



IBM System Storage N series
**Data ONTAP 8.0 7-Mode Block Access
Management Guide for iSCSI and FC**

Contents

Copyright information	11
Trademark information	13
About this guide	15
Audience	16
Supported features	16
Getting information, help, and services	16
Before you call	17
Using the documentation	17
Web sites	17
Accessing online technical support	17
Hardware service and support	17
Supported servers and operating systems	18
Firmware updates	18
Accessing Data ONTAP man pages	18
Terminology	19
Where to enter commands	20
Keyboard and formatting conventions	21
Special messages	22
How to send your comments	22
Introduction to block access	23
How hosts connect to storage systems	23
What Host Utilities are	23
What ALUA is	24
About SnapDrive for Windows and UNIX	24
How Data ONTAP implements an iSCSI network	25
What iSCSI is	25
What iSCSI nodes are	26
Supported configurations	26
How iSCSI nodes are identified	27
How the storage system checks initiator node names	28
Default port for iSCSI	28
What target portal groups are	28

What iSNS is	29
What CHAP authentication is	29
How iSCSI communication sessions work	30
How iSCSI works with HA pairs	30
Setting up the iSCSI protocol on a host and storage system	30
How Data ONTAP implements a Fibre Channel SAN	31
What FC is	32
What FC nodes are	32
How FC target nodes connect to the network	32
How FC nodes are identified	32
Unified Ethernet network management	34
Fibre Channel over Ethernet overview	35
What data center bridging is	35
Displaying DCB settings	36
Storage provisioning	39
Storage units for managing disk space	39
What autodelete is	40
What space reservation is	41
What fractional reserve is	42
Methods of provisioning storage in a SAN environment	43
Guidelines for provisioning storage in a SAN environment	44
Estimating how large a volume needs to be when using autodelete	45
Estimating how large a volume needs to be when using fractional reserve	46
Configuring volumes and LUNs when using autodelete	49
Retaining thinly provisioned LUNs online	54
About LUNs, igroups, and LUN maps	55
Information required to create a LUN	56
What igroups are	59
Required information for creating igroups	60
What LUN mapping is	62
Required information for mapping a LUN to an igroup	62
Guidelines for mapping LUNs to igroups	63
Mapping read-only LUNs to hosts at SnapMirror destinations	63
How to make LUNs available on specific FC target ports	64
Guidelines for LUN layout and space allocation	64

LUN alignment in virtual environments	65
Ways to create LUNs, create igroups, and map LUNs to igroups	65
Creating LUNs, creating igroups, and mapping LUNs with the LUN setup program	65
Creating LUNs, creating igroups, and mapping LUNs using individual commands	66
Creating LUNs on vFiler units for MultiStore	67
Displaying vFiler LUNs	68
LUN management	71
Displaying command-line Help for LUNs	71
Controlling LUN availability	72
Bringing LUNs online	72
Taking LUNs offline	73
Unmapping LUNs from igroups	73
Moving LUNs	74
Modifying LUN descriptions	74
Enabling and disabling space reservations for LUNs	75
Removing LUNs	75
Accessing LUNs with NAS protocols	76
Checking LUN, igroup, and FC settings	76
Displaying LUN serial numbers	78
Displaying LUN statistics	78
Displaying LUN mapping information	79
Displaying detailed LUN information	80
Displaying hidden staging area LUNs	80
igroup management	83
Creating igroups	83
Creating FCP igroups on UNIX hosts using the sanlun command	84
Deleting igroups	85
Adding initiators to an igroup	86
Removing initiators from an igroup	86
Displaying initiators	87
Renaming igroups	87
Setting the operating system type for an igroup	87
Enabling ALUA	88
When ALUA is automatically enabled	88

Manually setting the alua option to yes	88
Creating igroups for a non-default vFiler unit	89
Fibre Channel initiator request management	90
How Data ONTAP manages Fibre Channel initiator requests	90
How to use igroup throttles	90
How failover affects igroup throttles	91
Creating igroup throttles	91
Destroying igroup throttles	91
Borrowing queue resources from the unreserved pool	91
Displaying throttle information	92
Displaying igroup throttle usage	92
Displaying LUN statistics on exceeding throttles	93
iSCSI network management	95
Enabling multi-connection sessions	95
Enabling error recovery levels 1 and 2	96
iSCSI service management	97
Verifying that the iSCSI service is running	97
Verifying that iSCSI is licensed	97
Enabling the iSCSI license	98
Starting the iSCSI service	98
Stopping the iSCSI service	98
Displaying the target node name	98
Changing the target node name	99
Displaying the iSCSI target alias	100
Adding or changing the iSCSI target alias	100
iSCSI service management on storage system interfaces	101
Displaying iSCSI interface status	101
Enabling iSCSI on a storage system interface	102
Disabling iSCSI on a storage system interface	102
Displaying the storage system's target IP addresses	103
iSCSI interface access management	103
iSNS server registration	105
What an iSNS server does	105
How the storage system interacts with an iSNS server	105
About iSNS service version incompatibility	106
Setting the iSNS service revision	106

Registering the storage system with an ISNS server	107
Immediately updating the ISNS server	108
Disabling ISNS	108
Setting up vFiler units with the ISNS service	108
Displaying initiators connected to the storage system	109
iSCSI initiator security management	109
How iSCSI authentication works	110
Guidelines for using CHAP authentication	111
Defining an authentication method for an initiator	111
Defining a default authentication method for initiators	112
Displaying initiator authentication methods	113
Removing authentication settings for an initiator	113
iSCSI RADIUS configuration	113
Target portal group management	119
Range of values for target portal group tags	120
Important cautions for using target portal groups	120
Displaying target portal groups	121
Creating target portal groups	121
Destroying target portal groups	122
Adding interfaces to target portal groups	122
Removing interfaces from target portal groups	123
Configuring iSCSI target portal groups	123
Displaying iSCSI statistics	124
Definitions for iSCSI statistics	126
Displaying iSCSI session information	128
Displaying iSCSI connection information	129
Guidelines for using iSCSI with HA pairs	130
Simple HA pairs with iSCSI	130
Complex HA pairs with iSCSI	132
iSCSI problem resolution	132
LUNs not visible on the host	132
System cannot register with iSNS server	134
No multi-connection session	134
Sessions constantly connecting and disconnecting during takeover	134
Resolving iSCSI error messages on the storage system	135
FC SAN management	137

How to manage FC with HA pairs	137
How controller failover works	137
How to use port sets to make LUNs available on specific FC target ports	140
How port sets work in HA pairs	141
How upgrades affect port sets and igroups	141
How port sets affect igroup throttles	141
Creating port sets	142
Binding igroups to port sets	142
Unbinding igroups from port sets	143
Adding ports to port sets	143
Removing ports from port sets	144
Destroying port sets	144
Displaying the ports in a port set	145
Displaying igroup-to-port-set bindings	145
FC service management	145
Verifying that the FC service is running	146
Verifying that the FC service is licensed	146
Licensing the FC service	146
Disabling the FC license	147
Starting and stopping the FC service	147
Taking target expansion adapters offline and bringing them online	148
Changing the adapter speed	148
How WWPN assignments work with FC target expansion adapters	150
Changing the system's WWNN	152
WWPN aliases	153
Obtaining fabric zone server data	155
Obtaining a physical topology of the FC fabric	156
Obtaining fabric nameserver data	156
Checking connectivity of the initiators	157
Managing systems with onboard Fibre Channel adapters	158
Configuring onboard adapters for target mode	158
Configuring onboard adapters for initiator mode	160
Reconfiguring onboard FC adapters	161
Commands for displaying adapter information	162
Disk space management	173
Commands to display disk space information	173

Examples of disk space monitoring using the df command	174
Monitoring disk space on volumes with LUNs that do not use Snapshot copies	174
Monitoring disk space on volumes with LUNs that use Snapshot copies . .	176
How Data ONTAP can automatically provide more free space for full volumes ...	178
Configuring a FlexVol volume to grow automatically	179
Configuring automatic free space preservation for a FlexVol volume	179
Moving your volumes nondisruptively	180
Ways to use volume move	180
Requirements for performing a volume move	181
How the setup phase of volume move works	182
How the data copy phase of volume move works	182
How the cutover phase of volume move works	183
Performing the volume move operation	184
Pausing the volume move operation	185
Resuming the volume move operation	185
Monitoring the volume move status	186
Performing manual cutover of the volume move operation	186
Canceling the volume move operation	187
Working with VMware VAAI features for ESX hosts	187
Requirements for using the VAAI environment	188
Methods for determining whether VAAI features are supported	188
Statistics collected for VAAI features	189
Viewing statistics for the VAAI features	190
Data protection with Data ONTAP	193
Data protection methods	193
LUN clones	195
Reasons for cloning LUNs	196
Differences between FlexClone LUNs and LUN clones	196
Cloning LUNs	197
LUN clone splits	198
Displaying the progress of a clone-splitting operation	198
Stopping the clone-splitting process	199
Deleting Snapshot copies	199
Deleting backing Snapshot copies of deleted LUN clones	199
Deleting busy Snapshot copies	203

Restoring a Snapshot copy of a LUN in a volume 206

Restoring a single LUN 208

Backing up SAN systems to tape 209

Using volume copy to copy LUNs 212

Index 213

Copyright and trademark information

Copyright information

Copyright ©1994 - 2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

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

The product described in this manual may be protected by one or more U.S.A. 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

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp; the NetApp logo; the Network Appliance logo; Bycast; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StorageGRID; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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

NetApp is a licensee of the CompactFlash and CF Logo trademarks.

NetApp NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide describes how to use a storage system as Internet SCSI (iSCSI) and Fibre Channel (FC) protocol targets in a storage network. Specifically, this guide describes how to calculate the size of volumes containing logical units (LUNs), how to create and manage LUNs and initiator groups (igroups), and how to monitor iSCSI and FC traffic.

Note: This guide applies to systems, including systems with gateway functionality, running Data ONTAP 8.x 7-Mode. In the Data ONTAP 8.x 7-Mode product name, the term *7-Mode* signifies that the 8.x release has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Note: In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Next topics

[Audience](#) on page 16

[Supported features](#) on page 16

[Getting information, help, and services](#) on page 16

[Accessing Data ONTAP man pages](#) on page 18

[Terminology](#) on page 19

[Where to enter commands](#) on page 20

[Keyboard and formatting conventions](#) on page 21

[Special messages](#) on page 22

[How to send your comments](#) on page 22

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This guide is for system and storage administrators who are familiar with operating systems, such as Microsoft Windows 2003 and UNIX, that run on the hosts that access your storage systems. It also assumes that you know how block access protocols are used for block sharing or transfers.

This guide does not cover basic system or network administration topics, such as IP addressing, routing, and network topology.

Supported features

IBM® System Storage™ N series storage systems are driven by NetApp® Data ONTAP® software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details. Information about supported features can also be found at the following Web site:

www.ibm.com/storage/support/nas/

A listing of currently available N series products and features can be found at the following Web site:

www.ibm.com/storage/nas/

Getting information, help, and services

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Next topics

[Before you call](#) on page 17

[Using the documentation](#) on page 17

[Web sites](#) on page 17

[Accessing online technical support](#) on page 17

[Hardware service and support](#) on page 17

[Supported servers and operating systems](#) on page 18

[Firmware updates](#) on page 18

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected properly.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

Using the documentation

Information about N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Data ONTAP software publications are available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Web sites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- For NAS product information, go to the following Web site:
www.ibm.com/storage/nas/
- For NAS support information, go to the following Web site:
www.ibm.com/storage/support/nas/
- For AutoSupport information, go to the following Web site:
www.ibm.com/storage/support/nas/
- For the latest version of publications, go to the following Web site:
www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

www.ibm.com/storage/support/nas/

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following Web site for support telephone numbers:

www.ibm.com/planetwide/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, follow the link to the Interoperability Matrices from the following Web site:

www.ibm.com/storage/support/nas/

Firmware updates

As with all devices, it is recommended that you run the latest level of firmware, which can be downloaded by visiting the following Web site:

www.ibm.com/storage/support/nas/

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support. See the *Data ONTAP Upgrade Guide* for your version of Data ONTAP for more information on updating firmware.

Accessing Data ONTAP man pages

You can use the Data ONTAP manual (man) pages to access technical information.

About this task

Data ONTAP manual pages are available for the following types of information. They are grouped into sections according to standard UNIX naming conventions.

Types of information	Man page section
Commands	1
Special files	4
File formats and conventions	5
System management and services	8

Step

1. View man pages in the following ways:

- Enter the following command at the console command line:
`man command_or_file_name`
- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.

Note: All Data ONTAP 8.x 7-Mode man pages are stored on the system in files whose names are prefixed with the string "na_" to distinguish them from other man pages. The

prefixed names sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or service.

Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

Storage terms

array LUN	The storage that third-party storage arrays provide to storage systems running Data ONTAP software. One array LUN is the equivalent of one disk on a native disk shelf.
LUN (logical unit number)	A logical unit of storage identified by a number.
native disk	A disk that is sold as local storage for storage systems that run Data ONTAP software.
native disk shelf	A disk shelf that is sold as local storage for storage systems that run Data ONTAP software.
storage controller	The component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
storage system	The hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>gateways</i> , or <i>systems</i> .

Note: The term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

third-party storage The back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

Cluster and high-availability terms

cluster

- In Data ONTAP 8.x, a group of connected nodes (storage systems) that share a global namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- In the Data ONTAP 7.1 release family, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.

HA (high availability) In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.

HA pair In Data ONTAP 8.x, a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the FilerView graphical user interface.
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the client graphical user interface.
Your product documentation provides details about how to use the graphical user interface.
- You can enter commands either at the switch console or from any client that can obtain access to the switch using a Telnet session.

In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Keyboard conventions

Convention	What it means
The IBM NAS support site	Refers to www.ibm.com/storage/support/nas/ .
<i>Enter, enter</i>	<ul style="list-style-type: none"> Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic font</i>	<ul style="list-style-type: none"> Words or characters that require special attention. Placeholders for information that you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.
Monospaced font	<ul style="list-style-type: none"> Command names, option names, keywords, and daemon names. Information displayed on the system console or other computer monitors. Contents of files. File, path, and directory names.

Convention	What it means
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

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

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

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com. Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Introduction to block access

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

Next topics

[How hosts connect to storage systems](#) on page 23

[How Data ONTAP implements an iSCSI network](#) on page 25

[How Data ONTAP implements a Fibre Channel SAN](#) on page 31

[Unified Ethernet network management](#) on page 34

How hosts connect to storage systems

Hosts can connect to block storage using Internet small computer systems interface (iSCSI) or Fibre Channel (FC) protocol networks.

To connect to iSCSI networks, hosts can use standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs), or dedicated iSCSI host bus adapters (HBAs).

To connect to FC networks, hosts require Fibre Channel HBAs or CNAs.

Next topics

[What Host Utilities are](#) on page 23

[What ALUA is](#) on page 24

[About SnapDrive for Windows and UNIX](#) on page 24

What Host Utilities are

Host Utilities includes support software and documentation for connecting a supported host to an iSCSI or FC network.

The support software includes programs that display information about storage, and programs to collect information needed by Customer Support to diagnose problems. It also includes software to help tune and optimize the host settings for use in a IBM N series storage infrastructure.

Separate Host Utilities are offered for each supported host operating system. In some cases, different versions of the Host Utilities are available for different versions of the host operating system.

The documentation included with the Host Utilities describes how to install and use the Host Utilities software. It includes instructions for using the commands and features specific to your host operating system.

You must use the Host Utilities documentation along with this guide to set up and manage your iSCSI or FC network.

Related information

IBM Support Site - www.ibm.com/storage/support/nas/

What ALUA is

Data ONTAP 7.2 added support for the Asymmetric Logical Unit Access (ALUA) features of SCSI, also known as SCSI Target Port Groups or Target Port Group Support.

ALUA is an industry standard protocol for identifying optimized paths between a storage system and a host. ALUA enables the initiator to query the target about path attributes, such as primary path and secondary path. It also allows the target to communicate events back to the initiator. It is beneficial because multipathing software can be developed to support any array; proprietary SCSI commands are no longer required.

Attention: Ensure your host supports ALUA before enabling it. Enabling ALUA for a host that does not support it can cause host failures during cluster failover.

Related tasks

[Enabling ALUA](#) on page 88

About SnapDrive for Windows and UNIX

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

SnapDrive is a server-based software solution that provides advanced storage virtualization and management capabilities for Microsoft Windows environments. It is tightly integrated with Microsoft NTFS and provides a layer of abstraction between application data and physical storage associated with that data. SnapDrive runs on Windows Server hosts and complements native NTFS volume management with virtualization capabilities. It allows administrators to easily create virtual disks from pools of storage that can be distributed among several storage systems.

SnapDrive for UNIX provides simplified storage management, reduces operational costs, and improves storage management efficiency. It automates storage provisioning tasks and simplifies the process of creating host-consistent data Snapshot copies and clones from Snapshot copies.

Related information

IBM Support Site - www.ibm.com/storage/support/nas/

How Data ONTAP implements an iSCSI network

This section contains important concepts that are required to understand how Data ONTAP implements an iSCSI network.

Next topics

[What iSCSI is](#) on page 25

[What iSCSI nodes are](#) on page 26

[Supported configurations](#) on page 26

[How iSCSI nodes are identified](#) on page 27

[How the storage system checks initiator node names](#) on page 28

[Default port for iSCSI](#) on page 28

[What target portal groups are](#) on page 28

[What iSNS is](#) on page 29

[What CHAP authentication is](#) on page 29

[How iSCSI communication sessions work](#) on page 30

[How iSCSI works with HA pairs](#) on page 30

[Setting up the iSCSI protocol on a host and storage system](#) on page 30

What iSCSI is

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

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The iSCSI protocol is implemented over the storage system's standard gigabit Ethernet interfaces using a software driver.

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

Related information

[RFC 3720 - *www.ietf.org/*](#)

What iSCSI nodes are

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

Supported configurations

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

Next topics

[How iSCSI is implemented on the host](#) on page 26

[How iSCSI target nodes connect to the network](#) on page 26

Related information

[Fibre Channel and iSCSI Configuration Guide - www.ibm.com/storage/support/nas/](http://www.ibm.com/storage/support/nas/)

How iSCSI is implemented on the host

iSCSI can be implemented on the host in hardware or software.

You can implement iSCSI in one of the following ways:

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

How iSCSI target nodes connect to the network

You can implement iSCSI on the storage system using software or hardware solutions, depending on the model.

Target nodes can connect to the network:

- Over the system's Ethernet interfaces using software that is integrated into Data ONTAP. iSCSI can be implemented over multiple system interfaces, and an interface used for iSCSI can also transmit traffic for other protocols, such as CIFS and NFS.
- On the N3300 and N3600, N5000 series, and N7600, N7700, N7800, or N7900 systems, using an iSCSI target expansion adapter, to which some of the iSCSI protocol processing is offloaded. You can implement both hardware-based and software-based methods on the same system.
- Using a Fibre Channel over Ethernet (FCoE) unified target adapter (UTA).

How iSCSI nodes are identified

Every iSCSI node must have a node name.

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

Next topics

[*iqn-type designator*](#) on page 27

[*Storage system node name*](#) on page 28

[*eui-type designator*](#) on page 28

iqn-type designator

The iqn-type designator is a logical name that is not linked to an IP address.

It is based on the following components:

- The type designator, such as iqn
- A node name can contain alphabetic characters (a to z), numbers (0 to 9), and three special characters:
 - Period (“.”)
 - Hyphen (“-”)
 - Colon (“:”)
- The date when the naming authority acquired the domain name, followed by a period
- The name of the naming authority, optionally followed by a colon (:)
- A unique device name

Note: Some initiators might provide variations on the preceding format. Also, even though some hosts do support under-scores in the host name, they are not supported on IBM N series systems. For detailed information about the default initiator-supplied node name, see the documentation provided with your iSCSI Host Utilities.

The format is:

iqn.yyyy-mm.backward naming authority:unique device name

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

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

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

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

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

Storage system node name

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

The node name is displayed in the following format:

`iqn.1992-08.com.ibm:sn.serial-number`

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

`iqn.1992-08.com.ibm:sn.12345678`

eui-type designator

The eui-type designator is based on the type designator, eui, followed by a period, followed by sixteen hexadecimal digits.

The format is:

`eui.0123456789abcdef`

How the storage system checks initiator node names

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

Default port for iSCSI

The iSCSI protocol is configured in Data ONTAP to use TCP port number 3260.

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

What target portal groups are

A target portal group is a set of network portals within an iSCSI node over which an iSCSI session is conducted.

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

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

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

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

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

Note: Although this configuration is supported, it is not recommended for IBM N series storage systems. For more information, see the *Technical Report* on iSCSI Multipathing.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Related information

iSCSI Multipathing Possibilities on Windows with Data ONTAP-www.ibm.com/storage/support/nas/

What iSNS is

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

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

If you do not have an iSNS server on your network, you must manually configure each target to be visible to the host.

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

What CHAP authentication is

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

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

storage system verifies the response and authenticates the initiator. The CHAP password is used to compute the response.

How iSCSI communication sessions work

During an iSCSI session, the initiator and the target communicate over their standard Ethernet interfaces, unless the host has an iSCSI HBA or a FCoE CNA.

The storage system appears as a single iSCSI target node with one iSCSI node name. For storage systems with a MultiStore license enabled, each vFiler unit is a target with a different iSCSI node name.

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

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

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

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

How iSCSI works with HA pairs

HA pairs provide high availability because one system in the HA pair can take over if its partner fails. During failover, the working system assumes the IP addresses of the failed partner and can continue to support iSCSI LUNs.

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

Setting up the iSCSI protocol on a host and storage system

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

About this task

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

Steps

1. Install the initiator HBA and driver or software initiator on the host and record or change the host's iSCSI node name. It is recommended that you use the host name as part of the initiator node name to make it easier to associate the node name with the host.
2. Configure the storage system, including:
 - Licensing and starting the iSCSI service
 - Optionally configuring CHAP
 - Creating LUNs, creating an igroup that contains the host's iSCSI node name, and mapping the LUNs to that igroup

Note: If you are using SnapDrive, do not manually configure LUNs. Configure them using SnapDrive after it is installed.
3. Configure the initiator on the host, including:
 - Setting initiator parameters, including the IP address of the target on the storage system
 - Optionally configuring CHAP
 - Starting the iSCSI service
4. Access the LUNs from the host, including:
 - Creating file systems on the LUNs and mounting them, or configuring the LUNs as raw devices
 - Creating persistent mappings of LUNs to file systems

How Data ONTAP implements a Fibre Channel SAN

This section contains important concepts that are required to understand how Data ONTAP implements a Fibre Channel SAN.

Next topics

[*What FC is*](#) on page 32

[*What FC nodes are*](#) on page 32

[*How FC target nodes connect to the network*](#) on page 32

[*How FC nodes are identified*](#) on page 32

Related concepts

[*FC SAN management*](#) on page 137

What FC is

FC is a licensed service on the storage system that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over a Fibre Channel fabric.

Related concepts

[FC SAN management](#) on page 137

What FC nodes are

In a FC network, nodes include targets, initiators, and switches.

Targets are storage systems, and initiators are hosts. Nodes register with the Fabric Name Server when they are connected to a FC switch.

How FC target nodes connect to the network

Storage systems and hosts have adapters so they can be directly connected to each other or to FC switches with optical cable. For switch or storage system management, they might be connected to each other or to TCP/IP switches with Ethernet cable.

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

How FC nodes are identified

Each FC node is identified by a worldwide node name (WWNN) and a worldwide port name (WWPN).

Next topics

[How WWPNs are used](#) on page 32

[How storage systems are identified](#) on page 33

[How hosts are identified](#) on page 33

[How switches are identified](#) on page 33

How WWPNs are used

WWPNs identify each port on an adapter.

WWPNs are used for the following purposes:

- Creating an initiator group
The WWPNs of the host's HBAs are used to create an initiator group (igroup). An igroup is used to control host access to specific LUNs. You create an igroup by specifying a collection of WWPNs of initiators in an FC network. When you map a LUN on a storage system to an igroup, you grant all the initiators in that group access to that LUN. If a host's WWPN is not in an igroup that is mapped to a LUN, that host does not have access to the LUN. This means that the LUNs do not appear as disks on that host.

You can also create port sets to make a LUN visible only on specific target ports. A port set consists of a group of FC target ports. You bind a port set to an igroup. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

- Uniquely identifying a storage system's HBA target ports

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

Related concepts

[*Required information for mapping a LUN to an igroup*](#) on page 62

[*How to make LUNs available on specific FC target ports*](#) on page 64

How storage systems are identified

When the FCP service is first initialized, it assigns a WWNN to a storage system based on the serial number of its NVRAM adapter. The WWNN is stored on disk.

Each target port on the HBAs installed in the storage system has a unique WWPN. Both the WWNN and the WWPN are a 64-bit address represented in the following format:

nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value.

You can use commands such as `fcv show adapter`, `fcv config`, `sysconfig -v`, or `fcv nodename` to see the system's WWNN as `FC Nodename` or `nodename`, or the system's WWPN as `FC portname` or `portname`.

How hosts are identified

You use the `fcv show initiator` command to see all of the WWPNs, and any associated aliases, of the FC initiators that have logged on to the storage system. Data ONTAP displays the WWPN as `Portname`.

To know which WWPNs are associated with a specific host, see the FC Host Utilities documentation for your host. These documents describe commands supplied by the Host Utilities or the vendor of the initiator, or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, use the `lputilnt`, `HBAnywhere`, or `SANsurfer` applications, and for UNIX hosts, use the `sanlun` command.

How switches are identified

Fibre Channel switches have one WWNN for the device itself, and one WWPN for each of its ports.

For example, the following diagram shows how the WWPNs are assigned to each of the ports on a 16-port Brocade switch. For details about how the ports are numbered for a particular switch, see the vendor-supplied documentation for that switch.

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

Port numbers:															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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

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

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

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

Unified Ethernet network management

A unified Ethernet network entails running data and storage traffic, including iSCSI, CIFS, NFS, and Fibre Channel, over your existing Ethernet infrastructure.

Unified target adapters (UTAs) are 10-Gb Ethernet adapters that you install on your storage systems and converged network adapters (CNAs) are 10-Gb Ethernet adapters that you install on your hosts. These adapters are required for running Fibre Channel over Ethernet (FCoE) traffic, IP traffic, or both over your Ethernet network.

Note: UTAs and CNAs are configured and managed just like any other FC or Ethernet port; there are no unique configuration commands. See the *Data ONTAP 7-Mode File Access and Protocols Management Guide* for information about managing file system protocols.

In addition to the hardware components, Data ONTAP also supports the Data Center Bridging Exchange (DCBX) protocol, which is required for negotiating operating parameters that control transfers of both FC and Ethernet traffic over the Ethernet infrastructure.

Next topics

[Fibre Channel over Ethernet overview](#) on page 35

[What data center bridging is](#) on page 35

[Displaying DCB settings](#) on page 36

Related concepts

[iSCSI network management](#) on page 95

[FC SAN management](#) on page 137

Related information

[Technical Report: Fibre Channel over Ethernet \(FCoE\) End-to-End Deployment Guide -
www.ibm.com/storage/support/nas/](#)

[Fibre Channel and iSCSI Configuration Guide -
www.ibm.com/storage/support/nas/](#)

[IBM NAS documentation and support site -
www.ibm.com/storage/support/nas](#)

Fibre Channel over Ethernet overview

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. FCoE is very similar to traditional Fibre Channel (FC), as it maintains existing FC management and controls, but the hardware transport is a lossless 10 Gb Ethernet network.

Setting up an FCoE connection requires one or more supported converged network adapters (CNAs) in the host, connected to a supported data center bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

The CNA is presented to the host as both a FC HBA and a 10 Gb Ethernet adapter. The FC HBA portion of the CNA operates on all of the FC traffic when the FC traffic is sent and received as FC frames mapped into Ethernet packets (FC over Ethernet). The Ethernet adapter portion of the CNA operates on the host IP traffic, such as iSCSI, CIFS, NFS, and HTTP. Both the FCoE and IP portions of the CNA communicate over the same Ethernet port, which connects to the DCB switch.

Note: Unified target adapters (UTAs) are 10 Gb Ethernet adapters that you install on your storage systems. Starting with Data ONTAP 8.0.1, you can use UTAs for non-FCoE IP traffic such as NFS, CIFS, or iSCSI. This is *not* supported for Data ONTAP 8.0 and earlier.

In general, you configure and use FCoE connections just like traditional FC connections.

Note: For detailed information about how to set up and configure your host to run FCoE, see your appropriate host documentation.

What data center bridging is

Data center bridging (DCB) is a collection of extensions to the existing Ethernet standard that provides a lossless transport layer for FCoE traffic.

FC provides a reliable, dedicated fabric with no packet loss. Ethernet, however, is inherently lossy, which poses problems for transmitting FC traffic. The DCB standards solve this problem by implementing the following technologies:

Per-priority pause (priority-based flow control)	Enables a device to only inhibit the transmission of frames based on user-defined priorities.
Enhanced transmission selection	Allows administrators to allocate bandwidth on a percentage basis to different priorities.
Congestion notification	Transmits congestion information.
DCB Exchange (DCBX) protocol	Exchanges connection information with directly connected peers and detects misconfigurations.

Although these technologies possess their own independent functions, they operate together to provide an enhanced Ethernet standard that eliminates packet loss due to traffic congestion.

Related information

Technical Report: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide - www.ibm.com/storage/support/nas/

Data Center Bridging task group - www.ieee802.org/1/pages/dcbridges.html

Fibre Channel and iSCSI Configuration Guide - www.ibm.com/storage/support/nas/

Displaying DCB settings

When you install one or more UTAs, you can display the DCB settings associated with the adapters.

About this task

Note that these settings are configured at the switch level, and the storage system simply discovers and displays those pre-configured settings.

Steps

1. Enter the following command to include the bandwidth allocation:

```
dcb show interface
```

2. Enter the following command to display whether flow control is enabled for each priority:

```
dcb priority show interface
```

Result

The DCB settings are displayed as shown in the following examples.

```
system1> dcb show e2b
```

Interface	PGID	Priority	Applications	Bandwidth
-----	----	-----	-----	-----
e2b	0	0	unassigned	10%
	1	1 2 4 5 6 7	unassigned	0%
	2	3	FCoE	90%

```
system1>dcb priority show e2b
```

Interface	Priority	Applications	Flow Control	PGID
-----	-----	-----	-----	----
e2b	0	IP	enabled	0
	1	unassigned	disabled	1
	2	unassigned	disabled	1
	3	FCoE	enabled	2
	4	unassigned	disabled	1
	5	unassigned	disabled	1
	6	unassigned	disabled	1

	7	unassigned	disabled	1
Priority	The relative priorities for frames that have similar traffic handling requirements, such as latency and frame loss. The available priorities, from lowest to highest priority, are 0 to 7. The default priorities are 3 for FCoE traffic and 0 for IP traffic.			
Priority group	A collection of priorities bound together for the purpose of bandwidth allocation. A priority group can be associated with multiple priorities.			
Priority group ID (PGID)	A numerical ID from 0 to 15 that identifies each priority group.			
Bandwidth	The percentage of available bandwidth allocated to each priority group.			
Applications	Activities for which bandwidth and priorities are assigned, such as FCoE and IP traffic.			
Flow control	The flow control setting (<code>enabled</code> or <code>disabled</code>) for each priority. If priority-based flow control is enabled, then traffic at that priority might be paused to prevent frame loss due to congestion. Enabling priority-based flow control for one priority has no impact on traffic for a different priority.			

Storage provisioning

When you create a volume, you must estimate the amount of space you need for LUNs and Snapshot copies. You must also determine the amount of space you want to reserve so that applications can continue to write data to the LUNs in the volume.

Next topics

[*Storage units for managing disk space*](#) on page 39

[*What autodelete is*](#) on page 40

[*What space reservation is*](#) on page 41

[*What fractional reserve is*](#) on page 42

[*Methods of provisioning storage in a SAN environment*](#) on page 43

[*About LUNs, igroups, and LUN maps*](#) on page 55

[*Ways to create LUNs, create igroups, and map LUNs to igroups*](#) on page 65

[*Creating LUNs on vFiler units for MultiStore*](#) on page 67

Storage units for managing disk space

To properly provision storage, it is important to define and distinguish between the different units of storage.

The following list defines the various storage units:

Plaxes

A *plex* is a collection of one or more Redundant Array of Independent Disks (RAID) groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system aggregates or traditional volumes.

Data ONTAP uses plaxes as the unit of RAID-level mirroring when the SyncMirror software is enabled.

Aggregates

An *aggregate* is the physical layer of storage that consists of the disks within the RAID groups and the plaxes that contain the RAID groups.

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

Traditional or flexible volumes

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

A *FlexVol volume* is a volume that is loosely coupled to its containing aggregate. A FlexVol volume can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate can be the shared source of all the storage used by all the FlexVol volumes contained by that aggregate.

You use either traditional or FlexVol volumes to organize and manage system and user data. A volume can hold qtrees and LUNs.

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

Qtrees A *qtree* is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume in order to group LUNs.

LUNs A *LUN* is a logical unit of storage that represents all or part of an underlying physical disk.

You create LUNs in the root of a volume (traditional or flexible) or in the root of a qtree.

Note: Do not create LUNs in the root volume because it is used by Data ONTAP for system administration. The default root volume is /vol/vol0.

For detailed information about storage units, see the *Data ONTAP 7-Mode Storage Management Guide*.

Related information

IBM NAS documentation and support site - www.ibm.com/storage/support/nas

What autodelete is

Autodelete is a volume-level option that allows you to define a policy for automatically deleting Snapshot copies based on a definable threshold.

You can set that threshold, or *trigger*, to automatically delete Snapshot copies when:

- The volume is nearly full
- The snap reserve space is nearly full
- The overwrite reserved space is full

Using autodelete is recommended in most SAN configurations.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide* for more information about using autodelete to automatically delete Snapshot copies. Also see the Technical Report on thin provisioning below for additional details.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Related tasks

[Configuring volumes and LUNs when using autodelete](#) on page 49

[Estimating how large a volume needs to be when using autodelete](#) on page 45

Related information

[Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - media.netapp.com/documents/tr3483.pdf](http://media.netapp.com/documents/tr3483.pdf)

What space reservation is

When space reservation is enabled for one or more LUNs, Data ONTAP reserves enough space in the volume (traditional or FlexVol) so that writes to those LUNs do not fail because of a lack of disk space.

Note: LUNs in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

For example, if you create a 100-GB space reserved LUN in a 500-GB volume, that 100 GB of space is immediately allocated, leaving 400 GB remaining in the volume. In contrast, if space reservation is disabled on the LUN, all 500 GB in the volume remain available until writes are made to the LUN.

Space reservation is an attribute of the LUN; it is persistent across storage system reboots, takeovers, and givebacks. Space reservation is enabled for new LUNs by default, but you can create a LUN with space reservations disabled or enabled. After you create the LUN, you can change the space reservation attribute by using the `lun set reservation` command.

When a volume contains one or more LUNs with space reservation enabled, operations that require free space, such as the creation of Snapshot copies, are prevented from using the reserved space. If these operations do not have sufficient unreserved free space, they fail. However, writes to the LUNs with space reservation enabled will continue to succeed.

Related tasks

[Configuring volumes and LUNs when using autodelete](#) on page 49

What fractional reserve is

Fractional reserve is a volume option that enables you to determine how much space Data ONTAP reserves for Snapshot copy overwrites for LUNs, as well as for space-reserved files when all other space in the volume is used.

The fractional reserve setting defaults to 100%, but you can use the `vol options` command to set fractional reserve to any percentage from zero to 100.

It is best to use the autodelete function, but there may occasionally be circumstances under which fractional reserve can be used, including:

- When Snapshot copies cannot be deleted
- When preserving existing Snapshot copies is more important than creating new ones

Fractional reserve can be used on the following types of volumes:

- Traditional volumes
- FlexVol volumes with a space guarantee of `volume`
- FlexVol volumes with a space guarantee of `none`

You can only set fractional reserve for a volume with a space guarantee of `none` with Data ONTAP version 7.3.3 and later and version 8.0.1 and later.

Note: If the `guarantee` option for a FlexVol volume is set to `none` or `volume`, then fractional reserve for that volume can be set to the desired value. For the vast majority of configurations, you should set fractional reserve to zero when the `guarantee` option is set to `none` because it greatly simplifies space management. If the `guarantee` option for a FlexVol volume is set to `file`, then fractional reserve for that volume is set to 100 percent and is not adjustable.

If fractional reserve is set to 100%, when you create space-reserved LUNs, you can be sure that writes to those LUNs will always succeed without deleting Snapshot copies, even if all of the space-reserved LUNs are completely overwritten.

Setting fractional reserve to less than 100 percent causes the fractional reservation held for all space-reserved LUNs in that volume to be reduced to that percentage. Writes to the space-reserved LUNs in that volume are no longer unequivocally guaranteed, which is why you should use `snap autodelete` or `vol autogrow` for these volumes.

Fractional reserve is generally used for volumes that hold LUNs with a small percentage of data overwrite.

Note: If you are using fractional reserve in environments in which write errors due to lack of available space are unexpected and you are not using `snap autodelete` or `volume autosize`, you must monitor your free space and take corrective action to avoid write errors. Data ONTAP provides tools for monitoring available space in your volumes.

Note: Reducing the space reserved for overwrites (by using fractional reserve) does not affect the size of the space-reserved LUN. You can write data to the entire size of the LUN. The space reserved for overwrites is used only when the original data is overwritten.

Example

If you create a 500-GB space-reserved LUN, then Data ONTAP ensures that 500 GB of free space always remains available for that LUN to handle writes to the LUN.

If you then set fractional reserve to 50 for the LUN's containing volume and then take a Snapshot copy which locks the used blocks in the LUN, then Data ONTAP reserves 250 GB, or half of the space it was previously reserving for overwrites with fractional reserve set to 100. If more than half of the LUN is overwritten, then subsequent writes to the LUN could fail due to insufficient free space in the volume.

Note: When more than one LUN in the same volume has space reservations enabled, and fractional reserve for that volume is set to less than 100 percent, Data ONTAP does not limit any space-reserved LUN to its percentage of the reserved space. In other words, if you have two 100-GB LUNs in the same volume with fractional reserve set to 30, one of the LUNs could use up the entire 60 GB of reserved space for that volume.

For detailed information about using fractional reserve, see the Technical Report on thin provisioning.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Related tasks

[Configuring volumes and LUNs when using autodelete](#) on page 49

[Estimating how large a volume needs to be when using fractional reserve](#) on page 46

Related information

[Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - media.netapp.com/documents/tr3483.pdf](http://media.netapp.com/documents/tr3483.pdf)

Methods of provisioning storage in a SAN environment

When provisioning storage in a SAN environment, there are three primary methods to consider: using the autodelete feature, using the volume autosize feature, and using fractional reserve.

In Data ONTAP, fractional reserve is set to 100 percent and autodelete is disabled by default. However, in a SAN environment, it usually makes more sense to use autodelete (or sometimes autosize). In addition, this method is far simpler than using fractional reserve.

When using fractional reserve, you need to reserve enough space for the data inside the LUN, fractional reserve, and Snapshot copy, or: $X + Y + \text{delta}$. For example, you might need to reserve 50

GB for the LUN, 50 GB when fractional reserve is set to 100 percent, and 50 GB for Snapshot copy, or a volume of 150 GB. If fractional reserve is set to a percentage other than 100 percent, then the calculation becomes more complex.

In contrast, when using autodelete, you need only calculate the amount of space required for the LUN and Snapshot copy, or $X + \text{delta}$. Because you can configure the autodelete setting to automatically delete older Snapshot copies when space is required for data, you need not worry about running out of space for data.

For example, if you have a 100 GB volume, 50 GB is used for a LUN, and the remaining 50 GB is used for Snapshot copy. Or in that same 100 GB volume, you might reserve 30 GB for the LUN, and 70 GB is then allocated for Snapshot copies. In both cases, you can configure Snapshot copies to be automatically deleted to free up space for data, so fractional reserve is unnecessary.

Note: For detailed guidelines on using fractional reserve, see the technical report on thin provisioning.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Next topics

[Guidelines for provisioning storage in a SAN environment](#) on page 44

[Estimating how large a volume needs to be when using autodelete](#) on page 45

[Estimating how large a volume needs to be when using fractional reserve](#) on page 46

[Configuring volumes and LUNs when using autodelete](#) on page 49

[Retaining thinly provisioned LUNs online](#) on page 54

Related information

[Data ONTAP documentation on the NAS support site - \[www.ibm.com/storage/support/nas\]\(http://www.ibm.com/storage/support/nas\)](#)

[Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - \[media.netapp.com/documents/tr3483.pdf\]\(http://media.netapp.com/documents/tr3483.pdf\)](#)

Guidelines for provisioning storage in a SAN environment

When provisioning storage in a SAN environment, there are several best practices you should follow to ensure your systems run smoothly.

Follow these guidelines when creating traditional or FlexVol volumes that contain LUNs, regardless of which provisioning method you choose:

- Do not create any LUNs in the system's root volume.
Data ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.
- Ensure that no other files or directories exist in a volume that contains LUNs.

If this is not possible and you are storing LUNs and files in the same volume, use a separate qtree to contain the LUNs.

- If multiple hosts share the same volume, create a qtree on the volume to store all LUNs for the same host.

This is a recommended best practice that simplifies LUN administration and tracking.

- Ensure that the volume option `create_ucose` is set to on.
- Make the required changes to the snap reserve default settings.

Change the `snapreserve` setting for the volume to 0, set the `snap schedule` so that no controller-based Snapshot copies are taken, and delete all Snapshot copies after you create the volume.

- To simplify management, use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

For more information about creating volumes, see the *Data ONTAP 7-Mode Storage Management Guide*.

Related information

[Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas](http://www.ibm.com/storage/support/nas)
[Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - media.netapp.com/documents/tr3483.pdf](http://media.netapp.com/documents/tr3483.pdf)

Estimating how large a volume needs to be when using autodelete

Before you create a volume for use with autodelete, you can estimate how large it needs to be.

Steps

1. Calculate the Rate of Change (ROC) of your data per day. This value depends on how often you overwrite data. It is expressed as GB per day.
2. Calculate the amount of space you need for Snapshot copies by multiplying your ROC by the number of days of Snapshot copies you intend to keep.

Space required for Snapshot copies = ROC x number of days of Snapshot copies.

Example

You need a 200-GB LUN, and you estimate that your data changes at a rate of about 10 percent, or 20 GB each day. You want to take one Snapshot copy each day and want to keep three weeks' worth of Snapshot copies, for a total of 21 Snapshot copies. The amount of space you need for Snapshot copies is 21×20 GB, or 420 GB.

3. Calculate the required volume size by adding together the total data size and the space required for Snapshot copies.

Volume size calculation example

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

- You need to create two 200-GB LUNs.
The total LUN size is 400 GB.
- Your data changes at a rate of 10 percent of the total LUN size each day.
Your ROC is 40 GB per day (10 percent of 400 GB).
- You take one Snapshot copy each day and you want to keep the days of Snapshot copies for 10 days.
You need 400 GB of space for Snapshot copies (40 GB ROC × 10 Snapshot copies).
- You want to ensure that you can continue to write to the LUNs through the weekend, even after you take the last Snapshot copy and you have no more free space.

You would calculate the size of your volume as follows:

Volume size = Total data size + Space required for Snapshot copies.

The size of the volume in this example is 800 GB (400 GB + 400 GB).

After you finish

See the *Data Protection Online Backup and Recovery Guide* for more information about the autodelete function, and see the *Storage Management Guide* for more information about working with traditional and FlexVol volumes.

Related information

IBM NAS documentation and support site - www.ibm.com/storage/support/nas

Estimating how large a volume needs to be when using fractional reserve

Before you create a volume using fractional reserve, you can estimate how large it needs to be. The method you use to estimate the volume size depends on whether you need to create Snapshot copies of the volume.

1. *Calculating the total data size* on page 47
2. *Determining the volume size and fractional reserve setting when you need Snapshot copies* on page 47
3. *Determining the volume size when you do not need Snapshot copies* on page 49

Calculating the total data size

Determining the total data size—the sum of the sizes of all of the space-reserved LUNs in the volume—helps you estimate how large a volume needs to be.

Steps

1. Add up all of the space-reserved LUNs.

Example

If you know your database needs two 20-GB disks, you must create two 20-GB space-reserved LUNs. The total LUN size in this example is 40 GB.

2. Add in whatever amount of space you want to allocate for the non-space-reserved LUNs.

Note: This amount can vary, depending on the amount of space you have available and how much data you expect these LUNs to contain.

Determining the volume size and fractional reserve setting when you need Snapshot copies

The required volume size for a volume when you need Snapshot copies depends on several factors, including how much your data changes, how long you need to keep Snapshot copies, and how much data the volume is required to hold.

Steps

1. Calculate the Rate of Change (ROC) of your data per day.

This value depends on how often you overwrite data. It is expressed as GB per day.

2. Calculate the amount of space you need for Snapshot copies by multiplying your ROC by the number of days you want to keep Snapshot copies.

Space required for Snapshot copies = ROC × number of days the Snapshot copies will be kept

Example

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

3. Determine how much space you need for overwrites by multiplying your ROC by the amount of time, in days, you want to keep Snapshot copies before deleting.

Space required for overwrites = ROC × number of days you want to keep Snapshot copies before deleting

Example

You have a 20-GB LUN and your data changes at a rate of 2 GB each day. You want to ensure that write operations to the LUNs do not fail for three days after you take the last Snapshot copy. You need $2 \text{ GB} \times 3$, or 6 GB of space reserved for overwrites to the LUNs.

4. Calculate the required volume size by adding together the total data size, the space required for Snapshot copies, and the space required for overwrites.

Volume size = Total data size + space required for Snapshot copies + space required for overwrites

5. Calculate the fractional reserve value you must use for this volume by dividing the size of the space required for overwrites by the total size of the space-reserved LUNs in the volume.

Fractional reserve = space required for overwrites \div total data size.

Example

You have a 20-GB LUN. You require 6 GB for overwrites. Thirty percent of the total LUN size is 6 GB, so you must set your fractional reserve to 30.

Volume size calculation example

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

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

You would calculate the size of your volume as follows:

Volume size = Total data size + Space required for Snapshot copies + Space for overwrites.

The size of the volume in this example is 220 GB ($100 \text{ GB} + 100 \text{ GB} + 20 \text{ GB}$).

Note: This volume size requires that you set the fractional reserve setting for the new volume to 20. If you leave fractional reserve at 100 to ensure that writes could never fail, then you need to increase the volume size by 80 GB to accommodate the extra space needed for overwrites (100 GB rather than 20 GB).

Determining the volume size when you do not need Snapshot copies

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

Before you determine that you do not need Snapshot copies, verify the method for protecting data in your configuration. Most data protection methods, such as SnapRestore, SnapMirror, SnapManager for Microsoft Exchange or Microsoft SQL Server, SyncMirror, dump and restore, and `ndmcopy` methods rely on Snapshot copies. If you are using any of these methods, you cannot use this procedure to estimate volume size.

Note: Host-based backup methods do not require Snapshot copies.

Step

1. Use the following method to determine the required size of your volume, depending on your volume type.

If you are estimating a... Then...	
FlexVol volume	The FlexVol volume should be at least as large as the size of the data to be contained by the volume.
Traditional volume	The traditional volume should contain enough disks to hold the size of the data to be contained by the volume.

Example

If you need a traditional volume to contain two 200-GB LUNs, you should create the volume with enough disks to provide at least 400 GB of storage capacity.

Configuring volumes and LUNs when using autodelete

After you estimate how large your volumes should be, you can create your volumes, configure them with the necessary options, and create your LUNs.

1. [When to use the autodelete configuration](#) on page 50
2. [Setting volume options for the autodelete configuration](#) on page 50
3. [Required changes to Snapshot copy default settings](#) on page 51
4. [Verifying the create_ucose volume option](#) on page 53
5. [Enabling the create_ucose volume option](#) on page 54

Related tasks

[Creating LUNs, creating igroups, and mapping LUNs using individual commands](#) on page 66

When to use the autodelete configuration

Before implementing the autodelete configuration, it is important to consider the conditions under which this configuration works best.

The autodelete configuration is particularly useful under the following circumstances:

- You do not want your volumes to affect any other volumes in the aggregate.
For example, if you want to use the available space in an aggregate as a shared pool of storage for multiple volumes or applications, use the `autosize` option instead. Autosize is disabled under this configuration.
- Ensuring availability of your LUNs is more important to you than maintaining old Snapshot copies.

Setting volume options for the autodelete configuration

When implementing the autodelete configuration, you need to set the required volume space guarantee, `autosize`, fractional reserve, `try_first`, and Snapshot copy options.

Ensure you have created your volumes according to the guidelines in the *Data ONTAP 7-Mode Storage Management Guide*.

Note: For information about options related to Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide* and for information about volume options, see the *Data ONTAP 7-Mode Storage Management Guide*.

Steps

1. Set the space guarantee on the volumes by entering the following command:

```
vol options vol_name guarantee volume
```

2. Ensure that `autosize` is disabled by entering the following command:

```
vol autosize disable vol_name
```

Note: This option is disabled by default.

3. Set fractional reserve to zero percent, if it is not already, by entering the following command:

```
vol options vol_name fractional_reserve 0
```

4. Set the Snapshot copy reserve to zero percent by entering the following command:

```
snap reserve vol_name 0
```

The Snapshot copy space and application data is now combined into one large storage pool.

5. Configure Snapshot copies to begin being automatically deleted when the volume reaches the capacity threshold percentage by entering the following command:

```
snap autodelete vol_name trigger volume
```

Note: The capacity threshold percentage is based on the size of the volume. For more details, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

6. Set the `try_first` option to `snap_delete` by entering the following command:

```
vol options vol_name try_first snap_delete
```

This enables Data ONTAP to begin deleting Snapshot copies, starting with the oldest first, to free up space for application data.

7. Activate the snap autodelete settings by entering the following command:

```
snap autodelete vol_name on
```

After you set the volume options for the autodelete configuration, you must create your space-reserved LUNs.

Related tasks

[Creating LUNs, creating igroups, and mapping LUNs using individual commands](#) on page 66

Related information

[IBM NAS documentation and support site - www.ibm.com/storage/support/nas](http://www.ibm.com/storage/support/nas)

Required changes to Snapshot copy default settings

When you create a volume, Data ONTAP automatically schedules Snapshot copies and reserves space for them. You must modify these default settings to ensure that overwrites to LUNs in the volume do not fail.

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

When you create a volume, Data ONTAP automatically:

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

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

- Turn off the automatic Snapshot copy schedule.
- Delete all existing Snapshot copies.
- Set the percentage of space reserved for Snapshot copies to zero.

When finished, ensure the `create_ucode` volume is enabled.

Next topics

[Turning off the automatic Snapshot copy schedule](#) on page 52

[Deleting all existing Snapshot copies in a volume](#) on page 52

[*Setting the percentage of snap reserve space to zero*](#) on page 52

Turning off the automatic Snapshot copy schedule

When creating volumes that contain LUNs, turn off the automatic Snapshot copy schedule and verify that setting.

Steps

1. To turn off the automatic Snapshot copy schedule, enter the following command:

```
snap sched volname 0 0 0
```

Example

```
snap sched vol1 0 0 0
```

This command turns off the Snapshot copy schedule because there are no weekly, nightly, or hourly Snapshot copies scheduled. You can still take Snapshot copies manually by using the snap command.

2. To verify that the automatic Snapshot copy schedule is off, enter the following command:

```
snap sched [volname]
```

Example

```
snap sched vol1
```

The following output is a sample of what is displayed:

```
Volume vol1: 0 0 0
```

Deleting all existing Snapshot copies in a volume

When creating volumes that contain LUNs, delete all existing Snapshot copies in the volume.

Step

1. Enter the following command:

```
snap delete -a volname
```

Setting the percentage of snap reserve space to zero

When creating volumes that contain LUNs, set the percentage of space reserved for Snapshot copies to zero.

Steps

1. To set the percentage, enter the following command:

```
snap reserve volname percent
```

Example

```
snap reserve voll 0
```

2. To verify what percentage is set, enter the following command:

```
snap reserve [volname]
```

Example

```
snap reserve voll
```

The following output is a sample of what is displayed:

Volume voll: current snapshot reserve is 0% or 0 k-bytes.

Verifying the create_ucose volume option

You can use the `vol status` command to verify that the `create_ucose` volume option is enabled.

Step

1. To verify that the `create_ucose` option is enabled (on), enter the following command:

```
vol status [volname] -v
```

Example

```
vol status voll -v
```

Note: If you do not specify a volume, the status of all volumes is displayed.

The following output example shows that the `create_ucose` option is on:

Volume	State	Status	Options
voll	online	normal	nosnap=off, nosnapdir=off, minra=off, no_atime_update=off, raidsize=8, nvfail=off, snapmirrored=off, resyncsnaptime=60,create_ucose=on convert_ucose=off, maxdirsize=10240, fs_size_fixed=off, create_reserved=on raid_type=RAID4
Plex /vol/voll/plex0: online, normal, active			
RAID group /vol/voll/plex0/rg0: normal			

If necessary, enable the `create_ucose` volume option.

Enabling the create_unicode volume option

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

Step

1. To enable the `create_unicode` option, enter the following command:

```
vol options volname create_unicode on
```

Example

```
vol options vol1 create_unicode on
```

Retaining thinly provisioned LUNs online

When a LUN runs out of space and the containing volume cannot automatically grow further, the LUN goes offline. To retain the LUN online in an out-of-space condition, you should set the LUN option `-e space_alloc` to enable.

About this task

The LUN option `-e space_alloc` is set to `disable` by default. If this option is set to `disable`, then the LUN goes offline when it encounters the out-of-space condition.

For information about thin provisioning, see the technical report on thin provisioning.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Step

1. Enter the following command to retain the LUN online:

```
lun set space_alloc /vol/vol0/lun_name enable
```

Example

```
system1> lun set space_alloc /vol/vol0/lun1 enable
system1> lun set space_alloc /vol/vol0/lun1
Reporting of provisioning threshold events is enabled
```

Related information

[Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment - media.netapp.com/documents/tr3483.pdf](http://media.netapp.com/documents/tr3483.pdf)

About LUNs, igroups, and LUN maps

This section outlines the requirements for successfully provisioning storage and provides instructions for completing this process.

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

- Entering the `lun setup` command
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
- Using FilerView
This method provides a LUN Wizard that steps you through the process of creating and mapping new LUNs.
- Entering a series of individual commands (such as `lun create`, `igroup create`, and `lun map`)
Use this method to create one or more LUNs and igroups in any order.

Next topics

[Information required to create a LUN](#) on page 56

[What igroups are](#) on page 59

[Required information for creating igroups](#) on page 60

[What LUN mapping is](#) on page 62

[Required information for mapping a LUN to an igroup](#) on page 62

[Guidelines for mapping LUNs to igroups](#) on page 63

[Mapping read-only LUNs to hosts at SnapMirror destinations](#) on page 63

[How to make LUNs available on specific FC target ports](#) on page 64

[Guidelines for LUN layout and space allocation](#) on page 64

[LUN alignment in virtual environments](#) on page 65

Information required to create a LUN

When you create a LUN, you must specify the path name of the LUN, name of the LUN, LUN Multiprotocol Type, LUN size, LUN description, LUN identification number, and space reservation setting.

Next topics

Path name of the LUN on page 56

Name of the LUN on page 56

LUN Multiprotocol Type on page 56

LUN size on page 58

LUN description on page 58

LUN identification number on page 59

Space reservation setting on page 59

Path name of the LUN

The path name of a LUN must be at the root level of the qtree or volume in which the LUN is located.

Do not create LUNs in the root volume. The default root volume is /vol/vol0.

For clustered storage system configurations, it is recommended that you distribute LUNs across the cluster.

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

Name of the LUN

The name of the LUN is case-sensitive and can contain 1 to 255 characters. You cannot use spaces. LUN names must use only specific letters and characters.

LUN names can contain only the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), left brace (“{”), right brace (“}”), and period (“.”).

LUN Multiprotocol Type

The LUN Multiprotocol Type, or operating system type, specifies the OS of the host accessing the LUN. It also determines the layout of data on the LUN, the geometry used to access that data, and the minimum and maximum size of the LUN.

The LUN Multiprotocol Type values are `solaris`, `solaris_efi`, `windows`, `windows_gpt`, `windows_2008`, `hpux`, `aix`, `linux`, `netware`, `xen`, `hyper_v`, and `vmware`.

The following table describes the guidelines for using each LUN Multiprotocol Type:

LUN Multiprotocol Type	When to use
solaris	If your host operating system is Solaris and you are not using Solaris EFI labels.
solaris_efi	If you are using Solaris EFI labels. Note that using any other LUN Multiprotocol Type with Solaris EFI labels may result in LUN misalignment problems. Refer to your Solaris Host Utilities documentation and release notes for more information.
windows	If your host operating system is Windows 2000 Server, Windows XP, or Windows Server 2003 using the MBR partitioning method.
windows_gpt	If you want to use the GPT partitioning method and your host is capable of using it. Windows Server 2003, Service Pack 1 and later are capable of using the GPT partitioning method, and all 64-bit versions of Windows support it.
windows_2008	If your host operating system is Windows Server 2008; both MBR and GPT partitioning methods are supported.
hpux	If your host operating system is HP-UX.
aix	If your host operating system is AIX.
linux	If your host operating system is Linux.
netware	Your host operating system is Netware.
vmware	<p>If you are using ESX Server and your LUNs will be configured with VMFS.</p> <p>Note: If you configure the LUNs with RDM, use the guest operating system as the LUN Multiprotocol Type.</p>
xen	<p>If you are using Xen and your LUNs will be configured with Linux LVM with Dom0.</p> <p>Note: For raw LUNs, use the type of guest operating system as the LUN Multiprotocol Type.</p>
hyper_v	<p>If you are using Windows Server 2008 Hyper-V and your LUNs contain virtual hard disks (VHDs).</p> <p>Note: For raw LUNs, use the type of child operating system as the LUN Multiprotocol Type.</p>

Note: If you are using SnapDrive for Windows, the LUN Multiprotocol Type is automatically set.

When you create a LUN, you must specify the LUN type. Once the LUN is created, you cannot modify the LUN host operating system type.

See the N series Service and Support Web site for information about supported hosts.

Related information

IBM NAS support site - www.ibm.com/storage/support/nas/

LUN size

You specify the size of a LUN in bytes or by using specific multiplier suffixes.

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

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

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

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

LUN description

The LUN description is an optional attribute you use to specify additional information about the LUN.

You can edit this description at the command line or with FilerView.

LUN identification number

A LUN must have a unique identification number (ID) so that the host can identify and access the LUN. You map the LUN ID to an igroup so that all the hosts in that igroup can access the LUN.

If you do not specify a LUN ID, Data ONTAP automatically assigns one.

Space reservation setting

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

Note: You should keep space reservation on.

What igroups are

Initiator groups (igroups) are tables of FCP host WWPNs or iSCSI host nodenames. You define igroups and map them to LUNs to control which initiators have access to LUNs.

Typically, you want all of the host's HBAs or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA or software initiator of each clustered host needs redundant paths to the same LUN.

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

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator.

Note: An initiator cannot be a member of igroups of differing otypes. Also, a given igroup can be used for FCP or iSCSI, but not both.

Related concepts

[*igroup management*](#) on page 83

igroup example

You can create multiple igroups to define which LUNs are available to your hosts. For example, if you have a host cluster, you can use igroups to ensure that specific LUNs are visible to only one host in the cluster.

The following table illustrates how four igroups give access to the LUNs for four different hosts accessing the storage system. The clustered hosts (Host3 and Host4) are both members of the same igroup (aix-group2) and can access the LUNs mapped to this igroup. The igroup named aix-group3 contains the WWPNs of Host4 to store local information not intended to be seen by its partner.

Host with HBA WWPNs	igroups	WWPNs added to igroups	LUNs mapped to igroups
Host1, single-path (one HBA) 10:00:00:00:c9:2b:7c:0f	aix-group0	10:00:00:00:c9:2b:7c:0f	/vol/vol2/lun0
Host2, multipath (two HBAs) 10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	aix-group1	10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	/vol/vol2/lun1
Host3, multipath, clustered (connected to Host4) 10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02	aix-group2	10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/qtrees1/ lun2
Host4, multipath, clustered (connected to Host3) 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	aix-group3	10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/qtrees1/ lun3 /vol/vol2/qtrees1/ lun4

Required information for creating igroups

There are a number of attributes required when creating igroups, including the name of the igroup, type of igroup, ostype, iSCSI node name for iSCSI igroups, and WWPN for FCP igroups.

Next topics

[igroup name](#) on page 60

[igroup type](#) on page 61

[igroup ostype](#) on page 61

[iSCSI initiator node name](#) on page 61

[FCP initiator WWPN](#) on page 61

igroup name

The igroup name is a case-sensitive name that must satisfy several requirements.

The igroup name:

- Contains 1 to 96 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), colon (“:”), and period (“.”).
- Must start with a letter or number.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or Domain Name Service (DNS). If you name an igroup `aix1`, for example, it is not mapped to the actual IP host name (DNS name) of the host.

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

igroup type

The igroup type can be either `-i` for iSCSI or `-f` for FC.

igroup ostype

The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostyles of initiators are `solaris`, `windows`, `hpux`, `aix`, `netware`, `xen`, `hyper_v`, `vmware`, and `linux`.

You must select an ostyle for the igroup.

iSCSI initiator node name

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

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

FCP initiator WWPN

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

To know which WWPNs are associated with a specific host, see the Host Utilities documentation for your host. These documents describe commands supplied by the Host Utilities or the vendor of the initiator or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, use the `lputilnt`, `HBAnywhere`, and `SANsurfer` applications, and for UNIX hosts, use the `sanlun` command.

Related tasks

[Creating FCP igroups on UNIX hosts using the `sanlun` command](#) on page 84

What LUN mapping is

LUN mapping is the process of associating a LUN with an igroup. When you map the LUN to the igroup, you grant the initiators in the igroup access to the LUN.

Required information for mapping a LUN to an igroup

You must map a LUN to an igroup to make the LUN accessible to the host. Data ONTAP maintains a separate LUN map for each igroup to support a large number of hosts and to enforce access control.

Next topics

LUN name on page 62

igroup name on page 62

LUN identification number on page 62

LUN name

Specify the path name of the LUN to be mapped.

igroup name

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

LUN identification number

Assign a number for the LUN ID, or accept the default LUN ID.

Typically, the default LUN ID begins with 0 and increments by 1 for each additional LUN as it is created. The host associates the LUN ID with the location and path name of the LUN. The range of valid LUN ID numbers depends on the host.

Note: For detailed information, see the documentation provided with your Host Utilities.

If you are attempting to map a LUN when the cluster interconnect is down, you must not include a LUN ID, because the partner system will have no way of verifying that the LUN ID is unique. Data ONTAP reserves a range of LUN IDs for this purpose and automatically assigns the first available LUN ID in this range.

- If you are mapping the LUN from the primary system, Data ONTAP assigns a LUN in the range of 193 to 224.
- If you are mapping the LUN from the secondary system, Data ONTAP assigns a LUN in the range of 225 to 255.

For more information about HA pairs, refer to the *Data ONTAP 7-Mode High-Availability Configuration Guide*.

Guidelines for mapping LUNs to igroups

There are several important guidelines you must follow when mapping LUNs to an igroup.

- You can map two different LUNs with the same LUN ID to two different igroups without having a conflict, provided that the igroups do not share any initiators or only one of the LUNs is online at a given time.
- Make sure the LUNs are online before mapping them to an igroup. Do not map LUNs that are in the offline state.
- You can map a LUN only once to an igroup or a specific initiator.
- You can add a single initiator to multiple igroups, but the initiator can be mapped to a LUN only once. You cannot map a LUN to multiple igroups that contain the same initiator.
- You cannot use the same LUN ID for two LUNs mapped to the same igroup.
- You cannot map a LUN to both FC and iSCSI igroups if ALUA is enabled on one of the igroups. Run the `lun config_check` command to determine if any such conflicts exist.

Mapping read-only LUNs to hosts at SnapMirror destinations

When a qtree or volume containing LUNs is used as a SnapMirror source, the LUNs copied to the SnapMirror destination appear as read-only LUNs to the destination storage system. However, in prior versions of Data ONTAP, you could not manage these LUNs as long as the SnapMirror relationship was intact. As of Data ONTAP 7.2, there is limited ability to manage LUNs on the SnapMirror destination, even while the SnapMirror relationship is intact. In addition, you can manage LUN maps for LUNs on mirrored qtrees and volumes.

In prior versions of Data ONTAP, LUN maps created at the source location were copied to the destination storage system. In Data ONTAP 7.2, the LUN maps are stored in a separate database table, so they are no longer copied to the destination during the SnapMirror process.

As a result, the LUNs appear as unmapped and read-only. Therefore, you must explicitly map these read-only LUNs to the hosts at the destination. Once you map the LUNs to the host, the LUNs remain online, even after the SnapMirror relationship is broken.

You map these LUNs to the host in the same way that you map any other LUNs to a host.

The destination LUN is also assigned a new serial number. The online/offline status is inherited from the source LUN and cannot be changed on the destination LUN. The only operations allowed on read-only LUNs are `lun map`, `lun unmap`, `lun show`, `lun stats`, and changes to SCSI-2 reservations and SCSI-3 persistent reservations.

You can create new igroups on the destination, map the destination LUN to those igroups, or use any existing igroups. Once you set up the LUN maps for the destination LUN, you can continue to use the LUN, regardless of the current mirror relationship.

Once the mirror is broken, the LUN transparently migrates to a read/write state. Hosts may need to remount the device to notice the change.

Attention: Any attempt to write to read-only LUNs will fail, and might cause applications and hosts to fail as well. Before mapping read-only LUNs to hosts, ensure the operating system and application support read-only LUNs.

Also note that you cannot create LUNs on read-only qtrees or volumes. The LUNs that display in a mirrored destination inherit the read-only property from the container.

For more information about read-only LUNs and SnapMirror, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

How to make LUNs available on specific FC target ports

When you map a LUN to a Fibre Channel igroup, the LUN is available on all of the storage system's FC target ports if the igroup is not bound to a port set. A port set consists of a group of FC target ports.

By binding a port set to an igroup, you make the LUN available on a subset of the system's target ports. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

You define port sets for FC target ports only. You do not use port sets for iSCSI target ports.

Related concepts

[How to use port sets to make LUNs available on specific FC target ports](#) on page 140

Guidelines for LUN layout and space allocation

When you create LUNs, follow these guidelines for LUN layout and space allocation.

- Group LUNs according to their rates of change.
If you plan to take Snapshot copies, do not create LUNs with high rate of change in the same volumes as LUNs with a low rate of change. When you calculate the size of your volume, the rate of change of data enables you determine the amount of space you need for Snapshot copies. Data ONTAP takes Snapshot copies at the volume level, and the rate of change of data in all LUNs affects the amount of space needed for Snapshot copies. If you calculate your volume size based on a low rate of change, and you then create LUNs with a high rate of change in that volume, you might not have enough space for Snapshot copies.
- Keep backup LUNs in separate volumes.
Keep backup LUNs in separate volumes because the data in a backup LUN changes 100 percent for each backup period. For example, you might copy all the data in a LUN to a backup LUN and then move the backup LUN to tape each day. The data in the backup LUN changes 100 percent each day. If you want to keep backup LUNs in the same volume, calculate the size of the volume based on a high rate of change in your data.
- Quotas are another method you can use to allocate space.
For example, you might want to assign volume space to various database administrators and allow them to create and manage their own LUNs. You can organize the volume into qtrees with

quotas and enable the individual database administrators to manage the space they have been allocated.

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

LUN alignment in virtual environments

LUN alignment problems, which can lead to lower performance for your storage system, are common in virtualized server environments. In order to avoid LUN alignment problems, it is essential to follow the best practices for proper LUN alignment.

Refer to the following information for detailed guidelines and background information on provisioning storage in virtualized server environments.

Related information

[Storage Block Alignment with VMware Virtual Infrastructure and IBM System Storage N series - ftp://service.boulder.ibm.com/storage/isv/NS3593-0.pdf](http://service.boulder.ibm.com/storage/isv/NS3593-0.pdf)

Ways to create LUNs, create igroups, and map LUNs to igroups

The basic sequence for provisioning storage is to create the LUNs, create the igroups, and map the LUNs to the igroups. You can use the LUN setup program or individual commands to complete these tasks.

Next topics

[Creating LUNs, creating igroups, and mapping LUNs with the LUN setup program](#) on page 65

[Creating LUNs, creating igroups, and mapping LUNs using individual commands](#) on page 66

Creating LUNs, creating igroups, and mapping LUNs with the LUN setup program

LUN setup is a guided program that prompts you for the information needed to create a LUN and an igroup, and to map the LUN to the igroup. When a default is provided in brackets in the prompt, press Enter to accept it.

Before you begin

If you did not create volumes for storing LUNs before running the `lun setup` program, terminate the program and create volumes. If you want to use qtrees, create them before running the `lun setup` program.

Step

1. On the storage system command line, enter the following command:

```
lun setup
```

Result

The `lun setup` program displays prompts that lead you through the setup process.

Creating LUNs, creating igroups, and mapping LUNs using individual commands

Rather than use LUN setup, you can use individual commands to create LUNs, create igroups, and map the LUNs to the appropriate igroups.

Steps

1. Create a space-reserved LUN by entering the following command on the storage system command line:

```
lun create -s size -t ostype lun_path
```

`-s size` indicates the size of the LUN to be created, in bytes by default.

`-t ostype` indicates the LUN type. The LUN type refers to the operating system type, which determines the geometry used to store data on the LUN.

`lun_path` is the LUN's path name that includes the volume and qtree.

Example

The following example command creates a 5-GB LUN called `/vol/vol2/qtree1/lun3` that is accessible by a Windows host. Space reservation is enabled for the LUN.

```
lun create -s 5g -t windows /vol/vol2/qtree1/lun3
```

2. Create an igroup by entering the following command on the storage system command line:

```
igroup create {-i | -f} -t ostype initiator_group [node ...]
```

`-i` specifies that the igroup contains iSCSI node names.

`-f` specifies that the igroup contains FCP WWPNs.

`-t ostype` indicates the operating system type of the initiator. The values are `solaris`, `windows`, `2008`, `hpux`, `aix`, `netware`, `vmware`, and `linux`.

`initiator_group` is the name you specify as the name of the igroup.

`node` is a list of iSCSI node names or FCP WWPNs, separated by spaces.

Example

iSCSI example:

```
igroup create -i -t windows win_host5_group2 ign.
1991-05.com.microsoft:host5.domain.com
```

FCP example:

```
igroup create -f -t aix aix-igroup3 10:00:00:00c:2b:cc:92
```

3. Map the LUN to an igroup by entering the following command on the storage system command line:

```
lun map lun_path initiator_group [lun_id]
```

lun_path is the path name of the LUN you created.

initiator_group is the name of the igroup you created.

lun_id is the identification number that the initiator uses when the LUN is mapped to it. If you do not enter a number, Data ONTAP generates the next available LUN ID number.

Example

The following command maps /vol/vol1/qtrees1/lun3 to the igroup win_host5_group2 at LUN ID 0.

```
lun map /vol/vol1/qtrees1/lun3 win_host5_group2 0
```

Related concepts

[LUN size](#) on page 58

[LUN Multiprotocol Type](#) on page 56

[What igroups are](#) on page 59

Creating LUNs on vFiler units for MultiStore

The process for creating LUNs on vFiler units is slightly different from creating LUNs on other storage systems.

Before you begin

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

Note: SnapDrive can create, connect to, and manage LUNs on the vFiler units in the same way it does on the physical storage system.

Use the following guidelines when creating LUNs on vFiler units:

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

- As vFiler unit ownership of storage changes, so does ownership of the storage's LUNs.

About this task

You can issue LUN subcommands using the following methods:

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


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

Step

1. Enter the `lun create` command in the vFiler unit 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 unit at `/vol/vfstore/vf1_0`:

```
vfiler run vf1 lun create -s 2g -t windows /vol/vfstore/vf1_0/lun0
```

See the *Data ONTAP Multistore Management Guide* for more information.

Related information

Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas

Displaying vFiler LUNs

You might need to display all LUNs owned by a vFiler context. The command for displaying vFiler LUNs is slightly different from the command used on other storage systems.

Step

1. Enter the following command from the vFiler unit that contains the LUNs:

```
vfiler run * lun show
```

Result

The following information shows sample output:

```
==== 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)
```


LUN management

After you create your LUNs, you can manage them in a number of ways. For example, you can control LUN availability, unmap a LUN from an igroup, and remove, and rename a LUN.

You can use the command-line interface to manage LUNs.

Next topics

- [*Displaying command-line Help for LUNs*](#) on page 71
- [*Controlling LUN availability*](#) on page 72
- [*Unmapping LUNs from igroups*](#) on page 73
- [*Moving LUNs*](#) on page 74
- [*Modifying LUN descriptions*](#) on page 74
- [*Enabling and disabling space reservations for LUNs*](#) on page 75
- [*Removing LUNs*](#) on page 75
- [*Accessing LUNs with NAS protocols*](#) on page 76
- [*Checking LUN, igroup, and FC settings*](#) on page 76
- [*Displaying LUN serial numbers*](#) on page 78
- [*Displaying LUN statistics*](#) on page 78
- [*Displaying LUN mapping information*](#) on page 79
- [*Displaying detailed LUN information*](#) on page 80
- [*Displaying hidden staging area LUNs*](#) on page 80

Displaying command-line Help for LUNs

Use the `lun help` command to display online Help for all LUN commands and sub-commands.

Steps

1. On the storage system's command line, enter the following command:

```
lun help
```

A list of all LUN sub-commands is displayed:

```
lun help          - List LUN (logical unit of block storage) commands
lun config_check  - Check all lun/igroup/fcp settings for correctness
lun clone         - Manage LUN cloning
lun comment       - Display/Change descriptive comment string
lun create        - Create a LUN
lun destroy       - Destroy a LUN
lun map           - Map a LUN to an initiator group
lun maxsize       - Show the maximum possible size of a LUN on a given
volume or qtrees
```

<code>lun move</code>	- Move (rename) LUN
<code>lun offline</code>	- Stop block protocol access to LUN
<code>lun online</code>	- Restart block protocol access to LUN
<code>lun resize</code>	- Resize LUN
<code>lun serial</code>	- Display/change LUN serial number
<code>lun set</code>	- Manage LUN properties
<code>lun setup</code>	- Initialize/Configure LUNs, mapping
<code>lun share</code>	- Configure NAS file-sharing properties
<code>lun show</code>	- Display LUNs
<code>lun snap</code>	- Manage LUN and snapshot interactions
<code>lun stats</code>	- Displays or zeros read/write statistics for LUN
<code>lun unmap</code>	- Remove LUN mapping

2. To display the syntax for any of the subcommands, enter the following command:

```
lun help subcommand
```

Example

```
lun help show
```

Controlling LUN availability

Use the `lun online` and `lun offline` commands to control the availability of LUNs while preserving the LUN mappings.

Next topics

[Bringing LUNs online](#) on page 72

[Taking LUNs offline](#) on page 73

Bringing LUNs online

Use the `lun online` command to bring one or more LUNs back online, as described in the following step.

Before you begin

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

Step

1. Enter the following command:

```
lun online lun_path [lun_path ...]
```

Example

```
lun online /vol/vol1/lun0
```


Taking LUNs offline

Taking a LUN offline makes it unavailable for block protocol access. Use the `lun offline` command to take the LUN offline.

Before you begin

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

About this task

Taking a LUN offline makes it unavailable for block protocol access.

Step

1. To take a LUN offline, enter the following command:

```
lun offline lun_path [lun_path ...]
```

Example

```
lun offline /vol/vol1/lun0
```

Unmapping LUNs from igroups

You may need to occasionally unmap a LUN from an igroup. After you take the LUN offline, you can use the `lun unmap` command to unmap the LUN.

Steps

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
```

3. Bring the LUN back online:

```
lun online lun_path [lun_path ...]
```

Example

```
lun online /vol/vol1/lun1
```

Moving LUNs

You can use the `lun move` command to rename or move a LUN.

About this task

If you are organizing LUNs in qtrees, the existing path (`lun_path`) and the new path (`new_lun_path`) must be either in the same qtree or in another qtree in that same volume.

Note: This process is completely non-disruptive; it can be performed while the LUN is online and serving data.

Step

1. Enter the following command:

```
lun move lun_path new_lun_path
```

Example

```
lun move /vol/vol1/mylun /vol/vol1/mynewlun
```

Modifying LUN descriptions

You may have added a LUN description when creating the LUN. Use the `lun comment` command to modify that description or add a new one.

About this task

If you use spaces in the comment, enclose the comment in quotation marks.

Step

1. Enter the following command:

```
lun comment lun_path [comment]
```

Example

```
lun comment /vol/vol1/lun2 "10GB for payroll records"
```

Enabling and disabling space reservations for LUNs

Use the `lun set reservation` command to enable and disable space reservations for a LUN.

About this task

Attention: If you disable space reservations, write operations to a LUN might fail due to insufficient disk space, and the host application or operating system might crash. When write operations fail, Data ONTAP displays system messages (one message per file) on the console, or sends these messages to log files and other remote systems, as specified by its `/etc/syslog.conf` configuration file.

Steps

1. Enter the following command to display the status of space reservations for LUNs in a volume:

```
lun set reservation lun_path
```

Example

```
lun set reservation /vol/lunvol/hpux/lun0
```

Space Reservation for LUN /vol/lunvol/hpux/lun0 (inode 3903199): enabled

2. 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 LUNs

Use the `lun destroy` command to remove one or more LUNs.

About this task

Without the `-f` parameter, you must first take the LUN offline and unmap it, and then enter the `lun destroy` command.

Step

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.

Accessing LUNs with NAS protocols

When you create a LUN, you can only access it with the iSCSI or FC protocol by default. However, you can use NAS protocols to make a LUN available to a host if the NAS protocols are licensed and enabled on the storage system.

About this task

The usefulness of accessing a LUN over NAS protocols depends on the host application. For example, the application must be equipped to understand the format of the data within the LUN and be able to traverse any file system the LUN may contain. Access is provided to the LUN's raw data, but not to any particular piece of data within the LUN.

If you want to write to a LUN using a NAS protocol, you must take the LUN offline or unmap it to prevent an iSCSI or FCP host from overwriting data in the LUN.

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

Steps

1. Determine whether you want to read, write, or do both to the LUN over the NAS protocol and take the appropriate action:
 - If you want read access, the LUN can remain online.
 - If you want write access, ensure that the LUN is offline or unmapped.

2. Enter the following command:

```
lun share lun_path {none|read|write|all}
```

Example

```
lun share /vol/vol1/mtree1/lun2 read
```

The LUN is now readable over NAS.

Checking LUN, igroup, and FC settings

You can use the `lun config_check` command to verify a number of LUN, igroup, and FC settings.

About this task

The command performs the following actions:

- Checks whether any FC target interfaces are down.

- Verifies that the ALUA igroup settings are valid.
- Checks for nodename conflicts.
- Checks for igroup and LUN map conflicts.
- Checks for igroup ALUA conflicts.

Step

1. Enter the following command:

```
lun config_check [-v] [-S] [-s]
```

- Use the `-v` option for verbose mode, which provides detailed information about each check.
- Use the `-s` to only check the single_image cfmode settings.
- Use the `-s` option for silent mode, which only provides output if there are errors.

Example

```
system1> lun config_check -v
Checking for down fcp interfaces
=====
                No Problems Found

Checking initiators with mixed/incompatible settings
=====
                No Problems Found

Checking igroup ALUA settings
=====
                No Problems Found

Checking for nodename conflicts
=====
Checking for initiator group and lun map conflicts
=====
                No Problems Found

Checking for igroup ALUA conflicts
=====
                No Problems Found
```

Related concepts

[What ALUA is](#) on page 24

[igroup ostype](#) on page 61

[How Data ONTAP avoids igroup mapping conflicts during cluster failover](#) on page 138

Displaying LUN serial numbers

A LUN serial number is a unique, 12-byte, ASCII string generated by the storage system. Many multipathing software packages use this serial number to identify redundant paths to the same LUN.

About this task

Although the storage system displays the LUN serial number in ASCII format by default, you can display the serial number in hexadecimal format as well.

Step

1. Enter the following command:

```
lun show [-v] lun_path
or
```

```
lun serial [-x] lun_path new_lun_serial
```

Use `-v` option to display the serial numbers in ASCII format.

Use `-x` option to display the serial numbers in hexadecimal format.

Use `new_lun_serial` to change the existing LUN serial number to the specified serial number.

Note: Under normal circumstances, you should not change the LUN serial number. However, if you do need to change it, ensure the LUN is offline before issuing the command. Also, you can not use the `-x` option when changing the serial number; the new serial number must be in ASCII format.

Example

```
lun serial -x /vol/blocks_fvt/ncmds_lun2
```

```
Serial (hex)#: 0x4334656f476f424f594d2d6b
```

Displaying LUN statistics

You use the `lun stats` command to display the number of read and write operations and the number of operations per second for LUNs.

Step

1. Enter the following command:

```
lun stats -z -i interval -c count -o [-a | lun_path]
```

`-z` resets the statistics on all LUNs or the LUN specified in the `lun_path` option.

`-i interval` is the interval, in seconds, at which the statistics are displayed.

`-c count` is the number of intervals. For example, the `lun stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.

`-o` displays additional statistics, including the number of QFULL messages the storage system sends when its SCSI command queue is full and the amount of traffic received from the partner storage system.

`-a` shows statistics for all LUNs.

`lun_path` displays statistics for a specific LUN.

Example

```
system1>lun stats -o -i 1
```

Read Ops	Write Ops	Other Ops	QFull	Read kB	Write kB	Average Latency	Queue Length	Partner Ops	Partner kB	Lun
0	351	0	0	0	44992	11.35	3.00	0	0	/vol/tpcc/log_22
0	233	0	0	0	29888	14.85	2.05	0	0	/vol/tpcc/log_22
0	411	0	0	0	52672	8.93	2.08	0	0	/vol/tpcc/log_22
2	1	0	0	16	8	1.00	1.00	0	0	/vol/tpcc/ctrl_0
1	1	0	0	8	8	1.50	1.00	0	0	/vol/tpcc/ctrl_1
0	326	0	0	0	41600	11.93	3.00	0	0	/vol/tpcc/log_22
0	353	0	0	0	45056	10.57	2.09	0	0	/vol/tpcc/log_22
0	282	0	0	0	36160	12.81	2.07	0	0	/vol/tpcc/log_22

Displaying LUN mapping information

Use the `lun show -m` command to display a list of LUNs and the hosts to which they are mapped.

Step

1. On the storage system's command line, enter the following command:

```
lun show -m
```

Example

LUN path	Mapped to	LUN ID	Protocol
/vol/tpcc/ctrl_0	host5	0	iSCSI
/vol/tpcc/ctrl_1	host5	1	iSCSI
/vol/tpcc/crash1	host5	2	iSCSI
/vol/tpcc/crash2	host5	3	iSCSI
/vol/tpcc/cust_0	host6	4	iSCSI

/vol/tpcc/cust_1	host6	5	iSCSI
/vol/tpcc/cust_2	host6	6	iSCSI

Displaying detailed LUN information

Use the `lun show -v` command to show additional LUN details, such as the serial number, Multiprotocol type, and maps.

Step

1. On the storage system's command line, enter the following command to display LUN status and characteristics:

```
lun show -v
```

Example

```
/vol/tpcc_disks/cust_0_1 382m (400556032) (r/w, online, mapped)
  Serial#: VqmOVYoe3BUf
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: aix
  SnapValidator Offset: 1m (1048576)
  Maps: host5=0
  Cluster Shared Volume Information: 0x3
/vol/tpcc_disks/cust_0_2 382m (400556032) (r/w, online, mapped)
  Serial#: VqmOVYoe3BV6
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: aix
  SnapValidator Offset: 1m (1048576)
  Maps: host6=1
  Cluster Shared Volume Information: 0x0
```

Displaying hidden staging area LUNs

You can use the `lun show staging` command to obtain a list of all the hidden staging area LUNs. If you want to destroy an igroup to which the staging LUN is mapped, the `lun show staging` command indicates the reason for not being able to destroy an igroup.

About this task

The staging area luns are temporarily stored in `/vol/volnam/Staging_XXXX/lun_name` path when a nondisruptive restore is in progress and are automatically cleared when the restore completes.

successfully. If the nondisruptive restore fails, you should destroy the temporary LUNs manually using the `lun destroy` command.

Step

1. Obtain the list of hidden staging area LUNs by entering the following command:

```
lun show staging
```

Example: Hidden staging area LUNs

```
system1> lun show -v staging
/vol/volz/Staging_123/lun0      10m (10485760)  (r/w, online, mapped)
    Comment: "staging lun"
    Serial#: C4ab0o9vfayw
    Share: none
    Space Reservation: enabled
    Multiprotocol Type: linux
    Maps: gaston=1
/vol/volz/Staging_123/lun1      10m (10485760)  (r/w, online, mapped)
    Comment: "rrrrrrrrrrr"
    Serial#: C4ab0o9vppgp
    Share: none
    Space Reservation: enabled
    Multiprotocol Type: linux
    Maps: gaston=2
```


igroup management

To manage your initiator groups (igroups), you can perform a range of tasks, including creating igroups, destroying them, and renaming them.

Next topics

[Creating igroups](#) on page 83

[Creating FCP igroups on UNIX hosts using the sanlun command](#) on page 84

[Deleting igroups](#) on page 85

[Adding initiators to an igroup](#) on page 86

[Removing initiators from an igroup](#) on page 86

[Displaying initiators](#) on page 87

[Renaming igroups](#) on page 87

[Setting the operating system type for an igroup](#) on page 87

[Enabling ALUA](#) on page 88

[Creating igroups for a non-default vFiler unit](#) on page 89

[Fibre Channel initiator request management](#) on page 90

Related concepts

[What igroups are](#) on page 59

Creating igroups

Initiator groups, or igroups, are tables of host identifiers such as Fibre Channel WWPNs and iSCSI node names. You can use igroups to control which hosts can access specific LUNs.

Step

1. To create an igroup, enter the following command:

```
igroup create [-i | -f] -t ostype initiator_group [nodename ... /
WWPN ...] [wwpn alias ...] [-a portset]
```

-i indicates that it is an iSCSI igroup.

-f indicates that it is a FC igroup.

-t ostype indicates the operating system of the host. The values are solaris, windows 2008, hpux, aix, network, vmware, xen, hyper_v, and linux.

initiator_group is the name you give to the igroup.

nodename is an iSCSI node name. You can specify more than one node name.

WWPN is the FC worldwide port name. You can specify more than one WWPN.

wwpn alias is the name of the alias you created for a WWPN. You can specify more than one alias.

-a portset applies only to FC igroups. This binds the igroup to a port set. A port set is a group of target FC ports. When you bind an igroup to a port set, any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

Example

```
igroup create -i -t windows win-group0 ign.1991-05.com.microsoft:eng1
```

To create an iSCSI igroup called win-group0 that contains the node name of the Windows host associated with that node name.

Related concepts

[How to use port sets to make LUNs available on specific FC target ports](#) on page 140

[What igroups are](#) on page 59

Creating FCP igroups on UNIX hosts using the sanlun command

If you have a UNIX host, you can use the `sanlun` command to create FCP igroups. The command obtains the host's WWPNs and prints out the `igroup create` command with the correct arguments. Then you can copy and paste this command into the storage system command line.

Steps

1. Ensure that you are logged in as root on the host.
2. Change to the `/opt/ontap/santools/bin` directory.
3. Enter the following command to print a command to be run on the storage system that creates an igroup containing all the HBAs on your host:

```
./sanlun fcp show adapter -c
```

`-c` prints the full igroup create command on the screen.

The relevant igroup create command is displayed:

```
Enter this filer command to create an initiator group for this system:
igroup create -f -t aix "hostA" 10000000AA11BB22 10000000AA11EE33
```

In this example, the name of the host is **hostA**, so the name of the igroup with the two WWPNs is **hostA**.

4. Create a new session on the host and use the `telnet` command to access the storage system.

- 5. Copy the `igroup create` command from Step 3, paste the command on the storage system's command line, and press Enter to run the `igroup` command on the storage system.

An `igroup` is created on the storage system.

- 6. On the storage system's command line, enter the following command to verify the newly created `igroup`:

`igroup show`

Example

```
systemX> igroup show
hostA (FCP) (ostype: aix):
    10:00:00:00:AA:11:BB:22
    10:00:00:00:AA:11:EE:33
```

The newly-created `igroup` with the host's WWPNs is displayed.

Deleting igroups

When deleting `igroups`, you can use a single command to simultaneously remove the LUN mapping and delete the `igroup`. You can also use two separate commands to unmap the LUNs and delete the `igroup`.

Step

- 1. To delete one or more `igroups`, complete one of the following steps.

If you want to...	Then enter this command...
Remove LUN mappings before deleting the <code>igroup</code>	<pre>lun unmap lun-path igroup then igroup destroy igroup1 [igroup2, igroup3...]</pre>
Remove all LUN maps for an <code>igroup</code> and delete the <code>igroup</code> with one command	<pre>igroup destroy -f igroup1 [igroup2, igroup3...]</pre>

Example

```
lun unmap /vol/vol2/qtree/LUN10 win-group5
then

igroup destroy win-group5
```

Example

```
igroup destroy -f win-group5
```

Adding initiators to an igroup

Use the `igroup add` command to add initiators to an igroup.

About this task

An initiator cannot be a member of two igroups of differing types. For example, if you have an initiator that belongs to a Solaris igroup, Data ONTAP does not allow you to add this initiator to an AIX igroup.

Step

1. Enter the following command:

```
igroup add igroup_name [nodename|WWPN|WWPN alias]
```

Example

For Windows:

```
igroup add win-group2 iqn.1991-05.com.microsoft:eng2
```

For AIX:

```
igroup add aix-group2 10:00:00:00:c9:2b:02:1f
```

Removing initiators from an igroup

You can use the `igroup remove` command to remove an initiator from an igroup.

Step

1. Enter the following command:

```
igroup remove igroup_name [nodename|WWPN|WWPN alias]
```

Example

For Windows:

```
igroup remove win-group1 iqn.1991-05.com.microsoft:eng1
```

For AIX:

```
igroup remove aix-group1 10:00:00:00:c9:2b:7c:0f
```

Displaying initiators

Use the `igroup show` command to display all initiators belonging to a particular igroup.

Step

1. Enter the following command:

```
igroup show igroup_name
```

Example

```
igroup show win-group3
```

Renaming igroups

Use the `igroup rename` command to rename an igroup.

Step

1. Enter the following command:

```
igroup rename current_igroup_name new_igroup_name
```

Example

```
igroup rename win-group3 win-group4
```

Setting the operating system type for an igroup

When creating an igroup, you must set the operating system type, or `ostype`, to one of the following supported values: `solaris`, `windows`, `hpux`, `aix`, `linux`, `netware`, or `vmware`.

Step

1. Enter the following command:

```
igroup set [-f] igroup ostype value
```

`-f` overrides all warnings.

igroup is the name of the igroup.

value is the operating system type of the igroup.

Example

```
igroup set aix-group3 ostype aix
```

The ostype for igroup aix-group3 is set to aix.

Enabling ALUA

You can enable ALUA for your igroups, as long as the host supports the ALUA standard.

About this task

1. [When ALUA is automatically enabled](#) on page 88
2. [Manually setting the alua option to yes](#) on page 88

Related concepts

[What ALUA is](#) on page 24

Related tasks

[Configuring iSCSI target portal groups](#) on page 123

[Checking LUN, igroup, and FC settings](#) on page 76

When ALUA is automatically enabled

There are a number of circumstances under which an igroup is automatically enabled for ALUA.

When you create a new igroup or add the first initiator to an existing igroup, Data ONTAP checks whether that initiator is enabled for ALUA in an existing igroup. If so, the igroup being modified is automatically enabled for ALUA as well. Otherwise, you must manually set ALUA to `yes` for each igroup, unless the igroup ostype is `AIX`, `HP-UX`, or `Linux`. ALUA is automatically enabled for these operating systems.

Finally, if you map multiple igroups to a LUN and you enable one of the igroups for ALUA, you must enable all of the igroups for ALUA.

Related concepts

[What ALUA is](#) on page 24

Related tasks

[Configuring iSCSI target portal groups](#) on page 123

[Checking LUN, igroup, and FC settings](#) on page 76

Manually setting the alua option to yes

If ALUA is not automatically enabled for an igroup, you must manually set the `alua` option to `yes`.

Steps

1. Check whether ALUA is enabled by entering the following command:


```
igroup show -v igroup_name
```

Example

```
igroup show -v
```

```
system1> igroup show -v
linuxgrp (FCP):
OS Type: linux
Member: 10:00:00:00:c9:6b:76:49 (logged in on: vtic, 0a)
ALUA: No
```

2. If ALUA is not enabled, enter the following command to enable it:

```
igroup set igroup alua yes
```

Related concepts

[What ALUA is](#) on page 24

Related tasks

[Configuring iSCSI target portal groups](#) on page 123

[Checking LUN, igroup, and FC settings](#) on page 76

Creating igroups for a non-default vFiler unit

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

Steps

1. Change the context to the desired vFiler unit by entering the following command:

```
vfiler context vf1
```

The vFiler unit's prompt is displayed.

2. Create the igroup on vFiler unit determined in step 1 by entering the following command:

```
igroup create -i vf1_iscsi_group iqn.1991-05.com.microsoft:server1
```

3. Display the igroup by entering the following command:

```
igroup show
```

The following information is displayed:

```
vf1_iscsi_group (iSCSI) (ostype: windows):
iqn.1991-05.com.microsoft:server1
```

After you finish

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

Fibre Channel initiator request management

Data ONTAP implements a mechanism called igroup throttles, which you can use to ensure that critical initiators are guaranteed access to the queue resources and that less-critical initiators are not flooding the queue resources.

This section contains instructions for creating and managing igroup throttles.

Next topics

[How Data ONTAP manages Fibre Channel initiator requests](#) on page 90

[How to use igroup throttles](#) on page 90

[How failover affects igroup throttles](#) on page 91

[Creating igroup throttles](#) on page 91

[Destroying igroup throttles](#) on page 91

[Borrowing queue resources from the unreserved pool](#) on page 91

[Displaying throttle information](#) on page 92

[Displaying igroup throttle usage](#) on page 92

[Displaying LUN statistics on exceeding throttles](#) on page 93

How Data ONTAP manages Fibre Channel initiator requests

When you use igroup throttles, Data ONTAP calculates the total amount of command blocks available and allocates the appropriate number to reserve for an igroup, based on the percentage you specify when you create a throttle for that igroup.

Data ONTAP does not allow you to reserve more than 99 percent of all the resources. The remaining command blocks are always unreserved and are available for use by igroups without throttles.

How to use igroup throttles

You use igroup throttles to specify what percentage of the queue resources they can reserve for their use.

For example, if you set an igroup's throttle to be 20 percent, then 20 percent of the queue resources available at the storage system's ports are reserved for the initiators in that igroup. The remaining 80 percent of the queue resources are unreserved. In another example, if you have four hosts and they are in separate igroups, you might set the igroup throttle of the most critical host at 30 percent, the least critical at 10 percent, and the remaining two at 20 percent, leaving 20 percent of the resources unreserved.

Use igroup throttles to perform the following tasks:

- Create one igroup throttle per igroup, if desired.

Note: Any igroups without a throttle share all the unreserved queue resources.

- Assign a specific percentage of the queue resources on each physical port to the igroup.
- Reserve a minimum percentage of queue resources for a specific igroup.
- Restrict an igroup to a maximum percentage of use.
- Allow an igroup throttle to exceed its limit by borrowing from these resources:
 - The pool of unreserved resources to handle unexpected I/O requests
 - The pool of unused reserved resources, if those resources are available

How failover affects igroup throttles

Throttles manage physical ports, so during a takeover, their behavior is important to understand. Throttles apply to all ports and are divided by two when the HA pair is in takeover mode.

Creating igroup throttles

You can use igroup throttles to limit the number of concurrent I/O requests an initiator can send to the storage system, prevent initiators from flooding a port, prevent other initiators from accessing a LUN, and ensure that specific initiators have guaranteed access to the queue resources.

Step

1. Enter the following command:

```
igroup set igroup_name throttle_reserve percentage
```

Example

```
igroup set aix-igroup1 throttle_reserve 20
```

The igroup throttle is created for aix-igroup1, and it persists through reboots.

Destroying igroup throttles

You destroy an igroup throttle by setting the throttle reserve to zero.

Step

1. Enter the following command:

```
igroup set igroup_name throttle_reserve 0
```

Borrowing queue resources from the unreserved pool

If queue resources are available in the unreserved pool, you can borrow resources from the pool for a particular igroup.

About this task

To define whether an igroup can borrow queue resources from the unreserved pool, complete the following step with the appropriate option. The default when you create an igroup throttle is no.

Step

1. Enter the following command:

```
igroup set igroup_name throttle_borrow [yes|no]
```

Example

```
igroup set aix-igroup1 throttle_borrow yes
```

When you set the `throttle_borrow` setting to `yes`, the percentage of queue resources used by the initiators in the `igroup` might be exceeded if resources are available.

Displaying throttle information

You can use the `igroup show -t` command to display important information about the throttles assigned to `igroups`.

Step

1. Enter the following command:

```
igroup show -t
```

Example

```
system1>igroup show -t
```

name	reserved	exceeds	borrow
aix-igroup1	20%	0	N/A
aix-igroup2	10%	0	0

The *exceeds* column displays the number of times the initiator sends more requests than the throttle allows. The *borrow* column displays the number of times the throttle is exceeded and the storage system uses queue resources from the unreserved pool. In the *borrow* column, *N/A* indicates that the `igroup throttle_borrow` option is set to `no`.

Displaying igroup throttle usage

You can display real-time information about how many command blocks the initiator in the `igroup` is using, as well as the number of command blocks reserved for the `igroup` on the specified port.

Step

1. Enter the following command:

```
igroup show -t -i interval -c count [igroup|-a]
```

`-t` displays information on `igroup` throttles.

`-i interval` displays statistics for the throttles over an interval in seconds.

`-c count` determines how many intervals are shown.

igroup is the name of a specific `igroup` for which you want to show statistics.

-a displays statistics for all igroups, including idle igroups.

Example

```
igroup show -t -i 1
```

name	reserved	4a	4b	5a	5b
igroup1	20%	45/98	0/98	0/98	0/98
igroup2	10%	0/49	0/49	17/49	0/49
unreserved		87/344	0/344	112/344	0/344

The first number under the port name indicates the number of command blocks the initiator is using. The second number under the port name indicates the number of command blocks reserved for the igroup on that port.

In this example, the display indicates that igroup1 is using 45 of the 98 reserved command blocks on adapter 4a, and igroup2 is using 17 of the 49 reserved command blocks on adapter 5a.

igroups without throttles are counted as unreserved.

Displaying LUN statistics on exceeding throttles

Statistics are available about I/O requests for LUNs that exceed the igroup throttle. These statistics can be useful for troubleshooting and monitoring performance.

Steps

1. Enter the following command:

```
lun stats -o -i time_in_seconds
```

-i *time_in_seconds* is the interval over which performance statistics are reported. For example, -i 1 reports statistics each second.

-o displays additional statistics, including the number of QFULL messages, or "QFULLS".

Example

```
lun stats -o -i 1 /vol/vol1/lun2
```

The output displays performance statistics, including the QFULL column. This column indicates the number of initiator requests that exceeded the number allowed by the igroup throttle, and, as a result, received the SCSI Queue Full response.

2. Display the total count of QFULL messages sent for each LUN by entering the following command:

```
lun stats -o lun_path
```


iSCSI network management

This section describes how to manage the iSCSI service, as well as manage the storage system as a target in the iSCSI network.

Next topics

[Enabling multi-connection sessions](#) on page 95
[Enabling error recovery levels 1 and 2](#) on page 96
[iSCSI service management](#) on page 97
[iSNS server registration](#) on page 105
[Displaying initiators connected to the storage system](#) on page 109
[iSCSI initiator security management](#) on page 109
[Target portal group management](#) on page 119
[Displaying iSCSI statistics](#) on page 124
[Displaying iSCSI session information](#) on page 128
[Displaying iSCSI connection information](#) on page 129
[Guidelines for using iSCSI with HA pairs](#) on page 130
[iSCSI problem resolution](#) on page 132

Enabling multi-connection sessions

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

About this task

The `iscsi.max_connections_per_session` option specifies the number of connections per session allowed by the storage system. You can specify between 1 and 32 connections, or you can accept the default value.

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

Steps

1. Verify the current option setting by entering the following command on the system console:

```
options iscsi.max_connections_per_session
```

The current setting is displayed.

2. If needed, change the number of connections allowed by entering the following command:

```
options iscsi.max_connections_per_session [connections |  
use_system_default]
```

connections is the maximum number of connections allowed for each session, from 1 to 32.

use_system_default equals 1 for Data ONTAP 7.1 and 7.2, 4 for Data ONTAP 7.2.1 and subsequent 7.2 maintenance releases, and 32 starting with Data ONTAP 7.3. The meaning of this default might change in later releases.

Enabling error recovery levels 1 and 2

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

About this task

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

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

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

Steps

1. Verify the current option setting by entering the following command on the system console:

```
options iscsi.max_error_recovery_level
```

The current setting is displayed.

2. If needed, change the error recovery levels allowed by entering the following command:

```
options iscsi.max_error_recovery_level [level | use_system_default]
```

level is the maximum error recovery level allowed, 0, 1, or 2.

use_system_default equals 0 for Data ONTAP 7.1 and 7.2. The meaning of this default may change in later releases.

iSCSI service management

You need to ensure the iSCSI service is licensed and running on your system, as well as properly manage the target node name and target alias.

Next topics

[Verifying that the iSCSI service is running](#) on page 97
[Verifying that iSCSI is licensed](#) on page 97
[Enabling the iSCSI license](#) on page 98
[Starting the iSCSI service](#) on page 98
[Stopping the iSCSI service](#) on page 98
[Displaying the target node name](#) on page 98
[Changing the target node name](#) on page 99
[Displaying the iSCSI target alias](#) on page 100
[Adding or changing the iSCSI target alias](#) on page 100
[iSCSI service management on storage system interfaces](#) on page 101
[Displaying iSCSI interface status](#) on page 101
[Enabling iSCSI on a storage system interface](#) on page 102
[Disabling iSCSI on a storage system interface](#) on page 102
[Displaying the storage system's target IP addresses](#) on page 103
[iSCSI interface access management](#) on page 103

Verifying that the iSCSI service is running

You can use the `iscsi status` command to verify that the iSCSI service is running.

Step

1. On the storage system console, enter the following command:

```
iscsi status
```

A message is displayed indicating whether iSCSI service is running.

Verifying that iSCSI is licensed

Use the `license` command to verify that iSCSI is licensed on the storage system.

Step

1. On the storage system console, enter the following command:

```
license
```

A list of all available licenses is displayed. An enabled license shows the license code.

Enabling the iSCSI license

Use the `license add` command to enable the iSCSI license on the storage system.

About this task

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

- `volume option create_ucose to on`
- `cf.takeover.on_panic to on`

Step

1. On the storage system console, enter the following command:

```
license add license_code
```

license_code is the license code provided to you.

Starting the iSCSI service

Use the `iscsi start` command to start the iSCSI service on the storage system.

Step

1. On the storage system console, enter the following command:

```
iscsi start
```

Stopping the iSCSI service

Use the `iscsi stop` command to stop the iSCSI service on the storage system.

Step

1. On the storage system console, enter the following command:

```
iscsi stop
```

Displaying the target node name

Use the `iscsi nodename` command to display the storage system's target node name.

Step

1. On the storage system console, enter the following command:

```
iscsi nodename
```

Example

```
iscsi nodename
iSCSI target nodename: iqn.1992-08.com.ibm:sn.12345678
```

Changing the target node name

You may need to change the storage system's target node name.

About this task

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

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

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

Step

1. On the storage system console, enter the following command:

```
iscsi nodename iqn.1992-08.com.ibm:unique_device_name
```

Example

```
iscsi nodename iqn.1992-08.com.ibm:filerhq
```

Displaying the iSCSI target alias

The target alias is an optional name for the iSCSI target consisting of a text string with a maximum of 128 characters. It is displayed by an initiator's user interface to make it easier for someone to identify the desired target in a list of targets.

About this task

Depending on your initiator, the alias may or may not be displayed in the initiator's user interface.

Step

1. On the storage system console, enter the following command:

```
iscsi alias
```

Example

```
iscsi alias  
iSCSI target alias: Filer_1
```

Adding or changing the iSCSI target alias

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

Step

1. On the storage system console, enter the following command:

```
iscsi alias [-c | string]
```

-c clears the existing alias value

string is the new alias value, maximum 128 characters

Examples

```
iscsi alias Storage-System_2
New iSCSI target alias: Storage-System_2
```

```
iscsi alias -c
Clearing iSCSI target alias
```

iSCSI service management on storage system interfaces

Use the `iscsi interface` command to manage the iSCSI service on the storage system's Ethernet interfaces.

You can control which network interfaces are used for iSCSI communication. For example, you can enable iSCSI communication over specific gigabit Ethernet (GbE) interfaces.

By default, the iSCSI service is enabled on all Ethernet interfaces after you enable the license. Do not use 10/100 megabit Ethernet interfaces for iSCSI communication. The `e0m` management interface on storage systems is a 10/100 interface.

Displaying iSCSI interface status

Use the `iscsi interface show` command to display the status of the iSCSI service on a storage system interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface show [-a | interface]
```

`-a` specifies all interfaces. This is the default.

`interface` is list of specific Ethernet interfaces, separated by spaces.

Example

The following example shows the iSCSI service enabled on two storage system Ethernet interfaces:

```
iscsi interface show
Interface e0 disabled
```

```
Interface e9a enabled
Interface e9b enabled
```

Enabling iSCSI on a storage system interface

Use the `iscsi interface enable` command to enable the iSCSI service on an interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface enable [-a | interface ...]
```

`-a` specifies all interfaces.

interface is list of specific Ethernet interfaces, separated by spaces.

Example

The following example enables the iSCSI service on interfaces e9a and e9b:

```
iscsi interface enable e9a e9b
```

Disabling iSCSI on a storage system interface

You can use the `iscsi interface disable` command to disable the iSCSI service on an interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface disable [-f] {-a | interface ...}
```

`-f` forces the termination of any outstanding iSCSI sessions without prompting you for confirmation. If you do not use this option, the command displays a message notifying you that active sessions are in progress on the interface and requests confirmation before terminating these sessions and disabling the interface.

`-a` specifies all interfaces.

interface is a list of specific Ethernet interfaces, separated by spaces.

Displaying the storage system's target IP addresses

Use the `iscsi portal show` command to display the target IP addresses of the storage system. The storage system's target IP addresses are the addresses of the interfaces used for the iSCSI protocol.

Step

1. On the storage system console, enter the following command:

```
iscsi portal show
```

Result

The IP address, TCP port number, target portal group tag, and interface identifier are displayed for each interface.

Example

```
system1> iscsi portal show
Network portals:
IP address          TCP Port  TPGroup  Interface
10.60.155.105       3260     1000     e0
fe80::2a0:98ff:fe00:fd81 3260     1000     e0
10.1.1.10           3260     1003     e10a
fe80::200:c9ff:fe44:212b 3260     1003     e10a
```

iSCSI interface access management

Although you can use the `iscsi interface enable` command to enable the iSCSI service on an iSCSI interface, this command enables access for all initiators. As of Data ONTAP 7.3, you can use access lists to control the interfaces over which an initiator can access the storage system.

Access lists are useful in a number of ways:

- Performance: in some cases, you may achieve better performance by limiting the number of interfaces an initiator can access.
- Security: you can gain finer control over access to the interfaces.
- Controller failover: rather than contact all interfaces advertised by the storage system during giveback, the host will only attempt to contact the interfaces to which it has access, thereby improving failover times.

By default, all initiators have access to all interfaces, so access lists must be explicitly defined. When an initiator begins a discovery session using an iSCSI `SendTargets` command, it will only receive those IP addresses associated with network interfaces on its access list.

Next topics

[Creating iSCSI interface access lists](#) on page 104

[Removing interfaces from iSCSI interface access lists](#) on page 104

[Displaying iSCSI interface access lists](#) on page 105

Creating iSCSI interface access lists

You can use iSCSI interface access lists to control which interfaces an initiator can access. An access list ensures that an initiator only logs in with IP addresses associated with the interfaces defined in the access list.

Access list policies are based on the interface name, and can include physical interfaces, VIFs, and VLANs.

Note: For vFiler contexts, all interfaces can be added to the vFiler unit's access list, but the initiator will only be able to access the interfaces that are bound to the vFiler unit's IP addresses.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist add initiator name [-a | interface...]
```

-a specifies all interfaces. This is the default.

interface lists specific Ethernet interfaces, separated by spaces.

Example

```
iscsi interface accesslist add iqn.1991-05.com.microsoft:ms e0b
```

Related concepts

[Guidelines for using iSCSI with HA pairs](#) on page 130

Removing interfaces from iSCSI interface access lists

If you created an access list, you can remove one or more interfaces from the access list.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist remove initiator name [-a | interface...]
```

-a specifies all interfaces. This is the default.

interface lists specific Ethernet interfaces, separated by spaces.

Example

```
iscsi interface accesslist remove iqn.1991-05.com.microsoft:ms e0b
```


Displaying iSCSI interface access lists

If you created one or more access lists, you can display the initiators and the interfaces to which they have access.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist show
```

Example

```
system1> iscsi interface accesslist show
Initiator Nodename          Access List
iqn.1987-05.com.cisco:redhat    e0a, e0b
iqn.1991-05.com.microsoft:ms    e9
```

Only initiators defined as part of an access list are displayed.

iSNS server registration

You must ensure that your storage systems are properly registered with an Internet Storage Name Service server.

Next topics

[*What an iSNS server does*](#) on page 105

[*How the storage system interacts with an iSNS server*](#) on page 105

[*About iSNS service version incompatibility*](#) on page 106

[*Setting the iSNS service revision*](#) on page 106

[*Registering the storage system with an iSNS server*](#) on page 107

[*Immediately updating the iSNS server*](#) on page 108

[*Disabling iSNS*](#) on page 108

[*Setting up vFiler units with the iSNS service*](#) on page 108

What an iSNS server does

An iSNS server uses the Internet Storage Name Service 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 the storage system interacts with an iSNS server

The storage system automatically registers its IP address, node name, and portal groups with the iSNS server when the iSCSI service is started and iSNS is enabled. After iSNS is initially configured,

Data ONTAP automatically updates the iSNS server any time the storage system's configuration settings change.

There can be a delay of a few minutes between the time of the configuration change and the update being sent; you can use the `iscsi isns update` command to send an update immediately.

About iSNS service version incompatibility

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

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

If your Data ONTAP version is...	And your iSNS server version is...	Then you should...
7.1	Prior to 3.0	Set <code>iscsi.isns.rev</code> option to 18 or upgrade to iSNS server 3.0.
7.1	3.0	Verify that the <code>iscsi.isns.rev</code> option is set to 22.

Note: When you upgrade to a new version of Data ONTAP, the existing value for the `iscsi.isns.rev` option is maintained. This reduces the risk of a draft version problem when upgrading. If necessary, you must manually change `iscsi.isns.rev` to the correct value when upgrading Data ONTAP.

Setting the iSNS service revision

You can configure Data ONTAP to use a different iSNS draft version by changing the `iscsi.isns.rev` option on the storage system.

Steps

1. Verify the current iSNS revision value by entering the following command on the system console:

```
options iscsi.isns.rev
```

The current draft revision used by the storage system is displayed.

2. If needed, change the iSNS revision value by entering the following command:

```
options iscsi.isns.rev draft
```

draft is the iSNS standard draft revision, either 18 or 22.

Registering the storage system with an iSNS server

Use the `iscsi isns` command to configure the storage system to register with an iSNS server. This command specifies the information the storage system sends to the iSNS server.

About this task

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

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

Steps

1. Make sure the iSCSI service is running by entering the following command on the storage system console:

```
iscsi status
```

2. If the iSCSI service is not running, enter the following command:

```
iscsi start
```

3. On the storage system console, enter the following command to identify the iSNS server that the storage system registers with:

```
iscsi isns config [ip_addr/hostname]
```

ip_addr is the IP address of the iSNS server.

hostname is the hostname associated with the iSNS server.

Note: As of Data ONTAP 7.3.1, you can configure iSNS with an IPv6 address.

Note: As of Data ONTAP 8.0.1, you can configure iSNS with an IPv6 address.

4. Enter the following command:

```
iscsi isns start
```

The iSNS service is started and the storage system registers with the iSNS server.

Note: iSNS registration is persistent across reboots if the iSCSI service is running and iSNS is started.

Immediately updating the iSNS server

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

Step

1. On the storage system console, enter the following command:

```
iscsi isns update
```

Disabling iSNS

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

Step

1. On the storage system console, enter the following command:

```
iscsi isns stop
```

Setting up vFiler units with the iSNS service

Use the `iscsi isns` command on each vFiler unit to configure which iSNS server to use and to turn iSNS registration on or off.

About this task

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *Data ONTAP 7-Mode MultiStore Management Guide*.

Steps

1. Register the vFiler unit with the iSNS service by entering the following command:

```
iscsi isns config -i ip_addr
```

ip_addr is the IP address of the iSNS server.

2. Enter the following command to enable the iSNS service:

```
iscsi isns start
```

Examples for vFiler units

The following example defines the iSNS server for the default vFiler unit (vfiler0) on the hosting storage system:

```
iscsi isns config -i 10.10.122.101
```

The following example defines the iSNS server for a specific vFiler unit (vf1). The vfiler context command switches to the command line for a specific vFiler unit.

```
vfiler context vf1
vf1> iscsi isns config -i 10.10.122.101
```

Related information

Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas

Displaying initiators connected to the storage system

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

Step

1. On the storage system console, enter the following command:

```
iscsi initiator show
```

The initiators currently connected to the storage system are displayed.

Example

```
system1> iscsi initiator show
Initiators connected:
  TSIH  TPGroup  Initiator/ISID/IGroup
    1    1000    iqn.1991-05.com.microsoft:hual-lxp.hq.ibm.com /
40:00:01:37:00:00 / windows_ig2; windows_ig
    2    1000    vanclibern (iqn.1987-05.com.cisco:vanclibern /
00:02:3d:00:00:01 / linux_ig)
    4    1000    iqn.1991-05.com.microsoft:cox / 40:00:01:37:00:00 /
```

iSCSI initiator security management

Data ONTAP provides a number of features for managing security for iSCSI initiators. You can define a list of iSCSI initiators and the authentication method for each, display the initiators and their associated authentication methods in the authentication list, add and remove initiators from the authentication list, and define the default iSCSI initiator authentication method for initiators not in

the list. In addition, you can configure your storage systems to use Remote Authentication Dial-in User Service (RADIUS) for centralized password management.

Next topics

[How iSCSI authentication works](#) on page 110

[Guidelines for using CHAP authentication](#) on page 111

[Defining an authentication method for an initiator](#) on page 111

[Defining a default authentication method for initiators](#) on page 112

[Displaying initiator authentication methods](#) on page 113

[Removing authentication settings for an initiator](#) on page 113

[iSCSI RADIUS configuration](#) on page 113

How iSCSI authentication works

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

The authentication methods are:

- Challenge Handshake Authentication Protocol (CHAP)—The initiator logs in using a CHAP user name and password.
You can specify a CHAP password or generate a random password. There are two types of CHAP user names and passwords:
 - Inbound—The storage system authenticates the initiator.
Inbound settings are required if you are using CHAP authentication without RADIUS.
 - Outbound—This is an optional setting to enable the initiator to authenticate the storage system.
You can use outbound settings only if you defined an inbound user name and password on the storage system.
- deny—The initiator is denied access to the storage system.
- none—The storage system does not require authentication for the initiator.

You can define a list of initiators and their authentication methods. You can also define a default authentication method that applies to initiators that are not on this list.

The default iSCSI authentication method is `none`, which means any initiator not in the authentication list can log into the storage system without authentication. However, you can change the default method to `deny` or `CHAP`.

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

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

Note: For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *Data ONTAP 7-Mode MultiStore Management Guide*.

Related information

Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas

Guidelines for using CHAP authentication

You should follow these guidelines when using CHAP authentication.

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

Defining an authentication method for an initiator

You can define a list of initiators and their authentication methods. You can also define a default authentication method that applies to initiators that are not on this list.

About this task

You can generate a random password, or you can specify the password you want to use.

Steps

1. To generate a random password, enter the following command:

```
iscsi security generate
```

The storage system generates a 128-bit random password.

2. For each initiator, enter the following command:

```
iscsi security add -i initiator -s [chap | deny | none] [-f radius | -p  
inpassword -n inname] [-o outpassword -m outname]
```

initiator is the initiator name in the iSCSI nodename format.

The `-s` option takes one of several values:

- `chap`—Authenticate using a CHAP user name and password.
- `none`—The initiator can access the storage system without authentication.
- `deny`—The initiator cannot access the storage system.

radius indicates that RADIUS is used for authentication. Use the `-f` option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the `-f` option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

inpassword is the inbound password for CHAP authentication. The storage system uses the inbound password to authenticate the initiator. An inbound password is required if you are using CHAP authentication and you are not using RADIUS.

iname is a user name for inbound CHAP authentication. The storage system uses the inbound user name to authenticate the initiator.

outpassword is a password for outbound CHAP authentication. It is stored locally on the storage system, which uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Note: If you generated a random password, you can use this string for either *inpassword* or *outpassword*. If you enter a string, the storage system interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

Defining a default authentication method for initiators

Use the `iscsi security default` command to define a default authentication method for all initiators not specified with the `iscsi security add` command.

Step

1. On the storage system console, enter the following command:

```
iscsi security default -s [chap | none | deny] [-f radius | -p
inpassword -n inname] [-o outpassword -m outname]
```

The `-s` option takes one of three values:

- `chap`—Authenticate using a CHAP user name and password.
- `none`—The initiator can access the storage system without authentication.
- `deny`—The initiator cannot access the storage system.

radius indicates that RADIUS authentication is used. Use the `-f` option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the `-f` option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

inpassword is the inbound password for CHAP authentication. The storage system uses the inbound password to authenticate the initiator.

iname is a user name for inbound CHAP authentication. The storage system uses the inbound user name to authenticate the initiator.

outpassword is a password for outbound CHAP authentication. The storage system uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Displaying initiator authentication methods

Use the `iscsi security show` command to view a list of initiators and their authentication methods.

Step

1. On the storage system console, enter the following command:

```
iscsi security show
```

Removing authentication settings for an initiator

Use the `iscsi security delete` command to remove the authentication settings for an initiator and use the default authentication method.

Step

1. On the storage system console, enter the following command:

```
iscsi security delete -i initiator
```

-i initiator is the initiator name in the iSCSI node name format.

The initiator is removed from the authentication list and logs in to the storage system using the default authentication method.

iSCSI RADIUS configuration

You can configure your storage systems to use RADIUS for centrally managing iSCSI initiator authentication.

RADIUS uses CHAP to authenticate iSCSI initiators, but it enables you to manage the authentication process from a central RADIUS server, rather than manage it manually on each storage system. In larger SAN environments, this can greatly simplify iSCSI initiator management, CHAP password management, and provide added security.

RADIUS also reduces the load on your storage system because most of the authentication processing is handled by the RADIUS server.

Follow these steps to configure your storage systems for iSCSI RADIUS:

1. Define RADIUS as the authentication method for each initiator.
2. Start the RADIUS client service.
3. Add the RADIUS server.
4. Enable the storage system to use RADIUS for CHAP authentication.

Next topics

[Defining RADIUS as the authentication method for initiators](#) on page 114

[Starting the RADIUS client service](#) on page 115

[Adding a RADIUS server](#) on page 116

[Enabling the storage system to use RADIUS for CHAP authentication](#) on page 116

[Displaying the RADIUS service status](#) on page 117

[Stopping the RADIUS client service](#) on page 117

[Removing a RADIUS server](#) on page 118

[Displaying and clearing RADIUS statistics](#) on page 118

Defining RADIUS as the authentication method for initiators

You can define RADIUS as the authentication method for one or more initiators, as well as make it the default authentication method that applies to initiators that are not on this list.

You can generate a random password, or you can specify the password you want to use. Inbound passwords are saved on the RADIUS server and outbound passwords are saved on the storage system.

Steps

1. To generate a random password, enter the following command:

```
iscsi security generate
```

The storage system generates a 128-bit random password.

Note: If you generate a random inbound password, you must add this password to the RADIUS server.

2. For each initiator, enter the following command:

```
iscsi security add -i initiator -s chap -f radius [-o outpassword -m outname]
```

initiator is the initiator name in the iSCSI nodename format.

Use the `-f` option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the `-f` option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

outpassword is a password for outbound CHAP authentication. It is stored locally on the storage system, which uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Note: If you generated a random password, you can use this string for *outpassword*. If you enter a string, the storage system interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

3. To define RADIUS as the default authentication method for all initiators not previously specified, enter the following command:

```
iscsi security default -s chap -f radius [-o outpassword -m outname]
```

Examples

```
system1> iscsi security add -i iqn.1992-08.com.microsoft:system1 -s
chap -f radius
system1> iscsi security show
Default sec is CHAP RADIUS Outbound password: **** Outbound username:
init: iqn.1994-05.com.redhat:10ca21e21b75 auth: CHAP RADIUS Outbound
password: **** Outbound username: icroto
```

```
system1> iscsi security default -s chap -f radius
```

After enabling RADIUS authentication for the initiators, start the RADIUS client service on the storage system.

Starting the RADIUS client service

Once you enable RADIUS authentication for the appropriate initiators, you must start the RADIUS client.

Step

1. Enter the following command:

```
radius start
```

Example

```
system1> radius start
RADIUS client service started
```

After the RADIUS service is started, ensure you add one or more RADIUS servers with which the storage system can communicate.

Adding a RADIUS server

After you start the RADIUS client service, add a RADIUS server with which the storage system can communicate. You can add up to three RADIUS servers.

Step

1. Enter the following command:

```
radius add [-d] RADIUS_server_hostname or ip_address [-p port_number]
```

You can use the `-d` option to make the RADIUS server you are adding the default server. If there is no default server, the one you add becomes the default.

You can use the `-p` option to specify a port number on the RADIUS server. The default port number is 1812.

Example

```
system1> radius add 10.60.155.58 -p 1812
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58  UDPPort=1812
```

After adding the necessary servers, you must enable the storage system to use the RADIUS server for CHAP authentication.

Enabling the storage system to use RADIUS for CHAP authentication

Once RADIUS authentication is enabled for the initiators and the RADIUS client service is started, you must set the `iscsi.auth.radius.enable` option to on. This ensures the storage system uses RADIUS for CHAP authentication.

This option is set to off by default, and you must set it to on, regardless of whether you used the `-f` option when enabling RADIUS for the initiators.

Step

1. Enter the following command:

```
options iscsi.auth.radius.enable on
```

Your system is now enabled for RADIUS authentication.

```
system1> options iscsi.auth.radius.enable on
system1> options iscsi
iscsi.auth.radius.enable      on
iscsi.enable                  on
iscsi.isns.rev                22
iscsi.max_connections_per_session use_system_default
```

```
iscsi.max_error_recovery_level use_system_default
iscsi.max_ios_per_session      128
iscsi.tcp_window_size          131400
```

Displaying the RADIUS service status

You can use the `radius show` command to display important RADIUS information, including whether the service is running and the default RADIUS server.

Step

1. Enter the following command:

```
radius show
```

Example

```
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58  UDPPort=1812
```

You can also run the `radius status` command to see if the client service is running.

Example

```
system1> radius status
RADIUS client service is running
```

Stopping the RADIUS client service

You can use the `radius stop` command to stop the RADIUS client service.

Step

1. Enter the following command:

```
radius stop
```

```
system1> radius stop
RADIUS client service stopped
```

Removing a RADIUS server

You can use the `radius remove` command to ensure a RADIUS server is no longer used for RADIUS authentication.

Step

1. Enter the following command:

```
radius remove RADIUS_server_hostname or ip_address[-p port_number]
```

If the server is using a port other than 1812, use the `-p` option to specify the port number.

```
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58 UDPPort=1812

system1> radius remove 10.60.155.58
system1> radius show
RADIUS client service is running
```

Displaying and clearing RADIUS statistics

You can use the `radius stats` command to display important details about the RADIUS service, including packets accepted, packets rejected, and the number of authentication requests. You can also clear the existing statistics.

Step

1. Enter the following command:

```
radius stats [-z]
```

You can use the `-z` option to clear the statistics.

```
system1> radius stats
RADIUS client statistics
  RADIUS access-accepted-packets:    121
  RADIUS access-challenged-packets:   3
  RADIUS access-rejected-packets:     0
  RADIUS authentication-requests:    124
  RADIUS denied-packets:              0
  RADIUS late-packets:                0
```

```
RADIUS retransmitted-packets:    14
RADIUS short-packets:           0
RADIUS timed-out-packets:       0
RADIUS unknown-packets:         0
RADIUS unknown-server-packets:  0
```

```
system1> radius stats -z
```

```
system1> radius stats
RADIUS client statistics
RADIUS access-accepted-packets:    0
RADIUS access-challenged-packets:  0
RADIUS access-rejected-packets:    0
RADIUS authentication-requests:    0
RADIUS denied-packets:              0
RADIUS late-packets:                0
RADIUS retransmitted-packets:       0
RADIUS short-packets:               0
RADIUS timed-out-packets:           0
RADIUS unknown-packets:             0
RADIUS unknown-server-packets:      0
```

Target portal group management

A target portal group is a set of one or more storage system network interfaces that can be used for an iSCSI session between an initiator and a target. A target portal group is identified by a name and a numeric tag. If you want to have multiple connections per session across more than one interface for performance and reliability reasons, then you must use target portal groups.

Note: If you are using MultiStore, you can also configure non-default vFiler units for target portal group management based on IP address.

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

Starting with Data ONTAP, you can explicitly create target portal groups and assign tag values. If you want to increase performance and reliability by using multi-connections per session across more than one interface, you must create one or more target portal groups.

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

When an interface is added to the storage system, each network interface is automatically assigned to its own target portal group.

In addition, some storage systems support the use of an iSCSI Target expansion adapter, which contains special network interfaces that offload part of the iSCSI protocol processing. You cannot

combine these iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same target portal group.

Next topics

[Range of values for target portal group tags](#) on page 120

[Important cautions for using target portal groups](#) on page 120

[Displaying target portal groups](#) on page 121

[Creating target portal groups](#) on page 121

[Destroying target portal groups](#) on page 122

[Adding interfaces to target portal groups](#) on page 122

[Removing interfaces from target portal groups](#) on page 123

[Configuring iSCSI target portal groups](#) on page 123

Range of values for target portal group tags

If you create target portal groups, the valid values you can assign to target portal group tags vary depending on the type of interface.

The following table shows the ranges values for target portal group tags:

Target portal group type	Allowed values
User defined groups	1 - 256
Default groups for physical devices	1000 - 1511
Default groups for VLANs and VIFs	2000 - 2511
Default groups for IP-based vFiler units	4000 - 65535

Important cautions for using target portal groups

You must heed these important cautions when using target portal groups.

- Some initiators, including those used with Windows, HP-UX, Solaris, and Linux, create a persistent association between the target portal group tag value and the target. If the target portal group tag changes, the LUNs from that target will be unavailable.
- Adding or removing a NIC might change the target portal group assignments. Be sure to verify the target portal group settings are correct after adding or removing hardware, especially in HA pairs.
- When used with multi-connection sessions, the Windows iSCSI software initiator creates a persistent association between the target portal group tag value and the target interfaces. If the tag value changes while an iSCSI session is active, the initiator will be able to recover only one connection for a session. To recover the remaining connections, you must refresh the initiator's target information.

Displaying target portal groups

Use the `iscsi tpgroup show` command to display a list of existing target portal groups.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup show
```

Example

```
iscsi tpgroup show
TPGTag  Name                Member Interfaces
1000    e0_default            e0
1001    e5a_default           e5a
1002    e5b_default           e5b
1003    e9a_default           e9a
1004    e9b_default           e9b
```

Creating target portal groups

If you want to employ multi-connection sessions to improve performance and reliability, you must use target portal groups to define the interfaces available for each iSCSI session.

About this task

Create a target portal group that contains all of the interfaces you want to use for one iSCSI session. However, note that you cannot combine iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same target portal group.

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

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup create [-f] tpgroup_name [-t tag] [interface ...]
```

`-f` forces the new group to be created, even if that terminates an existing session using one of the interfaces being added to the group.

`tpgroup_name` is the name of the group being created (1 to 60 characters, no spaces or non-printing characters).

`-t tag` sets the target portal group tag to the specified value. In general you should accept the default tag value. User-specified tags must be in the range 1 to 256.

`interface ...` is the list of interfaces to include in the group, separated by spaces.

Example

The following command creates a target portal group named `server_group` that includes interfaces `e8a` and `e9a`:

```
iscsi tpgroup create server_group e8a e9a
```

Destroying target portal groups

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

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup destroy [-f] tpgroup_name
```

`-f` forces the group to be destroyed, even if that terminates an existing session using one of the interfaces in the group.

`tpgroup_name` is the name of the group being destroyed.

Adding interfaces to target portal groups

You can add interfaces to an existing target portal group. The specified interfaces are removed from their current groups and added to the new group.

About this task

Any iSCSI sessions using the specified interfaces are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group are not able to reconnect.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup add [-f] tpgroup_name [interface ...]
```

`-f` forces the interfaces to be added, even if that terminates an existing session using one of the interfaces being added to the group.

`tpgroup_name` is the name of the group.

interface ... is the list of interfaces to add to the group, separated by spaces.

Example

The following command adds interfaces e8a and e9a to the portal group named `server_group`:

```
iscsi tpgroup add server_group e8a e9a
```

Removing interfaces from target portal groups

You can remove interfaces from an existing target portal group. The specified interfaces are removed from the group and returned to their individual default target portal groups.

About this task

Any iSCSI sessions with the interfaces being removed are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group are not able to reconnect.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup remove [-f] tpgroup_name [interface ...]
```

`-f` forces the interfaces to be removed, even if that terminates an existing session using one of the interfaces being removed from the group.

tpgroup_name is the name of the group.

interface ... is the list of interfaces to remove from the group, separated by spaces.

Example

The following command removes interfaces e8a and e9a from the portal group named `server_group`, even though there is an iSCSI session currently using e8a:

```
iscsi tpgroup remove -f server_group e8a e9a
```

Configuring iSCSI target portal groups

When you enable ALUA, you can set the priority of your target portal groups for iSCSI to optimized or non-optimized. The optimized path becomes the preferred path and the non-optimized path becomes the secondary path.

About this task

When you first enable ALUA, all target portal groups are set to optimized by default.

Some storage systems support the use of an iSCSI Target HBA, which contains special network interfaces that offload part of the iSCSI protocol processing. You might want to set the target portal

groups that contain these iSCSI hardware-accelerated interfaces to optimized and the standard iSCSI storage system interfaces to non-optimized. As a result, the host uses the iSCSI hardware-accelerated interface as the primary path.

Attention: When setting the path priority for target portal groups on clustered storage systems, make sure that the path priority setting is identical for the target portal group on the primary storage system and the target portal group on its partner interface on the secondary storage system.

To change the path priority to a target portal group, complete the following step.

Step

1. Enter the following command:

```
iscsi tpgroup alua set target_portal_group_name [optimized | non-optimized]
```

Example

```
iscsi tpgroup alua set tpgroup1 non-optimized
```

Related concepts

[What ALUA is](#) on page 24

Related tasks

[Enabling ALUA](#) on page 88

Displaying iSCSI statistics

Use the `iscsi stats` command to display important iSCSI statistics.

Step

1. On the storage system console, enter the following command:

```
iscsi stats [-a | -z | ipv4 | ipv6]
```

`-a` displays the combined IPv4 and IPv6 statistics followed by the individual statistics for IPv4 and IPv6.

`-z` resets the iSCSI statistics.

`ipv4` displays only the IPv4 statistics.

`ipv6` displays only the IPv6 statistics.

Entering the `iscsi stats` command without any options displays only the combined IPv4 and IPv6 statistics.

```
system1> iscsi stats -a
```

iSCSI stats(total)

iSCSI PDUs Received

SCSI-Cmd:	1465619		Nop-Out:	4		SCSI
TaskMgtCmd:	0					
LoginReq:	6		LogoutReq:	1		Text
Req:	1					
DataOut:	0		SNACK:	0		
Unknown:	0					
Total:	1465631					

iSCSI PDUs Transmitted

SCSI-Rsp:	733684		Nop-In:	4		SCSI
TaskMgtRsp:	0					
LoginRsp:	6		LogoutRsp:	1		
TextRsp:	1					
Data_In:	790518		R2T:	0		
Asyncmsg:	0					
Reject:	0					
Total:	1524214					

iSCSI CDBs

DataIn Blocks:	5855367		DataOut Blocks:	0
Error Status:	1		Success Status:	1465618
Total CDBs:	1465619			

iSCSI ERRORS

Failed Logins:	0		Failed TaskMgt:	0
Failed Logouts:	0		Failed TextCmd:	0
Protocol:	0			
Digest:	0			
PDU discards (outside CmdSN window):	0			
PDU discards (invalid header):	0			
Total:	0			

iSCSI Stats(ipv4)

iSCSI PDUs Received

SCSI-Cmd:	732789		Nop-Out:	1		SCSI
TaskMgtCmd:	0					
LoginReq:	2		LogoutReq:	0		Text
Req:	0					
DataOut:	0		SNACK:	0		
Unknown:	0					
Total:	732792					

iSCSI PDUs Transmitted

SCSI-Rsp:	366488		Nop-In:	1		SCSI
TaskMgtRsp:	0					
LoginRsp:	2		LogoutRsp:	0		
TextRsp:	0					
Data_In:	395558		R2T:	0		
Asyncmsg:	0					
Reject:	0					
Total:	762049					

iSCSI CDBs

DataIn Blocks:	2930408		DataOut Blocks:	0
Error Status:	0		Success Status:	732789
Total CDBs:	732789			

```

iSCSI ERRORS
  Failed Logins:          0 | Failed TaskMgt:          0
  Failed Logouts:         0 | Failed TextCmd:           0
  Protocol:               0
  Digest:                 0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header):      0
  Total: 0

iSCSI Stats(ipv6)
iSCSI PDUs Received
  SCSI-Cmd:      732830 | Nop-Out:      3 | SCSI
  TaskMgtCmd:     0
  LoginReq:       4 | LogoutReq:    1 | Text
  Req:            1
  DataOut:        0 | SNACK:        0 |
  Unknown:        0
  Total: 732839
iSCSI PDUs Transmitted
  SCSI-Rsp:      367196 | Nop-In:      3 | SCSI
  TaskMgtRsp:     0
  LoginRsp:       4 | LogoutRsp:    1 |
  TextRsp:        1
  Data_In:       394960 | R2T:         0 |
  Asyncmsg:       0
  Reject:         0
  Total: 762165
iSCSI CDBs
  DataIn Blocks:  2924959 | DataOut Blocks:      0
  Error Status:   1 | Success Status:    732829
  Total CDBs: 732830
iSCSI ERRORS
  Failed Logins:          0 | Failed TaskMgt:          0
  Failed Logouts:         0 | Failed TextCmd:           0
  Protocol:               0
  Digest:                 0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header):      0
  Total: 0

```

Definitions for iSCSI statistics

The following tables define the iSCSI statistics that are displayed when you run the `iscsi stats` command. For vFile contexts, the statistics displayed refer to the entire storage system, not the individual vFile units.

iSCSI PDUs received

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

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

iSCSI PDUs transmitted

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

Field	Description
SCSI-Rsp	SCSI response messages.
LoginRsp	Responses to login requests during session setup.
DataIn	Messages containing data requested by SCSI read operations.
Nop-In	Responses to initiator Nop-Out messages.
Logout-Rsp	Responses to Logout-Req messages.
R2T	Ready to transfer messages indicating that the target is ready to receive data during a SCSI write operation.
SCSI TaskMgtRsp	Responses to task management requests.
TextRsp	Responses to Text-Req messages.
Asyncmsg	Messages the target sends to asynchronously notify the initiator of an event, such as the termination of a session.
Reject	Messages the target sends to report an error condition to the initiator, for example:

Field	Description
	<ul style="list-style-type: none"> • Data Digest Error (checksum failed) • Target does not support command sent by the initiator • Initiator sent a command PDU with an invalid PDU field

iSCSI CDBs

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

iSCSI Errors

This section lists login failures and other SCSI protocol errors.

Displaying iSCSI session information

Use the `iscsi session show` command to display iSCSI session information, such as TCP connection information and iSCSI session parameters.

About this task

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

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

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

Step

1. On the storage system console, enter the following command:

```
iscsi session show [-v | -t | -p | -c] [session_tsih ...]
```

`-v` displays all information and is equivalent to `-t -p -c`.

`-t` displays the TCP connection information for each session.

-p displays the iSCSI session parameters for each session.

-c displays the iSCSI commands in progress for each session.

session_tsih is a list of session identifiers, separated by spaces.

```
system1> iscsi session show -t
Session 2
  Initiator Information
    Initiator Name: iqn.1991-05.com.microsoft:legbreak
    ISID: 40:00:01:37:00:00
  Connection Information
  Connection 1
    Remote Endpoint: fe80::211:43ff:fece:ccce:1135
    Local Endpoint: fe80::2a0:98ff:fe00:fd81:3260
    Local Interface: e0
    TCP recv window size: 132480
  Connection 2
    Remote Endpoint: 10.60.155.31:2280
    Local Endpoint: 10.60.155.105:3260
    Local Interface: e0
    TCP recv window size: 131400
```

Displaying iSCSI connection information

Use the `iscsi connection show` command to display iscsi connection parameters.

Step

1. On the storage system console, enter the following command:

```
iscsi connection show [-v] [{new | session_tsih} conn_id]
```

-v displays all connection information.

newconn_id displays information about a single connection that is not yet associated with a session identifier. You must specify both the keyword `new` and the connection identifier.

session_tsih conn_id displays information about a single connection. You must specify both the session identifier and the connection identifier.

Example

The following example shows the -v option.

```
system1> iscsi connection show -v
No new connections
Session connections
Connection 2/1:
  State: Full_Feature_Phase
  Remote Endpoint: fe80::211:43ff:fece:ccce:1135
```

```

Local Endpoint: fe80::2a0:98ff:fe00:fd81:3260
Local Interface: e0
Connection 2/2:
State: Full_Feature_Phase
Remote Endpoint: 10.60.155.31:2280
Local Endpoint: 10.60.155.105:3260
Local Interface: e0

```

Guidelines for using iSCSI with HA pairs

To ensure that the partner storage system successfully takes over during a failure, you need to make sure that the two systems and the TCP/IP network are correctly configured.

Of special concern are the target portal group tags configured on the two storage systems.

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

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

Next topics


[Simple HA pairs with iSCSI](#) on page 130

[Complex HA pairs with iSCSI](#) on page 132

Simple HA pairs with iSCSI

The following example describes how to implement the best practices for using iSCSI with HA pairs.

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



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

In the HA pair, interface e9a on Storage System A is the partner of e9a on Storage System B. Likewise, e9b on System A is the partner of e9b on system B. For more information on configuring interfaces for an HA pair, see the *Data ONTAP 7-Mode High-Availability Configuration Guide*.

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

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

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

Related information

Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas

Complex HA pairs with iSCSI

If your cluster has a more complex networking configuration, including VIFs and VLANs, follow the same best practice of making the configurations identical.

For example, if you have a vif on storage system A, create the same vif on storage system B. Make sure the target portal group tag assigned to each vif is the same. The name of the target portal group does not have to be the same; only the tag value matters.

iSCSI problem resolution

This section contains tips for resolving common problems that occur with iSCSI networks.

Next topics

LUNs not visible on the host on page 132

System cannot register with iSNS server on page 134

No multi-connection session on page 134

Sessions constantly connecting and disconnecting during takeover on page 134

Resolving iSCSI error messages on the storage system on page 135

LUNs not visible on the host

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

Configuration setting	What to do
Cabling	Verify that the cables between the host and the storage system are properly connected.
Network connectivity	<p>Verify that there is TCP/IP connectivity between the host and the storage system.</p> <ul style="list-style-type: none"> From the storage system command line, ping the host interfaces that are being used for iSCSI. From the host command line, ping the storage system interfaces that are being used for iSCSI.
System requirements	<p>Verify that the components of your configuration are qualified. Verify that you have the correct host operating system (OS) service pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up to date system requirements in the N series Service and Support Web site at www.ibm.com/storage/support/nas/.</p>

Configuration setting	What to do
Jumbo frames	If you are using jumbo frames in your configuration, ensure that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system.
Initiator login	Verify that the initiator is logged in to the storage system. If the command output shows no initiators are logged in, check the initiator configuration on the host. Verify that the storage system is configured as a target of the initiator.
iSCSI node names	Verify that you are using the correct initiator node names in the igroup configuration. For the storage system, see “Managing igroups” on page 94. On the host, use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host must match.
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system console, use one of the following commands: <ul style="list-style-type: none"> • <code>lun show -m</code> Displays all LUNs and the igroups to which they are mapped. • <code>lun show -g igroup-name</code> Displays the LUNs mapped to a specific igroup. Or, using FilerView, Click LUNs > Manage—Displays all LUNs and the igroups to which they are mapped.

Related concepts

[igroup management](#) on page 83

[About LUNs, igroups, and LUN maps](#) on page 55

Related tasks

[Verifying that the iSCSI service is running](#) on page 97

[Displaying initiators connected to the storage system](#) on page 109

System cannot register with iSNS server

Different iSNS server versions follow different draft levels of the iSNS specification.

If there is a mismatch between the iSNS draft version used by the storage system and by the iSNS server, the storage system cannot register.

Related concepts

[About iSNS service version incompatibility](#) on page 106

No multi-connection session

All of the connections in a multi-connection iSCSI session must go to interfaces on the storage system that are in the same target portal group.

If an initiator is unable to establish a multi-connection session, check the portal group assignments of the initiator.

If an initiator can establish a multi-connection session, but not during a cluster failover (CFO), the target portal group assignment on the partner storage system is probably different from the target portal group assignment on the primary storage system.

Related concepts

[Target portal group management](#) on page 119

[Guidelines for using iSCSI with HA pairs](#) on page 130

Sessions constantly connecting and disconnecting during takeover

An iSCSI initiator that uses multipath I/O will constantly connect and disconnect from the target during cluster failover if the target portal group is not correctly configured.

The interfaces on the partner storage system must have the same target portal group tags as the interfaces on the primary storage system.

Related concepts

[Guidelines for using iSCSI with HA pairs](#) on page 130

Resolving iSCSI error messages on the storage system

There are a number of common iSCSI-related error messages that might display on your storage system console. The following table contains the most common error messages, and instructions for resolving them.

Message	Explanation	What to do
ISCSI: network interface <i>identifier</i> disabled for use; incoming connection discarded	The iSCSI service is not enabled on the interface.	Use the <code>iscsi</code> command to enable the iSCSI service on the interface. For example: <code>iscsi interface enable e9b</code>
ISCSI: Authentication failed for initiator <i>nodename</i>	CHAP is not configured correctly for the specified initiator.	Check CHAP settings. <ul style="list-style-type: none"> • Inbound credentials on the storage system must match outbound credentials on the initiator. • Outbound credentials on the storage system must match inbound credentials on the initiator. • You cannot use the same user name and password for inbound and outbound settings on the storage system.
ifconfig: <i>interface</i> cannot be configured: Address does not match any partner interface. or Cluster monitor: takeover during ifconfig_2 failed; takeover continuing...	A single-mode VIF can be a partner interface to a standalone, physical interface on a cluster partner. However, the partner statement in the <code>ifconfig</code> command must use the name of the partner interface, not the partner's IP address. If the IP address of the partner's physical interface is used, the interface will not be successfully taken over by the storage system's VIF interface.	<ol style="list-style-type: none"> 1. Add the partner's interface using the <code>ifconfig</code> command on each system in the HA pair. For example: <pre>system1> ifconfig vif0 partner e0a system2> ifconfig e0a partner vif0</pre> 2. Modify the <code>/etc/rc</code> file on both systems to contain the same interface information.

Related concepts

[Guidelines for using CHAP authentication](#) on page 111

FC SAN management

This section contains critical information required to successfully manage your FC SAN.

Next topics

How to manage FC with HA pairs on page 137

How to use port sets to make LUNs available on specific FC target ports on page 140

FC service management on page 145

Managing systems with onboard Fibre Channel adapters on page 158

How to manage FC with HA pairs

Data ONTAP provides important functionality that allows your system to continue running smoothly when one of the devices in your HA pairs fail. This section provides an overview of how this functionality works.

See the *Fibre Channel and iSCSI Configuration Guide* for additional configuration details.

Related information

Fibre Channel and iSCSI Configuration Guide - www.ibm.com/storage/support/nas


How controller failover works

Data ONTAP is equipped with functionality called controller failover that allows the partner system to assume responsibility for the failed system's LUNs without interruption.

A storage system with an HA pair has a single global WWNN, and both systems in the HA pair function as a single FC node. Each node in the HA pair shares the partner node's LUN map information.

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

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

**Next topics**

[How Data ONTAP avoids igroup mapping conflicts during cluster failover](#) on page 138

[Multipathing requirements for controller failover](#) on page 140

How Data ONTAP avoids igroup mapping conflicts during cluster failover

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

Next topics

[igroup ostype conflicts](#) on page 139

Reserved LUN ID ranges on page 139

Bringing LUNs online on page 139

When to override possible mapping conflicts on page 139

Related tasks

Checking LUN, igroup, and FC settings on page 76

igroup ostype conflicts

When you add an initiator WWPN to an igroup, Data ONTAP verifies that there are no igroup ostype conflicts.

An ostype conflict occurs, for example, when an initiator with the WWPN 10:00:00:00:c9:2b:cc:39 is a member of an AIX igroup on one node in the HA pair and the same WWPN is also a member of an group with the default ostype on the partner.

Reserved LUN ID ranges

By reserving LUN ID ranges on each storage system, Data ONTAP provides a mechanism for avoiding mapping conflicts.

If the cluster interconnect is down, and you try to map a LUN to a specific ID, the `lun map` command fails. If you do not specify an ID in the `lun map` command, Data ONTAP automatically assigns one from a reserved range.

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

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

Bringing LUNs online

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

When to override possible mapping conflicts

When the cluster interconnect is down, Data ONTAP cannot check for LUN mapping or igroup ostype conflicts.

The following commands fail unless you use the `-f` option to force these commands. The `-f` option is only available with these commands when the cluster interconnect is down.

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

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

```
lun map -f /vol/vol12/qtree1/lun3 aix_host5_group2 5
```

Multipathing requirements for controller failover

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

How to use port sets to make LUNs available on specific FC target ports

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

If an igroup is not bound to a port set, the LUNs mapped to the igroup are available on all of the storage system's FC target ports. The igroup controls which initiators LUNs are exported to. The port set limits the target ports on which those initiators have access.

You use port sets for LUNs that are accessed by FC hosts only. You cannot use port sets for LUNs accessed by iSCSI hosts.

Next topics

- [How port sets work in HA pairs](#) on page 141
- [How upgrades affect port sets and igroups](#) on page 141
- [How port sets affect igroup throttles](#) on page 141
- [Creating port sets](#) on page 142
- [Binding igroups to port sets](#) on page 142
- [Unbinding igroups from port sets](#) on page 143
- [Adding ports to port sets](#) on page 143
- [Removing ports from port sets](#) on page 144
- [Destroying port sets](#) on page 144
- [Displaying the ports in a port set](#) on page 145
- [Displaying igroup-to-port-set bindings](#) on page 145

How port sets work in HA pairs

All ports on both systems in the HA pairs are visible to the hosts. You use port sets to fine-tune which ports are available to specific hosts and limit the amount of paths to the LUNs to comply with the limitations of your multipathing software.

When using port sets, make sure your port set definitions and igroup bindings align with the cabling and zoning requirements of your configuration. See the *Fibre Channel and iSCSI Configuration Guide* for additional configuration details.

Related concepts

[How controller failover works](#) on page 137

Related information

[Fibre Channel and iSCSI Configuration Guide - *www.ibm.com/storage/support/nas*](#)

How upgrades affect port sets and igroups

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

How port sets affect igroup throttles

Port sets enable you to control queue resources on a per-port basis.

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

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

Example: port sets and igroup throttles

igroup_1 is bound to portset_1, which includes ports 4a and 4b on each system in the HA pair (SystemA:4a, SystemA:4b, SystemB:4a, SystemB:4b). The throttle setting of igroup is 40 percent.

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

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

Creating port sets

Use the `portset create` command to create portsets for FCP.

About this task

For HA pairs, when you add local ports to a port set, also add the partner system's corresponding target ports to the same port set.

For example, if you have local systems's target port 4a port in the port set, then make sure to include the partner system's port 4a in the port set as well. This ensures that the takeover and giveback occurs without connectivity problems.

Step

1. Enter the following command:

```
portset create -f portset_name [port...]
```

`-f` creates an FCP port set.

`portset_name` is the name you specify for the port set. You can specify a string of up to 95 characters.

`port` is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty port set. You can add as many as 18 target FCP ports.

You specify a port by using the following formats:

- `slotletter` is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.
- `filename:slotletter` adds only a specific port on a storage system—for example, SystemA:4b.

Binding igroups to port sets

Once you create a port set, you must bind the port set to an igroup so the host knows which FC ports to access.

About this task

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

Note: You cannot bind an igroup to an empty port set, as the initiators in the igroup would have no ports by which to access the LUN.

Step

1. Enter the following command:

```
igroup bind igroup_name portset_name
```

Example

```
igroup bind aix-igroup1 portset4
```

Unbinding igroups from port sets

You can use the `igroup unbind` command to unbind an igroup from a port set.

About this task

If you unbind or unmap an igroup from a port set, then all the host initiator ports in the igroup can access LUNs on all target ports.

Step

1. Enter the following command:

```
igroup unbind igroup_name
```

Example

```
igroup unbind aix-igroup1
```

Adding ports to port sets

After you create a port set, you can use the `portset add` command to add ports to the port set.

Step

1. Enter the following command:

```
portset add portset_name [port...]
```

portset_name is the name you specify for the port set. You can specify a string of up to 95 characters.

port is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty port set. You can add as many as 18 target FCP ports.

You specify a port by using the following formats:

- *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.

- *filename:slotletter* adds only a specific port on a storage system—for example, SystemA:4b.

Removing ports from port sets

Once you create a port set, use the `portset remove` command to remove ports from the portset.

About this task

Note that you cannot remove the last port in the port set if the port set is bound to an igroup. To remove the last port, first unbind the port set from the igroup, then remove the port.

Step

1. Enter the following command:

```
portset remove portset_name [port...]
```

portset_name is the name you specify for the port set. You can specify a string of up to 95 characters.

port is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty port set. You can add as many as 18 target FCP ports.

You specify a port by using the following formats:

- *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.
- *filename:slotletter* adds only a specific port on a storage system—for example, SystemA:4b.

Destroying port sets

Use the `portset destroy` command to delete a port set.

Steps

1. Unbind the port set from any igroups by entering the following command:

```
igroup unbind igroup_name portset_name
```

2. Enter the following command:

```
portset destroy [-f] portset_name...
```

You can specify a list of port sets.

If you use the `-f` option, you destroy the port set even if it is still bound to an igroup.

If you do not use the `-f` option and the port set is still bound to an igroup, the `portset destroy` command fails.

Example

```
portset destroy portset1 portset2 portset3
```

Displaying the ports in a port set

Use the `portset show` command to display all ports belonging to a particular port set.

Step

1. Enter the following command:

```
portset show portset_name
```

If you do not supply `portset_name`, all port sets and their respective ports are listed. If you supply `portset_name`, only the ports in the port set are listed.

Example

```
portset show portset1
```

Displaying igroup-to-port-set bindings

Use the `igroup show` command to display which igroups are bound to port sets.

Step

1. Enter the following command:

```
igroup show igroup_name
```

Example

```
igroup show aix-igroup1
```

FC service management

Use the `fc` commands for most of the tasks involved in managing the Fibre Channel service and the target and initiator adapters.

Enter `fc help` at the command line to display the list of available commands.

Next topics

[*Verifying that the FC service is running*](#) on page 146

[*Verifying that the FC service is licensed*](#) on page 146

[*Licensing the FC service*](#) on page 146

[*Disabling the FC license*](#) on page 147

[*Starting and stopping the FC service*](#) on page 147

[*Taking target expansion adapters offline and bringing them online*](#) on page 148

[*Changing the adapter speed*](#) on page 148

[How WWPN assignments work with FC target expansion adapters](#) on page 150

[Changing the system's WWNN](#) on page 152

[WWPN aliases](#) on page 153

[Obtaining fabric zone server data](#) on page 155

[Obtaining a physical topology of the FC fabric](#) on page 156

[Obtaining fabric nameserver data](#) on page 156

[Checking connectivity of the initiators](#) on page 157

Verifying that the FC service is running

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

Step

1. Enter the following command:

```
fcp status
```

A message is displayed indicating whether FC service is running.

Note: If the FC service is not running, verify that FC is licensed on the system.

Related tasks

[Licensing the FC service](#) on page 146

Verifying that the FC service is licensed

If you cannot start the FC service, verify that the service is licensed on the system.

Step

1. Enter the following command:

```
license
```

A list of all available services displays, and those services that are enabled show the license code; those that are not enabled are indicated as *not licensed*.

Licensing the FC service

The FC service must be licensed on the system before you can run the service on that system.

Step

1. Enter the following command:

```
license add license_code
```

license_code is the license code you received when you purchased the FC license.

Related concepts

Managing systems with onboard Fibre Channel adapters on page 158

Disabling the FC license

Use the `license delete` command to disable the FC license.

Step

1. Enter the following command:

```
license delete service
```

service is any service you can license.

Example

```
license delete fcp
```

Starting and stopping the FC service

Once the FC service is licensed, you can start and stop the service.

About this task

Stopping the FC service disables all FC ports on the system, which has important ramifications for HA pairs during cluster failover. For example, if you stop the FC service on System1, and System2 fails over, System1 will be unable to service System2's LUNs.

On the other hand, if System2 fails over, and you stop the FC service on System2 and start the FC service on System1, System1 will successfully service System2's LUNs.

Use the `partner fcp stop` command to disable the FC ports on the failed system during takeover, and use the `partner fcp start` command to re-enable the FC service after the giveback is complete.

Step

1. Enter the following command:

```
fcp [start|stop]
```

Example

```
fcp start
```

The FC service is enabled on all FC ports on the system. If you enter `fcp stop`, the FC service is disabled on all FC ports on the system.

Taking target expansion adapters offline and bringing them online

Use the `fcv config` command to take a target expansion adapter offline and to bring it back online.

Step

1. Enter the following command:

```
fcv config adapter [up|down]
```

Example

```
fcv config 4a down
```

The target adapter 4a is offline. If you enter `fcv config 4a up`, the adapter is brought online.

Changing the adapter speed

You can use the `fcv config` command to change the FC adapter speed.

About this task

The available speeds are dependent on the HBA being used. The following are list of the supported speeds available to the controllers:

- Autonegotiate (default)
- 1 Gb
- 2 Gb
- 4 Gb
- 8 Gb

Steps

1. Set the adapter to down using the following command:

```
fcv config adapter down
```

Example

```
: system1> fcv config 2a down
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineStart:notice]:
: Offlining Fibre Channel target adapter 2a.
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineComplete:notice]: Fibre Channel
: target adapter
: 2a offlined.
```

Adapter 2a is taken down, and the FC service might be temporarily interrupted on the adapter.

2. Enter the following command:

```
fcv config adapter speed [auto|1|2|4|8]
```

Example

```
: system1> fcp config 2a speed 2
```

The speed for adapter 2a is changed to 2.

3. Enter the following command:

```
fcp config adapter up
```

Example

```
: device1> fcp config 2a up
: Wed Jun 15 14:05:04 GMT [device1: scsitarget.ispfct.onlining:notice]:
: Onlining Fibre Channel target adapter 2a.

: device1> fcp config
: 2a:    ONLINE [ADAPTER UP]  Loop  No Fabric
:       host address 0000da
:       portname 50:0a:09:81:96:97:a7:f3  nodename
: 50:0a:09:80:86:97:a7:f3
mediatype auto speed 2Gb
```

Adapter 2a is brought back up and the speed is 2Gb.

After you finish

Although the `fcp config` command displays the current adapter speed setting, it does not necessarily display the actual speed at which the adapter is running. For example, if the speed is set to auto, the actual speed may be 1 Gb, 2 Gb, 4 Gb, and so on.

You can use the `show adapter -v` command to view the following:

- Actual speed at which the adapter is running and examine the Data Link Rate value
- Switchname and port number

```
nixon*> fcp show adapter -v 4a
Slot:                4a
Description:         Fibre Channel Target Adapter 4a (Dual-channel,
QLogic CNA 8112 (8152) rev. 2)
Status:              ONLINE
Host Port Address:   0x98d601
Firmware Rev:        5.3.4
MPI Firmware Rev:    1.38.0
PHY Firmware Rev:    1.7.0
FC VLAN ID:          5
FC Nodename:         50:0a:09:80:87:69:68:5a (500a09808769685a)
FC Portname:         50:0a:09:81:87:69:68:5a (500a09818769685a)
Cacheline Size:      16
FC Packet Size:      2048
SRAM Parity:         Yes
External GBIC:        No
Data Link Rate:    10 GBit
Adapter Type:         Local
Fabric Established:    Yes
Connection Established: PTP
Mediatype:            auto
```

```

Partner Adapter:      None
Standby:              No
Target Port ID:       0x1
Switch Port:          brcddcx_rtp02:214
Physical Link Rate:    10 GBit
Physical Link Status:  LINK UP

```

How WWPN assignments work with FC target expansion adapters

It is important to understand how WWPN assignments work with FC target expansion adapters so that your systems continue to run smoothly in the event of head swaps and upgrades, new adapter installations, and slot changes for existing adapters.

When the FC service is initially licensed and enabled on your storage system, the FC target expansion adapters are assigned WWPNs, which persist through head upgrades and replacements. The assignment information is stored in the system's root volume.

The WWPN is associated with the interface name. For example, a target expansion adapter installed in slot 2 may have the interface name of 2a and a WWPN of 50:0a:09:81:96:97:c3:ac. Since the WWPN assignments are persistent, a WWPN will never be automatically re-used, even if the port is disabled or removed. However, there are some circumstances under which you may need to manually change the WWPN assignments.

The following examples explain how WWPN assignments work under the most common circumstances:

- Swapping or upgrading a head
- Adding a new FC target expansion adapter
- Moving an existing adapter to a different slot

Swapping or upgrading a head

As long as the existing root volume is used in the head swap or upgrade, the same port-to-WWPN mapping applies. For example, port 0a on the replacement head will have the same WWPN as the original head. If the new head has different adapter ports, the new ports are assigned new WWPNs.

Adding new FC target expansion adapters

If you add a new adapter, the new ports are assigned new WWPNs. If you replace an existing adapter, the existing WWPNs are assigned to the replacement adapter.

For example, the following table shows the WWPN assignments if you replace a dual-port adapter with a quad-port adapter.

Original configuration	New configuration	WWPN assignments
2a - 50:0a:09:81:96:97:c3:ac	2a - 50:0a:09:81:96:97:c3:ac	No change
2b - 50:0a:09:83:96:97:c3:ac	2b - 50:0a:09:83:96:97:c3:ac	No change

Original configuration	New configuration	WWPN assignments
	2c - 50:0a:09:82:96:97:c3:ac	New
	2d - 50:0a:09:84:96:97:c3:ac	New

Moving a target expansion adapter to a different slot

If you move an adapter to a new slot, then adapter is assigned new WWPNs.

Original configuration	New configuration	WWPN assignments
2a - 50:0a:09:81:96:97:c3:ac	4a - 50:0a:09:85:96:97:c3:ac	New
2b - 50:0a:09:83:96:97:c3:ac	4b - 50:0a:09:86:96:97:c3:ac	New

Related tasks

[Changing the WWPN for a target adapter](#) on page 151

Changing the WWPN for a target adapter

Data ONTAP automatically sets the WWPNs on your target adapters during initialization. However, there are some circumstances in which you might need to change the WWPN assignments on your target expansion adapters or your onboard adapters.

There are two scenarios that might require you to change the WWPN assignments:

- **Head swap:** after performing a head swap, you might not be able to place the target adapters in their original slots, resulting in different WWPN assignments. In this situation it is important to change the WWPN assignments because many of the hosts will bind to these WWPNs. In addition, the fabric may be zoned by WWPN.
- **Fabric re-organization:** you might want to re-organize the fabric connections without having to physically move the target adapters or modify your cabling.

In some cases, you will need to set the new WWPN on a single adapter. In other cases, it will be easier to swap the WWPNs between two adapters, rather than individually set the WWPNs on both adapters.

Steps

1. Take the adapter offline by entering the following command:

```
fcg config adapter down
```

Example

```
fcg config 4a down
```

Note: If you are swapping WWPNs between two adapters, make sure that you take both adapters offline first.

2. Display the existing WWPNs by entering the following command:

```
fcpx portname show [-v]
```

If you do not use the `-v` option, all currently used WWPNs and their associated adapters are displayed. If you use the `-v` option, all other valid WWPNs that are not being used are also shown.

3. Set the new WWPN for a single adapter or swap WWPNs between two adapters.

Note: If you do not use the `-f` option, initiators might fail to reconnect to this adapter if the WWPN is changed. If you use the `-f` option, it overrides the warning message of changing the WWPNs.

If you want to...	Then...
Set the WWPN on a single adapter	Enter the following command: <code>fcpx portname set [-f] adapter wwpn</code>
Swap WWPNs between two adapters.	Enter the following command: <code>fcpx portname swap [-f] adapter1 adapter2</code>

Example

```
fcpx portname set -f 1b 50:0a:09:85:87:09:68:ad
```

Example

```
fcpx portname swap -f 1a 1b
```

4. Bring the adapter back online by entering the following command:

```
fcpx config adapter up
```

Example

```
fcpx config 4a up
```

Related concepts

[How WWPN assignments work with FC target expansion adapters](#) on page 150

Changing the system's WWNN

The WWNN of a storage system is generated by a serial number in its NVRAM, but it is stored on disk. If you ever replace a storage system chassis and reuse it in the same Fibre Channel SAN, it is

possible, although extremely rare, that the WWNN of the replaced storage system is duplicated. In this unlikely event, you can change the WWNN of the storage system.

About this task

Attention: You must change the WWNN on both systems. If both systems do not have the same WWNN, hosts cannot access LUNs on the same HA pair.

Step

1. Enter the following command:

```
fcpx nodename [-f]nodename
```

nodename is a 64-bit WWNN address in the following format: 50:0a:09:80:8X:XX:XX:XX, where X is a valid hexadecimal value.

Use `-f` to force the system to use an invalid nodename. You should not, under normal circumstances, use an invalid nodename.

Example

```
fcpx nodename 50:0a:09:80:82:02:8d:ff
```

WWPN aliases

A WWPN is a unique, 64-bit identifier displayed as a 16-character hexadecimal value in Data ONTAP. However, SAN Administrators may find it easier to identify FC ports using an alias instead, especially in larger SANs.

You can use the `wwpn-alias` sub-command to create, remove, and display WWPN aliases.

Next topics

[Creating WWPN aliases](#) on page 153

[Removing WWPN aliases](#) on page 154

[Displaying WWPN alias information](#) on page 154

Creating WWPN aliases

You use the `fcpx wwpn-alias set` command to create a new WWPN alias.

You can create multiple aliases for a WWPN, but you cannot use the same alias for multiple WWPNs. The alias can consist of up to 32 characters and can contain only the letters A through Z, a through z, numbers 0 through 9, hyphen ("-"), underscore ("_"), left brace ("{"), right brace ("}"), and period (".").

Step

1. Enter the following command:

```
fcpx wwpn-alias set [-f] alias wwpn
```

-f allows you to override a WWPN associated with an existing alias with the newly specified WWPN.

Example

```
fcpx wwpn-alias set my_alias_1 10:00:00:00:c9:30:80:2f
```

Example

```
fcpx wwpn-alias set -f my_alias_1 11:11:00:00:c9:30:80:2e
```

Removing WWPN aliases

You use the `fcpx wwpn-alias remove` command to remove an alias for a WWPN.

Step

1. Enter the following command:

```
fcpx wwpn-alias remove [-a alias ... | -w wwpn]
```

-a *alias* removes the specified aliases.

-w *wwpn* removes all aliases associated with the WWPN.

Example

```
fcpx wwpn-alias remove -a my_alias_1
```

Example

```
fcpx wwpn-alias remove -w 10:00:00:00:c9:30:80:2
```

Displaying WWPN alias information

You use the `fcpx wwpn-alias show` command to display the aliases associated with a WWPN or the WWPN associated with an alias.

Step

1. Enter the following command:

```
fcpx wwpn-alias show [-a alias | -w wwpn]
```

-a *alias* displays the WWPN associated with the alias.

-w *wwpn* displays all aliases associated with the WWPN.

Example

```
fcpx wwpn-alias show -a my_alias_1
```

Example

```
fcpx wwpn-alias show -w 10:00:00:00:c9:30:80:2
```

Example

```
fcpx wwpn-alias show
```

WWPN	Alias
----	-----
10:00:00:00:c9:2b:cb:7f	temp
10:00:00:00:c9:2b:cc:39	lrrr_1
10:00:00:00:c9:4c:be:ec	alias_0
10:00:00:00:c9:4c:be:ec	alias_0_temp
10:00:00:00:c9:2b:cc:39	lrrr_1_temp

Note: You can also use the `igroup show`, `igroup create`, `igroup add`, `igroup remove`, and `fcpx show initiator` commands to display WWPN aliases.

Obtaining fabric zone server data

You can use the zone server to access zone membership as well as port information. The `fcpx zone show` command enables you to view the active zone set on the fabric connected to the target port and to verify the zoning information on the fabric zone server.

About this task

Note: You should understand that not all FC switch vendors support the necessary fabric commands that are used to obtain zoning information.

Step

1. Obtain the fabric zone server data by entering the following command:

```
fcpx zone show
```

Example: Fabric zone server data

```
system1> fcpx zone show 4a
Active Zone Set on adapter 4a:
Zone Set Name: sanset (1 zones)
Zone Name: testzone
  Member Port Name: 10:00:00:00:c9:2d:60:dc
    Member Port Name: 50:0a:09:82:87:09:2b:7d
```

```
Member Port ID: 0x650003
Member Fabric Port Name: 20:07:00:0d:ec:00:22:80
```

Obtaining a physical topology of the FC fabric

The fabric configuration server provides information about the switches and their ports. This information can be used to generate a physical topology of the fabric.

Step

1. Obtain the physical topology of the fabric by entering the following command:

```
fcf show topology
```

Example

```
system1>fcf show topology
Port  Port WWPN                State  Type      Attached WWPN
=====
 0  20:01:00:0d:ec:00:22:80  Offline none
 1  20:02:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:82:87:39:7c:83
 2  20:03:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:81:87:c9:68:5a
 3  20:04:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:82:97:39:7c:83
 4  20:05:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:80:00:02:88:e2
 5  20:06:00:0d:ec:00:22:80  Online  F-Port    10:00:00:00:c9:2d:60:dc
 6  20:07:00:0d:ec:00:22:80  Offline none
 7  20:08:00:0d:ec:00:22:80  Offline none
 8  20:09:00:0d:ec:00:22:80  Offline none
 9  20:0a:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:80:00:02:8f:da
10  20:0b:00:0d:ec:00:22:80  Offline none
11  20:0c:00:0d:ec:00:22:80  Offline none
12  20:0d:00:0d:ec:00:22:80  Offline none
13  20:0e:00:0d:ec:00:22:80  Online  F-Port    20:00:00:e0:8b:09:89:59
14  20:0f:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:81:87:39:7c:83
15  20:10:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:81:97:39:7c:83
16  20:11:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:80:00:00:e1:66
17  20:12:00:0d:ec:00:22:80  Online  F-Port    50:0a:09:81:87:19:30:47
18  20:13:00:0d:ec:00:22:80  Online  F-Port    10:00:00:00:c9:58:46:58
19  20:14:00:0d:ec:00:22:80  Online  F-Port    10:00:00:00:c9:58:46:59
```

Obtaining fabric nameserver data

The *fabric nameserver* is the entity on the fabric that holds all information about devices in the fabric. The FC target sends a variety of defined FC commands to the nameserver to collect the fabric nameserver data.

Step

1. Obtaining the fabric nameserver data by entering the following command:

```
fcf nameserver show
```

Example

```

system1> fcp nameserver show
Name Server database connected on adapter 0c:No entries found.

Name Server database connected on adapter 0d:No entries found.

Name Server database connected on adapter 1a:

Port ID                :0xe60c00
Port Type              :N-Port
Port Name              :50:0a:09:81:87:19:66:26
Node Name              :50:0a:09:80:87:19:66:26
Symbolic Port Name     : FC Target Adapter (2532) system1:1a
Symbolic Node Name     :
                       N5300 (system1)
Fabric Port Name       :20:0c:00:05:1e:0f:7f:a5
Class of Service       :3
FC4 Type               :FCP

```

Checking connectivity of the initiators

You can use the `fcp ping` command to check the connectivity of the initiators and to verify the correctness of zoning. This command can also be used to check fabric latency between the initiator and target by using the `-s` option.

Step

1. Check the connectivity and latency by using the following command:

fcp ping

Example

```

system1> fcp ping 0c 10:00:00:00:c9:46:dc:6d
10:00:00:00:c9:46:dc:6d (0xe71100) is alive

system1> fcp ping -s 0c 10:00:00:00:c9:46:dc:6d
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=0 time=0.203 ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=1 time=0.438 ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=2 time=0.414 ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=3 time=0.246 ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=4 time=0.196 ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=5 time=0.305 ms

```

```
--- 10:00:00:00:c9:46:dc:6d ping statistics ---
6 frames transmitted, 6 frames received, 0% frame loss
```

Managing systems with onboard Fibre Channel adapters

Most systems have onboard FC adapters that you can configure as initiators or targets. Initiators connect to back-end disk shelves and targets connect to FC switches or other storage controllers.

Follow the instructions in this section to configure your onboard FC adapters as initiators or targets.

See the *Fibre Channel and iSCSI Configuration Guide* for additional configuration details.

Next topics

[Configuring onboard adapters for target mode](#) on page 158

[Configuring onboard adapters for initiator mode](#) on page 160

[Reconfiguring onboard FC adapters](#) on page 161

[Commands for displaying adapter information](#) on page 162

Related information

[Fibre Channel and iSCSI Configuration Guide - *www.ibm.com/storage/support/nas*](#)

Configuring onboard adapters for target mode

You can configure the onboard adapters for target mode to connect the adapters to the FC fabric or to another storage controller.

Before you begin

The FCP service must be licensed on the system.

About this task

If you are installing target expansion adapters, or if you exceed the allowed number of adapter ports, you must set the onboard adapters to unconfigured before installing the expansion adapters.

Note: For detailed information about the number of target adapters supported on each hardware platform, see the *iSCSI and Fibre Channel Configuration Guide*.

Steps

1. If you have already connected the port to a switch or fabric, take it offline by entering the following command:

```
fcp config adapter down
```

adapter is the port number. You can specify more than one port.

Example

```
fcp config 0c 0d down
```

Ports 0c and 0d are taken offline.

Note: If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

2. Set the onboard ports to operate in target mode by entering the following command:

```
fcadmin config -t target adapter...
```

adapter is the port number. You can specify more than one port.

Example

```
fcadmin config -t target 0c 0d
```

Ports 0c and 0d are set to target mode.

3. Run the following command to see the change in state for the ports:

```
fcadmin config
```

Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED	online
0b	initiator	CONFIGURED	online
0c	target	PENDING	online
0d	target	PENDING	online

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the fcadmin MAN page for detailed descriptions of each value.

Ports 0c and 0d are now in the PENDING state.

4. Reboot each system in the HA pair by entering the following command:

```
reboot
```

5. Start the FCP service by entering the following command:

```
fcps start
```

6. Verify that the FC ports are online and configured in the correct state for your configuration by entering the following command:

```
fcadmin config
```

Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED	online
0b	initiator	CONFIGURED	online
0c	target	PENDING	online
0d	target	PENDING	online

0a	initiator	CONFIGURED	online
0b	initiator	CONFIGURED	online
0c	target	CONFIGURED	online
0d	target	CONFIGURED	online

The preceding output displays for a four-port SAN configuration.

Related tasks

[Licensing the FC service](#) on page 146

[Reconfiguring onboard FC adapters](#) on page 161

Configuring onboard adapters for initiator mode

You can configure the onboard adapters for initiator mode to connect the adapters to back-end disk shelves.

Steps

1. If you have already connected the port to a switch or fabric, take it offline by entering the following command:

```
fcpx config adapter down
```

adapter is the port number. You can specify more than one port.

Example

```
fcpx config 0c 0d down
```

Ports 0c and 0d are taken offline.

Note: If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

2. Set the onboard ports to operate in initiator mode by entering the following command:

```
fcadmin config -t initiator adapter
```

adapter is the port number. You can specify more than one port.

Example

```
fcadmin config -t initiator 0c 0d
```

Ports 0c and 0d are set to initiator mode.

3. Reboot each system in the HA pair by entering the following command:

```
reboot
```

4. Verify that the FC ports are online and configured in the correct state for your configuration by entering the following command:

```
fcadmin config
```


Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED	online
0b	initiator	CONFIGURED	online
0c	target	CONFIGURED	online
0d	target	CONFIGURED	online

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the fcadmin MAN page for detailed descriptions of each value.

The preceding output displays for a four-port SAN configuration.

Reconfiguring onboard FC adapters

In some situations, you might need to set your onboard target adapters to unconfigured. Failure to do so could result in lost data or a system panic.

About this task

You must reconfigure the onboard adapters under the following circumstances:

- You are upgrading from a 2-Gb onboard adapter to a 4-Gb target expansion adapter or 4-Gb onboard adapter to a 8-Gb target expansion adapter. Because you cannot mix 2-Gb and 4-Gb adapters or 4-Gb and 8-Gb adapters on the same system, or on two systems in an HA pair, you must set the onboard adapters to unconfigured before installing the target expansion adapter.
- You have exceeded 16 target adapters, the maximum number of allowed adapters, on a N7600, N7700, N7800, or N7900 controller.

Steps

1. Stop the FCP service by entering the following command:

```
fcp stop
```

The FCP service is stopped and all target adapters are taken offline.

2. Set the onboard adapters to unconfigured by entering the following command:

```
fcadmin config -t unconfig ports
```

Example

```
fcadmin config -t unconfig 0b 0d
```

The onboard adapters are unconfigured.

3. Shut down the storage system.
4. If you are installing a 4-Gb or 8-Gb expansion adapters, install the adapters according to the instructions provided with the product.

5. Power on the system.

Commands for displaying adapter information

The following table lists the commands available for displaying information about adapters. The output varies depending on the storage system model.

If you want to display...	Use this command...
Information for all initiator adapters in the system, including firmware level, PCI bus width and clock speed, node name, cache size, FC packet size, link data rate, SRAM parity, and various states	storage show adapter
All adapter (HBAs, NICs, and switch ports) configuration and status information	sysconfig [-v] [adapter] <i>adapter</i> is a numerical value only. -v displays additional information about all adapters.
Disks, disk loops, and options configuration information that affects coredumps and takeover	sysconfig -c
FCP traffic information	sysstat -f
How long FCP has been running	uptime
Initiator HBA port address, port name, port name alias, node name, and igroup name connected to target adapters	fcv show initiator [-v] [adapter&portnumber] -v displays the Fibre Channel host address of the initiator. <i>adapter&portnumber</i> is the slot number with the port number, a or b; for example, 5a.
Service statistics	availtime
Target adapter configuration information	fcv config

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

Next topics

[Displaying the status of onboard FC adapters](#) on page 163

[Displaying information about all adapters](#) on page 164

[Displaying brief target adapter information](#) on page 165

[Displaying detailed target adapter information](#) on page 166

[Displaying the WWNN of a target adapter](#) on page 167

[Displaying HBA information](#) on page 168

[Displaying target adapter statistics](#) on page 168

[Displaying FC traffic information](#) on page 169

[Displaying information about FCP traffic from the partner](#) on page 170

[Displaying how long the FC service has been running](#) on page 170

[Displaying FCP service statistics](#) on page 170

Displaying the status of onboard FC adapters

Use the `fcadmin config` command to determine the status of the FC onboard adapters.

This command also displays other important information, including the configuration status of the adapter and whether it is configured as a target or initiator.

Note: Onboard FC adapters are set to initiator mode by default.

Step

1. Enter the following command:

```
fcadmin config
```

Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	CONFIGURED	online
0b	initiator	CONFIGURED	online
0c	target	CONFIGURED	online
0d	target	CONFIGURED	online

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the fcadmin MAN page for detailed descriptions of each value.

Displaying information about all adapters

You can use the `sysconfig -v` command to display system configuration and adapter information for all adapters in the system.

Step

1. Enter the following command:

```
sysconfig -v
```

Example

```
system1>sysconfig -v
slot 2: Fibre Channel Target Host Adapter 2a
          (Dual-channel, QLogic 2532 (2562) rev. 2, 32-bit,
[ONLINE])
          Firmware rev: 4.6.2
          Host Port Addr: 011200
          Cacheline size: 16
          SRAM parity: Yes
          FC Nodename: 50:0a:09:80:87:29:2a:42
(500a098087292a42)
          FC Portname: 50:0a:09:85:97:29:2a:42
(500a098597292a42)
          Connection: PTP, Fabric
          SFP Vendor Name: AVAGO
          SFP Vendor P/N: AFBR-57D5APZ
          SFP Vendor Rev: B
          SFP Serial No.: AD0820EA06W
          SFP Connector: LC
          SFP Capabilities: 2, 4, 8 Gbit/Sec
                        I/O base 0x00000000000008000, size 0x100
                        memory mapped I/O base 0xfe500000, size 0x4000
slot 2: Fibre Channel Target Host Adapter 2b
          (Dual-channel, QLogic 2532 (2562) rev. 2, 32-bit,
```

```
[ONLINE] )
          Firmware rev:    4.6.2
          Host Port Addr:  011300
          Cacheline size:  16
          SRAM parity:     Yes
          FC Nodename:     50:0a:09:80:87:29:2a:42
(500a098087292a42)
          FC Portname:     50:0a:09:86:97:29:2a:42
(500a098697292a42)
          Connection:      PTP, Fabric
          SFP Vendor Name:  AVAGO
          SFP Vendor P/N:   AFBR-57D5APZ
          SFP Vendor Rev:   B
          SFP Serial No.:   AD0820EA0ES
          SFP Connector:    LC
          SFP Capabilities: 2, 4, 8 Gbit/Sec
                        I/O base 0x00000000000008400, size 0x100
                        memory mapped I/O base 0xfe504000, size 0x4000
```

System configuration information and adapter information for each slot that is used is displayed on the screen. Look for *Fibre Channel Target Host Adapter* to get information about target HBAs.

Note: In the output, in the information about the Dual-channel QLogic HBA, the value 2532 does not specify the model number of the HBA; it refers to the device ID set by QLogic. Also, the output varies according to storage system model.

Displaying brief target adapter information

You can use the `fcpl config` command to display information about target adapters in the system, as well as to quickly detect whether the adapters are active and online.

The output of the `fcpl config` command depends on the storage system model.

Step

1. Enter the following command:

```
fcpl config
```

Example

The `fcpl config` command displays the following output:

```
7a:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 170900
      portname 50:0a:09:83:86:87:a5:09  nodename 50:0a:
09:80:86:87:a5:09
      mediatype ptp  partner adapter 7a

7b:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 171800
      portname 50:0a:09:8c:86:57:11:22  nodename 50:0a:
```

```
09:80:86:57:11:22
mediatype ptp partner adapter 7b
```

Example

The following example shows output for the N5000 series. The `fcv config` command displays information about the onboard ports connected to the SAN:

```
0c:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 010900
      portname 50:0a:09:81:86:f7:a8:42  nodename 50:0a:
09:80:86:f7:a8:42
      mediatype ptp partner adapter 0d

0d:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 010800
      portname 50:0a:09:8a:86:47:a8:32  nodename 50:0a:
09:80:86:47:a8:32
      mediatype ptp partner adapter 0c
```

Displaying detailed target adapter information

You can use the `fcv show adapter` command to display the node name, port name, and link state of all target adapters in the system.

Notice that the port name and node name are displayed with and without the separating colons. For Solaris hosts, you use the WWPN without separating colons when you map adapter port names (or these target WWPNs) to the host.

Step

1. Enter the following command:

```
fcv show adapter -v
```

Example

```
system1> fcv show adapter -v 4a
Slot: 4a
Description: Fibre Channel Target Adapter 4a (Dual-channel,
QLogic CNA 8112 (8152) rev. 2)
Status: ONLINE
Host Port Address: 0x98d601
Firmware Rev: 5.3.4
MPI Firmware Rev: 1.38.0
PHY Firmware Rev: 1.7.0
FC VLAN ID: 5
FC Nodename: 50:0a:09:80:87:69:68:5a (500a09808769685a)
FC Portname: 50:0a:09:81:87:69:68:5a (500a09818769685a)
Cacheline Size: 16
FC Packet Size: 2048
SRAM Parity: Yes
External GBIC: No
Data Link Rate: 10 GBit
Adapter Type: Local
Fabric Established: Yes
Connection Established: PTP
```

```

Mediatype:          auto
Partner Adapter:    None
Standby:            No
Target Port ID:     0x1
Switch Port:      brcddcx_rtp02:214
Physical Link Rate:  10 GBit
Physical Link Status: LINK UP

```

The information about the adapter in slot 1 displays.

Note: In the output, in the information about the Dual-channel QLogic HBA, the value 2312 does not specify the model number of the HBA; it refers to the device ID set by QLogic. Also, the output varies according to storage system model.

Note: Refer to the following table for definitions of the possible values in the Status field:

Status	Definition
Uninitialized	The firmware has not yet been loaded and initialized.
Link not connected	The driver has finished initializing the firmware. However, the link is not physically connected so the adapter is offline.
Online	The adapter is online for FC traffic.
Link disconnected	The adapter is offline due to a Fibre Channel link offline event.
Offline	The adapter is offline for FC traffic.
Offlined by user/system	A user manually took the adapter offline, or the system automatically took the adapter offline.

Displaying the WWNN of a target adapter

Use the `fcpl nodename` command to display the WWNN of a target adapter in the system.

Step

1. Enter the following command:

```
fcpl nodename
```

Example

```
Fibre Channel nodename: 50:a9:80:00:02:00:8d:b2 (50a9800002008db2)
```

Displaying HBA information

HBAs are adapters on the host machine that act as initiators. Use the `fcpl show initiator` command to display the port names, aliases, and igroup names of HBAs connected to target adapters on the storage system.

Step

1. Enter the following command:

```
fcpl show initiator
```

Example

```
fcpl show initiator
Portname                Alias      Group
10:00:00:00:c9:32:74:28 calculon0  calculon
10:00:00:00:c9:2d:60:dc gaston0    gaston
10:00:00:00:c9:2b:51:1f
Initiators connected on adapter 0b: None connected.
```

Displaying target adapter statistics

Use the `fcpl stats` command to display important statistics for the target adapters in your system.

Step

1. Enter the following command:

```
fcpl stats -i interval [-c count] [-a | adapter]
```

`-i interval` is the interval, in seconds, at which the statistics are displayed.

`-c count` is the number of intervals. For example, the `fcpl stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.

`-a` shows statistics for all adapters.

`adapter` is the slot and port number of a specific target adapter.

Example

```
fcpl stats -i 1
r/s    w/s    o/s    ki/s    ko/s    asvc_t    qlen hba
0       0       0       0       0       0.00     0.00 7a
110    113     0    7104    12120     9.64     1.05 7a
146     68     0    6240    13488    10.28     1.05 7a
106     92     0    5856    10716    12.26     1.06 7a
136    102     0    7696    13964     8.65     1.05 7a
```

Each column displays the following information:

r/s—The number of SCSI read operations per second.

w/s—The number of SCSI write operations per second.

o/s—The number of other SCSI operations per second.

ki/s— Kilobytes per second of received traffic

ko/s—Kilobytes per second send traffic.

asvc_t—Average time in milliseconds to process a request

qlen—The average number of outstanding requests pending.

hba—The HBA slot and port number.

To see additional statistics, enter the `fcv stats` command with no variables.

Displaying FC traffic information

You can use the `sysstat -f` command to display FC traffic information, such as operations per second and kilobytes per second.

Step

1. Enter the following command:

```
sysstat -f
```

Example

CPU	NFS	CIFS	FCP	Net in	kB/s out	Disk read	kB/s write	FCP in	kB/s out	Cache age
81%	0	0	6600	0	0	105874	56233	40148	232749	1
78%	0	0	5750	0	0	110831	37875	36519	237349	1
78%	0	0	5755	0	0	111789	37830	36152	236970	1
80%	0	0	7061	0	0	107742	49539	42651	232778	1
78%	0	0	5770	0	0	110739	37901	35933	237980	1
79%	0	0	5693	0	0	108322	47070	36231	234670	1
79%	0	0	5725	0	0	108482	47161	36266	237828	1
79%	0	0	6991	0	0	107032	39465	41792	233754	1
80%	0	0	5945	0	0	110555	48778	36994	235568	1
78%	0	0	5914	0	0	107562	43830	37396	235538	1

The following columns provide information about FCP statistics:

CPU—The percentage of the time that one or more CPUs were busy.

FCP—The number of FCP operations per second.

FCP KB/s—The number of kilobytes per second of incoming and outgoing FCP traffic.

Displaying information about FCP traffic from the partner

If you have an HA pair, you might want to obtain information about the amount of traffic coming to the system from its partner.

Step

1. Enter the following command:

```
sysstat -b
```

The following show the columns information about partner traffic:

Partner—The number of partner operations per second.

Partner KB/s—The number of kilobytes per second of incoming and outgoing partner traffic.

Related concepts

[How to manage FC with HA pairs](#) on page 137

Displaying how long the FC service has been running

Use the `uptime` command to display how long the FC service has been running on the system.

Step

1. Enter the following command:

```
uptime
```

Example

```
12:46am up 2 days, 8:59 102 NFS ops, 2609 CIFS ops, 0 HTTP ops, 0 DAFS
ops, 1933084 FCP ops, 0 iSCSI ops
```

Displaying FCP service statistics

Use the `availtime` command to display the FCP service statistics.

Step

1. Enter the following command:

```
availtime
```

Example

```
Service statistics as of Mon Jul 1 00:28:37 GMT 2002
System (UP). First recorded (3894833) on Thu May 16 22:34:44 GMT 2002
  P 28, 230257, 170104, Mon Jun 10 08:31:39 GMT 2002
  U 24, 131888, 121180, Fri Jun 7 17:39:36 GMT 2002
NFS (UP). First recorded (3894828) on Thu May 16 22:34:49 GMT 2002
  P 40, 231054, 170169, Mon June 10 08:32:44 GMT 2002
```

FCP	U	36, 130363, 121261, Fri Jun 7 17:40:57 GMT 2002
	P	19, 1417091, 1222127, Tue Jun 4 14:48:59 GMT 2002
	U	6, 139051, 121246, Fri Jun 7 17:40:42 GMT 2002

Disk space management

Data ONTAP is equipped with a number of tools for effectively managing disk space.

This section describes how to complete these tasks:

- Monitor available disk space
- Configure Data ONTAP to automatically grow a FlexVol volume
- Configure Data ONTAP to automatically delete Snapshot copies when a FlexVol volume begins to run out of free space

Note: For more in-depth discussions of disk space management, refer to the *Data ONTAP 7-Mode Storage Management Guide*.

Next topics

[Commands to display disk space information](#) on page 173

[Examples of disk space monitoring using the `df` command](#) on page 174

[How Data ONTAP can automatically provide more free space for full volumes](#) on page 178

[Configuring automatic free space preservation for a FlexVol volume](#) on page 179

[Moving your volumes nondisruptively](#) on page 180

[Working with VMware VAAI features for ESX hosts](#) on page 187

Related information

[Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas](http://www.ibm.com/storage/support/nas)

Commands to display disk space information

You can see information about how disk space is being used in your aggregates and volumes and their Snapshot copies.

Use this Data ONTAP command...	To display information about...
<code>aggr show_space</code>	Disk space usage for aggregates
<code>df</code>	Disk space usage for volumes or aggregates
<code>snap delta</code>	The estimated rate of change of data between Snapshot copies in a volume
<code>snap reclaimable</code>	The estimated amount of space freed if you delete the specified Snapshot copies

For more information about the `snap` commands, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*. For more information about the `df` and `aggr show_space` commands, see the appropriate man page.

Examples of disk space monitoring using the `df` command

You can use the `df` command to monitor disk space on a volume in which you created LUNs.

Note: These examples are written with the assumption that the storage system and host machine are already properly configured.

Next topics

[Monitoring disk space on volumes with LUNs that do not use Snapshot copies](#) on page 174

[Monitoring disk space on volumes with LUNs that use Snapshot copies](#) on page 176

Monitoring disk space on volumes with LUNs that do not use Snapshot copies

This example illustrates how to monitor disk space on a volume when you create a LUN without using Snapshot copies.

About this task

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

For simplicity, assume the LUN requires only three GB of disk space. For a traditional volume, the volume size must be approximately three GB plus 10 percent. The recommended volume size is approximately 2*3 GB plus the rate of change of data.

Steps

1. From the storage system, create a new traditional volume named `volspace` that has approximately 67 GB, and observe the effect on disk space by entering the following commands:

```
vol create volspace aggr1 67g
```

```
df-r/vol/volspace
```

The following sample output is displayed. There is a snap reserve of 20 percent on the volume, even though the volume is used for LUNs, because snap reserve is set to 20 percent by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/
volspace/.snapshot					

2. Set the percentage of snap reserve space to 0 and observe the effect on disk space by entering the following commands:

```
snap reserve volspace 0
```

```
df -r /vol/volspace
```

The following sample output is displayed. The amount of available Snapshot copy space becomes zero, and the 20 percent of Snapshot copy space is added to available space for /vol/volspace.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

3. Create a LUN named /vol/volspace/lun0 and observe the effect on disk space by entering the following commands:

```
lun create -s 3g -t aix /vol/volspace/lun0
```

```
df -r /vol/volspace
```

The following sample output is displayed. Three GB of space is used because this is the amount of space specified for the LUN, and LUN space reservation is enabled by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	3150268	59499640	0	/vol/volspace/
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

4. Create an igroup named aix_host and map the LUN to it by entering the following commands (assuming that the host node name is iqn.1996-04.aixhost.host1). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space.

```
igroup create-i -taiaix_hostiqn.1996-04.aixhost.host1
```

```
lun map /vol/volspace/lun0aix_host 0
```

5. From the host, discover the LUN, format it, make the file system available to the host, and write data to the file system. For information about these procedures, see your Host Utilities documentation. These commands have no effect on disk space.
6. From the storage system, ensure that creating the file system on the LUN and writing data to it has no effect on space on the storage system by entering the following command:

```
df -r /vol/volspace
```

The following sample output is displayed. From the storage system, the amount of space used by the LUN remains 3 GB.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	3150268	59499640	0	/vol/
volspace/					
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

7. Turn off space reservations and see the effect on space by entering the following commands:

```
lun set reservation /vol/volspace/lun0 disable
df -r /vol/volspace
```

The following sample output is displayed. The 3 GB of space for the LUN is no longer reserved, so it is not counted as used space; it is now available space. Any other requests to write data to the volume can occupy all of the available space, including the 3 GB that the LUN expects to have. If the available space is used before the LUN is written to, write operations to the LUN fail. To restore the reserved space for the LUN, turn space reservations on.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	144	62649584	0	/vol/volspace/
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

Monitoring disk space on volumes with LUNs that use Snapshot copies

This example illustrates how to monitor disk space on a volume when taking Snapshot copies.

About this task

Assume that you start with a new volume, and the LUN requires three GB of disk space, and fractional overwrite reserve is set to 100 percent.

Steps

1. From the storage system, create a new FlexVol volume named volspace that has approximately 67 GB, and observe the effect on disk space by entering the following commands:

```
vol create volspace aggr1 67g
df -r /vol/volspace
```

The following sample output is displayed. There is a snap reserve of 20 percent on the volume, even though the volume will be used for LUNs, because snap reserve is set to 20 percent by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/
volspace/.snapshot					

2. Set the percentage of snap reserve space to zero by entering the following command:

```
snap reserve volspace 0
```

3. Create a LUN (/vol/volspace/lun0) by entering the following commands:

```
lun create -s 6g -t aix /vol/volspace/lun0
df -r /vol/volspace
```

The following sample output is displayed. Approximately six GB of space is taken from available space and is displayed as used space for the LUN:

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	6300536	56169372	0	/vol/


```
volSPACE/
/vol/volSPACE/.snapshot      0      0      0      0 /vol/
volSPACE/.snapshot
```

4. Create an igroup named `aix_host` and map the LUN to it by entering the following commands (assuming that the host node name is `iqn.1996-04.aixhost.host1`). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space.

```
igroup create -i -t aix aix_host iqn.1996-04.aixhost.host1
lun map/vol/volSPACE/lun0aix_host 0
```

5. From the host, discover the LUN, format it, make the file system available to the host, and write data to the file system. For information about these procedures, refer to your Host Utilities documentation. These commands have no effect on disk space.
6. From the host, write data to the file system (the LUN on the storage system). This has no effect on disk space.
7. Ensure that the active file system is in a quiesced or synchronized state.
8. Take a Snapshot copy of the active file system named `snap1`, write one GB of data to it, and observe the effect on disk space by entering the following commands:

```
snap create volSPACE snap1
df -r /vol/volSPACE
```

The following sample output is displayed. The first Snapshot copy reserves enough space to overwrite every block of data in the active file system, so you see 12 GB of used space, the 6-GB LUN (which has 1 GB of data written to it), and one Snapshot copy. Notice that 6 GB appears in the reserved column to ensure write operations to the LUN do not fail. If you disable space reservation, this space is returned to available space.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volSPACE/	62649908	12601072	49808836	6300536	/vol/
/vol/volSPACE/.snapshot	0	180	0	0	/vol/
volSPACE/.snapshot					

9. From the host, write another 1 GB of data to the LUN. Then, from the storage system, observe the effect on disk space by entering the following commands:

```
df -r /vol/volSPACE
```

The following sample output is displayed. The amount of data stored in the active file system does not change. You just overwrote 1 GB of old data with 1 GB of new data. However, the Snapshot copy requires the old data to be retained. Before the write operation, there was only 1 GB of data, and after the write operation, there was 1 GB of new data and 1 GB of data in a Snapshot copy. Notice that the used space increases for the Snapshot copy by 1 GB, and the available space for the volume decreases by 1 GB.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volSPACE/	62649908	12601072	47758748	0	/vol/
volSPACE/					

/vol/volospace/.snapshot volospace/.snapshot	0	1050088	0	0	/vol/
---	---	---------	---	---	-------

10. Ensure that the active file system is in a quiesced or synchronized state.
11. Take a Snapshot copy of the active file system named snap2 and observe the effect on disk space by entering the following command:

```
snap create volospace snap2
```

The following sample output is displayed. Because the first Snapshot copy reserved enough space to overwrite every block, only 44 blocks are used to account for the second Snapshot copy.

Filesystem	kbytes	used	avail	reserved	Mounted
on					
/vol/volospace/	62649908	12601072	47758748	6300536	/vol/
volospace/					
/vol/volospace/.snapshot	0	1050136	0	0	/vol/
volospace/.snapshot					

12. From the host, write 2 GB of data to the LUN and observe the effect on disk space by entering the following command:

```
df -r /vol/volospace
```

The following sample output is displayed. The second write operation requires the amount of space actually used if it overwrites data in a Snapshot copy.

Filesystem	kbytes	used	avail	reserved	Mounted
on					
/vol/volospace/	62649908	12601072	4608427	6300536	/vol/
volospace/					
/vol/volospace/.snapshot	0	3150371	0	0	/vol/
volospace/					
.snapshot					

How Data ONTAP can automatically provide more free space for full volumes

Data ONTAP can automatically make more free space available for a FlexVol volume when that volume is nearly full. You can choose to make the space available by first allowing the volume size to increase, or by first deleting Snapshot copies.

Data ONTAP can automatically provide more free space for the volume by using one of the following methods:

- Increase the size of the volume when it is nearly full.
This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can increase the size in increments and set a maximum size for the volume.
- Delete Snapshot copies when the volume is nearly full.
For example, you can automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want to delete first—your oldest or newest Snapshot copies. You can also determine when to begin deleting Snapshot

copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.

For more information about deleting Snapshot copies automatically, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

You can choose which method (increasing the size of the volume or deleting Snapshot copies) you want Data ONTAP to try first. If the first method does not provide sufficient extra free space to the volume, Data ONTAP will try the other method next.

Configuring a FlexVol volume to grow automatically

You configure FlexVol volumes to grow automatically to ensure that space in your aggregates is used efficiently, and to reduce the likelihood that your volumes will run out of space.

Step

1. Enter the following command:

```
vol autosize vol_name [-m size] [-I size] on
```

-m size is the maximum size to which the volume will grow. Specify a size in **k** (KB), **m** (MB), **g** (GB) or **t** (TB).

-I size is the increment by which the volume's size increases. Specify a size in **k** (KB), **m** (MB), **g** (GB) or **t** (TB).

Result

If the specified FlexVol volume is about to run out of free space and is smaller than its maximum size, and if there is space available in its containing aggregate, its size will increase by the specified increment.

Configuring automatic free space preservation for a FlexVol volume

When you configure a FlexVol volume for automatic free space preservation, the FlexVol volume attempts to provide more free space when it becomes nearly full. It can provide more free space by increasing its size or by deleting Snapshot copies, depending on how you have configured the volume.

Step

1. Enter the following command:

```
vol options vol-name try_first [volume_grow|snap_delete]
```

If you specify `volume_grow`, Data ONTAP attempts to increase the volume's size before deleting any Snapshot copies. Data ONTAP increases the volume size based on specifications you provided using the `vol autosize` command.

If you specify `snap_delete`, Data ONTAP attempts to create more free space by deleting Snapshot copies, before increasing the size of the volume. Data ONTAP deletes Snapshot copies based on the specifications you provided using the `snap autodelete` command.

Moving your volumes nondisruptively

IBM N series DataMotion for Volumes enables you to nondisruptively move a volume from one aggregate to another within the same controller for capacity utilization, improved performance, and to satisfy service-level agreements. In a SAN environment, FlexVol volumes and the LUNs in the volumes are moved nondisruptively from one aggregate to another.

In a volume move, SCSI applications accessing different LUNs in the volume can continue to run during the move. Applications that use FC and iSCSI to access a LUN in the volume that is being moved do not see any I/O disruptions during the volume move. You can continue to access data in the volume during and after the volume move.

The volume move occurs in three phases: setup phase, data copy phase, and cutover phase.

Next topics

[*Ways to use volume move*](#) on page 180

[*Requirements for performing a volume move*](#) on page 181

[*How the setup phase of volume move works*](#) on page 182

[*How the data copy phase of volume move works*](#) on page 182

[*How the cutover phase of volume move works*](#) on page 183

[*Performing the volume move operation*](#) on page 184

[*Pausing the volume move operation*](#) on page 185

[*Resuming the volume move operation*](#) on page 185

[*Monitoring the volume move status*](#) on page 186

[*Performing manual cutover of the volume move operation*](#) on page 186

[*Canceling the volume move operation*](#) on page 187

Ways to use volume move

You can move a volume nondisruptively in different scenarios, such as moving it from a busy aggregate to a less busy aggregate or from a high-speed disk to a lower-speed disk.

You can move the volume in the following scenarios:

- From a high-speed disk to a lower-speed disk or from a lower-speed disk to a high-speed disk, to satisfy SLA requirements.
- From a full aggregate to an aggregate that has space for growth.

- From an aggregate laid out on third-party disks to an aggregate laid out on IBM N series disks by using gateways.
- Between different RAID types, such as RAID-DP and RAID4.
- Between different types of disk drives, such as array LUNs, SSDs, FC, SATA, or SAS.

Requirements for performing a volume move

Before you move a volume nondisruptively, you must be aware of the type of volumes you can move and the operations that might conflict with the volume move. The volume move does not start if the volume has unsupported settings or if there are conflicting operations.

- Your filer or gateway must be running Data ONTAP 8.0.1 7-Mode or later.
 - You can move only one 7-Mode FlexVol volume at a time.
 - The volume must be online.
 - You cannot move the following types of volumes:
 - A root volume
 - A FlexClone volume
 - A FlexCache volume
 - A volume that is the destination of any replication relationship, such as volume SnapMirror or qtree SnapMirror
 - A volume that is a SnapVault destination
- Note:** During a volume move, you must not initiate qtree SnapMirror or SnapVault relationships from the destination volume.
- A read-only volume
 - A volume in a nondefault vFiler unit
 - A volume from a 32-bit aggregate to a 64-bit aggregate, or from a 64-bit aggregate to a 32-bit aggregate
 - The data in the volume must not be compressed using the data compression feature.
 - The source volume should not be exported to NFS or CIFS clients when the volume move operation is in progress.

There is a small window of time when you can export the source volume over NFS or CIFS before the volume move enters the cutover phase. However, if you do so, the cutover phase might not be successfully completed. If the cutover phase is not completed, there is no disruption to SCSI clients because the volume move rolls back to continue with the data copy phase.

- The source volume must be consistent.
- The volume guarantee option must not be set to `file`.
- Deduplication operations must not be running on the source volume.

If deduplication is active, the volume move is paused and the cutover phase is not initiated. For more information about deduplication operation, see the *Data ONTAP 7-Mode Storage Management Guide*.

- The following conflicting operations must not be running:
 - SnapRestore of the source volume or the containing aggregate

- WAFLIron operation on the source or the destination aggregate
- Active LUN clone split operations on the source volume
- Revert operation on the storage system

Note: FlexClone volumes, fingerprint database, and change logs in the source volume are not moved along with the source volume.

Related concepts

[How the setup phase of volume move works](#) on page 182

[How the data copy phase of volume move works](#) on page 182

[How the cutover phase of volume move works](#) on page 183

Related information

[Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas](http://www.ibm.com/storage/support/nas)

How the setup phase of volume move works

The setup phase creates a temporary destination volume in the destination aggregate and initiates data transfer from the source volume to the destination volume.

During the setup phase, the system checks if the volume you plan to move meets the specified requirements. If any of these checks fail, then the volume move is terminated and an error message is displayed. You should follow the guidance of the error message before you can manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 181

[How the data copy phase of volume move works](#) on page 182

[How the cutover phase of volume move works](#) on page 183

Related tasks

[Resuming the volume move operation](#) on page 185

How the data copy phase of volume move works

The data copy phase follows the setup phase of a volume move operation. In the data copy phase, data is transferred automatically from the source volume to the destination volume, after which the cutover phase can begin.

After each block of data is transferred, the volume move determines if the cutover phase can be initiated.

If a SnapRestore or a WAFLIron operation is started on the source volume, the destination volume, or the containing aggregate, the volume move is canceled and an appropriate error message is recorded in the log file.

If the volume move finds any unsupported settings or conflicting operations before entering the cutover phase, the volume move operation is paused and the reason for the pause is displayed. You must resolve the issue before you can manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 181

[How the setup phase of volume move works](#) on page 182

[How the cutover phase of volume move works](#) on page 183

Related tasks

[Resuming the volume move operation](#) on page 185

How the cutover phase of volume move works

The cutover phase is the final phase of the volume move. During the cutover phase, the data in the source volume and the destination volume is synchronized. I/O operations are redirected to the destination volume and the volume move is complete.

Note: The host application might encounter I/O disruptions if storage system reboot, nondisruptive upgrade (NDU), shutdown, takeover, or giveback occurs during the volume move.

If the volume move is not completed within the specified cutover period (default 60 seconds), then the cutover phase is timed out, logging the appropriate error messages, and the volume move reverts to the data copy phase.

If the cutover phase is successful, it results in the following:

- The contents of the destination volume are identical to the source volume.
- The destination volume takes the identity of the source volume.
- After the volume is moved, the LUN at the destination starts processing I/O operations.
- The source volume is destroyed, unless you choose to retain it.

Depending on the number of cutover attempts, the volume move tries to enter the cutover phase again. If cutover is not completed within the specified number of cutover attempts, then the volume move is paused and an appropriate error message is recorded in the log file. You can then manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 181

[How the setup phase of volume move works](#) on page 182

[How the data copy phase of volume move works](#) on page 182

Related tasks

[Performing the volume move operation](#) on page 184

[Resuming the volume move operation](#) on page 185

Performing the volume move operation

You can nondisruptively move a volume from one aggregate to another within a storage system. You can continue to access data in the LUNs during the volume move.

Before you begin

Before the volume move enters the cutover phase, you must ensure that any existing synchronous SnapMirror relationships established on the source volume are destroyed. You can resynchronize the SnapMirror relationships after the volume move is completed.

About this task

A temporary volume is created at the beginning of the volume move. You should not change the contents, state, or attributes of the destination volume, or create any replication, disaster recovery, SnapVault, or qtree SnapMirror relationship with other volumes for the duration of the move.

MetroCluster relationships are not affected by the volume move.

Step

1. Start the volume move by entering the following command:

```
vol move start srcvol dstaggr [-k] [-m | -r num_cutover_attempts] [-w cutover_window] [-o] [-d]
```

srcvol specifies the source volume.

dstaggr specifies the destination aggregate.

-k retains the source volume after a successful move. The source volume remains offline.

-m specifies that the volume move does not initiate automatic cutover. The system continuously runs updates and you can initiate manual cutover at any point during the volume move.

num_cutover_attempts specifies the number of cutover attempts. The minimum number of cutover attempts is one and the default number of attempts is three. If cutover cannot be completed in the specified number of attempts, then the volume move is paused.

cutover_window specifies the duration of the cutover window. The default and minimum value is 60 seconds.

-o displays warning messages on the console and the operation continues.

-d runs all the data copy phase checks. If any of the checks fail, error messages are displayed on the console and the operation is terminated.

Result

If the volume move is successful, the destination volume retains the following:

- Snapshot copies of the source volume
- Attributes of the LUNs from the source volume in the corresponding LUNs in the destination volume

Related concepts

[How the setup phase of volume move works](#) on page 182

[How the data copy phase of volume move works](#) on page 182

[How the cutover phase of volume move works](#) on page 183

Pausing the volume move operation

You can manually pause the volume move during the setup phase or the data copy phase to complete any high priority I/O operations.

Step

1. Pause the volume move by entering the following command:

```
vol move pause srcvol
```

Example

```
system1> vol move pause voll
Wed Aug 29 08:11:40 GMT [system1: replication.src.err:error]:
SnapMirror: source transfer from voll to system1:
ndm_dstvol_1188375081 : transfer failed.
Wed Aug 29 08:11:41 GMT [system1: replication.dst.err:error]:
SnapMirror: destination transfer from 127.0.0.1:voll to
ndm_dstvol_1188375081 : replication transfer failed to complete.
Wed Aug 29 08:11:41 GMT [system1: vol.move.paused:info]:
Move of volume voll to aggregate aggr1 paused : User initiated
```

Resuming the volume move operation

When the volume move is manually or automatically paused, you can resume it by running the `vol move resume` command. On resuming, the volume move runs the same set of checks that were run during the data copy phase. You can add to or change the options you specified when you started the volume move.

Step

1. Resume the volume move operation by entering the following command:

```
vol move resume srcvol [-k] [-m | -r num_cutover_attempts] [-w  
cutover_window] [-o]
```

Example

```
system1> vol move resume voll -k -r 8 -w 120
Wed Aug 29 08:15:14 GMT [system1: vol.move.resume:info]:
```

```

Move of volume voll to aggregate aggr1 was resumed.
system1> Wed Aug 29 08:15:14 GMT [system1:
vol.move.transferStart:info]: Baseline transfer from volume voll
to ndm_dstvol_1188375081 started.

```

Monitoring the volume move status

You can use the `vol move status` command to display information about the volume that is moved.

Step

1. Obtain the status of the volume move operation by entering the following command:

```
vol move status srcvol [-v]
```

`-v` provides additional information about the destination volume name, amount of data transferred, the time taken for the data transfer, and the amount of data that is currently being transferred.

Example

```

system1> vol move status voll -v
Source           : voll
Destination      : aggr1:ndm_dstvol_1188375081
State            : move
Cutover Attempts : 3
Cutover Time     : 60
Last Completed Transfer:
    Data Transferred = 324 KB      Time Taken = 1 s
Current Transfer Size = 0 KB

```

Performing manual cutover of the volume move operation

If the volume move is unable to complete automatic cutover in the specified number of cutover attempts, you can initiate manual cutover. You can specify the `-m` option when starting or resuming the volume move to initiate cutover and increase the probability of completing the volume move within the cutover period.

Step

1. Manually cut over the volume move operation by entering the following command:

```
vol move cutover srcvol [-w cutover_window]
```

Canceling the volume move operation

You can cancel the volume move if you want to complete any high priority operations.

Step

1. Cancel the volume move operation by entering the following command:

```
vol move abort srcvol
```

Working with VMware VAAI features for ESX hosts

Data ONTAP 8.0.1 and later supports certain VMware vStorage APIs for Array Integration (VAAI) features when the ESX host is running ESX 4.1 or later. These features help offload operations from the ESX host to the storage system and increase the network throughput. The ESX host enables the features automatically in the correct environment. You can determine the extent to which your system is using the VAAI features by checking the statistics contained in the VAAI counters.

The VAAI feature set consists of the following:

- **Extended copy**
This feature offloads the work of certain copy operations (repeated reads and writes) from the host to the storage system, which results in saving ESX CPU cycles and increasing the network throughput. The extended copy feature is used in scenarios such as cloning a virtual machine. When invoked by the ESX host, the extended copy feature copies the data within the N series storage system rather than going through the host network. If this feature cannot be invoked, the ESX host automatically uses the standard ESX copy operation.
- **WRITE SAME**
This feature offloads the work of writing a repeated pattern, such as all zeros, to a storage array. The ESX host uses this feature in scenarios such as zero-filling a file.
- **VERIFY AND WRITE**
This feature bypasses certain file access concurrency limitations, which speeds up operations such as booting up a virtual machine.

Next topics

[Requirements for using the VAAI environment](#) on page 188

[Methods for determining whether VAAI features are supported](#) on page 188

[Statistics collected for VAAI features](#) on page 189

[Viewing statistics for the VAAI features](#) on page 190

Requirements for using the VAAI environment

The VAAI features are part of the ESX operating system and are automatically invoked by the ESX host when you have set up the correct environment.

The environment requirements are as follows:

- The ESX host must be running ESX 4.1 or later.
- The N series storage system that is hosting the VMware datastore must be running Data ONTAP 8.0.1 or later.
- (Extended copy only) Both the LUNs and the igroups must specify `VMware` as the OS type.
- (Extended copy only) The source and the destination of the VMware copy operation must be hosted on the same storage system.

It does not matter whether the VMware datastores are on different LUNs or volumes within that storage system.

Note: The extended copy feature currently does not support copying data between VMware datastores that are hosted on different storage systems.

Methods for determining whether VAAI features are supported

To confirm whether the ESX operating system supports the VAAI features, you can check either the Virtual Storage Console (VSC) or the statistics produced by the VAAI counters.

- When you are at the VSC, you can look at the VAAI Capable option. If it is displayed as Enabled, then the storage system is capable of using the VAAI features.
- To view the statistics on the VAAI features, you can use the `stats show vstorage` command. When you enter this command without an option, it displays all the counters associated with the VAAI features. When you enter it with the name of a counter as an option (`stats show vstorage:counter_name`), it displays information for only that counter.

By checking the requests counter for a feature, you can determine whether the ESX host is using that feature. This counter specifies how many requests for that feature have been sent to the storage system. The counter value increases as the ESX host invokes the feature.

The following table lists the requests counters for each feature:

Feature	Counter
Extended copy	<code>xcopy_copy_reqs</code>
WRITE SAME	<code>writesame_reqs</code>
VERIFY AND WRITE	<code>vaw_reqs</code>

Statistics collected for VAAI features

The VAAI counters supply numerous statistics that provide information such as which features the ESX host is using, how they are performing, and how much data is being operated on by the features.

Each of the following counters supplies information for a single vFiler unit.

xcopy_copy_reqs	The number of requests for the extended copy feature.
xcopy_abort_reqs	The number of requests to abort the extended copy feature commands.
xcopy_status_reqs	The number of requests for status information about the extended copy feature commands.
xcopy_total_data	<p>The sum of the kilobytes of data that was successfully copied using extended copy.</p> <p>This is a measurement of data copied at the N series storage system rather than through the network.</p>
xcopy_invalid_parms	The number of extended copy requests that had invalid parameters.
xcopy_authorization_failures	The number of unauthorized requests for the extended copy feature.
xcopy_authentication_failures	The number of requests for the extended copy feature that could not be authenticated.
xcopy_copy_failures	The total number of extended copy requests that failed during copy operations.
xcopy_copyErr_isDir	The number of extended copy requests that were sent to a directory instead of a file.
xcopy_copyErr_data_unrecov	The number of extended copy requests received that failed due to an unrecoverable RAID error.
xcopy_copyErr_offline	The number of extended copy requests that failed because the volume was offline.
xcopy_copyErr_staleFH	The number of extended copy requests that failed because the request referenced an invalid file handle.
xcopy_copyErr_IO	The number of extended copy requests that failed because there was no I/O available on the storage system.
xcopy_copyErr_noSpace	The number of extended copy requests that failed because of an internal I/O error.
xcopy_copyErr_diskQuota	The number of extended copy requests that failed because the disk quota on the storage system was exceeded.

xcopy_copyErr_readOnly	The number of extended copy requests that failed because the copy destination was read-only.
xcopy_copyErr_other	The number of extended copy requests that failed due to a generic copy operation failure.
xcopy_intravol_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were within the same volume.
xcopy_intervol_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were on different volumes.
xcopy_one2one_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were within the same virtual machine disk (VMDK).
xcopy_one2many_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were on different VMDKs.
writesame_reqs	The sum of the WRITE SAME requests.
writesame_holepunch_reqs	The number of requests for WRITE SAME operations that were used to perform hole punching (freeing of blocks).
writesame_total_data	The sum of the kilobytes of data that was successfully written using the WRITE SAME requests.
vaw_reqs	The sum of VAW requests.
vaw_miscompares	The sum of VAW requests that resulted in a miscompare (contention for resource).

Viewing statistics for the VAAI features

You can use the `stats show` command with the option `vstorage` to display the statistics that the counters collected about the VAAI features extended copy, WRITE SAME, and VERIFY AND WRITE.

Step

1. To view the statistics for the VAAI features, complete the appropriate action:

To view...	Enter...
All the statistics	The command: <code>stats show vstorage</code>

To view...	Enter...
A specific statistic	The <code>stats show vstorage</code> command with the name of the counter that contains the statistics you want to see:
	<code>stats show vstorage:counter_name</code>

Example

The following example uses the `stats show vstorage` command to display information from all the counters for the VAAI features:

```
TESTER1*> stats show vstorage
vstorage:vfiler0:xcopy_copy_reqs:1139
vstorage:vfiler0:xcopy_abort_reqs:0
vstorage:vfiler0:xcopy_status_reqs:0
vstorage:vfiler0:xcopy_total_data:4046848
vstorage:vfiler0:xcopy_invalid_parms:0
vstorage:vfiler0:xcopy_authorization_failures:0
vstorage:vfiler0:xcopy_authentication_failures:0
vstorage:vfiler0:xcopy_copy_failures:73
vstorage:vfiler0:xcopy_copyErr_isDir:0
vstorage:vfiler0:xcopy_copyErr_data_unrecov:0
vstorage:vfiler0:xcopy_copyErr_offline:0
vstorage:vfiler0:xcopy_copyErr_staleFH:0
vstorage:vfiler0:xcopy_copyErr_IO:0
vstorage:vfiler0:xcopy_copyErr_noSpace:0
vstorage:vfiler0:xcopy_copyErr_diskQuota:0
vstorage:vfiler0:xcopy_copyErr_readOnly:0
vstorage:vfiler0:xcopy_copyErr_other:0
vstorage:vfiler0:xcopy_intravol_moves:530
vstorage:vfiler0:xcopy_intervol_moves:536
vstorage:vfiler0:xcopy_one2one_moves:0
vstorage:vfiler0:xcopy_one2many_moves:0
vstorage:vfiler0:writesame_reqs:0
vstorage:vfiler0:writesame_holepunch_reqs:0
vstorage:vfiler0:writesame_total_data:0
vstorage:vfiler0:vaw_reqs:0
vstorage:vfiler0:vaw_miscompares:0
TESTER1*>
```

In the following example, the command displays only the information collected by the `xcopy_abort_reqs` counter:

```
TESTER1*> stats show vstorage:vfiler0:xcopy_abort_reqs
vstorage:vfiler0:xcopy_abort_reqs:0
TESTER1*>
```


Data protection with Data ONTAP

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

Snapshot copies are a standard feature of Data ONTAP. A Snapshot copy is a frozen, read-only image of the entire Data ONTAP file system, or WAFL (Write Anywhere File Layout) volume, that reflects the state of the LUN or the file system at the time the Snapshot copy is created. The other data protection methods listed in the table below rely on Snapshot copies or create, use, and destroy Snapshot copies, as required.

Next topics

[Data protection methods](#) on page 193

[LUN clones](#) on page 195

[Deleting busy Snapshot copies](#) on page 203

[Restoring a Snapshot copy of a LUN in a volume](#) on page 206

[Restoring a single LUN](#) on page 208

[Backing up SAN systems to tape](#) on page 209

[Using volume copy to copy LUNs](#) on page 212

Data protection methods

The following table describes the various methods for protecting your data with Data ONTAP.

Method	Used to...
Snapshot copy	Make point-in-time copies of a volume.
SnapRestore	<ul style="list-style-type: none"> Restore a LUN or file system to an earlier preserved state in less than a minute without rebooting the storage system, regardless of the size of the LUN or volume being restored. Recover from a corrupted database or a damaged application, a file system, a LUN, or a volume by using an existing Snapshot copy.

Method	Used to...
SnapMirror	<ul style="list-style-type: none"> Replicate data or asynchronously mirror data from one storage system to another over local or wide area networks (LANs or WANs). Transfer Snapshot copies taken at specific points in time to other storage systems or near-line 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	<ul style="list-style-type: none"> Back up data by using Snapshot copies on the storage system and transferring them on a scheduled basis to a destination storage system. Store these Snapshot copies on the destination storage system for weeks or months, allowing recovery operations to occur nearly instantaneously from the destination storage system to the original storage system.
SnapDrive for Windows or UNIX	<ul style="list-style-type: none"> Manage storage system Snapshot copies directly from a Windows or UNIX host. Manage storage (LUNs) directly from a host. Configure access to storage directly from a host. <p>SnapDrive for Windows supports Windows 2000 Server and Windows Server 2003. SnapDrive for UNIX supports a number of UNIX environments.</p> <p>Note: For more information about SnapDrive, see the <i>SnapDrive for Windows Installation and Administration Guide</i> or <i>SnapDrive for UNIX Installation and Administration Guide</i>.</p>
Native tape backup and recovery	<p>Store and retrieve data on tape.</p> <p>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.</p>

Method	Used to...
NDMP	<p>Control native backup and recovery facilities in storage systems and other file servers. Backup application vendors provide a common interface between backup applications and file servers.</p> <p>Note: NDMP is an open standard for centralized control of enterprise-wide data management. For more information about how NDMP-based topologies can be used by storage systems to protect data, see the <i>Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide</i>.</p>

Related information

Data ONTAP documentation on the NAS support site - www.ibm.com/storage/support/nas

LUN clones

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

A LUN clone shares space with the LUN in the backing Snapshot copy. When you clone a LUN, and new data is written to the LUN, the LUN clone still depends on data in the backing Snapshot copy. The clone does not require additional disk space until changes are made to it.

You cannot delete the backing Snapshot copy until you split the clone from it. When you split the clone from the backing Snapshot copy, the data is copied from the Snapshot copy to the clone, thereby removing any dependence on the Snapshot copy. After the splitting operation, both the backing Snapshot copy and the clone occupy their own space.

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

Next topics

[Reasons for cloning LUNs](#) on page 196

[Differences between FlexClone LUNs and LUN clones](#) on page 196

[Cloning LUNs](#) on page 197

[LUN clone splits](#) on page 198

[Displaying the progress of a clone-splitting operation](#) on page 198

[Stopping the clone-splitting process](#) on page 199

[Deleting Snapshot copies](#) on page 199

[Deleting backing Snapshot copies of deleted LUN clones](#) on page 199

Reasons for cloning LUNs

Use LUN clones to create multiple read/write copies of a LUN.

You might want to do this for the following reasons:

- You need to create a temporary copy of a LUN for testing purposes.
- You need to make a copy of your data available to additional users without giving them access to the production data.
- You want to create a clone of a database for manipulation and projection operations, while preserving the original data in unaltered form.
- You want to access a specific subset of a LUN's data (a specific logical volume or file system in a volume group, or a specific file or set of files in a file system) and copy it to the original LUN, without restoring the rest of the data in the original LUN. This works on operating systems that support mounting a LUN and a clone of the LUN at the same time. SnapDrive for UNIX allows this with the `snap connect` command.

Differences between FlexClone LUNs and LUN clones

Data ONTAP provides two LUN cloning capabilities—LUN clone with the support of a Snapshot copy and FlexClone LUN. However, there are a few differences between these two LUN cloning techniques.

The following table lists the key differences between the two LUN cloning features.

FlexClone LUN	LUN clone
To create a FlexClone LUN, you should use the <code>clone start</code> command.	To create a LUN clone, you should use the <code>lun clone create</code> command.
You need not create a Snapshot copy manually.	You need to create a Snapshot copy manually before creating a LUN clone, because a LUN clone uses a backing Snapshot copy
A temporary Snapshot copy is created during the cloning operation. The Snapshot copy is deleted immediately after the cloning operation. However, you can prevent the Snapshot copy creation by using the <code>-n</code> option of the <code>clone start</code> command.	A LUN clone is coupled with a Snapshot copy.
A FlexClone LUN is independent of Snapshot copies. Therefore, no splitting is required.	When a LUN clone is split from the backing Snapshot copy, it uses extra storage space. The amount of extra space used depends on the type of clone split.

FlexClone LUN	LUN clone
You can clone a complete LUN or a sub-LUN. To clone a sub-LUN, you should know the block range of the parent entity and clone entity.	You can only clone a complete LUN.
FlexClone LUNs are best for situations where you need to keep the clone for a long time.	LUN clones are best when you need a clone only for a short time.
No Snapshot copy management is required.	You need to manage Snapshot copies if you keep the LUN clones for a long time.

Cloning LUNs

Use LUN clones to create multiple readable, writable copies of a LUN.

Before you begin

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

About this task

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

Steps

1. Create a LUN by entering the following command:

```
lun create -s size -t lun type lun_path
```

Example

```
lun create -s 100g -t solaris /vol/vol1/lun0
```

2. Create a Snapshot copy of the volume containing the LUN to be cloned by entering the following command:

```
snap create volume_name snapshot_name
```

Example

```
snap create vol1 mysnap
```

3. Create the LUN clone by entering the following command:

```
lun clone create clone_lun_path -bparent_lun_path parent_snap
```

clone_lun_path is the path to the clone you are creating, for example, /vol/vol1/lun0clone.

parent_lun_path is the path to the original LUN.

parent_snap is the name of the Snapshot copy of the original LUN.

Example

```
lun clone create /vol/vol1/lun0clone -b /vol/vol1/lun0 mysnap
```

Result

The LUN clone is created.

LUN clone splits

After you clone a LUN, you can split the clone from the backing Snapshot copy.

The LUN clone split technology was significantly improved to create greater space efficiency. However, note that you must wait until the LUN clone split is complete before you can take additional Snapshot copies.

Splitting the clone from the backing Snapshot copy

If you want to delete the backing Snapshot copy, you can split the LUN clone from the backing Snapshot copy without taking the LUN offline. Any data from the Snapshot copy that the LUN clone depended on is copied to the LUN clone.

Note that you cannot delete the backing Snapshot copy or create a new Snapshot copy until the LUN clone split is complete.

Step

1. Begin the clone split operation by entering the following command:

```
lun clone split start lun_path
```

lun_path is the path to the cloned LUN.

The Snapshot copy can be deleted.

Displaying the progress of a clone-splitting operation

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

Step

1. Enter the following command:

```
lun clone split status lun_path
```

lun_path is the path to the cloned LUN.

Stopping the clone-splitting process

Use the `lun clone split stop` command to stop a clone split that is in progress.

Step

1. Enter the following command:

```
lun clone split stop lun_path
```

lun_path is the path to the cloned LUN.

Deleting Snapshot copies

Once you split the LUN clone from the backing Snapshot copy, you have removed any dependence on that Snapshot copy so it can be safely deleted.

Step

1. Delete the Snapshot copy by entering the following command:

```
snap delete vol-name snapshot-name
```

Example

```
snap delete vol2 snap2
```

Result

The Snapshot copy is deleted.

Deleting backing Snapshot copies of deleted LUN clones

Prior to Data ONTAP 7.3, the system automatically locked all backing Snapshot copies when Snapshot copies of LUN clones were taken. Starting with Data ONTAP 7.3, you can enable the system to only lock backing Snapshot copies for the active LUN clone. If you do this, when you delete the active LUN clone, you can delete the base Snapshot copy without having to first delete all of the more recent backing Snapshot copies.

About this task

This behavior is not enabled by default; use the `snapshot_clone_dependency` volume option to enable it. If this option is set to `off`, you will still be required to delete all subsequent Snapshot copies before deleting the base Snapshot copy.

If you enable this option, you are not required to rediscover the LUNs. If you perform a subsequent volume `snap restore` operation, the system restores whichever value was present at the time the Snapshot copy was taken.

Step

1. Enable this behavior by entering the following command:

```
vol options volume_name snapshot_clone_dependency on
```

Examples of deleting backing Snapshot copies of deleted LUN clones

Use the `snapshot_clone_dependency` option to determine whether you can delete the base Snapshot copy without deleting the more recent Snapshot copies after deleting a LUN clone. This option is set to `off` by default.

Example with `snapshot_clone_dependency` set to `off`

The following example illustrates how all newer backing Snapshot copies must be deleted before deleting the base Snapshot copy when a LUN clone is deleted.

Set the `snapshot_clone_dependency` option to `off` by entering the following command:

```
vol options volume_name snapshot_clone_dependency off
```

Create a new LUN clone, `lun_s1`, from the LUN in Snapshot copy `snap1`. Run the `lun show -v` command to show that `lun_s1` is backed by `snap1`.

```
system1> lun clone create /vol/vol1/lun_s1 -b /vol/vol1/lun snap1

system1> lun show -v
/vol/vol1/lun_s1      47.1m (49351680) (r/w, online)
  Serial#: C4e6SJI0ZqoH
  Backed by: /vol/vol1/.snapshot/snap1/lun
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: windows
  Cluster Shared Volume Information: 0x1
```

Run the `snap list` command to show that `snap1` is busy, as expected.

```
system1> snap list vol1
Volume vol1
working...

  %/used      %/total    date           name
  -----
  24% (24%)   0% ( 0%)   Dec 20 02:40   snap1          (busy, LUNs)
```

When you create a new Snapshot copy, `snap2`, it contains a copy of `lun_s1`, which is still backed by the LUN in `snap1`.

```
system1> snap create vol1 snap2
system1> snap list vol1
Volume vol1
working...
```

```
  %/used      %/total    date           name
  -----
```



```

24% (24%)    0% ( 0%)   Dec 20 02:41   snap2
43% (31%)    0% ( 0%)   Dec 20 02:40   snap1           (busy, LUNs)

```

Run the `lun snap usage` command to show this dependency.

```

system1> lun snap usage vol1 snap1
Active:
    LUN: /vol/vol1/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun
Snapshot - snap2:
    LUN: /vol/vol1/.snapshot/snap2/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun

```

Then delete the LUN clone `lun_s1`.

```

system1> lun destroy /vol/vol1/lun_s1
Wed Dec 20 02:42:23 GMT [wafl.inode.fill.disable:info]: fill
reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [wafl.inode.overwrite.disable:info]:
overwrite reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [lun.destroy:info]: LUN /vol/vol1/lun_s1
destroyed

```

```

system1> lun show
/vol/vol1/lun           30m (31457280)      (r/w,
online)

```

Run the `lun snap usage` command to show that `snap2` still has a dependency on `snap1`.

```

system1> lun snap usage vol1 snap1
Snapshot - snap2:
    LUN: /vol/vol1/.snapshot/snap2/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun

```

Run the `snap list` command to show that `snap1` is still busy.

```

system1> snap list vol1
Volume vol1
working...

```

%/used	%/total	date	name
39% (39%)	0% (0%)	Dec 20 02:41	snap2
53% (33%)	0% (0%)	Dec 20 02:40	snap1 (busy, LUNs)

Since `snap1` is still busy, you cannot delete it until you delete the more recent Snapshot copy, `snap2`.

Example with `snapshot_clone_dependency` set to on

The following example illustrates how you can delete a base Snapshot copy without deleting all newer backing Snapshot copies when a LUN clone is deleted.

Set the `snapshot_clone_dependency` option to on by entering the following command:

vol options volume_name snapshot_clone_dependency on

Create a new LUN clone, `lun_s1`, from the LUN in Snapshot copy `snap1`. Run the `lun show -v` command to show that `lun_s1` is backed by `snap1`.

```
system1> lun clone create /vol/vol1/lun_s1 -b /vol/vol1/lun snap1

system1> lun show -v
/vol/vol1/lun_s1      47.1m (49351680) (r/w, online)
  Serial#: C4e6SJ10ZqoH
  Backed by: /vol/vol1/.snapshot/snap1/lun
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: windows
  Cluster Shared Volume Information: 0x1
```

Run the `snap list` command to show that `snap1` is busy, as expected.

```
system1> snap list vol1
Volume vol1
working...

  %/used      %/total    date           name
  -----
  24% (24%)    0% ( 0%)    Dec 20 02:40    snap1           (busy, LUNs)
```

When you create a new Snapshot copy, `snap2`, it contains a copy of `lun_s1`, which is still backed by the LUN in `snap1`.

```
system1> snap create vol1 snap2
system1> snap list vol1
Volume vol1
working...

  %/used      %/total    date           name
  -----
  24% (24%)    0% ( 0%)    Dec 20 02:41    snap2
  43% (31%)    0% ( 0%)    Dec 20 02:40    snap1           (busy, LUNs)
```

Run the `lun snap usage` command to show this dependency.

```
system1> lun snap usage vol1 snap1
Active:
  LUN: /vol/vol1/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
Snapshot - snap2:
  LUN: /vol/vol1/.snapshot/snap2/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
```

Then delete the LUN clone `lun_s1`.

```
system1> lun destroy /vol/vol1/lun_s1
Wed Dec 20 02:42:23 GMT [waf1.inode.fill.disable:info]: fill
reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [waf1.inode.overwrite.disable:info]:
overwrite reservation disabled for inode 3087 (vol vol1).
```

```
Wed Dec 20 02:42:23 GMT [lun.destroy:info]: LUN /vol/vol1/lun_s1
destroyed
```

```
system1> lun show
          /vol/vol1/lun                30m (31457280)      (r/w,
online)
```

Run the `lun snap usage` command to show that `snap2` still has a dependency on `snap1`.

```
system1> lun snap usage vol1 snap1
Snapshot - snap2:
    LUN: /vol/vol1/.snapshot/snap2/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun
```

Run the `snap list` command to show that `snap1` is no longer busy.

```
system1> snap list vol1
Volume vol1
working...
```

%/used	%/total	date	name
39% (39%)	0% (0%)	Dec 20 02:41	snap2
53% (33%)	0% (0%)	Dec 20 02:40	snap1

Since `snap1` is no longer busy, you can delete it without first deleting `snap2`.

```
system1> snap delete vol1 snap1
Wed Dec 20 02:42:55 GMT [wafl.snap.delete:info]: Snapshot copy snap1
on volume vol1 was deleted by the Data ONTAP function snapcmd_delete.
The unique ID for this Snapshot copy is (1, 6).
```

```
system1> snap list vol1
Volume vol1
working...
```

%/used	%/total	date	name
38% (38%)	0% (0%)	Dec 20 02:41	snap2

Deleting busy Snapshot copies

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

About this task

Use the `lun snap usage` command to list all the LUNs backed by data in the specified Snapshot copy. It also lists the corresponding Snapshot copies in which these LUNs exist.

The `lun snap usage` command displays the following information:

- LUN clones that are holding a lock on the Snapshot copy given as input to this command
- Snapshots in which these LUN clones exist

Steps

1. Identify all Snapshot copies that are in a busy state, locked by LUNs, by entering the following command:

```
snap list vol-name
```

Example

```
snap list vol2
```

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 Snapshot copies that contain them by entering the following command:

```
lun snap usage [-s] vol_name snap_name
```

Use the `-s` option to only display the relevant backing LUNs and Snapshot copies that must be deleted.

Note: The `-s` option is particularly useful in making SnapDrive output more readable. For example:

```
lun snap usage -s vol2 snap0
You need to delete the following snapshots before deleting snapshot
"snap0":
/vol/vol1/.snapshot/snap1
/vol/vol2/.snapshot/snap2
```

Example

```
lun snap usage vol2 snap0
```

The following message is displayed:

```
active:
  LUN:          /vol/vol2/lunC
  Backed By:    /vol/vol2/.snapshot/snap0/lunA
snap2:
  LUN:          /vol/vol2/.snapshot/snap2/lunB
  Backed By:    /vol/vol2/.snapshot/snap0/lunA
```

```

snap1:
  LUN:          /vol/vol1/.snapshot/snap1/lunB
  Backed By:    /vol/vol2/.snapshot/snap0/lunA

```

Note: The LUNs are backed by lunA in the snap0 Snapshot copy.

In some cases, the path for LUN clones backed by a Snapshot copy cannot be determined. In those instances, a message is displayed so that those Snapshot copies can be identified. You must still delete these Snapshot copies in order to free the busy backing Snapshot copy. For example:

```
lun snap usage vol2 snap0
```

```

Snapshot - snap2:
  LUN: Unable to determine the path of the LUN
  Backed By: Unable to determine the path of the LUN
  LUN:          /vol/vol2/.snapshot/snap2/lunB
  Backed By:    /vol/vol2/.snapshot/snap0/lunA

```

3. Delete all the LUNs in the active file system that are displayed by the `lun snap usage` command by entering the following command:

```
lun destroy [-f] lun_path [lun_path ...]
```

Example

```
lun destroy /vol/vol2/lunC
```

4. Delete all the Snapshot copies that are displayed by the `lun snap usage` command in the order they appear, by entering the following command:

```
snap delete vol-name snapshot-name
```

Example

```
snap delete vol2 snap2
```

```
snap delete vol2 snap1
```

All the Snapshot copies containing lunB are now deleted and snap0 is no longer busy.

5. Delete the Snapshot copy by entering the following command:

```
snap delete vol-name snapshot-name
```

Example

```
snap delete vol2 snap0
```

Restoring a Snapshot copy of a LUN in a volume

You can use SnapRestore to restore a Snapshot copy of a LUN and the volume that contains it to its state when the Snapshot copy was taken. You can use SnapRestore to restore an entire volume or a single LUN.

Before you begin

Before using SnapRestore, you must perform the following tasks:

- Always unmount the LUN before you run the `snap restore` command on a volume containing the LUN or before you run a single file SnapRestore of the LUN. For a single file SnapRestore, you must also take the LUN offline.
- Check available space; SnapRestore does not revert the Snapshot copy if sufficient space is unavailable.

About this task

When restoring a volume using SnapRestore, you only need as much available space as the size of the volume you are restoring. For example, if you are restoring a 10 Gb volume, then you only need 10 Gb of available space to perform the SnapRestore.

Attention: When a single LUN is restored, it must be taken offline or be unmapped prior to recovery. Using SnapRestore on a LUN, or on a volume that contains LUNs, without stopping all host access to those LUNs, can cause data corruption and system errors.

Steps

1. From the host, stop all host access to the LUN.
2. From the host, if the LUN contains a host file system mounted on a host, unmount the LUN on that host.
3. From the storage system, unmap the LUN by entering the following command:

```
lun unmap lun_path initiator-group
```

4. Enter the following command:

```
snap restore [-f] [-t vol] volume_name [-s snapshot_name]
```

`-f` suppresses the warning message and the prompt for confirmation. This option is useful for scripts.

`-t vol volume_name` specifies the volume name to restore.

`volume_name` is the name of the volume to be restored. Enter the name only, not the complete path. You can enter only one volume name.

`-s snapshot_name` specifies the name of the Snapshot copy from which to restore the data. You can enter only one Snapshot copy name.

Example

```
snap restore -s payroll_lun_backup.2 -t vol /vol/payroll_lun
```

```
storage_system> 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
```

If you did not use the `-f` option, Data ONTAP displays a warning message and prompts you to confirm your decision to restore the volume.

5. Press **y** to confirm that you want to restore the volume.

Data ONTAP displays the name of the volume and the name of the Snapshot copy for the reversion. If you did not use the `-f` option, Data ONTAP prompts you to decide whether to proceed with the reversion.

6. Decide if you want to continue with the reversion.
 - If you want to continue the reversion, press **y**. The storage system reverts the volume from the selected Snapshot copy.
 - If you do not want to continue the reversion, press **n** or **Ctrl-C**. The volume is not reverted and you are returned to a storage system prompt.
7. Enter the following command to unmap the existing old maps that you do not want to keep.


```
lun unmap lun_path initiator-group
```
8. Remap the LUN by entering the following command:


```
lun map lun_path initiator-group
```
9. From the host, remount the LUN if it was mounted on a host.
10. From the host, restart access to the LUN.
11. From the storage system, bring the restored LUN online by entering the following command:


```
lun online lun_path
```

After you finish

After you use SnapRestore to update a LUN from a Snapshot copy, you also need to restart any applications you closed down and remount the volume from the host side.

Restoring a single LUN

You can use SnapRestore to restore a single LUN without restoring the volume that contains it.

Steps

1. Notify users that you are going to restore a LUN so that they know that the current data in the LUN will be replaced by that of the selected Snapshot copy.
2. Enter the following command:

```
snap restore [-f] [-t file] [-s snapshot_name] [-r restore_as_path]
path_and_LUN_name
```

`-f` suppresses the warning message and the prompt for confirmation.

`-t file` specifies that you are entering the name of a file to revert.

`-s snapshot_name` specifies the name of the Snapshot copy from which to restore the data.

`-r restore_as_path` restores the file to a location in the volume different from the location in the Snapshot copy. For example, if you specify `/vol/vol0/vol3/mylun` as the argument to `-r`, SnapRestore restores the file called `mylun` to the location `/vol/vol0/vol3` instead of to the path structure indicated by the path in `path_and_lun_name`.

`path_and_LUN_name` is the complete path to the name of the LUN to be restored. You can enter only one path name.

A LUN can be restored only to the volume where it was originally. The directory structure to which a LUN is to be restored must be the same as specified in the path. If this directory structure no longer exists, you must re-create it before restoring the file.

Unless you enter `-r` and a path name, only the LUN at the end of the `path_and_lun_name` is reverted.

If you did not use the `-f` option, Data ONTAP displays a warning message and prompts you to confirm your decision to restore the LUN.

3. Type the following character to confirm that you want to restore the file:

y

Data ONTAP displays the name of the LUN and the name of the Snapshot copy for the restore operation. If you did not use the `-f` option, Data ONTAP prompts you to decide whether to proceed with the restore operation.

4. Type the following character to continue with the restore operation:

y

Data ONTAP restores the LUN from the selected Snapshot copy.

Example of a single LUN restore

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

```
storage_system> 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
```

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

After a LUN is restored with SnapRestore, all data and all relevant user-visible attributes for that LUN in the active file system are identical to that contained in the Snapshot copy.

Backing up SAN systems to tape

In most cases, backup of SAN systems to tape takes place through a separate backup host to avoid performance degradation on the application host. It is imperative that you keep SAN and NAS data separated for backup purposes.


Before you begin

The following procedure assumes that you have already performed the following tasks:

- Created the production LUN
- Created the igroup to which the LUN will belong
The igroup must include the WWPN of the application server.
- Mapped the LUN to the igroup
- Formatted the LUN and made it accessible to the host

About this task

Configure volumes as SAN-only or NAS-only and configure qtrees within a single volume as SAN-only or NAS-only. From the point of view of the SAN host, LUNs can be confined to a single WAFL volume or qtree or spread across multiple WAFL volumes, qtrees, or storage systems.



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

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

Step 1 can be part of your SAN backup application's pre-processing script. Steps 5 and 6 can be part of your SAN backup application's post-processing script.

Steps

1. When you are ready to start the backup (usually after your application has been running for some time in your production environment), save the contents of host file system buffers to disk using the command provided by your host operating system, or by using SnapDrive for Windows or SnapDrive for UNIX.
2. Create a Snapshot copy by entering the following command:

```
snap create volume_name snapshot_name
```

Example

```
snap create vol1 payroll_backup
```

3. To create a clone of the production LUN, enter the following command:

```
lun clone create clone_lunpath -b parent_lunpath parent_snap
```

Example

```
lun clone create /vol/vol1/qtrees_1/payroll_lun_clone -b /vol/vol1/  
qtrees_1/payroll_lun payroll_backup
```

4. Create an igroup that includes the WWPN of the backup server by entering the following command:

```
igroup create -f -t ostype group [node ...]
```

Example

```
igroup create -f -t windows backup_server 10:00:00:00:d3:6d:0f:e1
```

Data ONTAP creates an igroup that includes the WWPN (10:00:00:00:d3:6d:0f:e1) of the Windows backup server.

5. To map the LUN clone you created in Step 3 to the backup host, enter the following command:

```
lun map lun_path initiator-group LUN_ID
```

Example

```
lun map /vol/vol1/qtrees_1/payroll_lun_clone backup_server 1
```

Data ONTAP maps the LUN clone (/vol/vol1/qtrees_1/payroll_lun_clone) to the igroup called backup_server with a SCSI ID of 1.

6. From the host, discover the new LUN and make the file system available to the host.
7. Back up the data in the LUN clone from the backup host to tape by using your SAN backup application.
8. Take the LUN clone offline by entering the following command: `lun offline /vol/vol_name/qtrees_name/lun_name`

Example

```
lun offline /vol/vol1/qtrees_1/payroll_lun_clone
```

9. Remove the LUN clone by entering the following command: `lun destroy lun_path`

Example

```
lun destroy /vol/vol1/qtrees_1/payroll_lun_clone
```

10. Remove the Snapshot copy by entering the following command:

```
snap delete volume_name lun_name
```

Example

```
snap delete vol1 payroll_backup
```

Using volume copy to copy LUNs

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

Before you begin

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

Note: The term *LUNs* in this context refer to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

About this task

The `vol copy` command enables you to copy data from one WAFL volume to another, either within the same storage system or to a different storage system. The result of the `vol copy` command is a restricted volume containing the same data that was on the source storage system at the time you initiate the copy operation.

Step

1. To copy a volume containing a LUN to the same or different storage system, enter the following command:

```
vol copy start -S source:source_volume dest:dest_volume
```

`-S` copies all Snapshot copies in the source volume to the destination volume. If the source volume has Snapshot copy-backed LUNs, you must use the `-S` option to ensure that the Snapshot copies are copied to the destination volume.

If the copying takes place between two storage systems, you can enter the `vol copy start` command on either the source or destination storage system. You cannot, however, enter the command on a third storage system that does not contain the source or destination volume.

Example

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

Index

A

- access lists 103–105
 - about 103
 - creating 104
 - displaying 105
 - removing interfaces from 104
- adapters 148, 151, 158, 160, 162, 164–166, 168
 - changing the speed for 148
 - changing the WWPN for 151
 - configuring for initiator mode 160
 - configuring for target mode 158
 - displaying brief target adapter information 165
 - displaying detailed target adapter information 166
 - displaying information about all 164
 - displaying information for FCP 162
 - displaying statistics for target adapters 168
- aggregate 39
 - defined 39
- aliases 153
 - for WWPNs 153
- ALUA 24, 88, 123
 - automatic enablement of 88
 - defined 24
 - enabling 88
 - igroup 88
 - manually enabling 88
 - setting the priority of target portal groups for 123
- authentication 110, 112
 - defining default for CHAP 112
 - using CHAP for iSCSI 110
- autodelete 43, 49, 50
 - configuring volumes and LUNs with 49
 - setting volume options for 50
 - when to use 50

B

- backing up SAN systems 209
- best practices 44
 - storage provisioning 44

C

- CHAP 29, 110, 112, 113, 116
 - and RADIUS 116

- authenticate 113
 - iSCSI initiator 113
- authentication for iSCSI 110
 - defined 29
 - defining default authentication 112
 - using with vFiler units 110
- cluster failover 137–140
 - avoiding igroup mapping conflicts with 138
 - multipathing requirements for 140
 - overriding mapping conflicts 139
 - understanding 137
- configure LUNs 49, 50
 - autodelete 49, 50
- configure volumes 49, 50
 - autodelete 49, 50
- create_ucose option 54
 - changing with the command line 54
- cutover phase 183
 - cutover attempts 183
 - volume move 183

D

- data center bridging 35
 - defined 35
- data copy phase 182
 - volume move 182
- Data ONTAP options 95, 96, 98, 106
 - automatically enabled 98
 - iscsi.isns.rev 106
 - iscsi.max_connections_per_session 95
 - iscsi.max_error_recovery_level 96
- DataMotion for Volumes 180
 - about 180
- DCB 35
 - defined 35
- DCB (data center bridging) switch 35
 - for FCoE 35
- df command 174
 - monitoring disk space using 174
- disk 174
 - space information, displaying 174
- disk space 174, 176
 - monitoring with Snapshot copies 176
 - monitoring without Snapshot copies 174

E

- enabling 88
 - ALUA 88
- error recovery level 96
 - enabling levels 1 and 2 96
- eui type designator 28
- extended copy feature 187–190
 - environment 188
 - invoked automatically 188
 - statistics collected 189
 - VAAI feature 187
 - viewing statistics 190
 - when the standard copy operation is used 187

F

- FC 76, 137, 148, 158
 - changing the adapter speed 148
 - checking interfaces 76
 - managing in HA pairs 137
 - managing systems with onboard adapters 158
- FCoE 34, 35
 - data center bridging 35
 - target adapters 34
- FCP 32, 33, 148, 152, 162
 - changing the WWNN 152
 - defined 32
 - displaying adapters 162
 - host nodes 33
 - how nodes are connected 32
 - how nodes are identified 32
 - noded defined 32
 - storage system nodes 33
 - switch nodes 33
 - taking adapters offline and online 148
- FCP commands 146–148, 151, 162
 - fc config 148, 162
 - fc nodename 162
 - fc portname set 151
 - fc show 162
 - fc start 147
 - fc stats 162
 - fc status 146
 - fc stop 147
 - license 146
 - license add 146
 - license delete 147
 - storage show adapter 162
- fc ping 157

- connectivity 157
 - fabric latency 157
- FCP service 146, 147, 169, 170
 - disabling 147
 - displaying how long running 170
 - displaying partner's traffic information 170
 - displaying statistics for 170
 - displaying traffic information about 169
 - licensing 146
 - starting and stopping 147
 - verifying the service is licensed 146
 - verifying the service is running 146
- Fibre Channel over Ethernet (FCoE) 35
 - overview 35
- FlexClone files and FlexClone LUNs 196
 - differences between FlexClone LUNs and LUN clones 196
- flexible volumes 39
 - described 39
- FlexVol volumes 178, 179
 - automatic free space preservation, configuring 179
 - automatically adding space for 178
 - automatically grow, configuring to 179
 - try_first volume option 178
- fractional reserve 42, 43
 - about 42
- free space 178
 - automatically increasing 178

H

- HA pairs 30, 130, 137
 - and controller failover 137
 - and iSCSI 30
 - using with iSCSI 130
- HBA 168
 - displaying information about 168
- head swap 151
 - changing WWPNs 151
- host bus adapters 168
 - displaying information about 168
- Host Utilities 23
 - defined 23

I

- igroup commands 66, 85–89
 - for vFiler units 89
 - igroup add 86
 - igroup create 66

- igroup destroy 85
- igroup remove 86
- igroup rename 87
- igroup set 87
- igroup set alua 88
- igroup show 87
- igroup commands for iSCSI 83
 - igroup create 83
- igroup mapping conflicts 138
 - avoiding during cluster failover 138
- igroup throttles 90–93, 141
 - borrowing queue resources 91
 - creating 91
 - defined 90
 - destroying 91
 - displaying information about 92
 - displaying LUN statistics for 93
 - displaying usage information 92
 - how Data ONTAP uses 90
 - how port sets affect 141
 - how to use 90
- igroups 91
 - borrowing queue resources for 91
- initiator groups 59–61, 73, 83, 84, 86, 87, 142, 145
 - adding 86
 - binding to portsets 142
 - creating for FCP using sanlun 84
 - creating for iSCSI 83
 - defined 59
 - displaying 87
 - name rules 60
 - naming 60
 - ostype of 61
 - renaming 87
 - requirements for creation 60
 - setting the ostype for 87
 - showing portset bindings 145
 - type of 61
 - unmapping LUNs from 73
- initiator, displaying for iSCSI 109
- initiators 160
 - configuring adapters as 160
- interface 102
 - disabling for iSCSI 102
 - enabling for iSCSI 102
- intiator groups 85, 86
 - destroying 85
 - removing initiators from 86
- IP addresses, displaying for iSCSI 103
- iqn type designator 27
- iSCSI 25–28, 30, 95–100, 102–105, 109, 110, 113, 119, 121, 122, 124, 128–130, 132
 - access lists 103
 - connection, displaying 129
 - creating access lists 104
 - creating target portal groups 121
 - default TCP port 28
 - destroying target portal groups 122
 - displaying access lists 105
 - displaying initiators 109
 - displaying statistics 124
 - enabling error recovery levels 1 and 2 96
 - enabling on interface 102
 - explained 25
 - how communication sessions work 30
 - how nodes are identified 27
 - implementation on the host 26
 - implementation on the storage system 26
 - iSNS 105
 - license 97
 - multi-connection sessions, enabling 95
 - node name rules 99
 - nodes defined 26
 - RADIUS 113
 - removing interfaces from access lists 104
 - security 110
 - service, verifying 97
 - session, displaying 128
 - setup procedure 30
 - supported configurations 26
 - target alias 100
 - target IP addresses 103
 - target node name 98
 - target portal groups defined 28, 119
 - troubleshooting 132
 - using with HA pairs 30
 - with HA pairs 130
- iscsi commands 97, 98, 100, 101, 103, 107, 109, 111, 121, 124, 128, 129
 - iscsi alias 100
 - iscsi connection 129
 - iscsi initiator 109
 - iscsi interface 101
 - iscsi isns 107
 - iscsi nodename 98
 - iscsi portal 103
 - iscsi security 111
 - iscsi session 128
 - iscsi start 98
 - iscsi stats 124

- iscsi status 97
- iscsi stop 98
- iscsi tpgroup 121
- iscsi.isns.rev option 106
- iscsi.max_connections_per_session option 95
- iscsi.max_error_recovery_level option 96
- iSNS 29, 105, 106, 108
 - defined 29
 - disabling 108
 - server versions 106
 - service for iSCSI 105
 - updating immediately 108
 - with vFiler units 108
- ISNS 107
 - and IPv6 107
 - registering 107

L

- license 97
 - iSCSI 97
- LUN clones 195–199
 - creating 197
 - defined 195
 - deleting Snapshot copies 199
 - displaying progress of split 198
 - reasons for using 196
 - splitting from Snapshot copy 198
 - stopping split 199
- lun commands 65, 66, 71–76, 78, 79
 - lun config_check 76
 - lun destroy 75
 - lun help 71
 - lun map 66
 - lun move 74
 - lun offline 73
 - lun online 72
 - lun set reservation 75
 - lun setup 65
 - lun share 76
 - lun show 79
 - lun stats 78
 - lun unmap 73
- LUN commands 85, 197, 198, 203
 - lun clone create 197
 - lun clone split 198
 - lun snap usage 203
 - lun unmap 85
- LUN creation 56, 58, 59
 - description attribute 58

- host operating system type 56
- LUN ID requirement 59
- path name 56
- size specifiers 58
- space reservation default 59
- LUN ID 62
 - ranges of 62
- LUN serial numbers 78
 - displaying 78
 - changing 78
- LUNs 50, 56, 63, 71–76, 78, 79, 93, 208
 - autosize 50
 - bringing online 72
 - checking settings for 76
 - controlling availability 72
 - displaying mapping 79
 - displaying reads, writes, and operations for 78
 - displaying serial numbers for 78
 - enabling space reservations 75
 - host operating system type 56
 - management task list 71
 - mapping guidelines 63
 - modifying description 74
 - multiprotocol type 56
 - read-only 63
 - removing 75
 - renaming 74
 - restoring 208
 - snapshot copies 50
 - snapshot copy 50
 - space reserved 50
 - statistics for igroup throttles 93
 - taking offline 73
 - unmapping from initiator group 73

M

- manually enabling ALUA 88
- mapping conflicts 139
 - overriding 139
- moving volumes 180
 - DataMotion for Volumes 180
- multi-connection sessions 95
 - enabling 95
- multipathing 140
 - requirements for cluster failover 140
- Multiprotocol type 56
- MultiStore 67
 - creating LUNs for vFiler units 67

N

- name rules 60, 99
 - igroups 60
 - iSCSI node name 99
- node name 28, 99
 - rules for iSCSI 99
 - storage system 28
- node type designator 27, 28
 - eui 28
 - iqn 27
- nodes 26, 32
 - FCP 32
 - iSCSI 26

O

- onboard adapters 158
 - configuring for target mode 158
- options 95, 96, 98, 106
 - automatically enabled 98
 - iscsi.isns.rev 106
 - iscsi.max_connections_per_session 95
 - iscsi.max_error_recovery_level 96
- ostype 87
 - setting 87

P

- plex 39
 - defined 39
- porset commands 142
 - portset create 142
- port sets 140
 - defined 140
- portset commands 143–145
 - portset add 143
 - portset destroy 144
 - portset remove 144
 - portset show 145
- portsets 141–145
 - adding ports 143
 - binding to igroups 142
 - creating 142
 - destroying 144
 - how they affect igroup throttles 141
 - how upgrades affect 141
 - removing 144
 - showing igroup bindings 145
 - unbinding igroups 143

- viewing ports in 145
- provisioning 43, 44
 - guidelines for 44
 - methods of 43

Q

- qtrees 39
 - defined 39

R

- RADIUS 113–118
 - adding a RADIUS server 116
 - clearing statistics for 118
 - defining as the authentication method 114
 - displaying statistics for 118
 - displaying the status of 117
 - enabling for CHAP authentication 116
 - overview 113
 - removing a RADIUS server 118
 - server 113
 - client service 113
 - starting the client service 115
 - stopping the service 117
- RAID-level mirroring 39
 - described 39
- restoring 208
 - LUNs 208
- resuming volume move 185
 - data copy phase 185

S

- SAN systems 209
 - backing up 209
- sanlun 84
 - creating igroups for FCP 84
- SCSI command 88
- serial numbers 78
 - for LUNs 78
- setup phase 182
 - volume move 182
- snap commands 206
 - snap restore 206
- snap reserve 52
 - setting the percentage 52
- SnapDrive 24
 - about 24

- SnapMirror destinations 63
 - mapping read-only LUNs to hosts at 63
- Snapshot copies 52, 203
 - deleting busy 203
 - schedule, turning off 52
- space reservations 41
 - about 41
- statistics 124, 189
 - collected for VAAI features 189
 - displaying for iSCSI 124
- stats command 190
 - viewing statistics for VAAI features 190
- storage system node name 28
 - defined 28
- storage units 39
 - types of 39
- SyncMirror 39
 - use of plexes in 39

T

- target adapter 167
 - displaying WWNN 167
- target adapters 34, 168
 - displaying statistics for 168
 - FCoE 34
- target alias for iSCSI 100
- target node name, iSCSI 98
- target port group support 24
- target port groups 24
- target portal groups 28, 119, 121–123
 - about 119
 - adding interfaces 122
 - creating 121
 - defined 28
 - destroying 122
 - removing interfaces 123
- targets 158
 - configuring adapters as 158
- TCP port 28
 - default for iSCSI 28
- traditional volumes 39
 - described 39
- troubleshooting iSCSI 132
- try_first volume option 178

U

- unified target adapters 34, 35
 - data center bridging 35

- managing 34

V

- VAAI features 187, 189, 190
 - extended copy feature 187
 - statistics collected 189
 - VERIFY AND WRITE feature 187
 - viewing statistics 190
 - WRITE SAME feature 187
- VERIFY AND WRITE feature 187–190
 - environment 188
 - invoked automatically 188
 - statistics collected 189
 - VAAI feature 187
 - viewing statistics 190
- vFiler units 67, 89, 108, 110
 - authentication using CHAP 110
 - creating LUNs for 67
 - using iSCSI igroups with 89
 - with iSNS 108
- volume move 180–187
 - abort 187
 - about 180
 - automatic cutover 186
 - cancel 187
 - conflicting operations 181
 - cutover phase 183
 - temporary destination volume 183
 - data transfer 186
 - DataMotion for Volumes 180
 - destination volume 180, 184
 - high priority 185
 - I/O operations 185
 - manual cutover 186
 - operations supported 181
 - operations unsupported 181
 - pausing 185
 - requirements 181, 182
 - resuming volume move 185
 - setup phase 182
 - SLA requirements 180
 - source volume 184
 - temporary volume 184
 - volume status 186
- volumes 43, 47, 50, 178, 180
 - automatically adding space for 178
 - estimating required size of 43, 47
 - moving nondisruptively 180
 - snap_delete 50

space reservation 50

W

WRITE SAME feature 187–190

environment 188

invoked automatically 188

statistics collected 189

VAAI feature 187

viewing statistics 190

WWNN 152, 167

changing 152

displaying for a target adapter 167

WWPN 32, 33, 151

changing for a target adapter 151

creating igroups with 32

how they are assigned 33

WWPN aliases 153, 154

about 153

creating 153

displaying 154

removing 154



NA 210-05012_A0, Printed in USA

GA32-0725-02

