Data ONTAP™ 6.5
System Administration
Block Access Management Guide for iSCSI

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 210-00571
December 2003

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

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

   This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

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

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

Portions of the software were created by Netscape Communications Corp.

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

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

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.
2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [$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/"

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

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

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

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

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

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

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

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

Trademark information

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

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

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Preface ix

Chapter 1 **How NetApp Implements an iSCSI Network** 1

Understanding NetApp storage appliances 2

Understanding how NetApp implements an iSCSI network 5

Related documents 11

Chapter 2 **Creating LUNs, igroups, and LUN maps** 13

Configuring disk space 14

Guidelines for creating volumes 17

Guidelines for calculating the size of a volume 23

Methods of creating LUNs, igroups, and LUN maps 28

Creating LUNs with the lun setup program 33

Creating LUNs and igroups with FilerView 38

Creating LUNs and igroups using individual commands 41

Creating LUNs on vFilers 45

Chapter 3 **Managing LUNs** 49

Managing LUNs and LUN maps 50

Displaying LUN information 56

Monitoring disk space 59

Chapter 4 **Managing igroups** 69

Managing igroups. 70

Using igroups on vFilers 73

Chapter 5 **Using Data Protection with iSCSI** 75

Data ONTAP protection methods 76

Using snapshots. 78

Using LUN clones 80
Deleting busy snapshots 82
Using SnapRestore 88
Backing up data to tape 94
Using NDMP 98
Using volume copy 99
Using NVFAIL 100

Chapter 6 **Managing the iSCSI Network** 103
Registering the filer with an iSNS server. 104
Managing the iSCSI service 107
Managing security for iSCSI initiators 110
Displaying statistics for iSCSI sessions 114
Displaying information for iSCSI sessions 117
Managing the iSCSI service on filer interfaces. 119

Glossary 121

Index 125

Preface

About this guide

This guide describes how to use a NetApp® storage appliance as an iSCSI target in a TCP/IP storage network. Specifically, this guide describes how to calculate the size of volumes containing LUNs, how to create and manage LUNs and initiator groups (igroups), and how to monitor iSCSI traffic. The guide assumes you have completed the following tasks:

- Referred to the *FCP/iSCSI Configuration Guide* to ensure that your configuration is supported.
- Installed your storage appliance according to the instructions in the *Site Requirements Guide*, other installation documentation, such as the *System Cabinet Guide*, and the hardware and service guide for your specific storage appliance.
- Configured your filers according to the instructions in the following documents:
  - *Data ONTAP Software Setup Guide*
  - If you purchased a software initiator licence, see the NetApp documentation for your iSCSI Initiator Support Kit and the appropriate host-supplied documentation
  - If you purchased an iSCSI Host Bus Adapter, see the NetApp documentation for your SAN Host Attach Kit for iSCSI Protocol on your host, which you can find at http://now.corp.netapp.com/NOW/knowledge/docs/client_filer_index.shtml.

Audience

This guide is for system and storage administrators who are familiar with operating systems, such as Windows® 2000 and UNIX®, that run on the hosts that access storage managed by NetApp storage appliances. It also assumes that you know how block access protocols are used for block sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology.

Terminology

This guide uses the following terms:

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

Command conventions

In examples that illustrate commands executed on a UNIX workstation the command syntax and output might differ, depending on your version of UNIX.

Keyboard conventions

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

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
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 uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note

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

Caution

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

How NetApp Implements an iSCSI Network 1

About this chapter

This chapter provides a brief introduction to NetApp storage appliances, how they are administered, and an overview of how NetApp implements the iSCSI protocol in a NetApp iSCSI network.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding NetApp storage appliances" on page 2
- "Understanding how NetApp implements an iSCSI network" on page 5
- "Related documents" on page 11

Understanding NetApp storage appliances

What NetApp storage appliances are

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

In an iSCSI network, filers are targets that have storage target devices, which are referred to a s LUNs (logical unit numbers). With Data ONTAP™, you configure the filer's storage by creating LUNs that can be accessed by hosts, which are the initiators.

What Data ONTAP is

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

Data ONTAP supports a multiprotocol environment. You can configure a filer as a target device in an iSCSI network using the SCSI protocol over TCP/IP (using the iscsi service) or 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 filer as a storage device in a NAS network using NFS, CIFS, DAFS, HTTP, or FTP.

Ways to administer a filer

You can administer a filer by using the following methods:

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

Command-line administration: You can issue Data ONTAP commands at the filer's console or you can open a Telnet or Remote Shell (rsh) session from a host. An Ethernet network interface card (NIC) is preinstalled in the filer.

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

When using the command line, you can get command-line syntax help form 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 the `man` *`na_command_name`* command. For example, if you want to read the man page about the lun command, you would enter the following command: `man na_lun`.

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

FilerView administration: As an alternative to entering commands at the command line or using scripts or configuration files, you can use FilerView to perform many common tasks. FilerView is the graphical management interface for managing a filer from a Web browser or for viewing information about the filer, its storage units (such as volumes), LUNs, and adapters, and statistics about the storage units and iSCSI or FCP 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 filer, followed by /na_admin/ as the location for the URL. **Example:** If you have a filer named toaster, enter the following URL in the browser: http://toaster/na_admin. **Result:** The Network Appliance Online administrative window appears.
3	Click FilerView. **Result:** ◆ If the filer is password protected, you are prompted for a user name and password. ◆ Otherwise, FilerView is launched, and a screen appears with a list of topics in the left panel and the system status in the main panel.
4	Click any of the topics in the left panel to expand navigational links.

Understanding how NetApp implements an iSCSI network

What iSCSI is

iSCSI is a licensed service on the filer that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over TCP/IP.

What nodes are

In an iSCSI network, there are two types of nodes, targets and initiators. Targets are filers and initiators are hosts. Filers have storage devices, which are referred to as logical units of storage (LUNs,). Switches, routers, and ports are TCP/IP devices only and are not iSCSI nodes.

What LUNs are

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

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

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

What a LUN serial number is

A LUN serial number is a unique 12-byte filer-generated ASCII string. Many multipathing software packages use this serial number to identify redundant paths to the same LUN. You can display the LUN serial number with the `lun serial` *`lun_path`* or the `lun show -v` command.

How nodes are connected

Filers and hosts can be connected directly to each other, or they can be connected to TCP/IP switches with Ethernet cable and use a TCP/IP network for all connectivity.

iSCSI is implemented over the target and initiator's standard Ethernet interfaces, unless you purchased an iSCSI Host Attach kit, which includes a vendor-provided HBA that is installed on the host. In that case, iSCSI is implemented through the HBA on the host.

On the filer, you can select which interface will be used for iSCSI traffic. For example, you can specify that iSCSI traffic will run exclusively over GbE interfaces and not the integrated 10/100 interface. Or you can specify that iSCSI traffic will run exclusively on an interface connected to a specific dedicated iSCSI network. You use the `iswt interface` command to control which network interface is used for iSCSI connectivity.

Note

If you want to use the internal 10/100 connection on the filer for management, NetApp strongly recommends that you set the Ethernet NIC, the host HBA, and the switch (if one is being used) to full duplex.

How initiator nodes are connected to the network: There are two ways you can configure a host to connect to an iSCSI network, depending on your host's operating system.

- You can install the iSCSI initiator software on the host. For example, you can install the iSCSI Microsoft Windows Initiator software, the iSCSI Linux Initiator software, or the iSCSI NetWare Initiator software on your host.
- You can install a hardware iSCSI Host Bus Adapter (HBA) with the accompanying software in your host. For example, you can install the HBA for the iSCSI protocol on Windows in your Windows host.

How target nodes are connected to the network: The filer does not use a hardware iSCSI HBA to implement the iSCSI service. The iSCSI protocol on the filer is implemented using a software driver that is integrated into Data ONTAP as a virtual adapter called the iSCSI Software Target (*iswt*). This virtual adapter is pre-installed on the filer. This provides two logical iswt adapters: iswta and iswtb. The iswta adapter controls all iSCSI traffic for the filer. If the filer is in a cluster, the iswtb adapter is dedicated to iSCSI traffic for the partner filer. If a failover occurs and the filer is operating in takeover mode, the iswtb adapter becomes active.

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 target must use the iqn type designator. The initiator can use either the iqn 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, followed by a colon (:)
- A unique device name

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.netapp.

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

Filer node name: Each filer has a default node name based on the NetApp reverse domain name and the serial number of the filer's NVRAM card in the following format:

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

The following example shows the default node name for a filer with the serial number 123456789:

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

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 to display iSCSI initiators

You can use the `iscsi show initiator` command or FilerView (click LUNs > Initiator Groups > Manage) to display the node names of the iSCSI initiators that are registered with the iSNS service, or that you set up manually.

How to display an iSCSI target node name

You can use the `iscsi nodename` command to display the node name of the filer itself.

How to change a node name

You can use the `iscsi nodename` command to change the filer node name. If you change the node name, it must be unique and it should be consistent.

How node names are used

iSCSI node names are used for the following purposes.

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

Uniquely identifying a filer's iswt target ports: The node name of the filer is used to uniquely identify each iswt target port on a filer.

Understanding iSNS

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

If you have an iSNS server on your network, and it is configured and enabled for use by both the initiator and the filer, the filer will automatically register 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 filer 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.

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 filer.

During the initial stage of an iSCSI session, the initiator sends a login request to the filer to begin the session. The login request includes the initiator's CHAP user name and algorithm. The filer responds with a CHAP challenge. The initiator provides a CHAP response. The filer 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 a vendor-supplied HBA. The filer appears as a single iSCSI target node with one iSCSI node name. The filer supports one TCP connection per iSCSI session.

An iSCSI session can be conducted over a filer virtual network interface (vif) or a virtual LAN (VLAN) interface.

The iSCSI software adapter maps each Ethernet interface on the target to a single portal group. In this way, an initiator port can conduct simultaneous iSCSI sessions on the target, with one session for each portal group. The iswt driver allows 1024 simultaneous sessions, depending on the memory capacity of the filer. To determine whether your host's initiator software can take advantage of this functionality, see your host documentation.

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

Cluster configurations

The following options are automatically enabled when the iSCSI service is turned on.

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

You can use the NetApp Cluster Configuration Checker to detect inconsistent settings on clusters licensed for iSCSI.

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

How vFilers are used

If you purchased a MultiStore™ licence and created vFiler™ virtual filers, you can enable the iscsi license for each vFiler to manage LUNs and igroups on a per vFiler basis. For information about vFilers, see "Creating LUNs on vFilers" on page 45 and the sections on iSCSI service on vFilers or LUNs on vFilers in the *MultiStore Management Guide*.

Related documents

Where to go for more information

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

If you want...	Go to...
The most current system requirements for your host	*FCP/iSCSI Configuration Guide*
The latest information about the supported filer models for Data ONTAP licensed with iSCSI	◆ *FCP/iSCSI Configuration Guide* ◆ *System Configuration Guide*
Information about TCP/IP network features supported by Data ONAP	*Data ONTAP 6.5 Network Guide*
Information about how to install and configure iSCSI initiator software	◆ iSCSI host initiator Support Kit documentation from Network Appliance for your specific host at http://now.netapp.com/NOW/cgi-bin/software ◆ iSCSI host initiator documentation from the host at their web site, from a link on the NetApp NOW site at http://now.netapp.com/NOW/knowledge/docs/client_filer_index.shtml ◆ iSCSI Host Bus Adapter Attach Kit documentation for your specific host that was supplied with the adapter and at http://now.netapp.com/NOW/knowledge/docs/client_filer_index.shtml
The latest information about how to configure the iSCSI service on the filer	*Data ONTAP 6.5 Release Notes* (if available)

Creating LUNs, igroups, and LUN maps 2

About this chapter

This chapter describes the preparation process for creating LUNs and the methods for creating LUNs, igroups, and LUN maps. In addition, it describes how to create LUNs on vFilers.

This chapter assumes that your NetApp iSCSI network is set up and configured, and that the iSCSI service is licensed and enabled. If that is not the case, see "Managing the iSCSI Network" on page 103.

Topics in this chapter

This chapter discusses the following topics:

- "Configuring disk space" on page 14
- "Guidelines for creating volumes" on page 17
- "Guidelines for calculating the size of a volume" on page 23
- "Methods of creating LUNs, igroups, and LUN maps" on page 28
- "Creating LUNs on vFilers" on page 45

Configuring disk space

Filer storage units for managing disk space

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

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

The volume is the largest unit of storage. It can hold qtrees, files, and LUNs. A qtree is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume to group LUNs. You use LUNs to store data in a SAN environment, and you use files to store data in a NAS environment.

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

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

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

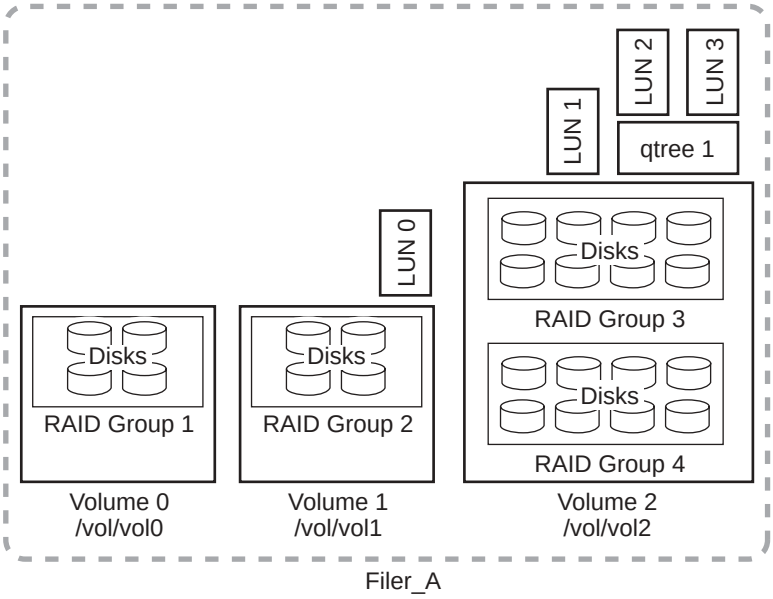

Why the space reservation option is important

LUNs appear as local disks attached to the host. As a result, it is imperative that LUNs behave as attached disks, which means that when you create a LUN of a specified size, the application storing data in the LUN must have guaranteed access to that amount of space. Data ONTAP achieves this with the space reservation option.

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

How space reservation persists

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

Why you need to change snapshot defaults

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

Snapshots cannot use reserved space, so snapshots can fail if there is not enough available space on the volume. When using snapshots, you need to calculate the amount of storage capacity required for the volume based on the information at "Guidelines for calculating the size of a volume" on page 23.

When you create a volume, Data ONTAP automatically

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

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

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

With Windows hosts, you use SnapDrive and SnapManager to ensure that applications accessing LUNs are quiesced or synchronized automatically before taking snapshots. With hosts that are not supported with SnapDrive, NetApp recommends that you ensure that the file system or application accessing the LUN is quiesced or synchronized before taking snapshots. For information about how to use snapshots, see "Using snapshots" on page 78.

Guidelines for creating volumes

Required information to create a volume

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

The capacity and the number of disks you specify determine the size of the volume. For example, if you have 72-GB disks and you use seven disks to create a volume, six disks are used for data, and one is used for parity; the volume capacity is 432 GB. Assuming that a 72-GB disk results in 67.9 GB of physical capacity, six 72-GB disks provide 407.4 GB of storage capacity. The actual amount of physical storage capacity might vary, depending on the disk manufacturer.

You calculate the size of the volume by determining the type and the amount of data that will be stored in the LUNs on the volume. For more information, see "Guidelines for calculating the size of a volume" on page 23.

Optional information to create a volume

In addition to the number of disks, you can specify the size of the RAID group. For information about disk allocation or RAID groups and how to configure them, see the *Data ONTAP 6.5 System Administration Storage Management Guide*. For information about disk configurations for your filer model, see the *System Configuration Guide*.

Guidelines to use when creating volumes

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

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

  When you first set up and configure the filer, Data ONTAP creates a /vol/vol0 root volume. This root volume is used by Data ONTAP to administer the filer.
- Create a volume large enough to accommodate the LUNs that will be stored on the volume, using the information at "Guidelines for calculating the size of a volume" on page 23. At a minimum, create volumes using seven disks.
- Ensure that the snapshot functionality is modified as follows:
  - Set snap reserve to 0%.
  - Turn off the automatic snapshot schedule.

- Ensure that no other files or directories exist in a volume that contains a LUN.

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

  Data ONTAP requires that the path of a volume or qtree containing a LUN is in the Unicode format. This option is On by default when you create a volume, but it is important to verify that any existing volumes still have this option enabled before creating LUNs in them.

Setting the percentage of snap reserve space using the command line

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

Step	Action
1	To set the percentage, enter the following command: **`snap reserve`** ***`volname percent`*** **Note** For volumes that contain LUNs and no snapshots, NetApp recommends 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.`

Setting the percentage of snap reserve space using FilerView

To use FilerView to set a percentage of snap reserve space (or *snapshot* reserve space, as it appears in FilerView) on a volume and to verify what percentage is set, complete the following steps.

Step	Action
1	Launch FilerView, as described in "Launching FilerView" on page 3.
2	Click Volumes > Snapshots > Configure.
3	Select the volume you want to configure.
4	At the Snapshot Reserve field, enter 0 as the percentage of space the volume reserves for snapshots. **Note** For volumes that contain LUNs and no snapshots, NetApp recommends that you set the percentage to zero.
5	Click the Apply button.

Turning off the automatic snapshot schedule using the command line

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

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

Turning off the automatic snapshot schedule using FilerView

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

Step	Action
1	Launch FilerView, as described in "Launching FilerView" on page 3.
2	Click Volumes > Snapshots > Configure.
3	Select the volume you want to configure.
4	At the Hourly Snapshot Schedule, ensure that no time slots are selected. For example, if a check appears at 8:00 AM, click it to deselect it.
5	Click the Apply button.

Verifying and modifying the volume option create_ucode using the command line

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

Step	Action

1	To verify that the `create_ucode` option is enabled (On), enter the following command:
	`vol status [`*`volname`*`] -v`
	Example: `vol status /vol1 -v`
	Result: The following output is a sample of what is displayed:
	`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, status on 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`

Verifying and modifying the volume option create_ucode using FilerView

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

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

Guidelines for calculating the size of a volume

What the volume size depends on

NetApp recommends that you use a minimum of seven disks to create a volume. so the size of your smallest volume should be six times your disk capacity (because one disk is for parity). In the example at "Required information to create a volume" on page 17, the disk size was 72 GB. A volume created with six of these disks provided a volume size of 407.4 GB.

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

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

  A backing store LUN means some or all of the unchanged data in the LUN is stored in the snapshot that was used to create the LUN.
- LUN clones

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

If you create a LUN and...	**Create a volume that...**	**Example (given that LUN = x bytes)**
Do not use snap-shots	Is at least 10% larger than the total size of the LUN (or the cumulative total of all the LUNs to be stored in the volume).	Volume size = x + 10%x or volume size = 1.10x. If LUN size = 10 GB, volume size = 1.10*10 GB or 11 GB

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Take one or more snapshots	Is at least two times the size of the LUN, plus delta. Delta is the rate of change, given as a percentage.	Volume size = 2x delta. If LUN size = 10 GB, and the rate of change since the snapshot is 15%, volume size = 2*10 GB+(.15*10 GB) = 21.5 GB. The first snapshot requires enough space to overwrite every block of data in the active file system; however, subsequent snapshots only require enough space to store what is overwritten in snapshot data. They do not require additional reserved space.
Create a backing store LUN using the `lun create -b` command	Has enough space to contain a backing store LUN, which is typically used for prototyping. This type of LUN only needs space for changed data, because the unchanged data it contains is stored in the snapshot from which it was created. You can create a backing store LUN with or without space reservation enabled, and you can create it as read-only, if desired.	With space reservation enabled (NOT using the `-o noreserve` option): volume size = original LUN size, plus snapshot requirements, plus space for this new LUN and its snapshot requirements. With space reservation disabled (using the `-o noreserve` option): Volume size = original LUN size, plus snapshot requirements.
Do a single file snap restore	Has enough space to store the original LUN and the restored LUN.	Volume size = 2x + 10%x or volume size = 2.1x.

If you create a LUN and...	Create a volume that...	Example (given that LUN = x bytes)
Create a clone of the backing store LUN using the `lun clone` command	Has enough space to store the original LUN and its snapshot, and the LUN clone and its snapshot. To clone a backing store LUN and make it readable-writable, enable space reservations on the LUN and treat it as you would any other LUN.	Volume size = original LUN size, plus its snapshot requirements, plus the LUN clone and its snapshot requirements. Typically, this is 4x + delta.

Creating a volume using the command line

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

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

Creating a volume using FilerView

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

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

Organizing LUNs by creating qtrees using the command line

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

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

Organizing LUNs by creating qtrees using FilerView

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

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

Methods of creating LUNs, igroups, and LUN maps

About the methods

You create LUNs using one of the following methods:

- Entering the `lun setup` command

  The 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 33.

- 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 38.

- Entering a series of individual commands (such as `lun create`, `igroup create`, and `lun map`)

  This method allows you to create one or more LUNs and igroups in any order. For information about this method, see "Creating LUNs and igroups using individual commands" on page 41.

- Using SnapDrive 3.0 or later for Windows

  SnapDrive furnishes a Windows-like interface to manage LUNs on a filer. This product provides online storage configuration, LUN expansion, and streamlined storage management. Integrated with NetApp Snapshot™ functionality, SnapDrive creates point-in-time images of data stored in LUNs that are connected to the filer, for backup, archiving, recovery, and asynchronous mirroring. For more information, see the *SnapDrive Installation and Administration Guide*.

  Note

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

What is required to create a LUN

Whichever method you choose, you create a LUN by specifying the following attributes:

The path name of the LUN: The path name must be at the root level of a qtree or a volume that specifies where the LUN is located. Do not create LUNs at /vol/vol0.

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

Note

You might find it useful to provide a meaningful path name for the LUN. For example, you might choose a name that describes how the LUN is used, such as the name of the application or the user using the data or the kind of data that it is.

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

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

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

The size of the LUN: When you create a LUN, you specify its size as raw disk space, depending on the filer and the host. You specify the size, in bytes (default), or by using the following multiplier suffixes.

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

Multiplier suffix	Size
t	terabytes

The usable space in the LUN is dependent on host or application overhead requirements. For example, partition tables and host file system metadata reduce usable space for applications. For information about the minimum and maximum sizes for LUNs, see the iSCSI host initiator documentation for your host.

The disk geometry used by the operating system determines the minimum and maximum values of LUNs. For example, Windows geometry is characterized by 63 sectors/track and 255 heads, which dictates a minimum LUN size of 31.5MB. For information about the maximum sizes for LUNs, see the iSCSI host initiator Support Kit documentation or the iSCSI Host Bus Adapter Attach Kit documentation for your host.

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

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

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

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

About igroups

Initiator groups are tables of iSCSI node names of hosts that you use to control access to LUNs. Typically, you want one host to have access to a LUN. However, if you are using multipathing software, or have clustered hosts, you need to have redundant paths to the same LUN.

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

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator.

What is required to create an igroup

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

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

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

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or DNS. If you name an igroup windows-1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

Note

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

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

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

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

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

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

Example of creating an igroup

Assume you have a Windows host whose name is Host1, and you installed the iSCSI initiator software on that host. You obtain the Windows' node name by going to the Start menu, choosing Settings > Control Panel > iSCSI Initiator, clicking on the Initiator Settings tab, and recording the node name from the "Change initiator node name" field.

You use the `igroup create` command or FilerView to create an igroup, giving it a name like *windows-group1* and adding the node name of the initiator, Host1, to it.

What is required to map a LUN to an igroup

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

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

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

LUN ID: You assign a number for the LUN ID or accept the default LUN ID. Typically, the default LUN ID begins with 0 and increments by 1 for each additional LUN as it is created. The range of valid LUN ID numbers are host-dependent. For example, the range for Windows is 0-254

For information about the ranges of values to use when assigning LUN IDs for your host, see the appropriate iSCSI host initiator Support Kit documentation for your host.

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

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

Guidelines for mapping LUNs

Use the following guidelines when mapping LUNs:

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

Creating LUNs with the lun setup program

What the lun setup program does

The `lun setup` program prompts you for information needed for creating a LUN and an igroup, and for mapping the LUN to the igroup. When a default is provided in brackets in the prompt, you can press Enter to accept it.

Prerequisites for running the lun setup program

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

Running the lun setup program

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

Step	Action

1	On the filer 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) [image]:` **Example:** `windows` For information about specifying the ostype of the LUN, see "The host operating system type" on page 29.

Step	Action

3	Specify where the LUN will be located by responding to the next prompt:
	`A LUN path must be absolute. A LUN can only reside in a volume or qtree root. For example, to create a LUN with the name “lun0” in the qtree root /vol/vol/q0, specify the path as “/vol/vol1/q0/lun0”.`
	`Enter LUN path:`
	Example: f you previously created /vol/vol1/ and want to create a LUN called lun0, you would enter, `/vol/vol1/lun0`.
	Caution
	Do not create LUNs in /vol/vol0. This is the root volume used for filer administration.
	Result: *lun0* is created in the root of /vol/vol1 if you accept the configuration information later in this program.

4	Specify whether you want the LUN created with space reservations enabled by responding to the prompt:
	`A LUN can be created with or without space reservations being enabled. Space reservation guarantees that data writes to that LUN will never fail.`
	`Do you want the LUN to be space reserved? [y]:`
	Example: Enter
	Caution
	If you choose n, space reservation is disabled. This might cause writes to the filer to fail, which can cause data corruption. NetApp strongly recommends that you enable space reservations.

5	Specify the size of the LUN by responding to the next prompt:
	`Size for a LUN is specified in bytes. You can use single-character multiplier suffixes: b(sectors), k(KB), m(MB), g(GB) or t(TB).`
	`Enter LUN size:`
	Example: `5g`
	Result: A LUN with 5 GB of raw disk space is created if you accept the configuration information later in this program. The amount of usable disk space by the host varies, depending on the operating system type and the application using the LUN.

Step	Action

6	Create a comment or a brief description about the LUN by responding to the next prompt:
	
	`You can add a comment string to describe the contents of the LUN. Please type a string (without quotes), or hit ENTER if you don't want to supply a comment.`
	`Enter comment string:`
	
	Example: 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 a description by using FilerView.

7	Create or use an existing igroup by responding to the next prompt:
	
	`The LUN will be accessible to an initiator group. You can use an existing group name, or supply a new name to create a new initiator group. Enter '?' to see existing initiator group names.`
	`Name of initiator group[]:`
	
	Result: If you have already created one or more igroups, you can enter ? to list them. The last igroup you used appears as the default. If you press Enter, that igroup will be used.
	
	If you have not created any igroups, enter a name of the igroup you want to create now. For information about naming an igroup, see "The name of the igroup" on page 31.

8	Specify which protocol will be used by the initiator in the igroup by responding to the next prompt: `Type of initiator group windows-igroup2 (FCP/iSCSI)[FCP]: iscsi` **Result:** The initiators in this igroup use the iSCSI protocol.

9	Add the iSCSI node names of the hosts that will be in the igroup by responding to the next prompt:
	`An iSCSI initiator group is a collection of initiator node names. Each node name can begin with either 'eui.' or 'iqn.' and should be in the following formats: eui.{EUI-64 address} or iqn.yyyy-mm.{reserved domain name}:{any string}.`
	`Eg. iqn.2001-04.com.acme:storage.tape.sys1.xyz or eui.02004567A25678D`
	`You can separate node names by commas. Enter '?' to display a list of connected initiators. Hit ENTER when you are done adding port names to this group.`
	`Enter comma separated nodenames:`

Step	Action
	Result: If you want to list the iSCSI node names, enter ? to list them. When you enter a node name, it is added to the igroup you specified in Step 7. **Note** To determine which node name is associated with a host, use the vendor-supplied commands or methods you downloaded from the iSCSI host initiator support kit on NetApp's NOW web site, or from the CD-ROM in the iSCSI Host Attach Kit for Windows.
10	Specify the operating system type initiators in the igroup use to access LUNs by responding to the next prompt: `The initiator group has an associated OS type. The following are currently supported: solaris, windows, hpux, aix, linux, or default OS type of initiator group "windows-igroup1" [windows]:` If the default is the correct os type, press Enter; otherwise, type the appropriate os type. In this example, it is `windows`. For information about specifying the ostype of an igroup, see "About igroups" on page 30.

Step	Action

11	Specify the LUN ID that the host will map to the LUN by responding to the next prompt:
	
	`The LUN will be accessible to all the initiators in the initiator group. Enter ‘?’ to display LUNs already in use by one or more initiators in group “windows-igroup1”.`
	`LUN ID at which initiator group “windows-igroup1” sees “/vol/vol1/lun0” [0]:`
	
	Result: If you press Enter to accept the default, Data ONTAP issues the lowest valid unallocated LUN ID to map it to the initiator, starting with zero. Alternatively, you can enter any valid number. For information about valid LUN IDs for your host initiator, see the documentation provided with your iSCSI host Initiator Support Kit or with your SAN Host Attach Kit for iSCSI Protocol on your host.
	
	Note
	Network Appliance recommends that you accept the default value for the LUN ID.
	
	After you press Enter, the lun setup program displays the information you entered:
	`LUN Path : /vol/vol1/q0/lun0`
	`OS Type : windows`
	`Size : 5g (5368709120)`
	`Comment : 5 GB Windows LUN for finance records`
	`Initiator Group : windows-igroup1`
	`Initiator Group Type : ISCSI`
	`Initiator Group Members : iqn.1991-05.com.microsoft:lab1.hq.XYZcompany.com`
	`Mapped to LUN-ID : 0`

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

Creating LUNs and igroups with FilerView

Methods of creating LUNs

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

- LUN Wizard
- Menu
  - Create LUN
  - Create igroup
  - Map LUN

Creating LUNs with the LUN Wizard

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

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

Step	Action
4	Click the Next button to continue. **Result:** The first window of fields in the LUN Wizard appears. 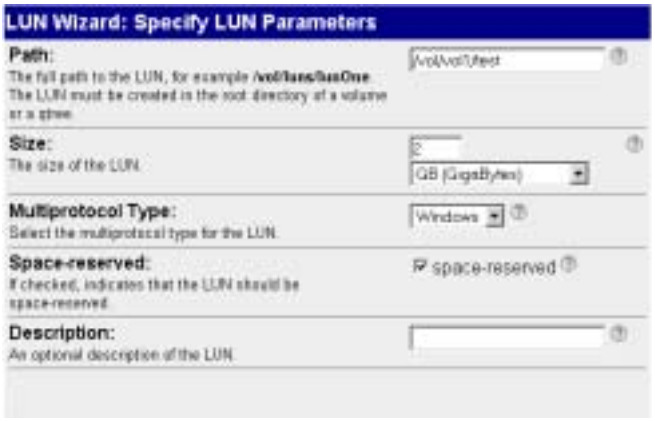
5	Follow the instructions in the LUN Wizard and enter information in the appropriate fields.
6	At the LUN Wizard Commit Changes window, review your input. If everything is correct, click Commit. **Result:** The LUN Wizard: Success! window appears and the LUN you created is mapped to the igroups you specified.

Creating LUNs and igroups with FilerView menus

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

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

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

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

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

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

Creating LUNs and igroups using individual commands

When 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 the initiators that are available to create igroups	`iscsi show initiator`
	Sample result:
	`Initiators connected on adapter iswta:`
	`  Tgt_PG  iSCSI Initiator Name / ISID`
	`  2  iqn.1986-03.com.intel:lab.hba-1 / 40:00:02:da:00:02`

Determine which hosts are associated with the node names	To determine the node name that is associated with a host, use the command or method provided by the vendor-supplied software that came with the SAN HBA kit, such as the Intel PRO/1000 T IP Storage Adapter for Windows hosts or the vendor-supplied iSCSI host initiator software, such as the iSCSI Microsoft Windows Initiator Support Kit.

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

Create an igroup	`igroup create -i -t` *`ostype initiator_group`* `[`*`node`*`]`
	
	`-i` indicates that the igroup contains iSCSI node names.
	
	`-t` *ostype* indicates the operating system of the initiator. The values are: `default`, `solaris`, `windows`, `hpux`, `aix`, or `linux`.
	
	For more information about specifying the ostype of an igroup, see "The name of the igroup" on page 31.
	
	initiator_group is the name you specify as the name of the igroup.
	
	node is an iSCSI nodename
	
	Example:
	
	`igroup create -i -t windows windows-igroup3 iqn.1991-05.com.microsoft:lab1.hq.XYZcompany.com`

Create a 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 29. `-t` *ostype* indicates the operating system type that determines the geometry used to store data on the LUN. The values are `image` (default), `solaris`, `windows`, `hpux`, `aix`, and `linux`. For information about specifying the ostype of the LUN, see "The host operating system type" on page 29. For information about specifying the LUN ostype, see "The host operating system type" on page 29. *lun_path* is the LUN's path name that includes the volume and qtree. **Example:** `lun create -s 4g -t windows /vol/vol1/qtree1/lun3` **Result:** A 4-GB LUN called /vol/vol1/qtree1/lun3 is accessible by a Windows host. In general, when you format and partition LUNs as a disk on a host, the size of the disk is smaller due to the overhead required by the host.

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

Map the LUN to an igroup	`lun map` *`lun_path igroup-name`* `[`*`LUN-ID`*`]`
	
	lun_path is the path name of the LUN you created.
	
	igroup-name is the name of the igroup you created.
	
	LUN-ID is the identification number that you specify the initiator should use when the LUN is mapped to it. The range of the number depends on the operating system of the initiator. For information about valid LUN ID ranges for your host initiator, see the documentation provided with your iSCSI host Initiator Support Kit or with your SAN Host Attach Kit for iSCSI Protocol on your host.
	
	Example 1: `lun map /vol/vol1/qtree1/lun3 windows-igroup3 0`
	
	Result: Data ONTAP maps /vol/vol1/qtree1/lun3 to the igroup windows-igroup3 at LUN ID 0.
	
	Example 2: `lun map /vol/vol1/lun0 windows-igroup3`
	
	Result: Data ONTAP assigns the next lowest valid LUN ID to map the LUN to the igroup.
	
	After entering the command in this example, Data ONTAP displays the following message:
	
	`lun map: auto-assigned windows-igroup3=0`

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: windows`
	`  Maps: windows-igroup3=0`

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

Display the LUN ID mapping	`lun show -m`
	Sample result:
	`LUN path  Mapped to  LUN ID`
	`---`
	`/vol/vol1/lun0  windows-igroup0  0`
	`/vol/vol1/lun1  windows-igroup1  0`
	`/vol/vol1/qtree1/lun2  windows-igroup2  0`
	`/vol/vol1/qtree1/lun3  windows-igroup3  0`
	In this example, each LUN is mapped to a different igroup, which was named to reflect the host whose node name was used when creating the igroup. Note that the LUN ID mapping begins with 0 for each igroup. If a second LUN is mapped to the same igroup, the LUN ID automatically increments by one.

Creating LUNs on vFilers

Prerequisite for creating LUNs on vFilers

You must purchase a MultiStore licence to create vFilers. 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 filer, not to vFilers.

Guidelines for creating LUNs on vFilers

Use the following guidelines when creating LUNs on vFilers.

- The vFiler access rights are enforced when the filer 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.

Setting up vFilers with the iSNS service

To set up vFilers with the iSNS service, complete the following steps for each vFiler.

Step	Action
1	Set the IP address of the iSNS server.
2	Enter the following command to enable the service: **`iscsi isns start`**

LUN subcommands available on vFilers

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` for vFiler LUNs.

Methods for issuing LUN subcommands on vFilers

You can issue LUN subcommands using the following methods:

- From the default vFiler (vfiler0) on the hosting filer, you can:
  - Enter the `vfiler run * lun` *`subcommand`*
  - Change the vfiler context by entering the following commands:

    filer> `vfiler context` *`vfiler_name`*

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

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)`

How to manage vFiler LUNs

You use the available LUN commands by preceding them with the `vfiler run` *`vfiler_name`* `*` command.

Managing LUNs 3

About this chapter

This chapter describes how to manage LUN maps, change LUN attributes, and display LUN statistics.

Topics in this chapter

This chapter discusses the following topics

- "Managing LUNs and LUN maps" on page 50
- "Displaying LUN information" on page 56
- "Monitoring disk space" on page 59

Managing LUNs and LUN maps

Tasks to manage LUNs and LUN maps

You can use the command-line interface or FilerView to

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

Controlling block protocol access

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

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

Step	Action
1	Enter the following command: **`lun online 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; however, you can write to an offline LUN using NAS protocols. To take a LUN offline, complete the following step.

Step	Action
1	Enter the following command: `lun offline` *`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** The existing lun_path and the 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, resizing is only supported on basic disks. Resizing is not supported on dynamic disks.
- On Solaris, decreasing the LUN size is not supported. Only increasing the LUN size is supported.

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

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

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

Caution

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

Step	Action
1	Take a LUN offline before resizing it by entering the following command: **`lun offline`** ***`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 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. **Example:** For Solaris hosts, use Solaris commands to erase the disk label and create a new one, and then use the Solaris `mkfs` command to increase the file system on the disk slice to fill the newly enlarged slice.

Modifying the LUN description

To modify the LUN description, complete the following step.

Step	Action
1	Enter the following command: **`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.

Modifying a LUN’s serial number

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

Step	Action
1	Take the LUN offline by entering the following command: **`lun offline`** ***`lun-path`*** **Example:** `lun offline /vol/vol1/qtree1/lun2`
2	Modify the LUN’s serial number by entering the following command: **`lun serial`** ***`lun-path new-serial-number`*** **Example:** `lun serial /vol/vol1/qtree1/lun2 OcmePIdQxjcH`

Enabling or disabling space reservations for LUNs

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

Caution

If you disable space reservations, you might experience write failures due to insufficient disk space. When this occurs, Data ONTAP displays system messages (one message per file) on the console, or sends these messages to log files and other remote systems, as specified by its /etc/syslog.conf configuration file.

Step	Action
1	Enter the following command: **`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 only be accessed with the iSCSI protocol by default. However, you can make a LUN available to a host using NAS protocols if NAS protocols are licensed and enabled on the filer. The usefulness of accessing a LUN over NAS protocols is host application-dependent.

Note

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

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

Step	Action
1	Determine whether you want to read or write to the LUN over the NAS protocol and take appropriate action: ◆ If you want read access, the LUN can remain online. ◆ If you want write access, ensure that the LUN is offline or unmapped.
2	Enter the following command: **`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
- Read, write, and operations per second statistics
- Statistics and characteristics
- Space reservations setting

Displaying command-line help

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

Step	Action

1	On the filer's command line, enter the following command:
	`lun help`
	Result: A list of all LUN subcommands is displayed:
	`lun help - List LUN (logical unit of block storage) commands`
	`lun clone - Manage LUN cloning`
	`lun comment - Display/Change descriptive comment string`
	`lun create - Create a LUN`
	`lun destroy - Destroy a LUN`
	`lun map - Map a LUN to an initiator group`
	`lun move - Move (rename) LUN`
	`lun offline - Stop block protocol access to LUN`
	`lun online - Restart block protocol access to LUN`
	`lun resize - Resize LUN`
	`lun serial - Display/change LUN serial number`
	`lun set - Manage LUN properties`
	`lun setup - Initialize/Configure LUNs, mapping`
	`lun share - Configure NAS file-sharing properties`
	`lun show - Display LUNs`
	`lun 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 reads, data writes, and operations per second for LUNs, complete the following step.

Step	Action

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

Displaying status and characteristics

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

Step	Action

1	On the filer's command line, enter the following command to display LUN status and characteristics:
	`lun show -v`
	Sample display:
	`/vol/vol1/lun0 5g (5368709120)  (r/w, online, mapped)`
	`Comment: "First LUN"`
	`Serial#: OcmePIdQwRb1`
	`Share: none`
	`Space Reservations: enabled`
	`Multiprotocol Type: windows`
	`Maps: windows-igroup0=0`
	`/vol/vol1/lun1 100m (104857600)  (r/w, online, mapped)`
	`Comment: "Second LUN"`
	`Serial#: OcmePIdQuEtX`
	`Share: none`
	`Space Reservations: enabled`
	`Multiprotocol Type: windows`
	`Maps: windows-igroup1=0`

2	On the filer's command line, enter the following command to display LUN mapping information:
	`lun show -m`
	Result:
	`LUN path  Mapped to  LUN ID`
	`---`
	`/vol/vol1/lun0  windows-igroup0  0`
	`/vol/vol1/lun1  windows-igroup1  0`
	`/vol/vol1/qtree1/lun2  windows-igroup2  0`
	`/vol/vol1/qtree1/lun3  swindows-igroup2  1`

Displaying space reservation status for LUNs

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

Step	Action
1	Enter the following command: **`lun set reservation lun-path`**

Monitoring disk space

Command to use

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

What the command output displays

The output of the `df` command displays four columns of statistics about the active file system in the volume and the snapshot directory for that volume. The statistics displayed are as follows:

- Amount of total space on the volume, in the size of bytes you specific

  Total space = used space + available space
- Amount of used space

  Used space = space storing data + space storing snapshots + space storing reserved space
- Amount of available space

  Available space = space that is not used or reserved, it is free space
- Percentage of the volume capacity being used

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

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

The effect of LUN and snapshot operations on disk space

The following table illustrates the effect on disk space when you create a volume, create a LUN, write data to the LUN, take snapshots of the LUN, and expand the size of the volume.

For this example, assume that space reservation is enabled and snap reserve is set to 0%.

Action	Result	Comment
Create a 100-GB volume.	Used space = 0 GB Reserved space = 0 GB Available space = 100 GB Volume Total: 100 GB Snapshot creation is allowed.	N/A
Create a 40-GB LUN on that volume.	Used space = 40 GB Reserved space = 0 GB Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	Used space is 40 GB for the LUN. If the LUN size was limited to accommodate at least one snapshot when it was created, it will always be less than one half of the volume size.
Write 40 GB of data to the LUN.	Used space = 40 GB Reserved space = 0 GB Available space = 60 GB Volume Total: 100 GB Snapshot creation is allowed.	The amount of used space does not change because with space reservations set to On, the same amount of space is used when you write to the LUN as when you created the LUN.
Create a snap-shot of the LUN.	Used space = 80 GB Reserved space = 40 GB Available space = 20 GB Volume Total: 100 GB Snapshot succeeds.	The snapshot locks all the data on the LUN so that even if that data is later deleted, it remains in the snapshot until the snapshot is deleted. As soon as a snapshot is created, the reserved space must be large enough to ensure that any future writes to the LUN succeed. Reserved space is now 40 GB, the same size of the LUN. Data ONTAP always displays the amount of reserved space required for successful writes to LUNs. Because reserved space is also counted as used space, used space is 80 GB.

Action	Result	Comment
Overwrite all 40 GB of data on the LUN with new data.	Used space = 100 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 100 GB Snapshot creation is blocked.	Data ONTAP manages the space so that the overwrite increases used space to 100 GB and decreases available space to 0. The 40 GB for reserved space is still displayed. You cannot take another snapshot because no space is available. That is, all space is used by data or held in reserve so that any and all changes to the content of the LUN can be written to the volume.
Expand the volume by 100 GB.	Used space = 120 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	After you expand the volume, the amount of used space displays the amount needed for the 40 GB LUN, the 40 GB snapshot, and 40 GB of reserved space. Free space becomes available again, so snapshot creation is no longer blocked.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 120 GB Reserved space = 40 GB Available space = 80 GB Volume Total: 200 GB Snapshot creation is allowed.	Because none of the overwritten data belongs to a snapshot, it disappears when the new data replaces it. As a result, the used total remains unchanged.
Create a snapshot of the LUN.	Used space = 160 GB Reserved space = 40 GB Available space = 40 GB Volume Total: 200 GB Snapshot creation is allowed.	The snapshot locks all 40 GB of data currently on the LUN. The used space is the sum of 40 GB for the LUN, 40 GB for each snapshot, and 40 GB for reserved space.
Overwrite all 40 GB of data on the LUN with new data.	Used space = 160 GB Reserved space = 40 GB Available space = 40 GB Volume Total: 200 GB Snapshot creation is allowed.	Because the data being replaced belongs to a snapshot, it remains on the volume.

Action	Result	Comment
Expand the LUN by 40 GB.	Used space = 200 GB Reserved space = 40 GB Available space = 0 GB Volume Total: 200 GB Snapshot creation is blocked.	The amount of used space increases by the amount of LUN expansion. The amount of reserved space remains at 40 GB. Because the available space has decreased to 0, snapshot creation is blocked.
Delete both snapshots of the volume.	Used space = 80 GB Reserved space = 0 GB Available space = 120 GB Volume Total: 200 GB Snapshot creation is allowed.	The 80 GB of data locked by the two snapshots disappears from the used total when the snapshots are deleted. Because there are no more snapshots of this LUN, the reserved space decreases to 0 GB. Snapshot creation is once again allowed.
Delete the LUN.	Used space = 0 GB Reserved space = 0 GB Available space = 200 GB Volume Total: 200 GB	Because no snapshots exist for this volume, deletion of the LUN causes the used space to decrease to 0 GB.

Example of monitoring disk space

The following example illustrates how to use the `df -r` command to monitor disk space on a volume when you perform the following tasks:

- Create a volume
- Create a LUN
- Write data to the LUN
- Take a snapshot
- Create a backing store LUN
- Create a clone of the backing store LUN

For simplicity, assume that the amount of disk space required for the original LUN is 3 GB. The information from "Guidelines for calculating the size of a volume" on page 23 recommends that the volume size be approximately four times 3 GB plus the amount of change expected between snapshots. Assuming this rate is approximately 3 GB plus 10%. If you plan to use 72 GB disks (which typically results in 67.9 GB of physical capacity, depending on the manufacturer), two disks provide more than enough space, one for data and one for parity.

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

In the example, the filer is named *toaster*, and the node name for the Windows host is iqn.1991-05.com.microsoft:eng1.

To work through the example, complete the following steps.

Step	Action

1	From the filer, 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% on the volume even though the volume will be used for iSCSI 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 spa`
	`ce/.snapshot`

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 space becomes zero, and the 20% of snapshot 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/vol spa`
	`ce/.snapshot`

Step	Action

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 windows /vol/volspace/lun0`**
	
	toaster> **`df -r /vol/volspace`**
	
	Result: The following sample output is displayed. Approximately 3GB of space is taken from available space and is displayed as reserved space for the LUN:
	
	`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 *win-igroup1* and map the LUN to the igroup by entering the following commands (assuming the node name is iqn.1991-05.com.microsoft:eng1). These commands have no effect on disk space. toaster> **`igroup create -i -t windows win-igroup1 iqn.1991-05.com.microsoft:eng1`** toaster> **`lun map /vol/volspace/lun0 win-igroup1 0`**
5	From the host, discover the LUNs, format them, and make the file system available to the host. For information about these procedures for your host initiator, see the documentation provided with your iSCSI host Initiator Support Kit or with your SAN Host Attach Kit for iSCSI Protocol on your host.

6	From the filer, see that there is a minimum effect on used and available space after creating the file system on the LUN by entering the following command.
	
	`toaster> df -r /vol/volspace`
	
	Result: The following sample output is displayed:
	
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 3150324 59499584 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volspace/`
	`.snapshot`

7	From the host, write data to the file system (the LUN on the filer).

Step	Action

8	From the filer, see the effect of writing to the LUN by entering the following command:
	`toaster> df -r /vol/volspace`
	Result: The following sample output is displayed:
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 3150324 59499584 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volspace/.snapshot`

9	From the host, continue writing data to the file system until the LUN is full.
	Again, there is no effect.
	Result: The following sample output is displayed:
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 3150324 59499584 0 /vol/volspace/`
	`/vol/volspace/.snapshot 0 0 0 0 /vol/volspace/.snapshot`

10	Take a snapshot named *snap1* of the active file system and observe the effect on disk space.
	Caution
	From the host or the application, ensure that the active file system is in a quiesced or synchronized state prior to taking a snapshot.
	Enter the following commands:
	toaster> **`snap create volspace snap1`**
	toaster> **`df -r /vol/volspace`**
	Result: The following sample output is displayed. The amount of reserved space increases to approximately 3 GB to ensure that in addition to being able to write data to the LUN, data in the snapshot can be overwritten as well because when Data ONTAP writes to a block of data in a snapshot, both the old data and the new data must be stored. In this example, 722 KB of space is used for the snapshot.
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 6036168 56613740 2884368 /vol/volspace/`
	`/vol/volspace/.snapshot 0 180 0 0 /vol/volspace/.snapshot`

Step	Action

11	Create a backing store LUN0 named *snaplun* and observe the effect on disk space by entering the following commands:
	
	toaster> **`lun create -b /vol/volspace/.snapshot/snap1/lun0 /vol/volspace/snaplun`**
	
	Note
	After you create this LUN, you must go through the same procedure on the host to discover the new LUN, format it, and so on, except for creating a new file system on it.
	
	toaster> **`df -r /vol/volspace`**
	
	Result: The following sample output is displayed. The amount of reserved space increases to approximately 6 GB to ensure that both LUNs can be written to, and a small amount of space is increased for the snapshot.
	
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 9185024 53464884 2884392 /vol/volspace/`
	`/vol/volspace/.snapshot 0 150 0 0 /vol/volspace/`
	`.snapshot`

12	Display the status of both LUNs by entering the following command:
	
	toaster> **`lun show -v`**
	
	Result: The following sample output is displayed:
	
	`/vol/volspace/lun0 3g (3221225472) (r/w, online, mapped)`
	`Serial#: 0dDICn1A-Z0J`
	`Share: none`
	`Space Reservation: enabled`
	`Multiprotocol Type: windows`
	`Maps: win-igroup1=0`
	`/vol/volsapce/snaplun 3g (3221225472) (r/w, online)`
	`Serial#: 0dDICn1RRA0J`
	`Backed by: /vol/volspace/.snapshot/snap1/lun0`
	`Share: none`
	`Space Reservation: enabled`
	`Multiprotocol Type: windows`
	`Maps: win-igroup1=0`

Step	Action

13	Create a clone of the backing store LUN you created in step 11 and observe the effect on disk space by entering the following commands:
	
	toaster> **`lun clone start /vol/volspace/snaplun`**
	
	toaster> **`df -r /vol/volspace`**
	
	Result: The following sample output is displayed. The amount of used space increases to 45 GB, available space decreases, and again, a small amount of space is used by snapshot.
	
	`Filesystem kbytes used avail reserved Mounted on`
	`/vol/volspace/ 62649908 9290128 53359780 2989448 /vol/volspace/`
	`/vol/volspace/.snapshot 0 196 0 0 /vol/volspace/.s`
	`napshot`

Managing igroups

4

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 70
- "Using igroups on vFilers" on page 73

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 [node_name ...]`** `-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. *node_name* is an iSCSI node name. **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 node name of the Windows host associated with that node name.

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.

Step	Action
1	Enter the following command: **`igroup add`** ***`igroup iscsi_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 node_name`*** **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, or linux), 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 vFilers

How igroups work on vFilers

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 vFilers.

Creating an igroup in a non-default vFiler

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

Step	Action
1	Change the context to the desired vFiler by entering the following command: `myfiler*> vfiler context vf1` **Result:** The vFiler's prompt is displayed.
2	Create the igroup on vFiler determined by step 1 by entering the following command: `vf1@myfiler*> igroup create -i vf1_iscsi_group iqn.2002-06.xom.microsoft.tin1`

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.2002-06.xom.microsoft.tin1`

Mapping LUNs to igroups

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

Using Data Protection with iSCSI 5

About this chapter

This chapter provides information about how to use Data ONTAP data protection features using the iSCSI protocol in a TCP/IP Network.

Topics in this chapter

This chapter discusses the following topics:

- "Data ONTAP protection methods" on page 76
- "Using snapshots" on page 78
- "Using LUN clones" on page 80
- "Deleting busy snapshots" on page 82
- "Using SnapRestore" on page 88
- "Backing up data to tape" on page 94
- "Using NDMP" on page 98
- "Using volume copy" on page 99
- "Using NVFAIL" on page 100

Data ONTAP protection methods

Data protection methods

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

Snapshots are a standard feature of Data ONTAP. A snapshot is a frozen, read-only image of the entire Data ONTAP file system (or WAFL® volume) that reflects the state of the LUN or the file system at the time the snapshot was created. All the other data protection methods rely on snapshots or create, use, and destroy snapshots, as required.

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

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

Method	Used to...
SnapDrive™	◆ Manage a filer's LUNs that serve as virtual storage devices for application data in Windows 2000 Server and Windows 2003 Server environments in an integrated environment with the Windows Volume Manager. ◆ Perform online storage configuration, LUN expansion, and streamlined management. **Note** For more information about SnapDrive, see the *SnapDrive 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.
NDMP	◆ Control native backup and recovery facilities in NetApp filers and other file servers. Backup application vendors provide a common interface between backup applications and file servers. **Note** The Network Data Management Protocol (NDMP) is an open standard for centralized control of enterprise-wide data management. For more information about how NDMP-based topologies can be used by filers to protect data, see the *Data Protection Solutions Overview*, Technical Report TR3131 at http://www.netapp.com/tech_library/3131.html.

Using snapshots

How Data ONTAP snapshots work in an FC network

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

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

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

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

How snapshots are used in the SAN environment

Data ONTAP cannot ensure that the data within a LUN is in a consistent state with regard to the application accessing the data inside the LUN. Therefore, prior to creating a snapshot, you must quiesce the application or file system using the LUN (which flushes the host filer system buffers to disk) so that the snapshot is consistent. For example, you can use batch files and scripts on a host that has administrative access to the filer. You use these scripts to perform the following tasks:

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

Note

On Windows hosts, you can use the Windows Task Scheduler service to execute this script at specified intervals.
In addition, you can use SnapDrive 2.0 or later to save the contents of the host file system buffers to disk and create snapshots. See the *SnapDrive Installation and Administration Guide*.

The relationship between a LUN and a snapshot

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

What you can do with a snapshot

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

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

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

What snapshots require

In Data ONTAP 6.5, LUNs are space reserved by default. This means that enough space is reserved so that writes to the LUNs are guaranteed. The more space that is reserved, the less free space is available. If free space within the volume is below a certain threshold, snapshots cannot be taken. For information about how to manage available space, see “Monitoring disk space” on page 59.

Using LUN clones

What a LUN clone is

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

What LUN clones are used for

You can use LUN clones to create a permanent copy of the LUN at the time the snapshot of the original LUN was taken. You can use it when you need to delete a snapshot of a LUN without taking the LUN offline or losing its contents. You can clone the contents of the writable snapshot of the LUN without any disruption to an application using the LUN. After the clone operation is complete, you can delete the snapshot.

Note

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

Creating a writable snapshot of a LUN

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

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

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

Starting the cloning process

To clone a LUN, complete the following step.

Step	Action
1	Begin the clone operation by entering the following command: **`lun clone start lun_clone_path`**

Displaying the progress of the clone operation

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

Step	Action
1	Enter the following command: **`lun clone show lun_clone_path`**

Stopping the cloning process

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

Step	Action
1	Enter the following the command: **`lun clone stop lunpath`**

Deleting busy snapshots

What a snapshot in a busy state means

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

Process of a snapshot becoming busy

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

Stage	Busy snapshot scenario
1	Create a LUN. **Example:** `lun create -s 10m /vol/vol2/lunA` **Result:** lunA exists in the active file system.

2	Display LUN status to see lunA.
	
	Example: `lun show -v`
	
	Result: The following message is displayed:
	`/vol/vol2/lunA  10m (10485760)  (r/w, online)`
	`  Serial#: BnBt/YiART3H`
	`  Share: none`
	`  Space Reservation: enabled`
	`  Multiprotocol Type: image`

3	Take a snapshot of the volume containing lunA. **Example:** `snap create vol2 snap0` **Result:** snap0 of vol2 is created, but it is not busy.

Stage	Busy snapshot scenario

4	Display snapshot status to see that snap0 is not busy.
	
	Example: `snap list vol2`
	
	Result: The following message is displayed:
	`Volume vol2`
	`working ...`
	`%/used  %/total  date  name`
	`---------  ---------  ---------  -----------`
	` 0% ( 0%)  0%  ( 0%)  Mar 13 02:47  snap0`

5	Create another LUN, backed by lunA in snap0. **Example:** `lun create -b /vol/vol2/.snapshot/snap0/lunA /vol/vol2/lunB` **Result:** lunB is a writable snapshot of lunA (a backing store LUN). snap0 is now marked "(busy, LUNs)."

6	Display snapshot status to see that snap0 is marked busy.
	
	Example: `snap list vol2`
	
	Result: The following message is displayed:
	`Volume vol2`
	`working ...`
	`%/used  %/total  date  name`
	`---------  ---------  ---------  -----------`
	`19% (19%)  0%  ( 0%)  Mar 13 02:47  snap0  (busy,LUNs)`

7	Take another snapshot of vol2. **Example:** `snap create vol2 snap1` **Result:** snap1 now exists, but it is not busy. snap0 is still busy.

Stage	Busy snapshot scenario

8	Display snapshot status to see that snap1 is not busy and snap0 is still busy.
	
	Example: `snap list vol2`
	
	Result: The following message is displayed:
	`Volume vol2`
	`working ...`
	`%/used  %/total  date  name`
	`---  ---  ---  ---`
	`0% ( 0%)  0%  ( 0%)  Mar 13 02:50  snap1`
	`19% (19%)  0%  ( 0%)  Mar 13 02:47  snap0  (busy,LUNs)`

9	Delete the LUN that was backed by lunA in snap0, lunB. **Example:** `lun destroy /vol/vol2/lunB` **Result:** lunB is deleted, but snap0 is still busy because snap1 contains a copy of lunB. If you attempt to delete snap0, the following message is displayed: `vol2: Snapshot in use by dump, snap restore, SnapMirror, a CIFS share, RAID mirroring, LUNs or volume clone.`

10	Identify the LUNs in existing snapshots that are backed by data in the specified snapshot.
	
	Example: `lun snap usage vol2 snap0`
	
	Result: The following message is displayed:
	`snap1:`
	`LUN: /vol/vol2/.snapshot/snap1/lunB`
	`Backed By: /vol/vol2/.snapshot/snap0/lunA`

11	Delete snap1. **Example:** `snap delete vol2 snap1` **Result:** snap1 is deleted, and snap0 is no longer marked busy.

Stage	Busy snapshot scenario

12	Display snapshot status to see that snap0 is no longer busy.
	
	Example: `snap list vol2`
	
	Result: The following message is displayed:
	`Volume vol2`
	`working ...`
	`%/used  %/total  date  name`
	`---------  ---------  ---------  -----------`
	`19% (19%)  0%  ( 0%)  Mar 13 02:47  snap0`

13	Delete snap0. **Example:** `snap delete vol2 snap0` **Result:** snap0 is deleted.

14	Display snapshot status to see that there are no more snapshots.
	
	Example: `snap list vol2`
	
	Result: The following message is displayed:
	`Volume vol2`
	`working ...`
	
	`No snapshots exist.`

Command to use to find snapshots in a busy state

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

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

Deleting snapshots in a busy state

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

Step	Action

1	Identify all snapshots that are in a busy state, locked by LUNs, by entering the following command:
	snap list *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)`

2	Identify the LUNs and the snapshots that contain them by entering the following command:
	lun snap usage *vol_name snap_name*
	Example:
	`lun snap usage vol2 snap0`
	Result: The following message is displayed:
	`active:`
	`LUN: /vol/vol2/lunC`
	`Backed By: /vol/vol2/.snapshot/snap0/lunA`
	`snap2:`
	`LUN: /vol/vol2/.snapshot/snap2/lunB`
	`Backed By: /vol/vol2/.snapshot/snap0/lunA`
	`snap1:`
	`LUN: /vol/vol1/.snapshot/snap1/lunB`
	`Backed By: /vol/vol2/.snapshot/snap0/lunA`
	Note
	The preceding LUNs are backed by lunA in the snap0 snapshot.

Step	Action
3	Delete all the LUNs in the active file system that are displayed by the `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 snapshots 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 snapshots containing lunB are now deleted and snap0 is no longer busy.

5	Delete the snapshot 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 to revert an entire volume or a LUN to the state it was in when the snapshot was taken, within seconds, regardless of size or utilization. You can use SnapRestore to restore an entire volume or a LUN, or you can perform a single file SnapRestore on a LUN.

Requirements for using SnapRestore

Before using SnapRestore, you must perform the following tasks:

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

Caution

When restoring a single LUN, it must be taken offline or be unmapped prior to recovery if the LUN has not been destroyed.
Using SnapRestore on a LUN, or on volume that contains LUNs, without stopping all host access to those LUNs, can cause data corruption and system errors.

Restoring a snapshot of a LUN

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

Step	Action
1	From the host, stop all host access to the LUN.
2	From the host, if the LUN contains a host file system mounted on a host, unmount the LUN on that host.
3	From the filer, unmap the LUN by entering the following command: **`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 revert.
	volume_name is the name of the volume to be reverted. Enter the name only, not the complete path. You can enter only one volume name.
	`-s` *snapshot_name* specifies the name of the snapshot from which to revert the data. You can enter only one snapshot 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 revert the volume.

5	Press **y** to confirm that you want to revert the volume. **Result:** Data ONTAP displays the name of the volume and the name of the snapshot for the reversion. If you did not use the `-f` option, Data ONTAP prompts you to decide whether to proceed with the reversion.

Step	Action	
6	**If...**	**Then...**
	You want to continue with the reversion	Press **y**. **Result:** The filer reverts the volume from the selected snapshot. If you are reverting the root volume, the filer reboots the system.
	You do not want to proceed with the reversion	Press **n** or press **Ctrl-C**. **Result:** The volume is not reverted and you are returned to a filer prompt.
7	If the old maps aren't the ones you want, unmap them (if any old maps exist) by entering the following command: **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 filer, 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, you also need to restart any database applications you closed down and remount the volume from the host side.

Restoring an online LUN from tape

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

```
RESTORE: Inode XXX: file creation failed.
```

Reverting a file

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

Note

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

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

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

Example:

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

filer> WARNING! This will restore a file from a snapshot into the
active filesystem.  If the file already exists in the active
filesystem, it will be overwritten with the contents from the
snapshot.
Are you sure you want to do this? y

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

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

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

Backing up data to tape

Structure of SAN backups

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

Note

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

From the point of view of the SAN host, LUNs can be confined to a single WAFL volume or qtree or spread across multiple WAFL volumes, qtrees, or filers.

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

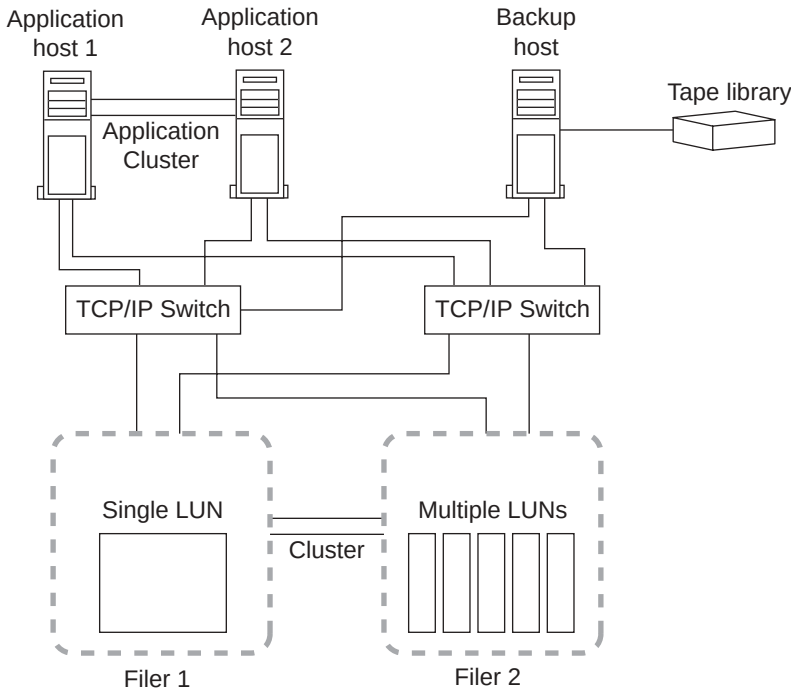

Volumes on the iSCSI host can consist of a single LUN mapped from the filer or multiple LUNs using a volume manager.

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

Backing up a single LUN to tape

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

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

Note

Steps 4, 5, and 6 can be part of your SAN backup application's pre-processing script.
Steps 9 and 10 can be part of your SAN backup application's post-processing script.

Step	Action
1	Enter the following command to create an igroup: **`igroup create -i [-t ostype] group [node ...]`** **Example:** `igroup create -i -t windows payroll_server iqn.1991-05.com.microsoft:eng1` **Result:** Data ONTAP creates an igroup with `-i`, the iSCSI option, the operating system Windows, the name payroll, and the node name `iqn.1991-05.com.microsoft:eng1`.
2	Enter the following command to create a 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.

Step	Action
3	Enter the following command to map the LUN to the igroup with a SCSI ID: **`lun map lun-path initiator-group scsi-lun`** **Example:** `lun map /vol/vol1/qtree_1/payroll_lun payroll_server 1` **Result:** Data ONTAP maps the /vol/vol_name/qtree_1/payroll_lun LUN to the payroll_server igroup with a LUN ID of 1.
4	Save the contents of host file system buffers to disk using the command provided by your host operating system, or by using SnapDrive 2.0 or later for Windows systems.
5	Create a snapshot by using SnapDrive 2.0 for Windows systems, or enter the following command for non-Windows systems: **`snap create volume_name snapshot_name`** **Example:** `snap create vol1 payroll_backup`
6	Enter the following command to create a new LUN in the active file system so that it can be mapped to the backup server from the existing snapshot: **`lun create -b snapshot lun-path lun-path`** **Example:** `lun create -b vol/vol1/.snapshot/payroll_backup/qtree_1/payroll_lun /vol/vol1/qtree_1/payroll_backup_lun`
7	Enter the following command to map the new LUN you created to the backup host: (*backup_server* is an igroup that already exists for the backup host.) **`lun map lun-path initiator-group scsi-lun`** **Example:** `lun map /vol/vol_name/qtree_1/payroll_lun payroll_server 1` **Result:** Data ONTAP maps the LUN */vol/vol_name/qtree_1/payroll_lun* to the igroup *payroll_server* with a SCSI ID of 1.

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

Using NDMP

When to use native or NDMP backup

Tape backup and recovery operations of LUNs should generally only be performed on the filer for disaster recovery scenarios, applications with transaction logging, or when combined with other filer-based protection elements, such as SnapMirror and SnapVault. For information about these features, see the *Data ONTAP 6.5 Data Protection Online Backup and Recovery* guide.

All tape operations local to the filer operate on the entire LUN and cannot interpret the data or file system within the LUN. Thus, you can only recover LUNs to a specific point-in-time unless transaction logs exist to roll forward. When finer granularity is required, use host-based backup and recovery methods.

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

When to use the ndmpcopy command

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

Using volume copy

Command to use

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

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

Copying a volume

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

Caution

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

Step	Action
1	Enter the following command: **`vol copy start -S source:`*`source_volume`*` dest:`*`dest_volume`*** **Note** If the copying takes place between two filers, you can enter the `vol copy start` command on either the source or destination filer. You cannot, however, enter the command on a third filer that does not contain the source or destination volume. **Example:** `vol copy start -S /vol/vol0 filerB:/vol/vol1`

Using NVFAIL

How NVFAIL works with LUNs

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

- Offlines the LUNs in the volumes that had the NVRAM failure
- Stops exporting LUNs over iSCSI
- 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 51

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

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

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

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

Enabling the NVFAIL option

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

Step	Action
1	Enter the following command: **`vol options`** ***`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 filer's /etc directory.
2	List the full path and file name, one file per line, within the nvfail_rename file. **Example:** `/vol/vol1/home/dbs/oracle-WG73.dbf`
3	Save the file.

Managing the iSCSI Network 6

About this chapter

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

Topics in this section

This section discusses the following topics

- "Registering the filer with an iSNS server" on page 104
- "Managing the iSCSI service" on page 107
- "Managing security for iSCSI initiators" on page 110
- "Displaying statistics for iSCSI sessions" on page 114
- "Displaying information for iSCSI sessions" on page 117
- "Managing the iSCSI service on filer interfaces" on page 119

Registering the filer 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.

How you manage the iSNS server

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

How the filer interacts with an iSNS server

The filer automatically registers its IP address, node name, and portal groups with the iSNS server when the iSCSI service is started and iSNS is enabled. In addition, Data ONTAP automatically updates the iSNS server any time the filer's configuration settings change.

Command to register the filer

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

Configuring the filer to register with an iSNS server

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

Step	Action
1	On the filer console, enter the following command to identify the iSNS server that the filer registers with: **`iscsi isns config -i`** ***`ip_addr`*** *ip_addr* is the IP address of the iSNS server.

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

Reregistering the filer with the iSNS server

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

To reregister the filer, complete the following step.

Step	Action
1	On the filer console, enter the following command: **`iscsi isns update`**

Disabling iSNS

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

To stop the iSNS service, complete the following step.

Step	Action
1	Enter the following command: **`iscsi isns stop`**

Configuring iSNS for vFilers

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

- Configure which iSNS server to use
- Configure which iSNS Entity ID to use

- Turn iSNS registration on or off

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

Managing the iSCSI service

Verifying that the iSCSI service is running

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

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

Verifying that iSCSI is licensed

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

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

Enabling the iSCSI license

To enable the iSCSI license, complete the following step.

Step	Action
1	Enter the following command: **`license add license_code`** *license_code* is the license code you received from NetApp when you purchased the iSCSI license.

Starting and stopping the iSCSI service

To start or stop the iSCSI service, complete the following step.

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

Displaying the filer's node name

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

Step	Action

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

Changing the filer's node name

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

Step	Action
1	Enter the following command: **`iscsi nodename iqn.1992-08.com.netapp:`*`unique_device_name`*** *unique_device_name* is the unique name for the filer. **Example:** `iscsi nodename iqn.1992-08.com.netapp:filer_hq`

Displaying initiators connected to the filer

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

Step	Action

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

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 filer to begin an iSCSI session. The filer permits or denies the login request according to one of the following authentication methods:

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

  There are two types of CHAP user names and passwords:

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

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

Guidelines for using CHAP authentication

The following guidelines apply to CHAP authentication:

- If you define an inbound user name and password on the filer, you must use the same user name and password for outbound CHAP settings on the initiator.

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

Defining an authentication method for an initiator

To define an authentication method for initiator that are 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 filer generates a 128-bit random password. **2.** Proceed to Step 2.
	◆ Use CHAP authentication and specify a password ◆ Use another security method	Proceed to Step 2.

Step	Action
2	For each initiator, enter the following command: **`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 filer without authentication. ◆ *deny*—The initiator cannot access the filer. *inpassword* is the inbound password for CHAP authentication. The filer uses the inbound password to authenticate the initiator. *inname* is a user name for inbound CHAP authentication. The filer uses the inbound user name to authenticate the initiator. *outpassword* is a password for outbound CHAP authentication. The filer uses this password for authentication by the initiator. *outname* is a user name for outbound CHAP authentication. The filer uses this user name for authentication by the initiator. **Note** If you generated a random password in Step 1, you can use this string for either *inpassword* or *outpassword*. If you enter a string, the filer interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

Displaying initiator authentication methods

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

Step	Action
1	Enter the following command: **`iscsi security show`**

Defining a default authentication method

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

Step	Action
1	Enter the following command: **`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 filer. ◆ *deny*—Initiators that are not on the list are denied access to the filer. *inpassword* is the inbound password for CHAP authentication. The filer uses the inbound password to authenticate the initiator. *inname* is a user name for inbound CHAP authentication. The filer uses the inbound user name to authenticate the initiator. *outpassword* is a password for outbound CHAP authentication. The filer uses this password for authentication by the initiator. *outname* is a user name for outbound CHAP authentication. The filer uses this user name for authentication by the initiator.

Removing an initiator from the authentication list

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

Step	Action
1	Enter the following command: **`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 filer using the default authentication method.

Displaying statistics for iSCSI sessions

Displaying iSCSI statistics

To display iSCSI statistics, complete the following step.

Step	Action

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

Interpreting iSCSI statistics

The `iscsi stats` command displays the following statistics:

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

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

  Note
  This release of Data ONTAP does not support SNACK messages.

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

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

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

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

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

Displaying information for iSCSI sessions

Types of session information

You can display all of the information about iSCSI sessions using the `iswta session show -v` command, or you can use the `-t` or `-p` option to specify which type of information you want to display.

- TCP connections
- iSCSI session parameters

Displaying TCP connection information

To display TCP connection information, complete the following step.

Step	Action

1	Enter the following command:
	`iscsi session show -t iswta`
	Result: The following information is displayed:
	`Session Initiator`
	`Sessions connected on adapter iswta:`
	`Session 1`
	`  Initiator Information`
	`  Initiator Name: iqn.1992-08.com.initvendor1:host1`
	`  ISID: 00:00:00:00:00:00`
	`  TCP Connection Information`
	`  Remote Endpoint: 10.60.155.58:32771`
	`  Local Endpoint: 10.60.155.31:3260`
	`  Local Interface: e5`

Displaying iSCSI session parameters

Display session parameters is a convenient way to confirm whether parameters are turned on or off. For example, keys such as HeaderDigest and DataDigest can dramatically affect performance. Rather than taking a network trace of the login request, you can display the current session parameters and confirm that they are turned off. To display iSCSI session parameters, complete the following step.

Step	Action

1	To view the current session on the istwa driver, enter the following command:
	`iscsi session show -p iswta`
	Result: The following information is displayed:
	`Session Initiator`
	`Sessions connected on adapter iswta:`
	`Session 1`
	`  Initiator Information`
	`  Initiator Name: iqn.1992-08.com.netapp.host1`
	`  ISID: 00:00:00:00:00:00`
	`  Session Parameters`
	`  InitiatorName=iqn.1992-08.com.netapp.host1`
	`  InitiatorAlias=`
	`  SessionType=Normal`
	`  ImmediateData=Yes`
	`  InitialR2T=Yes`
	`  MaxOutstandingR2T=1`
	`  HeaderDigest=None`
	`  DataDigest=None`
	`  MaxRecvDataSegmentLength=16384 (Initiator)`
	`  MaxRecvDataSegmentLength=65536 (Target)`
	`  FirstBurstLength=65536`
	`  MaxBurstLength=65536`
	`  MaxConnections=1`
	`  ErrorRecoveryLevel=0`
	`  DefaultTime2Wait=2`
	`  DefaultTime2Retain=20`
	`  DataPDUInOrder=Yes`
	`  DataSequenceInOrder=Yes`

Managing the iSCSI service on filer interfaces

Command to use

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

Displaying iSCSI status on filer interfaces

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

Step	Action

1	Enter the following command:
	`iswt interface show [-a|`*`interface`*`]`
	-a specifies all interfaces.
	interface is a specific Ethernet interface.
	Example: The following example shows the iSCSI service enabled on two filer Ethernet interfaces:
	`iswt interface -a`
	`Interface e0 enabled`
	`Interface e5 enabled`

Enabling iSCSI on a filer interface

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

Step	Action
1	Enter the following command: **`iswt interface enable {-a | interface ... }`** -*a* specifies all interfaces. *interface* is a specific Ethernet interface. **Example:** The following example enables the iSCSI service on interface e5: `iswt interface enable e5`

Disabling iSCSI on a filer interface

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

Step	Action
1	Enter the following command: **`iswt 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 `iswt` 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 specific Ethernet interface.

Enabling iSCSI on vifs

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

Glossary

client
A computer that shares files on a filer. See also host.

FCP
Fibre Channel Protocol. A licensed service on the filer that enables you to export LUNs to hosts using the SCSI protocol over a Fibre Channel fabric.

filer
A NetApp storage appliance, including the R100, FAS200, FAS800, and FAS900 series models.

HBA
Host bus adapter. An I/O adapter that connects a host I/O bus to a computer's memory system in SCSI environments. The HBA might be an FCP adapter or an iSCSI adapter.

host
Any computer system that accesses data on a filer as blocks using the iSCSI or FCP protocols, or is used to administer a filer.

igroup
Initiator group. A collection of unique identifiers, either iSCSI node names of initiators (hosts) in an IP network, or FCP WWPNs in a SCSI network, that are given access to LUNs when they are mapped to those LUNs.

initiator
The system component that originates an I/O command over an I/O bus or network; a host that has iSCSI initiator software installed on it, or a host that has a host bus adapter (HBA) installed in it, which is connected to the iSCSI or FCP network with the appropriate license enabled.

initiator group
See igroup.

iSCSI
A licensed service on the filer that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

iSCSI node name

A logical name to identify an iSCSI node, with the format 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.netapp.

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.

iswt

The iSCSI Software Target driver that provides the ability for filers to connect to initiators using the iSCSI protocol over TCP/IP. In a filer cluster, iswta refers to the local filer and iswtb refers to the partner. If the filer is not clustered, iswta refers to the single filer.

LUN

A logical unit of storage.

LUN clone

A complete copy of a LUN, which was initially created to be backed by a LUN or a file in a snapshot. The clone creates a complete copy of the LUN and frees the snapshot, which you can then delete.

LUN ID

The numerical identifier that the filer exports for a given LUN. The LUN ID is mapped to an igroup to enable host access.

LUN path

The path to a LUN on the filer. 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 filer.

map

To create an association between a LUN and an igroup. A LUN mapped to an igroup is exported to the nodes in the igroup (iqn or eui) when the LUN is online. LUN maps are used to secure access relationships between LUNs and the host.

online

Signifies that a LUN is exported to its mapped igroups. A LUN can be online only if it is enabled for read/write access.

offline

Disables the export of the LUN to its mapped igroups. The LUN is not available to hosts.

qtree

A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes. Qtrees can be used to group LUNs.

SAN

Storage Area Network. A storage network composed of one or more filers connected to one or more hosts in either a direct-attached or network-attached configuration using the iSCSI protocol over TCP/IP or the SCSI protocol over FCP.

share

An entity that allows the LUN's data to be accessible through multiple file protocols such as NFS and iSCSI. You can share a LUN for read or write access, or all permissions.

space reservations

An option that determines whether disk space is reserved for a specified LUN or file remains available for writes to any LUNs, files, or snapshots. Required for guaranteed space availability for a given LUN with or without snapshots.

storage appliance

Hardware and software-based storage systems, such as filers, that serve and protect data using protocols for both SAN and NAS networks.

target

The system component that receives a SCSI I/O command. A filer with the iSCSI or FCP license enabled and serving the data requested by the initiator.

volume

A file system. Volume refers to a functional unit of filer storage, based on one or more RAID groups, that is made available to the host. LUNs are stored in volumes.

WWN

World Wide Number. A unique 48- or 64-bit number assigned by a recognized naming authority (often through block assignment to a manufacturer) that identifies a connection for an FCP node to the storage network. A WWN is assigned for the life of a connection (device).

WWNN

World Wide Node Name. A unique 64-bit address represented in the following format: nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value. NetApp assigns a WWNN to a filer based on the serial number of its NVRAM. The WWNN is stored on disk. Data ONTAP refers to this number as a *Fibre Channel Nodename*, or simply, a node name.

WWPN

World Wide Port Name. A unique 64-bit address represented in the following format: nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value. Each Fibre Channel device has one or more ports that are used to connect to a SCSI network. Each port has a unique WWPN, which Data ONTAP refers to as an *FC Portname*, or simply, a port name.

Index

Symbols

/etc/nvfail_rename, database protection 100

B

backing store LUN 23
backup
    data to tape 94
    single LUNs to tape 95
    tape, when to use 98

C

CHAP authentication, description of 8
cluster configurations, required options for 9
command levels, description of 2
communication session, how they work 9

D

Data ONTAP, description of 2
database 100
database protection
    using /etc/nvfail_rename 100
    using vol options nvfail 100

E

eui type designator 7

F

filer administration
    using FilerView 3
    using the command line 2
filer, defined as target 2
finer node name, defined 7

I

initiator group, defined 8
initiator groups
    unmapping LUNs from 51
iqn node names,creating igroups with 8
iqn type designator 6
iSCSI
    nodes defined 5
    nodes, how connected 5
    nodes, how identified 6
iSCSI node names, how used 8
iSCSI, description of 5
iSNS, description of 8
iswt target port, identifying 8

L

lun commands
    lun online 50
    lun serial 53
    lun unmap 51
LUNs
    backing store 23
    bringing online 50
    changing serial number 53
    defined 5
    resizing restrictions 51
    serial number 5
    unmapping from initiator group 51

M

man page command 3

N

node name, of filer 7
node type designator
    eui 7
    iqn 6
nvfail option, of vol options command 100

R

restoring snapshots of LUNs 88

S

serial number, of LUNs, changing 53
Single File SnapRestore, using with LUNs 91
snapshot schedule, turning off 19
snapshots, using with SnapRestore 88

V

vFilers, how used 10
vol option nvfail, using with LUNs 100