Data ONTAP® 7.3 MultiStore Management Guide

NetApp, Inc. 495 East Java Drive Sunnyvale, CA 94089 USA Telephone: +1 (408) 822-6000 Fax: +1 (408) 822-4501 Support telephone: +1 (888) 4-NETAPP Documentation comments: doccomments@netapp.com Information Web: http://www.netapp.com

Part number: 210-04855_A0 Updated for Data ONTAP 7.3.3 on 4 March 2010

Contents

Trademark information 11 About this guide 13 Audience 13 Audience 13 Accessing Data ONTAP man pages 13 Terminology 14 Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for disaster recovery and data migration 22 Number of vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 Hosting storage system tasks 24 MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 <tr< th=""><th>Copyright information</th><th></th></tr<>	Copyright information	
About this guide 13 Audience 13 Audience 13 Accessing Data ONTAP man pages 13 Terminology 14 Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 Benefits of or consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 SAN guidelines 29<	Trademark information	
Audience 13 Accessing Data ONTAP man pages 13 Terminology 14 Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore for creating vFiler units 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 Active/active configu	About this guide	
Accessing Data ONTAP man pages 13 Terminology 14 Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for disaster recovery and data migration 22 Number of vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 Hosting the MultiStore license 26 Disabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The default vFiler unit 31 Manually setting up a vFiler unit 31	Audience	
Terminology 14 Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The default the network interface is ready 30 Creating a vFiler unit 31	Accessing Data ONTAP man pages	
Where to enter commands 15 Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 SAN guidelines 25 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31	Terminology	
Keyboard and formatting conventions 16 Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 Hosting storage system tasks 24 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 25 SAN guidelines 25 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32 <	Where to enter commands	
Special messages 17 How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31	Keyboard and formatting conventions	
How to send your comments 17 What MultiStore is 19 Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31	Special messages	
What MultiStore is	How to send your comments	
Benefits of using MultiStore 19 MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	What MultiStore is	
MultiStore for consolidating servers 20 MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 31 Manually setting up a vFiler unit 32	Benefits of using MultiStore	19
MultiStore for service providers and enterprises 21 MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	MultiStore for consolidating servers	
MultiStore for disaster recovery and data migration 22 The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	MultiStore for service providers and enterprises	
The default vFiler unit 22 Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	MultiStore for disaster recovery and data migration	
Number of vFiler units allowed 22 Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	The default vFiler unit	
Data access from the hosting storage system 23 Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Number of vFiler units allowed	
Hosting storage system tasks 23 MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Data access from the hosting storage system	
MultiStore management 25 Enabling the MultiStore license 26 Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Hosting storage system tasks	
Enabling the MultiStore license26Disabling the MultiStore license27Prerequisites for creating vFiler units27Storage guidelines28Language guidelines28Quota guidelines29Active/active configuration guidelines29SAN guidelines29The vFiler commands30Ensuring that the network interface is ready30Creating a vFiler unit31Manually setting up a vFiler unit32	MultiStore management	
Disabling the MultiStore license 27 Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Enabling the MultiStore license	
Prerequisites for creating vFiler units 27 Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Disabling the MultiStore license	
Storage guidelines 28 Language guidelines 28 Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Prerequisites for creating vFiler units	
Language guidelines28Quota guidelines29Active/active configuration guidelines29SAN guidelines29The vFiler commands29The vFiler commands30Ensuring that the network interface is ready30Creating a vFiler unit31Manually setting up a vFiler unit32	Storage guidelines	
Quota guidelines 29 Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Language guidelines	
Active/active configuration guidelines 29 SAN guidelines 29 The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	Quota guidelines	
SAN guidelines	Active/active configuration guidelines	
The vFiler commands 30 Ensuring that the network interface is ready 30 Creating a vFiler unit 31 Manually setting up a vFiler unit 32	SAN guidelines	
Ensuring that the network interface is ready	The vFiler commands	
Creating a vFiler unit	Ensuring that the network interface is ready	
Manually setting up a vFiler unit	Creating a vFiler unit	
	Manually setting up a vFiler unit	

vFiler unit storage management from the hosting storage system	. 34
Effects of adding, removing, and moving vFiler unit resources	. 34
Adding resources to a vFiler unit	. 35
Requirements for moving and removing resources	. 35
Removing resources from a vFiler unit	. 35
Moving resources between vFiler units	. 36
Default limits on the number of vFiler units	. 37
Maximum vFiler units allowed	. 37
Viewing the current limit on the number of vFiler units	. 38
Increasing the vFiler unit limit	. 38
Decreasing the vFiler unit limit	. 39
What the vfiler rename command does	. 39
Renaming a vFiler unit	. 40
Stopping a vFiler unit	. 40
Destroying a vFiler unit	. 41
Restoring a vFiler unit	. 42
Starting a vFiler unit	. 42
Protocols supported by a vFiler unit	. 43
Allowing a protocol on a vFiler unit	. 44
Effects of disallowing protocols on a vFiler unit	. 44
Disallowing a protocol on a vFiler unit	. 45
Displaying the vFiler unit status	. 45
Viewing commands that can be executed from a vFiler unit	. 46
Executing commands from a vFiler unit	. 46
Executing commands from the hosting storage system	. 47
Executing RSH commands for a vFiler unit	. 47
Executing SSH commands for a vFiler unit	. 48
List of RSH and SSH commands	. 49
Effects of storage system reboot on a vFiler unit	. 50
Volumes and qtrees on a vFiler unit	. 50
Effects of taking a vFiler unit volume offline	. 51
Changes required after volumes are renamed	. 51
Who can change qtree security styles and oplock settings	. 51
Differences in qtree command output	. 51
Viewing all qtrees and the owner vFiler units	. 52
Backup of vFiler units	. 52

NDMP support	. 53
Available NDMP options	. 53
Support for the ndmpcopy command	. 53
NDMP command support	. 53
NDMP password support	. 54
LUNs on a vFiler unit	. 54
iSCSI LUNs and igroups on a vFiler unit	. 54
The iSCSI service on a vFiler unit	. 55
LUN and igroup limitations on vFiler units	. 56
Networking guidelines	. 57
The routed daemon on vFiler units	. 57
Command for changing the routing table in the default IPspace	. 57
The /etc/dgateways file	. 57
IPsec on a vFiler unit	. 57
SnapMirror on the hosting storage system	. 58
Guidelines for using SnapMirror	. 58
Determining the status of SnapMirror relationships	. 59
Deduplication support on vFiler units	. 59
Running deduplication commands on a vFiler unit	. 60
How MultiStore works with FlexClone files and FlexClone LUNs	. 61
SnapVault on the hosting storage system	. 62
Where to enter SnapVault commands	. 62
Features and limitations of the snapvault command	. 63
Determining the status of SnapVault relationships	. 63
SNMP support on vFiler units	. 63
vFiler unit data from MIBs	. 64
Monitoring performance and statistics	. 64
Viewing storage system statistics	. 64
Viewing uptime statistics	. 65
Viewing NFS statistics	. 65
Viewing CIFS statistics	. 65
What an IPspace is	67
Guidelines for vFiler unit participation in an IPspace	. 67
IPspace application scenario	. 68
Interface participation in an IPspace	. 69
Routing in an IPspace	. 70

Advantages of using VLAN tagging for IPspaces	70
Active/active configuration and IPspaces	71
IPspace naming requirement	71
IPspace assignment requirement	71
Asymmetric active/active configuration setup	71
Specifying partners in an asymmetric active/active configuration setup	72
Creating an IPspace	73
IPspace and the routed daemon	73
Listing IPspaces on a storage system	73
Removing an IP address from an interface	74
Assigning an interface to an IPspace	75
Destroying IPspaces	75
Creating a vFiler unit in a nondefault IPspace	76
File system access using NFS and CIFS	79
Path name specification for NFS exports or CIFS shares	79
vFiler unit preparation for NFS	80
Starting the NFS protocol	80
Exporting all file systems in /etc/exports	81
vFiler unit preparation for CIFS	81
Commands run from the hosting storage system	81
Local user accounts for vFiler units	82
Virus protection for CIFS	82
Virus scanner registration	83
Virus scanning on vFiler units	83
Effect of virus scanner availability on CIFS access	83
Configuring virus scanning for a vFiler unit	83
Disaster recovery using MultiStore	85
Checking and preparing the storage system	85
Storage checklist	88
Checking the network	88
Network checklist	92
Secure communication for setting up disaster recovery	93
Creating a disaster recovery vFiler unit	93
Deleting the disaster recovery vFiler unit	95
The vfiler dr configure command	96
Activating the disaster recovery vFiler unit	97

What activating the disaster recovery vFiler unit does	98
Reactivation of the original vFiler unit	
Resynchronizing the vFiler unit	
Handling resynchronization failures	101
Reactivating the original vFiler unit by using SnapMirror commands	102
Reactivating the original vFiler unit by using vfiler dr commands	105
Re-creating the vFiler unit on a replacement storage system	107
Data migration using MultiStore	109
Secure communication for setting up data migration	109
How migrating a vFiler unit affects clients	110
Offline migration of vFiler units	110
The vfiler migrate commands	111
Migrating a vFiler unit by copying data	111
Adjusting client and network configurations if migrating to a differe	nt
subnet	113
vFiler unit migration without copying data	114
Prerequisites for vFiler unit migration between the nodes of an activ	e/
active configuration	114
Enabling SnapMover vFiler unit migration	115
Guidelines for setting up volumes to support SnapMover vFiler unit	
migration	116
SnapMover vFiler unit migration	117
Migrating a vFiler unit using SnapMover	117
Disabling SnapMover vFiler unit migration	118
What Data Motion is	120
Online migration of vFiler units	120
Licenses required for online migration	121
Stages of a vFiler unit migration	121
How to perform online migration of vFiler units	122
Viewing the status of a vFiler unit migration	123
Commands not allowed during online migration	123
Target portal group management for online migration of vFiler units	126
Data migration implications for IP-based target portal group	
management	127
Enabling IP-based target portal group management	128
Creating IP-based target portal groups	130

Adding IP addresses to IP-based target portal groups	131
Removing IP addresses from IP-based target portal groups	131
Destroying IP-based target portal groups	132
Displaying IP-based target portal group information	132
Disk space management using quotas	135
Allowing or disallowing quotas for a volume	135
Quota specification management	136
Turning on or turning off quotas from a vFiler unit	137
When quota thresholds and soft quotas are exceeded	137
How you can resize quotas	138
How the quotas file works	138
Displaying the quota status	138
Displaying a quota report	139
Index	141

Copyright information

Copyright [©] 1994–2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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

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

NetApp; the NetApp logo; the Network Appliance logo; 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; 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.

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.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

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

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

NetApp, Inc. NetCache is certified RealSystem compatible.

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 document describes how to administer vFiler units with the MultiStore software available by license with Data ONTAP.

Next topics

Audience on page 13 Accessing Data ONTAP man pages on page 13 Terminology on page 14 Where to enter commands on page 15 Keyboard and formatting conventions on page 16 Special messages on page 17 How to send your comments on page 17

Audience

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

This document is for system administrators who are familiar with operating systems that run on clients (such as Windows XP, Mac OS X and Linux) of the storage systems.

You should be familiar with how to configure the storage system and how the NFS, CIFS, and HTTP protocols are used for file sharing or transfers. This document does not cover basic system or network administration topics, such as IP addressing, routing, and network topology.

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

Types of information	Man page section
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 storage system command line:

man command_or_file_name

- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.
- Use the *Commands: Manual Page Reference*, Volumes 1 and 2 (which can be downloaded or ordered through the NOW site).

Note: All Data ONTAP man pages are stored on the storage system in files whose names are prefixed with the string "na_" to distinguish them from client man pages. The prefixed names are used to distinguish storage system man pages from other man pages and sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or services.

Terminology

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

Storage terms

array LUN	Refers to 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)	Refers to a logical unit of storage identified by a number.
native disk	Refers to a disk that is sold as local storage for storage systems that run Data ONTAP software.
native disk shelf	Refers to a disk shelf that is sold as local storage for storage systems that run Data ONTAP software.

storage controller	Refers to 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, storage appliances, appliances, storage engines, heads, CPU modules,</i> or <i>controller modules.</i>
storage system	Refers to 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, appliances, storage appliances, V-Series systems</i> , or <i>systems</i> .
third-party storage	Refers to 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

active/active	In the Data ONTAP 7.2 and 7.3 release families, refers to a pair of storage
configuration	systems (sometimes called <i>nodes</i>) configured to serve data for each other if one
	of the two systems stops functioning. Also sometimes referred to as active/active
	pairs. In the Data ONTAP 7.1 release family and earlier releases, this
	functionality is referred to as a <i>cluster</i> .

clusterIn the Data ONTAP 7.1 release family and earlier releases, refers to a pair of
storage systems (sometimes called *nodes*) configured to serve data for each other
if one of the two systems 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. For information about accessing your system with FilerView, see the *Data ONTAP System Administration Guide*.
- 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 NOW site	Refers to NetApp On the Web at http://now.netapp.com/.
Enter, enter	 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</i> font	 Words or characters that require special attention. Placeholders for information that you must supply. For example, if the guide says to enter the arp -d hostname command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.

Convention	What it means
Monospaced font	 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.
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

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

What MultiStore is

MultiStore enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. MultiStore is optional software that is available by license with Data ONTAP.

Each storage system created as a result of the partitioning is called a *vFiler unit*. A vFiler unit, using the resources assigned, delivers file services to its clients the same way that a storage system does.

The storage system on which you create vFiler units is called the hosting storage system. The storage and network resources used by the vFiler units exist on the hosting storage system.

The storage resource assigned to a vFiler unit can be one or more qtrees or volumes. The network resource assigned can be one or more base IP addresses or IP aliases associated with network interfaces. You can add or remove resources at any time.

You can use IPv6 addresses as network resources that can be assigned to the vFiler units. To use IPv6 addresses on the vFiler units, you must enable the IPv6 protocol on the hosting storage system. You must not disable the IPv6 protocol on a hosting storage system that has vFiler units. If you do so, you see a warning message similar to the following: vfilers are configured with IPv6 addresses. This option cannot be disabled.

For more information about enabling or disabling IPv6 on the hosting storage system, see the *Data ONTAP Network Management Guide*.

Next topics

Benefits of using MultiStore on page 19 MultiStore for consolidating servers on page 20 MultiStore for service providers and enterprises on page 21 MultiStore for disaster recovery and data migration on page 22 The default vFiler unit on page 22 Number of vFiler units allowed on page 22 Data access from the hosting storage system on page 23 Hosting storage system tasks on page 23

Benefits of using MultiStore

You can use MultiStore features such as virtualization, consolidation and management of storage requirements, security, disaster recovery, and data migration for your provisioning needs.

Virtualization

MultiStore enables you to manage tasks such as storage administration, provisioning, and management.

- Consolidation and ease of management Application service providers can consolidate the storage requirements of their customers. You can reduce management costs while offering independent, domain-specific storage management.
- Security
 Security is one of the key concerns when storage is consolidated either within an organization or
 by an application service provider. Using vFiler units enables you to have different security
 domains within the same storage system.
- Delegation of management

Administrators of vFiler unit can manage all vFiler units that they are authorized to access. However, vFiler unit administrators have access rights different from those of storage system administrators.

Disaster recovery and data migration

MultiStore enables you to migrate or back up data from one storage system to another without extensive reconfiguration on the destination storage system.

Related concepts

MultiStore for disaster recovery and data migration on page 22

MultiStore for consolidating servers

If you manage multiple application or operating system servers, you can store all of the data on one storage system for easier administration. You can consolidate the servers by partitioning the storage system into vFiler units and then copying the data from the servers to the vFiler units.

If the vFiler units are for CIFS users, you can set up the vFiler units to use the same computer names as the servers. This enables CIFS clients to share resources without having to remap their drives or search for the new server in Network Neighborhood. If the vFiler units are for NFS users, the NFS clients might need to remount the file systems, or they can use the automounter to automatically mount the file systems from the new locations.

The following illustration shows how multiple servers can be consolidated and can share the same resources.



A storage system that does not have MultiStore enabled can participate in only one security domain. Therefore, if your environment requires that different groups of CIFS users be in different domains, you must use multiple storage systems. MultiStore enables you to install each vFiler unit in the appropriate domain while keeping all of the data on the same physical storage system.

Because you can set up NIS and DNS servers for individual vFiler units, after you consolidate the servers on one storage system, network clients of the vFiler units can continue to use the same NIS and DNS servers as before.

MultiStore for service providers and enterprises

Service providers, such as ISPs and SSPs, can partition the resources of a storage system to create many vFiler units for client companies. Similarly, the information technology (IT) department of an enterprise can create vFiler units for various organizations, or customers, within the enterprise.

The administrator for each customer can manage and view files only on the assigned vFiler unit, not on other vFiler units that reside on the same storage system. In addition, there is no data flow between vFiler units. A customer using a vFiler unit is assured that no sensitive information is exposed to other customers that store data on the same storage system.

For example, an SSP can create the following vFiler units on a storage system:

- A vFiler unit named vFilerA It uses the /vol/vol1 volume and the e0 interface on the storage system. It is leased to CompanyA.
- A vFiler unit named vFilerB It uses the /vol/vol2 volume and the e1 interface on the storage system.

It is leased to CompanyB.

Although both CompanyA and CompanyB store data on the same storage system, network traffic for each company is restricted to the specified interface. The administrator at CompanyA (that uses NFS to access data) cannot use the showmount command on a UNIX client to view directories on the storage system that are outside the /vol/vol1 volume. Similarly, the administrator at CompanyB (that uses CIFS to access data) cannot browse any shared directories that are outside the /vol/vol2 volume.

MultiStore for disaster recovery and data migration

MultiStore enables easier migration and mirroring of data because all the information about users, CIFS shares, NFS exports, LUNs, igroups, ACLs, and so on, is encapsulated in the vFiler unit. If a disaster occurs, you can activate this vFiler unit on the destination storage systems with minimum reconfiguration.

You do not have to edit the files' ACLs, local user group definitions, user mapping information, and so on, before users can access the data.

Note: The static routing information is not carried to the destination storage system.

NFS and iSCSI users experience minimum disruption in service when the vFiler unit on the destination storage system starts serving data instead of the vFiler unit on the source storage system.

The default vFiler unit

When you license MultiStore, Data ONTAP automatically creates a default vFiler unit on the hosting storage system that is named vfiler0. The vfiler0 unit owns all the resources of the storage system. When you create vFiler units and assign resources to them, the resources are assigned from vfiler0. Therefore, vfiler0 owns all resources that are not owned by nondefault vFiler units.

The default vFiler unit exists as long as the MultiStore license is enabled. On a storage system with the MultiStore license enabled, you cannot destroy vfiler0.

All information provided about the vFiler units is applicable to vfiler0, unless noted otherwise.

Number of vFiler units allowed

There are limits to the number of vFiler units allowed in a storage system that has the MultiStore license enabled. You can usually have a maximum of 65 vFiler units on a storage system. However, the maximum limit depends on the memory capacity of the hosting storage system.

You can create 64 vFiler units on a storage system. The 65th vFiler unit is vfiler0, which is created automatically when MultiStore is licensed on the storage system. The default vFiler unit exists as long as MultiStore is licensed.

In an active/active configuration, you can create up to 64 vFiler units on each node of the active/ active configuration, for a maximum of 130 vFiler units in the active/active configuration.

Note:

These limits can be exceeded only during a takeover scenario, when one storage system takes over the resources of a vFiler unit in another storage system.

You can create a maximum of 16 vFiler units in FAS2040 systems.

Related tasks

Viewing the current limit on the number of vFiler units on page 38

Related references

Maximum vFiler units allowed on page 37

Data access from the hosting storage system

As the hosting storage system administrator, you can access all the data contained in a vFiler unit by using the vfiler context or the vfiler run commands. However, after you assign a qtree or volume to a vFiler unit, you no longer have access to the data in that qtree or volume.

For example, if you create a vFiler unit with the /vol/vol1 volume, you can configure the /etc/ exports file to mount the /vol/vol1 volume. However, after you create the vFiler unit, an attempt to mount the /vol/vol1 volume from the hosting storage system results in the following error message:

```
".../vol/vol1 belongs to vFiler unit A, cannot mount from vfiler0."
```

Related concepts

vFiler unit storage management from the hosting storage system on page 34

Related tasks

Executing commands from the hosting storage system on page 47

Hosting storage system tasks

You can perform tasks related to managing the resources on the hosting storage system in the same way that you perform them on a storage system without a MultiStore license.

You can use either the command line or FilerView to perform the following tasks:

- Manage volumes, disks, and RAID groups
- Increase data availability through Snapshot management, SnapMirror management, and volume copy

24 | Data ONTAP 7.3 MultiStore Management Guide

• Back up and recover data

These tasks are covered in detail in the *Data ONTAP File Access and Protocols Management Guide*, the *Data ONTAP Storage Management Guide*, and the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

MultiStore management

You can manage MultiStore from the hosting storage system by using the command line or FilerView. You can perform tasks such as creating, starting or stopping, and destroying vFiler units. You can also manage resources and protocols and monitor the status of vFiler units.

Next topics

Enabling the MultiStore license on page 26 Disabling the MultiStore license on page 27 Prerequisites for creating vFiler units on page 27 Creating a vFiler unit on page 31 Manually setting up a vFiler unit on page 32 vFiler unit storage management from the hosting storage system on page 34 Default limits on the number of vFiler units on page 37 What the vfiler rename command does on page 39 *Renaming a vFiler unit* on page 40 Stopping a vFiler unit on page 40 Destroying a vFiler unit on page 41 Restoring a vFiler unit on page 42 Starting a vFiler unit on page 42 Protocols supported by a vFiler unit on page 43 Allowing a protocol on a vFiler unit on page 44 Effects of disallowing protocols on a vFiler unit on page 44 Disallowing a protocol on a vFiler unit on page 45 *Displaying the vFiler unit status* on page 45 Viewing commands that can be executed from a vFiler unit on page 46 Executing commands from a vFiler unit on page 46 Executing commands from the hosting storage system on page 47 Executing RSH commands for a vFiler unit on page 47 Executing SSH commands for a vFiler unit on page 48 List of RSH and SSH commands on page 49 Effects of storage system reboot on a vFiler unit on page 50 Volumes and gtrees on a vFiler unit on page 50 Backup of vFiler units on page 52 LUNs on a vFiler unit on page 54 Networking guidelines on page 57 SnapMirror on the hosting storage system on page 58 Deduplication support on vFiler units on page 59

How MultiStore works with FlexClone files and FlexClone LUNs on page 61 SnapVault on the hosting storage system on page 62 SNMP support on vFiler units on page 63 Monitoring performance and statistics on page 64

Enabling the MultiStore license

To begin using MultiStore, you must enable the MultiStore license key.

Step

1. Enter the following command:

license add license_key

Result

Enabling the MultiStore license has the following effects on the storage system:

- You can use the vfiler and ipspace commands.
- Data ONTAP starts logging the status of vFiler units and sends the information to technical support using the AutoSupport feature.
- The routed daemon is enabled, but only in vfiler0.
- The ip.match_any_ifaddr option is set to off.

Note: You can turn the ip.match_any_ifaddr option to on if this option was enabled before licensing MultiStore.

• The vFiler unit limit (the number of vFiler units you can create on this storage system, including vfiler0) is set to a default value between 3 and 11, depending on the memory capacity of the hosting storage system.

Related tasks

Viewing the current limit on the number of vFiler units on page 38

Related references

Maximum vFiler units allowed on page 37

Disabling the MultiStore license

If you are not using vFiler units, you can disable the MultiStore license.

Before you begin

You can disable the MultiStore license only when there are no vFiler units other than vfiler0 on the storage system. If there are other vFiler units on the storage system, you must destroy them before disabling the MultiStore license.

Step

1. Enter the following command:

license delete multistore

Result

Disabling the MultiStore license has the following effects:

- MultiStore becomes unavailable immediately.
- You can no longer use the vfiler and ipspace commands.

Related tasks

Destroying a vFiler unit on page 41

Prerequisites for creating vFiler units

Before you create vFiler units, you must ensure that you have created at least one unit of storage (qtree, traditional volume, or FlexVol volume). Also, the IP address for the vFiler unit must not be configured when you create the vFiler unit.

The storage unit that contains information about configuring the vFiler unit must be writable. It must not be a read-only file system, such as the destination volume or qtree in a SnapMirror relationship.

Next topics

Storage guidelines on page 28 Language guidelines on page 28 Quota guidelines on page 29 Active/active configuration guidelines on page 29 SAN guidelines on page 29 The vFiler commands on page 30 Ensuring that the network interface is ready on page 30

Storage guidelines

When you assign storage units to vFiler units, the first storage unit (qtree, a traditional volume or FlexVol volume) assigned to the vFiler unit should not be removed as long as the vFiler unit exists. You cannot assign aggregates to a vFiler unit.

The first storage unit assigned to a vFiler unit is called the primary storage unit. The primary storage unit contains information about configuring the vFiler unit. Although you can remove storage units from a vFiler unit at any time after the vFiler unit is created, the primary storage unit must remain for as long as the vFiler unit exists. The primary storage unit has the same security characteristics it had before it was transferred to the vFiler unit.

When you create a new vFiler unit, C\$ share-level permissions are restricted to administrators only, but file-level security is not modified. The vFiler unit administrator can set more restrictive file-level permissions.

If the qtree or volume to be used as the primary storage unit contains an /etc directory, the data in the directory is lost after you add the qtree or volume to a vFiler unit. Data in qtrees and volumes that are used as nonprimary storage units is preserved. A volume assigned to a vFiler unit must not be the storage system's root volume. However, you can assign qtrees in the root volume to a vFiler unit. A volume assigned to a vFiler unit can be a traditional volume or a FlexVol volume.

For information about traditional volumes, FlexVol volumes, and aggregates, see the *Data ONTAP Storage Management Guide*.

FlexCache volumes can be created on the default vFiler unit, vfiler0, but cannot be assigned or moved to any other vFiler unit.

A qtree is assigned to a vFiler unit, owned by a vFiler unit, or associated with a vFiler unit only if that qtree is added as a resource to a vFiler unit. If a volume containing a qtree is added as a resource to a vFiler unit, then the qtree implicitly becomes a resource of that vFiler unit. If the vFiler unit administrator needs to create qtrees on the vFiler unit, the administrator assigns volumes instead of qtrees to the vFiler unit when creating the vFiler unit. This is because qtrees can be created only at the root of a volume.

If you anticipate that you might have to move the disks that are used for the vFiler unit's storage from one storage system to another, you should assign volumes, instead of qtrees, to the vFiler unit.

When managing NFS exports, CIFS shares, quotas, and options, vFiler unit administrators need to enter the complete path names of the storage resources used by the vFiler units in commands and configuration files. Therefore, the storage system administrator should choose volume and qtree names appropriately so that the complete path names beginning with filer_name:/vol/vol_name can be shared with the vFiler unit administrators.

Language guidelines

Administrators of vFiler units need to edit the /etc/quotas and /etc/usermap.cfg files for their vFiler units. These files support Unicode and root volume UNIX encoding. To ensure that vFiler unit

administrators can edit these files without requiring Unicode-capable editors, you should create vFiler units on a storage system whose root volume language can be used for editing.

Quota guidelines

When you create a vFiler unit, the ownership of a volume or qtree is changed from the hosting storage system to the vFiler unit that is created. This change requires that quotas be turned off for the affected volume before you create the vFiler unit. You can turn on the quotas for the volume after the vFiler unit is created.

Active/active configuration guidelines

When you set up vFiler units in an active/active configuration, both nodes must have the MultiStore license enabled to take over the partner. IPspaces created in one node in an active/active configuration must also be created in the partner system for failover to function correctly.

- You can create up to 64 vFiler units on each member of an active/active configuration, depending on the memory capacity of the hosting storage systems.
- The vFiler units hosted by the storage systems of the active/active configuration are created and configured independently.

This means each storage system can host a different number of vFiler units, and the vFiler unit configurations on the storage systems can be different from each other.

• In takeover mode, the functioning storage system takes over all the vFiler units created on the failed storage system.

These vFiler units include the vFiler units you have created and the default unit vfiler0. Therefore, for vFiler units on the failed storage system to work correctly after takeover, each network interface used by a vFiler unit in an active/active configuration must have a partner interface.

Related tasks

Viewing the current limit on the number of vFiler units on page 38

SAN guidelines

When you create vFiler units in a SAN environment, note that FC LUNs and igroups are not supported on vFiler unit. However, iSCSI LUNs and igroups are supported on all vFiler units managed separately for each vFiler unit.

Keep in mind the following:

- FC LUNs and igroups are supported only on the hosting storage system.
- When you create a vFiler unit on a storage system on which iSCSI is licensed, the iSCSI service is automatically started on the vFiler unit.

Note: If iSCSI is licensed on the hosting storage system and the storage to be allocated to the vFiler unit contains LUNs, you should unmap the LUNs.

For more information about LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FC.*

- Starting a vFiler unit starts iSCSI packet processing for that vFiler unit.
- Stopping a vFiler unit stops iSCSI packet processing for that vFiler unit.

The vFiler commands

The vfiler commands, which are supported only on the hosting storage system, enable the hosting storage system administrator to set up vFiler units and manage vFiler unit resources and Data ONTAP features on individual vFiler units. Each vfiler command has a different syntax.

The general vfiler command syntax is as follows:

vfiler command vfilertemplate options...

Some vfiler commands support the vfilertemplate option. vfilertemplate can be any of the following.

- A vFiler unit name
- A comma-separated list of vFiler unit names
- An IPspace declaration
- An asterisk (*) used as a wildcard

For example, you can run the setup command for the vFiler unit vfiler1:

vfiler run vfiler1 setup

If you use the asterisk, the command takes effect on all vFiler units, including vfiler0 (the hosting storage system), unless the command cannot be applied to vfiler0. See the na_vfiler(1) man page for more information.

Some vfiler commands include a complete path name for the qtree or volume that is assigned to the specified vFiler unit.

Related concepts

LUNs on a vFiler unit on page 54 File system access using NFS and CIFS on page 79 Disk space management using quotas on page 135 Virus protection for CIFS on page 82

Ensuring that the network interface is ready

Before you create a vFiler unit, you must ensure that the network interface is ready.

Steps

1. If the IP address for the vFiler unit is a base IP address for an interface, enter the following command to change the state of the interface for the IP address to down:

ifconfig interface down

Example

The following command changes the state of the e0 interface to down.

ifconfig e0 down

2. If the IP address for the vFiler unit is an IP alias for an interface, enter the following command to remove the alias:

ifconfig interface -alias address

Example

The following command removes the IP alias from the e0 interface:

ifconfig e0 -alias 123.123.123.123

3. If the IP alias is currently assigned to an interface, enter the following command to remove the alias:

ifconfig interface -alias address

If the IP alias is currently unassigned, the network interface is ready.

4. If the base IP address for the vFiler unit is assigned to an interface in the up state, enter the following command to change the state of the interface to down:

ifconfig interface down

If the base IP address for the vFiler unit is assigned to an interface in the down state, the network interface is ready.

Creating a vFiler unit

You can create a vFiler unit when you want to partition the storage and network resources of a single storage system. You can create a vFiler unit by using the vfiler create command from the command-line interface. You can also use FilerView to create a vFiler unit.

Steps

1. From the CLI, enter the following command:

```
vfiler create vfiler_name [-s ipspace] -i ip_address [ -i
ip_address ] ... path [ path ] ...
```

vfiler_name is the name of the vFiler unit.

ipspace is the IPspace the vFiler unit must belong to.

ip_address is an IP address of the vFiler unit.

path is the complete path name to an existing volume or qtree.

The first path name is the storage unit that contains the /etc directory. The /etc directory contains the configuration information about the vFiler unit.

Example

The following command creates a vFiler unit with two IP addresses, one volume, and one qtree:

vfiler create vfiler1 -i 123.123.123.123 -i 123.123.123.124 /vol/vol1 / vol/vol2/qtree2

For more information about the vfiler create command, see the na_vfiler(1) man page.

2. Respond to the prompts to set up the storage system, and to set up CIFS if necessary.

Result

The setup process does the following on the new vFiler unit:

- Starts NFS (if NFS is licensed on the hosting storage system) and configures the vFiler unit's primary storage (root volume) to be exported to the vFiler unit's administration host (using an entry in the /etc/exports file)
- Configures the vFiler unit's IP addresses and adds the appropriate entries to /etc/rc
- Creates a "pseudo-root" that allows CIFS clients to see all the storage that has been assigned to the vFiler unit as a single tree
- Starts the iSCSI service (if iSCSI is licensed on the hosting storage system)

As the vFiler unit administrator, you can now mount the root volume of the vFiler unit, edit /etc/ exports to suit your needs, and rerun the exports command.

Related concepts

File system access using NFS and CIFS on page 79

Related tasks

Manually setting up a vFiler unit on page 32

Manually setting up a vFiler unit

If you want to change the network configuration of your vFiler unit, you can manually set up the vFiler unit.

About this task

If you use the -n option of the vfiler create command, no automatic setup is performed, and no protocol servers are run on the vFiler unit until you set it up manually.

The procedure for setting up a vFiler unit manually is similar to that for setting up a storage system.

• Starting with Data ONTAP 7.3.1, the IPv6 address options are supported for the setup command:

The general setup command syntax, when used in the vfiler context, is as follows:

setup [-e ifname:ipv4address|[ipv6 address]:netmask|prefixlen ...] [-d
DNS domain name:DNS server ipv4 address|[DNS server ipv6 address] ...]

```
[-n NIS domain name:NIS server ipv4 address|[NIS server ipv6
address] ...] [-a ipv4 address|[ipv6 address]| name:ipv4 address|[ipv6
address ] [-p root password]
```

Starting with Data ONTAP 7.3.1, the IPv6 address options mentioned above are available.

Note: Ensure that you enter the IPv6 address in [] brackets.

You can use any of these options with the setup command:

• The -e option creates the bindings of the vFiler unit and provides netmask with IPv4 address and prefixlen with IPv6 address.

Netmask is of the form a.b.c.d, where a, b, c, and d should be between 0 and 255. Prefixlen is an integer between 0 and 127.

- The -d option specifies a DNS domain name and the IP addresses of one or more DNS servers.
- The -n option specifies a NIS domain name and the IP addresses of one or more NIS servers.
- The -a option specifies the administrator host name and IP address.
- The -p option sets the password of the vFiler unit's root user.

Note: The setup command does not prompt you for the time zone. All vFiler units are in the same time zone as the hosting storage system.

If the vFiler unit is licensed to deliver CIFS service, you must run cifs setup, as you would for a storage system, in addition to running the setup command.

Steps

1. To run the setup command from the default vFiler unit, enter the following command:

vfiler run vfiler_name setup

The setup command displays prompts for you to configure the vFiler unit. After you respond to all the prompts, configuration files, such as the /etc/exports file, are created in the /etc directory for the vFiler unit.

Note: Unlike the setup command for the storage system, the setup command for a vFiler unit does not cause NFS to start running.

If the vFiler unit runs the CIFS protocol, go to Step 2. Otherwise, the setup is complete.

2. Enter the following command:

vfiler run vfiler_name cifs setup

The cifs setup command displays prompts for you to configure CIFS on the vFiler unit. After you respond to all the prompts, CIFS starts running.

Related concepts

File system access using NFS and CIFS on page 79

vFiler unit storage management from the hosting storage system

As the physical storage system administrator, if you need to manage storage resources that belong to a vFiler unit but you do not have administrative access to the vFiler unit, you can temporarily move the vFiler unit's resources, or temporarily destroy the vFiler unit.

Note: Before taking either of the following actions, you should unmap any LUNs that have been created in the affected storage resources. For instructions, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

- Temporarily move the resources to the hosting storage system. However, you cannot move the vFiler unit's primary /etc volume.
- Temporarily destroy the vFiler unit. This returns ownership of all resources to the hosting storage system. No user data is modified when you destroy a vFiler unit.

Depending on which action you take, you can later move the storage resources back to the vFiler unit, or restore the vFiler unit.

Next topics

Effects of adding, removing, and moving vFiler unit resources on page 34 *Adding resources to a vFiler unit* on page 35 *Requirements for moving and removing resources* on page 35 *Removing resources from a vFiler unit* on page 35 *Moving resources between vFiler units* on page 36

Effects of adding, removing, and moving vFiler unit resources

Adding, removing, or moving vFiler unit resources affects only the association between the vFiler unit and the resources. It does not have any effect on user data in the vFiler unit.

- After you add storage resources to a vFiler unit, the resources are moved from the hosting storage system to a vFiler unit.
- After you remove storage resources from a vFiler unit, the resources are removed from the vFiler unit to the hosting storage system.
- After you add an IP address to a vFiler unit, you can assign the address as an IP alias of an interface or assign the address to a network interface that has not been configured.
- After you remove an IP address from a vFiler unit, the IP address becomes an unassigned IP address.
- After you move resources from one vFiler unit to another, the resources are moved from the resource vFiler unit and added to the destination vFiler unit.

Adding resources to a vFiler unit

To partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network, you must add resources to a vFiler unit.

Step

1. Enter the following command:

```
vfiler add vfiler_name [ -f ] [ -I ip_address [ -I ip_address ] ...]
[ path [ path ... ] ]
```

You can use the -f option to skip warning messages.

Example

The following command adds an IP address and a volume to an existing vFiler unit:

vfiler add vfiler1 -I 123.123.123.125 /vol/vol3

Example

The following example adds the IPv4 address 69.2.2.2 and the IPv6 address 3000::B9:23 in vfiler2.

```
vfiler add vfiler2 -i 3000::B9:23 -i 69.2.2.2
```

Requirements for moving and removing resources

When you move or remove vFiler unit resources, both the source and destination vFiler units must be in the same IPspace. If the resource that is being moved or removed is a storage unit, then the storage unit must not contain the vFiler unit's /etc directory.

- If a storage unit is to be moved or removed, and it contains any CIFS shares, home directories, or open files and directories, you must remove the CIFS shares, remove the home directories from the list of home directories, or close open files and directories.
- If the IP address is an IP alias, the alias must be removed. If the IP address is not an IP alias, the network interface associated with the address must not be configured.

Removing resources from a vFiler unit

You can remove resources you have added to the vFiler unit. For example, you might have to remove resources when an SSP wants to reduce the amount of storage used, to reduce operating costs.

Before you begin

If a storage unit you want to remove contains LUNs, you must first unmap the LUNs. See the *Data* ONTAP Block Access Management Guide for iSCSI and FC.

You must ensure that there are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vfiler context.

Step

1. Enter the following command:

```
vfiler remove vfiler_name [ -f ] [ -i ip_address [ -i ip_address ] ...]
[ path [ path ... ] ]
```

You can use the -f option to skip warning messages.

Example

The following command removes an IP address and a volume from an existing vFiler unit:

vfiler remove vfiler1 -i 123.123.123.125 /vol/vol3

Moving resources between vFiler units

You can move storage resources from one vFiler unit to another. For example, if the data on a qtree or volume changes ownership, that data can be moved to another vFiler unit that is owned by another company or business unit.

Before you begin

If a storage unit you want to move from one vFiler unit to another contains LUNs, you must first unmap the LUNs. See the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

You must ensure that there are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vfiler context.

About this task

- After you move a storage unit from one vFiler unit to another, the security information associated with the files in the storage unit is retained. As a result, users might be unable to access files properly.
- If you reassign a volume from one vFiler unit to another, Data ONTAP turns off quotas for the volume. After the volume is moved, you can turn quotas on again for the volume from the destination vFiler unit.
- If you reassign a qtree from one vFiler unit to another, Data ONTAP turns off quotas for the volume containing the qtree on both the source vFiler unit and the destination vFiler unit. After the qtree is moved, you can turn on the quotas again for the volume.
- When resources are being moved, all network connections to the vFiler units are terminated.

Step

1. Enter the following command:

```
vfiler move source_vfiler destination_vfiler [ -f ] [ -i ip_address [ -i
ip_address ] ...] [ path [ path ... ] ]
```

You can use the -f option to skip warning messages.
Example

The following command moves an IP address and a volume from one vFiler unit to another:

```
vfiler move vfiler1 vfiler2 -i 123.123.123.125 /vol/vol3
```

Default limits on the number of vFiler units

By default, there are limits on the number of vFiler units in storage systems that have the MultiStore license enabled.

The limits are as follows:

- Three vFiler units for storage systems with less than 1 GB (1,024 MB) of memory
- Five vFiler units for storage systems with 1 GB through 3 GB of memory
- Eleven vFiler units for storage systems with more than 3 GB of memory

These limits include vfiler0. Therefore, a limit of 11 vFiler units means that you can create a maximum of 10 vFiler units on each node of an active/active configuration.

Next topics

Maximum vFiler units allowed on page 37 Viewing the current limit on the number of vFiler units on page 38 Increasing the vFiler unit limit on page 38 Decreasing the vFiler unit limit on page 39

Maximum vFiler units allowed

There are limits on the number of vFiler units that you can create, depending on the available memory of the storage system.

The following table lists the number of vFiler units that can be created, depending on the available memory.

Memory	Number of vFiler units
Less than 1 GB	11
1 GB or more	26
2 GB or more	65

You can use the sysconfig -v command to verify the memory size of your system.

Viewing the current limit on the number of vFiler units

To determine whether you want to increase or decrease the current limit on the number of vFiler units that you can have on a hosting storage system, you can view the current limit.

About this task

The vFiler unit limit specifies the maximum number of vFiler units that can exist on the hosting storage system. Because the limit includes the hosting storage system, vfiler0 (which always exists if the MultiStore license is enabled), the number of vFiler units you can create on a storage system is one less than the vFiler unit limit set on a storage system.

In an active/active configuration, the same limit applies to each of the nodes of the active/active configuration. During a takeover, the hosting storage system can take over only the number of vFiler units that were set using the vfiler limit command.

Step

1. Enter the following command:

vfiler limit

Increasing the vFiler unit limit

If you need more partitions on your hosting storage system, you can increase the vFiler unit limit. The maximum number of vFiler units you can have on a storage system depends on the memory capacity of the hosting storage system.

Step

1. Enter the following command:

vfiler limit number

Example

To increase the number of vFiler units that you can create to 15, enter the following command on the hosting storage system:

vfiler limit 16

The limit is set to a number that must be one more than the number you create, because one vFiler unit is created automatically when you enable MultiStore.

Result

In an active/active configuration, this sets a limit of 16 vFiler units on each of the nodes of the active/ active configuration.

After you finish

For the change to take effect, you must reboot the storage system (or each storage system in an active/active configuration).

Decreasing the vFiler unit limit

You can decrease the vFiler unit limit when you need fewer partitions on your hosting storage system.

About this task

When you decrease the limit, the change is effective immediately and does not require a reboot of the storage system.

Step

1. Enter the following command:

vfiler limit

Example

To reduce the number of vFiler units that you can create from 15 to 10, enter the following command on the hosting storage system:

vfiler limit 11

The limit is set to a number that must be one more than the number you create, because one vFiler unit is created automatically when you enable MultiStore.

What the vfiler rename command does

The vfiler rename command renames the vFiler unit. You can rename the vFiler unit when you want the vFiler unit to have a unique name. The command changes the name of the vFiler unit only within Data ONTAP.

The vfiler rename command does not re-broadcast the new name to the CIFS domain controllers or the NetBIOS name servers because these protocols might be using a different name for the vFiler unit from the name that Data ONTAP uses. To change the name mapping in the CIFS domain controllers, you should run CIFS setup for each of these protocols.

Renaming a vFiler unit

You can use the vfiler rename command to rename the vFiler unit, for example, when you want the vFiler unit to have a unique name.

Before you begin

You should not rename a vFiler unit while it is being migrated. If you rename a vFiler unit that is being migrated, the migrate command on the remote system fails.

About this task

The new name for the vFiler unit should not exist on the storage system or on the partner storage system in an active/active configuration.

Although Data ONTAP allows the storage system and its partner to have vFiler units with identical names, it is easier to administer the storage systems if each vFiler unit has a unique name.

Step

1. Enter the following command to rename the vFiler unit:

```
vfiler rename old_vfiler_name new_vfiler_name
```

Example

The following command renames the vFiler unit vfiler1 as vfiler2:

vfiler rename vfiler1 vfiler2

Stopping a vFiler unit

You can stop a vFiler unit if you need to troubleshoot or destroy a vFiler unit.

About this task

After you stop a vFiler unit, the vFiler unit can no longer receive packets from clients.

Note: You cannot stop vfiler0.

The stopped state is not persistent across reboots. When you reboot the storage system, the vFiler unit that was stopped before the reboot operation resumes automatically.

Step

1. Enter the following command to stop the vFiler unit:

vfiler stop vfilertemplate

Example

Assume the storage system supports two vFiler units: vfiler1 and vfiler2. The following command stops all vFiler units, except vfiler0:

vfiler stop *

The following message appears after you enter the command:

```
vfiler stop *
vfiler1 stopped
vfiler2 stopped
```

Related concepts

The vFiler commands on page 30

Destroying a vFiler unit

If you want to return storage resources back to the hosting storage system (and the storage administrator's domain), you should destroy the vFiler unit that owns the storage resources.

Before you begin

Before you destroy the vFiler unit, you must unmap any LUNs that are mapped to the vFiler unit's storage. For information about how to unmap any LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

You should also stop the vFiler unit by using the vfiler stop command.

You must ensure that there are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vfiler context.

Step

1. To destroy the vFiler unit, enter the following command:

vfiler destroy [- f] vfiler_name

The -f (force) option in the vfiler destroy command destroys the vFiler unit immediately. Without the -f option, the command displays a confirmation prompt.

Result

Destroying a vFiler unit has the following effects:

- Clients using LUNs experience an interruption in service.
- All resources associated with the vFiler unit are released to the hosting storage system.
- There is no loss of data as the data can be accessed from the hosting storage system.

- All the vFiler unit's IP addresses are not configured and the corresponding entries in the storage system's /etc/rc file are removed.
- If the vFiler unit that is destroyed was not in the same IPspace as the hosting storage system, the IP addresses previously owned by the vFiler unit are not available for use after you destroy the vFiler unit.
- The effects on the quotas specified in the destroyed vFiler unit are the same as the effects on the quotas when you move resources from a vFiler unit to the hosting storage system.

Related tasks

Moving resources between vFiler units on page 36

Restoring a vFiler unit

If you accidentally destroy a vFiler unit, or if you want to reallocate storage resources from a storage system to a vFiler unit, then you can restore any vFiler unit that you have destroyed.

Step

1. Enter the following command to restore the destroyed vFiler unit:

```
vfiler create oldname -r oldfirstpath
```

Example

The following command restores vfiler2 that was destroyed:

vfiler create vfiler2

oldname is the name of the original vFiler unit.

oldfirstpath is the first path that was specified in the original vfiler create command that created the vFiler unit.

For more information about the vfiler create command, see the na_vfiler(1) man page.

Related tasks

Manually setting up a vFiler unit on page 32

Starting a vFiler unit

You can start a vFiler unit that is in the stopped state by using the vfiler start command.

About this task

After a vFiler unit starts, it is in a running state and can receive packets from clients. For example, if iSCSI is licensed on the storage system, starting or stopping a vFiler unit starts or stops iSCSI packet processing for that vFiler unit.

Note: You cannot start vfiler0.

Step

1. Enter the following command to start a vFiler unit:

```
vfiler start vfilertemplate
```

Example

Assume the storage system supports two vFiler units: vfiler1 and vfiler2. The following command starts all vFiler units, except vfiler0:

vfiler start *

The following message appears after you enter the vfiler start command:

```
vfiler start *
The default vfiler cannot be stopped or started.
Vfiler vfiler1 is already running.
Vfiler vfiler2 is already running.
vfiler0 running
vfiler1 running
vfiler2 running
```

Related concepts

The vFiler commands on page 30

Protocols supported by a vFiler unit

All protocols that are supported by the hosting storage system are supported by vFiler units. However, you can select the protocols that you want to allow on each of the vFiler units by using the vfiler allow command.

The following protocols are supported by vFiler units:

- CIFS
- NFS
- RSH
- SSH
- iSCSI
- FTP
- HTTP

The maximum number of FTP connections to a vFiler unit is determined by the ftpd.max_connections option set on the hosting storage system. The value set for this option is shared among the vFiler units on a storage system.

Allowing a protocol on a vFiler unit

You can select the protocols that you want to allow on the vFiler units by using the vfiler allow command.

Before you begin

To allow CIFS, NFS, FTP, or HTTP on a vFiler unit, each of the protocols must have an active license on the hosting storage system.

Step

1. Enter the following command to allow a protocol on a vFiler unit:

```
vfiler allow vfilertemplate proto=protocol ...
```

protocol can be any one of the following supported protocols: nfs, cifs, iscsi, rsh, ssh, ftp, or http.

Example

The following command allows the NFS and RSH protocols on the vFiler unit named vfiler1:

vfiler allow vfiler1 proto=nfs proto=rsh

Related concepts

The vFiler commands on page 30

Effects of disallowing protocols on a vFiler unit

When you disallow protocols on a vFiler unit, it has certain effects, depending on the protocol that is disabled. For example, if you disable HTTP or FTP, no new FTP or HTTP connections are allowed on the vFiler unit.

Effects of disallowing iSCSI: After iSCSI is disallowed on a vFiler unit, the following conditions apply on that vFiler unit:

- You cannot start iSCSI.
- No new iSCSI sessions are allowed.
- iSCSI commands on existing sessions are rejected.

Effects of disallowing FTP: After FTP is disallowed on a vFiler unit, no new FTP connections are allowed on the vFiler unit. However, transfers that started before FTP was disallowed are completed.

Effects of disallowing HTTP: After HTTP is disallowed on a vFiler unit, no new HTTP connections are allowed on the vFiler unit. Each new request receives a 503 HTTP server is disabled message. Any existing connections remain active.

Disallowing RSH and SSH: Although the Data ONTAP rsh.enable and ssh.enable option values (On or Off) determine whether the RSH or SSH server is enabled or disabled on a storage

system, disallowing RSH or SSH on a vFiler unit is independent of the value for that option. A vFiler unit can be configured to restrict RSH or SSH even when the corresponding enable option is set to On.

Note: To allow RSH on a vFiler unit, you must have the rsh.enable option set to On. To allow SSH on a vFiler unit, you must have the ssh.enable option set to On.

Disallowing a protocol on a vFiler unit

You can disallow protocols on a vFiler unit by using the vfiler disallow command. For example, for security reasons you might want to disallow access to the vFiler unit by SSH or RSH.

About this task

If the CIFS, NFS, iSCSI, FTP, or HTTP protocol is running and you disallow it, the protocol continues to run until the storage system reboots. However, packets destined for the vFiler unit are ignored.

Step

1. Enter the following command to disallow a protocol for a vFiler unit:

```
vfiler disallow vfilertemplate proto=protocol ...
```

protocol can be any one of the following supported protocols: nfs, cifs, iscsi, rsh, ssh, ftp, or http.

Example

The following command disallows the NFS and RSH protocols on the vFiler unit named vfiler1:

vfiler disallow vfiler1 proto=nfs proto=rsh

Related concepts

The vFiler commands on page 30

Displaying the vFiler unit status

You can use the vfiler status command to check the status of a vFiler unit. For example, you might want to check if a vFiler unit is in the stopped or running state.

Step

1. Enter the following command:

vfiler status

Result

The vfiler status command shows the following information about the vFiler unit:

- The state of the vFiler unit, whether stopped or running.
- Whether the IP addresses that have been assigned are configured, and the interfaces that they are bound to.
- Which protocols are allowed or disallowed.

See the na_vfiler(1) man page for more information.

Viewing commands that can be executed from a vFiler unit

You cannot execute all Data ONTAP commands from a vFiler unit. You can view the commands you can execute from a vFiler unit by running the vfiler run vfilertemplate ? command.

About this task

You need to be aware of the special considerations and restrictions that apply to some of the commands that can be executed from a vFiler unit. For more information about the special considerations and restrictions, see the vFiler considerations section of each command's man pages.

Step

1. Enter the following command to view commands that can be executed from a vFiler unit:

```
vfiler run vfilertemplate ?
```

Executing commands from a vFiler unit

To work with the data on a vFiler unit, you can execute commands directly from the vFiler unit. These commands are run in the same way that they are run from the vFiler unit's console. These commands can be run only on storage resources owned by the vFiler unit.

Steps

1. Enter the following command to switch to a particular vfiler context:

```
vfiler context vfiler_name
```

Example

vfiler context vfiler1

You are now in the context of vfiler1.

2. Enter the command you want to run from that particular vfiler context.

Example

The following command shows all the commands that are available from the context of the vFiler unit:

?

3. To return to the context of the hosting storage system, enter the following command:

```
vfiler context vfiler0
```

Executing commands from the hosting storage system

The storage system administrator can execute commands for any vFiler unit from the hosting storage system. However, a vFiler unit administrator cannot execute any commands from the hosting storage system.

Step

1. Enter the following command:

vfiler run vfilertemplate command

Example

vfiler run vfiler1 setup

You can now run the setup command for vfiler1 from the vfiler0 context.

Related tasks

Executing commands from a vFiler unit on page 46

Executing RSH commands for a vFiler unit

You can execute commands for a vFiler unit by using RSH after switching to a vfiler context by using the vfiler context command. You can also execute commands from a vFiler unit by using the vfiler run command. The RSH protocol is supported on vFiler units when they are configured to use IPv6 addresses.

Before you begin

The RSH protocol must be allowed for the vFiler unit. By default, RSH is allowed.

To enable the RSH protocol for the vFiler unit, the rsh.enable option for the vFiler unit must be set to on.

You must enter the command on a client of the vFiler unit that is permitted to have RSH access to the vFiler unit. The client must be one of the hosts specified by the rsh.access option for the vFiler unit.

You cannot launch RSH as an interactive shell or issue a vFiler command that requires user interaction through RSH.

Step

1. Enter the following command to execute commands for a vFiler unit by using RSH:

```
rsh vfiler_IP_address command
```

Example

The following command displays all options on the vFiler unit with the IP address 123.123.123.1:

```
rsh 123.123.123.1 options
```

Executing SSH commands for a vFiler unit

You can execute commands for a vFiler unit by using SSH after switching to a vfiler context by using the vfiler context command. You can also execute commands from a vFiler unit by using the vfiler run command.

Before you begin

The SSH protocol must be allowed for the vFiler unit. By default, SSH is allowed. Also, you must ensure that you enable the SSH protocol by setting the ssh.enable option to on.

You must enter the command on a client of the vFiler unit that is permitted to have SSH access to the vFiler unit. The client must be one of the hosts specified by the ssh.access option for the vFiler unit.

You cannot launch SSH as an interactive shell or issue a vFiler command that requires user interaction through SSH.

Step

1. Enter the following command to execute commands for a vFiler unit by using SSH:

ssh vfiler_IP_address command

Example

The following command displays all options on the vFiler unit with the IP address 123.123.123.1:

ssh 123.123.123.1 options

List of RSH and SSH commands

All the commands that can be executed on a vFiler unit are not available when you use RSH or SSH.

The following table lists the RSH and SSH commands that you can execute on a vFiler unit.

?	nfsstat	
cifs	nis	
clone	options	
config	passwd	
df	qtree	
dns	quota	
echo	route	
exportfs	sectrace	
filestats	secureadmin	
fpolicy	sftp	
fsecurity	sis	
help	snap	
hostname	snapmirror	
igroup	snapvault	
ipsec	useradmin	
iscsi	vol	
keymgr	vscan	
lock	wcc	
lun	ypcat	
nbtstat	ypgroup	
ndmpcopy	ypmatch	
ndmpd	ypwhich	
nfs		

Note: The hostname command is only for displaying, not changing, the name of the vFiler unit.

Effects of storage system reboot on a vFiler unit

Rebooting a storage system can affect the state of a protocol or a vFiler unit located in the storage system. If you allow or disallow a protocol on a vFiler unit, the allowed or disallowed protocol persists across reboots.

If you want to allow a protocol again after a reboot, you must reenable the protocol. For example, if you disallow NFS for a vFiler unit and then reboot the storage system, NFS remains disallowed for the vFiler unit after the reboot.

When you stop a vFiler unit and then reboot the storage system, the stopped vFiler unit starts running again after the reboot.

Volumes and qtrees on a vFiler unit

A volume assigned to a vFiler unit can be a traditional volume or a FlexVol volume. You cannot assign an aggregate to a vFiler unit. You can create qtrees in a volume only if the vFiler unit owns that volume.

For information about traditional volumes, FlexVol volumes, and aggregates, see the *Data ONTAP Storage Management Guide*.

Example: Creating a qtree in the default vFiler unit

Assume the /vol/vol0 volume is owned by vfiler0. You can use the following command to create a qtree in the /vol/vol0 volume:

qtree create /vol/vol0/qtree1

Example: Creating a qtree in a nondefault vFiler unit

Assume the /vol/vol1 volume is owned by vfiler1. The administrator for vfiler1 can use the following command to create a qtree in the /vol/vol1 volume:

rsh vfiler1 qtree create /vol/vol1/qtree2

Next topics

Effects of taking a vFiler unit volume offline on page 51 *Changes required after volumes are renamed* on page 51 *Who can change qtree security styles and oplock settings* on page 51 *Differences in qtree command output* on page 51 *Viewing all qtrees and the owner vFiler units* on page 52

Effects of taking a vFiler unit volume offline

Taking a volume that is used for vFiler unit storage offline affects protocols used by the vFiler unit and LUNs in the volume.

The effects are as follows:

- All CIFS shares and NFS exports in the volume are deactivated.
- If the volume contains the /etc directory for a vFiler unit, the vFiler unit stops running. The vFiler unit starts running again only after you bring the volume back online.
- All LUNs become unavailable.

Changes required after volumes are renamed

After you rename a volume that is used for vFiler unit storage, you should change the path names used in the vFiler unit's /etc/exports file accordingly. You should also verify that CIFS shares and quotas are configured properly.

Who can change qtree security styles and oplock settings

You can change the security style and oplock setting for a qtree or volume only if you are the owner of that qtree or volume.

- If a vFiler unit owns a volume, you can change the security styles or oplock settings for the volume and all qtrees on the volume from the vFiler unit.
- If a vFiler unit owns qtrees on a volume owned by the hosting storage system, you can change the security styles or oplock settings from the vFiler unit only for the qtrees the vFiler unit owns.
- If the hosting storage system owns a volume that contains qtrees assigned to vFiler units, you can change the security styles or oplock settings from the hosting storage system only for the qtrees the hosting storage system owns.

Differences in qtree command output

The gtree command output changes, depending on whether you enter the command from the vFiler unit or the hosting storage system.

- If you enter the <code>qtree</code> command from the hosting storage system, the command displays information about all qtrees on the storage system, irrespective of whether the qtrees are owned by the hosting storage system or vFiler units.
- If you enter the qtree command from a vFiler unit, the command displays information about qtrees on that vFiler unit only.
- If you enter the qtree command without arguments from a vFiler unit, a qtree that is the destination qtree for SnapMirror is shown as read_only in the Status column.

Viewing all qtrees and the owner vFiler units

To view a list of qtrees grouped by the vFiler units that own the qtrees, you must run the vfiler run command from the hosting storage system.

Step

1. Enter the following command to view all qtrees and the owner vFiler units:

```
vfiler run * qtree status
```

Backup of vFiler units

You can back up vFiler unit data from the hosting storage system if you want to back up all vFiler units at the same time. If you want to back up individual vFiler unit's data separately, you should back up from a vFiler unit's client.

Keep in mind the following points when you plan vFiler unit backups:

• From the hosting storage system, you can back up storage units owned by vFiler units—for example, by using the dump command.

You can back up all vFiler units at the same time. This method does not separate the data by vFiler unit; therefore, it is not suitable if each vFiler unit's data must be backed up separately.

- From a client of a vFiler unit, you can back up only that vFiler unit's data, but not any other vFiler unit's data.
 - A CIFS client can mount "/" from a vFiler unit and see a virtual tree comprising all of that vFiler unit's storage units.
 - A CIFS client can back up the entire data, including both CIFS and NFS data of a vFiler unit
 - An NFS client cannot see a virtual tree for the vFiler unit.
 - An NFS client can back up all of the vFiler unit's NFS data, but not its CIFS data.

If you want to back up an individual vFiler unit's data separately, a good way is to back up from a client (particularly a CIFS client). This backup method does not allow you to back up all vFiler units at the same time.

Next topics

NDMP support on page 53 Available NDMP options on page 53 Support for the ndmpcopy command on page 53 NDMP command support on page 53 NDMP password support on page 54

NDMP support

NDMP supports vFiler units. Because each vFiler unit has its own NDMP server, NDMP enables you to back up or restore each vFiler unit independently, and you can set NDMP options on each vFiler unit.

NDMP support for a vFiler unit is identical to NDMP support for a storage system, except in the following areas:

- Local tape backup and restore commands are not supported in individual vFiler units. Commands that access physical tape drive resources must be executed in the default vFiler (vfiler0) context.
- NDMP SAN management commands are not supported in the individual vFiler unit context. These commands must be executed in the default vFiler (vfiler0) context.
- VERITAS NDMP management commands are not supported in the individual vFiler unit context. These commands must be executed on the storage system.
- There is a hard limit of 160 concurrent NDMP sessions per storage system. Therefore, an NDMP server running on a vFiler unit might return an All sessions used up message even when there are no active sessions running on the vFiler unit.

Available NDMP options

All the NDMP options from the options command are available on the default vFiler unit (vfiler0). They are ndmpd.access, ndmpd.authtype, ndmpd.connectlog.enabled, ndmpd.enable, ndmpd.ignore_ctime.enabled, ndmpd.offset_map.enable, ndmpd.password.length and ndmpd.preferred_interface.

These NDMP options are also available on nondefault vFiler units, except the ndmpd.preferred_interface option.

Support for the ndmpcopy command

The ndmpcopy command uses NDMP over external IP interfaces. Therefore, you must first ensure that you have network connectivity, name resolution, and NDMP services configured properly at the source and destination locations before attempting to use the ndmpcopy command.

You can use the ndmpcopy command to copy data from one vFiler unit to another vFiler unit, or between different locations on the same vFiler unit.

NDMP command support

You can use the ndmpd command to manage and monitor the NDMP service for individual vFiler units. When you enable the NDMP service on a vFiler unit, the service is enabled only for that vFiler unit and not for all vFiler units.

Also, when you use the ndmpd command to monitor NDMP services and sessions in an individual nondefault vFiler unit context, it displays information only about the vFiler unit you are currently monitoring.

NDMP password support

When you use the NDMP commands on the storage system, you should use the storage system's root user's password in the ndmpcopy command. For enhanced security, the NDMP root user for individual nondefault vFiler units has a separate user name and password.

To view a nondefault vFiler unit's root user or nonroot password on any vFiler unit, you should use the ndmpd password command with that user name. This command lists the NDMP user password required by the ndmpdcopy command.

LUNs on a vFiler unit

LUNs are portions in a storage system that, when exported, look and act like local disks on the importing host. Data on a LUN can be managed at the block level (for example, by a database manager) as well as at the file level. A LUN is the basic unit of storage in a SAN.

Next topics

iSCSI LUNs and igroups on a vFiler unit on page 54 *The iSCSI service on a vFiler unit* on page 55 *LUN and igroup limitations on vFiler units* on page 56

iSCSI LUNs and igroups on a vFiler unit

Data ONTAP allows you to create and manage a separate set of iSCSI LUNs and igroups on each vFiler unit.

From the point of view of a host importing LUNs, a vFiler unit looks and acts just like a storage system. Administrators of those hosts do not need to be aware that the LUN resides on a storage unit owned by a vFiler unit. However, as the vFiler unit or hosting storage system administrator, you need to be aware of certain considerations.

You must remember the following points when you manage LUNs on a storage system on which a MultiStore license is enabled:

- You must create and manage LUNs from the vFiler unit that owns the storage containing the LUNs.
- A vFiler unit is aware only of those LUNs in the storage unit that the vFiler unit owns. When executed on a vFiler unit, the lun show command displays only that vFiler unit's LUNs.
- Ownership of LUNs changes with the ownership of the storage unit that contains the LUNs.
- LUNs must be unmapped before you can move the storage unit containing the LUNs. Therefore, you must unmap all affected LUNs before performing any of the following tasks:
 - Assigning storage that contains LUNs to a vFiler unit, either when you create the vFiler unit or later
 - Destroying a vFiler unit that owns storage containing LUNs

• Moving storage that contains LUNs from one vFiler unit to another, or between a vFiler unit and the hosting storage system

If you try to move a storage unit without unmapping the LUNs it contains, the operation fails. For instructions on unmapping LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Note: You do not have to unmap LUNs when you migrate a vFiler unit or replace it for disaster recovery purposes.

- igroups are owned by the vFiler unit on which they are created.
- Like LUNs, a vFiler unit is aware only of those igroups that it owns. When executed on a vFiler unit, the igroup show command displays only that vFiler unit's igroups.
- LUNs must be mapped to igroups owned by the vFiler unit that owns the LUNs.
- Each vFiler unit has its own namespace for LUNs and igroups:
 - igroups on different vFiler units can have the same initiator.
 - LUNs on different vFiler units can have the same LUN ID.
- When you migrate a vFiler unit or replicate it for disaster recovery purposes, LUNs owned by the vFiler unit are also migrated or replicated, along with their maps, igroups, and iSCSI configuration (the node names and the state of the iSCSI service). However, iSCSI authentication is not migrated or replicated.

Related concepts

Disaster recovery using MultiStore on page 85

Related tasks

Executing commands from a vFiler unit on page 46

The iSCSI service on a vFiler unit

In general, the iSCSI service operates on individual vFiler units, treating each of the vFiler units like a physical storage system. But the iSCSI software adapter (iSCSI software target) and the commands that manage and report on it, and the underlying NICs, operate on the hosting storage system.

As an iSCSI adapter on a storage system has only one identity (there are no vFiler unit-specific adapter names), there is only one set of iSCSI sessions and statistics.

For more information about the iSCSI service, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Keep in mind the following considerations about the iSCSI service on the hosting storage system:

- The iswt command, which you use to manage the iSCSI service on the storage system's NICs, operates on the hosting storage system, not on individual vFiler units.
- The iSCSI software target driver allows the storage system to be accessed as an iSCSI target device over the storage system's standard network interfaces.

- If iSCSI is licensed on the hosting storage system, the iSCSI service is available by default during the vFiler unit setup.
- The iscsi stats command displays statistics by hosting storage system and iSCSI adapter.

Keep in mind the following considerations about the iSCSI service on individual vFiler unit:

- The iSCSI protocol can be allowed or disallowed, and the iSCSI service can be started or stopped, for each vFiler unit.
- The iSCSI software adapter is online or offline for each vFiler unit, depending on whether the iSCSI service is running or stopped on that vFiler unit.
- Each vFiler unit has its own iSCSI node name, which includes the vFiler unit's UUID.
- Portal groups are defined for each vFiler unit.
- iSCSI subcommands operate specifically on each vFiler unit on which they are executed, except for the iscsi stats command.
- You should configure iSCSI security separately for each vFiler unit. This includes setting the default authentication mode: none, deny, or CHAP. For more information about CHAP, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.
- In the case of CHAP, there is a separate list of initiators and passwords for each vFiler unit.
- You can configure iSNS separately for each vFiler unit. You can use iscsi isns subcommands on each vFiler unit to do the following:
 - Configure which iSNS server to use
 - Turn on or turn off iSNS registration

When created, a vFiler unit's iSNS configuration is in the "not configured" state, regardless of its state on the hosting storage system.

For more information, see the na_iscsi(1) man page.

LUN and igroup limitations on vFiler units

FC LUNs are supported only on the hosting storage system, not on a vFiler unit. You can create only iSCSI igroups on vFiler units. You cannot create FC igroups on vFiler units.

Keep in mind the following limitations when you create LUNs:

- You can create FC igroups only on the hosting storage system.
- FC-connected hosts can access only those LUNs that are owned by the hosting storage system.
- The fcp command does not recognize vFiler units.

For detailed information about FC LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

Networking guidelines

To understand how vFiler units function, you must know how vFiler units operate with routing tables and gateways.

Next topics

The routed daemon on vFiler units on page 57 Command for changing the routing table in the default IPspace on page 57 The /etc/dgateways file on page 57 IPsec on a vFiler unit on page 57

The routed daemon on vFiler units

When you enable the MultiStore license, the routed daemon is enabled, but only on the default vFiler unit vfiler0.

When vFiler units are licensed and the routed daemon is on, the console displays the following message:

```
routed on Fri Nov 4 22:42:10 GMT
[ip.drd.vfiler.info:info]:Although vFiler units are licensed, the
routing daemon runs in the default IP space only.
```

Command for changing the routing table in the default IPspace

As vFiler units in the same IPspace share one routing table, you can change the routing table by entering the route command from the hosting storage system.

The route command has the following syntax:

```
route [-fn] add delete [host net] destination [gateway metric]
```

For more information about the route command and options, see the na_route(1) man page.

You can include the route command in the storage system /etc/rc file so that the routes are added automatically each time the storage system is rebooted.

Related concepts

What an IPspace is on page 67

The /etc/dgateways file

Only the hosting storage system contains the /etc/dgateways file. vFiler units do not maintain their own /etc/dgateways file.

IPsec on a vFiler unit

You can configure a vFiler unit as an IPsec end point. You can configure each vFiler unit separately with unique security configurations. When you migrate the vFiler unit from one hosting storage

system to another (or replicate it for disaster recovery purposes), the IPsec configuration is preserved as long as the vFiler unit's IP address does not change.

For more information, see the *Data ONTAP Network Management Guide* and the na_ipsec(1) man page.

SnapMirror on the hosting storage system

The SnapMirror feature for mirroring volumes and qtrees works with MultiStore after SnapMirror is licensed on the source and destination storage systems. You can enter SnapMirror commands from the default vFiler unit vfiler0 or nondefault vFiler units.

SnapMirror commands entered from the default vFiler unit can be used to make changes on and display information about all the nondefault vFiler units on the hosting storage system. snapmirror commands entered from a nondefault vFiler unit makes changes on or displays information only about that specific vFiler unit.

For backward compatibility, the default vFiler unit can operate snapmirror commands on all volumes and qtrees, even if they are owned by other vFiler units.

If vFiler unit storage volumes and qtrees are mirrored by vfiler0, the SnapMirror relationship is reflected only on vfiler0.

Next topics

Guidelines for using SnapMirror on page 58 *Determining the status of SnapMirror relationships* on page 59

Guidelines for using SnapMirror

You can use the snapmirror command on a nondefault vFiler unit in the same way as you do on a storage system. However, there are some exceptions.

The exceptions are as follows:

- Qtree SnapMirror is only supported for qtrees inside volumes owned by a vFiler unit.
- Qtree SnapMirror is only supported if a vFiler unit is rooted on a volume.
- Tape devices are not supported.
- SnapMirror sources and destinations cannot be qtrees in shared volumes.
- Synchronous SnapMirror is not supported.

Additionally, SnapMirror in a MultiStore context has the following features:

- The feature can be turned on and off independently on each vFiler unit.
- The snapmirror.access, snapmirror.checkip.enable, and snapmirror.enable options can be managed independently on each vFiler unit.
- Each vFiler unit has its own snapmirror.conf file in the /etc directory.
- A nondefault vFiler unit can only operate on the volumes and qtrees the vFiler unit owns.

- vFiler units do not require additional SnapMirror licenses and they use the same license as the storage system.
- SnapMirror relationships established between vFiler units are maintained during vFiler unit migration.
- SnapMirror destination updates are supported on both the hosting storage system and the vFiler unit.

Note: SnapMirror relationships between vFiler units and all the Snapshot copies in vFiler units are destroyed before a revert operation.

When specifying a path name in the /etc/snapmirror.conf file, ensure that you use the storage system name, and not the vFiler unit name. For more information, see the na_snapmirror.conf(5) man page.

Determining the status of SnapMirror relationships

On a vFiler unit, you can display active transfer entries related only to that vFiler unit. On the physical storage unit, you can display active transfer entries from all vFiler units. Inactive transfers are displayed only on the relevant vFiler unit.

Step

1. To display a comprehensive and readable list of SnapMirror transfers, enter the following command:

vfiler run * snapmirror status

This command iterates through all vFiler units and lists their transfers.

Deduplication support on vFiler units

Deduplication enables you to save space in vFiler units by reducing redundant data blocks within the vFiler units. You can run deduplication commands from the default and nondefault vfiler contexts.

Storage owned by a vFiler unit is not accessible and cannot be discovered by any other vFiler unit by using deduplication commands. Therefore, you cannot execute any deduplication command on a FlexVol volume that is not owned by the current vfiler context. Deduplication is used to deduplicate redundant data only on the FlexVol volumes owned by a vFiler unit (this is called FlexVol volume granularity). Currently, deduplication is supported at the FlexVol volume level. It is not supported at the qtree level, although a vFiler unit can own volumes and qtrees.

vfiler0 can execute deduplication commands on volumes owned by any vFiler unit in the storage system.

The maximum SIS session limit can be configured per vFiler unit.

Note: You can use the sis.max_vfiler_active_ops option to limit the number of active deduplication instances on a vFiler unit.

You can have a maximum of eight SIS sessions, which is also the default limit. The minimum SIS session limit is one. The hosting storage system allows a maximum of eight concurrent deduplication operations and they are shared among all hosted vFiler units.

During a vFiler unit migration, deduplicated volumes in the vFiler unit are also migrated. The FlexVol volumes on the destination vFiler unit inherit the deduplication attributes of the source vFiler unit.

Deduplicated volumes in a vFiler unit can be recovered during disaster recovery. All FlexVol volumes on the destination vFiler unit inherit the deduplication attributes of the source vFiler unit.

Note:

- Deduplication must be licensed on the primary storage system.
- Deduplication does not need to be licensed on the secondary storage system. However, if there is a situation in which the primary storage system is down and the secondary storage system becomes the new primary storage system, deduplication needs to be licensed for deduplication to continue. The best practice is to have deduplication licensed on both locations.

For more information about deduplication, see the *Data ONTAP Data Protection Online Backup and Recovery Guide.*

Running deduplication commands on a vFiler unit

You can run deduplication commands from interfaces such as CLI, RSH, or SSH. When using RSH or SSH, the user request is routed to the network stack, with the IP address and IPspace of the destination vFiler unit. The vfiler context is set for further processing of the command.

Steps

1. Enter the following command to switch to a particular vfiler context:

```
vfiler context vfiler_name
Example
```

vfiler context vfiler1

2. Run sis commands from the particular vfiler context:

sis on vol_name

Example

sis on /vol/vola

Note: All sis commands are supported.

For more information about sis commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

You can run deduplication commands on a nondefault vFiler unit by using the vfiler run command from the vfiler0 context. You should enter the following command: vfiler run [-q] vfilertemplate sis command.

The output of these commands is specific to the vfiler context, which means that it shows information about all the volumes in the current vFiler unit only.

The vfiler run command runs the command on the vFiler unit or vFiler units specified in *vfilertemplate*. If more than one vFiler unit is specified, the command runs separately for each vFiler unit.

How MultiStore works with FlexClone files and FlexClone LUNs

Starting with Data ONTAP 7.3.3, the FlexClone files and LUNs commands are available in the default and nondefault vfiler contexts. You can use FlexClone files and LUNs feature to create writable, space-efficient clones of parent files and parent LUNs within a vFiler unit. Storage owned by a vFiler unit cannot be accessed or discovered from other vFiler units by using the FlexClone file or LUN commands.

The following are considerations for creating FlexClone files and LUNs on vFiler units:

- Both MultiStore and FlexClone licenses must be enabled on the storage system.
- If you are a vFiler unit administrator, you can perform all FlexClone file and LUN operations only on vFiler units that you are authorized to manage.
- If you are a storage system administrator, you can perform FlexClone file operations on storage resources owned by all vFiler units from the default vfiler context.
- A storage system administrator cannot run FlexClone LUN operations from a default vfiler context on a LUN owned by a nondefault vFiler unit.
- If you are running a FlexClone LUN operation in the default vfiler context and if the volume or qtree on which the FlexClone LUN operation is running is moved to a nondefault vFiler unit, then the FlexClone LUN operation fails.
- A storage system administrator can see all clone operations running on different vFiler units from the default vfiler context.
- FlexClone file and LUN operations are visible only from the vfiler context on which the operations are being run and from the default vFiler context. You cannot view clone operations being run on other vFiler units.
- During reboot or takeover, if the storage is moved between vFiler units, the clone operation fails. However, this does not happen if the file clone operation was started from the default vfiler context.
- If a storage system reboots, then all the clone operations are restarted on the same vFiler unit after reboot.
- You can run a maximum of 500 FlexClone file and LUN operations on a storage system.

- All FlexClone file and LUNs commands are supported on vFiler units. The FlexClone file or LUN operations do not interrupt the offline vfiler migration and disaster recovery operations.
- The logs of all clone operations performed in vFiler units are stored at /vol/vol0/etc/log/ clone.
- You can run the following FlexClone file and LUN commands using RSH, SSH, or CLI in a vfiler context:

Note: You must switch to the vfiler context of the vFiler unit that owns the FlexVol volume or qtree.

- clone start
- clone status
- clone stop
- clone clear

Note: During online migration of a vFiler unit, you cannot use the clone start command on volumes that are owned by vFiler units.

For more information about FlexClone files and FlexClone LUNs, see *Data ONTAP Storage Management Guide*.

SnapVault on the hosting storage system

The SnapVault feature works with MultiStore and enables you to back up volumes and qtrees from the default vFiler unit or from nondefault vFiler units. You must license SnapVault on the source and destination storage systems before you can use the SnapVault feature.

Next topics

Where to enter SnapVault commands on page 62 Features and limitations of the snapvault command on page 63 Determining the status of SnapVault relationships on page 63

Where to enter SnapVault commands

You can enter SnapVault commands either from the default storage system (vfiler0) or from any nondefault vFiler unit.

Any command entered on the default vFiler unit makes changes on or displays information from all the vFiler units on the hosting storage system. Commands entered on a nondefault vFiler unit makes changes on or displays information only about that specific vFiler unit.

Features and limitations of the snapvault command

The snapvault command has some features and limitations when used in the MultiStore context.

The features of the snapvault command when used in a MultiStore context are as follows:

- Additional SnapVault licenses are not required. vFiler units use the same source and destination licenses as the physical storage systems.
- The SnapVault feature can be turned on and off independently on each vFiler unit.
- The snapvault.access and snapvault.enable options can be changed independently on each vFiler unit.
- Each vFiler unit has its own snapvault.conf file in the /etc directory.
- SnapVault relationships established between vFiler units are maintained across vFiler unit migration.

The limitations of the snapvault command when used in a MultiStore context are as follows:

- The SnapVault feature cannot be used in nondefault vFiler units that are rooted in a qtree.
- Nondefault vFiler units can operate only on the volumes and qtrees they individually own.

Determining the status of SnapVault relationships

On a vFiler unit, the status command displays active transfer entries related only to that vFiler unit. On the storage system, the status command displays active transfer entries from all vFiler units. Inactive transfers are displayed only on the relevant vFiler unit.

Step

1. To display a comprehensive and readable list of SnapVault transfers, enter the following command:

vfiler run * snapvault status

This command iterates through all vFiler units and lists their transfers.

SNMP support on vFiler units

SNMP is supported only on the hosting storage system and is not supported on individual vFiler units. You can enable SNMP on the hosting storage system to collect data about vFiler units.

vFiler unit data from MIBs

Data about vFiler units can be collected from the standard MIB and from the Data ONTAP custom MIB.

In the standard MIB, all vFiler unit data is global. It pertains to the sum of data from all vFiler units on the storage system, with the following exceptions:

- Statistics related to network interfaces are for the interfaces in the default IPspace.
- TCP statistics include data only from the connections and listen sockets in the default vFiler unit.
- UDP statistics include data only from sockets in the default vFiler unit.
- Quota information is gathered for each volume.

If the hosting storage system or a vFiler unit owns a volume with quotas, quota information is provided for the hosting storage system or the vFiler unit owning the volume. If a vFiler unit owns qtrees in a volume that it does not own, no quota information is provided for the vFiler unit.

In the Data ONTAP custom MIB, a group named vFiler is included. It provides information about each vFiler unit, such as the MultiStore license, IP address, protocols allowed, and so on.

Monitoring performance and statistics

You can view storage system statistics, NFS statistics, and CIFS statistics to determine how well your vFiler units are performing.

Next topics

Viewing storage system statistics on page 64 Viewing uptime statistics on page 65 Viewing NFS statistics on page 65 Viewing CIFS statistics on page 65

Viewing storage system statistics

You can view the storage system statistics only for the sum of statistics generated by all vFiler units, including vfiler0. You cannot view the statistics of a particular vFiler unit.

Step

1. Enter the following command to view the storage system statistics:

sysstat

Viewing uptime statistics

You can view the uptime statistics only for the storage system. You cannot view the uptime statistics for specific vFiler units.

Step

1. Enter the following command to view the uptime statistics:

uptime

Viewing NFS statistics

You can view the NFS statistics for the entire storage system or for specified vFiler units by using the nfsstat command.

Step

1. Enter the nfsstat command required for the statistics you want to view:

If you want to view the statistics of	Enter the following command from the hosting storage system	
All vFiler units together	nfsstat	
Specified vFiler units	vfiler run vfilertemplate nfsstat	
The hosting storage system	vfiler run vfiler0 nfsstat	

Viewing CIFS statistics

You can view the CIFS statistics for the entire storage system or for specified vFiler units by using the cifs stat command.

Step

1. Enter the cifs stat command required for the statistics you want to view:

If you want to view statistics of	Enter the following command from the hosting storage system	
All vFiler units together	cifs stat	
Specified vFiler units	vfiler run vfilertemplate cifs stat	
The hosting storage system	vfiler run vfiler0 cifs stat	

What an IPspace is

An IPspace defines a distinct IP address space in which vFiler units can participate. IP addresses defined for an IPspace are applicable only within that IPspace. A distinct routing table is maintained for each IPspace. No cross-IPspace traffic is routed.

Note: IPspaces support IPv6 addresses on its routing domains.

Each IPspace has a unique loopback interface assigned to it. The loopback traffic of each IPspace is completely isolated from the other IPspaces.

Next topics

Guidelines for vFiler unit participation in an IPspace on page 67 IPspace application scenario on page 68 Interface participation in an IPspace on page 69 Routing in an IPspace on page 70 Advantages of using VLAN tagging for IPspaces on page 70 Active/active configuration and IPspaces on page 71 Creating an IPspace on page 73 IPspace and the routed daemon on page 73 Listing IPspaces on a storage system on page 73 Removing an IP address from an interface on page 74 Assigning an interface to an IPspace on page 75 Destroying IPspaces on page 75 Creating a vFiler unit in a nondefault IPspace on page 76

Guidelines for vFiler unit participation in an IPspace

When you assign an IPspace to a vFiler unit, ensure that the vFiler unit has a unique IP address within that IPspace. After you assign an IPspace to a vFiler unit, you cannot change the IPspace without destroying the vFiler unit.

An IPspace can contain multiple vFiler units. However, a vFiler unit can belong only to one IPspace. A vFiler unit in one IPspace can have the same IP address as a vFiler unit in a different IPspace.

Each vFiler unit must have one IP address on the interface that leads to the default gateway of the assigned IPspace. This requirement ensures that the vFiler unit is reachable from within the IPspace.

IPspace application scenario

A typical application of an IPspace is when an SSP needs to connect customers of company A and company B to a storage system in the SSP's premises.

The SSP creates two vFiler units on the physical storage system—one per customer—and provides a dedicated network path from one vFiler unit to company A's network and one from the other vFiler unit to company B's network.

This deployment should work if both companies are using non-private IP address ranges. However, the following illustration shows both companies using the same private address ranges.



Both companies use the private IP address subnet 10.0.0.0, causing the following problems:

- The two vFiler units on the storage system at the SSP location have conflicting IP addresses if both companies decide to use the same IP address for their respective vFiler units.
- Even when the two companies agree on using different IP addresses for their vFiler units, problems arise: if any client in Company A's network has the same IP address as a client in Company B's network, packets destined for a client in A's address space might get routed to a client in B's address space, and vice versa.

• Assume the two companies decide to use mutually exclusive address spaces (for example, Company A uses 10.0.0.0 with a network mask of 255.128.0.0 and Company B uses 10.128.0.0 with a network mask of 255.128.0.0). The SSP needs to configure static routes on the storage system to route traffic appropriately to A's and B's networks. This solution is neither scalable (because of static routes) nor secure (broadcast traffic is sent to all interfaces of the storage system).

To overcome these problems, two IPspaces are defined on the storage system—one per vFiler unit. Because a distinct routing table is maintained for each IPspace and no cross-IPspace traffic is routed, the data for each company is securely routed to its respective network even if the two vFiler units are configured in the 10.0.0.0 address space, as shown in the following illustration:



Additionally, the IP addresses referred to by the various configuration files, such as the /etc/hosts file, the /etc/hosts.equiv file, and the /etc/rc file, are relative to that IPspace. Therefore, the IPspaces allow the SSP to configure the same IP address for the configuration and authentication data for both vFiler units, without conflict.

Interface participation in an IPspace

If MultiStore is licensed on the storage system, all its IP-addressable interfaces, including interfaces such as vifs and VLAN, belong to the default IPspace. The default IPspace exists automatically and cannot be destroyed.

When you create a new IPspace, you assign interfaces to the new IPspace from the default IPspace. An interface can belong only to one IPspace.

Routing in an IPspace

A distinct routing table is maintained for each IPspace. All vFiler units participating in an IPspace share its routing table. The IP address of the interface and the IPspace identifier are used to identify the vFiler unit for which the packet is intended.

All packets coming in through an interface are tagged with the IPspace identifier of the IPspace to which the interface belongs.

All outgoing traffic uses the IPspace identifier of the vFiler unit that is generating the traffic to determine the routing table to use. Data ONTAP ensures that packets generated by the vFiler units of an IPspace are transmitted through the interfaces that belong to that IPspace.

Note: Broadcast packets are restricted to the vFiler units within the destination IPspace.

Advantages of using VLAN tagging for IPspaces

You can use VLAN tagging for IPspaces to provide traffic separation for customers, to set up more IPspaces, and to securely deliver packets to a vFiler unit in an IPspace.

Traffic separation

VLAN tagging for IPspaces provides traffic separation from each customer to the storage system without incurring the cost of additional network devices, such as switches.

Without VLANs, you must provide physically separate network connections to ensure that the traffic from each customer is forwarded securely to and from the storage system. This solution is neither cost-effective nor scalable.

With VLAN tagging, you can set up distinct VLANs for each customer on a single switch. Thus, VLAN tagging provides an alternative to physically separate networks.

More IPspaces with VLAN tagging

Dedicating at least one physical interface per IPspace limits the number of IPspaces that can be set up on a storage system to the number of physical interfaces available on the storage system. VLAN tagging enables you to overcome this limitation.

VLAN tags can be used to forward traffic to appropriate IPspaces in cases where more than one IPspace shares the same physical interface.

Secure delivery of packets to a vFiler unit in an IPspace

VLANs inherently confine the broadcast domains. Therefore, only vFiler units belonging to a VLAN receive broadcasts intended for that VLAN, even if multiple vFiler units share a physical network interface.

Active/active configuration and IPspaces

You can use IPspaces in an active/active configuration. However, you should keep in mind the requirements for naming and assigning IPspaces.

Next topics

IPspace naming requirement on page 71 IPspace assignment requirement on page 71 Asymmetric active/active configuration setup on page 71 Specifying partners in an asymmetric active/active configuration setup on page 72

IPspace naming requirement

The names of IPspaces to which the partner interfaces are assigned must be the same on both storage systems.

For example, in an active/active configuration of storage system A and storage system B, if IPspaceA is created on storage system A, an IPspaceA must also exist on storage system B.

IPspace assignment requirement

The partner interfaces on both partners must be assigned to IPspaces with the same name on their respective storage systems.

For example, in an active/active configuration of storage system A and storage system B, interface e4 of storage system B is the takeover interface of interface e0 of storage system A, and vice versa. If interface e0 belongs to IPspaceA on storage system A, interface e4 must belong to IPspaceA on storage system B.

Asymmetric active/active configuration setup

In an asymmetric active/active configuration setup, the vFiler unit-IPspace configuration in one active/active configuration is different from that of its partner. For example, each partner might have a different number of vFiler units configured in a specific IPspace, or one partner might have no vFiler units.

A "standby" active/active configuration is an example of an asymmetric active/active configuration setup in which one of the hosts is connected to minimal storage of its own. This host takes over its partner's storage and vFiler units if the partner fails. In such a configuration, the standby host might not have enough storage to support the number of vFiler units that the primary host has.

Specifying partners in an asymmetric active/active configuration setup

You can use the interface name of the partner instead of the IP address to specify the partner interface when setting up an active/active configuration.

About this task

In this example, two storage systems, storage system1 and storage system2, are configured as shown in the following table.

Storage system	vFiler unit (associated IP address)	IPspace	Interface name
storage system1	vfiler0 (1.1.1.1)	default	e0
	vFiler1 (2.1.1.1)	ips1	e1
	vFiler2 (3.1.1.1)	ips2	e2
storage system2	vfiler0 (1.1.1.2)	default	e0
			e4a
			e4b

Steps

1. Specify partner interfaces on storage system1 by creating a host-partner relationship using the following commands:

ifconfig e0 1.1.1.1 netmask 255.255.255.0 partner e0 ifconfig e1 2.1.1.1 netmask 255.255.255.0 partner e4a ifconfig e2 3.1.1.1 netmask 255.255.255.0 partner e4b

- 2. Specify partner interfaces on storage system2 by completing the following steps:
 - **a.** Create two IPspaces: ips1 and ips2.
 - **b.** Assign interface e4a to ips1 and interface e4b to ips2.
 - c. Create a host-partner relationship on storage system2 by entering the following commands: ifconfig e0 1.1.1.2 netmask 255.255.0 partner e0 ifconfig e4a partner e1 ifconfig e4b partner e2
Creating an IPspace

IPspaces are distinct IP address spaces in which vFiler units reside. You create IPspaces when you need your vFiler units to have their own secure storage, administration, and routing.

About this task

You can have a maximum of 101 IPspaces per storage system. Of the 101 IPspaces, one is created by default when you install the MultiStore license on your storage system. You can create the remaining 100 IPspaces on the storage system.

You can use an alphanumeric string, 1 to 31 characters long, as the IPspace name.

All the IPspace names you create on a storage system must be unique. However, the IPspace names on active/active configuration partners must be the same.

Step

1. Enter the following command:

ipspace create ipspacename

ipspacename is the IPspace name that you want to create.

Example

To create IPspace1 on a storage system, enter the following command:

ipspace create ipspace1

IPspace and the routed daemon

When MultiStore is licensed, the routed daemon runs only in the default IPspace. The routed daemon is disabled when you create a new IPspace. It is enabled after the next reboot. Optionally, you can enable the routed daemon by using the route on command, without rebooting the storage system.

Listing IPspaces on a storage system

When you want to see which interfaces are assigned to each IPspace, you can use the ipspace list command.

Step

1. Enter the following command:

ipspace list
Example
If you enter the ipspace list command on a storage system that has three nondefault IPspaces,
you see the following output:
Number of ipspaces configured: 4 default-ipspace (e3) ipspace1 (e2d)
ipspace2 (e2c) ipspace3 (e10 e2b sf_vif)

Removing an IP address from an interface

You must remove an IP address from an interface before you can assign the IP address to an IPspace.

Steps

1. To remove an IP address from an interface, complete the following step:

If the interface is	Then
Currently configured with an IP address that belongs to a vFiler unit in a nondefault IPspace	Remove the IP address configured for the interface from the vFiler unit as you would when changing resources for a vFiler unit.
In the default IPspace	Go to Step 2.
Configured with an autoconfigured	Enter the following command to remove the IPv6 address:
IPv6 address	ifconfig interface inet6 -alias IPv6_address
	<i>interface</i> is the name of the interface.
	<i>IPv6_address</i> is the IPv6 address configured on the interface.
Configured with an IP alias	Enter the following command to remove the IP alias:
	ifconfig interface -alias address
	<i>interface</i> is the name of the interface.
	address is the IP address configured for the alias.
	Go to Step 2.
Not configured with an IP alias	Go to Step 2.

2. Enter the following command:

ifconfig interface 0.0.0.0

interface is the name of the interface from which you want to remove the IP address.

Related concepts

Effects of adding, removing, and moving vFiler unit resources on page 34

Assigning an interface to an IPspace

To assign an interface to an IPspace, ensure that the interface does not have a configured IP address. You can assign only one interface at a time. If IPv6 is enabled on your storage system and you bring up an interface, an IPv6 address is automatically configured on the interface.

Step

1. Enter the following command:

ipspace assign ipspacename interface_name

ipspacename is the IPspace name to which the interface is assigned.

interface_name is the name of the interface to be assigned.

Destroying IPspaces

If you no longer need an IPspace, you can destroy it.

Before you begin

Ensure there are no network interfaces or vFiler units associated with the IPspace you want to destroy.

Step

1. Enter the following command:

ipspace destroy ipspace_name

ipspace_name is the IPspace name that you want to destroy.

Related tasks

Listing IPspaces on a storage system on page 73 *Assigning an interface to an IPspace* on page 75 *Displaying the vFiler unit status* on page 45 *Destroying a vFiler unit* on page 41

Creating a vFiler unit in a nondefault IPspace

If you are a SSP, you might have different clients that require complete network isolation. Creating vFiler units in separate, or nondefault, IPspaces ensures that the data on these vFiler units is not routed to other vFiler units or the hosting storage system in other IPspaces.

Before you begin

The following conditions must be met before you can create a vFiler unit:

- You must create at least one unit of storage (qtrees, traditional or flexible volumes) before creating the vFiler unit.
- The storage unit that contains information about configuring the vFiler unit must be writable. It must not be a read-only file system, such as the destination volume or qtree in a SnapMirror relationship.
- The IP address used by the vFiler unit must not be configured when you create the vFiler unit.
- Create an IPspace.
- Verify that each interface to be used by the vFiler unit is ready to be configured.

About this task

When creating a vFiler unit in a nondefault IPspace, you need to meet the same prerequisites and follow the same guidelines as those for creating a vFiler unit in the default IPspace.

Steps

- 1. Assign an interface to be used by the vFiler unit in the newly created IPspace.
- 2. To create the vFiler unit, enter the following command:

```
vfiler create vfiler_name -n -s ipspace -i ip_address [ -i
ip_address ] ... path [ path ] ...
```

ipspace is the IPspace in which the vFiler unit's IP addresses reside.

ip_address is an IP address.

path is the complete path name to an existing volume or qtree.

The first path name is the storage unit that contains the /etc directory, which contains the configuration information about the vFiler unit.

Note: Ensure you use the -n option of the vfiler create command, and do not use the setup command to specify IP addresses for interfaces assigned to different IPspaces. The setup command (which runs automatically after the vfiler create command unless you use the -n option) does not allow duplicate IP addresses even if they are for interfaces in different IPspaces.

Example

The following command creates a vFiler unit named vFiler1 in the IPspace named ipspace1, using the IP address 123.123.123.123 and the /vol/vol1 volume as resources:

vfiler create vfiler1 -n -s ipspace1 -i 123.123.123.123 /vol/vol1

3. Perform the following step:

For each IP address used by the vFiler that is	Go to
An IP alias	Step 6.
A base IP address	Step 7.

4. Enter the following command to create the IP alias:

ifconfig interface alias address

If the interface to which you are assigning the alias is currently down, go to Step 7. Otherwise, go to Step 8.

- 5. Use the ifconfig command to configure the interface as Up with the IP address you specified in Step 4.
- 6. To modify the routing table for the vFiler unit, enter the following command:

vfiler run vfiler_name route add [host | net] [prefixlen prefixlen] destination gateway metric

Example

The following example adds a route to the routing table used by the vFiler unit named vFiler1:

vfiler run vfiler1 route add 1.2.3/24 1.2.3.1 1

- 7. For each interface used by the vFiler unit, add the following information to the /etc/rc file on the hosting storage system:
 - An ifconfig command to configure the interface as Up, with the address you specified in Step 4
 - The route command you entered in Step 8

This configures the interfaces as Up and enforces the route commands across reboots.

8. Optionally, if the hosting storage system is part of an active/active configuration, edit the /etc/ rc file in each partner of the active/active configuration to define a partner interface for each interface that the vFiler unit uses.

Example

The following command ensures that interface e10 has a partner e10 interface: ifconfig e10 partner e10

Alternatively, you can run the setup command from the storage system and enter the partner interface for each interface.

Attention: If you use the storage system setup command to automatically define a partner interface, it clears all existing information about vFiler unit configuration in /etc/rc file.

9. The IPspace that the vFiler unit is assigned to must have a default gateway. If the IPspace already has a default gateway, skip this step. To establish a route to the default gateway in the IPspace that the vFiler unit can use, enter the following command:

vfiler run vfiler_name route add default gateway metric

Example

The following command adds a route to the default gateway for the IPspace used by the vFiler unit named vFiler1:

vfiler run vfiler1 route add default 1.2.3.1. 1

Related concepts

Prerequisites for creating vFiler units on page 27

Related tasks

Creating an IPspace on page 73 *Assigning an interface to an IPspace* on page 75 *Ensuring that the network interface is ready* on page 30 *Creating an IPspace* on page 73

File system access using NFS and CIFS

To access a vFiler unit's file system using NFS and CIFS, you must prepare the vFiler unit by using the NFS and CIFS protocol, respectively.

For CIFS clients, the root of the primary storage qtree is the root ("/") of a vFiler unit's file system. If "/" is shared, a CIFS client mapping to it can browse all of the vFiler unit's storage in a single tree. This mechanism is called the vFiler unit's *pseudo-root*.

NFS clients must import discrete storage units as they are defined on the hosting storage systems. Pseudo-root directories are not available to NFS clients.

Next topics

Path name specification for NFS exports or CIFS shares on page 79 vFiler unit preparation for NFS on page 80 vFiler unit preparation for CIFS on page 81 Virus protection for CIFS on page 82

Path name specification for NFS exports or CIFS shares

When you specify a path to export to NFS clients or to share with CIFS clients, you should use the complete path name.

Example of a path for an NFS export

Assume a vFiler unit named vfiler1 uses the /vol/vol1 volume for storage. To export the home directory at the root of this volume to the clients of vfiler1, you should use /vol/vol1/ home in the /etc/exports file or in the exports command.

Example of a path for a CIFS share

Assume a vFiler unit named vfiler1 uses the hosting storage system's /vol/vol1 volume as its primary storage. To share the entire volume, and all other storage owned by the vFiler unit, in a single tree, you should specify / as the share path. To offer the home directory at the root of this volume as the home share, specify /home as the path name for the home share. The vFiler unit mechanism that makes this possible is known as *pseudo-root*.

vFiler unit preparation for NFS

To prepare the vFiler unit for NFS, you must start the NFS protocol and export discrete storage units to NFS clients in the same way they are defined on the hosting storage system.

However, you must do this only if you want to start NFS manually. If you used the default form of the vfiler create command, you are prompted for CIFS setup information as soon as you create the vFiler unit.

Next topics

Starting the NFS protocol on page 80 Exporting all file systems in /etc/exports on page 81

Related tasks

Executing commands from a vFiler unit on page 46

Starting the NFS protocol

To start the NFS protocol on the vFiler unit, you should use the nfs on command.

Step

- 1. Complete one of the following steps:
 - From the vFiler unit, enter the following command:

nfs on

• From the hosting storage system, enter the following command:

vfiler run vfiler_name nfs on

Note: When you use an IPv6 address, you must enter the following command:

vfiler run vfiler_name options nfs.ipv6.enable on

Result

The NFS protocol server starts running on the vFiler unit.

Exporting all file systems in /etc/exports

As pseudo-root directories are not available to NFS clients, to start using NFS on a vFiler unit, NFS clients must import discrete storage units in the same way they are defined on the hosting storage system.

Step

- 1. Complete one of the following steps:
 - From the vFiler unit, enter the following command:

```
exportfs -a
```

• From the hosting storage system, enter the following command:

```
vfiler run vfiler_name exportfs -a
```

• From a vFiler unit client that is allowed to connect to the vFiler unit through RSH, enter the following command:

rsh vfiler_name exportfs -a

vFiler unit preparation for CIFS

From the hosting storage system, you can use the vfiler run command to issue CIFS commands for vFiler units. You can use User Manager or Server Manager to manage user accounts and shares. The cifs command is not available on RSH.

Note: When you use an IPv6 address, you should enter the following command: vfiler run vfiler_name options cifs.ipv6.enable on.

Data ONTAP does not limit the number of users, shares, open files, and locked files on a per vFiler unit basis.

Next topics

Commands run from the hosting storage system on page 81 *Local user accounts for vFiler units* on page 82

Related tasks

Executing commands from a vFiler unit on page 46

Commands run from the hosting storage system

Server Manager does not perform all the functions of the cifs shares -add and cifs shares - change commands. You can execute the commands that Server Manager cannot run from the vFiler

unit's CLI (through the vfiler context command) or from the hosting storage system (through the vfiler run command).

The following commands cannot be run from Server Manager:

- cifs shares -add-forcegroup group_name
- cifs shares -add share_name pathname -nosymlink_strict_security
- cifs shares -add- widelink
- cifs shares -add -novscan
- cifs shares -add -novscanread
- cifs shares -change share_name {-forcegroup group_name | -noforcegroup }
- cifs shares -change share_name { -symlink_strict_security | nosymlink_strict_security }
- cifs shares -change share_name { -widelink | -nowidelink }
- cifs shares -change share_name { -vscan | -novscan }
- cifs shares -change share_name { -vscanread | -novscanread }

Local user accounts for vFiler units

From the hosting storage system, you can use the useradmin command to create local accounts for CIFS users of each vFiler unit. Each vFiler unit supports up to 96 local user accounts.

The maximum number of vFiler unit user accounts per storage system is 96 times the maximum number of vFiler units for that storage system.

For additional information about using local accounts for CIFS authentication, see the *Data ONTAP File Access and Protocols Management Guide*.

For detailed information about managing local user accounts by using the useradmin command, see the *Data ONTAP Storage Management Guide*.

Related tasks

Viewing the current limit on the number of vFiler units on page 38

Virus protection for CIFS

You can perform virus scanning for vFiler units that run the CIFS protocol. The hosting storage system administrator can configure virus scanning on files owned by the hosting storage system and files owned by vFiler units. Administrators of vFiler unit can configure virus scanning only for vFiler units they administer.

Next topics

Virus scanner registration on page 83 *Virus scanning on vFiler units* on page 83 *Effect of virus scanner availability on CIFS access* on page 83 *Configuring virus scanning for a vFiler unit* on page 83

Virus scanner registration

Virus scanners can be registered with the hosting storage system or any vFiler unit. A virus scanner that is registered with a vFiler unit always takes precedence over the virus scanner registered with the hosting storage system.

You can determine whether files on a vFiler unit are scanned by virus scanners registered with the hosting storage system or the vFiler unit. If a virus scanner is registered with a vFiler unit and is functional, the files on the vFiler unit are scanned by the scanner that is registered with the vFiler unit.

If the scanner that is registered with a vFiler unit becomes unavailable, then the hosting storage system's virus scanner scans the files on the vFiler unit on the following conditions:

- The vscan options use_host_scanners command is set to On
- A scanner is registered with the hosting storage system

When the scanner local to the vFiler unit becomes available, it takes over from the hosting storage system's scanner.

Virus scanning on vFiler units

To scan files on nondefault vFiler units, virus scanning must be enabled on these vFiler units. Files on the default vFiler unit can be scanned only by a virus scanner registered with the hosting storage system. Also, only hosting storage system administrators can run virus scanning on vfiler0.

Before you run virus scanning on nondefault vFiler units, ensure the following requirements are met:

- Virus scanning must be enabled for a nondefault vFiler unit.
- A virus scanner must be registered either with the vFiler unit or hosting storage system, and the vFiler unit must be allowed to use it.

Effect of virus scanner availability on CIFS access

Although virus scanning is enabled and the mandatory_scan option for the vscan command is set to On, CIFS clients of the vFiler unit are not allowed to open any files on the vFiler unit if no virus scanner is available.

Configuring virus scanning for a vFiler unit

You can configure virus scanning either to use or not use the virus scanner registered with the hosting storage system, if the virus scanner registered with the vFiler unit is unavailable.

Steps

1. Enter the following command to specify the virus scanner for scanning files on a vFiler unit:

vfiler run vfilertemplate vscan options use_host_scanners on | off

When you set the use_host_scanners option to On, the vFiler unit uses the virus scanner registered with the hosting storage system. When this option is set to On, the hosting storage system and its vFiler units share the virus scanner. However, the vFiler unit uses the virus scanner registered with the hosting storage system only when the virus scanner registered with the vFiler unit is unavailable.

You can set the use_host_scanners option to Off if you do not want to allow the vFiler unit to use the virus scanner registered with the hosting storage system.

Note: The use_host_scanners option is applicable only to a vFiler unit you created. You cannot set it on vfiler0 or a storage system.

2. Enter the following command to enable or disable virus scanning:

vfiler run vfilertemplate vscan on | off

Disaster recovery using MultiStore

MultiStore supports disaster recovery, and you can prepare for disaster recovery by creating a backup vFiler unit to prevent loss of data if a disaster occurs.

You can safeguard information by creating vFiler units on the destination storage system, which remain inactive unless a disaster occurs. You should perform checks to ensure that the storage system and network are ready for disaster recovery.

Next topics

Checking and preparing the storage system on page 85 Checking the network on page 88 Secure communication for setting up disaster recovery on page 93 Creating a disaster recovery vFiler unit on page 93 Deleting the disaster recovery vFiler unit on page 95 The vfiler dr configure command on page 96 Activating the disaster recovery vFiler unit on page 97 What activating the disaster recovery vFiler unit does on page 98 Reactivation of the original vFiler unit on page 98 Resynchronizing the vFiler unit on page 98 Handling resynchronization failures on page 101 Reactivating the original vFiler unit by using SnapMirror commands on page 102 Reactivating the original vFiler unit by using vfiler dr commands on page 105 Re-creating the vFiler unit on a replacement storage system on page 107

Related concepts

What an IPspace is on page 67 *Data migration using MultiStore* on page 109

Checking and preparing the storage system

You must ensure that the destination storage system can support the disaster recovery vFiler unit.

Steps

- 1. Verify that the destination storage system has enough storage space to hold the source vFiler unit's volumes.
 - **a.** On the source storage system, enter the vfiler status -r command to see the volumes that the vFiler unit is using.

b. Enter the df command on each of those volumes to check the disk space being used. The destination volumes must have at least the same amount of space that is being used on the source volumes. You can run the df command on the destination storage system to check this.

Note: If the source and destination storage systems use different-sized disks and have different block sizes, you should adjust the df numbers accordingly.

2. Perform the appropriate action based on what you learned about space in Step 1:

Note: The destination volumes must have the same amount of space as the source volumes.

If the destination volumes	Then
Have enough space	Go to the Step 3.
Do not have enough space	a. Install new disk shelves.b. Use the aggr add command to add new disks to the destination volumes.

3. Ensure that the destination storage system has the same volume structure as the source and the volumes to be used by the destination vFiler unit are not used by any other vFiler unit.

The volumes to be used by the destination vFiler unit must have the same path names as those used by the source vFiler unit.

If the destination storage system	Then
Has a volume whose path name matches the path name used by the source vFiler unit and the volumes to be used by the destination vFiler unit are not used by any other vFiler unit	Go to Step 4.
Has a volume whose path name matches the path name of the source vFiler unit, but the volume is used by another vFiler unit	 Perform one of the following steps: If the volume is the root volume of the vFiler unit, use the vfiler destroy command to destroy the vFiler unit. If the volume can be removed, use the vfiler remove command to disassociate the volume from that vFiler unit. If the volume cannot be destroyed or removed, use the vol rename command to rename the volume. Then, create a new volume with the old name of the volume you just renamed.

If the destination storage system	Then
Does not have any path name that matches the name used by the source vFiler unit or by volumes to be used by the destination	 Perform one of the following steps: For traditional and FlexVol volumes, use the vol create command on the destination storage system to create volumes whose names match those being used by the source vFiler unit. Use the vol rename to rename a volume.
	For more information, see the <i>Data ONTAP Storage Management Guide</i> .

4. Ensure that the destination volumes do not contain any qtree with the same name as that used by a qtree in the source volumes.

For disaster recovery, SnapMirror is used, and SnapMirror replicates qtree names from the source to the destination volume. Therefore, ensure that the qtree names on the source do not exist on the destination.

If there are	Then
Qtrees in the destination volumes that have	Rename the matching qtrees in the destination volumes.
source volumes	To rename a qtree, move it from a client the same way you move a directory or folder. For more information, see the <i>Data ONTAP Storage Management Guide</i> .

No matching qtree names in the destination Go to Step 5. volumes

5. Check whether quotas are being enforced from the hosting storage system.

To check where quotas are being enforced from, enter the following command from the hosting storage system:

quota report

Quotas enforced from the vFiler unit are copied to the new vFiler unit, but quotas enforced from the hosting storage system are not copied.

6. Depending on the result of Step 5, perform the following:

If quotas for qtrees used by the vFiler unit	Then
Are being enforced	Go to the Step 7.
Are not being enforced	You have completed checking the storage system.

- 7. Keep a record of the storage system's /vol/vol0/etc/quotas file for future reference.
- 8. Copy the relevant entries into the destination storage system's /vol/vol0/etc/quotas file.

Storage checklist

You can use the storage checklist to record storage system information and to ensure that your systems are ready to use for disaster recovery.

- How much disk space is used on the source storage system's volumes?
 df of source storage system's volumes:
- How much disk space is free on the destination storage system's volumes? df of destination storage system's volumes:
- Have you added enough disks to the destination volumes, if required?
- Do the path names of the source and destination volumes match?______
- If you are managing qtree-based vFiler units, do any destination volume qtree names match those on the source volume? ______
- Have you copied storage system-based quota information from the source to the destination storage system's /etc/quotas file?_____

Checking the network

Before setting up a disaster recovery vFiler unit, you must check whether the source and destination storage systems are on the same subnet, the IPspace used by the vFiler unit, and whether the destination vFiler unit can access the same NIS and DNS servers as the source vFiler unit.

Steps

1. Check whether the destination vFiler unit can take over the source vFiler unit's IP addresses by displaying information about all the network interfaces of the destination vFiler unit.

You must use the *ifconfig* -a command on the source vFiler unit and the destination storage system to display information about all the network interfaces.

You can reuse the source IP addresses and aliases on the destination vFiler unit if the destination vFiler unit is on the same subnet as the source vFiler unit.

2. Depending on the result of Step 1, perform the following step:

If the source and destination storage systems	Then
Are on the same subnet	Go to Step 3.
Note: This is the default for the vfiler migrate command.	

If the source and destination storage systems	Then
Are on different subnets	a. Obtain all new IP addresses that are in use on the source vFiler unit for the destination vFiler unit.
	Note: You might need to replicate subnet-separation arrangements that exist on the source vFiler unit. For example, the source vFiler unit might use one IP address for a service network and another for an administration network.
	b. Make a note of the new IP addresses on a worksheet. The vfiler dr configure command prompts you for these addresses when you create the disaster recovery vFiler unit. The vfiler migrate command also prompts you for these addresses, but you might then need to run the setup command on the destination vFiler unit.

3. Check whether the source vFiler unit is using the default IPspace.

To display information about IPspaces and the interfaces assigned to them, use the ipspace list command on the source vFiler unit.

If the ipspace list command reports	Then
Default-ipspace	Go to Step 5.
Something other than default-ipspace	a. Use the ipspace create command to create a corresponding IPspace with the same name on the destination storage system.
	 b. Use the ipspace assign command to assign physical interfaces to the IPspace. These interfaces should be attached to the same physical network.

4. Depending on the result of Step 3, perform the following step:

5. Check whether the destination vFiler unit has access to the same NIS servers as the source.

Note: You can skip this check if the source and destination vFiler units are on the same subnet.

To see the NIS servers that are available to the source vFiler unit, use the nis info command.

Note: The ypwhich command shows the server to which the storage system is currently bound.

6. Depending on the result of Step 5, perform the following step:

If	Then
The destination vFiler unit can use the same NIS servers as the source vFiler unit	Go to Step 7.
Note: This is the default for the vfiler migrate command.	
The destination vFiler unit cannot use the same NIS servers	a. Find the NIS servers that are available for the destination storage system.
	b. Make a note of the IP addresses of those servers on the network checklist.
	The vfiler dr configure command prompts you for these addresses when you create the disaster recovery vFiler unit.
	The vfiler migrate command does not prompt you for these addresses. If you move a vFiler unit to a different subnet, you might need to run the setup command on the destination vFiler unit.

7. Check whether the destination vFiler unit has access to the same DNS servers as the source.

Note: You can skip this check if the source and destination vFiler units are on the same subnet.

To see what DNS servers are available to the source vFiler unit, use the dns $\,$ info $\,$ on command.

8. Depending on the result of Step 7, perform the following step:

If	Then
The destination vFiler unit can use the same DNS servers as the source vFiler unit	Go to Step 9.
Note: This is the default for the vfiler migrate command.	

If	Then
The destination vFiler unit cannot use the same DNS servers	a. Find the DNS servers that are available for the destination storage system.
	b. Make a note of the IP addresses of those servers on a worksheet.
	The vfiler dr configure command prompts you for these addresses.
	The vfiler migrate command does not prompt you for these addresses. If you move a vFiler unit to a different subnet, you might need to run the setup command on the destination vFiler unit.

- **9.** Check whether the destination vFiler unit has access to the same WINS servers and the same Windows security network as the source.
- **10.** Depending on the result of Step 9, perform the following step:

If the destination vFiler unit	Then
Can use the same WINS servers and Windows security network as the source vFiler unit	Go to Step 11.
Cannot use the same WINS servers and Windows security network	a. Find the name and type (Windows NT 4 or Windows 2000) of the domain the destination vFiler unit is in.
	b. Note this information on the network checklist.
	When you activate the disaster recovery vFiler unit, you have to configure it into the new domain.
	If you move a vFiler unit into a different domain, you have to configure it into the new domain.

- **11.** Check whether the destination vFiler unit can use the same trusted host for vFiler unit administration as the source vFiler unit.
- **12.** Depending on the result of Step 11, perform the following step:

If the destination vFiler unit	Then	
Can use the same trusted host as the source vFiler unit	You have completed this task.	

If the destination vFiler unit	Then
Cannot use the same trusted host	 a. Find the name of the new trusted host. b. Note this information on the network checklist. You must configure the new trusted-host information after configuring the disaster recovery vFiler unit, or after moving the vFiler unit.

Related tasks

Creating a disaster recovery vFiler unit on page 93 Adjusting client and network configurations if migrating to a different subnet on page 113 Activating the disaster recovery vFiler unit on page 97

Related references

Storage checklist on page 88

Network checklist

You can use the network checklist to record network information and to ensure that your systems are ready to use for disaster recovery.

• Are there enough IP addresses available for the vFiler unit on the destination network?

 old interface:
 new interface:

 old interface:
 new interface:

 old interface:
 new interface:

 old interface:
 new interface:

Note: Check syntax carefully. Interface names are case-sensitive.

- · Have you created the number of nondefault IPspaces, if any are required?
- Have you gathered all the authority servers?
 - old NIS domain: ______ new NIS domain: _____
 - old NIS IP address: ______ new NIS IP address: _____
 - old NIS IP address :______ new NIS IP address:_____
 - old DNS domain:_____ new DNS domain:_____
 - old DNS IP address :______ new DNS IP address:_____
 - old DNS IP address : _______ new DNS IP address: ______
 - old DNS IP address :______ new DNS IP address:______
 - old WINS IP address : ______ new WINS IP address: ______
 - old WINS IP address :______ new WINS IP address:_____
 - old NT domain type: NT4 W2K

old domain name (FQDN and NetBIOS):_____

new NT domain type: NT4 W2K

new domain name (FQDN and NetBIOS):___

Related tasks

Checking the network on page 88

Secure communication for setting up disaster recovery

Disaster recovery implemented over SSL provides both authentication and confidentiality of user data.

RSH is an insecure command channel in which passwords and other confidential information are sent across the network in plain text. RSH is disabled in some deployments for security reasons. Therefore, the disaster recovery feature is not available in those deployments.

Note: When you use SSL in your setup, ensure the following:

- SSL should be enabled on the source vFiler unit.
- The httpd.enable and httpd.admin.enable options should be turned on at the source vFiler unit.

Creating a disaster recovery vFiler unit

You can create a disaster recovery vFiler unit on a destination storage system that has the storage capacity, characteristics, and the network connectivity to host an identical copy of the vFiler unit, if the original vFiler unit is not able to serve data.

Before you begin

Ensure the following:

- You have prepared the destination storage system.
- SnapMirror is licensed and enabled on both the source and the destination storage system.
- The source and destination storage systems can communicate with each other over the network (for example, by means of DNS lookup or entries in the /etc/hosts file).
- · The destination volumes are online.
- You know the source storage system's administrative ID and password.

About this task

Attention: On the disaster recovery storage system, protect any volumes that have the same names as the volumes on the original vFiler unit. Otherwise, data in those volumes will be lost.

Steps

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

```
vfiler dr configure source_vfiler@source_filer
```

Note:

- If you want to set up synchronous SnapMirror between the source and destination storage systems, you should use the -s option of the vfiler dr configure command.
 For more information about the -s option, see the na_vfiler(1) man page.
- If you want to set up multiple paths from the source to the destination storage system, you should use the -a option of the vfiler dr configure command.
 For more information, see the section on using SnapMirror over multiple paths in the *Data ONTAP Data Protection Online Backup and Recovery Guide*.
- If you want to use the SSL command channel for disaster recovery, you should use the -c option of the vfiler dr configure command.

Note: Ensure that when you are using IPv6 addresses, the httpd.ipv6.enable option is set to on.

For more information about the -c option, see the na_vfiler(1) man page.

- **2.** Respond to the login prompt with a valid administrative ID and password for the source storage system.
- 3. Respond to the IP address and binding prompts.
- 4. Respond to the NIS and DNS server prompts.
- 5. Optionally, monitor the progress of the disaster recovery by using the following command:

vfiler dr status source_vfiler@source_filer

When the vfiler dr status command output shows that all the storage units of the source vFiler unit are mirrored, the disaster recovery vFiler unit has been created. However, the disaster recovery vFiler unit has not been started.

Note: The vfiler dr configure command might take some time to complete, especially if a source qtree has many millions of inodes.

6. If you copied quota information to the destination storage system's /etc/quotas file, activate the quotas on that storage system. For activating quota on each volume, use the following command:

quota on volume_name

7. Edit the disaster recovery vFiler unit's /etc/hosts.equiv file by adding the name of the trusted host for administering the disaster recovery vFiler unit.

Note: If the trusted host is either a Windows or a UNIX system and the trusted user is not the root user, you need to add the user name as well. For example:

adminhost joe_smith

8. Add the path to the root volume and the name of the trusted host to the disaster recovery vFiler unit's /etc/exports file.

Example

/vol/vf1_root access=adminhost, root=adminhost

9. If the vFiler unit's storage units contain iSCSI LUNs, reconfigure iSCSI authentication.

For instructions, see the Data ONTAP Block Access Management Guide for iSCSI and FC.

Related concepts

Disaster recovery using MultiStore on page 85 *The vfiler dr configure command* on page 96

Related references

Storage checklist on page 88 *Network checklist* on page 92

Deleting the disaster recovery vFiler unit

You can delete the disaster recovery vFiler unit at any time after setting it up.

Step

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

```
vfiler dr delete source_vfiler@source_filer
```

Example

To remove a disaster recovery vFiler unit on the destination storage system "StorageSystem 2" created for a vFiler unit "vfiler1" on the source storage system "StorageSystem1", even if SnapMirror errors exists, enter the following command on "StorageSystem 2":

vfiler dr delete -f vfiler1@StorageSystem1

Result

Before removing the disaster recovery vFiler unit, the vfiler dr delete command removes all SnapMirror relationships, and any other configuration information related to the disaster recovery vFiler unit, from the source vFiler unit.

If any errors are detected in the SnapMirror relationships, the deletion of the vFiler unit is aborted. To ignore SnapMirror errors and remove the disaster recovery vFiler unit, you can use the -f option available in the vfiler dr delete command.

The vfiler dr configure command

The vfiler dr configure command uses the Data ONTAP SnapMirror feature as its underlying technology. Multiple paths can be set from the source to the destination storage systems in SnapMirror. The -a option of the vfiler dr configure command enables you to set multiple paths for the configuration operation.

The vFiler dr configure command does the following tasks:

- · Checks whether the destination storage system is capable of receiving the source data.
- Configures and runs SnapMirror to copy the data from the source to the destination vFiler unit.
 - iSCSI LUNs (including the LUN maps) are copied from the source vFiler unit to the destination vFiler unit.
 - igroups and the iSCSI configuration, including node names and the iSCSI service state, are copied to the destination vFiler unit.
 - iSCSI authentication is not copied to the destination vFiler unit.
- Saves the IP configuration and binding information you supplied when you created the disaster recovery vFiler unit.
- Saves the NIS and DNS server information you supply.
- Saves the quota information from the source vFiler unit's /etc/quotas file.
- Causes a baseline transfer to occur from the source to the destination.
- Sets the incremental update interval from the source to the destination to be once every three minutes.
 - If you want to change the default setting, you should edit the etc/snapmirror.conf file as described in the *Data ONTAP Data Protection Online Backup and Recovery Guide*.
 - The vfiler dr configure command automatically configures everything that SnapMirror requires for regular updates. No other SnapMirror configuration is necessary. If any SnapMirror configuration requirements are missing from your system (for example, a missing volume or license), the vfiler dr configure command returns errors.
- Overwrites all data on the volumes of the destination vFiler unit. You must protect any volumes on the destination storage system that have the same name as the volumes on the source vFiler unit. Otherwise, data in the volumes of the destination vFiler unit will be lost.
- Creates a special type of vFiler unit, known as a disaster recovery backup vFiler unit, on the destination storage system.

This vFiler unit is stopped and cannot be started except when a disaster occurs. Before activation, the vFiler unit responds only to the vfiler dr delete, vfiler dr status, and vfiler dr resync commands. You should not use ifconfig to configure its addresses.

Related tasks

Creating a disaster recovery vFiler unit on page 93 Activating the disaster recovery vFiler unit on page 97 Creating a disaster recovery vFiler unit on page 93

Activating the disaster recovery vFiler unit

After a disaster, you can keep serving data by switching to the disaster recovery vFiler unit while trying to recover the original vFiler unit that was damaged during the disaster.

Steps

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

vfiler dr activate source_vfiler@source_filer

2. Configure the DNS or NIS servers:

If you specified	Then
A different set of DNS or NIS servers	Copy the /etc/hosts.equiv.bak file to the /etc/ hosts.equiv file.
The same set of DNS or NIS servers	Go to Step 3.

When you activate the vFiler unit in the event of a disaster, the /etc/hosts.equiv file can be overwritten. If you specified a different set of DNS or NIS servers for the disaster recovery location when you created the disaster recovery vFiler unit, the existing /etc/hosts.equiv file is overwritten and the old file is copied to an /etc/hosts.equiv.bak file.

- 3. To change the name of the Windows domain controller, use the cifs prefdc command.
- 4. To change the Windows WINS server, use the cifs setup command.

Note: If the Windows domain has changed, you might have to change the permissions on the Windows data files to allow your users the same access they had in the old domain.

- 5. Make adjustments on the clients, such as remounting volumes and qtrees.
- **6.** Add static route entry if required, because static routing information is not carried to the destination storage system.

What activating the disaster recovery vFiler unit does

When you activate a disaster recovery vFiler unit, the original vFiler unit is replaced with the disaster recovery vFiler unit.

The vfiler dr activate command does the following:

- Breaks the SnapMirror relationships between the source and destination storage systems
- Activates the disaster recovery vFiler unit which then starts and responds to all commands that vFiler units support
- Brings LUNs online
- Configures IP bindings on the destination vFiler unit according to the information you provided in the vfiler dr configure command, adding the destination IP information to the destination /etc/rc file

Any IP information that pertains only to the source vFiler unit is removed from the destination / etc/rc file.

- Configures the NIS and DNS servers according to the information you provided to the vfiler dr configure command
- Configures any quota information saved by the vfiler dr configure command

Reactivation of the original vFiler unit

By resynchronizing the disaster recovery vFiler unit with the original vFiler unit, you can re-create the disaster recovery configuration. The original vFiler unit is now reactivated. You can do this, provided the system is not damaged and has suffered only a temporary failure. You might have to use a different storage system if the original storage system is severely damaged.

Resynchronizing the vFiler unit

You can resynchronize the original vFiler unit with the currently activated disaster recovery vFiler unit before reactivating the original vFiler unit. If you resynchronize, you do not have to delete the original vFiler unit and create a new vFiler unit. A baseline transfer is also not required.

Before you begin

You can resynchronize the vFiler unit if the following prerequisites are met:

- The storage elements for the vFiler unit are only volumes (traditional or FlexVol volumes) and not qtrees.
- The source and destination vFiler units contain identical volumes.
- The size of the volumes on the source and destination vFiler units is the same.

- The vFiler unit from which you are updating is activated.
- The original vFiler unit is not in the process of migration.
- If new storage elements have been added to a disaster recovery activated vFiler unit, the newly added storage elements exist on the original storage system as well.
- Ensure that the original vFiler unit you are resynchronizing is in a stopped state.

About this task

If you do not resynchronize the original vFiler unit, baseline transfer occurs between the new vFiler unit and the disaster recovery vFiler unit.

The vfiler dr resync command performs the following:

- The vfiler dr resync command resynchronizes all storage elements that belong to the disaster recovery vFiler unit, including the volumes that were added to the disaster recovery vFiler unit after it was activated.
- The vfiler dr resync command also sets the incremental update interval from the source to the destination to be once every three minutes. Three minutes is the default setting. If you want to change the default setting, you should edit the etc/snapmirror.conf file as described in the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Attention:

On the storage system on which you are resynchronizing the original vFiler unit, protect any volumes that have the same names as the volumes on the disaster recovery vFiler unit.

If a volume with the same name exists, the volume is automatically added and initialized for SnapMirror transfers from the disaster recovery vFiler unit. Any existing data on the newly added volume is lost.

If you want to delete old Snapshot copies after executing the vfiler dr resync command, you must first release them by using the snapmirror release command. If you fail to do this and Snapshot copies are deleted, then further SnapMirror updates will fail to synchronize.

See the na_snapmirror(1) man page for more information on snapmirror commands.

Steps

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

vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure][s] dr_vfilername@disaster_recovery_filer

authinfo is the authentication information specified in the username:password format, where username is the login name of the administration host on the disaster recovery storage system and password is the password for that user name. If you do not specify the authentication information in the vfiler dr resync command, you are prompted for it when you run the command. alt-remote is the alternate host name or IP address of the source (the disaster recovery storage system, in this case).

alt-local is the alternate host name or IP address of the destination (the original storage system, in this case).

-c enables you to use the secure command channel.

This option is used only in the vfiler dr commands that interact with the remote storage system.

-s enables you to set up a synchronous SnapMirror relationship between the source and destination storage systems.

dr-vfilername is the name of the disaster recovery vFiler unit that is currently in the activated state.

disaster_recovery_filer is the name of the storage system on which the currently activated disaster recovery vFiler unit exists.

Note: When you run the vfiler dr resync command on the disaster recovery vFiler unit to resynchronize it with the original vFiler unit, you do not have to specify the -1, -a, and -s options. The values stored when the vfiler dr configure command was run (for a baseline transfer) to back up the original vFiler unit are used.

2. After the resynchronization operation is complete, enter the following command on the storage system on which the original vFiler unit exists to check the status of the vFiler unit that was resynchronized:

vfiler status -r original_vfilername

The original vFiler unit is now in a stopped, DR backup state. This is because the vfiler dr resync command does not activate the vFiler unit on resynchronizing. The vFiler unit continues to behave like a backup disaster recovery vFiler unit until you use the vfiler dr activate command to reactivate it.

3. Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

If the storage elements for the vFiler unit are	Then
Volumes	On the disaster recovery storage system, enter the following command:
	vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] original_vfilername@original_filer

If the storage elements for the vFiler unit are	Then
Qtrees	a. On the disaster recovery storage system, destroy the disaster recovery vFiler unit by entering the following command on the disaster recovery hosting storage system:
	vfiler destroy vfilername vfilername is the name of the original vFiler unit.
	b. On the disaster recovery storage system, re-create the disaster recovery vFiler unit.

Related tasks

Activating the disaster recovery vFiler unit on page 97 Creating a disaster recovery vFiler unit on page 93

Handling resynchronization failures

You can take some corrective action if the resynchronization of the original vFiler unit with the disaster recovery vFiler unit is interrupted or not completed.

Steps

1. Depending on the error message you get, perform the following step:

If you get	Then
A "volume offline or does not exist" error message	a. Make the volume online or create it.b. Go to Step 4.
Volume resync error messages	Reactivate the original vFiler unit by using SnapMirror commands.
No volume resync error messages	Go to Step 2.

2. Enter the following command to check whether the vFiler unit you were resynchronizing exists on the storage system on which you were running the vfiler dr resync command:

vfiler status

3. Depending on whether the vFiler unit exists, perform the following step:

If	Then
The vFiler unit exists	Enter the following command to destroy it:
	vfiler destroy vfilername

If	Then
The vFiler unit does not exist	Go to Step 4.

4. Enter the following command to re-create the vFiler unit you destroyed:

vfiler create -r vfilername pathname

5. Enter the following command to resynchronize the vFiler unit:

```
vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] dr_vfilername@disaster_recovery_filer
```

Related tasks

Reactivating the original vFiler unit by using SnapMirror commands on page 102 *Resynchronizing the vFiler unit* on page 98

Reactivating the original vFiler unit by using SnapMirror commands

If the storage system is not damaged but has failed temporarily and if the storage element associated with the vFiler unit is a qtree, you can reactivate the vFiler unit by using SnapMirror commands.

Steps

1. Boot the original storage system and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress.

Note: If you do not press the Del or Esc key in time, you can press Ctrl-C when prompted later during the boot, choose option 5 (maintenance mode), and enter the halt command.

2. At the loader prompt, set the no-vfiler-ips? parameter as follows:

setenv no-vfiler-ips? true

This ensures that the storage system does not try to bind IP addresses already being used by the disaster recovery vFiler unit. When the storage system boots, the original vFiler unit starts running. However, it does not accept any read or write requests because its interfaces are not configured.

3. Resynchronize the mirrored volumes and qtrees.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system that you are trying to activate:

snapmirror resync -S disaster_recovery_filer:/pathname original_filer:/ pathname

Example

snapmirror resync -S drfiler:/vol/vfiler1/qtree1 prfiler:/vol/vfiler1/ qtree

If the snapmirror resync command fails with a message that there are no matching Snapshot copies, you might have accidentally deleted the Snapshot copies that SnapMirror depends on. You must then initialize SnapMirror by using the snapmirror initialize command, and then go to Step 4 but skip Step 7.

For more information about the snapmirror initialize command, see the na_snapmirror(1) man page.

4. To verify that the snapmirror.access option on the disaster recovery storage system is set to legacy, enter the following command on the disaster recovery storage system:

options snapmirror

5. Depending on the result of the options snapmirror command, perform the following step:

If the options snapmirror command returns	Then
snapmirror.access legacy	Edit the /etc/snapmirror.allow file and add the host name of the original storage system if it is not already added to the file.
A list of host names that does not contain the name of the original	Use the options snapmirror command to add the host name of the original storage system. For example:
storage system	options snapmirror.access host=fridge,toaster,prfiler

- 6. Run the setup command on the disaster recovery vFiler unit and unconfigure its IP addresses.
- 7. Update the data on the original vFiler unit.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system:

snapmirror update -S disaster_recovery_filer:/pathname original_filer:/ pathname

Example

snapmirror update -S drfiler:/vol/vfiler1/qtree1 prfiler:/vol/vfiler1/ qtree1

8. Stop SnapMirror transfers to the disaster recovery vFiler unit.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system:

snapmirror quiesce pathname

Example

snapmirror quiesce /vol/vfiler1/qtree1

Note: This operation can take a long time. Use Ctrl-C to interrupt it, if required.

9. Check that all the paths are quiesced by entering the following command:

snapmirror status

The status column in the output should show each path as Quiesced.

10. Break the SnapMirror relationship.

For each volume and qtree owned by the original vFiler unit, enter the following command on the original storage system:

snapmirror break pathname

Example

snapmirror break /vol/vfiler1/qtree1

- **11.** On the original vFiler, run the setup command to configure the vFiler unit's IP addresses and the NIS and DNS servers.
- **12.** If the storage units that have been copied contain iSCSI LUNs, check that the iSCSI configuration on the original vFiler unit is not affected.

You might have to remap the LUNs and re-create initiator groups (igroups).

- **13.** If the storage units that have been copied contain iSCSI LUNs, bring the LUNs back online on the original vFiler unit.
- **14.** To stop the disaster recovery vFiler unit, enter the following command on the disaster recovery storage system:

vfiler stop vfilername

vfilername is the name of the disaster recovery vFiler unit.

You have completed reactivating the original storage system.

15. From the disaster recovery storage system, resynchronize the original vFiler unit and the disaster recovery vFiler unit.

Related tasks

Resynchronizing the vFiler unit on page 98

Reactivating the original vFiler unit by using vfiler dr commands

If the storage system is severely damaged, or you are not familiar with using SnapMirror commands, you can reactivate the vFiler unit by using vfiler dr commands.

Steps

1. Boot the original storage system and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress.

Note: If you do not press the Del or Esc key in time, you can press Ctrl-C when prompted later during the boot, choose option 5 (maintenance mode), and enter halt.

2. At the loader prompt, set the no-vfiler-ips? parameter as follows:

setenv no-vfiler-ips? true

This ensures that the storage system does not try to bind IP addresses already being used by the disaster recovery vFiler unit.

- 3. At the loader prompt, enter the boot command.
- 4. To destroy the original vFiler unit, enter the following command on the original storage system:

vfiler destroy vfilername

vfilername is the name of the original vFiler unit.

- 5. Stop the disaster recovery vFiler unit by using the vfiler stop command.
- 6. Create the disaster recovery vFiler unit.
- 7. Update the data on the original storage system.

For each volume and qtree owned by the new vFiler unit, enter the following commands on the original storage system (or its replacement) to update the data:

snapmirror break name

snapmirror resync -S disaster_recovery_filer:name production_filer:name

disaster_recovery_filer is the name of the disaster recovery storage system.

production_filer is the name of the original storage system.

name is the volume name or path name of the qtree.

Example

For the volume vol1, enter the following commands:

snapmirror break drfiler:vol1

snapmirror resync -S drfiler:vol1 prfiler:vol1

Example

For the qtree qtree1, enter the following commands:

snapmirror break drfiler:/vol/vol2/qtree1

snapmirror resync -S drfiler:/vol/vol2/qtree1 prfiler:/vol/vol2/qtree1

8. Create a new vFiler unit on the original storage system (or its replacement).

The original vFiler unit is now reactivated on the original storage system or its replacement.

9. Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

If the storage elements for the vFiler unit are	Th	ien
Volumes	On the disaster recovery storage system, enter the following command:	
	vfiler dr resync [-l authinfo] [-a alt-remote, alt-local] [-c secure] [-s] original_vfilername@original_filer	
Qtrees	a.	On the disaster recovery storage system, destroy the disaster recovery vFiler unit by entering the following command on the disaster recovery hosting storage system:
		vfiler destroy vfilername <i>vfilername</i> is the name of the original vFiler unit.
	b.	On the disaster recovery storage system, re-create the disaster recovery vFiler unit.

Related concepts

Disaster recovery using MultiStore on page 85

Related tasks

Creating a disaster recovery vFiler unit on page 93 Re-creating the vFiler unit on a replacement storage system on page 107 Checking and preparing the storage system on page 85 Checking the network on page 88

Re-creating the vFiler unit on a replacement storage system

You can re-create the original vFiler unit on a replacement storage system if the original storage system is damaged beyond repair.

Steps

- 1. Boot the replacement storage system.
- 2. Stop the disaster recovery vFiler unit by using the vfiler stop command.
- 3. Prepare the new vFiler unit on the original storage system (or its replacement).
- **4.** To update the data on the original storage system for each volume and qtree owned by the new vFiler unit, enter the following commands on the original storage system (or its replacement):

snapmirror break name

```
snapmirror resync -S disaster_recovery_filer:nameproduction_filer:name
```

disaster_recovery_filer is the name of the disaster recovery storage system.

production_filer is the name of the original storage system.

name is the volume name or path name of the qtree.

Example

For the volume vol1, enter the following commands:

snapmirror break drfiler:vol1

snapmirror resync -S drfiler:vol1 prfiler:vol1

Example

For the qtree qtree1, enter the following commands:

snapmirror break drfiler:/vol/vol2/qtree1

snapmirror resync -S drfiler:/vol/vol2/qtree1 prfiler:/vol/vol2/qtree1

5. Create the new vFiler unit on the original storage system (or its replacement).

The original vFiler unit is now reactivated on the original storage system or its replacement.

6. Reinstate the disaster recovery vFiler unit on the disaster recovery storage system:

108 | Data ONTAP 7.3 MultiStore Management Guide

If the storage elements for the vFiler unit are	Tł	ien	
Volumes	Or	On the disaster recovery storage system, enter the following command:	
	vf [-	iler dr resync [-l authinfo] [-a alt-remote, alt-local] c secure] [-s] original_vfilername@original_filer	
Qtrees	a.	On the disaster recovery storage system, destroy the disaster recovery vFiler unit by entering the following command on the disaster recovery hosting storage system:	
		vfiler destroy vfilername vfilername is the name of the original vFiler unit.	
	b.	On the disaster recovery storage system, re-create the disaster recovery vFiler unit.	

Related concepts

Disaster recovery using MultiStore on page 85

Related tasks

Reactivating the original vFiler unit by using vfiler dr commands on page 105 Creating a disaster recovery vFiler unit on page 93 Resynchronizing the vFiler unit on page 98
Data migration using MultiStore

MultiStore enables you to migrate data from one storage system to another without extensive reconfiguration on the destination storage system. Migration moves a specified vFiler unit from a remote storage system to a local one. Migration is initiated on the destination storage system that hosts the vFiler unit after the migration.

Migrating data across storage systems enables you to manage the workload efficiently. Migration automatically destroys the source vFiler unit and activates the destination vFiler unit. The destination then starts serving data to its clients automatically. Only the vFiler unit configuration is destroyed on the source, not the data contained in the vFiler unit.

When performing offline migration, static route entry can be added, if required, because static routing information is not carried to the destination storage system.

However, when performing online migration, static routing information is carried to the destination storage system.

MultiStore supports both online migration and offline migration. Offline migration is the default method of migration.

Next topics

Secure communication for setting up data migration on page 109 How migrating a vFiler unit affects clients on page 110 Offline migration of vFiler units on page 110 What Data Motion is on page 120 Target portal group management for online migration of vFiler units on page 126

Related tasks

Checking and preparing the storage system on page 85 *Checking the network* on page 88

Secure communication for setting up data migration

Data migration implemented over SSL provides both authentication and confidentiality of user data.

RSH is an insecure command channel in which passwords and other confidential information are sent across the network in plain text. RSH is disabled in some deployments for security reasons. Therefore, the data migration feature is not available in those deployments.

How migrating a vFiler unit affects clients

When you migrate a vFiler unit to another storage system on the same subnet, you have to reconnect CIFS clients and remount NFS qtree-level mounts. However, you do not have to remount NFS volume-level mounts. You can continue to access LUNs without any interruption.

Although an iSCSI host is briefly disconnected from the source vFiler unit, an initiator hides this brief disruption from applications accessing the LUNs.

Offline migration of vFiler units

Data ONTAP supports two methods of offline migration. Data from a source vFiler unit is copied to a destination vFiler unit by using SnapMirror. Alternatively, vFiler units can be migrated from the source storage system to the destination by using SnapMover. During offline migration, users cannot access data from the vFiler units.

The period of time from when the source vFiler unit becomes inaccessible and the destination vFiler unit becomes available is called the cutover period.

The cutover period can take from a few minutes to several minutes, depending on the number of volumes owned by the vFiler unit. Because users and other applications cannot access the vFiler unit that is being migrated during the cutover period, offline migration is disruptive.

Next topics

The vfiler migrate commands on page 111 Migrating a vFiler unit by copying data on page 111 Adjusting client and network configurations if migrating to a different subnet on page 113 vFiler unit migration without copying data on page 114 Prerequisites for vFiler unit migration between the nodes of an active/active configuration on page 114 Enabling SnapMover vFiler unit migration on page 115 Guidelines for setting up volumes to support SnapMover vFiler unit migration on page 116 SnapMover vFiler unit migration on page 117 Migrating a vFiler unit using SnapMover on page 117 Disabling SnapMover vFiler unit migration on page 118

Related concepts

Online migration of vFiler units on page 120

The vfiler migrate commands

You use two commands to migrate a vFiler unit by using SnapMirror: vfiler migrate start and vfiler migrate complete.

The vfiler migrate start command does the following:

- Checks if the destination storage system is capable of receiving the source data.
- Configures and runs SnapMirror to copy the data from the source to the destination vFiler unit.
- Saves the quota information from the /etc/quotas file of the source vFiler unit.

The vfiler migrate complete command does the following:

- Stops the source vFiler unit.
- Updates the data on the destination vFiler unit.
- Breaks the SnapMirror relationships.
- Configures IP bindings on the destination vFiler unit according to the information you provided when running the vfiler migrate start command, and adds the destination IP information to the destination /etc/rc file.

Any IP information that pertains only to the source vFiler unit is removed from the destination / etc/rc file.

- Configures any quota information saved by the vfiler migrate start command.
- Destroys the source vFiler unit.
- Brings LUNs online using migrated LUN maps and igroups.

Migrating a vFiler unit by copying data

You might want to copy data from one storage system to another for moving the workload from an older vFiler unit that is to be replaced.

Before you begin

Ensure that the following prerequisites are met:

- Your storage systems and network are ready for migration.
- There are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vfiler context.
- You have prepared the destination storage system.
- SnapMirror is licensed and enabled on both the source and the destination storage systems.
- The source and destination storage systems can communicate with each other over the network (for example, by means of DNS lookup or entries in the /etc/hosts file).
- The destination volumes are online.
- To log in to the source storage system, you know the administrative user ID and password.

About this task

Attention: This procedure destroys the original vFiler unit after replicating it on the destination storage system. You can retain the original vFiler unit and create its backup copy for disaster recovery purposes.

Steps

1. On the destination storage system, enter one of the following commands:

If the time by when the console is locked	Then
Is important (for example, if you want to lock the console for a minimum amount of time)	Enter the following command:
	vfiler migrate start source_vfiler@source_filer
Is not important (for example, if you	Enter the following command:
overnight)	vfiler migrate source_vfiler@source_filer
	Note: If you use this command, skip Step 4 and Step 5. For more information, see the na_vfiler(1) man page.

- **2.** Respond to the login prompt with a valid administrative ID and password for the source storage system.
- 3. Respond to the IP address and binding prompts.
- 4. Monitor the progress of the migration by using the following command:

```
vfiler migrate status source_vfiler@source_filer
```

Note: The vfiler migrate command might take some time to complete, especially if a source qtree has many millions of inodes.

5. When the status command reports that SnapMirror has replicated all the storage units of the source vFiler unit, you can either complete the migration or cancel the migration:

If you want to	Then
Complete the migration	Enter the following command:
	vfiler migrate complete [-l user:passwd] [-c secure] source_vfiler@source_filer
	Note: You have to use the user name and password when the password is changed at the remote storage system after vfiler migrate start is started.

If you want to	Then
Cancel the	Enter the following command:
migration	vfiler migrate cancel source_vfiler@source_filer
	This destroys the destination vFiler unit and removes the SnapMirror and other migration-related configuration information from the source vFiler unit.

6. If you copied quota information to the destination storage system's /etc/quotas file when you prepared the destination storage system, activate the quotas on that storage system. For activating quotas on each of the volumes, use the following command:

quota on volume_name

7. Remount the qtrees. However, volume-level mounts remain unchanged.

If you have moved the vFiler unit to a different subnet, CIFS domain, or Windows domain, you must rerun the CIFS setup. Also, you have to make adjustments on the clients and modify data-file security attributes.

8. Reconfigure iSCSI authentication if the vFiler unit's storage units contain iSCSI LUNs.

For instructions, see the Data ONTAP Block Access Management Guide for iSCSI and FC.

Related tasks

Checking and preparing the storage system on page 85 Checking the network on page 88 Adjusting client and network configurations if migrating to a different subnet on page 113 Creating a disaster recovery vFiler unit on page 93

Related references

Storage checklist on page 88 *Network checklist* on page 92

Adjusting client and network configurations if migrating to a different subnet

If you have moved a vFiler unit to a different subnet, you might have to configure the network servers and make adjustments on the clients. If you have moved the vFiler unit to a different Windows domain, you might also have to modify data-file security attributes, and run CIFS setup.

About this task

You must perform this task on the destination vFiler unit.

Steps

1. To configure NIS and DNS servers, run setup.

- 2. To change the name of the Windows domain controller, use the cifs prefdc command.
- 3. To change the Windows WINS server, run the cifs setup command.

Note: If the Windows domain has changed, you might have to change the permissions on the Windows data files to allow your users the same access they had in the old domain.

- 4. To change the trusted host, perform the following steps:
 - **a.** Edit the vFiler unit's /etc/hosts.equiv file, adding the name of the trusted host for administering the vFiler unit.
 - **b.** Add the path to the root volume and the name of the trusted host to the vFiler unit's /etc/ exports file.

Example

/vol/vol0 access=adminhost, root=adminhost

5. Remount volumes and qtrees on the clients.

Related references

Network checklist on page 92

vFiler unit migration without copying data

SnapMover vFiler unit migration in an active/active configuration is the no-copy transfer of a volume-level vFiler unit from the source node to the other. The SnapMover feature uses softwarebased disk ownership to transfer ownership of the aggregate that contains the vFiler unit from the original source node to the destination node in an active/active configuration.

The SnapMover feature uses software-based disk ownership to transfer ownership of the aggregate that contains the vFiler unit from the original source node to its destination node in an active/active configuration. Because the SnapMover feature performs the vFiler unit migration through transfer of disk ownership rather than by copying data from one set of disks to another, the migration operation is quickly completed.

For V-Series systems, you can use SnapMover vFiler unit migration between any of the nodes in a V-Series neighborhood. The procedures are the same as in an active/active configuration.

For more information about V-Series neighborhoods, see the V-Series Installation Requirements and Reference Guide.

Prerequisites for vFiler unit migration between the nodes of an active/ active configuration

Before you use SnapMover to migrate a vFiler unit between the nodes of an active/active configuration, ensure that the MultiStore license is enabled on each node of the configuration. Also, if any license is enabled on the source node, it must also be enabled on the destination node.

The prerequisites are as follows:

• Other licenses used by the vFiler unit must match.

For example, if CIFS is licensed on the source node of the active/active configuration, CIFS must also be licensed on the destination node of the active/active configuration. Otherwise, moving the vFiler unit causes CIFS to be unavailable for that vFiler unit after it is moved.

• Both the nodes of the active/active configuration must be connected to the same storage subsystem.

The disks must be visible to the nodes of the source and destination active/active configuration.

- To ensure that software-based disk ownership changes are transparent to NFS users, the destination node must have an Ethernet connection to the same subnet that the source node uses.
- The volumes assigned to the vFiler unit can be either traditional volumes or FlexVol volumes. If the volumes are FlexVol volumes, the containing aggregate must contain only volumes belonging to the migrating vFiler unit. For information about traditional and FlexVol volumes, see the *Data ONTAP Storage Management Guide*.
- The vFiler unit's storage units must all be composed of complete volumes; that is, the vFiler unit's paths must use the form /vol/volname. SnapMover migration of storage units that name specific volume subdirectories—for example, /vol/volname/directory or /vol/volname/qtree is not supported.
- The volume containing the configuration information for the vFiler unit (/vol/volname) must be writable.
- Ensure that the destination storage system does not contain any aggregate or volume with the same name as the aggregates in the source storage system.

Enabling SnapMover vFiler unit migration

You must reconfigure the storage system configured as an active/active configuration before migrating a vFiler unit.

Steps

1. On both the source and destination nodes in the active/active configuration, confirm that MultiStore is licensed by entering the following command:

license

2. Temporarily disable the active/active configuration by entering the following command:

cf disable

3. On both the source and destination nodes in the active/active configuration, confirm that the software-based disk ownership is enabled by entering the following command in the command-line interface of one of the nodes in the active/active configuration:

disk show

• If the system displays disks in the active/active configuration to which disk ownership is assigned, then software-based disk ownership is enabled. Go to Step 10.

- If the system does not displays disks in the active/active configuration to which disk ownership is assigned, then go to Step 4.
- **4.** Reboot both the source and destination nodes in the active/active configuration. During the boot process, press Ctrl-C to display the boot menu options.
- 5. Enter the choice for booting in maintenance mode.
- **6.** On both the source and destination nodes in the active/active configuration in maintenance mode, enter the following command:

disk upgrade-ownership

This command writes software-based disk ownership information to enable SnapMover.

7. Enter the following command on both the source and destination nodes in the active/active configuration to confirm the new software-based disk ownership scheme:

disk show

The system displays all the disks that have ownership assigned in the active/active configuration. Disks assigned to both the source and destination nodes in the active/active configuration are visible.

8. Halt both the source and destination nodes in the active/active configuration to exit from the maintenance mode by entering the following command:

halt

- 9. Reboot both the source and destination nodes in the active/active configuration in normal mode.
- 10. Reenable active/active configuration by entering the following command:

cf enable

Related tasks

Disabling SnapMover vFiler unit migration on page 118

Guidelines for setting up volumes to support SnapMover vFiler unit migration

When a vFiler unit is migrated, all volumes associated with that vFiler unit are moved. Note that SnapMover cannot migrate a subset of the volumes that are managed by the vFiler unit it is migrating.

When you create vFiler units, keep in mind the following points:

- The destination node of the active/active configuration must be able to accommodate the vFiler unit with all its associated volumes.
- The names of the vFiler units and volumes being moved from the source to the destination must be unique on the destination.

Although you can rename a volume at the destination, it is best not to do so. This is because if you have NFS clients, renaming of the volume is not transparent to the NFS clients. When the storage system uses NFS to export a file system, the volume name is part of the exported path

name. NFS clients try to mount by using the old path name. Therefore, to access the data after the vFiler unit has been migrated, clients must remount by using the new path name.

SnapMover vFiler unit migration

To perform a vFiler unit migration using SnapMover, you can use the vfiler migrate -m nocopy command.

The vfiler migrate -m nocopy command performs the following:

- Verifies that no vFiler unit with the same name exists on the destination node of the active/active configuration.
- Verifies that both the source and destination nodes in the active/active configuration run the same version of Data ONTAP.
- Saves the IP configuration and binding information that you supplied when you created the vFiler unit.
- Saves the quota information from the source vFiler unit's /etc/quotas file.
- Stops the source vFiler unit.
- Destroys the source vFiler unit.
- Rewrites the disk ownership information so that the ownership of the vFiler unit volumes is transferred from the source node in the active/active configuration to the destination node in the active/active configuration.
- Re-creates the vFiler unit on the destination node in the active/active configuration.

Migrating a vFiler unit using SnapMover

If many clients are using the same vFiler unit, thereby affecting the performance of its hosting node, and the other node of the active/active configuration is lightly loaded, you can transfer ownership of that vFiler unit to the hosting node's active/active configuration partner to balance the load processing on the two nodes.

Before you begin

You must ensure that there are no FlexClone file and FlexClone LUN operations running on the storage unit from the nondefault vfiler context.

Steps

1. On the destination node in the active/active configuration, enter the following command:

vfiler migrate -m nocopy vfilername@source_cl_partner

vfilername is the name of the vFiler unit that you are migrating.

source_cl_partner is the active/active configuration from which you are moving the vFiler unit.

For detailed information, see the na_vfiler(1) man page.

- 2. Answer the prompts, including the following information:
 - A valid administrative login ID and password
 - The IP address and binding information for the destination vFiler unit

The vFiler unit is migrated from the source node active/active configuration to the destination node in the active/active configuration.

3. Verify that the vFiler unit was moved by entering the following command on the destination node in the active/active configuration:

vfiler status -r vfilername

Disabling SnapMover vFiler unit migration

If the load on the storage system in an active/active configuration is too high and you want to migrate vFiler units to another storage system, you have to disable the SnapMover vFiler unit migration.

Steps

1. On one of the nodes in the active/active configuration, enter the following command to verify whether the current software-based disk ownership assignments align with the active/active configuration's A loop B loop topology:

aggr status -r

For backward compatibility, you can also enter the vol status -r command.

For more information about ensuring disk ownership within disk shelves and loops, see the *Data ONTAP Active/Active Configuration Guide*.

In the output of this command, all disks assigned to the current node of the active/active configuration are listed as volume disks or as spare disks. All disks assigned to the partner node in the active/active configuration are listed as partner disks.

All disks on the same shelf must be assigned to the same node in the active/active configuration and all disk shelf assignments must be consistent with the A loop B loop topology of the active/ active configuration.

2. Change the disk ownership if it deviates for the A loop B loop topology:

If the disk ownership	Then
Deviates from the A loop B loop topology because you migrated a vFiler unit with the vfiler migrate -m nocopy command	 Perform the following: Change it to the A loop B loop by using the vfiler migrate -m nocopy command. Go to Step 3.

If the disk ownership	Then
Deviates from the A loop B loop topology because you split ownership of a single expansion disk shelf between the two nodes of an active/active configuration	 Perform the following: Modify the disk assignment by using the disk assign command. This command assigns all disks on the shelf to the node attached to the shelf's A loop connection. Go to Step 3.
Deviates from the A loop B loop topology because you reassigned a spare disk with the disk assign command	 Perform the following: Modify the disk assignment by using the disk assign command to change it to the A loop B loop topology. Go to Step 3.
Reflects the active/active configuration's A loop B loop topology	Go to Step 3.

3. Disable the active/active configuration by entering the following command:

cf disable

- **4.** Reboot each node in the active/active configuration, and when prompted press Ctrl-C to display the boot menu options.
- **5.** On each node in the active/active configuration, enter the choice for booting in maintenance mode.
- **6.** In the maintenance mode, on each node in the active/active configuration, enter the following command:

disk remove_ownership

This command disables software-based disk ownership.

7. Halt each node in the active/active configuration to exit maintenance mode by entering the following command:

halt

8. Reboot each node in the active/active configuration in normal mode.

The system reverts to ownership based on A loop and B loop connections.

9. Reenable active/active configuration by entering the following command:

cf enable

Related tasks

Migrating a vFiler unit using SnapMover on page 117

What Data Motion is

Data Motion is a data migration solution that integrates virtual storage, mirroring, and provisioning software technologies so that you can perform migrations nondisruptively in both physical and virtual environments.

By using the Provisioning Manager interface, you can migrate data from one storage system to another, as long as the data is contained in vFiler units and associated with datasets. Migration operations are performed transparently, so users are unaware of the migration operation being performed, and nondisruptively, so users retain access to migrated data, and the hosts and applications that access the migrated data do not require reconfiguration.

The application interfaces and documentation commonly refer to the Data Motion capability as "online migration," "online dataset migration," or "online vFiler unit migration."

Next topics

Online migration of vFiler units on page 120 Licenses required for online migration on page 121 Stages of a vFiler unit migration on page 121 How to perform online migration of vFiler units on page 122 Viewing the status of a vFiler unit migration on page 123 Commands not allowed during online migration on page 123

Online migration of vFiler units

You must use Provisioning Manager to perform online migration of vFiler units. During online migration, users of vFiler units experience only a delay or pause while accessing or writing data. This is because the cutover period is short and can be a maximum of 120 seconds.

You do not have to stop and restart applications running on vFiler units during online migration. Therefore, users and client applications of vFiler units do not experience any disruption in data availability. This method of migration is useful when users require uninterrupted access to the vFiler units.

Online migration is nondisruptive only for vFiler units configured to use NFS or iSCSI protocols. If the vFiler unit is configured to use CIFS, the CIFS connection will be terminated by Provisioning Manager before starting online migration. Provisioning Manager reconnects all terminated CIFS connections after completion of online migration.

For more information about online migration, see the *Provisioning Manager and Protection Manager Printable Help for use with DataFabric Manager Server 4.0* and the *Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators.*

Related concepts

What data migration is on page 0

Offline migration of vFiler units on page 110 *Data migration using MultiStore* on page 109

Related tasks

Checking and preparing the storage system on page 85 *Checking the network* on page 88

Related references

Storage checklist on page 88 *Network checklist* on page 92

Licenses required for online migration

You must license HTTP on both the source and the destination storage systems to perform online migration of vFiler units. You must also enable the httpd.enable and the httpd.admin.enable options in the source and the destination storage systems.

Stages of a vFiler unit migration

Online migration of vFiler units consists of three stages: vFiler unit migrate start, vFiler unit migrate prepare, and vFiler unit migrate complete. You can view the different states of the migrating vFiler unit by using the vfiler status command or by viewing messages on the console of the hosting storage system.

vFiler unit migrate start stage

A SnapMirror relationship is established and baseline transfer is initialized during the vFiler unit migrate start stage.

The vFiler unit in the source storage system is in the running state and is accessible by users connected to the vFiler unit.

Also, a temporary vFiler unit is created in the destination storage system. This temporary vFiler unit is in the stopped, migrating state.

vFiler unit migrate prepare stage

During the vFiler unit migrate prepare stage, the source and destination vFiler units are prepared for the final stage of migration. The vFiler unit in the source storage system is accessible by users connected to the vFiler unit and is in the running, migrate_prepare state.

The temporary vFiler unit that was created in the destination storage system during the previous stage is in the stopped, migrating state.

In the vFiler unit migrate prepare stage, MultiStore verifies the following:

• The SnapMirror relationship between the source and the destination is synchronized.

- The aggregates that own the volumes assigned to the vFiler unit are not in degraded or reconstructing state at the source or at the destination storage system.
- The destination or source storage system is in takeover mode. If either of the storage systems is in takeover mode, then the migrate complete stage returns an error message.
- The destination or source storage system is in the process of being taken over.
- Volumes are added or removed from source storage system after the vFiler unit migrate start stage.
- The volumes owned by the vFiler unit being migrated are online.
- The source storage system has no other vFiler units in the cutover stage of the migration.

vFiler unit migrate complete stage

In the vFiler unit migrate complete stage, MultiStore verifies the following:

- The SnapMirror relationship between the source and the destination is synchronized.
- The aggregates that own the volumes assigned to the vFiler unit are not in degraded or reconstructing state at the source or at the destination storage system.
- The destination or source storage system is in takeover mode. If either of the storage systems is in takeover mode, then the migrate complete stage returns an error message.
- The destination or source storage system is in the process of being taken over.
- Volumes are added or removed from source storage system after the vFiler unit migrate start stage.
- The volumes owned by the vFiler unit being migrated are online.
- The source storage system has no other vFiler units in the cutover stage of the migration.

The vFiler unit in the source storage system is now in the stopped, migrating_source state and is not accessible by users connected to the vFiler unit.

The temporary vFiler unit that was created in the destination storage system is now in the running state. This vFiler unit is now accessible by the users who were connected to the source vFiler unit. Users do not have to reconfigure or restart any applications they were running on the original vFiler unit.

How to perform online migration of vFiler units

You can perform online migration of vFiler units only from Provisioning Manager. Provisioning Manager performs premigration storage checks, postmigration tasks, and automates features such as destination volume provisioning, network provisioning, rapid rollback, and rapid restart of online migration.

When you start online migration from Provisioning Manager, all SnapMirror, SnapVault, NDMP, dump and restore operations that are currently running on the vFiler unit that is being migrated are aborted before cutover is initiated by Provisioning Manager. If there are any currently running deduplication, FlexClone file and LUN, or LUN clone split operations, online migration fails. You must retry online migration after these operations are complete. New SnapMirror and SnapVault operations are not allowed.

If online migration fails, error messages are displayed on the console of the hosting storage system. You can also view the log of error messages in the /etc/messages file of the hosting storage system.

Before you start online migration from Provisioning Manager, it is best to configure static routes on the destination storage system, when the source storage system has more than 20 routes.

For more information about using Provisioning Manager to perform online migration, premigration storage checks, and postmigration tasks, see the *Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators* and the *Provisioning Manager and Protection Manager Printable Help for use with DataFabric Manager Server 4.0.*

Viewing the status of a vFiler unit migration

You can view the messages on the CLI of the hosting storage system or the default vFiler unit vfiler0 to check the status of the migrating vFiler unit. You can also the use the vfiler status command to view the status of the source or destination vFiler units.

Step

1. Enter the following command in the CLI of the default vFiler unit vfiler0:

vfiler status vfiler_name

vfiler_name is the name of the vFiler unit.

Example

Assume a vFiler unit vfilerA is being migrated. The following command shows the status of the vFiler unit vfilerA:

vfiler status vfilerA

The output of this command shows the migration stage of the vFiler unit vfilerA at that point in time. For example, if the vFiler unit is in the migrate prepare stage, the vfiler status vfilerA command shows the following output:

```
vfiler status -r vfilerA
vfilerA running, migrate_prepare
ipspace: default-ipspace IP address: 12.21.1.1 [e0a]
Path: /vol/vol_test1 [/etc]
Path: /vol/smvol4
UUID: 679b86c0-be3a-11de-aee6-00a09807609b
```

Commands not allowed during online migration

During online migration of a vFiler unit, you must not run Data ONTAP commands from the CLI of the hosting storage system or the default vFiler unit vfiler0.

MultiStore commands that are not allowed during online migration

• vfiler add

- vfiler allow
- vfiler context
- vfiler disallow
- vfiler destroy
- vfiler dr resync
- vfiler move
- vfiler remove
- vfiler rename
- vfiler run
- vfiler start
- vfiler stop

Volume and aggregate commands that are not allowed during online migration

- vol copy
- vol rename
- vol online
- vol offline
- vol restrict
- vol options
- vol size
- aggr split

For more information about volume and aggregate commands, see the *Data ONTAP Storage Management Guide*.

Snapshot commands that are not allowed during online migration

- snap create
- snap delete
- snap list
- snap delta
- snap rename
- snap sched
- snap restore
- snap reclaimable
- snap reserve
- snap autodelete

For more information about Snapshot commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide.*

SnapMirror commands that are not allowed during online migration

- snapmirror resync
- snapmirror initialize

For more information about SnapMirror commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide.*

FlexClone file and FlexClone LUN commands that are not allowed during online migration

clone start

For more information about FlexClone file and FlexClone LUN commands, see the *Data ONTAP Storage Management Guide*.

Deduplication commands that are not allowed during online migration

- sis on
- sis off
- sis config -s
- sis start
- sis start -s

For more information about deduplication commands, see the *Data ONTAP Storage Management Guide*.

Tape backup and NDMP commands that are not allowed during online migration

- dump
- restore
- ndmpcopy

For more information about NDMP commands, see the *Data ONTAP Data Protection Tape Backup* and *Recovery Guide*.

Related concepts

Data access from the hosting storage system on page 23

Related tasks

Moving resources between vFiler units on page 36 Adding resources to a vFiler unit on page 35 Destroying a vFiler unit on page 41 Renaming a vFiler unit on page 40 Starting a vFiler unit on page 42 Stopping a vFiler unit on page 40 Allowing a protocol on a vFiler unit on page 44 Disallowing a protocol on a vFiler unit on page 45 Resynchronizing the vFiler unit on page 98

Target portal group management for online migration of vFiler units

Target portal groups enable you to efficiently manage iSCSI sessions between initiators and targets. Although Data ONTAP manages target portal groups by network interface by default, you can also manage target portal groups by IP address starting with Data ONTAP 7.3.3. This is required if you want to perform an online migration of vFiler units, which allows you to nondisruptively migrate data from one storage system to another.

Note: Provisioning Manager is required for performing online migrations of vFiler units.

When you migrate data, the target portal group tag on the destination network interface must be identical to the target portal group tag on the source network interface. This is problematic in a MultiStore environment because the source and destination storage systems may be of different hardware platforms. Changing the target portal group tags after migration is not sufficient because some hosts, such as HPUX and Solaris, do not support dynamic iSCSI target discovery, resulting in a disruption of service to those hosts in the process.

Note: If offline (disruptive) migrations are not problematic in your environment, or if all of your hosts support dynamic iSCSI target discovery, then IP-based target portal group management is unnecessary.

If you do choose to implement IP-based target portal groups by enabling the iscsi.ip_based_tpgroup option, interface-based target portal groups are automatically converted to IP-based target portal groups, and any future target portal group assignments will be IPbased as well. However, note that if you are migrating between a system with IP-based target portal groups and a system with interface-based target portal groups, the target portal group information is lost and the iSCSI service may be disrupted.

Note: ALUA is not supported with IP-based target portal groups.

For more information on Provisioning Manager, see the *Provisioning Manager and Protection Manager Administration Guide*.

Next topics

Data migration implications for IP-based target portal group management on page 127 Enabling IP-based target portal group management on page 128 Creating IP-based target portal groups on page 130 Adding IP addresses to IP-based target portal groups on page 131 Removing IP addresses from IP-based target portal groups on page 131 *Destroying IP-based target portal groups* on page 132 *Displaying IP-based target portal group information* on page 132

Related information

Documentation on NOW - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

Data migration implications for IP-based target portal group management

It is important to understand how target portal groups are affected before beginning the migration.

The following table describes the various online migration scenarios and how they affect your target portal group configurations.

Online migration scenario	Impact to target portal groups
Migration between two systems running Data ONTAP 7.3.2 and earlier.	No change - existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between a system running Data ONTAP 7.3.2 or earlier and a system running 7.3.3 or later.	No change - existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between two systems running Data ONTAP 7.3.3 or later with interface-based target portal groups.	No change - existing interface-based target portal groups are not converted to IP-based target portal groups.
Migration between two systems running Data ONTAP 7.3.3 or later with IP-based target portal groups.	IP-based target portal group assignments are preserved and there is no disruption in the iSCSI service.
Migration between a system running Data ONTAP 7.3.3 or later with interface-based target portal groups and a system running 7.3.2 or earlier with interface-based target portal groups.	No change - target portal group assignments are preserved.
Migration between a system running Data ONTAP 7.3.3 or later with IP-based target portal groups and any system with interface-based target portal groups.	Target portal group information is lost and the iSCSI service may be disrupted.

Related information

Data ONTAP documentation on NOW - now.netapp.com/NOW/knowledge/docs/ontap/ ontap_index.shtml

Enabling IP-based target portal group management

If you want to perform online migrations in a MultiStore environment, you must enable IP-based target portal groups on your vFiler units.

About this task

When you enable IP-based target portal groups, the existing interface-based target portal groups are automatically converted to IP-based target portal groups. However, note that the interface-based target portal groups remain intact for the default vFiler unit.

Step

1. Enter the following command:

vfiler run vFiler unit options iscsi.ip_based_tpgroup on

The existing interface-based target portal groups are converted to IP-based target portal groups with no disruption in service to the host.

Example

Before enabling IP-based target portal groups, the target port group information for vFiler unit 2 (vf2) looks like this:

system1:	vfiler run vf2 iscs	i tpgroup show
TPGTag	Name M	lember Interfaces
32	user_defined32	(none)
1000	e0_default	e0
1002	ellb_default	ellb
1003	ellc_default	ellc
1004	elld_default	elld
1005	e9a_default	e9a
1006	e9b_default	e9b
1007	e10a_default	el0a
1008	e10b_default	e10b
2000	vif_e0-1_default	vif_e0-1
2001	vif_e0-2_default	vif_e0-2
2002	vif_e0-3_default	vif_e0-3
2003	vif_ella-1_default	vif_ella-l
2004	vif_ella-2_default	vif_ella-2
2005	vif_ella-3_default	vif_ella-3

Each interface is associated with various IP addresses, and some of those are assigned to vFiler unit vf2. For example:

```
system1> vfiler run vf2 iscsi portal show
Network portals:
IP address TCP Port TPGroup Interface
10.60.155.104 3260 1000 e0
192.168.11.100 3260 2003 vif_ella-1
192.168.11.101 3260 2003 vif_ella-1
```

192.168.13.10032602005vif_ella-3192.168.13.10132602005vif_ella-3

After enabling IP-based target portal groups for vf2, the relevant interface-based target portal groups for vf2 are nondisruptively converted to IP-based target portal groups.

system1> vfiler run vf2 options iscsi.ip_based_tpgroup on system1> vfiler run -q vf2 iscsi ip_tpgroup show TPGTag Name Member IP Addresses 1000 e0 default 10.60.155.104 2003 vif_ella-1_default 192.168.11.100, 192.168.11.101 2005 vif_ella-3_default 192.168.13.100, 192.168.13.101 system1> vfiler run -q vf2 iscsi portal show Network portals: TCP Port TPGroup Interface IP address 10.60.155.104 1000 3260 e0 192.168.11.100 3260 2003 vif_ella-1 192.168.11.101 3260 2003 vif ella-1 192.168.13.100 3260 2005 vif_e11a-3 3260 2005 vif_ella-3 192.168.13.101 If you configure another IP address for vf2, then a new default IP-based target portal group (4000) is automatically created. For example: system1> vfiler add vf2 -i 192.168.13.102 system1> ifconfig vif_e11a-3 alias 192.168.13.102 system1> vfiler run vf2 iscsi ip_tpgroup show TPGTag Name Member IP Addresses 10.60.155.104 1000 e0_default 192.168.11.100, 192.168.11.101 2003 vif_ella-1_default 2005 vif_ella-3_default 192.168.13.100, 192.168.13.101 4000 192.168.13.102_default 192.168.13.102 system1> vfiler run vf2 iscsi portal show Network portals: IP address TCP Port TPGroup Interface 10.60.155.104 1000 e0 3260 192.168.11.100 3260 2003 vif_ella-1 192.168.11.101 3260 2003 vif_ella-1 2005 vif_e11a-3 192.168.13.100 3260 2005 vif_ella-3 192.168.13.101 3260 192.168.13.102 3260 4000 vif_e11a-3

After you finish

After you enable IP-based target portal group management, it is recommended to leave it enabled. However, if you must disable IP-based target portal groups for some reason, enter the following command:

options iscsi.ip_based_tpgroup off

As a result, any IP-based target portal group information will be discarded, and the interface-based target portal group information is re-enabled. Note that this process might disrupt the iSCSI service to the hosts.

Also note that if an IP address is unassigned from a vFiler unit or unconfigured from the network interface, that IP address is no longer a valid iSCSI portal. However, the IP-based target portal group to which that IP address belonged remains intact so that if you add the IP address back at some point, it is automatically assigned back to the original target portal group.

Creating IP-based target portal groups

You can create new IP-based target portal groups in which to add and remove existing IP addresses.

Before you begin

Before creating the target portal groups, make sure you enable IP-based target portal group management by entering the following command:

options iscsi.ip_based_tpgroup on

Step

1. Enter the following command:

```
vfiler run vFiler unit ip_tpgroup create [-f] [-t | tag] tpgroup_name IP address...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

-t tag sets the target portal group tag to the specified value. In general you should accept the default tag value.

tpgroup_name is the target portal group name.

IP address is the list of IP addresses to include in the group, separated by spaces.

xample	
filer run vfiler2 iscsi ip_tpgroup create -t 233 vfiler2_t	:pg1
0.1.3.5	

After you finish

After you create a new IP-based target portal group, you can add and remove IP addresses from the new group.

Adding IP addresses to IP-based target portal groups

Use the iscsi ip_tpgroup add command to add an IP address to an existing IP-based target portal group.

Before you begin

Ensure you have enabled IP-based target portal group management and that there is at least one existing IP-based target portal group.

Step

1. Enter the following command:

```
vfiler run vFiler unit iscsi ip_tpgroup add [-f] tpgroup_name IP address ...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

tpgroup_name is the target portal group name.

IP address is the list of IP addresses to include in the group, separated by spaces.

Example

vfiler run vfiler2 iscsi ip_tpgroup add vfiler2_tpg1 192.168.2.1 192.112.2.1

Removing IP addresses from IP-based target portal groups

In the course of reconfiguring your network, you might need to remove one or more IP addresses from an IP-based target portal group.

Step

1. Enter the following command:

vfiler run vFiler unit iscsi ip_tpgroup remove [-f] tpgroup_name IP address ...

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

tpgroup_name is the target portal group name.

IP address is the list of IP addresses to remove from the group, separated by spaces.

Example

vfiler run vfiler2 iscsi ip_tpgroup remove vfiler2_tpg1 192.112.2.1

Destroying IP-based target portal groups

If necessary, you can destroy IP-based target portal groups.

About this task

If there are active iSCSI sessions when you destroy the group, those sessions will be lost.

Step

1. Enter the following command:

```
vfiler run vFiler unit iscsi ip_tpgroup destroy [-f] tpgroup_name
```

-f forces the group to be destroyed, even if that terminates an existing session using one of the IP addresses in the group.

tpgroup_name is the target portal group name.

The target portal group is destroyed, and if there are active iSCSI sessions, a warning message displays indicating that those connections will be lost.

Example

vfiler run vfiler2 iscsi ip_tpgroup destroy vfiler2_tpg1

Displaying IP-based target portal group information

Use the iscsi ip_tpgroup show command to display important information about your IP-based target portal groups, including target portal group tags, target portal group names, and the IP addresses that belong to each group.

Step

1. Enter the following command:

vfiler run vFiler unit iscsi ip_tpgroup show

```
Example

system1> vfiler run vfiler2 iscsi ip_tpgroup show

TPGTag Name Member IP Addresses

1 vfiler2_migrate_test0 (none)

2 vfiler2_migrate_test1 (none)

3 vfiler2_migrate_test3 (none)

100 user_defined_tp1 (none)
```

128	vfiler2_ui_review	1.1.1.1
1007	e10a_default	10.1.1.8
1008	e10b_default	1.1.1.2
4000	10.1.1.5_default	10.1.1.5
4001	10.60.155.104_default	10.60.155.104
4002	192.168.1.1_default	192.168.1.1

Disk space management using quotas

You can apply user, group, and qtree quotas on a vFiler unit in the same way that you apply on a storage system. When you create a vFiler unit, quotas are automatically turned off on both the hosting storage system and the new vFiler unit.

However, quotas are turned off for all the volumes that you assign to the new vFiler unit, and for all the volumes from which you assign qtrees to the new vFiler unit.

Quotas are also turned off for volumes that you move from one vFiler unit to another. To activate quotas again, you must allow them and turn them on.

On a hosting storage system licensed for vFiler units, the hosting storage system administrator must allow quotas for a volume before you can turn on quotas or turn off quotas for the volume. By default, quotas are allowed on all volumes. Only hosting storage system administrators can allow or disallow quotas for a volume.

Next topics

Allowing or disallowing quotas for a volume on page 135 Quota specification management on page 136 Turning on or turning off quotas from a vFiler unit on page 137 When quota thresholds and soft quotas are exceeded on page 137 How you can resize quotas on page 138 How the quotas file works on page 138 Displaying the quota status on page 138 Displaying a quota report on page 139

Allowing or disallowing quotas for a volume

As a hosting storage system administrator, you must allow quotas for a volume on the hosting storage system before you can turn on or turn off quotas. By default, quotas are allowed on all volumes.

Step

1. To allow or disallow quotas for a volume, enter the following command:

quota allow | disallow volume

volume is the name of the volume for which you want to allow or disallow the quota.

After you enter the quota allow command, you can turn on quotas for the specified volume from a vFiler unit.

After you enter the quota disallow command, vFiler units are prevented from turning quotas on for the specified volume. If quotas are currently turned on for any volume in vFiler units, they are turned off immediately.

Result

If you disallow quotas on a volume, the following effects occur on all vFiler units that have storage units in the volume:

- If quotas are currently turned off, you or the vFiler unit administrator cannot turn on quotas for that volume.
- If quotas are currently turned on, they are turned off immediately and cannot be turned back on.

Quota specification management

The vFiler unit administrator specifies the size of each quota in the vFiler unit's /etc/quotas file. The vFiler unit administrator tracks and limits the amount of disk space and the number of files each user, group, or qtree uses.

If a qtree owned by the vFiler unit resides on a volume owned by the hosting storage system, then the hosting storage system administrator can also specify a quota for the qtree in the hosting storage system's /etc/quotas file. The following example shows how qtree quota on the hosting storage system affects a vFiler unit qtree.

How qtree quota on the hosting storage system affects a vFiler unit qtree

Assume that the /vol/voll/qtree1 qtree is a storage unit of the vFiler unit, and the /vol/ voll volume is owned by the hosting storage system.

In the /etc/quotas file of the vFiler unit, the vFiler unit administrator specifies that this qtree is limited to 20 GB of disk space. In the /etc/quotas file of the hosting storage system, the storage system administrator can specify can specify the disk space limit for the qtree as 10 GB. Therefore, if quotas are turned on from the hosting storage system for the / vol/vol1 volume, the qtree cannot exceed the limit in either of the /etc/quotas files, whichever is lower. In this example, the qtree cannot exceed 10 GB.

The hosting storage system administrator controls the usage of quotas on each volume that the storage system owns by using the quota allow and quota disallow commands. When the hosting storage system administrator allows quotas, the vFiler unit administrator can turn quotas on or off on the vFiler units by using the quota on and quota off commands, respectively.

For more information about quotas, see the Data ONTAP Storage Management Guide.

Turning on or turning off quotas from a vFiler unit

You can turn on quotas by using the quota on volume command, and you can turn off quotas by using the quota off volume command. The quota on volume command activates quotas on the specified volume based on the contents of /etc/quotas.

About this task

Changes made to /etc/quotas do not take effect the next time the quota on or quota resize command is executed. Turning quotas off by using the quota off volume command deactivates quotas on the specified volume. You can turn quotas on and off on a per-volume basis for a vFiler unit. After you turn on quotas for a particular volume, Data ONTAP initializes quotas for the storage units residing on the volume that is owned by the vFiler unit. The on or off states of quotas are persistent and stay set after reboots.

Step

1. To turn quotas on or off for a volume owned by a vFiler unit, follow the instructions appropriate to your situation:

If you manage the vFiler unit from	Then
The hosting storage system	Enter the following command:
	vfiler run vfilertemplate quota on off volume
The vFiler unit	Enter the following command through an RSH connection to the vFiler unit:
	quota on off volume

Note: Whenever a qtree is explicitly reassigned to a vFiler unit, you must reenable the quota manually if quotas are used. Qtrees are explicitly reassigned to vFiler units when you create vFiler units (using the vfiler create command) or when you move qtrees between vFiler units (using the vfiler move, vfiler add, or vfiler remove commands).

When quota thresholds and soft quotas are exceeded

When a threshold or soft quota defined on a vFiler unit is exceeded, a warning message is logged on the storage system console.

You see a warning message similar to the following:

```
[vfiler1@quota.softlimit.exceeded:notice]: Threshold exceeded for tree 3 on
volume vol1 for vfiler "vfiler1"
```

How you can resize quotas

When you use the quota resize command, Data ONTAP rereads the quotas file for the specified volume. You can resize quotas only for certain types of changes to the quotas file, otherwise, you have to reinitialize quotas.

For more information about resizing quotas, see the Data ONTAP Storage Management Guide.

How the quotas file works

The quotas file, found in the /etc directory, contains one or more entries specifying limit or tracking quotas for qtrees, groups, and users. The file can contain default (general) and specific entries.

Displaying the quota status

You can display the quota status for any volume on which your vFiler unit owns storage space.

Step

1. To display the quota status, complete the following step:

If you manage the vFiler unit from	Then
The hosting storage system	Enter the following command:
	vfiler run vfilertemplate quota
The vFiler unit	Enter the following command through an RSH connection to the vFiler unit:
	quota

The command displays quota status information about all the volumes in which the vFiler unit owns storage space. The following is a sample message in the command output:

vol0: quotas are on. vol1: quotas are off. vol2: quotas are disabled.

For more information about quotas, see the Data ONTAP Storage Management Guide.

Displaying a quota report

You display a quota report using the quota report command. You can display a quota report for all quotas or for a specific file, directory, qtree or volume by specifying a pathname.

Step

1. To display a quota report, enter the following command:

quota report [path]

You can display a quota report for all quotas or for a specific file, directory, qtree or volume by specifying a path.

You can control the format and fields displayed using the quota report command options. For more information on the available options, see the na_quota(1) man page.

Index

/etc/dgateways file 57 /etc/exports file exporting all file systems in 81

A

adding resources to a vfiler unit 35 aggr add command adding new disks 86 aggr status command 118

B

backing up a vFiler unit 52 backup vFiler unit 93

С

CIFS support for a vFiler unit 43 local user accounts 82 path names for shares 79 statistics 65 virus scanning 83 cifs setup command 32 cifs stat command 65 clients, effect of vFiler unit move on 110 commands halt 119 cf disable 119 cf enable 119 ipspace assign 89 ipspace create 89 ipspace list 89 options snapmirror 103 options snapmirror.access 103 remove ownership 119 vfiler dr configure 89, 94 vfiler migrate 90 vpwhich 88 commands for a vFiler unit vfiler command, purpose of 30 entering through rsh 49 vfiler context command 46 vfiler dr activate command 97

vfiler migrate start command 96, 111 vfiler move command 36 vfiler remove command 35 vfiler setup command 32 vfiler start command 42 vfiler status command 45, 85 configuring virus scanning 82, 83 consolidating multiple servers 19 context of vFiler unit, switching to 46 create a vFiler unit prerequisites 27 creating a vFiler unit in nondefault IPspaces 76 required status of network interface 30 Creating a vFiler unit 31 creating volumes 87 cutover period online migration 120

D

daemon enabling the routed 26 data migration 109 Data Motion online migration component 120 deduplication 59 default IPspace, interfaces 69 destination storage system, preparing for migration 85 destroying a vFiler unit 41 destroying IPspaces 75 disabling MultiStore license 27 disabling SnapMover vFiler unit migration 118 disaster recovery quotas 87 vFiler unit 93 disk upgrade-ownership command 116 displaying quota status 87, 138 distinct IP address space 67 domains DNS servers and vFiler unit 20 NIS servers 20

F

FCP LUNs 56 FilerView configuring resources 31 FlexClone files and FlexClone LUNs interoperability with MultiStore 61 FlexVol volume 50 FTP support for a vFiler unit 43

H

hosting storage system access to data on a vFiler unit 23 manage a vFiler unit 34 quotas not copied to new vFiler unit 87 HTTP support for a vFiler unit 43

I

IP addresses configuring for vFiler unit creation 30 IPsec, configuring vFiler unit for 57 IPspace secure routing 68 typical applications 68 IPspaces guidelines 67 interfaces 69 naming requirement 71 routing tables incoming packets for an IPspace, outgoing packets for an IPspace 70 iSCSI and a vFiler unit 55 iSCSI support for a vFiler unit 43

L

language, guidelines for /etc/quotas file etc/usermap.cfg ,language for encoding 28 language for encoding 28 license for MultiStore, enabling or disabling 26 local user accounts 82 LUNs on a vFiler unit 54

Μ

managing storage system

volumes, disks, and RAID groups backups and data recovery 23 mandatory_scan option 83 maximum number of vFiler units 38 migrate by copying data 110 storage system data 22 Migration of vFiler units 109 moving a vFiler unit 87 moving resources, about 35 multiple security domains 20 multiple server consolidation 19

N

NDMP 53 network checks for migration to destination storage system 88 Network Data Management Protocol 53 network interfaces configuring down 30 for a vFiler unit 22 network resources base IP address IP alias 30 requirements for moving and removing 35 network resources, storage resources 35 NFS support for a vFiler unit 43 path names for exporting 79 starting the protocol 80 statistics 65 nfsstat command 65 NIS servers and migration, disaster recovery 90 no-vfiler-ips? variable, setting 105 not supported commands MultiStore commands ndmp commands 123-125 SnapMirror commands 123-125

0

offline migration 110 online migration about 120 licenses required 121 options 121 supporting Data Motion 120

P

partitioning system resources 21 path names, for NFS exports and CIFS shares 79 performance, monitoring 64 primary unit 28 protocols allowing or disallowing, CIFS or NFS 44 supported for a vFiler unit 43

Q

qtree command output (how it differs for a vFiler unit and hosting storage system) 51 qtrees who can create on a vFiler unit 51 quota report command 87 quota reports displaying 139 quotas effects of destroying a vFiler unit 41 guideline for 29 prerequisite for turning on and off 135, 137 resizing 138 types supported for a vFiler unit 135 who specifies 136

R

rebooting storage system, effects on a vFiler unit 50 removing IP addresses from an interface 74 removing IP aliases from an interface IP alias 30 resizing quotas 138 resources assigning 31 guidelines for assigning 28 restricting storage system traffic 21 routed daemon and IPspace 73 effects of disabling 57 enabling 26 routing table for the storage system 57 vFiler unit in default IPspace 57 rsh access to vFiler unit from clients 47 enable option 47 RSH support for a vFiler unit 43

S

SAN guidelines 29

Server Manager 81 setting up a vFiler unit 32 setup command for a vFiler unit 32 sis 59 SnapMirror options snapmirror.access command 103 using to reactivate vFiler unit 102 snapmirror commands snapmirror break command 104 snapmirror quiesce 104 snapmirror status 104 snapmirror update command 103 snapmirror update command 107 snapmirror.allow file 103 **SnapMover** disabling 118 SnapMover: enabling 115 software-based disk ownership disk upgrade-ownership command 116 enabling 115 SSH 43, 48 SSL 93, 109 starting a vFiler unit 42 storage checklist 88 storage checks for migration 85 storage checks for migration to destination storage system 92 storage resources assigning volumes 28 requirements for moving and removing 35 storage system partitioning 19 storage system reboot, effects on a vFiler unit 50 storage system resources partitioning 21 subnet, moving vFiler unit 113

Т

target portal groups adding IP addresses to IP-based groups 131 data migration implications for 127 deleting IP-based groups 132 displaying information about IP-based groups 132 enabling IP-based 128 traditional volume for a vFiler unit 50 trusted host changing after migration 114

U

uptime command 65 User Manager 81

V

vfiler allow 44 vfiler rename 39 vFiler unit increasing the limit 37 decreasing the limit 39 default 22 disaster recovery checking storage space 92 disaster recovery: 98 displaying status 45 maximum number 37 migration disabling 118 SnapMover, enabling 115 migration stages 121, 122 moving to different subnet 113 to different Windows domain 113 protocols supported 22 reactivating using vfiler dr commands 105 reactivating via SnapMirror commands 102 rename 40 replacing on original storage system 107

resources, IP addresses, moving resources, adding and removing vFiler unit resources 34 resynchronization (resync) handling failures 101 setup 32 states, stopped or running 40 viewing status 123 vFiler unit limit 38 vFiler unit migration using the vfiler migrate command 117 by copying data 111 vfiler0 included in vFiler unit limit 37 vfilertemplate, defined 30 virus scanning registering scanners 83 requirements for 83 VLAN tagging for traffic separation tagging for more IPspaces 70 vol create command 87 volumes in a vFiler unit effects of renaming 51 taking offline 51

W

WINS server and migration, disaster recovery 91 changing after migration 114