amount of time you want to retain them (retention period).
- Rate at which data in the volume changes.
- Amount 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 49
- "Determining the volume size when you do not need Snapshot copies" on page 50
- "Calculating the amount of space for Snapshot copies" on page 50
- "Calculating the fractional reserve" on page 51



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 space-reserved LUNs. The total LUN size in this example is 40 GB. The total LUN size does not include LUNs that do not have space reservation enabled.

Determining the volume size when you do not need Snapshot copies

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

- Traditional volumes

  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.

- FlexVol volumes

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

ONTAP data protection methods and Snapshot copies: Before you determine that you do not need Snapshot copies, verify the method for protecting data in your configuration. Most data protection methods, such as SnapRestore, SnapMirror, SnapManager for Microsoft Exchange or Microsoft SQL Server, SyncMirror®, dump and restore, and ndmpcopy methods rely on Snapshot copies. If you are using these methods, calculate the amount of space required for these Snapshot copies.

Note

Host based backup methods do not require additional space.

Calculating the amount of space for Snapshot copies

The amount of space you need for Snapshot copies 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 Snapshot copies and fractional overwrite reserve. The ROC depends on how often you overwrite data.

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

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

Space for Snapshot copies = ROC in bytes per day * number of Snapshot copies

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 copy each day and want to keep three weeks' worth of Snapshot copies, for a total of 21 Snapshot copies. The amount of space you need for Snapshot copies 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 Snapshot copies 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 Snapshot copies. You want to ensure that write operations to the LUNs do not fail for three days after you take the last Snapshot copy. You need 2 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 copy each day and you want to keep the Snapshot copies for 10 days.

  You need 100 GB of space for Snapshot copies (10 GB ROC * 10 Snapshot copies).
- You want to ensure that you can continue to write to the LUNs through the weekend, even after you take the last Snapshot copy 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 GB = 20 percent of 100 GB).

Calculate the size of your volume as follows:

Volume size = Total LUN size + Amount of space for Snapshot copies + 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 copy. You must monitor available

space and take corrective action by increasing the size of your volume or aggregate or deleting Snapshot copies to ensure you can continue to write to the LUNs.

CAUTION

If you do not actively monitor available space and the volume becomes full, writes to the LUN fail, the LUN goes offline, and your application might crash.

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 Snapshot copies

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

Space requirements for LUN clones

A space-reserved LUN clone requires as much space as the space-reserved parent LUN. If the clone is not space-reserved, make sure the volume has enough space to accommodate changes to the clone.

Changing the size of a FlexVol volume

After you calculate the initial size of a FlexVol volume and create LUNs, you can monitor available disk space to confirm that you correctly estimated your volume size or increase the volume size depending on your application requirements. You can also define space management policy to perform the following tasks:

- Automatically increase the size of the FlexVol volume when it begins to run out of space
- Automatically delete Snapshot copies when the FlexVol volume begins to run out of space

For detailed information, see Chapter 9, “Managing Disk Space”.

Guidelines for creating volumes that contain LUNs

Changing Snapshot copy defaults

Data ONTAP Snapshot copies are required for many optional features, such as the SnapMirror feature, SyncMirror feature, dump and restore, and ndmpcopy.

When you create a volume, Data ONTAP automatically:

- Reserves 20 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for taking Snapshot copies within Data ONTAP has no means of ensuring that the data within a LUN is in a consistent state, it is recommended that you change these Snapshot copy settings by performing the following tasks:

- Turn off the automatic Snapshot copy schedule.
- Delete all existing Snapshot copies.
- Set the percentage of snap reserve to zero.

For information about how to change Snapshot copy defaults, see "Changing Snapshot copy defaults" on page 53.

For information about how to use Snapshot copies, see "Using Snapshot copies" on page 163.

Other guidelines to use when creating volumes

It is strongly recommended that you use the following guidelines to create traditional or FlexVol volumes that contain LUNs:

- Do not create any LUNs in the system's root volume. Data ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.
- 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. This is a recommended best practice that simplifies LUN administration and tracking.
- Ensure that the volume option `create_ucode` is set to On.

Data ONTAP requires that the path of a volume or qtree containing a LUN is in the Unicode format. This option is Off by default when you create a volume. It is important to enable this option for volumes that will contain LUNs.

For detailed procedures, see "Verifying and modifying the volume option create_ucode" on page 56.

- To simplify management, 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 copy defaults

Turning off the automatic Snapshot copy schedule: To turn off the automatic Snapshot copy 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 copy schedule, enter the following command: **`snap sched volname 0 0 0`** **Example:** `snap sched vol1 0 0 0` **Result:** This command turns off the Snapshot copy schedule because there are no weekly, nightly, or hourly Snapshot copies scheduled. You can still take Snapshot copies manually by using the `snap` command.
2	To verify that the automatic Snapshot copy schedule is off, enter the following command: **`snap sched [volname]`** **Example:** `snap sched vol1` **Result:** The following output is a sample of what is displayed: `Volume vol1: 0 0 0`

Deleting all existing Snapshot copies: To delete all Snapshot copies, complete the following step.

Step	Action
1	Enter the following command: **`snap delete -a volname`**

Setting the percentage of snap reserve space: 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: **`snap reserve volname percent`** **Note** For volumes that contain LUNs and no Snapshot copies, it is recommended that you set the percentage to zero. **Example:** `snap reserve vol1 0`
2	To verify what percentage is set, enter the following command: **`snap reserve [volname]`** **Example:** `snap reserve vol1` **Result:** The following output is a sample of what is displayed: `Volume vol1: current snapshot reserve is 0% or 0 k-bytes.`

Verifying and modifying the volume option create_ucode

To verify that the create_ucode volume option is enabled, or to enable the option, complete the following steps.

Step	Action

1	To verify that the create_ucode option is enabled (On), enter the following command:
	vol status [*volname*] -v
	Example: vol status vol1 -v
	Result: The following output example shows that the create_ucode option is on:
	`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_ucode=on`**
	`convert_ucode=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`
	Note
	If you do not specify a volume, the status of all volumes is displayed.

2	To enable the create_ucode option, enter the following command: **vol options *volname* create_ucode on** **Example:** vol options vol1 create_ucode on

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 65.

- 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 70.

- 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 73.

Note about using SnapDrive

For Windows and some UNIX hosts, you can use SnapDrive for Windows or SnapDrive for UNIX to create and manage LUNs and igroups. 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 environment, see the *NetApp iSCSI S Support Matrix* or *NetApp FCP SAN Compatibility Matrix* and the *SnapDrive & SnapManager 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 the section labeled "NetApp SnapDrive" or "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 storage system configurations, it is recommended that you distribute LUNs across the 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*, *netware*, *vmware*, 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 appropriate support matrix, specify *image*.

For information about supported hosts, see the *Fibre Channel Host Support Matrices* or the *iSCSI Solutions Support Matrices* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

The size of the LUN: When you create a LUN, you specify its size as raw disk space, depending on the storage system 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

Multiplier suffix	Size
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

It is recommended that you keep space reservation on.

About igroups

Initiator groups (igroups) are tables of host identifiers (FCP WWPNs or iSCSI node names) that are used to control access to LUNs. Typically, you want all of the host's HBAs or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA or software initiator 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 ostypes. Also, a given igroup can be used for FCP or iSCSI, but not both.

FCP example: The following table illustrates how four igroups give access to the LUNs for four different hosts accessing the storage system. The clustered hosts (Host3 and Host4) are both members of the same igroup (aix-group2) and can access the LUNs mapped to this igroup. The igroup named aix-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	aix-group0	10:00:00:00:c9:2b:7c:0f	/vol/vol2/lun0
Host2, multipath (two HBAs) 10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	aix-group1	10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	/vol/vol2/lun1

Host with HBA WWPNs	igroups	WWPNs added to igroups	LUNs mapped to igroups
Host3, multipath, clustered (connected to Host4) 10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02	aix-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/vol2/qtree1/lun2
Host4, multipath, clustered (connected to Host3) 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	aix-group3	10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/qtree1/lun3 /vol/vol2/qtree1/lun4

iSCSI example: The following table shows two hosts and their igroups and LUNs.

Host with node name	igroup	LUNs
Host 5 iqn.1991-05.com.microsoft:host5.domain.com	win_host5_group1	/vol/vol3/lun0 /vol/vol3/lun1
Host 6 iqn.1987-05.com.cisco:host6.domain.com	linux_host6_group1	/vol/vol3/lun2 /vol/vol3/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 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), colon (“:”), and period (“.”).
- Must start with a letter or number.

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 aix1, 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 either iSCSI (`-i`) or FCP (`-f`).

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 ostypes of initiators are *solaris*, *windows*, *hpux*, *aix*, *netware*, *vmware*, and *linux*. If your host OS is not one of these values but it is listed as a supported OS in the appropriate support matrix, specify *default*.

iSCSI node names of iSCSI initiators: You can specify the node names of the initiators when you create an igroup. You can also add them or remove them later.

To know which node names are associated with a specific host, see the Host Support Kit documentation for your host. These documents describe commands that display the host's iSCSI node name.

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

To know which WWPNs are associated with a specific host, see the Host Attach Kit documentation for your host. These documents describe commands supplied by the Attach Kit 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 101.

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 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.

Making LUNs available on specific FCP target ports

When you map a LUN to a Fibre Channel igroup, the LUN is available on all of the storage system's FCP target ports if the igroup is not bound to a portset. A portset consists of a group of FCP target ports. By binding a portset to an igroup, you make the LUN available on a subset of the system's target ports. Any host in the igroup can access the LUNs only by connecting to the target ports in the portset.

You define portsets for FCP target ports only. You do not use portsets for iSCSI target ports.

For detailed information about creating portsets and binding them to igroups, see "Making LUNs available on specific FCP target ports" on page 133.

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 Snapshot copies, 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 Snapshot copies. Data ONTAP takes Snapshot copies at the volume level, and the rate of change of data in all LUNs affects the amount of space needed for Snapshot copies. 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 Snapshot copies.

- Keep backup LUNs in separate volumes.

  It is recommended 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 host support or 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 iSCSI in a Windows environment.

Step	Action

1	On the storage system command line, enter the following command.
	`lun setup`
	Result: The `lun setup` program displays the following instructions. Press Enter to continue or `n` to terminate the program.
	`This setup will take you through the steps needed to create LUNs and to make them accessible by initiators. You can type ^C (Control-C) at any time to abort the setup and no unconfirmed changes will be made to the system.`
	`Do you want to create a LUN? [y]:`

2	Specify the operating system that will be accessing the LUN by responding to the next prompt: `OS type of LUN (image/solaris/windows/hpux/aix/linux/netware/vmware) [image]:` **Example:** `windows` For information about specifying the ostype of the LUN, see "The host operating system type" on page 58.

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 system 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 storage system to fail, which can cause data corruption. It is strongly recommended 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: `5 GB Windows LUN for finance records`
	If you choose not to provide a comment at this time, you can add a comment later with the `lun comment` command or fill in the description field 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[win_host5_group1]:`
	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 igroup you want to create now. For information about naming an igroup, see "The name of the igroup" on page 61.

8	If you entered a new igroup name, specify which protocol will be used by the hosts in the igroup by responding to the next prompt: `Type of initiator group win_host5_group2 (FCP/iSCSI)[FCP]:` **Example:** `iscsi` **Result:** The initiators in this igroup use the iSCSI protocol. Be sure to specify `fcp` or `iscsi` as needed.
9	If you specified an iSCSI igroup, add the iSCSI node names of the initiators that can access LUNs in the igroup by responding to the next prompt: `Enter comma separated nodenames:` **Example:** `iqn.1991-05.com.microsoft:host5.domain.com`
10	If you specified an FCP igroup, add the WWPNs of the hosts that will be in the igroup by responding to the next prompt: `Enter comma separated portnames:`

Step	Action

	Example a: Enter `?` to display initiators that are logged in to the storage system.
	
	Result: The following output is an example of what is displayed:
	`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:`

	Example b: 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 24.

11	If you entered a new igroup name, 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, netware, vmware or default.`
	`OS type of initiator group "win_host5_group2"[windows]:`
	
	For information about specifying the ostype of an igroup, see "About igroups" on page 60.

Step	Action

12	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 "win_host5_group2".`
	`LUN ID at which initiator group "win_host5_group2" sees "/vol/finance/records" [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 Host Attach Kit or Host Support Kit documentation, or host operating system documentation, for information about valid LUN ID numbers.
	
	Note
	It is recommended that you accept the default value for the LUN ID.
	
	After you press Enter, the `lun setup` program displays the information you entered:
	`LUN Path : /vol/finance/records`
	`OS Type : windows`
	`Size : 5g (5368709120)`
	`Comment : 5 GB Windows LUN for finance records`
	`Initiator Group : win_host5_group2`
	`Initiator Group Type : iSCSI`
	`Initiator Group Members : iqn.1991-05.com.microsoft:host5.domain.com`
	`Mapped to LUN-ID : 0`

13	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.
14	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

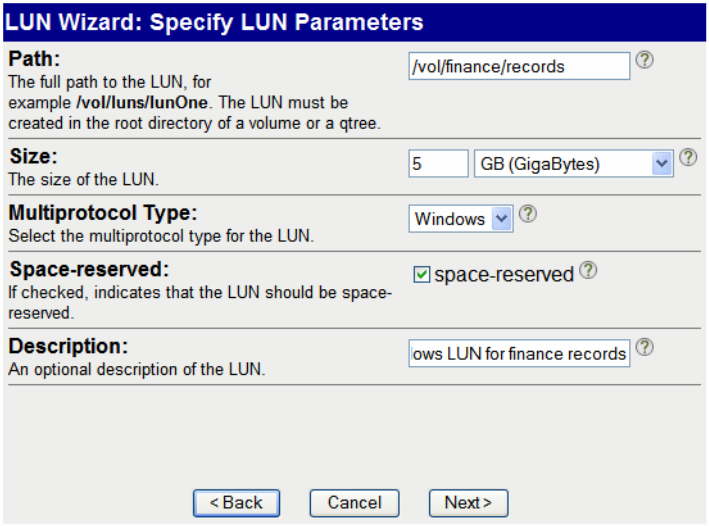
You can use FilerView to create LUNs and igroups with the following methods:

- 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	In the left panel of the FilerView screen, click LUNs. **Result:** The management tasks you can perform on LUNs are displayed.
2	Click Wizard. **Result:** The LUN Wizard window appears.

Step	Action
3	Click the Next button to continue. **Result:** The first window of fields in the LUN Wizard appears. 
4	Enter LUN information in the appropriate fields and click Next.
5	Specify the following information in the next windows: ◆ Whether you want to add an igroup. ◆ Whether you want to use an existing igroup or create a new one. ◆ The iSCSI node names or FCP WWPNs of the initiators in the igroup. ◆ LUN mapping.
6	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	Click LUNs > Add.
2	Fill in the fields.
3	Click Add to commit changes.

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

Step	Action
1	Click Initiator Groups > Add.
2	Fill in the fields.
3	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	Click LUNs > Manage.
2	If the maps are not displayed, click the Show Maps link.
3	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.
4	Click Add Groups to Map.
5	Select an igroup name from the list on the right side of the window.
6	To commit your changes click Add.

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 node names of iSCSI initiators connected to the storage system	**`iscsi initiator show`**
	Sample result:
	`Initiators connected:`
	`  TSIH  TPGroup  Initiator`
	`    64       1   iqn.1991-05.com.microsoft:host5.domain.com / 40:01:37:00:06:00`
	`    66       1   iqn.1991-05.com.microsoft:host6.domain.com / 40:01:37:00:00:00`

Display the WWPNs of FCP initiators connected to the storage system	**`fcp show initiator`**
	Sample result:
	`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`

To do this...	Use this command...
Determine which iSCSI hosts are associated with the initiator node names	To determine the iSCSI node name associated with a host, use the command or method provided by the initiator software. For example, the Windows software initiator applet lists the node name on the General tab or Initiator Settings tab, depending on version. Linux systems typically store the node name in the /etc/initiatorname.iscsi file. For more information, see the iSCSI Host Support Kit or Attach Kit for your host, or the documentation provided by the iSCSI initiator or HBA vendor.
Determine which FCP 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 24.

Create an igroup	**`igroup create {-i | -f} -t ostype initiator_group [node ...]`**
	
	`-i` specifies that the igroup contains iSCSI node names.
	
	`-f` specifies that the igroup contains FCP WWPNs.
	
	`-t` *ostype* indicates the operating system type of the initiator. The values are: `default`, `solaris`, `windows`, `hpux`, `aix`, `linux`, `netware`, `or vmware`.
	
	For information about specifying the ostype of an igroup, see "About igroups" on page 60.
	
	initiator_group is the name you specify as the name of the igroup.
	
	node is a list of iSCSI node names or FCP WWPNs, separated by spaces.
	
	iSCSI example:
	
	`igroup create -i -t windows win_host5_group2 iqn.1991-05.com.microsoft:host5.domain.com`
	
	FCP example:
	
	`igroup create -f -t aix aix-igroup3 10:00:00:00c:2b:cc:92`

To do this...	Use this command...
Create a space-reserved LUN	**`lun create -s size -t ostype lun_path`** `-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 58. `-t` *ostype* 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 58. *lun_path* is the LUN's path name that includes the volume and qtree. **Example:** `lun create -s 5g -t windows /vol/vol2/qtree1/lun3` **Result:** A 5-GB LUN called /vol/vol2/qtree1/lun3 is accessible by a Windows host. Space reservation is enabled for the LUN.
Map the LUN to an igroup	**`lun map lun_path initiator_group [lun_id]`** *lun_path* is the path name of the LUN you created. *initiator_group* is the name of the igroup you created. *lun_id* 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:** `lun map /vol/vol2/qtree1/lun3 win_host5_group2 0` **Result:** Data ONTAP maps /vol/vol1/qtree1/lun3 to the igroup win_host5_group2 at LUN ID 0. **Example 2:** `lun map /vol/vol2/lun4 aix-igroup0` **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: `lun map: auto-assigned aix-igroup0=0`

To do this...	Use this command...
Display the LUNs you created	**lun show -v** -v provides additional information, such as the comment string, serial number, and LUN mapping. **Example:** lun show -v **Sample result:** /vol/vol1/qtree1/lun3 4g (4294967296) (r/w, online, mapped) Serial#: 0dCfh3bgaBTU Share: none Space Reservation: enabled Multiprotocol Type: aix Maps: aix-igroup0=0
Display the LUN ID mapping	**lun show -m** -m provides mapping information in a tabular format. **Sample result:** LUN path Mapped to LUN ID Protocol --- /vol/tpcc_disks/ctrl_0 aix_cluster 0 FCP /vol/tpcc_disks/ctrl_1 aix_cluster 1 FCP /vol/tpcc_disks/crash1 aix_cluster 2 FCP /vol/tpcc_disks/crash2 aix_cluster 3 FCP /vol/tpcc_disks/cust_0 aix_cluster 4 FCP /vol/tpcc_disks/cust_1 aix_cluster 5 FCP /vol/tpcc_disks/cust_2 aix_cluster 6 FCP

To do this...	Use this command...

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 Snapshot copies.
	
	Sample result:
	
	`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)`

Creating iSCSI LUNs on vFiler units for MultiStore

Prerequisite for creating LUNs on vFiler units

MultiStore vFiler technology is supported for the iSCSI protocol only. You must purchase a MultiStore license to create vFiler units. Then you can enable the iSCSI license for each vFiler to manage LUNs (and igroups) on a per vFiler basis.

Note

SnapDrive can only connect to and manage LUNs on the hosting storage system (vfiler0), not to vFiler units.

Guidelines for creating LUNs on vFiler units

Use the following guidelines when creating LUNs on vFiler units.

- The vFiler access rights are enforced when the storage system processes iSCSI host requests.
- LUNs inherit vFiler ownership from the storage unit on which they are created. For example, if /vol/vfstore/vf1_0 is a qtree owned by vFiler vf1, all LUNs created in this qtree are owned by vf1.
- As vFiler ownership of storage changes, so does ownership of the storage's LUNs.

LUN subcommands available on vFiler units

You can use the following LUN subcommands on vFiler LUNs:

`attribute`	`help`	`online`	`show`
`clone`	`map`	`resize`	`snap`
`comment`	`maxsize`	`serial`	`stats`
`create`	`geometry`	`set`	`unmap`
`destroy`	`move`	`setup`	
`df`	`offline`	`share`	

Note

You cannot use the `lun rescan` command for vFiler LUNs.

Methods for issuing LUN subcommands on vFiler units

You can issue LUN subcommands using the following methods:

- From the default vFiler (vfiler0) on the hosting storage system, you can do the following:
  - Enter the `vfiler run * lun` *`subcommand`*., which runs the `lun` *`subcommand`* on all vFiler units.
  - Run a LUN subcommand on a specific vFiler. To access a specific vfiler, you change the vfiler context by entering the following commands:

    `filer> vfiler context` *`vfiler_name`*

    `vfiler_name@filer> lun` *`subcommand`*
- From non-default vFiler units, you can:
  - Enter the `vfiler run * lun` command

Creating LUNs on a vFiler

To create LUNs on a vFiler, complete the following step.

Step	Action

1	Enter the `lun create` command in the vFiler context that owns the storage, as follows:
	`vfiler run vfiler_name lun create -s 2g -t os_type /vol/vfstore/vf1_0/lun0`
	Example: The following command creates a LUN on a vFiler at /vol/vfstore/vf1_0:
	`vfiler run vf1 lun create -s 2g -t windows /vol/vfstore/vf1_0/lun0`
	Note
	If you omit the `vfiler` command and the context, an error message is displayed.
	Example: The following command omits the `vfiler run` command and the storage context (vf1).
	`lun create -s 2g -t windows /vol/vfstore/vf1_0/lun0`
	Result: The following error message is displayed:
	`lun create: Requested LUN path is inaccessible.`

Displaying vFiler LUNs

To display LUNs owned by the vFiler context, complete the following step.

Step	Action

1	Enter the following command from the vFiler that contains the LUNs:
	`vfiler run * lun show`
	Result: The following information is an example of what is displayed.
	`==== vfiler0`
	`/vol/vfstore/vf0_0/vf0_lun0  2g  (21437483648)  (r/w, online)`
	`/vol/vfstore/vf0_0/vf0_lun1  2g  (21437483648)  (r/w, online)`
	`==== vfiler1`
	`/vol/vfstore/vf0_0/vf1_lun0  2g  (21437483648)  (r/w, online)`
	`/vol/vfstore/vf0_0/vf1_lun1  2g  (21437483648)  (r/w, online)`

Managing LUNs 5

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 82
- "Displaying LUN information" on page 88

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 AIX hosts, you use the `cfgmgr` command. For detailed procedures, see the documentation for your Host Support Kit or 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: **lun online *lun_path* [*lun_path* ...]** **Example:** `lun online /vol/vol1/lun0`

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: **lun offline *lun_path* [*lun_path* ...]** **Example:** `lun offline /vol/vol1/lun0`

Unmapping a LUN from an igroup

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

Step	Action
1	Enter the following command: **lun offline *lun_path*** **Example:** `lun offline /vol/vol1/lun1`
2	Enter the following command: **lun unmap *lun_path igroup* LUN_ID** **Example:** `lun unmap /vol/vol1/lun1 solaris-igroup0 0`

Renaming a LUN

To rename a LUN, completing the following step.

Step	Action
1	Enter the following command: **lun move *lun_path new_lun_path*** **Example:** `lun move /vol/vol1/mylun /vol/vol1/mynewlun` **Note** If you are organizing LUNs in qtrees, the existing path (`lun_path`) and the new path (`new_lun_path`) 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 of ten times the original size of the LUN. Data ONTAP also 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 Host Support Kit or 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: **`lun offline lun_path`** **Example:** `lun offline /vol/vol1/qtree/lun2`
2	Change the size of the LUN by entering the following command: **`lun resize [-f] lun_path new_size`** `-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) `lun resize /vol/vol1/qtree1/lun2 10g`
3	From the host, rescan or rediscover the LUN so that the new size is recognized.

Modifying the LUN description

To modify the LUN description, complete the following step.

Step	Action
1	Enter the following command: **`lun comment lun_path [comment]`** **Example:** `lun comment /vol/vol1/lun2 "10GB for payroll records"` **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. 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: **`lun set reservation lun_path [enable | disable]`** *lun_path* 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: **`lun destroy [-f] lun_path [lun_path ...]`** `-f` forces the `lun destroy` command to execute even if the LUNs specified by one or more *lun_paths* are mapped or are online. Without the `-f` parameter, you must first take the LUN offline and unmap it, and then enter the `lun destroy` command.

Accessing a LUN with NAS protocols

When you create a LUN, it can be accessed only with the iSCSI or FCP protocol 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 storage system. 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 using a NAS protocol, you must take the LUN offline or unmap it to prevent an iSCSI or FCP 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: **`lun share` *`lun_path`* `{none|read|write|all}`** **Example:** `lun share /vol/vol1/qtree1/lun2 read` **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 storage system's command line, enter the following command:
	`lun help`
	Result: A list of all LUN subcommands is displayed:
	`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: **lun help *subcommand*** **Example:** `lun help show`

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:
	`lun stats -z -i interval -c count -o [-a | lun_path]`
	`-z` zeros statistics
	Note
	The statistics start at zero at boot time.
	`-i` interval is the interval, in seconds, at which the statistics are displayed.
	`-c` *count* is the number of intervals. For example, the `lun stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.
	`-o` *displays* additional statistics, including the number of QFULL messages the storage system sends when its SCSI command queue is full and the amount of traffic received from the partner storage system.
	`-a` shows statistics for all LUNs.
	lun_path displays statistics for a specific LUN.
	Example:
	`lun stats -o -i 1`
	`Read Write Other QFull Read Write Average Queue Partner Lun`
	`Ops Ops Ops kB kB Latency Length Ops kB`
	`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 storage system's command line, enter the following command: **lun show -m** **Result:** LUN path Mapped to LUN ID Protocol ---------- /vol/tpcc/ctrl_0 host5 0 iSCSI /vol/tpcc/ctrl_1 host5 1 iSCSI /vol/tpcc/crash1 host5 2 iSCSI /vol/tpcc/crash2 host5 3 iSCSI /vol/tpcc/cust_0 host6 4 iSCSI /vol/tpcc/cust_1 host6 5 iSCSI /vol/tpcc/cust_2 host6 6 iSCSI

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 storage system's command line, enter the following command to display LUN status and characteristics:
	`lun show -v`
	Example:
	`/vol/tpcc_disks/cust_0_1 382m (400556032)  (r/w, online, mapped)`
	`Serial#: VqmOVYoe3BUf`
	`Share: none`
	`Space Reservation: enabled`
	`Multiprotocol Type: aix`
	`SnapValidator Offset: 1m (1048576)`
	`Maps: host5=0`
	`/vol/tpcc_disks/cust_0_2 382m (400556032)  (r/w, online, mapped)`
	`Serial#: VqmOVYoe3BV6`
	`Share: none`
	`Space Reservation: enabled`
	`Multiprotocol Type: aix`
	`SnapValidator Offset:  1m (1048576)`
	`Maps: host6=1`

Managing iSCSI igroups 6

About this chapter

This chapter explains how to create and manage igroups.

Topics in this chapter

This chapter discusses the following topics:

- "Managing igroups" on page 94
- "Using igroups on vFiler units" on page 97

Managing igroups

Tasks to manage igroups

You can use the command-line interface or FilerView to

- Create igroups
- Destroy igroups
- Add initiators (through their node names) to igroups
- Remove initiators (through their node names) from igroups
- Display all the initiators in an igroup
- Set the operating system type (ostype) for an igroup

Creating an igroup

To create an igroup, complete the following step.

Step	Action
1	Enter the following command: **`igroup create -i [-t ostype] initiator_group [nodename ...]`** `-i` indicates that it is an iSCSI igroup. `-t` *ostype* indicates the operating system of the host. The values are *default*, *solaris, windows, hpux, aix,* or *linux*. Use *default* if you are using another operating system. *initiator_group* is the name of the igroup you specify. *nodename* is an iSCSI nodename. You can specify more than one nodename. **Example:** `igroup create -i -t windows win-group0 iqn.1991-05.com.microsoft:eng1` **Result:** You created an igroup called win-group0 that contains the nodename of the Windows host associated with that nodename.

Destroying an igroup

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

Step	Action	
1	**If you want to...**	**Then enter this command...**
	Remove LUNs mapped to an igroup before deleting the igroup	**lun unmap *lun-path igroup*** **Example:** `lun unmap /vol/vol2/qtree/LUN10 win-group5`
	Delete one or more igroups	**igroup destroy igroup [*igroup,...*]** **Example:** `igroup destroy win-group5`
	Remove all LUN maps for an igroup and delete the igroup with one command	**igroup destroy -f *igroup* [*igroup ...*]** **Example:** `igroup destroy -f win-group5`

Adding an initiator

To add an initiator to an igroup, complete the following step.

Note

An initiator cannot be a member of igroups of differing types. For example, if you have an initiator that already belongs to a solaris igroup, you cannot add this initiator to an aix igroup.

Step	Action
1	Enter the following command: **igroup add *igroup nodename*** **CAUTION** When adding initiators to an igroup, ensure that each initiator sees, at most, one LUN at a given LUN ID. **Example:** `igroup add win-group2 iqn.1991-05.com.microsoft:eng2` **Result:** `You added the host associated with node name iqn.1991-05.com.microsoft:eng2 to the initiator group win-group2.`

Removing an initiator

To remove an initiator from an igroup, complete the following step.

Step	Action
1	Enter the following command: **igroup remove *igroup nodename*** **Example:** `igroup remove win-group1 iqn.1991-05.com.microsoft:eng1`

Displaying initiators

To display all the initiators in the specified igroup, complete the following step.

Step	Action
1	Enter the following command: **igroup show [*igroup*]** **Example:** `igroup show win-group3`

Setting the ostype

To set the operating system type (ostype) for an igroup to one of the values (`default`, `solaris`, `windows`, `hpux`, `aix`, `linux`, `netware`, and `vmware`), complete the following step.

Step	Action
1	Enter the following command. **igroup set *igroup ostype*** **Example:** `igroup set win-group3 windows`

Getting command-line syntax help

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

Step	Action
1	Enter the following command: **igroup help *subcommand***

Using igroups on vFiler units

How igroups work on vFiler units

igroups are owned by vFiler contexts. The vFiler ownership of igroups is determined by the vFiler context in which the igroup is created. You can create iSCSI igroups in non-default vFiler units.

Creating an igroup in a non-default vFiler unit

To create an igroup in a non-default vFiler unit, complete the following steps.

Step	Action
1	Change the context to the desired vFiler unit by entering the following command: **`myfiler> vfiler context vf1`** **Result:** The vFiler unit's prompt is displayed.
2	Create the igroup on vFiler unit determined by step 1 by entering the following command: **`vf1@myfiler> igroup create -i vf1_iscsi_group iqn.1991-05.com.microsoft:server1`**

3	Display the igroup by entering the following command:
	`vf1@myfiler> igroup show`
	Result: The following is a sample display.
	`vf1_iscsi_group (iSCSI) (ostype: default):`
	`iqn.1991-05.com.microsoft:server1`

Mapping LUNs to igroups

You must map LUNs to igroups that are in the same vFiler unit.

Managing FCP Initiator Groups 7

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 Fibre Channel 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 Data ONTAP command line

To create an igroup, complete the following step.

Step	Action
1	Enter the following command: **`igroup create -f [-t ostype] initiator_group [node_name...] [-a portset]`** `-f` indicates that it is an FCP igroup. `-t` *ostype* indicates the operating system of the host. The values are *solaris, windows, hpux, aix, linux*, *netware* or *vmware*. *initiator_group* is the name of the igroup you specify. *node_name* is an FCP WWPN. You can specify more than one WWPN. -a *portset* binds the igroup to a portset. A portset is a group of target FCP ports. When you bind an igroup to a portset, any host in the igroup can access the LUNs only by connecting to the target ports in the portset. For details about portsets, see "Making LUNs available on specific FCP target ports" on page 133. **Example:** `igroup create -f -t aix aix-group5 50:06:0b:00:00:10:a7:00 50:06:0b:00:00:10:a6:06`

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 storage system's command line.

To create an 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 /opt/NetApp/santools/bin directory.

3	Enter the following command to print a command to be run on the storage system that creates an igroup containing all the HBAs on your host:
	`./sanlun fcp show adapter -c`
	`-c` prints the full `igroup create` command on the screen.
	Result: An `igroup create` 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:`
	`igroup create -f -t aix "hostA" 10000000AA11BB22 10000000AA11EE33`
	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 `telnet` command to access the storage system.
5	Copy the `igroup create` command from Step 3, paste the command on the storage system's command line, and press Enter to run the igroup command on the storage system. **Result:** An igroup is created on the storage system.

Step	Action

6	On the storage system's command line, enter the following command to verify the newly created igroup:
	`igroup show`
	Result: The newly created igroup with the host's WWPNs is displayed.
	Example:
	`filerX> igroup show`
	`hostA (FCP) (ostype: aix):`
	`10:00:00:00:AA:11:BB:22`
	`10:00:00:00:AA:11:EE:33`

Destroying an igroup

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

Step	Action	
1	**If you want to...**	**Then enter this command...**
	Remove LUNs mapped to an igroup before deleting the igroup	**`lun unmap lun_path igroup`** **Example:** `lun unmap /vol/vol2/qtree/LUN10 aix-group5`
	Delete one or more igroups	**`igroup destroy igroup [igroup,...]`** **Example:** `igroup destroy aix-group5`
	Remove all LUN maps for an igroup and delete the igroup with one command	**`igroup destroy -f igroup [igroup ...]`** **Example:** `igroup destroy -f aix-group5`

Adding an initiator

To add an initiator to an igroup, complete the following step.

Step	Action
1	Enter the following command: **igroup add *igroup WWPN*** **CAUTION** When adding initiators to an igroup, ensure that each initiator sees only one LUN at a given LUN ID. **Example:** `igroup add aix-group2 10:00:00:00:c9:2b:02:1f` **Result:** You added the second port of Host2 to the igroup aix-group2.

Removing an initiator

To remove an initiator from an igroup, complete the following step.

Step	Action
1	Enter the following command: **igroup remove *igroup WWPN*** **Example:** `igroup remove aix-group1 10:00:00:00:c9:2b:7c:0f`

Displaying initiators

To display all the initiators in the specified igroup, complete the following step.

Step	Action
1	Enter the following command: **igroup show [*igroup*]** **Example:** `igroup show aix-group3`

Setting the ostype

To set the operating system type (ostype) for an igroup, complete the following step.

Step	Action
1	Enter the following command. **`igroup set igroup ostype value`** *igroup* is the name of the igroup. *value* is the ostype of the igroup. The ostypes of initiators are `solaris`, `windows`, `hpux`, `aix`, `linux`, `netware` and `vmware`. If your host OS is not one of these values but it is listed as a supported OS in the FCP compatibility matrix, specify `default`. For information about supported hosts and operating systems, see the *NetApp FCP SAN Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml. **Example:** `igroup set aix-group3 ostype aix`

Managing Fibre Channel initiator requests

Why you need to manage initiator requests

Each physical port on the target HBA in the storage system 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 storage system.
- 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 storage system'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

Any 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: **`igroup set` *`igroup_name`* `throttle_reserve` *`percentage`*** **Example:** `igroup set aix-igroup1 throttle_reserve 20` **Result:** The igroup throttle is created for aix-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: **`igroup set` *`igroup_name`* `throttle_reserve 0`**

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: **`igroup set igroup_name throttle_borrow {yes | no}`** **Example:** `igroup set aix-igroup1 throttle_borrow yes` **Result:** When you set the `throttle_borrow` setting to `yes`, 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:
	`igroup show -t`
	Sample output:
	`name reserved exceeds borrows`
	`aix-igroup1 20% 0 N/A`
	`aix-igroup2 10% 0 0`
	Explanation of output: The `exceeds` column displays the number of times the initiator sends more requests than the throttle allows. The `borrows` column displays the number of times the throttle is exceeded and the storage system uses queue resources from the unreserved pool. In the `borrows` column, "N/A" indicates that the `igroup throttle_borrow` option is set to `no`.

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:
	`igroup show -t -i interval -c count [igroup|-a]`
	-t displays information on igroup throttles.
	-i *interval* displays statistics for the throttles over an interval in seconds.
	-c *count* determines how many intervals are shown.
	igroup is the name of a specific group for which you want to show statistics.
	-a displays statistics for all igroups, including idle igroups.
	Example: `igroup show -t -i 1`
	Result: The following is a sample display:
	`name reserved 4a 4b 5a 5b`
	`igroup1 20% 45/98 0/98 0/98 0/98`
	`iqroup2 10% 0/49 0/49 17/49 0/49`
	`unreserved 87/344 0/344 112/344 0/344`
	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.

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: **`lun stats -o -i time_in_seconds`** `-i` *time_in_seconds* is the interval over which performance statistics are reported. For example, -i 1 reports statistics each second. `-o` displays additional statistics, including the number of QFULL messages. **Example:** `lun stats -o -i 1 /vol/vol1/lun2` **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: **`lun stats -o lun_path`**

How a cluster failover affects igroup throttles

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

FCP cfmode	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, dual_fabric, or single_image	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:
	`igroup show -t`
	Example: `igroup show -t`
	Result: The following is a sample display:
	`name reserved exceeds borrows`
	`aix-igroup1 20% 0 N/A (Reduced by takeover to 10%)`
	`aix-igroup2 10% 0 0 (Reduced by takeover to 5%)`

Managing FCP in a clustered environment 8

About this chapter

If your storage systems are in a cluster, Data ONTAP provides multiple modes of operation required to support homogeneous and heterogeneous host operating systems. 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.
- Provide access to local and partner LUNs in a cluster.

This chapter provides an overview of each cfmode setting and describes how to change the default cfmode for requirements of your configuration.

Note

The cfmode setting of your cluster and the number of paths available must align with your cabling, configuration limits, and zoning requirements. For information about different configurations, see the online *Fibre Channel Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPConfigurationGuide.pdf.

Topics in this chapter

This chapter discusses the following topics:

- "How FCP cfmode settings work" on page 112
- "Changing the cluster's cfmode setting" on page 131
- "Making LUNs available on specific FCP target ports" on page 133

How FCP cfmode settings work

Summary of cfmode settings and supported systems

The following table summarizes the cfmodes, supported systems, benefits and limitations.

cfmode	Supported Systems	Benefits and Limitations
partner	All systems except for FAS270c	◆ Supports all host OS types. ◆ Supports all switches.
single_image	All systems	◆ Supports all host OS types. ◆ Supports all switches. ◆ Makes all LUNs available on all target ports.
dual_fabric	FAS270c only	◆ Supports all host OS types. ◆ Requires fewer switch ports. ◆ Does not support all switches. Requires switches that support public loop.
standby	All systems, except FAS270c	◆ Requires more switch ports. ◆ Supports only Windows and Solaris hosts.
mixed	All systems, except FAS270c	◆ Supports all operating systems. ◆ Does not support all switches. Requires switches that support public loop.

cfmode settings on both cluster nodes

The FCP cfmode settings must be set to the same value for both nodes in a cluster. If the cfmode is not set to the same value, your hosts might not be able to access data stored on the system.

cfmode settings for new systems and upgrades

Partner is the default cfmode setting of a new system with a new installation of Data ONTAP 7.1 or later.

When you upgrade to Data ONTAP 7.1 or later, the cluster saves the cfmode setting from the previous version. For example, if the systems in your cluster were set to standby mode before the upgrade, the cfmode setting remains in standby mode after the upgrade.

Overview of partner mode

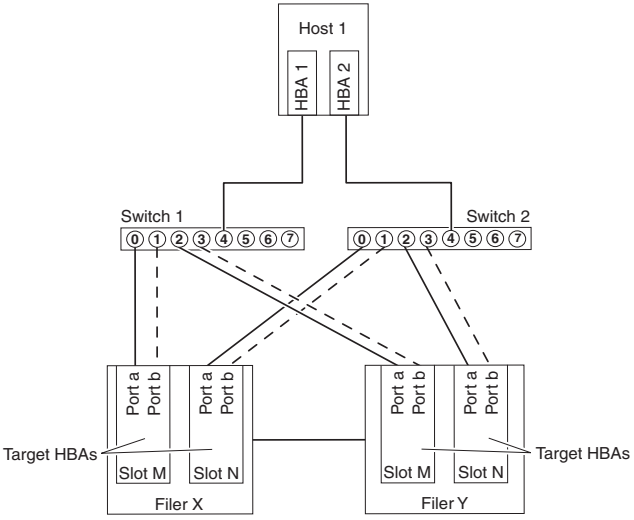
How target ports provide access to LUNs

The partner cfmode is the default setting for all new systems. It is supported on all FCP-licensed systems except for the FAS270c. It is also supported for all host OS types.

For systems with HBAs, 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 system. The target ports log into the fabric using a point-to-point topology.

If you have a FAS3000 series system with a new installation of DATA ONTAP, the state of the onboard Fibre Channel port depends on your configuration. In the default two-port configuration, ports 0c and 0d connect to the SAN. Port 0c provides access to local LUNs, and port 0d provides access to LUNs on the partner. In a four-port configuration in which all onboard ports connect to the SAN, port 0a and 0c on each node in the cluster provide access to local LUNs, and ports 0b and 0d provide access to LUNs on the partner.

The following figure shows a sample configuration with a multi-attached host connecting to a cluster with target HBAs. The solid lines represent paths to LUNs on the local storage system, the dotted lines represent paths to partner LUNs.



F8xx or FAS9xx cluster

Requirements for multipathing software

Partner mode requires hostside multipathing software. The multipathing policy is active/passive. The primary paths to the LUNs are always through the A ports. The B ports are secondary paths.

The following table shows the available paths between the host and the storage system cluster in the preceding configuration example.

	Type of path	Target FCP ports
Filer X LUN	Local/primary	Port a, Slot M on Filer X Port a, Slot N on Filer X
	Partner/secondary	Port b, Slot M on Filer Y Port b, Slot N on Filer Y
Filer Y LUN	Local/primary	Port a, Slot M on Filer Y Port a, Slot N on Filer Y
	Partner/secondary	Port b, Slot M on Filer X Port b, Slot N on Filer X

What happens when a link fails

When a link fails—for example, Switch 1 in the preceding example fails—the host loses a primary path (through the A port of the HBA in Slot M) to Filer X. The host fails over to the other primary paths. If there are no other primary paths available, the host can access LUNs through the secondary paths (B ports) on Filer Y.

The failover method depends on the host and multipathing software. For example, if you have VERITAS® Volume Manager (VxVM) with Dynamic Multipathing software and the Array Support Library (ASL) on a Solaris host, all LUNs that share active paths form a group. If all active paths fail for a LUN in a group, all LUNS in the group fail over to the secondary paths. For detailed information about how each host handles failover, see the ASL documentation for your FCP host attach kit.

What happens during a takeover

If Filer Y takes over for Filer X, the host continues to access LUNs on Filer X through the B ports on Filer Y. The WWNN and WWPN of the B ports on Filer Y do not change. This enables HP-UX and AIX hosts, which track target devices based on WWPN/WWNN and N_Port ID (the switch-assigned addresses), to maintain correct information about available paths.

How Data ONTAP displays target ports in 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 storage system in the cluster. For example, port B on Filer X presents the WWNN of Filer Y. This means that the WWNNs and WWPNs do not change during a takeover. The static address requirements for the N_Port ID, S_ID and D_ID are maintained.

The following `fcp config` command output shows how Data ONTAP displays WWNN and WWPN when the storage system'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
```

Overview of single_image mode

How cluster nodes are identified

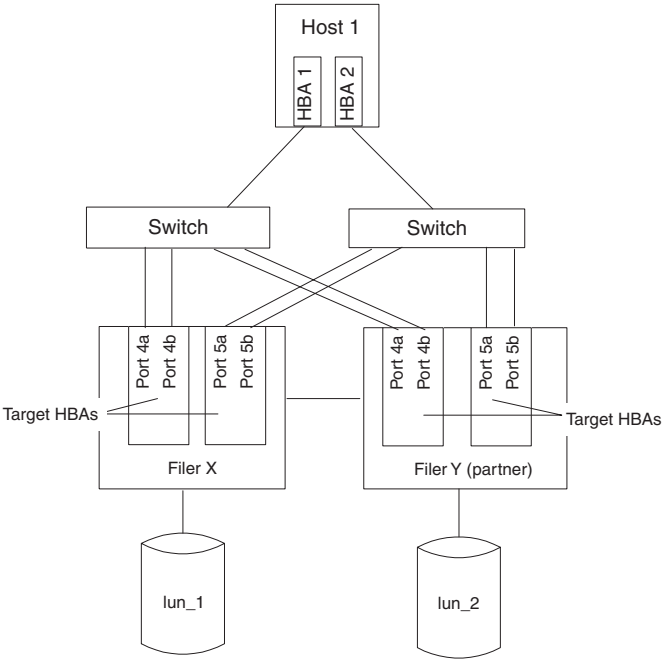
The single_image cfmode setting is available starting with Data ONTAP 7.1. In single_image mode, a cluster has a single global WWNN, and both systems in the cluster function as a single Fibre Channel node. Each node in the cluster shares the partner node's LUN map information.

How target ports provide access to LUNs

All LUNs in the cluster are available on all ports in the cluster by default. As a result, there are more paths to LUNs stored on the cluster, because any port on each node can provide access to both local and partner LUNs. You can specify the LUNs available on a subset of ports by defining portsets and binding them to an igroup. Any host in the igroup can access the LUNs only by connecting to the target ports in the portset.

For information about using portsets, see "Making LUNs available on specific FCP target ports" on page 133.

The following figure shows an example configuration with a multi-attached host. If the host accesses lun_1 through ports 4a, 4b, 5a, or 5b on Filer X, then Filer X recognizes that lun_1 is a local LUN. If the host accesses lun_1 through any of the ports on Filer Y, lun_1 is recognized as a partner LUN and Filer Y sends the SCSI requests to Filer X over the cluster interconnect.



LUN and igroup mapping conflicts

Each node in the cluster shares its partner's igroup and LUN mapping information. Data ONTAP uses the cluster interconnect to check igroup and LUN mapping information and also provides the mechanisms for avoiding mapping conflicts.

igroup ostypes: When you add an initiator WWPN to an igroup, Data ONTAP verifies that there are no igroup ostype conflicts. An example ostype conflict occurs when an initiator with the WWPN 10:00:00:00:c9:2b:cc:39 is a member of an AIX igroup on one node in the cluster and the same WWPN is also a member of an group with the default ostype on the partner.

Reserved LUN ID ranges: The LUN ID range on each storage system is divided into three areas:

- IDs 0 to 192 are shared between the nodes. You can map a LUN to an ID in this range on either node in the cluster.
- IDs 193 to 224 are reserved for one storage system.
- IDs 225 to 255 are reserved for the other storage system in the cluster.

By reserving LUN ID ranges on each storage system, Data ONTAP provides a mechanism for avoiding mapping conflicts. If the cluster interconnect is down, and you try to map a LUN to a specific ID, the `lun map` command fails. If you do not specify an ID in the `lun map` command, Data ONTAP automatically assigns one from a reserved range.

Bringing LUNs online: The `lun online` command fails when the cluster interconnect is down to avoid possible LUN mapping conflicts.

Overriding possible mapping conflicts: When the cluster interconnect is down, Data ONTAP cannot check for LUN mapping or igroup ostype conflicts. To avoid mapping conflicts, the following commands fail unless you use the `-f` option to force these commands:

- `lun map`
- `lun online`
- `igroup add`

You might want to override possible mapping conflicts in disaster recovery situations or situations in which the partner in the cluster cannot be reached and you want to regain access to LUNs. For example, the following command maps a LUN to an AIX igroup and assigns a LUN ID of 5, regardless of any possible mapping conflicts:

```
lun map -f /vol/vol2/qtree1/lun3 aix_host5_group2 5
```

Requirements for multipathing software

Multipathing software is required on the host so that SCSI commands fail over to alternate paths when links go down because of switch failures or cluster failovers. In the event of a failover, none of the adapters on the takeover storage system assume the WWPNs of the failed storage system.

How Data ONTAP displays target ports

The following `fcp config` output shows how Data ONTAP displays target ports when the cluster is in single_image mode and in normal operation. Each system has two adapters. Note that all ports show the same WWNN (node name), and the mediatype of all adapter ports is set to auto. This means that the ports log into the fabric using point-to-point mode. If point-to-point mode fails, then the ports try to log into the fabric in loop mode. You can use the `fcp config mediatype` command to change the default mediatype of the ports from auto to another mode according to the requirements of your configuration.

```
filer_a> fcp config
4a:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011f00
        portname 50:0a:09:81:82:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
4b:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011700
        portname 50:0a:09:82:82:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
5a:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011e00
        portname 50:0a:09:83:82:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
5b:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011400
        portname 50:0a:09:84:82:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
filer_b> fcp config
4a:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011e00
        portname 50:0a:09:81:92:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
4b:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011400
        portname 50:0a:09:82:92:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
5a:   ONLINE <ADAPTER UP>  PTP  Fabric
        host address 011f00
        portname 50:0a:09:83:92:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
5b:   ONLINE <ADAPTER UP>  Loop  Fabric
        host address 0117da
        portname 50:0a:09:84:92:00:96:d5  nodename 50:0a:09:80:82:00:96:d5
        mediatype auto
```

Overview of standby mode

How target ports provide access to LUNs

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 system in the cluster fails, Port B on the partner system becomes active and provides access to the LUNs on the failed system. The Port B assumes the WWPN of the Port A on the failed partner. The ports log in to the fabric in point-to-point mode.

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.

Limitations of standby mode

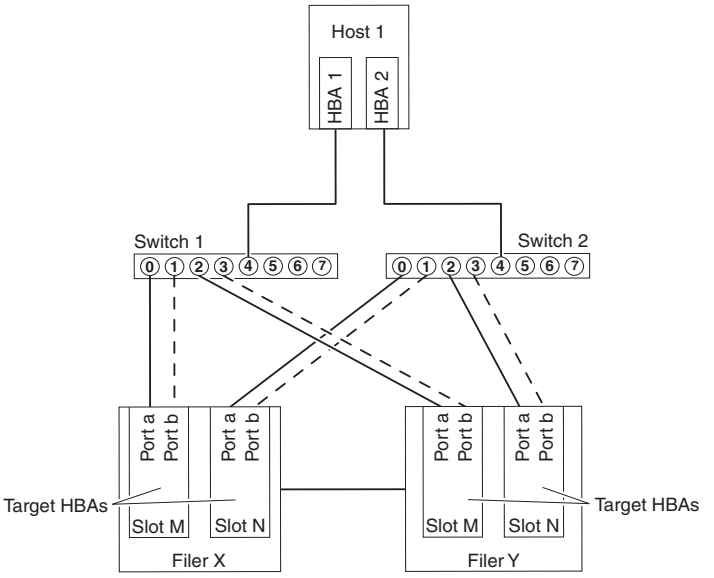
The standby cfmode setting is supported on all FCP-licensed systems except for the FAS270c.

Only Solaris and Windows igroup types can access a system in standby mode. The HP-UX, AIX, and default igroup types are not supported.

This setting also requires more switch ports because Port A and Port B on each HBA must connect to the switch, even though B becomes active only in the event of a takeover.

What happens during a failover

The following example shows a configuration in which Port B operates as a standby port. The storage system cluster pair shows each storage system with two target HBAs, in slots M and N. For each storage system, the slot-M HBA connects to Switch 1 and slot-N HBA connects to Switch 2. The solid lines indicate active connections. The dotted lines indicate standby connections.



If Filer X fails, then Filer Y takes over and the following occurs:

- Slot-M-port b on Filer Y takes over for slot-M-port a in Filer X
- Slot-N-port b on Filer Y takes over for slot-N-port a in Filer X

Port B on each HBA in Filer Y becomes active and enables the host to access the storage until Filer X is repaired and running. Each B port assumes the WWNN and WWPN of the corresponding A port on the failed storage system.

How Data ONTAP displays information about target ports

When the FCP cfmode setting is standby, the local WWNN and WWPN 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 WWPN. The standby WWNN and WWPN have a pattern of 20:01:00:*nn:nn:nn:nn:nn*.

The following `fcp config` output shows target port information for a storage system 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
```

Overview of dual_fabric mode

How target ports provide access to LUNs

The dual_fabric mode is supported on FAS270c clusters only. It is not supported on other systems.

The FAS270c cluster consists of two storage systems integrated into a disk shelf. Each storage system has two Fibre Channel ports. The orange-labeled port operates as a Fibre Channel target port after you license the FCP service and reboot the storage system. The blue-labeled port connects to the internal disks, enabling you to connect additional disk shelves to an FAS270c cluster. The Fibre Channel target port of each FAS270c storage system in the cluster supports three virtual ports:

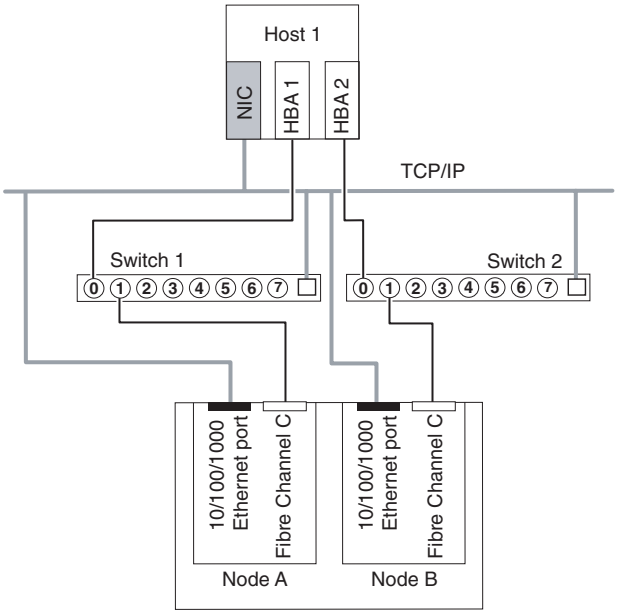
- Virtual local port, which provides access to LUNs on the local FAS270c
- Virtual standby port, which is not used
- Virtual partner port, which provides access to LUNs on the partner node

Note

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

Configurations with the FAS270c require that multipathing software be installed on the hosts.

The following figure shows the recommended production configuration, in which a multi-attached host accesses a FAS270c cluster.



What happens during a failover

The virtual ports enable the FAS270c cluster, which has only one physical target FCP port per node, to support traffic for both nodes in the cluster. For example, if Switch 1 fails, the multipathing software on the host uses HBA 2 and Switch 2 to access the partner virtual port on Node B. The partner virtual port forwards the requests to Node A. When the cluster is in a takeover state (Node B takes over for Node A), Node B uses the partner virtual port to directly serve data for LUNs on Node A.

How Data ONTAP displays target port information

State of Fibre Channel Port 1: When the FAS270c port labeled Fibre Channel 1 operates as a SAN target, the `sysconfig -v` command shows this port as a Fibre Channel Target Host Adapter installed in slot 0. The following example shows `sysconfig -v` output for a FAS270c in SAN target mode.

```
fas270> sysconfig
        NetApp Release 7.1: Mon June 25 02:20:04 PDT 2005
        System ID: 0084165669 (filer_1); partner ID: 0084165671 (filer_2)
        System Serial Number: 379589 (filer_1)
        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: Fibre Channel Target Host Adapter 0c
        slot 0: SB1250-Gigabit Dual Ethernet Controller
                e0a MAC Address:    00:a0:98:00:d5:90 (100tx-fd-up)
                e0b MAC Address:    00:a0:98:00:d5:91 (auto-unknown-cfg_down)
        slot 0: NetApp ATA/IDE Adapter 0a (0x00000000000001f0)
                0a.0                 122MB
```

Target virtual ports: FCP commands display three virtual ports for the one FAS270c physical target FCP port. For example, the `fcp show adapter` command shows the target FCP port as an adapter in slot 0c. The local virtual port is 0c_0, the standby virtual port is 0c_1, and the partner virtual port is 0c_2.

The following output example shows how the `fcp config` command displays virtual ports when the storage system is in dual_fabric mode and the cluster is in normal operation.

```
FAS270> fcp config
0c:   ONLINE <ADAPTER UP>  Loop  Fabric
        host address 0100da
        portname 50:0a:09:81:85:c4:45:91  nodename 50:0a:09:80:85:c4:45:91
        mediatype loop  partner adapter 0c
0c_0: ONLINE Local
        portname 50:0a:09:81:85:c4:45:91  nodename 50:0a:09:80:85:c4:45:91
        loopid 0x7 portid 0x0100da
0c_1: OFFLINE Standby
        portname 50:0a:09:81:85:c4:45:88  nodename 50:0a:09:80:85:c4:45:88
        loopid 0x0 portid 0x000000
0c_2: ONLINE Partner
        portname 50:0a:09:89:85:c4:45:88  nodename 50:0a:09:80:85:c4:45:88
        loopid 0x9 portid 0x0100d6
```

Overview of mixed mode

How target ports provide access to LUNs

The mixed mode setting is supported on all systems except for the FAS270c. Each FCP target port supports three virtual ports:

- Virtual local port, which provides access to LUNs on the local system.
- Virtual standby port, which provides access to LUNs on the failed system when a takeover occurs. The standby virtual port assumes the WWPN of the corresponding port on the failed partner.
- Virtual partner port, which provides access to LUNs on the partner system. 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.

AIX or HP-UX hosts connected to a cluster in mixed mode must have multipathing software installed. For information about the multipathing software supported for your host, see the documentation for your FCP Attach Kit.

The following output example shows how the `fcp config` command displays virtual ports when the storage system is in mixed mode and the cluster is in normal operation.

```
filer> fcp config
9a:   ONLINE <ADAPTER UP>  Loop  Fabric
        host address 021bda
        portname 50:a9:80:01:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
        mediatype loop  partner adapter 9a
9a_0: ONLINE Local
        portname 50:a9:80:01:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
        loopid 0x7 portid 0x021bda
9a_1: OFFLINE Standby
        portname 50:a9:80:01:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
        loopid 0x0 portid 0x000000
9a_2: ONLINE Partner
        portname 50:a9:80:09:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
        loopid 0x9 portid 0x021bd6
9b:   ONLINE <ADAPTER UP>  Loop  Fabric
        host address 021ada
        portname 50:a9:80:02:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
        mediatype loop  partner adapter 9b
9b_0: ONLINE Local
        portname 50:a9:80:02:03:00:e0:73  nodename 50:a9:80:00:03:00:e0:73
        loopid 0x7 portid 0x021ada
9b_1: OFFLINE Standby
        portname 50:a9:80:02:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
        loopid 0x0 portid 0x000000
9b_2: ONLINE Partner
        portname 50:a9:80:0a:03:00:e0:5f  nodename 50:a9:80:00:03:00:e0:5f
        loopid 0x9 portid 0x021ad6
```

Changing the cluster's cfmode setting

Reasons for changing the cfmode setting

Starting with version 7.1, Data ONTAP implements the use of the single_image mode for clustered systems. This cfmode setting provides the following advantages:

- The host can access all LUNs through any target port on the cluster.
- The single_image mode is supported on all storage systems.
- The single_image mode is compatible with all supported FCP hosts and clustered storage systems that support the FCP protocol. There are no switch limitations. You can connect a storage system cluster in single_image mode to any FCP switch supported by Data ONTAP.

You might also want to change your cfmode setting to for the following reasons:

- To increase the number of paths to the cluster— For example, you might want to use the single_image setting to increase the number of available paths in your configuration. All ports on each storage system are available to access local LUNs. In addition, single_image node is available for all supported SAN hosts.
- To support heterogeneous configurations—For example, you upgraded an existing SAN configuration with Solaris or Windows hosts and you want to add HP-UX or AIX hosts to the configuration.

  HP-UX and AIX hosts do not support failover when the storage systems in a cluster are in standby mode. If you want to add these hosts to an existing homogenous configuration that has Solaris-only or Windows-only hosts, you have to change the cfmode setting on the storage system to either single_image or partner.

  Solaris or Windows hosts in the same environment as HP-UX or AIX hosts must also have multipathing software installed when the cfmode setting is single_image or partner.

Impact of changing the cfmode setting

When you change the cfmode setting of the cluster to single_image, the following configuration components are affected.

- Host access to LUNs

  Hosts cannot access data on mapped LUNs. When you change the cfmode setting, you change the available paths between the host and the storage system cluster. Some previously available paths are no longer available and some new paths become available. The LUNs might be accessible but cannot

be used until you reconfigure the host to discover the new paths. You must reconfigure every host that is connected to the cluster to discover the new paths. The LUNs are not accessible until you reconfigure the host. The procedure depends on your host operating system.

- Multipathing software

  If you have multipathing software in your configuration, changing the cfmode setting might also affect the multipathing policy.

- Switch zoning

  When you change the cluster's cfmode setting to single_image cfmode both nodes in the cluster use the same WWNN. One node assumes the WWNN of its partner. If your filer connects to a switch by using soft zoning (zoning by WWPN), you must update your zones to accommodate the WWNN change. Systems that connect to switches using hard (port) zoning are not affected.

- Cabling

  The single_image setting makes more target ports available to the host. This means you might have to change your cabling configuration.

Scheduling downtime

Changing the cfmode setting on the storage system requires host reconfiguration and in some cases you might have to reboot the host. The procedures also require you to quiesce host I/O and take host applications offline. It is recommended that you schedule downtime for your configuration before you change cfmode settings.

For detailed procedures

For procedures that describe how to change your cluster's cfmode setting to single_image mode, see the following documents:

- FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.
- The documentation for your Fibre Channel Host Attach Kit.

Making LUNs available on specific FCP target ports

What portsets are

A portset consists of a group of FCP target ports. You bind a portset to an igroup, to make the LUN available only on a subset of the storage system's target ports. Any host in the igroup can access the LUNs only by connecting to the target ports in the portset.

If an igroup is not bound to a portset, the LUNs mapped to the igroup are available on all of the storage system's FCP target ports. By using portsets, you can selectively control which initiators can access LUNs and the ports on which they access LUNs.

You use portsets for LUNs that are accessed by FCP hosts only. You cannot use portsets for LUNs accessed by iSCSI hosts.

How portsets are used in storage system clusters

Portsets are supported only with the single_image cfmode setting. They are not supported with other cfmode settings. The single_image setting makes all ports on both systems in the cluster visible to the hosts. You use portsets to fine-tune which ports are available to specific hosts and limit the amount of paths to the LUNs to comply with the limitations of your multipathing software.

The single_image mode is not the default cfmode setting for a new system or for an upgrade. Before you use portsets, you must change your cfmode setting to single_image.

Note

Make sure your portset definitions and igroup bindings align with the cabling and zoning requirements of your configuration.

For detailed information about single_image mode, see "How Data ONTAP supports FCP with clustered systems" on page 25.

For procedures that describe how to change your cluster's cfmode setting to single_image mode, see the following documents:

- FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.
- The documentation for your Fibre Channel Host Attach Kit.

How upgrades affect igroups and portsets

When you upgrade to Data ONTAP 7.1, all ports are visible to all initiators in the igroups until you create portsets and bind them to the igroups.

How portsets affect igroup throttles

Portsets enable you to control queue resources on a per-port basis. If you assign a a throttle reserve of 40 percent to an igroup that is not bound to a portset, then the initiators in the igroup are guaranteed 40 percent of the queue resources on every target port. If you bind the same igroup to a portset, then the initiators in the igroup have 40 percent of the queue resources only on the target ports in the portset. This means that you can free up resources on other target ports for other igroups and initiators.

Before you bind new portsets to an igroup, verify the igroup's throttle reserve setting by using the `igroup show -t` command. It is important to check existing throttle reserves because you cannot assign more than 99 percent of a target port's queue resources to an igroup. When you bind more than one igroup to a portset, the combined throttle reserve settings might exceed 100 percent.

Example: igroup_1 is bound to portset_1, which includes ports 4a and 4b on each system in the cluster (FilerA:4a, FilerA:4b, FilerB:4a, FilerB:4b). The throttle setting of igroup is 40 percent.

You create a new igroup (igroup_2), with a throttle setting of 70 percent. You bind igroup_2 to portset_2, which includes ports 4b on each system in the cluster (FilerA:4b, FilerB:4b). The throttle setting of the igroup is 70 percent. In this case, ports 4b on each system are overcommitted. Data ONTAP prevents you from binding the portset and displays a warning message prompting you to change the igroup throttle settings.

It is also important to check throttle reserves before you unbind a portset from an igroup. In this case, you make the ports visible to all igroups that are mapped to LUNs. The throttle reserve settings of multiple igroups might exceed the available resources on a port.

For details about igroup throttles, see "Managing Fibre Channel initiator requests" on page 105.

Creating a portset

For clustered systems, it is recommended that when you add local ports to a portset, that you also add the partner system's corresponding target ports to the same portset. For example, if you have local systems's target port 4a port in the

portset, then make sure to include the partner system's port 4a in the portset as well. This ensures that cluster takeover and giveback occurs without connectivity problems.

To create a portset, complete the following step.

Step	Action
1	Enter the following command: **`portset create -f portset_name port...`** `-f` creates an FCP portset. *portset_name* is the name you specify for the portset. You can specify a string of up to 95 characters. *port* is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty portset. You can add as many as 18 target FCP ports. Check your configuration limits in the *Fibre Channel Configuration Guide* to understand the maximum number of ports supported for your setup. You specify a port by using the following formats: ◆ *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in a cluster, the port from both the local and partner storage system is added to the portset. ◆ *filername:slotletter* adds only a specific port on a storage system—for example, filerA:4b.

Binding an igroup to a portset

If you do not bind an igroup to a portset, and you map a LUN to the igroup, then the initiators in the igroup can access the LUN on any port on the storage system. To bind an igroup to a portset, complete the following step.

Step	Action
1	Enter the following command: **`igroup bind igroup_name portset_name`** *igroup_name* is the name of the igroup. *portset_name* is the name of the portset.

Adding or removing ports from a portset

To add a port to a portset, complete the following step.

Step	Action
1	Enter the following command. **`portset add portset_name port...`** *portset_name* is the name of the portset. *port* is the target FCP port. You can specify more than one port. You specify a port by using the following formats: ◆ *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in a cluster, the port from both the local and partner system is added to the portset. ◆ *filername:slotletter* adds only a specific port on a system—for example, filerA:4b.

To remove a port from a portset, complete the following step.

Step	Action
1	Enter the following command. **`portset remove portset_name port...`** *portset_name* is the name of the portset. *port* is the target FCP port. You can specify more than one port. You specify a port by using the following formats: ◆ *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in a cluster, the port from both the local and partner system is removed from the portset. ◆ *filername:slotletter* removes only a specific port on a system—for example, filerA:4b.

Unbinding an igroup from a portset

If you unbind or unmap an igroup from a portset, then all hosts in the igroup can access LUNs on all target ports. To unbind an igroup from a portset, complete the following step.

Step	Action
1	Enter the following command. **`igroup unbind`** ***`igroup_name`*** *igroup_name* is the name of the igroup.

Viewing the ports in a portset

To view the ports in a portset, complete the following step.

Step	Action
1	Enter the following command. **`portset show [`*`portset_name`*`]`** If you do not supply *portset_name*, all portsets and their respective ports are listed. If you supply *portset_name*, only the ports in the portset are listed.

Showing igroup-to-portset bindings

To show which igroups are bound to portsets, complete the following step.

Step	Action
1	Enter the following command. **`igroup show`** ***`igroup_name`***

Destroying a portset

Before you destroy a portset, make sure you unbind it from any igroups. To destroy a portset, complete the following step.

Step	Action
1	Unbind the portset from any igroups by entering the following command: **`igroup unbind` *`igroup_name portset_name`***
2	Enter the following command. **`portset destroy [-f]` *`portset_name...`*** You can specify a list of portsets. If you use the `-f` option, you destroy the portset even if it is still bound to an igroup. If you do not use the `-f` option and the portset is still bound to an igroup, the `portset destroy` command fails.

Managing Disk Space 9

About this chapter

This chapter describes how monitor available disk space and how to define a space management policy—how to configure Data ONTAP to automatically grow a FlexVol volume or delete Snapshot copies when the FlexVol volume begins to run out of free space.

Topics in this chapter

This chapter discusses the following topics

- "Monitoring disk space" on page 140
- "Defining a space management policy" on page 152

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 Snapshot copies in a volume. For detailed information, see "Estimating the data change rate between Snapshot copies" below.
- `snap reclaimable`—Estimates the amount of space freed if you delete the specified Snapshot copies. If space in your volume is scarce, you can reclaim free space by deleting a set of Snapshot copies. For detailed information, see "Estimating the amount of space freed by Snapshot copies" on page 142.
- `df`—Displays the statistics about the active file system and the Snapshot copy directory in a volume or aggregate. For detailed information, see "Displaying statistics about free space" on page 142.

Estimating the data change rate between Snapshot copies

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 change rate: To display the rate of change of data between Snapshot copies, complete the following steps:

Step	Action
1	Enter the following command: **`snap delta [-A] [vol_name [begin_snapshot] [end_snapshot]]`** `-A` displays the rate of change of data between Snapshot copies for all aggregates in the system. *vol_name* is the name of the volume. *begin_snapshot* is the name of the first Snapshot copy to report data for. *end_snapshot* is the names of the last Snapshot copy to report data for. If you do not specify an argument, the `snap delta` command displays the rate of change of data between Snapshot copies for all volumes in the system. **Example:** The following example displays the rate of change of data between all Snapshot copies in vol0.

```
filer_1> snap delta vol0
Volume vol0
working...
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

Summary...
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 copy and the active file system. The following rows provide the rate of change between successive Snapshot copies. Each row displays the names of the two Snapshot copies that

are compared, the amount of data that has changed between them, the time elapsed between the two Snapshot copies, and how fast the data changed between the two Snapshot copies.

If you do not specify any Snapshot copies 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 copy and the active file system.

Estimating the amount of space freed by Snapshot copies

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

Step	Action

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

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 storage system. 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 copy 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 Snapshot copies + 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 copy 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 Snapshot copies is also counted in the used space of the active file system.

How LUN and Snapshot copy 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 Snapshot copies 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 copy when it was created, the LUN will always be less than one-half of the volume size.

Action	Result	Comment
Write 40 GB of data to the LUN.	Used space = 40 GB Reserved space = 0 GB Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	The amount of used space does not change because with space reservations set to On, the same amount of space is used when you write to the LUN as when you created the LUN.
Create a Snapshot copy of the LUN.	Used space = 80 GB Reserved space = 40 GB Available space = 20 GB Volume Total: 100 GB Snapshot succeeds.	The Snapshot copy locks all the data on the LUN so that even if that data is later deleted, it remains in the Snapshot copy until the Snapshot copy is deleted. As soon as a Snapshot copy 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 copy 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 copy, and 40 GB of reserved space. Free space becomes available again, so Snapshot copy creation is no longer blocked.

Action	Result	Comment
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 copy, it disappears when the new data replaces it. As a result, the total amount of used space remains unchanged.
Create a Snapshot copy of the LUN.	Used space = 120 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	The Snapshot copy 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 copy, and 40 GB for reserved space. The used space does not increase because data has not changed.
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 copy, 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 copy creation is blocked.
Delete both Snapshot copies 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 Snapshot copies disappears from the used total when the Snapshot copies are deleted. Because there are no more Snapshot copies 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 Snapshot copies 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 storage system or to perform tasks on the host.

- Without using Snapshot copies
- Using Snapshot copies
- Using backing store LUNs and LUN clones

In the examples, assume that the storage system is named *toaster*.

Monitoring disk space without using Snapshot copies: The following example illustrates how to monitor disk space on a volume when you create a LUN without using Snapshot copies. For this example, assume that you require less than the minimum capacity based on the 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 storage system, create a new traditional volume named *volspace* that has approximately 67 GB, and observe the effect on disk space by entering the following commands:
	toaster> vol create volspace 2
	toaster> df -r /vol/volspace
	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.
	`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`

Step	Action

2	Set the percentage of snap reserve space to zero and observe the effect on disk space by entering the following commands:
	toaster> snap reserve volspace 0
	toaster> df -r /vol/volspace
	Result: The following sample output is displayed. The amount of available Snapshot copy space becomes zero, and the 20 percent of Snapshot copy space is added to available space for /vol/volspace.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 1440 62648468 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volspace/.snapshot`

3	Create a LUN (/vol/volspace/lun0) and observe the effect on disk space by entering the following commands:
	toaster> lun create -s 3g -t aix /vol/volspace/lun0
	toaster> df -r /vol/volspace
	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.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 3150268 59499640 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volspace/.snapshot`

4	Create an igroup named *aix_host* and map the LUN to it by entering the following commands (assuming that the host node name is iqn.1996-04.aixhost.host1). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space. **toaster> igroup create -i -t aix aix_host iqn.1996-04.aixhost.host1** **toaster> lun map /vol/volspace/lun0 aix_host 0**
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 documentation from your Host Support Kit or Host Attach Kit. These commands have no effect on disk space.

Step	Action

6	From the storage system, ensure that creating the file system on the LUN and writing data to it has no effect on space on the storage system by entering the following command:
	toaster> df -r /vol/volspace
	Result: The following sample output is displayed. From the storage system, the amount of space used by the LUN remains 3 GB.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 3150268 59499640 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/vol`
	`space/.snapshot`

7	Turn off space reservations and see the effect on space by entering the following commands:
	toaster> lun set reservation /vol/volspace/lun0 disable
	toaster> df -r /vol/volspace
	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.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 144 62649584 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/vols`
	`pace/.snapshot`

Monitoring disk space using Snapshot copies

The following example illustrates how to monitor disk space on a volume when taking Snapshot copies. 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 storage system, create a new volume named *volspace* that has approximately 67 GB and observe the effect on disk space by entering the following commands:
	toaster> vol create volspace 2
	toaster> df -r /vol/volspace
	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.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace 50119928 1440 50118488 0 /vol/volspace/`
	`/vol/volspace/.snapshot 12529980 0 12529980 0 /vol/vol s pace/.snapshot`

2	Set the percentage of snap reserve space to zero by entering the following command: **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:
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 6300536 56169372 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volsp ace/.snapshot`

4	Create an igroup named *aix_host* and map the LUN to it by entering the following commands (assuming that the host node name is iqn.1996-04.aixhost.host1). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space. **toaster> igroup create -i -t aix aix_host iqn.1996-04.aixhost.host1** **toaster> lun map /vol/volspace/lun0 aix_host 0**
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 documentation from your Host Support Kit or Host Attach Kit. These commands have no effect on disk space.

Step	Action
6	From the host, write data to the file system (the LUN on the storage system). This has no effect on disk space.

7	Take a Snapshot copy named *snap1* 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:
	toaster> snap create volspace snap1
	toaster> df -r /vol/volspace
	
	Result: The following sample output is displayed. The first Snapshot copy 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 copy. 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.
	
	`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 storage system, observe the effect on disk space by entering the following commands:
	toaster> df -r /vol/volspace
	
	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 copy 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 copy. Notice that the used space increases for the Snapshot copy by 1 GB, and the available space for the volume decreases by 1 GB.
	
	`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 copy named *snap2* 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 copy reserved enough space to overwrite every block, only 44 blocks are used to account for the second Snapshot copy.
	
	`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 copy.
	
	`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`

Defining a space management policy

What space management policies are

A space management policy enables you to automatically reclaim space for a FlexVol volume when that volume is nearly full. You can configure a FlexVol volume to automatically reclaim space by using the following policies:

- Grow a FlexVol volume automatically when it is nearly full.

  This policy is useful if the containing aggregate has enough space to grow the FlexVol volume. You can grow a volume in increments and set a maximum size for the volume.

- Automatically delete Snapshot copies when the FlexVol volume is nearly full.

  For example, you can automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want to delete first—your oldest or newest Snapshot copies. You can also determine when to begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot copy reserve is nearly full.

You can define the order in which you want to apply these policies when a FlexVol volume is running out of space. For example, you can automatically grow the volume first, and then begin deleting Snapshot copies, or you can reclaim space by first automatically deleting Snapshot copies, and then growing the volume.

Configuring a volume to grow automatically

To configure a volume to grow automatically, complete the following step.

Step	Action
1	Enter the following command: **`vol autosize vol-name [-m size] [-i size] [on|off|reset]`** *vol-name* is the name of the FlexVol volume. You cannot use this command on traditional volumes. `-m` *size* is the maximum size to which you can grow the volume. Specify a size in `k` (KB), `m` (MB), `g` (GB) or `t` (TB). The volume does not grow if its size is equal to or greater than the maximum size. `-i` *size* is the increment by which you grow the volume. Specify a size in `k` (KB), `m` (MB), `g` (GB) or `t` (TB). `on` enables the volume to automatically grow. `off` disables automatically growing the volume. By default, the `vol autosize` command is set to Off. `reset` restores the `autosize` settings of the volume to the default setting, which is Off.

Automatically deleting Snapshot copies

Defining a policy for deleting Snapshot copies: To define a policy, complete the following step:

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td>1</td><td colspan="2">Enter the following command:
<tt>snap autodelete vol-name option value</tt>
Define which Snapshot copies you delete by entering the following options and their values:</td></tr>
<tr><td></td><td>Option</td><td>Values</td></tr>
<tr><td></td><td><tt>commitment</tt></td><td>Specifies whether a Snapshot copy is linked to data protection utilities (SnapMirror, Dump, or NDMPcopy) or data backing mechanisms (volume or LUN clones).
◆ <tt>try</tt>—delete only Snapshot copies that are not linked to data protection utilities and data backing mechanisms.
◆ <tt>disrupt</tt>—delete only Snapshot copies that are not linked to data backing mechanisms.</td></tr>
<tr><td></td><td><tt>trigger</tt></td><td>Defines when to begin automatically deleting Snapshot copies.
◆ <tt>volume</tt>—begin deleting when the volume is nearly full.
◆ <tt>snap_reserve</tt>—begin deleting when the snapshot reserve is nearly full.
◆ <tt>space_reserve</tt>—begin deleting when the space reserved in the volume is nearly full.</td></tr>
<tr><td></td><td><tt>target_free_space</tt></td><td>Determines when to stop deleting Snapshot copies. Specify a percentage. For example, if you specify 30, then Snapshot copies are deleted until 30 percent of the volume is free space.</td></tr>
<tr><td></td><td><tt>delete_order</tt></td><td>◆ <tt>newest_first</tt>—delete the most recent Snapshot copies first.
◆ <tt>oldest_first</tt>—delete the oldest Snapshot copies first.</td></tr>
<tr><td></td><td><tt>defer_delete</tt></td><td>Delete one of the following types of Snapshot copies last:
◆ <tt>scheduled</tt>—scheduled Snapshot copies, which are identified by their naming convention.
◆ <tt>user_created</tt>—Snapshot copies that are not scheduled Snapshot copies
◆ <tt>prefix</tt>—Snapshot copies with a prefix matching the specified <tt>prefix_string</tt> value</td></tr>
<tr><td></td><td><tt>prefix</tt></td><td>Delete Snapshot copies with a specific prefix last. You can specify up to 15 characters (for example, sv_snap_week). Use this option only if you specify <tt>prefix</tt> for the <tt>defer_delete</tt> option, otherwise the value you set for prefix is not applied.</td></tr>
</table>

Enabling the Snapshot copy autodelete policy: To enable the Snapshot copy autodelete policy you define, complete the following step.

Step	Action
1	Enter the following command: **`snap autodelete vol-name on`**

Viewing current Snapshot copy autodelete settings: To view current autodelete settings, complete the following step:

Step	Action
1	Enter the following command: **`snap autodelete vol-name show`**

Restoring default Snapshot copy autodelete settings: To restore default Snapshot copy autodelete settings, complete the following step:

Step	Action
1	Enter the following command: **`snap autodelete vol-name reset`** *vol-name* is the name of the volume. **Result:** Snapshot autodelete settings revert to the following defaults: ◆ state—off ◆ commitment —try ◆ trigger—volume ◆ target_free_space— 20% ◆ delete_order—oldest_first ◆ defer_delete—user_created ◆ prefix— no prefix specified

Disabling a Snapshot copy autodelete policy: Complete the following step:

Step	Action
1	Enter the following command: **`snap autodelete vol-name off`** *vol-name* is the name of the volume. **Result:** Snapshots are not automatically deleted when the volume is nearly full.

Defining how you apply space management policies

You can configure Data ONTAP to apply space management policies in one of the following ways:

- Automatically grow the volume first, then automatically delete Snapshot copies.

  This approach is useful if you create smaller FlexVol volumes and leave enough space in the aggregate to increase the size of these volumes as needed. If you provision your data based on aggregates, you might want to automatically grow the volume when it is nearly full before you begin automatically deleting Snapshot copies.
- Automatically delete Snapshot copies first, then grow the volume.

  You might want to automatically delete Snapshot copies if you are maintaining a large number of Snapshot copies in your volume or you maintain older Snapshot copies that are no longer needed.

To determine how Data ONTAP applies space management policies, complete the following step:

Step	Action
1	Enter the following command. **`vol options vol-name try_first [volume_grow|snap_delete]`** *vol-name* is the name of the FlexVol volume `volume_grow`—Automatically increase the size of the FlexVol volume before automatically deleting Snapshot copies. `snap_delete`—Automatically delete Snapshot copies according to the policy you defined in "Automatically deleting Snapshot copies" on page 153 before automatically increasing the size of the volume.

Using Data Protection with iSCSI and FCP 10

About this chapter

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

Topics in this chapter

This chapter discusses the following topics:

- "Data ONTAP protection methods" on page 160
- "Using Snapshot copies" on page 163
- "Using LUN clones" on page 165
- "Deleting busy Snapshot copies" on page 168
- "Using SnapRestore" on page 171
- "Backing up data to tape" on page 176
- "Using NDMP" on page 180
- "Using volume copy" on page 181
- "Cloning FlexVol volumes" on page 182
- "Using NVFAIL" on page 187
- "Using SnapValidator" on page 189

Data ONTAP protection methods

Data protection methods

Data ONTAP provides a variety of methods for protecting data in an iSCSI or Fibre Channel SAN. These methods, described in the following table, are based on Snapshot™ technology in Data ONTAP, which enables you to maintain multiple read-only versions of LUNs online per volume.

Snapshots are a standard feature of Data ONTAP. A Snapshot copy is a frozen, read-only image of the entire Data ONTAP file system, or WAFL® (Write Anywhere File Layout) volume, that reflects the state of the LUN or the file system at the time the Snapshot copy is created. The other data protection methods listed in the table below rely on Snapshot copies or create, use, and destroy Snapshot copies, 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 storage system, regardless of the size of the LUN or volume being restored. ◆ Recover from a corrupted database or a damaged application, a file system, a LUN, or a volume by using an existing Snapshot copy.
SnapMirror®	◆ Replicate data or asynchronously mirror data from one storage system to another over local or wide area networks (LANs or WANs). ◆ Transfer Snapshot copies taken at specific points in time to other storage systems or 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.

Method	Used to...
SnapVault®	◆ Back up data by using Snapshot copies on the storage system and transferring them on a scheduled basis to a destination storage system. ◆ Store these Snapshot copies on the destination storage system for weeks or months, allowing recovery operations to occur nearly instantaneously from the destination storage system to the original storage system.
SnapDrive™ for Windows or UNIX	◆ Manage storage system Snapshot copies directly from a Windows or UNIX host. ◆ Manage storage (LUNs) directly from a host. ◆ Configure access to storage directly from a host. SnapDrive for Windows supports Windows 2000 Server and Windows Server 2003. SnapDrive for UNIX supports a number of UNIX environments. For a list of supported host environments, see the *NetApp iSCSI Support Matrix*, *NetApp FCP SAN Compatibility Matrix*, and *SnapDrive & SnapManager Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/. **Note** For more information about SnapDrive, see the *SnapDrive Installation and Administration Guide* or *SnapDrive for UNIX Installation and Administration Guide*.
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.

Method	Used to...
NDMP	◆ Control native backup and recovery facilities in storage systems and other file servers. Backup application vendors provide a common interface between backup applications and file servers. **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 storage systems to protect data, see the *Data ONTAP Data Protection Tape Backup and Recovery Guide*.

Using Snapshot copies

How Data ONTAP Snapshot copies work in an iSCSI or FCP network

Snapshots of applications running on a file system may result in the Snapshot copy containing inconsistent data unless measures are taken (such as quiescing the application prior to the Snapshot copy) to ensure the data on disk is logically consistent before you take the Snapshot copy. If you want to take a Snapshot copy 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 copy is taken. The Snapshot copy 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 copy or a backup to occur while the application operates in a degraded mode, with limited performance.

Data ONTAP cannot take Snapshot copies of applications that have the ability to work with raw device partitions. It is recommended 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 Snapshot copies 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 copy, 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 copy is consistent. For example, you can use batch files and scripts on a host that has administrative access to the storage system. 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 in the script to access the storage system and use the `snap` command create the Snapshot copy on the storage system (this takes only a few seconds, regardless of volume size or use).

- Return the application to normal operation.

Snapshot scripts can be scheduled to run at specified intervals. On Windows hosts, you can use the Windows Task Scheduler. On UNIX hosts, you can use `cron` or other utilities. In addition, you can use SnapDrive to save the contents of the host file system buffers to disk and to create Snapshot copies. See the *SnapDrive Installation and Administration Guide*.

The relationship between a LUN and a Snapshot copy

When you take a Snapshot copy of a LUN, it is initially backed by data in the Snapshot copy. After the Snapshot copy is taken, data written to the LUN is in the active file system.

After you have a Snapshot copy, 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 copy, you cannot delete the Snapshot copy until you split the clone from it.

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

What Snapshot copies require

In Data ONTAP 6.5 and later, space reservation is enabled by default 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, Snapshot copies cannot be taken. For information about how to manage available space, see Chapter 9, “Managing Disk Space.”

Using LUN clones

What a LUN clone is

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

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

Note

Cloning is not NVLOG protected, so if the storage system 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.
- You want to access a specific subset of a LUN's data (a specific logical volume or file system in a volume group, or a specific file or set of files in a file system) and copy it to the original LUN, without restoring the rest of the data in the original LUN. This works on operating systems that support mounting a LUN and a clone of the LUN at the same time. SnapDrive for UNIX allows this with the `snap connect` command.

Creating a Snapshot copy of a LUN

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

Step	Action
1	Create a LUN by entering the following command: **`lun create -s size lun_path`** **Example:** `lun create -s 100g /vol/vol1/lun0`
2	Create a Snapshot copy of the volume containing the LUN to be cloned by entering the following command: **`snap create volume_name snapshot_name`** **Example:** `snap create vol1 mysnap`

Creating a clone of the LUN

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

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

Splitting the clone from the backing Snapshot copy

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

Step	Action
1	Begin the clone operation by entering the following command: **`lun clone split start lun_path`** *lun_path* is the path to the parent LUN. **Result:** The clone does not share data blocks with the Snapshot copy of the original LUN. This means you can delete the Snapshot copy.

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: **`lun clone split status lun_path`** *lun_path* 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: **`lun clone split stop lun_path`** *lun_path* is the path to the parent LUN.

Deleting busy Snapshot copies

What a Snapshot copy in a busy state means

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

Command to use to find Snapshot copies in a busy state

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

- writable Snapshot copy LUNs (backing store LUNs) that are holding a lock on the Snapshot copy given as input to this command
- Snapshots in which these snapshot-backed LUNs exist

Deleting Snapshot copies in a busy state

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

Step	Action

1	Identify all Snapshot copies that are in a busy state, locked by LUNs, by entering the following command:
	`snap list` ***`vol-name`***
	Example:
	`snap list vol2`
	Result: The following message is displayed:
	`Volume vol2`
	`working...`
	
	`  %/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 Snapshot copies that contain them by entering the following command:
	lun snap usage *vol_name snap_name*
	Example:
	`lun snap usage vol2 snap0`
	Result: The following message is displayed:
	`active:`
	`  LUN:  /vol/vol2/lunC`
	`  Backed By: /vol/vol2/.snapshot/snap0/lunA`
	`snap2:`
	`  LUN:  /vol/vol2/.snapshot/snap2/lunB`
	`  Backed By: /vol/vol2/.snapshot/snap0/lunA`
	`snap1:`
	`  LUN:  /vol/vol1/.snapshot/snap1/lunB`
	`  Backed By: /vol/vol2/.snapshot/snap0/lunA`
	Note
	The LUNs are backed by lunA in the snap0 Snapshot copy.

3	Delete all the LUNs in the active file system that are displayed by the `lun snap usage` command by entering the following command: **lun destroy [-f] *lun_path* [*lun_path* ...]** **Example:** `lun destroy /vol/vol2/lunC`

4	Delete all the Snapshot copies that are displayed by the `lun snap usage` command in the order they appear, by entering the following command:
	snap delete *vol-name snapshot-name*
	Example:
	`snap delete vol2 snap2`
	`snap delete vol2 snap1`
	Result: All the Snapshot copies containing lunB are now deleted and snap0 is no longer busy.

Step	Action
5	Delete the Snapshot copy by entering the following command: **snap delete *vol-name snapshot-name*** **Example:** `snap delete vol2 snap0`

Using SnapRestore

What SnapRestore does

SnapRestore uses a Snapshot copy to revert an entire volume or a LUN to its state when the Snapshot copy 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 copy 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 copy of a LUN

To use SnapRestore to restore a Snapshot copy 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 storage system, unmap the LUN by entering the following command: **`lun unmap`** ***`lun_path initiator-group`***

Step	Action

4	Enter the following command:
	snap restore [-f] [-t vol] *volume_name* [-s *snapshot_name*]
	`-f` suppresses the warning message and the prompt for confirmation. This option is useful for scripts.
	`-t vol` *volume_name* specifies the volume name to restore.
	volume_name 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` *snapshot_name* specifies the name of the Snapshot copy from which to restore the data. You can enter only one Snapshot copy name.
	Example:
	`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`
	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 copy 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 storage system reverts the volume from the selected Snapshot copy.
	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 storage system prompt.
7	Enter the following command to unmap the existing old maps that you do not want to keep. **lun unmap *lun_path initiator-group***	
8	Remap the LUN by entering the following command: **lun map *lun_path initiator-group***	
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 storage system, bring the restored LUN online by entering the following command: **lun online *lun_path***	

Note

After you use SnapRestore to update a LUN from a Snapshot copy, 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 an 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 copy.
2	Enter the following command: **`snap restore [-f] [-t file] [-s snapshot_name] [-r restore_as_path] path_and_LUN_name`** `-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 copy 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 copy. 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 copy for the restore operation. If you did not use the `-f` 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 copy.

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 copy.

Backing up data to tape

Structure of SAN backups

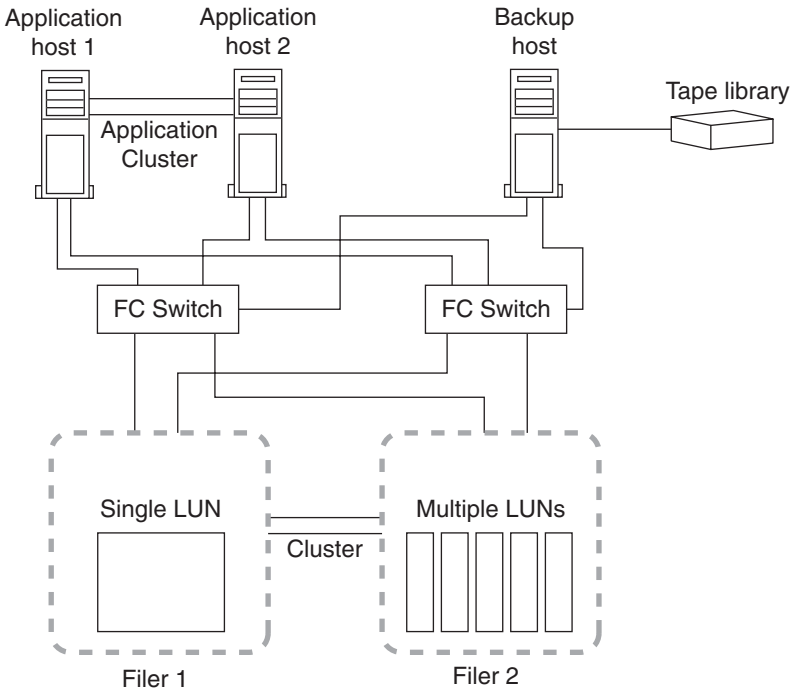
In most cases, backup of SAN systems to tape takes place through a separate backup host to avoid performance degradation on the application host.

Note

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

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 storage systems.

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



Volumes on a host can consist of a single LUN mapped from the storage system 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 copy 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: **`igroup create -f [-t ostype] group [node ...]`** **Example:** `igroup create -f -t windows payroll_server 10:00:00:00:c3:4a:0e:e1` **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: **`lun create -s size [-t type] lun_path`** **Example:** `lun create -s 48g -t windows /vol/vol1/qtree_1/payroll_lun` **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. **`lun map lun_path initiator-group LUN_ID`** **Example:** `lun map /vol/vol1/qtree_1/payroll_lun payroll_server 1` **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 documentation that came with your Host Support Kit or 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 copy by entering the following command: **`snap create volume_name snapshot_name`** **Example:** `snap create vol1 payroll_backup`
7	Enter the following command to create a clone of the production LUN: **`lun clone create clone_lunpath -b parent_lunpath parent_snap`** **Example:** `lun clone create /vol/vol1/qtree_1/payroll_lun_clone -b /vol/vol1/qtree_1/payroll_lun payroll_backup`
8	Create an igroup that includes the WWPN of the backup server. **`igroup create -f [-t ostype] group [node ...]`** **Example:** `igroup create -f -t windows backup_server 10:00:00:00:d3:6d:0f:e1` **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: **lun map *lun_path initiator-group LUN_ID*** **Example:** `lun map /vol/vol1/qtree_1/payroll_lun_clone backup_server 1` **Result:** Data ONTAP maps the LUN clone (/vol/vol1/qtree_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 documentation that came with your Host Support Kit or 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: **lun offline /vol/*vol_name/qtree_name/lun_name*** **Example:** `lun offline /vol/vol1/qtree_1/payroll_lun_clone`
13	Remove the LUN clone by entering the following command: **lun destroy *lun_path*** **Example:** `lun destroy /vol/vol1/qtree_1/payroll_lun_clone`
14	Remove the Snapshot copy by entering the following command: **snap delete *volume_name lun_name*** **Example:** `snap delete vol1 payroll_backup`

Using NDMP

When to use native or NDMP backup

Tape backup and recovery operations of LUNs should generally only be performed on the storage system for disaster recovery scenarios, applications with transaction logging, or when combined with other data protection methods, 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 storage system 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 copy when performing a native or NDMP backup operation, the storage system creates one before proceeding. This Snapshot copy is deleted when the backup is completed. When a file system contains FCP data, it is recommended that you specify a Snapshot copy 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 copy. After the Snapshot copy is created, normal application operation can resume and tape backup of the Snapshot copy 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 storage system or to a different storage system. The result of the `vol copy` command is a restricted volume containing the same data that was on the source storage system at the time you initiate the copy operation.

Copying a volume

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

CAUTION

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

Step	Action
1	Enter the following command: **`vol copy start -S source:`*`source_volume`* `dest:`*`dest_volume`* `-S` copies all Snapshot copies in the source volume to the destination volume. If the source volume has Snapshot copy-backed LUNs, you must use the `-S` option to ensure that the Snapshot copies are copied to the destination volume. **Note** If the copying takes place between two storage systems, you can enter the `vol copy start` command on either the source or destination storage system. You cannot, however, enter the command on a third storage system that does not contain the source or destination volume. **Example:** `vol copy start -S /vol/vol0 filerB:/vol/vol1`

Cloning FlexVol volumes

What clones volumes are

Clone volumes are writable, point-in-time copy of a parent FlexVol volume. They 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 FlexVol volumes

You can clone FlexVol volumes when you want a writable, point-in-time copy of a FlexVol volume. For example, you might want to clone FlexVol 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 clone volumes, 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 Snapshot copies, the clone also inherits the Snapshot copies.

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 copy, which is also inherited by the volume clone. The LUN clone's base Snapshot copy in the parent volume shares blocks with the LUN clone's base Snapshot copy in the volume clone. You cannot delete the LUN clone's base Snapshot copy in the parent volume until you delete the base Snapshot copy 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_ucode=off, convert_ucode=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
c1     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 FlexVol volumes

You use the following commands to clone FlexVol volumes:

- `vol clone create`—creates a volume clone and a base Snapshot copy of the parent volume.

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

Cloning a FlexVol volume

To clone a FlexVol volume, complete the following steps.

Step	Action
1	Enter the following command to clone the volume: **`vol clone create cl_vol_name [-s {volume|file|none}] -b f_p_vol_name [parent_snap]`** *cl_vol_name* is the name of the clone volume that you want to create. `-s {volume | file | none}` 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 183. **Note** For Data ONTAP 7.0, space guarantees are disabled for clone volumes until they are split from the parent volume. *f_p_vol_name* is the name of the flexible parent volume that you intend to clone. [*parent_snap*] is the name of the base Snapshot copy of the parent volume. If no name is specified, Data ONTAP creates a base Snapshot copy with the name clone_*cl_name_prefix.id*, where *cl_name_prefix* is the name of the new clone volume (up to 16 characters) and *id* is a unique digit identifier (for example 1,2, and so on). The base Snapshot copy cannot be deleted as long as the parent volume or any of its clones exists. **Example Snapshot copy name:** To create a clone *newclone* of the volume named flexvol1, the following command is entered: `vol clone create newclone -b flexvol1` The Snapshot copy created by Data ONTAP is named clone_newclone.1.

Step	Action
2	Verify the success of the clone creation by entering the following command: **vol status -v *cl_vol_name***

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: **df -A *aggr_name*** *aggr_name* is the name of the containing aggregate of the FlexVol volume clone that you want to split. The `avail` column tells you how much available space you have in your aggregate. When a volume clone is split from its parent, the resulting two FlexVol volumes occupy completely different blocks within the same aggregate.

Step	Action
2	Enter the following command to split the volume: **`vol clone split start cl_vol_name`** *cl_vol_name* 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: **`vol clone status cl_vol_name`**
4	If you want to stop the progress of an ongoing clone-splitting operation, enter the following command: **`vol clone stop cl_vol_name`** 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: **`vol status -v cl_vol_name`**

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 those LUNs over iSCSI or 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 84.

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: **`vol options volume-name nvfail on`**

Creating the nvfail_rename file

To create the nvfail_rename file, complete the following steps.

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

Using SnapValidator

What SnapValidator 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™ software implements Oracle H.A.R.D. checks to detect and reject invalid Oracle data before it is written to the storage system.

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 storage system or if you want to store a new Oracle database on the storage system.

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 a 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 57.
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 191.
7. Enable SnapValidator checks on your volumes. For more information, see "Enabling SnapValidator checks on volumes" on page 192.

   Make sure you set SnapValidator to log all errors to the storage system 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 storage system.
9. Set SnapValidator to reject invalid operations and return an error to the host and log that error to the storage system 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 storage system:

- License SnapValidator.

  For detailed information, see "Licensing SnapValidator" on page 191.
- Enable SnapValidator checks on the volume that contains the Oracle data.

  For detailed information, see "Enabling SnapValidator checks on volumes" on page 192.
- 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 193.

Licensing SnapValidator

To license SnapValidator complete the following steps.

Step	Action
1	Verify whether SnapValidator is licensed by entering the following command: **`license`** **Result:** 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. `snapvalidator not licensed`

Step	Action	
2	**If SnapValidator is...**	**Then...**
	Licensed	Proceed to "Enabling SnapValidator checks on volumes" on page 192.
	Not licensed	Enter the following command: **`license add license_code`** *license_code* is the license code you received 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	On the storage system command line, enable SnapValidator by entering the following command: **`vol options volume-name svo_enable on`** **Result:** All SnapValidator checks are enabled on the volume, with the exception of checksums.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td rowspan="5">2</td><td>If you want to...</td><td>Then enter the following command:</td></tr>
<tr><td>Enable data checksumming on the volume</td><td><tt>vol options volume-name svo_checksum on</tt></td></tr>
<tr><td>Disable block number checks because the volume contains Oracle Recovery Manager (RMAN) backup data.</td><td><tt>vol options volume-name svo_allow_rman on</tt></td></tr>
<tr><td>Set SnapValidator to return an error log to the host and storage system consoles for all invalid operations. You might want to do this when you are testing SnapValidator before you put your database into production.</td><td><tt>vol options volume-name svo_reject_errors off</tt>
When you set this option to Off, SnapValidator only logs errors but does not reject invalid operations.</td></tr>
<tr><td>Set SnapValidator to reject all invalid operations and return an error log to the host and storage system consoles.</td><td><tt>vol options volume-name svo_reject_errors on</tt>
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 storage system console:
<tt>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</tt></td></tr>
<tr><td>3</td><td colspan="2">If the volume contains LUNs, proceed to “Enabling SnapValidator checks on LUNs” in the next section.</td></tr>
</table>

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 storage system. The method you use to identify the offset depends on your host. For details see:

- "Identifying the disk offset for Solaris hosts" on page 194
- "Identifying the disk offset for other hosts"
- "Defining the disk offset on the storage system"

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

Step	Action

1	On the host, enter the following command:
	prtvtoc /dev/rdsk/*device_name*
	Result: The host console displays a partition map for the disk.
	Example: The following output example shows the partition map for disk c3t9d1s2:
	`prtvtoc /dev/rdsk/c3t9d1s2`
	`* /dev/rdsk/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`
	`*`
	`*  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`

2	Obtain the offset value by multiplying the value of the first sector of the partition you are using by the bytes/sector value listed under Dimensions. In the example shown in Step 1, which is using slice 6, the disk offset is 6144 * 512 = 3145728.

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

Step	Action
1	Prepare the LUN for storing Oracle data, for example, by setting up raw volumes.
2	On the host console, enter the following command: **`dd if=/dev/zero of=/dev/`*`path_to_storage`*` bs=4096 count=1 conv=notrunc`** *path_to_storage* is the path to the LUN on the host. **Result:** The host writes an Oracle 4K block of zeros to the storage system.
3	Check the SnapValidator error message displayed on the storage system console. The error message displays the offset. **Example:** The following error message example shows that the disk offset is 1,048,576 bytes. `filerA> Thu Mar 10 16:26:01 EST [filerA:wafl.svo.checkFailed:error]: SnapValidator: Validation error Zero Data:: v:9r2 vol:test inode:3184174 length:4096` **`Offset: 1048576`**

Defining the disk offset on the storage system: To define the disk offset on the storage system, complete the following step.

Step	Action
1	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.
2	On the storage system command line, enter the following command: **`lun set svo_offset `*`lun_path offset`* *offset* is specified in bytes, with an optional multiplier suffix: `c` (1), `w` (2), `b` (512), `k` (1,024), `m` (k*k), `g` (k*m), `t` (m*m).

Disabling SnapValidator on a volume

To disable SnapValidator, complete the following steps.

Step	Action
1	On the storage system command line, enter the following command: **`vol options volume-name svo_enable off`** **Result:** 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.
2	To disable a specific SnapValidator option, enter the following command: **`vol options volume-name option off`** *option* is one of the following: ◆ *svo_checksum*—disables data checksumming on the volume. ◆ *svo_allow_rman*—allows block number checks on the volume. You disable this option (set it to Off) if the volume does not contain RMAN data. ◆ *svo_reject_errors*—detects invalid operations but does not reject them. Invalid operations are only logged as errors.

Disabling SnapValidator checks on a LUN

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

Step	Action
1	On the storage system command line, enter the following command: **`lun set lun_path svo_offset disable`**

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 storage system 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: wafl.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 193.

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:

You enabled the SnapValidator checks on the volumes that contain your data files. For more information, see "Enabling SnapValidator checks on volumes" on page 192.

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 options on the storage system are correctly set but you still receive the above errors, you might have the following problems:

- Your host is writing invalid data to the storage system. 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 storage system.

Improving Read/Write Performance 11

About this chapter

This chapter describes commands and options that enable you to improve LUN and volume layout and improve the read/write performance of host applications that access data on the storage system.

Topics in this chapter

This chapter discusses the following topics

- "Reallocating LUN and volume layout" on page 200
- "Improving Microsoft Exchange read performance" on page 210

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 storage system.

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.

The rate at which the reallocation scan runs (the blocks reallocated per second) depends on CPU and disk loads. For example, if you have a high CPU load, the reallocation scan will run at a slower rate, so as not to impact system performance.

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 storage system 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 storage system before you run a scan or schedule regular scans.

To enable reallocation scans, complete the following step.

Step	Action
1	On the storage system's command line, enter the following command: **`reallocate on`**

Defining a reallocation scan

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

Step	Action

1	On the storage system's command line, enter the following command:
	`reallocate start [-t threshold] [-n] [-i interval] lun_path`
	`-t` *threshold* 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` *interval* 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.
	`lun create -s 100g /vol/vol2/lun0`
	`reallocate start /vol/vol2/lun0`

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 204.
	You do not want to define a schedule	Proceed to "Tasks for managing reallocation scans" on page 205.

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

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: **`reallocate schedule -d lun_path`** **Example:** `reallocate schedule -d /vol/myvol/lun1`

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: **`reallocate start -o -n lun_path`** `-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.

Caution

You should *not* perform a full reallocation on an entire volume that has Snapshot copies. 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 copy after the scan. For individual LUNs or files, try to avoid transferring large amounts of data from the Snapshot copy to the active filesystem unless absolutely necessary. The greater the differences between the LUN or file and the Snapshot copy, the more likely the full reallocation will be successful.

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

Step	Action
1	Enter the following command: **`reallocate start -f`** ***`lun_path`*** **`|`** ***`volume-path`***

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: **`reallocate quiesce lun_path`**

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: **`reallocate restart lun_path`** **Result:** The command restarts a quiesced scan. If there is a scheduled scan that is idle, the `reallocate restart` 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: **`reallocate status [-v] lun_path`** `-v` provides verbose output *lun_path* is the path to the LUN for which you want to see reallocation scan status. If you do not specify a value for *lun_path*, then the status for all scans is displayed. **Result:** The `reallocate status` 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 `reallocate status` 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: **`reallocate stop lun_path`** **Result:** The `reallocate stop` 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 storage system. 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 storage system's command line, enter the following command: **`reallocate off`**

Best practice recommendations

The following best practices are recommended 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.

Improving Microsoft Exchange read performance

How logical extents improve sequential read performance

A logical extent is a group of data blocks that are logically aligned and logically contiguous. When you enable logical extents, Data ONTAP processes write operations by creating groups of logically contiguous data blocks that are physically close to each other on the disk. Extents optimize sequential data block layout and improve the amount of time required for applications to perform sequential read operations, such as database scans.

In Microsoft Exchange environments, you use the Exchange `eseutil` tool to perform database scans for validation purposes. Exchange database scans access data by mostly using a sequential read pattern. By enabling logical extents, you improve Exchange sequential read performance and database validation time.

When to enable logical extents

You enable logical extents for volumes that contain Microsoft Exchange data only. The decision to use logical extents involves a trade-off between improved database validation performance and runtime performance. You use logical extents when you want to improve validation performance. If runtime performance is higher priority, you might not want to use extents.

Enabling logical extents

You can enable logical extents on a traditional or FlexVol volume. To enable logical extents, enter the following command:

`vol options` *`vol-name`* `extent [on | off]`

`on` enables logical extents for the volume.

`off` disables logical extents for the volume. By default, logical extents are disabled.

Managing the iSCSI Network 12

About this chapter

This chapter describes how to manage the iSCSI service and the storage system as a target in the iSCSI network.

Note

The commands and FilerView pages used to manage iSCSI on a NetApp storage system changed in Data ONTAP 7.1. This chapter includes an overview of the changes. The new commands are used in the relevant procedures.

Topics in this section

This section discusses the following topics

- "Management changes for iSCSI in Data ONTAP 7.1" on page 212
- "Managing the iSCSI service" on page 216
- "Registering the storage system with an iSNS server" on page 222
- "Displaying initiators connected to the storage system" on page 228
- "Managing security for iSCSI initiators" on page 229
- "Managing target portal groups" on page 236
- "Displaying statistics for iSCSI sessions" on page 243
- "Displaying information for iSCSI sessions and connections" on page 247
- "Managing the iSCSI service on storage system interfaces" on page 252
- "Using iSCSI on clustered storage systems" on page 256
- "Troubleshooting common iSCSI problems" on page 259

Management changes for iSCSI in Data ONTAP 7.1

Administrative model changes

The overall administrative model for iSCSI is changed in Data ONTAP 7.1. These changes are necessary to support new iSCSI target functionality. The new iSCSI functions are:

- Multi-connection sessions. Requires the ability to assign network interfaces to specific target portal groups. Sessions are no longer tied to adapters. Interfaces are managed by the standard networking commands (`ifconfig`, `ifstat`, `vlan`, and `vif`).
- Target alias. Adds the ability to assign an alternate name to identify the storage system.
- Error recovery level greater than zero. No administration needed. Requires an initiator that supports this function and has been qualified by NetApp.
- Virtual interfaces (vifs) and VLAN-tagged interfaces are managed by the iSCSI command directly. You no longer enable and disable iSCSI for a virtual interface using the `iswt` command on the underlying physical interfaces.

Command line changes

The `iscsi` and `iswt` commands are changed as follows. See the man pages or the *Data ONTAP 7.1 Commands: Manual Page Reference* for the complete command syntax and options.

- `iswt interface` replaced by `iscsi interface`
- `iswt session show` replaced by `iscsi session show`
- `iswt connection show` replaced by `iscsi connection show`
- `iscsi show initiator` replaced by `iscsi initiator show`
- `iscsi config` removed
- `iscsi show adapter` removed
- `iscsi alias` added
- `iscsi tpgroup` added
- `iscsi portal show` added
- `iscsi stats` modified to report statistics for the entire storage system instead of for an individual adapter

FilerView changes

The FilerView LUN pages are changed as follows. See the FilerView help for information on specific fields and pages.

- LUNs > iSCSI > Report page displays additional information
- LUNs > iSCSI > Adapters page removed
- LUNs > iSCSI > Initiator Security page renamed LUNs > iSCSI > Manage Initiator Security
- LUNs > iSCSI > Manage Names page added
- LUNs > iSCSI > Portal Addresses page added
- LUNs > iSCSI > Manage Interfaces page added
- LUNs > iSCSI > Initiators page added
- LUNs > iSCSI > iSNS page renamed Manage iSNS

Enabling multi-connection sessions

By default, Data ONTAP 7.1 is configured to use a single TCP/IP connection for each iSCSI session. If you are using an initiator that has been qualified for multi-connection sessions, you can specify the maximum number of connections allowed for each session on the storage system.

Check the NetApp iSCSI support matrix to verify if your initiator has been qualified for multi-connection sessions. The matrix is available at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml

The `iscsi.max_connections_per_session` option specifies the number of connections per session allowed by the storage system. You can specify between 1 and 16 connections, or you can accept the default value.

Note that this option specifies the maximum number of connections per session supported by the storage system. The initiator and storage system negotiate the actual number allowed for a session when the session is created; this is the smaller of the initiator's maximum and the storage system's maximum. The number of connection actually used also depends on how many connections the initiator establishes.

To view or change the setting of the `iscsi.max_connections_per_session` option, complete the following steps:

Step	Action
1	Verify the current option setting by entering the following command on the system console: **`options iscsi.max_connections_per_session`** **Result:** The current setting is displayed.
2	If needed, change the number of connections allowed by entering the following command: **`options iscsi.max_connections_per_session [`*`connections`*` | use_system_default]`** *connections* is the maximum number of connections allowed for each session, from 1 to 16. `use_system_default` equals 1 for Data ONTAP 7.1. The meaning of this default may change in later releases.

Enabling error recovery levels 1 and 2

By default, Data ONTAP 7.1 is configured to use only error recovery level 0 for iSCSI sessions. If you are using an initiator that has been qualified for error recovery levels 1 or 2, you can specify the maximum error recovery level allowed by the storage system.

Check the NetApp iSCSI support matrix to verify if your initiator has been qualified for error recovery levels greater than 0. The matrix is available at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml

There may be a minor performance reduction for sessions running error recovery level 1 or 2.

The `iscsi.max_error_recovery_level` option specifies the maximum error recovery level allowed by the storage system. You can specify 0, 1, or 2, or you can accept the default value.

Note that this option specifies the maximum error recovery level supported by the storage system. The initiator and storage system negotiate the actual error recovery level used for a session when the session is created; this is the smaller of the initiator's maximum and the storage system's maximum.

To view or change the setting of the `iscsi.max_error_recovery_level` option, complete the following steps:

Step	Action
1	Verify the current option setting by entering the following command on the system console: **`options iscsi.max_error_recovery_level`** **Result:** The current setting is displayed.
2	If needed, change the error recovery levels allowed by entering the following command: **`options iscsi.max_error_recovery_level [`*`level`*` | use_system_default]`** *level* is the maximum error recovery level allowed, 0, 1, or 2. `use_system_default` equals 0 for Data ONTAP 7.1. The meaning of this default may change in later releases.

Managing the iSCSI service

Verifying that the iSCSI service is running

Verifying that the iSCSI service is running using the command line:

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

Step	Action
1	On the storage system console, enter the following command: `iscsi status` **Result:** A message is displayed indicating whether iSCSI service is running.

Verifying that this iSCSI service is running using FilerView: To verify that the iSCSI service is running, complete the following step.

Step	Action
1	Click LUNs > Enable/Disable. **Result:** The status of the iSCSI service is displayed.

Note

If the iSCSI service is not running, verify that the iSCSI license is enabled and start the service.

Verifying that iSCSI is licensed

Verifying that iSCSI is licensed using the command line: To verify that the iSCSI service is licensed, complete the following step.

Step	Action
1	On the storage system console, enter the following command: `license` **Result:** A list of all available licenses is displayed. An enabled license shows the license code.

Verifying that iSCSI is licensed using FilerView: To verify that the iSCSI service is licensed, complete the following step.

Step	Action
1	Click Filer > Manage Licenses. **Result:** A list of all available licenses is displayed. An enabled license shows the license code.

Enabling the iSCSI license

Enabling the iSCSI license using the command line: To enable the iSCSI license, complete the following step.

Step	Action
1	On the storage system console, enter the following command: `license add` ***`license_code`*** *license_code* is the license code provided to you.

Enabling the iSCSI license using FilerView: To enable the iSCSI license, complete the following steps.

Step	Action
1	Click Filer > Manage Licenses.
2	In the iSCSI field, enter the license code provided to you.
3	Click Apply.

Starting and stopping the iSCSI service

Starting and stopping the iSCSI service using the command line: To start or stop the iSCSI service, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi {start|stop}`**

Starting and stopping the iSCSI service using FilerView: To start or stop the iSCSI service, complete the following steps.

Step	Action
1	Click LUNs > Enable/Disable.
2	To start the iSCSI service, check the Enable box. To stop the iSCSI service, clear the Enable box.
3	Click Apply.

Displaying the target node name

Displaying the target node name using the command line: To display the storage system's target node name, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi nodename`
	Example:
	`iscsi nodename`
	`iSCSI target nodename: iqn.1992-08.com.netapp:sn.12345678`

Displaying the target node name using FilerView: To display the storage system's target node name, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Manage Names. **Result:** The target node name is displayed in the "Change node name" field.

Changing the target node name

Changing the storage system's node name while iSCSI sessions are in progress does not disrupt the existing sessions. However, when you change the storage system's node name, you must reconfigure the initiator so that it recognizes the new target node name. If you do not reconfigure the initiator, subsequent initiator attempts to log in to the target will fail.

Node name rules: If you change the storage system's target node name, be sure the new name follows all of these rules:

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

Changing the target node name using the command line: To change the storage system's target node name, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi nodename iqn.1992-08.com.netapp:`*`unique_device_name`*** *unique_device_name* is the unique name for the storage system. **Example:** `iscsi nodename iqn.1992-08.com.netapp:filerhq`

Changing the target node name using FilerView: To change the storage system's target node name, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Names.
2	Enter the new target node name in the "Change node name" field.
3	Click Apply.

Displaying the target alias

The target alias is an optional human-readable name for the iSCSI target. For example, if your target node name was iqn.1992-08.com.netapp:sn.33604646, you might want an alias such as Filer_1. The alias is a text string with a maximum of 128 characters.

The alias is intended to be displayed by an initiator's user interface to make it easier for someone to identify the desired target in a list of targets. Depending on your initiator, the alias may or may not be displayed in the initiator's user interface.

Displaying the target alias using the command line: To display the storage system's target alias, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi alias`
	Example:
	`iscsi alias`
	`iSCSI target alias: Filer_1`

Displaying the target node name using FilerView: To display the storage system's target alias, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Manage Names. **Result:** The target alias is displayed in the Change Alias field.

Adding or changing the target alias

You can change the target alias or clear the alias at any time without disrupting existing sessions. The new alias will be sent to the initiators the next time they log in to the target.

Changing the target alias using the command line: To create, change, or clear the storage system's target alias, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi alias [-c	string]`** `-c` clears the existing alias value *string* is the new alias value, maximum 128 characters **Example 1:** `iscsi alias Filer_2` `New iSCSI target alias: Filer_2` **Example 2:** `iscsi alias -c` `Clearing iSCSI target alias`

Changing the target node name using FilerView: To create, change, or clear the storage system's target alias, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Names. **Result:** The current target node alias is displayed in the Change Alias field.
2	Enter the new target alias in the Change Alias field. To clear the target alias, delete all of the existing text in the field.
3	Click Apply.

Registering the storage system with an iSNS server

What an iSNS server does

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

NetApp does not supply or resell iSNS servers. You obtain these servers from a vendor supported by NetApp. Be sure to check the *NetApp iSCSI Support Matrix* to see which iSNS servers are currently supported.

Resolving iSNS service version incompatibility

The specification for the iSNS service is still in draft form. Some draft versions are different enough to prevent the storage system from registering with the iSNS server. Because the protocol does not provide version information to the draft level, iSNS servers and storage systems cannot negotiate the draft level being used.

By default, Data ONTAP versions prior to 7.1 used iSNS draft 18. This draft was also used by Microsoft iSNS server versions prior to 3.0.

In Data ONTAP 7.1, the default iSNS version is draft 22. This draft is also used by Microsoft iSNS server 3.0.

Choices for iSNS service: You can either use the iSNS server that matches your Data ONTAP version, or you can configure Data ONTAP to use a different iSNS draft version by changing the `iscsi.isns.rev` option on the storage system. Refer to the following table.

If your Data ONTAP version is...	And your iSNS server version is...	Then you should...
Prior to 7.1	Prior to 3.0	Verify that the `iscsi.isns.rev` option is set to 18.
7.1	Prior to 3.0	Set `iscsi.isns.rev` option to 18 or upgrade to iSNS server 3.0.

If your Data ONTAP version is...	And your iSNS server version is...	Then you should...
Prior to 6.5.4	3.0	Upgrade Data ONTAP or use a version of iSNS server prior to 3.0.
6.5.4 to 7.0.x	3.0	Set `iscsi.isns.rev` option to 22 or use a version of iSNS server prior to 3.0.
7.1	3.0	Verify that the `iscsi.isns.rev` option is set to 22.

Note

When you upgrade to a new version of Data ONTAP, the existing value for the `iscsi.isns.rev` option is maintained. This reduces the risk of a draft version problem when upgrading. For example, if you upgrade from the 7.0 family to 7.1, the default value 18 from the 7.0 family is also used for 7.1. You must change the `iscsi.isns.rev` value manually when upgrading Data ONTAP.

Setting the iSNS service revision: To set the iSNS service revision, complete the following steps.

Step	Action
1	Verify the current iSNS revision value by entering the following command on the system console: **`options iscsi.isns.rev`** **Result:** The current draft revision used by the storage system is displayed.
2	If needed, change the iSNS revision value by entering the following command: **`options iscsi.isns.rev`** ***`draft`*** *draft* is the iSNS standard draft revision, either 18 or 22.

How the storage system interacts with an iSNS server

The storage system automatically registers its IP address, node name, and portal groups with the iSNS server when the iSCSI service is started and iSNS is enabled. After iSNS is initially configured, Data ONTAP automatically updates the iSNS server any time the storage system's configuration settings change. There can be a delay of a few minutes between the time of the configuration change and the update being sent; you can use the `iscsi isns update` command to send an update immediately.

Command to register the storage system

You can use the `iscsi isns` command or FilerView to configure the storage system to register with an iSNS server. This command specifies the information the storage system sends to the iSNS server.

How you manage the iSNS server

The `iscsi isns` command only configures the storage system to register with the iSNS server. The storage system does not provide commands that enable you to configure or manage the iSNS server.

To manage the iSNS server, use the server administration tools or interface provided by the vendor of the iSNS server.

Configuring the storage system to register with an iSNS server

Registering with iSNS using the command line: To configure the storage system to register with the iSNS server, complete the following steps.

Step	Action
1	Make sure the iSCSI service is running by entering the following command on the storage system console: **`iscsi status`**
2	If the iSCSI service is not running, enter the following command: **`iscsi start`**
3	On the storage system console, enter the following command to identify the iSNS server that the storage system registers with: **`iscsi isns config -i ip_addr`** *ip_addr* is the IP address of the iSNS server.

Step	Action
4	Enter the following command: `iscsi isns start` **Result:** The iSNS service is started and the storage system registers with the iSNS server. **Note** iSNS registration is persistent across reboots if the iSCSI service is running and iSNS is started.

Registering with iSNS using FilerView: To configure the storage system to register with the iSNS server, complete the following steps.

Step	Action
1	Click LUNs > Enable/Disable. Verify that the iSCSI service is enabled.
2	If the iSCSI service is not enabled, check the Enable box and click Apply.
3	Click LUNs > iSCSI > Manage iSNS.
4	Check the Enable box.
5	Enter the IP address of the iSNS server.
6	Click Apply. **Result:** The iSNS service is started and the storage system registers with the iSNS server. **Note** iSNS registration is persistent across reboots if the iSCSI service is running and iSNS is started.

Updating the iSNS server immediately

Data ONTAP checks for iSCSI configuration changes on the storage system every few minutes and automatically sends any changes to the iSNS server. If you do not want to wait for an automatic update, you can update the iSNS server immediately.

To immediately update the iSNS server with iSCSI configuration changes, complete the following step:

Step	Action
1	On the storage system console, enter the following command: `iscsi isns update`

Disabling iSNS

Disabling iSNS using the command line: When you stop the iSNS service, the storage system stops registering its iSCSI information with the iSNS server.

To stop the iSNS service, complete the following step.

Step	Action
1	On the storage system console, enter the following command: `iscsi isns stop`

Disabling iSNS using FilerView: When you stop the iSNS service, the storage system stops registering its iSCSI information with the iSNS server.

To stop the iSNS service, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage iSNS.
2	Clear the Enable box.
3	Click Apply.

Setting up vFiler units with the iSNS service

You can set up iSNS separately for each vFiler by using the `iscsi isns` command on each vFiler to:

- Configure which iSNS server to use

- Turn iSNS registration on or off

To set up vFiler units with the iSNS service, complete the following steps for each vFiler. Configuring iSNS for each vFiler must be done using the command line.

Step	Action

1	Register the vFiler with the iSNS service by entering the following command:
	`iscsi isns config -i ip_addr`
	ip_addr is the IP address of the iSNS server.
	Examples:
	The following example defines the iSNS server for the default vFiler (vfiler0) on the hosting storage system:
	`iscsi isns config -i 10.10.122.101`
	The following example defines the iSNS server for a specific vFiler (vf1). The `vfiler context` command switches to the command line for a specific vFiler.
	`vfiler context vf1`
	`vf1> iscsi isns config -i 10.10.122.101`

2	Enter the following command to enable the iSNS service: **`iscsi isns start`**

Note

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *MultiStore Management Guide*.

Displaying initiators connected to the storage system

Initiator information displayed

You can display a list of initiators currently connected to the storage system. The information displayed for each initiator includes the target session identifier handle (TSIH) assigned to the session, the target portal group tag of the group the initiator is connected to, the iSCSI initiator alias (if provided by the initiator), and the initiator's iSCSI node name and initiator session identifier (ISID).

Displaying initiators using the command line

To display a list of iSCSI initiators connected to the storage system, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi initiator show`
	Result: The initiators currently connected to the storage system are displayed.
	Example:
	`toaster> iscsi initiator show`
	`Initiators connected:`
	`  TSIH  TPGroup  Initiator`
	`    19    1000   iqn.1991-05.com.microsoft:host1.fake.com / 40:01:37:00:06:00`
	`    21    1002   iqn.1991-05.com.microsoft:host2.fake.com / 40:01:37:00:00:00`

Displaying initiators using FilerView

To display a list of iSCSI initiators connected to the storage system, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Initiators. **Result:** The initiators currently connected to the storage system are displayed.

Managing security for iSCSI initiators

Ways to manage initiator security with authentication methods

You can manage the security for iSCSI initiators by performing the following tasks:

- Define iSCSI initiator authentication methods that are kept in an authentication list
- Display the authentication methods in the list
- Define iSCSI initiator authentication methods for initiators not in the list
- Add initiators to the authentication list
- Remove initiators from the authentication list

How iSCSI authentication works

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

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

  There are two types of CHAP user names and passwords:

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

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

Authentication with vFiler units

If you use iSCSI with vFiler units, the CHAP authentication settings are configured separately for each vFiler. Each vFiler has its own default authentication mode and list of initiators and passwords.

To configure CHAP settings for vFiler units, you must use the command line.

Note

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *MultiStore Management Guide*.

Guidelines for using CHAP authentication

The following guidelines apply to CHAP authentication:

- If you define an inbound user name and password on the storage system, you must use the same user name and password for outbound CHAP settings on the initiator.
- If you also define an outbound user name and password on the storage system to enable bidirectional authentication, you must use the same user name and password for inbound CHAP settings on the initiator.
- You cannot use the same user name and password for inbound and outbound settings on the storage system.
- CHAP user names can be 1 to 128 bytes. A null user name is not allowed.
- CHAP passwords (secrets) can be 1 to 512 bytes. Passwords can be hexadecimal values or strings. For hexadecimal values, enter the value with a prefix of “0x” or “0X”. A null password is not allowed.
- See the initiator’s documentation for additional restrictions. For example, the Microsoft iSCSI software initiator requires both the initiator and target CHAP passwords to be at least 12 bytes if IPsec encryption is not being used. The maximum password length is 16 bytes regardless of whether IPsec is used.

Upgrading from a previous release

If you upgrade from Data ONTAP 6.4.*x* to Data ONTAP 6.5 or later, and you have CHAP authentication configured, the CHAP configuration from the previous release is not saved. The CHAP configuration file in Data ONTAP 6.5 and later uses a new format that is not compatible with the CHAP configuration file format of the previous release. When you upgrade, you must use the `iscsi security` command to reconfigure CHAP settings.

If you do not reconfigure CHAP after the upgrade, Data ONTAP displays the following message when the initiator sends a login message to the storage system:

```
"ISCSI: Incorrect iSCSI configuration file version"
```

Defining an authentication method for an initiator

Defining an authentication method using the command line: To define an authentication method for an initiator that is in the authentication list, complete the following steps.

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

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

Defining an authentication method using FilerView: To define an authentication method for an initiator, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Initiator Security. **Result:** A list of initiators is displayed.

Step	Action	
2	**If you want to...**	**Then...**
	Define authentication for an initiator in the list	Check the box for the initiator and then click Modify. Proceed to Step 3.
	Define authentication for an initiator that is not in the list	**1.** Click Add New Initiator Security. **2.** Enter the initiator node name. Proceed to Step 3.
3	Select the security type.	
4	Click Next.	
5	If you selected CHAP, enter an inbound user name and password, and then click Next.	
6	If you selected CHAP, optionally enter an outbound user name and password, and then click Next.	
7	Click Commit.	

Displaying initiator authentication methods

Displaying initiator authentication methods using the command line: To view a list of initiators and their authentication methods, complete the following step.

Step	Action
1	On the storage system console, enter the following command: `iscsi security show`

Displaying initiator authentication methods using FilerView: To view a list of initiators and their authentication methods, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Manage Initiator Security.

Defining a default authentication method

To define a default authentication method for initiators that are not on the authentication list, complete the following step. You must use the command line to define a default method.

Step	Action
1	On the storage system console, enter the following command: **`iscsi security default -s method -p inpassword -n inname [-o outpassword -m outname]`** *method* is one of the following: ◆ `chap`—Authenticate using a CHAP user name and password. ◆ `none`—Initiators that are not on the list do not require authentication to access the storage system. ◆ `deny`—Initiators that are not on the list are denied access to the storage system. *inpassword* is the inbound password for CHAP authentication. The storage system uses the inbound password to authenticate the initiator. *inname* is a user name for inbound CHAP authentication. The storage system uses the inbound user name to authenticate the initiator. *outpassword* is a password for outbound CHAP authentication. The storage system uses this password for authentication by the initiator. *outname* is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Removing specific authentication settings for an initiator

Removing specific authentication settings using the command line:

To remove specific authentication settings for an initiator, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi security delete -i initiator`** `-i` *initiator* is the initiator name in the iSCSI node name format. **Result:** The initiator is removed from the authentication list and logs in to the storage system using the default authentication method.

Removing specific authentication settings using FilerView: To remove an initiator from the authentication list, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Initiator Security.
2	Check the box for the desired initiator.
3	Click Unset, and then click OK. **Result:** The specific security settings for the initiator are removed and the initiator logs in to the storage system using the default authentication method.

Managing target portal groups

About target portal groups

A target portal group is a set of one or more storage system network interfaces that can be used for an iSCSI session between an initiator and a target. A target portal group is identified by a name and a numeric tag.

For iSCSI sessions that use multiple connections, all of the connections must use interfaces in the same target portal group. Each interface belongs to one and only one target portal group. Interfaces can be physical interfaces or logical interfaces (VLANs and vifs). By default, each interface is in its own target portal group.

Prior to Data ONTAP 7.1, each interface was assigned to its own target portal group. The target portal group tag was assigned based on the interface location and could not be modified. This works fine for single-connection sessions.

Starting with Data ONTAP 7.1, you can explicitly create target portal groups and assign tag values. If you want to use multi-connection sessions, you need to create one or more target portal groups.

Because a session can use interfaces in only one target portal group, you may want to put all of your interfaces in one large group. However, some initiators are also limited to one session with a given target portal group. To support multipath I/O (MPIO), you need to have one session per path, and therefore more than one target portal group.

CAUTION

Some initiators, including those used with Windows, HP-UX, and Linux, create a persistent association between the target portal group tag value and the target. If the target portal group tag changes, the LUNs from that target will be unavailable.

When you migrate from a prior version of Data ONTAP to version 7.1, the target portal group tags will change. See the *Data ONTAP 7.1 Upgrade Guide* and the *Data ONTAP Release Notes* for information on how to migrate from an earlier release when you have iSCSI LUNs used by these operating systems.

CAUTION

When used with multi-connection sessions, the Windows iSCSI software initiator creates a persistent association between the target portal group tag value and the target interfaces. If the tag value changes while an iSCSI session is active, the initiator will be able to recover only one connection for a session. To recover the remaining connections, you must refresh the initiator's target information.

Listing target portal groups

You can view a list of the current target portal groups using the command line or FilerView. For each target portal group, the list includes the name, tag, and the interfaces that belong to the group.

Listing target portal groups using the command line: To list target portal groups, complete the following step.

Step	Action
1	On the storage system console, enter the following command: `iscsi tpgroup show`

Listing target portal groups using FilerView: To list target portal groups, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups.

Creating a target portal group

If you do not plan to use multi-connection iSCSI sessions, you do not need to create target portal groups.

If you do plan to use multi-connection sessions, create a target portal group that contains all of the interfaces you want to use for one iSCSI session.

When you create a target portal group, the specified interfaces are removed from their current groups and added to the new group. Any iSCSI sessions using the specified interfaces are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group will not be able to reconnect.

Creating a target portal group using the command line: To create a target portal group, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi tpgroup create [-f] tpgroup_name [-t tag] [interface ...]`** `-f` forces the new group to be created, even if that terminates an existing session using one of the interfaces being added to the group. *tpgroup_name* is the name of the group being created (1 to 32 characters, no spaces or non-printing characters). `-t` *tag* sets the target portal group tag to the specified value. In general you should accept the default tag value; see "About target portal groups" on page 236 for more information. User-specified tags must be in the range 1 to 256. *interface ...* is the list of interfaces to include in the group, separated by spaces. **Example:** The following command creates a target portal group named server_group that includes interfaces e8a and e9a: `iscsi tpgroup create server_group e8a e9a`

Creating a target portal group using FilerView: To create a target portal group, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups.
2	Click Create Portal Group.
3	Enter the Portal Group Name (1 to 32 characters, no spaces or non-printing characters).
4	Select the interfaces to include in the group. Use Ctrl-click to select multiple interfaces.

Step	Action
5	Optionally enter the User-Defined Tag. In general you should leave this field blank to accept the default tag value; see “About target portal groups” on page 236 for more information. User-specified tags must be in the range 1 to 256.
6	Click Create, and then click OK.

Destroying a target portal group

Destroying a target portal group removes the group from the storage system. Any interfaces that belonged to the group are returned to their individual default target portal groups. Any iSCSI sessions with the interfaces in the group being destroyed will be terminated.

Destroying a target portal group using the command line: To destroy a target portal group, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi tpgroup destroy [-f] tpgroup_name`** `-f` forces the group to be destroyed, even if that terminates an existing session using one of the interfaces in the group. *tpgroup_name* is the name of the group being destroyed.

Destroying a target portal group using FilerView: To destroy a target portal group, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups.
2	Select the Portal Group. Note that you cannot destroy the default target portal groups.
3	Click Destroy, and then click OK.

Adding interfaces to an existing target portal group

You can add interfaces to an existing target portal group. The specified interfaces are removed from their current groups and added to the new group. Any iSCSI sessions using the specified interfaces are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group will not be able to reconnect.

Adding interfaces using the command line: To add one or more interfaces to an existing target portal group, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi tpgroup add [-f] tpgroup_name [interface ...]`** `-f` forces the interfaces to be added, even if that terminates an existing session using one of the interfaces being added to the group. *tpgroup_name* is the name of the group. *interface ...* is the list of interfaces to add to the group, separated by spaces. **Example:** The following command adds interfaces e8a and e9a to the portal group named server_group: `iscsi tpgroup add server_group e8a e9a`

Adding interfaces using FilerView: To add one or more interfaces to an existing target portal group, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups.
2	Click the name of the Portal Group in the list. Note that you cannot add interfaces to the default target portal groups. **Result:** The Modify iSCSI Portal Group page is displayed, with the current interfaces highlighted in the Interfaces field.
3	In the Interfaces field, select all of the interfaces you want in the target portal group. Be sure the current interfaces are selected as well as the interfaces being added. Use Ctrl-click to select multiple interfaces.

Step	Action
4	Click Modify, and then click OK.

Removing interfaces from an existing target portal group

You can remove interfaces from an existing target portal group. The specified interfaces are removed from the group and returned to their individual default target portal groups. Any iSCSI sessions with the interfaces being removed are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group will not be able to reconnect.

Removing interfaces using the command line: To remove one or more interfaces from an existing target portal group, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi tpgroup remove [-f] tpgroup_name [interface ...]`** `-f` forces the interfaces to be removed, even if that terminates an existing session using one of the interfaces being removed from the group. *tpgroup_name* is the name of the group. *interface ...* is the list of interfaces to remove from the group, separated by spaces. **Example:** The following command removes interfaces e8a and e9a from the portal group named server_group, even though there is an iSCSI session currently using e8a: `iscsi tpgroup remove -f server_group e8a e9a`

Removing interfaces using FilerView: To remove one or more interfaces from an existing target portal group, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups.

Step	Action
2	Click the name of the Portal Group in the list. Note that you cannot add interfaces to the default target portal groups. **Result:** The Modify iSCSI Portal Group page is displayed, with the current interfaces highlighted in the Interfaces field.
3	In the Interfaces field, deselect the interfaces you want to remove. Be sure the interfaces you want to keep are still selected. Use Ctrl-click to select or deselect multiple interfaces.
4	Click Modify, and then click OK.

Displaying statistics for iSCSI sessions

Displaying iSCSI statistics

Displaying iSCSI statistics using the command line: To display iSCSI statistics, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi stats`
	Result: The following information is displayed:
	`iSCSI PDUs Received`
	`  SCSI-Cmd:  229  |  Nop-Out:        0  |  SCSI TaskMgtCmd:   0`
	`  LoginReq:    4  |  LogoutReq:      1  |  Text Req:          1`
	`  DataOut:     0  |  SNACK:          0  |  Unknown:           0`
	`  Total: 235`
	`iSCSI PDUs Transmitted`
	`  SCSI-Rsp:  221  |  Nop-In:         0  |  SCSI TaskMgtRsp:   0`
	`  LoginRsp:    4  |  LogoutRsp:      1  |  TextRsp:           1`
	`  Data_In:    17  |  R2T:            0  |  Asyncmsg:          0`
	`  Reject:      0`
	`  Total: 244`
	`iSCSI CDBs`
	`  DataIn Blocks:  127  |  DataOut Blocks:       0`
	`  Error Status:     3  |  Success Status:     226`
	`  Total CDBs: 229`
	`iSCSI ERRORS`
	`  Failed Logins:     1  |  Failed TaskMgt:      0`
	`  Failed Logouts:    0  |  Failed TextCmd:      0`
	`  Protocol:          0`
	`  Digest:            0`
	`  PDU discards (outside CmdSN window): 0`
	`  PDU discards (invalid header):       0`
	`  Total: 145`

Displaying iSCSI statistics using FilerView: To display iSCSI statistics, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Report. **Result:** The following information is displayed:

iSCSI Report

LUNs → iSCSI → Report

iSCSI status

```
iSCSI service is running
```

iSCSI nodename

```
iSCSI target nodename: iqn.1992-08.com.netapp:sn.33604646
```

iSCSI alias

```
iSCSI target alias: toaster
```

iSCSI portal show

```
Network portals:
IP address        TCP Port  TPGroup  Interface
10.60.128.100         3260     1000    e0
192.168.12.12         3260     1001    e9a
10.60.128.99          3260     1002    e9b
```

iSCSI Statistics

```
iSCSI PDUs Received
  SCSI-Cmd:        1726   | Nop-Out:            2   | SCSI TaskMgtCmd:          0
  LoginReq:          68   | LogoutReq:         15   | Text Req:                10
  DataOut:            0   | SNACK:              0   | Unknown:                  0
  Total: 1821
iSCSI PDUs Transmitted
  SCSI-Rsp:         425   | Nop-In:             2   | SCSI TaskMgtRsp:          0
  LoginRsp:          68   | LogoutRsp:         15   | TextRsp:                 10
  Data_In:         1396   | R2T:                0   | Asyncmsg:                 0
  Reject:             0
  Total: 1916
iSCSI CDBs
  DataIn Blocks:         6913   | DataOut Blocks:          0
  Error Status:            22   | Success Status:       1704
  Total CDBs: 1726
iSCSI ERRORS
  Failed Logins:            0   | Failed TaskMgt:          0
  Failed Logouts:           0   | Failed TextCmd:          0
  Protocol:                 0
  Digest:                   0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header):       0
  Total: 0
```

Interpreting iSCSI statistics

The `iscsi stats` command and the FilerView LUNs > iSCSI > Report page both display the following statistics:

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

- SCSI-CMD—SCSI-level command descriptor blocks.
- LoginReq—Login request PDUs sent by initiators during session setup.
- DataOut—PDUs containing write operation data that did not fit within the PDU of the SCSI command. The PDU maximum size is set by the storage system during the operation negotiation phase of the iSCSI login sequence.
- Nop-Out—A message sent by initiators to check whether the target is still responding.
- Logout-Req—A request sent by initiators to terminate active iSCSI sessions or to terminate one connection of a multi-connection session.
- SNACK—A PDU sent by the initiator to acknowledge receipt of a set of DATA_IN PDUs or to request retransmission of specific PDUs.
- SCSI TaskMgtCmd—SCSI-level task management messages, such as ABORT_TASK and RESET_LUN.
- Text-Req—Text request PDUs that initiators send to request target information and renegotiate session parameters.

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

- SCSI-Rsp—SCSI response messages.
- LoginRsp—Responses to login requests during session setup.
- DataIn—Messages containing data requested by SCSI read operations.
- Nop-In—Responses to initiator Nop-Out messages.
- Logout-Rsp—Responses to Logout-Req messages.
- R2T—Ready to transfer messages indicating that the target is ready to receive data during a SCSI write operation.
- SCSI TaskMgtRsp—Responses to task management requests.
- TextRsp—Responses to Text-Req messages.
- Asyncmsg—Messages the target sends to asynchronously notify the initiator of an event, such as the termination of a session.
- Reject—Messages the target sends to report an error condition to the initiator, for example:
  - Data Digest Error (checksum failed)
  - Target does not support command sent by the initiator
  - Initiator sent a command PDU with an invalid PDU field

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

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

Displaying information for iSCSI sessions and connections

Types of session and connection information

An iSCSI session can have zero or more connections. Typically a session has at least one connection. Connections can be added and removed during the life of the iSCSI session.

You can display information about all sessions or connections, or only specified sessions or connections. The `iscsi session show` command displays session information, and the `iscsi connection show` command displays connection information. The session information is also available using FilerView.

The command line options for these commands control the type of information displayed. For troubleshooting performance problems, the session parameters (especially HeaderDigest and DataDigest) are of particular interest.The `-v` option displays all available information. In FilerView, the iSCSI Session Information page has buttons that control which information is displayed.

Displaying session information

Displaying session information using the command line: To display iSCSI session information, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi session show [-v] [-t] [-p] [-c] [`*`session_tsih`*` ...]`** `-v` displays all information and is equivalent to `-t -p -c`. `-t` displays the TCP connection information for each session. `-p` displays the iSCSI session parameters for each session. `-c` displays the iSCSI commands in progress for each session. *session_tsih* is a list of session identifiers, separated by spaces.

Step	Action

	Example 1: The following is the output for one session using the `-t` option:
	`Session 27`
	`  Initiator Information`
	`    Initiator Name: iqn.1991-05.com.microsoft:host1`
	`    ISID: 40:01:37:00:00:00`
	
	`  Connection Information`
	`  Connection 2`
	`    Remote Endpoint: 10.60.8.60:4325`
	`    Local Endpoint: 10.60.128.99:3260`
	`    Local Interface: e9b`
	`    TCP recv window size: 132300`

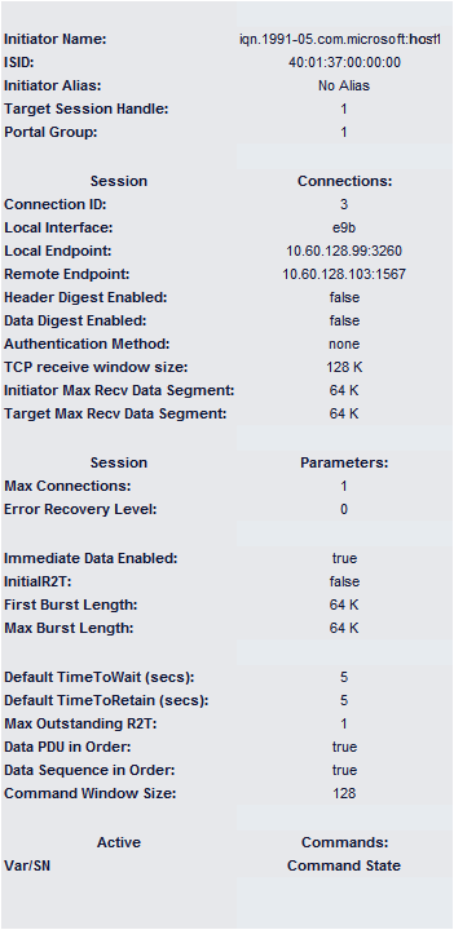
	Example 2: The following is the output for one session using the `-p` option:
	`Session 27`
	`  Initiator Information`
	`    Initiator Name: iqn.1991-05.com.microsoft:host1`
	`    ISID: 40:01:37:00:00:00`
	
	`  Session Parameters`
	`    SessionType=Normal`
	`    TargetPortalGroupTag=1`
	`    MaxConnections=4`
	`    ErrorRecoveryLevel=2`
	`    AuthMethod=None`
	`    HeaderDigest=None`
	`    DataDigest=None`
	`    ImmediateData=Yes`
	`    InitialR2T=No`
	`    FirstBurstLength=65536`
	`    MaxBurstLength=65536`
	`    Initiator MaxRecvDataSegmentLength=65536`
	`    Target MaxRecvDataSegmentLength=65536`
	`    DefaultTime2Wait=0`
	`    DefaultTime2Retain=20`
	`    MaxOutstandingR2T=1`
	`    DataPDUInOrder=Yes`
	`    DataSequenceInOrder=Yes`
	`    Command Window Size: 128`

Step	Action

	Example 3: The following is the output for one session using the `-c` option:
	`Session 27`
	`  Initiator Information`
	`    Initiator Name: iqn.1991-05.com.microsoft:host1`
	`    ISID: 40:01:37:00:00:00`
	
	`Command Information`
	`  Var/SN  State`
	`  Seq/164  Scsicdb_SR_Waiting_StatSN_ACK`

Displaying session information using FilerView: To display iSCSI session information, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Initiators.
2	Click the name of an initiator in the list. **Result:** The iSCSI Session Information page is displayed.

Step	Action
3	Click Show All, Show Commands, Show Session Parameters, or Show Connections to select the type of information to display. **Example:** The following page show the results of clicking Show All. 

Initiator Name:	iqn.1991-05.com.microsoft:host1
ISID:	40:01:37:00:00:00
Initiator Alias:	No Alias
Target Session Handle:	1
Portal Group:	1
Session	**Connections:**
Connection ID:	3
Local Interface:	e9b
Local Endpoint:	10.60.128.99:3260
Remote Endpoint:	10.60.128.103:1567
Header Digest Enabled:	false
Data Digest Enabled:	false
Authentication Method:	none
TCP receive window size:	128 K
Initiator Max Recv Data Segment:	64 K
Target Max Recv Data Segment:	64 K
Session	**Parameters:**
Max Connections:	1
Error Recovery Level:	0
Immediate Data Enabled:	true
InitialR2T:	false
First Burst Length:	64 K
Max Burst Length:	64 K
Default TimeToWait (secs):	5
Default TimeToRetain (secs):	5
Max Outstanding R2T:	1
Data PDU in Order:	true
Data Sequence in Order:	true
Command Window Size:	128
Active	**Commands:**
Var/SN	Command State

Displaying connection information

To display iSCSI connection parameters, complete the following step. You must use the command line to display this information.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi connection show [-v] [{new | session_tsih} conn_id]`
	`-v` displays all connection information.
	`new` *conn_id* displays information about a single connection that is not yet associated with a session identifier. You must specify both the keyword `new` and the connection identifier.
	session_tsih conn_id displays information about a single connection. You must specify both the session identifier and the connection identifier.
	Example: The following is the output for a single connection using the `-v` option.
	`Connection 38/1:`
	`  State: Full_Feature_Phase`
	`  Remote Endpoint: 10.60.8.60:3193`
	`  Local Endpoint: 10.60.128.99:3260`
	`  Local Interface: e9b`

Managing the iSCSI service on storage system interfaces

Command to use

You can use the `iscsi interface` command or FilerView to manage the iSCSI service on the storage system's Ethernet interfaces. You can control which network interfaces are used for iSCSI communication. For example, you can enable iSCSI communication over specific gigabit Ethernet (GbE) interfaces.

By default, the iSCSI service is enabled on all Ethernet interfaces after you enable the license. It is recommended that you do not use 10/100 megabit Ethernet interfaces for iSCSI communication. The e0 management interface on many storage systems is a 10/100 interface.

Displaying iSCSI status on storage system interfaces

Displaying iSCSI interface status using the command line: To display the status of the iSCSI service on storage system interfaces, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi interface show [-a|`*`interface`*`]`
	`-a` specifies all interfaces. This is the default.
	interface is list of specific Ethernet interfaces, separated by spaces.
	Example: The following example shows the iSCSI service enabled on two storage system Ethernet interfaces:
	`iscsi interface show`
	`Interface e0 disabled`
	`Interface e9a enabled`
	`Interface e9b enabled`

Displaying iSCSI interface status using FilerView: To display the status of the iSCSI service on storage system interfaces, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Manage Interfaces. **Result:** The iSCSI status of all storage system Ethernet interfaces is displayed.

Enabling iSCSI on a storage system interface

Enabling iSCSI on an interface using the command line: To enable the iSCSI service on an interface, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi interface enable {-a | interface ...}`** `-a` specifies all interfaces. *interface* is list of specific Ethernet interfaces, separated by spaces. **Example:** The following example enables the iSCSI service on interfaces e9a and e9b: `iscsi interface enable e9a e9b`

Enabling iSCSI on an interface using FilerView: To enable the iSCSI service on an interface, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Interfaces.
2	Select the check box for each interface you want to enable.
3	Click Enable, and then click OK.

Disabling iSCSI on a storage system interface

Disabling iSCSI on an interface using the command line: To disable the iSCSI service on an interface, complete the following step.

Step	Action
1	On the storage system console, enter the following command: **`iscsi interface disable [-f] {-a | interface ...}`** `-f` forces the termination of any outstanding iSCSI sessions without prompting you for confirmation. If you do not use this option, the command displays a message notifying you that active sessions are in progress on the interface and requests confirmation before terminating these sessions and disabling the interface. `-a` specifies all interfaces. *interface* is a list of specific Ethernet interfaces, separated by spaces.

Disabling iSCSI on an interface using FilerView: To disable the iSCSI service on an interface, complete the following steps.

Step	Action
1	Click LUNs > iSCSI > Manage Interfaces.
2	Select the check box for each interface you want to disable.
3	Click Disable, and then click OK.

Displaying the storage system's target IP addresses

The storage system's target IP addresses are the addresses of the interfaces used for the iSCSI protocol.

Displaying target IP addresses using the command line: To display the storage system's target IP addresses, complete the following step.

Step	Action

1	On the storage system console, enter the following command:
	`iscsi portal show`
	Result: The IP address, TCP port number, target portal group tag, and interface identifier are displayed for each interface.
	Example:
	`toaster> iscsi portal show`
	`Network portals:`
	`IP address  TCP Port  TPGroup  Interface`
	`10.60.128.99  3260  1  e9a`
	`10.60.128.100  3260  2  e9b`

Displaying target IP addresses using FilerView: To display the storage system's target IP addresses, complete the following step.

Step	Action
1	Click LUNs > iSCSI > Portal Addresses. **Result:** The IP address, TCP port number, target portal group tag, and interface identifier are displayed for each interface.

Using iSCSI on clustered storage systems

About clustered storage systems

A clustered storage system consists of two storage systems that are both connected to the same set of disks. If one storage system fails, its partner storage system can take over for the failed system and continue to make its data available.

About CFO

The takeover process is called cluster failover (CFO). During CFO, the surviving storage system responds to iSCSI hosts as if it were the original storage system. Specifically, the partner assumes the IP addresses, iSCSI target identities, LUNs, igroups, CHAP settings, and other settings of the failed system.

From the host's perspective, during CFO the target stops responding to the initiator and the iSCSI session is lost. Then the target (now running on the partner storage system) resumes responding to the initiator, and a new iSCSI session is established. If the initiator had outstanding SCSI commands, the initiator resends those commands and the new target processes them. The host is not aware that the CFO took place, only that the target did not respond for a short period of time and that the iSCSI session had to be reestablished.

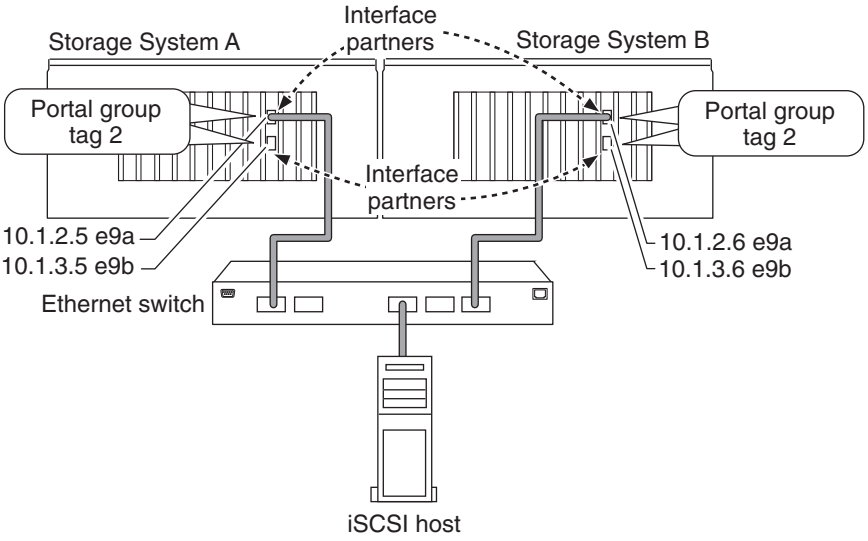
Requirements for clustered iSCSI systems

For CFO to work correctly, the two storage systems must be configured correctly, and the TCP/IP network must be configured correctly. Of special concern are the target portal group tags configured on the two storage systems.

The best practice is to configure the two partners of the cluster identically:

- Use the same network cards in the same slots.
- Create the same networking configuration with the matching pairs of ports connected to the same subnets.
- Put the matching pairs of interfaces into the matching target portal groups and assign the same tag values to both groups.

Simple configuration: Consider the following simplified example. Storage system A has a two-port Ethernet card in slot 9. Interface e9a has the IP address 10.1.2.5, and interface e9b has the IP address 10.1.3.5. The two interfaces belong to a user-defined target portal group with tag value 2.



Storage system B has the same Ethernet card in slot 9. Interface e9a is assigned 10.1.2.6, and e9b is assigned 10.1.3.6. Again, the two interfaces are in a user-defined target portal group with tag value 2.

In the cluster configuration, interface e9a on storage system A is the partner of e9a on storage system B. Likewise, e9b on system A is the partner of e9b on system B. For more information on configuring interfaces for a cluster, see the Data ONTAP *Cluster Installation and Administration Guide*.

Now assume that storage system B fails and its iSCSI sessions are dropped. Storage system A assumes the identity of storage system B. Interface e9a now has two IP addresses: its original address of 10.1.2.5, and the 10.1.2.6 address from storage system B. The iSCSI host that was using storage system B reestablishes its iSCSI session with the target on storage system A.

If the e9a interface on storage system A was in a target portal group with a different tag value than the interface on storage system B, the host might not be able to continue its iSCSI session from storage system B. This behavior varies depending on the specific host and initiator.

To ensure correct CFO behavior, both the IP address and the tag value must be the same as on the failed system. And because the target portal group tag is a property of the interface and not the IP address, the surviving interface cannot change the tag value during a CFO.

Complex configurations: If your cluster has a more complex networking configuration, including virtual interfaces (vifs) and VLANs, follow the same best practice of making the configurations the same. For example, if you have a vif on storage system A, create the same vif on storage system B. Make sure the target portal group tag assigned to each vif is the same. The name of the target portal group does not have to be the same, only the tag value matters.

Troubleshooting common iSCSI problems

LUNs are not visible on the host

The iSCSI LUNs appear as local disks to the host. If the storage system LUNs are not available as disks on the hosts, verify the following configuration settings.

Configuration setting	What to do
Cabling	Verify that the cables between the host and the storage system are properly connected.
Network connectivity	Verify that there is TCP/IP connectivity between the host and the storage system. ◆ From the storage system command line, ping the host interfaces that are being used for iSCSI. ◆ From the host command line, ping the storage system interfaces that are being used for iSCSI.
System requirements	Verify that the components of your configuration are qualified by NetApp. Verify that you have the correct host operating system (OS) service pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up to date system requirements in the NetApp iSCSI Support Matrix at the following URL: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml
Jumbo frames	If you are using jumbo frames in your configuration, ensure that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system according to the procedure described in “Managing the iSCSI service” on page 216.
Initiator login	Verify that the initiator is logged in to the storage system. See “Displaying initiators connected to the storage system” on page 228. If the command output shows no initiators are logged in, check the initiator configuration on the host. Verify that the storage system is configured as a target of the initiator.

Configuration setting	What to do
iSCSI node names	Verify that you are using the correct initiator node names in the igroup configuration. For the storage system, see “Managing igroups” on page 94. On the host, use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host must match.
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system console, use one of the following commands: `lun show -m`—Displays all LUNs and the igroups they are mapped to. `lun show -g` *`igroup-name`*—Displays the LUNs mapped to a specific igroup. Or, using FilerView, Click LUNs > Manage—Displays all LUNs and the igroups that are mapped to. For more information, see “Creating LUNs, igroups, and LUN maps” on page 57.

Storage system cannot register with iSNS server

Different iSNS server versions follow different draft levels of the iSNS specification. If there is a mismatch between the iSNS draft version used by the storage system and by the iSNS server, the storage system will not be able to register.

For more information, see “Resolving iSNS service version incompatibility” on page 222.

No multi-connection session

All of the connections in a multi-connection iSCSI session must go to interfaces on the storage system that are in the same target portal group. If an initiator is unable to establish a multi-connection session, check the portal group assignments of the initiator. For more information, see “Managing target portal groups” on page 236.

If an initiator can establish an multi-connection session, but not during a cluster failover (CFO), the target portal group assignment on the partner storage system is probably different from the target portal group assignment on the primary storage system. For more information, see “Requirements for clustered iSCSI systems” on page 256.

Sessions constantly connecting and disconnecting during CFO takeover

An iSCSI initiator that uses multipath I/O will constantly connect and disconnect from the target during cluster failover (CFO) if the target portal group is not configured correctly. The interfaces on the partner storage system must have the same target portal group tags as the interfaces on the primary storage system. For more information, see "Requirements for clustered iSCSI systems" on page 256.

Storage system iSCSI error messages

The following section describes some common iSCSI error messages, explanations of these messages, and what to do.

Message	Explanation	What to do
`ISCSI: Incorrect iSCSI configuration file version`	If you upgrade from Data ONTAP 6.4.*x* and you have CHAP authentication configured, the CHAP configuration from the previous release is not saved. The CHAP configuration file in Data ONTAP 6.5 uses a new format that is not compatible with the CHAP configuration file format of the previous release.	Use the `iscsi security` command to reconfigure CHAP settings. For detailed information, see "Managing security for iSCSI initiators" on page 229.
`ISCSI: network interface` *`identifier`* `disabled for use; incoming connection discarded`	The iSCSI service is not enabled on the interface.	Use the `iscsi` command or FilerView LUNs > iSCSI > Manage Interfaces page to enable the iSCSI service on the interface. **Example:** `iscsi interface enable e9b`

Message	Explanation	What to do
`ISCSI: Authentication failed for initiator` *`nodename`*	CHAP is not configured correctly for the specified initiator.	Check CHAP settings. ◆ Inbound credentials on the storage system must match outbound credentials on the initiator. ◆ Outbound credentials on the storage system must match inbound credentials on the initiator. ◆ You cannot use the same user name and password for inbound and outbound settings on the storage system. For detailed information, see "Managing security for iSCSI initiators" on page 229

Managing the Fibre Channel SAN 13

About this chapter

This chapter provides an overview of how to manage adapters, initiators, igroups, and traffic in a Fibre Channel environment.

Topics in this chapter

This chapter discusses the following topics:

- "Managing the FCP service" on page 264
- "Managing the FCP service on systems with onboard ports" on page 268
- "Displaying information about HBAs" on page 280

Managing the FCP service

Commands to use

You use the `fcp` 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 `fcp` commands, enter the `fcp help` command at the storage system prompt.

You can also use FilerView and go to

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

Verifying that FCP service is running

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

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

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

Verifying that the FCP service is licensed

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

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

Enabling the FCP service

To enable the FCP service, complete the following step.

Step	Action
1	Enter the following command: `license add` ***`license_code`*** *license_code* is the license code you received when you purchased the FCP license.

For FAS270 storage systems: After you license the FCP service on an FAS270 storage system, you must reboot. When the storage system boots up, the port labeled Fibre Channel 2 is in SAN target mode. When you enter Data ONTAP commands that display adapter statistics, this port is slot 0, so the virtual ports are shown as 0c_0, 0c_1, and 0c_2. For detailed information, see "Managing the FCP service on systems with onboard ports" on page 268.

Starting and stopping FCP service

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

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

Taking HBAs offline and bringing them online

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

Step	Action
1	Enter the following command: **`fcp config adapter [up|down]`** **Example:** `fcp config 4a down` **Result:** The target HBA 4a is offline. If you enter `fcp config 4a up`, the target HBA is brought online.

Disabling the FCP license

To disable the FCP license, complete the following step.

Step	Action
1	Enter the following command: **`license delete service`** *service* is any service you can license. **Example:** `license delete fcp`

Changing the system's WWNN

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

CAUTION

If you have a storage system cluster, and the cfmode of each system is single_image, then you must change the WWNN on both systems. If both systems do not have the same WWNN, hosts cannot access LUNs on the cluster.

To change the WWNN, complete the following step.

Step	Action
1	Enter the following command: **`fcp nodename`** ***`nodename`*** *nodename* is a 64-bit WWNN address. **Example:** `fcp nodename 50:a9:80:00:02:00:8d:ff`

Managing the FCP service on systems with onboard ports

Storage systems with onboard ports

The following systems have onboard FCP adapters, or ports, that you can configure to connect to disk shelves or to operate in SAN target mode:

- FAS270 models
- FAS3000 models

FAS270 storage systems

FAS270 onboard ports: A FAS270 storage system 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.

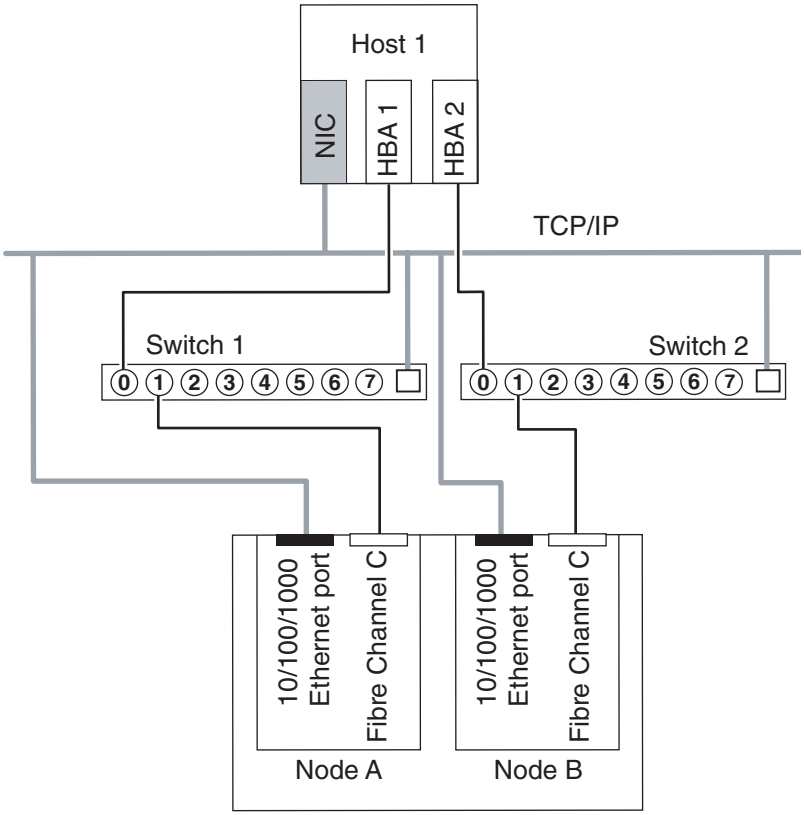
The Fibre Channel C port does not support mixed initiator/target mode. The default mode for this port is initiator mode. If you want to license the FCP service and connect the FAS270 to a SAN, you have to configure this port to operate in SAN target mode.

FAS270 cluster configuration example: FAS270 cluster configurations in dual_fabric mode 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.

FAS270 models also support single_image mode. If you upgrade your configuration to single_image mode, there are no switch restrictions. For information about changing your cfmode setting, see the *Fibre Channel Configuration Guide,* available from the following URL:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPConfigurationGuide.pdf

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/FCPConfigurationGuide.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	If the FCP protocol is not licensed, install the license by entering the following command:
	`license add FCP_code`
	FCP_code is the FCP service license code provided to you.
	Example:
	`fas270a> license add XXXXXXX`
	`A fcp site license has been installed.`
	`cf.takeover.on_panic is changed to on`
	`Run 'fcp start' to start the FCP service.`
	`Also run 'lun setup' if necessary to configure LUNs.`
	`A reboot is required for FCP service to become available.`
	`FCP enabled.`
	`fas270a> Fri Dec 5 14:54:24 EST [fas270a: rc:notice]: fcp licensed`

2	Reboot the FAS270 by entering the following command: **`reboot`**

Step	Action

3	Verify that the Fibre Channel C port is in target mode by entering the following command:
	`sysconfig`
	Example:
	`fas270a> 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: **Fibre Channel Target Host Adapter 0c**`
	`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 (0x00000000000001f0)`
	`0a.0 245MB`
	Note
	The Fibre Channel C port is identified as Fibre Channel Target Host Adapter 0c.

4	Start the FCP service by entering the following command:
	`fcp start`
	Example:
	`fas270a> fcp start`
	`FCP service is running.`
	`Wed Sep 17 15:17:04 GMT [fas270a: fcp.service.startup:info]: FCP service startup`

FAS3000Configuring the Fibre Channel port for initiator mode: To configure one or more onboard Fibre Channel ports to operate in initiator mode, complete the following steps.

series systems

Step	Action

1	Remove the FCP protocol license by entering the following command:
	`license delete fcp`
	adapter is the port number. You can specify more than one port.
	Example:
	`fas270a> license delete fcp`
	`Fri Dec 5 14:59:02 EST [fas270a: fcp.service.shutdown:info]: FCP service shutdown`
	`cf.takeover.on_panic is changed to off`
	`A reboot is required for TapeSAN service to become available.`
	`unlicensed fcp.`
	`FCP disabled.`
	`fas270a> Fri Dec 5 14:59:02 EST [fas270a: rc:notice]: fcp unlicensed`

2	Reboot the storage system by entering the following command: **`reboot`**

3	After the reboot, verify that the port 0c is in initiator mode by entering the following command:
	`sysconfig`
	Example:
	`fas270a> sysconfig`
	`NetApp RscrimshawN_030824_2300: Mon Aug 25 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:` **`Fibre Channel Target Host Adapter 0c`**
	`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 (0x00000000000001f0)`
	`0a.0 245MB`

Step	Action

4	Enable port 0c by entering the following command:
	`storage enable adapter 0c`
	Example:
	`fas270a> storage enable adapter 0c`
	`Mon Dec 8 08:55:09 GMT [rc:notice]: Onlining Fibre Channel adapter 0c.`
	`host adapter 0c enable succeeded`

FAS300 series onboard ports: The FAS3000 has four onboard Fibre Channel ports that have orange labels and are numbered 0a, 0b, 0c, 0d. Each port can be configured to operate in one of the following modes:

- SAN target mode, in which they connect to Fibre Channel switches or fabric.
- Initiator mode, in which they connect to disk shelves.

The operating mode of the Fibre Channel port depends on your configuration. See the following sections for information about the two recommended SAN configurations:

- "FAS3000 configuration with two Fibre Channel ports" below.
- "FAS3000 configuration using four onboard ports" on page 275

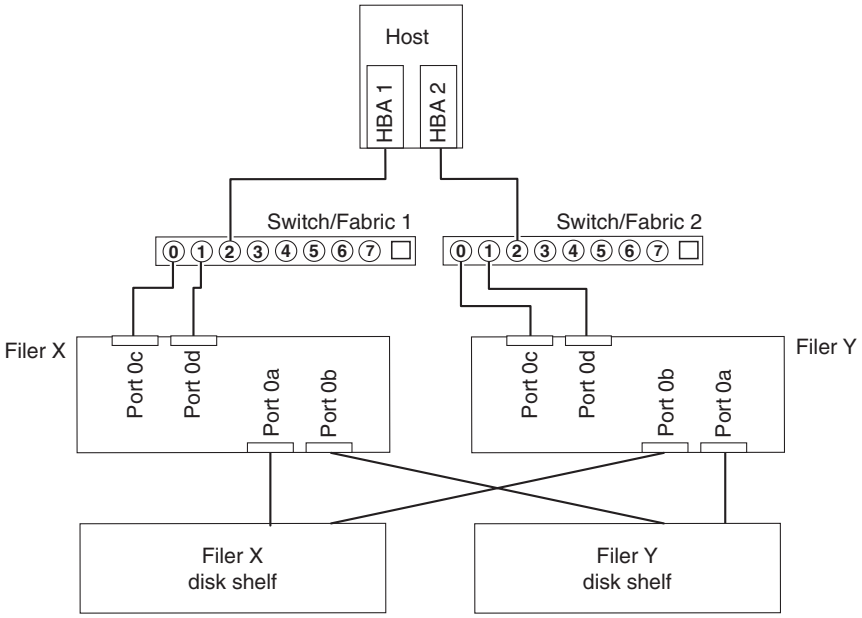
FAS3000 configuration with two Fibre Channel ports: The following figure shows the default SAN configuration in which a multi-attached host accesses a FAS3000 cluster. You cable the Fibre Channel ports as follows:

- Port 0a and 0b connect to the local and partner disk shelves.
- Port 0c and 0d connect to each FCP switch or fabric.

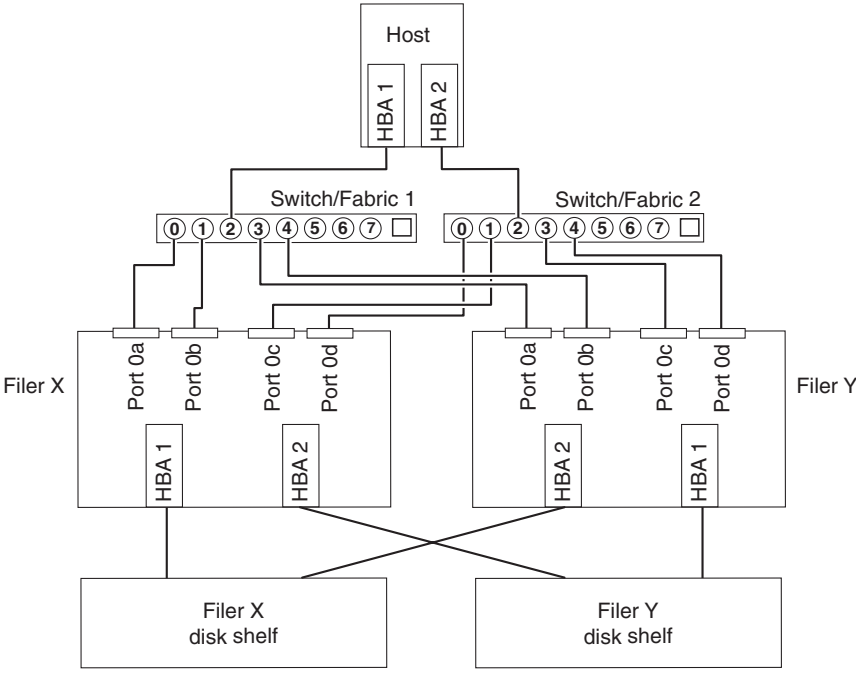
For detailed cabling instructions, see the *Installation and Setup Instructions* flyer that shipped with your system.

In this configuration, partner mode is the only supported cfmode of each node in the cluster. On each node in the cluster, port 0c provides access to local LUNs, and port 0d provides access to LUNs on the partner. This configuration requires that multipathing software is installed on the host.

If you order a FAS3000 system with the FCP license, ports 0a and 0b are preconfigured to operate in initiator mode. Ports 0c and 0d are preconfigured to operate in SAN target mode.


Host
HBA 1
HBA 2
Switch/Fabric 1
Switch/Fabric 2
0 1 2 3 4 5 6 7
Filer X
Port 0c
Port 0d
Port 0a
Port 0b
Filer Y
Port 0c
Port 0d
Port 0b
Port 0a
Filer X
disk shelf
Filer Y
disk shelf

FAS3000 configuration using four onboard ports: The following example shows a configuration that uses all four onboard Fibre Channel ports to connect to the SAN. On each storage system in the cluster, ports 0a and 0c connect to Switch/Fabric 1. Ports 0b and 0d connect to Switch/Fabric 2. Each storage system has two 64-bit Fibre Channel HBAs, which are used to connect to local and partner disk shelves.



In this configuration, the default cfmode of each node in the cluster is partner. On each node in the cluster, port 0a and 0c provide access to local LUNs, and ports 0b and 0d provide access to LUNs on the partner. This configuration requires that multipathing software is installed on the host.

Note

This configuration also supports the other cfmode settings. For information on changing the default cfmode from partner to another setting, see the online *NetApp Fibre Channel Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPConfigurationGuide.pdf.

If you ordered this configuration, then all four onboard ports are preconfigured to operate in target mode. If you have the two-port Fibre Channel configuration and want to upgrade to this configuration, then you have to configure ports 0c and 0d to operate in target mode by using the `fcadmin config` command.

Configuring the onboard ports for target mode: To configure the onboard ports to operate in target mode, complete the following steps.

Step	Action
1	If you have not licensed the FCP service, install the license by entering the following command: **`license add`** ***`license_code`*** *license_code* is the license code you received when you purchased the FCP license.
2	If you have already connected the port to a switch or fabric, take it offline by entering the following command: **`fcadmin config -d`** ***`adapter`*** *adapter* is the port number. You can specify more than one port. **Example:** The following example takes ports 0c and 0d offline. `fcadmin config -d 0c 0d`
3	Set the onboard ports to operate in target mode by entering the following command: **`fcadmin config -t target`** ***`adapter...`*** *adapter* is the port number. You can specify more than one port. **Example:** The following example sets onboard ports 0c and 0d to target mode. `fcadmin config -t target 0c 0d`
4	Reboot each system in the cluster by entering the following command: **`reboot`**

Step	Action

5	Start the FCP service by entering the following command:
	fcp start
	Example:
	`fas3050a> fcp start`
	`FCP service is running.`
	`Wed Mar 17 15:17:05 GMT`
	`[fas3050a:fcp.service.startup:info]: FCP service startup`

6	Verify that the Fibre Channel ports are online and configured in the correct state for your configuration by entering the following command:
	fcadmin config
	Example: The following output example shows the correct configuration of Fibre Channel ports for a four-port SAN configuration.
	Note
	The output might display the Local State of a target port as UNDEFINED on new systems. This is a default state for new systems. This state does not indicated that your port is misconfigured. It is still configured to operate in target mode.
	`fas3050a> fcadmin config`
	`Adapter Type Local State Status`
	`---------- ---------- ---------- ------`
	`0a target CONFIGURED online`
	`0b target CONFIGURED online`
	`0c target CONFIGURED online`
	`0d target CONFIGURED online`

Configuring the onboard ports for initiator mode: To configure one or more onboard Fibre Channel ports to operate in initiator mode, complete the following steps.

Step	Action
1	If you have not licensed the FCP service, install the license by entering the following command: **`license add license_code`** *license_code* is the license code you received from NetApp when you purchased the FCP license.
2	If you have already connected the port to a switch or fabric, take it offline by entering the following command: **`fcadmin config -d adapter`** *adapter* is the port number. You can specify more than one port. **Example:** The following example takes ports 0c and 0d offline. `fcadmin config -d 0c 0d`
3	Set the onboard ports to operate in target mode by entering the following command: **`fcadmin config -t target adapter...`** *adapter* is the port number. You can specify more than one port. **Example:** The following example sets onboard ports 0c and 0d to target mode. `fcadmin config -t target 0c 0d`
4	Reboot each system in the cluster by entering the following command: **`reboot`**

5	Start the FCP service by entering the following command:
	`fcp start`
	Example:
	`fas3050a> fcp start`
	`FCP service is running.`
	`Wed Mar 17 15:17:05 GMT [fas270a:`
	`fcp.service.startup:info]: FCP service startup`

Step	Action

6	Verify that the Fibre Channel ports are online and configured in the correct state for your configuration by entering the following command:
	`fcadmin config`
	Example: The following output example shows the correct configuration of Fibre Channel ports for a four-port SAN configuration.
	Note
	The output might display the Local State of a target port as UNDEFINED on new systems. This is a default state for new systems. This state does not indicated that your port is misconfigured. It is still configured to operate in target mode.
	`fas3050-1> fcadmin config`
	`Adapter Type Local State Status`
	`---`
	`0a target CONFIGURED online`
	`0b target CONFIGURED online`
	`0c target CONFIGURED online`
	`0d target CONFIGURED online`

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 cfmode setting and the storage system 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, Fibre Channel packet size, link data rate, SRAM parity, and various states	**storage show adapter**
All adapters (including HBAs, NICs, and switch port) configuration and status information	**sysconfig [-v] [*adapter*]** *adapter* is a numerical value only, for example, 5. -v displays additional information about all adapters.
Disks, disk loops, and options configuration information that affects coredumps and takeover	**sysconfig -c**
FCP cfmode setting	**fcp show cfmode**
FCP Traffic information	**sysstat -f**
How long FCP has been running	**uptime**
Initiator HBA port address, port name, node name, and igroup name connected to target HBAs	**fcp show initiator [-v] [*adapter&portnumber*]** -v displays the Fibre Channel host address of the initiator. *adapter&portnumber* is the slot number with the port number, a or b; for example, 5a.
Service statistics	**availtime**
Target HBAs configuration information	**fcp config**

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

Displaying information about all adapters

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

Step	Action

1	At the storage system, enter the following command to see information about all adapters.
	`sysconfig -v`
	Result: System configuration information and adapter information for each slot that is used is displayed on the screen. Look for `Fibre Channel Target Host Adapter` to get information about target HBAs.
	Note
	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.
	Note
	The output varies according to storage system model. For example, if you have a FAS270, the target port is displayed as `slot 0: Fibre Channel Target Host Adapter 0c`.
	Example: A partial display of information about a target HBA installed in slot 7 appears as follows:
	`slot 7: Fibre Channel Target Host Adapter 7a`
	`(Dual-channel, QLogic 2312 (2352) rev. 2, 64-bit, <ONLINE>)`
	`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, <ONLINE>)`
	`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`

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 `fcp config` command depends on the storage system model and cfmode setting. For examples, see Chapter 8, "Managing FCP in a clustered environment."

Step	Action

1	At the storage system, enter the following command.
	`fcp config`
	Sample output:
	`7a:  ONLINE <ADAPTER UP>  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 <ADAPTER UP>  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`
	Sample output for FAS270: For the FAS270, the `fcp config` command displays the target virtual local, standby, and partner ports.
	
	`0c:  ONLINE <ADAPTER UP>  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`
	Sample output for FAS3000: The fcp config command displays information about the onboard ports connected to the SAN:
	
	`0c:  ONLINE <ADAPTER UP>  PTP  Fabric`
	`  host address 010900`
	`  portname 50:0a:09:81:86:f7:a8:42  nodename 50:0a:09:80:86:f7:a8:42`
	`  mediatype ptp  partner adapter 0d`
	
	`0d:  ONLINE <ADAPTER UP>  PTP  Fabric`
	`  host address 010800`
	`  portname 50:0a:09:8a:86:47:a8:32  nodename 50:0a:09:80:86:47:a8:32`
	`  mediatype ptp  partner adapter 0c`

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	At the storage system, enter the following command:
	`fcp show adapter`
	Sample output for F8xx or FAS9xx series systems: The following sample output displays information for the HBA in slot 7:
	`Slot: 7a`
	`Description: Fibre Channel Target Adapter 7a (Dual-channel, QLogic 2312 (2352) rev. 2)`
	`Adapter Type: Local`
	`Status: ONLINE`
	`FC Nodename: 50:0a:09:80:86:87:a5:09 (500a09808687a509)`
	`FC Portname: 50:0a:09:83:86:87:a5:09 (500a09838687a509)`
	`Standby: No`
	
	`Slot: 7b`
	`Description: Fibre Channel Target Adapter 7b (Dual-channel, QLogic 2312 (2352) rev. 2)`
	`Adapter Type: Partner`
	`Status: ONLINE`
	`FC Nodename: 50:0a:09:80:86:57:11:22 (500a098086571122)`
	`FC Portname: 50:0a:09:8c:86:57:11:22 (500a098c86571122)`
	`Standby: No`
	Note
	In the display, the information about the Dual-channel QLogic HBA, the value 2312, does not specify the model number of the HBA; it refers to the device ID set by QLogic.
	Note
	For the FAS270, the `fcp show adapter` command displays the target virtual local (0c_0), standby (0c_1), and partner (0c_2) ports.

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	At the storage system, enter the following command:
	`fcp show initiator`
	Result: The following output is displayed:
	`Initiators connected on adapter 7a:`
	`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`
	
	`Initiators connected on adapter 7b:`
	`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	Enter the following command:
	`fcp stats -i interval [-c count] [-a | adapter]`
	`-i` interval is the interval, in seconds, at which the statistics are displayed.
	`-c` *count* is the number of intervals. For example, the `lun stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.
	-a shows statistics for all adapters.
	adapter is the slot and port number of a specific target HBA.
	Example output:
	`fcp 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`
	Explanation of output: Each column displays the following information:
	r/s—The number of SCSI read operations per second.
	w/s—The number of SCSI write operations per second.
	o/s—The number of other SCSI operations per second.
	ki/s— Kilobytes per second of received traffic
	ko/s—Kilobytes per second send traffic.
	asvc_t—Average time in milliseconds to process a request
	qlen—The average number of outstanding requests pending.
	hba—The HBA slot and port number.

Displaying FCP traffic information

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

Step	Action
1	Enter the following command: **sysstat -f** **Result:** The following output is displayed: CPU NFS CIFS FCP Net kB/s Disk kB/s FCP kB/s Cache in out read write in out 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 **Explanation of FCP statistics:** 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 system's cfmode setting is partner, mixed, or dual_fabric, you might want to obtain information about the amount of traffic coming to the system from its partner.

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

Step	Action
1	Enter the following command: **sysstat -b** **Result:** The following columns display information about partner traffic: ◆ Partner—The number of partner operations per second ◆ Partner KB/s—The number of kilobytes per second of incoming and outgoing partner partner traffic.

Displaying how long FCP has been running

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

Step	Action
1	Enter the following command: **uptime** **Result:** The following output is displayed: `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`

Displaying FCP service statistics

To display FCP service statistics, complete the following step.

Step	Action
1	Enter the following command: **availtime** **Result:** The following output is displayed: 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

Displaying the HBA's WWNN

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

Step	Action
1	Enter the following command: **fcp nodename** **Result:** Fibre Channel nodename: 50:a9:80:00:02:00:8d:b2 (50a9800002008db2)

Glossary

client A computer that shares files on a storage system. See also host.

HBA Host bus adapter. An I/O adapter that connects a host I/O bus to a computer's memory system in SCSI environments. The HBA might be an FCP adapter or an iSCSI adapter.

host Any computer system that accesses data on a storage system as blocks using the iSCSI protocol, or is used to administer a storage system.

igroup Initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to LUNs when they are mapped to those LUNs.

initiator The system component that originates an I/O command over an I/O bus or network; a host that has iSCSI initiator software installed on it, or a host that has a host bus adapter (HBA) installed in it, which is connected to the iSCSI or FCP network with the appropriate license enabled.

initiator group See igroup.

iSCSI A licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

iSCSI node name A logical name to identify an iSCSI node, with the format iqn.*yyyy-mm.backward_naming_authority*:sn.*unique_device_name.*

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

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

unique_device_name is a free-format unique name for this device assigned by the naming authority, preceded by sn. Typically, the *unique_device_name* is a serial number.

LUN A logical unit of storage.

LUN clone A complete copy of a LUN, which was initially created to be backed by a LUN or a file in a Snapshot copy. The clone creates a complete copy of the LUN and frees the Snapshot copy, which you can then delete.

LUN ID The numerical identifier that the storage system exports for a given LUN. The LUN ID is mapped to an igroup to enable host access.

LUN path The path to a LUN on the storage system. The following example shows a LUN path:

```
LUN path                Mapped to     LUN ID
-------------------------------------------
/vol/vol01/iscsidb.lun      igroup_1        6
```

LUN serial number The unique serial number for a LUN, as defined by the storage system.

map To create an association between a LUN and an igroup. A LUN mapped to an igroup is exported to the nodes in the igroup (iqn or eui) when the LUN is online. LUN maps are used to secure access relationships between LUNs and the host.

online Signifies that a LUN is exported to its mapped igroups. A LUN can be online only if it is enabled for read/write access.

offline Disables the export of the LUN to its mapped igroups. The LUN is not available to hosts.

qtree A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes. You can use qtrees to group LUNs.

SAN Storage Area Network. A storage network composed of one or more storage systems connected to one or more hosts in either a direct-attached or network-attached configuration using the iSCSI protocol over TCP/IP or the SCSI protocol over FCP.

share An entity that allows the LUN's data to be accessible through multiple file protocols such as NFS and iSCSI. You can share a LUN for read or write access, or all permissions.

space reservations An option that determines whether disk space is reserved for a specified LUN or file remains available for writes to any LUNs, files, or Snapshot copies. Required for guaranteed space availability for a given LUN with or without Snapshot copies.

storage system Hardware and software-based systems that serve and protect data using protocols for both SAN and NAS networks.

target The system component that receives a SCSI I/O command. A storage system with the iSCSI or FCP license enabled and serving the data requested by the initiator.

volume A file system. Volume refers to a functional unit of storage on a storage system, based on one or more RAID groups, that is made available to the host. LUNs are stored in volumes.

WWN World Wide Number. A unique 48- or 64-bit number assigned by a recognized naming authority (often through block assignment to a manufacturer) that identifies a connection for an FCP node to the storage network. A WWN is assigned for the life of a connection (device).

WWNN Worldwide 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. A WWNN is assigned to a storage system based on the serial number of its NVRAM. The WWNN is stored on disk. Data ONTAP refers to this number as a *Fibre Channel node name*, or simply, a node name.

WWPN Worldwide 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 *FC Portname*, or simply, a port name.

Index

Symbols

/etc/nvfail_rename, database protection 187

A

adapters
    displaying information about 280
administration
    iSCSI changes 212
aggregate
    defined 28
authentication
    defining default for CHAP 234
    using CHAP for iSCSI 229

B

backup
    data to tape 176
    hot backup mode 180
    native operation 180
    NDMP operation 180
    single LUNs to tape 177
    tape, when to use 180

C

changes for this release
    Data ONTAP 12
    iSCSI error recovery level 12
    iSCSI multi-connection sessions 12
    iSCSI target portal groups 12
CHAP
    authentication for iSCSI 229
    authentication, description of 17
    defining default authentication 234
    using with vFiler units 230
clones. *See* LUN clones
clustered storage systems
    about CFO with iSCSI 256
    options required 25
    using FCP on FAS270 268
    using iSCSI 18, 256
communication sessions, how they work 17
create_ucode option
    changing with the command line 56

D

Data ONTAP
    caution when upgrading 236
    changes for this release 12
    description of 2
Data ONTAP options
    automatically enabled 18, 25
    iscsi.isns.rev 222
    iscsi.max_connections_per_session 213
    iscsi.max_error_recovery_level 214
database protection
    using /etc/nvfail_rename 187
    using vol options nvfail 187
df command 140
disk space
    affected by snapshots 143
    displaying free 142
    monitoring 140
    monitoring with snapshots 148
    monitoring without snapshots 146
documentation
    downloading 6
    related 9
dual_fabric mode 131

E

error recovery level
    changes for this release 12
    enabling levels 1 and 2 214
eui type designator 15
Exchange
    performance 210
extents, logical 210

F

FAS270
- cluster for FCP 268
- dual-fabric mode 125
- onboard FCP ports 268
- switch requirement 125

FAS3000
- onboard FCP ports 273

FCP
- cfmode setting 25, 111
- changing WWNN 267
- displaying HBAs 280
- HBA online and offline 266
- licensed service 22
- nodes defined 22
- nodes, filer 23
- nodes, host 24
- nodes, how connected 22
- nodes, how identified 23
- nodes, switch 24
- systems with onboard ports 268, 273

fcp commands
- fcp config 266
- fcp nodename 267
- fcp show 281
- fcp show initiator 73
- fcp start 266
- fcp stats 281
- fcp status 264
- fcp stop 266

FCP service
- displaying how long running 289
- displaying traffic information about 288
- license 265
- starting and stopping 266

Fibre Channel ports
- onboard 268
- onboard FAS3000 273

filer administration
- using FilerView 3
- using the command line 2

filer node name, defined 15

filer, defined as target 2

FilerView
- adding interfaces to iSCSI target portal groups 240
- changes for iSCSI 213
- changing iSCSI target alias 221
- changing iSCSI target node name 219
- creating and mapping LUNs and igroups 72
- creating iSCSI target portal groups 238
- defining iSCSI authentication 232
- destroying iSCSI target portal groups 239
- disabling iSCSI on interfaces 254
- disabling iSNS 226
- displaying iSCSI authentication 233
- displaying iSCSI initiators 228
- displaying iSCSI interface status 253
- displaying iSCSI sessions 249
- displaying iSCSI statistics 244
- displaying iSCSI target addresses 255
- displaying iSCSI target alias 220
- displaying iSCSI target node name 218
- displaying iSCSI target portal groups 237
- enabling iSCSI license 217
- enabling iSCSI on interfaces 253
- launching 3
- registering with iSNS 225
- removing interfaces from iSCSI target portal groups 241
- removing iSCSI authentication 235
- starting and stopping iSCSI service 218
- verifying iSCSI license 217

flexible volumes
- described 28
- setting guarantees 45

fractional reserve
- 50 percent example 40
- affected by flexible volume guarantees 45
- calculating 51
- reducing 40
- setting to 0 44

free space, displaying for disks 142

G

guarantees, flexible volume 45

H

HBA

displaying information for FCP 280
displaying WWNN 290
FCP online and offline 266
host bus adapters
displaying information about 288, 289
initiator, displaying information about 286
Host Support Kit, defined 5

I

igroup commands for FCP
igroup add 103
igroup bind 135
igroup create 74, 100
igroup destroy 102
igroup remove 103
igroup set 104
igroup show 102, 103
igroup unbind 137
igroup commands for iSCSI
igroup add 95
igroup create 74, 94
igroup destroy 95
igroup remove 96
igroup set 96
igroup show 96
with vFiler units 97
initiator groups
adding for FCP 103
adding initiator for iSCSI 95
binding to portsets 135
creating for FCP 100
creating for FCP using sanlun 101
creating for iSCSI 94
creating with FilerView 72
defined 16, 23, 60
destroying for FCP 102
destroying for iSCSI 95
displaying for FCP 103
displaying for iSCSI 96
mapping to LUNs with FilerView 72
name rules 61
naming 61
ostype of 62
removing for FCP 103
removing initiator for iSCSI 96
requirements 61
requirements for creation 61
setting OS type for FCP 104
setting OS type for iSCSI 96
type of 62
unmapping LUNs from 83
with vFiler units 97
initiator HBAs, displaying information about 286
initiator, displaying for iSCSI 228
interface
disabling for iSCSI 254
enabling for iSCSI 253
managing for iSCSI 252
IP addresses
displaying for iSCSI 255
iqn node names, creating igroups with 16
iqn type designator 14
iSCSI
administration changes 212
command changes 212
connection, displaying 251
creating target portal groups 237
default TCP port 16
description of 13
destroying target portal groups 239
displaying initiators 228
displaying statistics 243
enabling error recovery levels 1 and 2 214
enabling on interface 253
Host Support Kit 5
iSNS 222
licence 217
listing target portal groups 237
managing storage system interfaces 252
multi-connection sessions, enabling 213
node name rules 219
node names, how used 16
nodes defined 13
nodes, how connected 13
nodes, how identified 14
security 229
service, start and stop 218
service, verifying 216
session, displaying 247

setup overview 19
target alias 220
target IP addresses 255
target node name 218
target portal groups defined 16, 236
troubleshooting 259
using with clustered storage systems 18
with clustered storage systems 256
iscsi commands
iscsi alias 220
iscsi connection 251
iscsi initiator 228
iscsi interface 252
iscsi isns 224
iscsi nodename 218
iscsi portal 255
iscsi security 231
iscsi session 247
iscsi start 218
iscsi stats 243
iscsi status 216
iscsi stop 218
iscsi tpgroup 237
iscsi.isns.rev option 222
iscsi.max_connections_per_session option 213
iscsi.max_error_recovery_level option 214
iSNS
description of 17
disabling 226
registering the storage system 224
server versions 222
service for iSCSI 222
updating immediately 226
with vFiler units 227
iswt command 212

L

license
FCP 265
iSCSI 217
logical extents 210
LUN clones
creating 166
creating snapshot 166
defined 165
displaying progress of split 167
splitting from snapshot 167
stopping split 167
lun commands
lun clone create 166
lun clone split 167
lun comment 85
lun destroy 86
lun help 88
lun map 75
lun move 84
lun offline 83
lun online 83
lun resize 85
lun set 195
lun set reservation 86
lun setup 65
lun share 87
lun show 91
lun snap usage 168
lun stats 90
lun unmap 83, 95
LUN creation
description attribute 59
host operating system type 58
LUN ID requirement 59
path name 58
size specifiers 58
space reservation default 59
LUN ID, ranges of 62, 63
LUNs
accessing with NAS protocols 86
bringing online 83
controlling availability 82
copying with vol copy command 181
creating with FilerView 72
creating with FilerView wizard 70
defined 13, 22
displaying mapping 91
displaying reads, writes, and operations for 90
enabling space reservations 86
management task list 82
mapping guidelines 63
mapping to igroups with FilerView 72

modifying description 85
reallocating to improve performance 200
removing 86
renaming 84
resizing restrictions 84
restoring 174
serial number 22
unmapping from initiator group 83
when using SnapDrive 57

M

man page command 3
Microsoft Exchange
performance 210
multi-connection sessions
changes for this release 12
enabling 213
MultiStore
creating iSCSI LUNs for vFiler units 78
vFiler units described 18

N

name rules
igroups 61
iSCSI node name 219
NDMP backup 180
node name
of initiator HBA, displaying 286
of storage system 15
rules for iSCSI 219
node type designator
eui 15
iqn 14
nvfail option, of vol options command 187
NVRAM failure 187

O

options
automatically enabled 18, 25
iscsi.isns.rev 222
iscsi.max_connections_per_session 213
iscsi.max_error_recovery_level 214
Oracle H.A.R.D., with SnapValidator 189

P

performance
improving 199
Microsoft Exchange read 210
plex, defined 28
portnames of initiator adapters, displaying 286
ports
resources, managing 25
used in clustered configurations 25, 111
portset commands
portset add 136
portset create 135
portset destroy 138
portset remove 136
portset show 137
portsets
adding ports 136
binding 135
creating 135
defined 133
destroying 138
how affect igroup throttles 134
in storage system clusters 133
removing ports 136
unbinding igroups 137
viewing ports in 137
when upgrading 134

Q

qtrees, defined 28

R

RAID-level mirroring
described 28
reallocate commands
reallocate off 209
reallocate on 201
reallocate quiesce 206
reallocate restart 207
reallocate schedule 204
reallocate start 202, 205
reallocate status 208
reallocate stop 208

reallocation
best practices 209
defining scans 202
deleting a scan 208
disabling scans 209
enabling scans 201
full 206
managing scans 201
quiescing scans 206
restarting scans 207
scans 200
scheduling scans 204
starting one-time scan 205
viewing scan status 207
with LUNs 200
restoring
LUNs 174
snapshots of LUNs 171

S

sanlun
creating igroups for FCP 101
fcp show adapter command 101
scans, reallocation 201
service
FCP 264
iSCSI 218
Single File SnapRestore, using with LUNs 173
snap commands
snap autodelete 154, 155
snap delta 140
snap reclaimable 140
snap restore 172, 174
snap reserve, setting the percentage 55
SnapDrive
cautions 57
for UNIX 7
for Windows 7
introduction 7
limitations 8
snapshots
deleting automatically 153
disk space 143
schedule, turning off 54
using with SnapRestore 171
SnapValidator
described 189
disabling checks on LUNs 196
disabling checks on volumes 196
disk offset for other hosts 195
disk offset for Solaris 194
disk offset on storage system 195
enabling checks on LUNs 193
enabling checks on volumes 192
error messges 197
licensing 191
preparing database files 190
when upgrading 196
Solaris hosts, with SnapValidator 194
space management policy
applying 156
defining 152
space reservation
how it persists 31
statistics
displaying for iSCSI 243
storage system
defined as target 2
definition 2
iSCSI node name 15
storage system administration
using FilerView 3
using the command line 2
storage units
types of 28
SyncMirror
use of plexes in 28

T

tape, backing up to 176
target alias for iSCSI 220
target node name, iSCSI 218
target portal groups
about 236
adding interfaces 240
caution when upgrading Data ONTAP 236
changes for this release 12
creating 237

defined 16
destroying 239
listing 237
removing interfaces 241
TCP, default port for iSCSI 16
traditional volumes
described 28
troubleshooting
iSCSI 259

V

vFiler units
authentication using CHAP 230
creating iSCSI LUNs for 78
how used 18
using iSCSI igroups 97
with iSNS 227
vol commands
vol copy 181
vol destroy 184, 185
vol options 188
vol options nvfail, using with LUNs 187
volume size
calculating 51
volumes
destroying (vol destroy) 184, 185
growing automatically 153
guidelines for creation 53

W

WWNN
changing for storage system 267
displaying for HBA 290
WWPN
creating igroups with 23
how assigned 24
identifying filer ports with 23