

Data ONTAP™ 6.4 System Administration **MultiStore™ Management Guide**

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

Part number 210-00326
January 2003

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademark information

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

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

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

Table of Contents

	Preface	.ix
Chapter 1	About the MultiStore Software Product	1
	Understanding MultiStore	2
	System administration for a filer running MultiStore	6
Chapter 2	Managing MultiStore	9
	Establishing a vfiler.	10
	Enabling or disabling the vfiler license	11
	Planning for vfiler creation.	12
	Creating a vfiler	14
	Understanding the vfiler command family.	18
	Setting up a vfiler manually	19
	Administering vfiler resources from the hosting filer	21
	Changing resources for vfilers	22
	Adjusting the vfiler limit	25
	Stopping and starting vfilers	27
	Destroying a vfiler	29
	Allowing or disallowing protocols for vfilers	30
	Displaying vfiler status	32
	Autosupport and syslogd messages about vfilers.	33
	Executing commands and setting options for vfilers	34
	Effects of filer reboot on vfilers	37
	Understanding volumes and qtrees on vfilers	38
	Networking considerations	40
	Understanding SnapMirror on the hosting filer	41
	Understanding SNMP when used with vfilers	42
	Performance monitoring and statistics	43

Chapter 3	Using Vfilers in Different IPspaces.	45
	Understanding IPspaces	46
	Advantages of using VLAN tagging for IPspaces	50
	Requirements for using clusters and IPspaces	51
	Case studies for IPspaces	53
	Managing IPspaces on a filer	54
	Creating an IPspace	55
	Listing IPspaces on a filer	56
	Assigning an interface to an IPspace	57
	Destroying an IPspace	59
	Creating a vfiler in a nondefault IPspace	60
 Chapter 4	 File Access Using NFS and CIFS	 63
	How to specify path names for NFS exports or CIFS shares.	64
	Preparing the vfiler for NFS	65
	Preparing the vfiler for CIFS	66
 Chapter 5	 Disk Space Management Using Quotas	 67
	Understanding quotas.	68
	Turning quotas on and off	69
	How to resize quotas	71
	Understanding the /etc/quotas file	72
	Displaying quota status	73
	Understanding quota reports	74
	Messages from exceeded quota thresholds and soft quotas	76
 Chapter 6	 Disaster Recovery and Data Migration	 77
	Understanding disaster recovery and data migration.	78
	How to create and activate a disaster-recovery vfiler	79
	How to move a vfiler	80
	Preparing the destination filer	81
	Storage and network checklists.	92

Creating the disaster-recovery vfiler	95
Activating the disaster-recovery vfiler	98
Reactivating the original vfiler	100
Moving the vfiler	107
Chapter 7	
Virus Protection for CIFS	113
Understanding virus scanning for vfilers	114
Configuring virus scanning for vfilers	115
Index	117

Preface

About this guide

This guide describes how to use the MultiStore™ software product on Network Appliance™ filer that run Data ONTAP™ 6.4 software and are Serving Data By Design™. It covers all NetApp—the Network Appliance Company™ filer models.

Audience

This guide is for system administrators who are familiar with operating systems that run on the filer's clients, such as UNIX®, Windows 95™, Windows NT®, and Windows® 2000. It also assumes that you are familiar with how to configure the filer and how the NFS, CIFS, and HTTP protocols are used for file sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the NetApp filer.

Terminology

When the term “filer” is used, the information applies to all NetApp® filer models.

This guide uses the term “type” to mean pressing one or more keys on the keyboard. It uses the term “enter” to mean pressing one or more keys and then pressing the Enter key.

Command conventions

You can enter filer commands on either the system console or from any client computer that can access the filer through a Telnet session.

In examples that illustrate commands executed on a UNIX workstation, this guide uses the command syntax of SunOS 4.1.x. The command syntax and output might differ, depending on your version of UNIX.

Keyboard conventions

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

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d <i>hostname</i></code> command, you enter the characters “arp -d” followed by the actual name of the host. Book titles in cross-references.
<code>Monospaced font</code>	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
<code>Bold monospaced font</code>	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note

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

Caution

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

WARNING

A warning contains instructions that you must follow to avoid personal injury.

About this chapter This chapter introduces you to the MultiStore software product.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Understanding MultiStore](#)” on page 2
- ◆ “[System administration for a filer running MultiStore](#)” on page 6

Understanding MultiStore

What MultiStore is

MultiStore is an optional software product that enables you to partition the storage and network resources of a single filer so that it appears as multiple filers on the network. Each “filer” created as a result of the partitioning is called a vfiler™ unit. A vfiler, using the resources assigned, delivers file services to its clients as a filer does.

The storage resource assigned to a vfiler 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.

Prerequisites for using MultiStore

Because MultiStore is an optional software product, you must enable the software license for it. The name of the license is vfiler.

Reasons for using MultiStore

MultiStore is typically used for the following reasons:

- ◆ You want to consolidate multiple servers to one filer.
- ◆ You use the filer to host data for multiple customers, such as clients of a service provider or different organizations within an enterprise.
- ◆ You use the SnapMirror™ feature to migrate data from one filer to another.

MultiStore for server consolidation: The following scenarios illustrate how you can use MultiStore to consolidate servers:

- ◆ Suppose you currently manage multiple servers and you want to store all data on one filer for easier administration.

To consolidate the server, you can partition the filer into vfilers and then copy the data from the servers to the vfilers. If the vfilers are for CIFS users, you can set up the vfilers 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 vfilers 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.
- ◆ A filer without MultiStore can participate in only one security domain, so if your environment requires that different groups of CIFS users be in different domains, you must use multiple filers. MultiStore, however, enables you to

install each vfiler in the appropriate domain while keeping all the data on the same physical filer.

- ◆ Because you can set up NIS and DNS servers for individual vfilers, after you consolidate the servers on one filer, network clients of the vfilers can continue to use the same NIS and DNS servers as before.

MultiStore for service providers and enterprises: Service providers, such as Internet service providers (ISPs) and storage service providers (SSPs), can partition the resources of a filer to create many vfilers for client companies. Similarly, the information technology (IT) department of an enterprise can create vfilers for various organizations, or customers, within the enterprise.

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

For example, an SSP can create the following vfilers on a filer:

- ◆ A vfiler named VfilerA. It uses the /vol/vol1 volume and the e0 interface on the filer. It is leased to CompanyA.
- ◆ A vfiler named VfilerB. It uses the /vol/vol2 volume and the e1 interface on the filer. It is leased to CompanyB.

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

See [Chapter 3](#) on page 45 for a more detailed discussion.

MultiStore for data migration: MultiStore enables you to migrate data from one filer to another without extensive reconfiguration on the destination filer.

Without MultiStore, after you use SnapMirror to copy data from one filer to another, you must edit the files' access control lists (ACLs), local user group definitions, user mapping information, and so on, before users can access the data. If the data being copied is stored on a vfiler, however, all user, group, and ACL information is encapsulated in the vfiler. You re-create a vfiler on the destination filer with the encapsulated information, and data can be served from the destination filer without further reconfiguration.

About the hosting filer

The filer on which you create vfilers is called the hosting filer. The storage and network resources used by the vfilers exist on the hosting filer.

About the default vfiler

Licensing MultiStore automatically causes Data ONTAP to create a default vfiler on the hosting filer. The name of the default vfiler is vfiler0.

Initially, vfiler0 owns all the resources on the filer. After you create vfilers, vfiler0 owns the resources that are not owned by the vfilers you create.

The default vfiler exists as long as the vfiler license is enabled. On a filer with the vfiler license enabled, you cannot destroy vfiler0.

All information provided in this guide about vfilers is applicable to vfiler0, unless noted otherwise.

Some limitations of MultiStore

Limitations of MultiStore include the following.

- ◆ You can create up to 32 vfilers on a filer, depending on the memory capacity of the hosting filer.

A filer containing the maximum number of vfilers actually has 33 vfilers: vfiler0 plus the 32 vfilers that you, the system administrator, can create.

See “[Adjusting the vfiler limit](#)” on page 25 for more information.
- ◆ Only Ethernet interfaces are supported for vfilers.
- ◆ Only the NFS and CIFS file protocols are supported for vfilers.
- ◆ Not all Data ONTAP commands are supported on vfilers.

Hosting filer's access to data contained in vfilers

You, as the hosting filer administrator, do not have CIFS or NFS access to the data contained in vfilers, except for the information in vfiler0. After you assign a qtree or volume to a vfiler, you lose access to the data in that qtree or volume.

Example: Before you create a vfiler with the /vol/vol1 volume, you can configure the /etc/exports file so that you can mount the /vol/vol1 volume. After you create the vfiler, an attempt to mount the /vol/vol1 volume from the hosting filer results in the “Stale NFS file handle” error message.

Sample vfiler

A filer named toaster has the vfiler license enabled and is configured with three volumes, which are /vol/vol0, /vol/vol1, and /vol/vol2.

Because of the vfiler license, a default vfiler, vfiler0, is created on toaster.

On toaster, you can create a vfiler named vfiler1. You can assign the IP address 123.123.123.1, the qtree /vol/vol0/qtree1, and the volume /vol/vol1 to vfiler1. Traffic destined for vfiler1 uses the IP address 123.123.123.1, and all data served by vfiler1 is stored in the assigned qtree and volume.

Because toaster hosts vfiler1, toaster is the hosting filer. It uses all the IP addresses created on toaster except 123.123.123.1. It uses the following resources for storage:

- ◆ Space in the /vol/vol0 volume that is not in /vol/vol0/qtree1 qtree
- ◆ The /vol/vol2 volume

System administration for a filer running MultiStore

Types of system administration tasks

The following list describes the types of system administration tasks that you perform when managing a filer running MultiStore:

- ◆ Hosting filer tasks, which are the same as those for a filer not supporting vfilers
- ◆ MultiStore management tasks, which are related to establishing vfilers and managing resources on vfilers
- ◆ Tasks related to using Data ONTAP features supported by vfilers

Hosting filer tasks

You can perform tasks related to managing the resources on the hosting filer in the same way as you perform them on a filer without a vfiler license. The following tasks are some examples:

- ◆ Managing volumes, disks, and RAID groups
- ◆ Increasing data availability through Snapshot™ management, SnapMirror management, and volume copy
- ◆ Backing up and recovering data

These tasks are covered in detail in the other volumes of the Data ONTAP System Administration set, particularly the *File Access Management Guide*, the *Storage Management Guide*, and the *Data Protection Guide*.

MultiStore management tasks

From the hosting filer, you can manage MultiStore by

- ◆ Enabling and disabling the vfiler license
- ◆ Allowing and disallowing protocols to be run on vfilers
- ◆ Creating vfilers
- ◆ Setting up vfilers
- ◆ Starting and stopping vfilers
- ◆ Destroying vfilers
- ◆ Moving resources to and from vfilers
- ◆ Monitoring the status of vfilers

For more information about these tasks, see Chapter 2, “[Managing MultiStore](#),” on page 9.

From a destination filer, you can perform the following additional tasks:

- ◆ Moving a vfiler from one hosting filer to another
- ◆ Creating a disaster-recovery vfiler

See Chapter 6, “[Disaster Recovery and Data Migration](#),” on page 77 for more information.

Tasks related to using Data ONTAP on vfilers

You can use Data ONTAP features on individual vfilers in the same way as you use them on a filer. Some of the tasks related to these features can be performed by vfiler administrators. These tasks are as follows:

- ◆ Managing users. Information about managing users (for example, information about using the `/etc/usermap.cfg` file for mapping users) is the same for filers and vfilers. For more information, see the *Storage Management Guide*.
- ◆ Managing NFS exports and CIFS shares. For more information, see “[File Access Using NFS and CIFS](#)” on page 63.
- ◆ Managing quotas. For more information, see “[Disk Space Management Using Quotas](#)” on page 67.
- ◆ Preparing vfilers for virus protection. For more information, see “[Virus Protection for CIFS](#)” on page 113.

About this chapter

This chapter provides information about MultiStore management tasks that can only be performed from the hosting filer.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Establishing a vfiler”](#) on page 10
- ◆ [“Administering vfiler resources from the hosting filer”](#) on page 21
- ◆ [“Changing resources for vfilers”](#) on page 22
- ◆ [“Adjusting the vfiler limit”](#) on page 25
- ◆ [“Stopping and starting vfilers”](#) on page 27
- ◆ [“Destroying a vfiler”](#) on page 29
- ◆ [“Allowing or disallowing protocols for vfilers”](#) on page 30
- ◆ [“Displaying vfiler status”](#) on page 32
- ◆ [“Autosupport and syslogd messages about vfilers”](#) on page 33
- ◆ [“Executing commands and setting options for vfilers”](#) on page 34
- ◆ [“Effects of filer reboot on vfilers”](#) on page 37
- ◆ [“Understanding volumes and qtrees on vfilers”](#) on page 38
- ◆ [“Networking considerations”](#) on page 40
- ◆ [“Understanding SnapMirror on the hosting filer”](#) on page 41
- ◆ [“Understanding SNMP when used with vfilers”](#) on page 42
- ◆ [“Performance monitoring and statistics”](#) on page 43

Establishing a vfiler

Tasks for establishing a vfiler

Before you can start protocol servers on a vfiler, you must establish a vfiler on your filer.

For detailed information

For detailed information about how to establish a vfiler, see the following topics:

- ◆ [“Enabling or disabling the vfiler license”](#) on page 11
- ◆ [“Planning for vfiler creation”](#) on page 12
- ◆ [“Creating a vfiler”](#) on page 14
- ◆ [“Understanding the vfiler command family”](#) on page 18
- ◆ [“Setting up a vfiler manually”](#) on page 19

Enabling or disabling the vfiler license

Effects of enabling the vfiler license

Enabling the vfiler license has the following effects on the filer:

- ◆ You can use the `vfiler` and `ipspace` commands.
- ◆ Data ONTAP starts logging vfiler status information and sends vfiler status information using the autosupport feature.
- ◆ The `routed` daemon is disabled.
- ◆ The `ip.match_any_ifaddr` option is set to Off.
- ◆ The “vfiler limit” (the number of vfilers you can create on this filer, including `vfiler0`) is set to a default value between 3 and 11, depending on the memory capacity of the hosting filer. You can change this value; see “[Adjusting the vfiler limit](#)” on page 25 for more information.

Prerequisite for disabling the vfiler license

You can disable the vfiler license only when no vfilers other than `vfiler0` are on the filer. If there are vfilers on the filer, you must destroy them before disabling the vfiler license.

Effects of disabling the vfiler license

Disabling the vfiler license has the following effects:

- ◆ MultiStore becomes unavailable immediately. You can no longer use the `vfiler` and `ipspace` commands.
- ◆ The `routed` daemon is enabled after the next reboot.

Enabling or disabling the vfiler license

To enable or disable the vfiler license, complete the following step.

Step	Action	
1	If you want to...	Enter this command:
	Enable the vfiler license	<code>license vfiler=code</code>
	Disable the vfiler license	<code>license vfiler=disable</code>

Planning for vfiler creation

Prerequisites for vfiler creation

The following prerequisites must be met before you can create a vfiler:

- ◆ You must create the units of storage (qtrees or volumes) before creating the vfiler.
- ◆ The storage unit for vfiler configuration information 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 must be unconfigured when you create the vfiler.

Storage guidelines

When deciding which storage units to assign to vfilers, remember the following guidelines:

- ◆ The first storage unit assigned to the vfiler contains the vfiler configuration information. It is called the primary storage unit. Although you can remove storage units from a vfiler at any time after the vfiler is created, the primary storage unit must remain for as long as the vfiler exists. Therefore, plan carefully the storage unit that will be the primary storage unit.
- ◆ 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. Data in qtrees and volumes used as other storage units is preserved.
- ◆ A volume assigned to a vfiler must not be the filer's root volume. You can, however, assign qtrees in the root volume to a vfiler.
- ◆ If the vfiler administrator needs to create qtrees on the vfiler, assign volumes instead of qtrees to the vfiler. This is because qtrees can be created only at the root of a volume.
- ◆ If, in the future, you want to move from one filer to another the disks that are used for the vfiler's storage, assign volumes, not qtrees, to the vfiler.
- ◆ When managing NFS exports, CIFS shares, quotas, and options, vfiler administrators need to enter the complete path names of the storage resources used by the vfilers in commands and configuration files. Therefore, choose volume and qtree names appropriately, so that you can share with the vfiler administrators the complete path names beginning with *filer_name:/vol/vol_name*.

Naming guidelines	When deciding the vfiler name, remember the following guidelines: ◆ The name can contain up to 31 alphanumeric ASCII characters. ◆ The name can contain the dash (-) and the underscore (_) characters, but the dash character must not be the first character. ◆ The name is case-insensitive. ◆ The name vfiler0 is reserved.
Language guideline	Vfiler administrators need to edit the /etc/quotas and /etc/usermap.cfg files for their vfilers. These files support Unicode and root volume UNIX encoding. To ensure that vfiler administrators who do not use Unicode-capable editors can edit the files, create vfilers on a filer whose root volume language can be used for editing.
Quotas guideline	Creating a vfiler involves changing the ownership of a volume or qtree from the hosting filer to a specified vfiler. This change requires that quotas be turned off for the affected volume. You can turn quotas back on for the volume after the vfiler is created.
Clustering considerations	Remember the following considerations for establishing vfilers on a cluster: ◆ You can create up to 32 vfilers on each member of a cluster, depending on the memory capacity of the hosting filers. See “Adjusting the vfiler limit” on page 25 for details. ◆ A member of a cluster that does not have a vfiler license can take over its partner with a vfiler license. ◆ The vfilers hosted by the filers of the cluster are created and configured independently. That is, each filer can host a different number of vfilers, and the vfiler configuration on the filers can be different from each other. ◆ In takeover mode, the functioning filer takes over all vfilers created on the failed filer. These vfilers include the vfilers you create and vfiler0. Therefore, for vfilers on the failed filer to work correctly after the takeover, each network interface used by a vfiler in a cluster must have a partner interface.

About this section

This section describes how to create a vfiler in the default IP address space (IPspace). For the definition of IPspace and information about creating vfilers in nondefault IPspaces, see Chapter 3, “[Using Vfilers in Different IPspaces](#),” on page 45.

State of the vfiler after creation

Immediately after you create the vfiler, the vfiler is running.

The vfiler is configured with a set of options with default settings. These settings are not inherited from the hosting filer. You or the vfiler administrator can modify the option settings.

The vfiler configuration information is stored in the vfiler’s /etc directory. It contains a subset of files that are in the hosting filer’s /etc directory:

- ◆ cifsconfig.cfg
- ◆ cifssec.cfg
- ◆ exports
- ◆ filersid.cfg
- ◆ group
- ◆ hosts
- ◆ hosts.equiv
- ◆ passwd
- ◆ quotas
- ◆ files related to the registry
- ◆ rmtab
- ◆ sm
- ◆ usermap.cfg

Major steps for creating a vfiler

The following list describes the major steps for creating a vfiler:

- ◆ Ensure that the network interface to be configured with a vfiler IP address is ready for vfiler creation.
- ◆ Assign and configure vfiler resources.

Ensuring that the network interface is ready

To ensure that the network interface to be configured with a vfiler IP address is ready for vfiler creation, complete the following steps.

Step	Action	
1	If the IP address for the vfiler is...	Then...
	A base IP address for an interface	Go to Step 3 .
	An IP alias for an interface	Go to Step 2 .
2	If the IP alias is currently...	Then...
	Assigned to an interface	Enter the following command to remove the alias: ifconfig interface -alias address Example: The following example removes the IP alias from the e0 interface: <pre>ifconfig e0 -alias 123.123.123.123</pre> You are done.
	Unassigned	You are done.

Step	Action	
3	If the base IP address for the vfiler is assigned to...	Then...
	An interface in the UP state	Enter the following command to change the state of the interface for the IP address to DOWN: ifconfig interface down Example: The following example changes the state of the e0 interface to DOWN: ifconfig e0 down
	An interface in the DOWN state	You are done.

Assigning and configuring vfiler resources

You can assign and configure resources for a vfiler either via the command line or via FilerView. To configure the vfiler via the command line, complete the following steps.

Note

The procedure that follows assumes that you want to set up the vfiler for networking (and CIFS if necessary) immediately after creating the vfiler. If you want to defer setup, use the `-n` option of the `vfiler create` command, and then follow directions under [“Setting up a vfiler manually”](#) on page 19 when you are ready to do the setup. For more information on the `vfiler create` command, see the `na_vfiler man` page.

Step	Action
1	Enter the following command: <pre>vfiler create vfiler_name -i ip_address [-i ip_address] ... path [path] ...</pre> <i>vfiler_name</i> is the name of the vfiler. <i>ip_address</i> is an IP address. <i>path</i> 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. Example: The following example creates a vfiler using two IP addresses, one volume, and one qtree: <pre>vfiler create vfiler1 -i 123.123.123.123 -i 123.123.123.124 /vol/vol1 /vol/vol2/mtree2</pre>
2	Respond to prompts to set up the filer, and to set up CIFS if necessary.

What has happened: By completing the setup, you have started NFS and the configured the vfiler’s root volume to be exported to the vfiler’s administrative host (by means of an entry in the /etc/exports file). The setup has also configured the vfiler’s IP addresses configured and added the appropriate entries to /etc/rc.

What to do next

You, or the vfiler administrator, can now mount the root volume of the vfiler, edit /etc/exports to suit your needs and rerun the `exportfs` command. See Chapter 4, “[File Access Using NFS and CIFS](#),” on page 63 for more information.

Understanding the vfiler command family

Purpose of the vfiler command

The `vfiler` command, which is supported only on the hosting filer, enables the hosting filer administrator to set up vfilers, manage vfiler resources, and manage Data ONTAP features on individual vfilers.

General syntax of the vfiler command

The `vfiler` command family has several commands, and each command has its own syntax. The following syntax is the general `vfiler` command syntax:

```
vfiler command vfilertemplate options ...
```

Meaning of vfilertemplate

Some `vfiler` commands support the *vfilertemplate* option. Wherever you see *vfilertemplate* in this manual or the `na_vfiler` man page, you can substitute a vfiler name, a comma-separated list of vfiler names, an IPspace declaration, or an asterisk as a wildcard. If you use the asterisk, the command takes effect on all vfilers, including `vfiler0` (the physical filer), unless the command is inapplicable to `vfiler0`. See the `na_vfiler` man page for more information.

Format of path names in the vfiler command

Some `vfiler` commands include a path name. It is the complete path name for the qtree or volume assigned to the specified vfiler.

Setting up a vfiler manually

When to use this section

If you use the default form of the `vfiler create` command, you are prompted for setup information (and CIFS setup information) as soon as the vfiler has been created.

In that case, you can skip this section unless you need to change the networking configuration.

If you use the `-n` option of the `vfiler create` command, no setup is done, and no protocol servers will run on the vfiler until you set it up manually.

The procedure for setting up a vfiler manually is similar to that for setting up a filer. You run the `setup` program, which enables you to

- ◆ Specify the administration host for the vfiler
- ◆ Specify the password for the root user on the vfiler
- ◆ Enable DNS and specify the DNS domain and servers
- ◆ Enable NIS client service and specify the NIS domain and servers

Note

The `setup` program does not prompt you for the time zone. All vfilers are in the same time zone as the hosting filer.

If the vfiler is licensed to deliver CIFS service, you must run `cifs setup`, as you would for a filer, in addition to running `setup`.

Setting up the vfiler

To set up the vfiler, complete the following steps.

Step	Action
1	Enter the following command: <pre>vfiler run vfiler_name setup</pre> Result: The <code>setup</code> command displays prompts for you to configure the vfiler. After you respond to all the prompts, configuration files, such as the <code>/etc/exports</code> file, are created in the <code>/etc</code> directory for the vfiler. Note _____ Unlike the <code>setup</code> command for the filer, the <code>setup</code> command for vfilers does not cause NFS to start running. For more information about starting NFS, see Chapter 4, “File Access Using NFS and CIFS,” on page 63. _____ If the vfiler runs the CIFS protocol, go to the next step. Otherwise, you are done.
2	Enter the following command: <pre>vfiler run vfiler_name cifs setup</pre> Result: The <code>cifs setup</code> command displays prompts for you to configure CIFS on the vfiler. After you respond to all the prompts, CIFS starts running.

Administering vfiler resources from the hosting filer

Managing vfiler storage from the filer

If you, as the physical filer administrator, need to manage storage resources that belong to a vfiler, but you do not have administrative access to the vfiler, you can do one of the following:

- ◆ Temporarily move the resources to the hosting filer (but you cannot move the vfiler's primary /etc volume)
- ◆ Temporarily destroy the vfiler.
This returns ownership of all resources to the hosting filer. No user data is modified when you destroy a vfiler.

Once you have done what you need to do, you can either move the storage resources back to the vfiler, or, if you destroyed the vfiler, restore the vfiler using the command

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

oldname is the name of the original vfiler and *oldfirstpath* is the first path that was specified in the original `vfiler create` command that created the vfiler. See the `na_vfiler` man page for more information.

Changing resources for vfilers

What you can change

You can add, move, or remove resources for vfilers. The resources can be network resources (that is, IP addresses) or storage resources (that is, volumes or qtrees).

Effects of changing resources

Adding, removing, and moving vfiler resources has the following effects:

- ◆ After you add storage resources to a vfiler, the resources are moved from the hosting filer to a vfiler.
- ◆ After you remove storage resources from a vfiler, the resources are moved from the vfiler to the hosting filer.
- ◆ After you add an IP address to a vfiler, you can assign the address as an IP alias of an interface or assign the address to an unconfigured network interface.
- ◆ After you remove an IP address from a vfiler, the IP address becomes an unassigned IP address.
- ◆ After you move resources from one vfiler to another, the resources are removed from the source vfiler and added to the destination vfiler.

Note

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

Requirements for moving and removing resources

Remember the following requirements when you move or remove vfiler resources:

- ◆ The storage unit to be moved or removed cannot be the one containing the vfiler's /etc directory.
- ◆ The following requirements are for moving an IP address between vfilers or removing an IP address from a vfiler:
 - ❖ Both the source vfiler and destination vfilers must be in the same IPspace.
 - ❖ If the address is an IP alias, the alias must be removed. If the address is not an IP alias, the network interface associated with the address must be unconfigured.

Considerations for moving resources between vfilers

Remember the following considerations when you move resources between vfilers:

- ◆ After you move a storage unit from one vfiler to another, the security information associated with the files in the storage unit is retained. This might cause users to be unable to access files properly. For example, before the move, a particular user ID (UID) was allowed access to a file; after the move, the UID corresponds to a different user on the destination vfiler, and the user previously represented by this UID can no longer gain access the file.
- ◆ If you reassign a volume from one vfiler 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.
- ◆ If you reassign a qtree from one vfiler to another, Data ONTAP turns off quotas for the volume containing the qtree on both the source vfiler and the destination vfiler. After the qtree is moved, you can turn quotas on again for the volume.
- ◆ All network connections to the vfilers are terminated when resources are being moved.

Adding resources to a vfiler

To add resources to a vfiler, complete the following step.

Step	Action
1	Enter the following command: <pre>vfiler add vfiler_name [-f] [-i ip_address [-i ip_address] ...] [path [path ...]]</pre> Use the <code>-f</code> option to skip warning messages. Example: The following example adds an IP address and a volume to an existing vfiler: <pre>vfiler add vfiler1 -i 123.123.123.125 /vol/vol3</pre>

Removing resources from a vfiler

To remove resources from a vfiler, complete the following step.

Step	Action
1	Enter the following command: <pre>vfiler remove vfiler_name [-f] [-i ip_address [-i ip_address] ...] [path [path ...]]</pre> Use the <code>-f</code> option to skip warning messages. Example: The following example removes an IP address and a volume from an existing vfiler: <pre>vfiler remove vfiler1 -i 123.123.123.125 /vol/vol3</pre>

Moving resources between vfilers

To move resources between vfilers, complete the following step.

Step	Action
1	Enter the following command: <pre>vfiler move source_vfiler destination_vfiler [-f] [-i ip_address [-i ip_address] ...] [path [path ...]]</pre> Use the <code>-f</code> option to skip warning messages. Example: The following example moves an IP address and a volume from one vfiler to another: <pre>vfiler move vfiler1 vfiler2 -i 123.123.123.125 /vol/vol3</pre>

Adjusting the vfiler limit

What the vfiler limit is

The vfiler limit sets the maximum number of vfilers that you can create on the hosting filer. In a cluster, the same limit applies to each cluster node.

Default limits

The default vfiler limits for systems on which you have just activated a MultiStore license are as follows:

- ❖ 3 for filers with less than 1 GB of memory
- ❖ 5 for filers with between 1 GB and 3 GB of memory
- ❖ 11 for filers with 3 GB of memory or more

The default is 11 for systems on which a MultiStore license has been in effect since a release of Data ONTAP earlier than 6.4.

These limits include vfiler0 (the hosting filer, which is created when the MultiStore license is enabled and persists as long as the license is in effect) so a limit of 11 means you can create a maximum of 10 vfilers. In a cluster, a limit of 11 means that you can create a maximum of 10 vfilers on each cluster node.

To see the current limit, enter the following command:

```
vfiler limit
```

Increasing the limit

You can use the `vfiler limit` command to increase the limit to a maximum of up to 33, depending on the memory capacity of the hosting filer. The maximum limits are as follows:

- ❖ 11 for filers with less than 1 GB of memory
- ❖ 26 for filers with between 1 GB and 2GB of memory
- ❖ 33 for filers with over 2 GB of memory

Because the limit includes the hosting filer, vfiler0, which always exists, the number of vfilers you can create is one less than the limit.

Example: To raise the number of vfilers you can create to 15, enter the following command on the hosting filer:

```
vfiler limit 16
```

In a cluster, this sets a limit of 16 vfilers on each cluster node.

Note

For the change to take effect, you must reboot the filer (or each filer in a cluster).

Decreasing the limit

You can use the `vfiler limit` command to decrease the limit, setting the maximum number of vfilers lower than it now is.

Because the limit includes the hosting filer `vfiler0`, which already exists, the number of vfilers you can create in each case is one less than the limit.

Example: To reduce the number of vfilers you can create from 15 to 10, enter the following command on the hosting filer:

```
vfiler limit 11
```

Note

For the change to take effect, you must reboot the filer (or each filer in a cluster).

Caution

Before you reboot, make sure that the number of active vfilers on this filer (and on any cluster partner) is less than the limit. On reboot, the number of vfilers started will be at most one less than the limit you specified (one less than the limit on each node in a cluster), even if more vfilers are configured.

Stopping and starting vfilers

Reasons to stop a vfiler

You stop a vfiler if you need to troubleshoot vfiler problems or destroy a vfiler.

Note _____
You cannot stop or start vfiler0.

What it means to stop and start a vfiler

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

The stopped state is not persistent across reboots: after you reboot the filer, the vfiler resumes automatically.

You can start vfilers that are in the stopped state; after a vfiler starts, it is in running state and can receive packets from clients.

Stopping vfilers

To stop vfilers, complete the following step.

Step	Action
1	Enter the following command: <pre>vfiler stop vfilertemplate</pre> Example: The filer supports two vfilers named vfiler1 and vfiler2. The following command stops all vfilers, except vfiler 0: <pre>vfiler stop *</pre> The following messages appear after you enter the command: <pre>vfiler1 stopped vfiler2 stopped</pre>

Starting vfilers

To start vfilers, complete the following step.

Step	Action
1	Enter the following command: vfiler start vfilertemplate Example: The filer supports two vfilers named vfiler1 and vfiler2. The following command starts all vfilers, except vfiler 0, which is already running: vfiler start * The following messages appear after you enter the command: vfiler1 running vfiler2 running

Destroying a vfiler

Effects of destroying a vfiler

Destroying a vfiler has the following effects:

- ◆ It releases to the hosting filer all resources that are associated with the vfiler. User data contained in the vfiler remains intact after you destroy the vfiler.
- ◆ All the vfiler’s IP addresses are unconfigured and the corresponding entries in the filer’s /etc/rc file are removed.
- ◆ If the vfiler destroyed was not in the same IPspace as the hosting filer, the IP addresses previously owned by the vfiler are not available for use after you destroy the vfiler.
- ◆ The effects on quotas are the same as when you move resources from a vfiler to the hosting filer. See “[Considerations for moving resources between vfilers](#)” on page 23 for more information about how moving the resources to the hosting filer affects quotas.

Prerequisites for destroying a vfiler

The vfiler must be stopped before you can destroy it.

Destroying a vfiler

To destroy a vfiler, complete the following steps.

Step	Action
1	Enter the following command to stop the vfiler: vfiler stop vfiler_name
2	Enter the following command: vfiler destroy [-f] vfiler_name Without the -f option, the command displays a confirmation prompt. With the -f option, the command destroys the vfiler immediately. Result: All resources previously owned by the vfiler are returned to the hosting filer.

Allowing or disallowing protocols for vfilers

Protocols that vfilers can use

By default, vfilers can use the protocols that are licensed for the hosting filer. For example, if the hosting filer is licensed for CIFS and NFS, the vfilers created can also use the CIFS and NFS protocols.

You can, however, select the protocols that can be run on the vfilers.

Vfilers support the following protocols:

- ◆ CIFS
- ◆ NFS
- ◆ rsh

Effects of disallowing and allowing CIFS or NFS

If CIFS or NFS is running before you disallow the protocol, it remains running until the filer reboots. Disallowing a protocol, however, causes packets for that protocol destined for that vfiler to be ignored.

After you allow the protocol again, the vfiler resumes processing the packets for that protocol.

If CIFS or NFS is not running and you try to enable it when it is disallowed, an error message appears.

After you allow the protocol again, you must enable the protocol to allow the packets for that protocol to be processed.

Effects of disallowing and allowing rsh

After you disallow rsh, attempts to connect to the vfiler using rsh fail, regardless of the `rsh.enable` option setting.

After you allow rsh again, you can connect to the vfiler using rsh if the `rsh.enable` option is On.

Allowing or disallowing rsh has no effect on the `rsh.enable` option. Also, you can enable or disable the `rsh.enable` option regardless of whether rsh is allowed.

Allowing or disallowing protocols for vfilers

To allow or disallow protocols for vfilers, complete the following steps.

Step	Action	
1	If you want to...	Then...
	Allow a protocol	Go to Step 2 .
	Disallow a protocol	Go to Step 3 .
2	To allow a protocol, enter the following command: <pre>vfiler allow vfiler_template proto=protocol ...</pre> <i>protocol</i> is nfs, cifs, or rsh. Example: The following example allows the NFS and rsh protocols on the vfiler named vfiler1: <pre>vfiler allow vfiler1 proto=nfs proto=rsh</pre> You are done.	
3	To disallow a protocol, enter the following command: <pre>vfiler disallow vfiler_template proto=protocol ...</pre> <i>protocol</i> is nfs, cifs, or rsh. Example: The following command disallows the NFS and rsh protocols on the vfiler named vfiler1: <pre>vfiler disallow vfiler1 proto=nfs proto=rsh</pre>	

Displaying vfiler status

Status information you can display

You can use the `vfiler status` command to display the following types of vfiler status information:

- ◆ Whether the vfiler is stopped or started
- ◆ IPspace
- ◆ IP addresses that have been assigned and, if they are configured, which interfaces they are bound to
- ◆ Path names to the qtrees or volumes assigned
- ◆ UUID, a globally unique user ID assigned to the vfiler
Technical support might need this ID when diagnosing problems related to vfilers.
- ◆ Protocols allowed
- ◆ Protocols disallowed

See the `na_vfiler` man page for more information.

Autosupport and syslogd messages about vfilers

About autosupport messages

A filer licensed for vfilers includes the output of the `vfiler status` command in the autosupport messages.

About syslogd messages

The output of the `vfiler status` command, which indicates whether the vfilers are stopped or running, is logged in the `/etc/messages` file on the hosting filer.

All messages generated by vfilers are logged to the hosting filer's `/etc/messages` file. These messages are preceded with the name of the vfiler. For example, if the vfiler named `vfiler1` exceeds a quota threshold, the following message is generated:

```
[vfiler1@quota.softlimit.exceeded:notice]: Threshold exceeded for  
tree 3 on volume vol1 for vfiler "vfiler1"
```

Executing commands and setting options for vfilers

Methods for executing commands for vfilers

You can enter commands for vfilers in the following ways:

- ◆ From the context of the vfiler, you can enter commands directly.
See “[Entering commands from the vfiler](#)” immediately following.
- ◆ From the hosting filer, you can use the `vfiler run` command to execute commands for the specified vfilers.
See “[Entering commands for vfilers from the hosting filer](#)” on page 35.
- ◆ From the client of a vfiler, you can establish an `rsh` connection with the vfiler and then enter the commands.
See “[Entering a command for a vfiler through rsh](#)” on page 36.

Entering commands from the vfiler

To enter commands for vfilers from the vfiler, complete the following step.

Step	Action
1	To run in the context of the vfiler, enter the following command: vfiler context <i>vfiler_name</i> Example: <code>vfiler context vfiler1</code>
2	Enter the command you want to run. Example: <code>?</code> This shows all the commands available to you from the context of the vfiler.
3	To return to the context of the hosting filer, enter the command vfiler context vfiler0

Note

Not all Data ONTAP commands can be executed for vfilers. To see what commands are supported, enter `?` from the context of a vfiler.

Entering commands for vfilers from the hosting filer

To enter commands for vfilers from the hosting filer, complete the following step.

Step	Action
1	Enter the following command: <code>vfiler run vfilertemplate command</code> Example: <code>vfiler run vfiler1 useradmin</code>

Note

Not all Data ONTAP commands can be executed for vfilers. To see what commands are supported, enter `?` from the context of a vfiler (see “[Entering commands from the vfiler](#)” on page 34) or enter the following command:

`vfiler run vfilertemplate ?`

Prerequisites for entering commands through rsh from a vfiler client

You must meet the following prerequisites if you want to enter commands for a vfiler through rsh:

- ◆ The rsh protocol must be allowed for the vfiler. By default, rsh is allowed.
- ◆ The rsh protocol must be enabled for the vfiler; that is, the `rsh.enable` option for the vfiler must be set to On. By default, rsh is enabled.
- ◆ You must enter the command on a client that is permitted to have rsh access to the vfiler. In other words, the client must be one of the hosts specified by the `rsh.access` option for the vfiler.

Commands that you can enter through rsh

You can enter the commands listed in the following table through an rsh connection to a vfiler.

<code>?</code>	<code>help</code>	<code>qtree</code>	<code>passwd</code>
<code>exportfs</code>	<code>quota</code>	<code>options</code>	<code>hostname</code>

Note

The `hostname` command is only for displaying, not changing, the name of the vfiler.

Entering a command for a vfiler through rsh

To enter a command for a vfiler through an rsh connection, complete the following step.

Step	Action
1	Enter the following command: rsh vfiler_IP_address command Example: The following command displays all options on the vfiler whose IP address is 123.123.123.1: rsh 123.123.123.1 options

Options that can be set for vfilers

Not all Data ONTAP options are supported on vfilers. You can use the options command for the following options:

- ◆ all cifs options
- ◆ all dns options
- ◆ all nfs options
- ◆ all nis options
- ◆ rsh.access
- ◆ all pcnfsd options
- ◆ wafl.default_nt_user
- ◆ wafl.default_unix_user
- ◆ wafl.nt_admin_priv_map_to_root
- ◆ all wafl.wcc_options

Exception: You can display the cifs.wins_servers option only from the hosting filer.

Path names specified by options

Path names specified by options are complete path names. For example, the path name specified by the cifs.home_dir option should begin with the /vol/vol_name prefix.

Effects of filer reboot on vfilers

Effects of filer reboot on protocols

The disallow or allow state of a vfiler is persistent across reboots. For example, if you disallow NFS for a vfiler and then reboot the filer, NFS remains disallowed for the vfiler after the reboot.

If you disallow CIFS or NFS for a vfiler when the protocol is enabled, rebooting the filer disables the protocol for the vfiler. If you allow the protocol again after a reboot, you must reenable the protocol.

Effect of filer reboot on the vfiler state

Rebooting the filer causes all stopped vfilers to enter the running state. For example, if you stop a vfiler and then reboot the filer, the vfiler starts running again after the reboot.

Understanding volumes and qtrees on vfilers

Effects of taking a vfiler volume offline

If you take a volume offline and that volume is used for vfiler storage:

- ◆ All CIFS shares and NFS exports in the volume are deactivated.
- ◆ If the volume contains the /etc directory for a vfiler, the vfiler enters the stopped state.

After you put the volume back online, the vfiler starts running again.

Changes required after volumes are renamed

After you rename a volume used for vfiler storage, the vfiler administrator should change the path names used in the vfiler's /etc/exports file accordingly. In addition, the vfiler administrator should verify that CIFS shares and quotas are configured properly.

Who can create qtrees

You can create qtrees only if your vfiler owns the volume containing the qtree. That is, as the hosting filer administrator, you can create a qtree in a volume owned by vfiler0, but not in volumes owned by other vfilers. A vfiler administrator can create a qtree in a volume that is a storage unit of the vfiler.

Example of creating a qtree on vfiler0: 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 of creating a qtree on a vfiler: 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
```

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 the qtree or volume. Therefore:

- ◆ If a vfiler owns a volume, you can change the security styles or oplock settings for the volume and all qtrees on the volume from the vfiler.
- ◆ If a vfiler owns qtrees on a volume owned by the hosting filer, you can change the security styles or oplock settings from the vfiler only for the qtrees the vfiler owns.

- ◆ If the hosting filer owns a volume that contains qtrees assigned to vfilers, you can change the security styles or oplock settings from the hosting filer only for the qtrees the hosting filer owns.

Differences in qtree command output

The `qtree` command output is different in the following ways depending on where you enter the command:

- ◆ If you enter the `qtree` command from the hosting filer, the command displays information about all qtrees on the filer, whether the qtrees are owned by the hosting filer or vfilers.
If you enter the `qtree` command from a vfiler, the command displays information about qtrees on that vfiler only. This applies also to `vfiler0`; that is, if you enter the `qtree` command from `vfiler0`, the command displays information about qtrees owned by `vfiler0`.
- ◆ If you enter the `qtree` command without arguments from a vfiler, a qtree that is the destination qtree for SnapMirror is shown as `read_only` in the Status column.

Command for displaying all qtrees and the owning vfilers

Although you can use the `qtree` command from the hosting filer to display all qtrees on the filer, the list of qtrees in the output is not grouped by the vfilers that own the qtrees. The following command displays all qtrees and their owning vfilers:

```
vfiler run * qtree status
```

Networking considerations

Effects of a disabled routed daemon

Enabling the vfiler license causes Data ONTAP to disable the `routed` daemon.

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

Command for changing the routing table in the default IPspace

Because vfilers in the same IPspace share one routing table, you can manipulate the routing table by entering the `route` command from the hosting filer. The `route` command has the following syntax:

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

You can include the command in the filer `/etc/rc` file. In this way, the routes are added automatically after each filer reboot.

Command for changing routing tables in nondefault IPspaces

For more information about manipulating routing tables in nondefault IPspaces, see Chapter 3, “[Using Vfilers in Different IPspaces](#),” on page 45.

About the `/etc/dgateways` file

Only the hosting filer has the `/etc/dgateways` file. Vfilers do not maintain their own `/etc/dgateways` file.

Understanding SnapMirror on the hosting filer

Where to enter the snapmirror command

You can enter the `snapmirror` command from the hosting filer only, not from a vfiler. However, you can perform asynchronous mirroring for any qtree or volume on the filer, whether the qtree or volume is owned by a vfiler or the hosting filer.

Considerations for using the snapmirror command

The procedure for using the `snapmirror` command for data on a vfiler is the same as that for data on a filer.

When typing a filer name in a path name in the `/etc/snapmirror.conf` file, be sure to use the name of the filer, not the name of a vfiler, even though the path name is for a storage unit in a vfiler.

Understanding SNMP when used with vfilers

About SNMP

SNMP is not supported on individual vfilers; it is supported only on the hosting filer. You can enable SNMP on the hosting filer to collect data about vfilers.

About vfiler data in the standard MIB and NetApp custom MIB

In the standard MIB, all vfiler data is global. It pertains to the sum of data from all vfilers on the filer, 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.
- ◆ UDP statistics include data only from sockets in the default vfiler.
- ◆ Quota information is gathered for each volume. If the hosting filer or a vfiler owns a volume with quotas, quota information is provided for the hosting filer or vfiler owning the volume. If a vfiler owns qtrees in a volume that it does not own, no quota information is provided for the vfiler.

In the NetApp custom MIB, a group named vfiler is included. It is for per-vfiler information such as the vfiler license, IP address, protocols allowed, and so on.

Performance monitoring and statistics

Commands for displaying filer statistics

The `sysstat` command displays filer statistics, which are the sum of statistics generated by all vfilers, including `vfiler0`. You cannot use the `sysstat` command to display statistics for particular vfilers.

The `uptime` command displays the uptime for the filer. You cannot use the `uptime` command to display the uptime for particular vfilers.

Commands for displaying NFS and CIFS statistics

You can use the `nfsstat` command and `cifs stat` command to display NFS statistics and CIFS statistics, respectively. These commands can display statistics for the entire filer, specified vfilers, or the hosting filer, depending on the command format, as explained in the following table.

Command entered from the hosting filer		Statistics displayed
For NFS statistics	For CIFS statistics	
<code>nfsstat</code>	<code>cifs stat</code>	Statistics from all vfilers combined
<code>vfiler run vfiler_template nfsstat</code>	<code>vfiler run vfiler_template cifs stat</code>	Statistics from the specified vfilers
<code>vfiler run vfiler0 nfsstat</code>	<code>vfiler run vfiler0 cifs stat</code>	Statistics from the hosting filer

About this chapter

This chapter explains what IPspaces are, why you may need them, and how you create and manage IPspaces and vfilers in those IPspaces.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Understanding IPspaces”](#) on page 46
- ◆ [“Advantages of using VLAN tagging for IPspaces”](#) on page 50
- ◆ [“Requirements for using clusters and IPspaces”](#) on page 51
- ◆ [“Case studies for IPspaces”](#) on page 53
- ◆ [“Managing IPspaces on a filer”](#) on page 54
- ◆ [“Creating a vfiler in a nondefault IPspace”](#) on page 60

Understanding IPspaces

About IPspaces

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

Guidelines for vfiler participation in an IPspace


The following guidelines apply to vfiler participation in an IPspace:

- ◆ An IPspace can contain multiple vfilers; however, a vfiler can belong only to one IPspace.
- ◆ All vfilers in an IPspace must have unique IP addresses; however, multiple vfilers can have the same IP address if they are in distinct IPspaces.

Typical application of IPspaces

Suppose an SSP needs to connect customers of company A and company B to a filer on its premises. The SSP creates two vfilers on the physical filer—one per customer—and provides a dedicated network path from one vfiler to A's network and from the other vfiler to B's network.


This deployment would work if both companies were using nonprivate IP address ranges; however, as shown in the following illustration, this is not the case.



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

- ◆ The two vfilers on the filer at the SSP site have conflicting IP addresses if both companies decide to use the same IP address for their respective vfilers.
- ◆ If the two companies agree on using different IP addresses for their vfilers, problems still 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.
- ◆ Suppose the two companies decide to use mutually exclusive address spaces. (For example, Company A uses 10.0.0.0 with a network mask of 255.127.0.0 and Company B uses 10.128.0.0 with a network mask of 255.127.0.0.) The SSP needs to configure static routes on the filer to appropriately route traffic 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 filer).

To alleviate these problems, two IPspaces are defined on the filer—one per vfiler. 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 vfilers 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, `/etc/hosts.equiv` file, and `/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 vfilers, without conflict.

How interfaces participate in an IPspace

If the filer is licensed for vfilers, 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 vfilers participating in an IPspace share its routing table.

Incoming packets: All packets coming in through an interface are tagged with the IPspace identifier (ID) of the IPspace to which the interface belongs. The IP address of the interface and the IPspace ID are used to identify the vfiler for which the packet is intended.

Outgoing packets: All outgoing traffic uses the IPspace ID of the vfiler that is generating the traffic to determine the routing table to use. Data ONTAP ensures that packets generated by the vfilers of an IPspace are transmitted through the interfaces that belong to that IPspace.

Note

Broadcast packets are restricted to the vfilers within the destination IPspace.

Advantages of using VLAN tagging for IPspaces

Traffic separation

VLAN tagging for IPspaces provides traffic separation from each customer to the filer 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 filer. 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 than interfaces allowed

Using VLAN tagging for IPspaces enables you to set up more IPspaces than there are physical interfaces on the filer.

Dedicating at least one physical interface per IPspace limits the number of IPspaces that can be set up on a filer to the number of physical interfaces available on the filer. VLAN tagging overcomes 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 in an IPspace

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

Requirements for using clusters and IPspaces

- IPspace naming requirement**

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

For example, in a cluster pair of Filer A and Filer B, if IPspaceA is created on Filer A, an IPspaceA must also exist on Filer B.
- IPspace assignment requirement**

The partner interfaces on both partners must be assigned to the same-name IPspaces on their respective filers.

For example, in a cluster pair of Filer A and Filer B, interface e4 of Filer B is the takeover interface of interface e0 of Filer A, and vice versa. If interface e0 belongs to IPspaceA on Filer A, interface e4 must belong to IPspaceA on Filer B.
- Specifying partners in an asymmetric cluster setup**

In an asymmetric cluster setup, vfiler-IPspace configuration on one cluster partner is different from the other; for example, one of the partners might not have any (or the same number of) vfilers configured in a specific IPspace, while the other partner does.

A “standby” cluster configuration is an example of an asymmetric cluster setup in which one of the hosts is connected to minimal storage of its own. This host takes over its partner’s storage and vfilers if the partner fails. In such a configuration, you might not have any or enough storage on this host to configure the same number of vfilers as the primary host. To specify the partner interface when setting up such a cluster configuration, you can use the interface name of the partner instead of the IP address.

For example, two filers, Filer1 and Filer2, are configured as shown in the following table.

Filer	vfiler (associated IP address)	IPspace	Interface name
Filer1	vfiler0 (1.1.1.1)	default	e0
	vfiler1 (2.1.1.1)	ips1	e1
	vfiler2 (3.1.1.1)	ips2	e2

Filer	vfiler (associated IP address)	IPspace	Interface name
Filer2	vfiler0 (1.1.1.2)	default	e0
			e4a
			e4b

To specify partner interfaces on Filer1, complete the following step.

Step	Action
1	Create a host-partner relationship in the following way: <pre>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</pre>

To specify partner interfaces on Filer2, complete the following steps.

Step	Action
1	Create two IPspaces: ips1 and ips2.
2	Assign interface e4a to ips1 and interface e4b to ips2.
3	Create a host-partner relationship on Filer2 in the following way: <pre>ifconfig e0 1.1.1.2 netmask 255.255.255.0 partner e0 ifconfig e4a partner e1 ifconfig e4b partner e2</pre>

Case studies for IPspaces

How SSPs use IPspaces

SSPs typically provide storage services to customers of different companies who have the following specifications:

- ◆ They use the private address space, such as 10.0.0.0, for their intranet.
- ◆ They want a dedicated connection from the storage appliance at the SSP premises to their intranet.
- ◆ They want to ensure that the storage appliance is accessible only to the users they authorize.

Using the IPspace solution available on the filers, the SSP sets up an IPspace per customer, thus creating independent routing domains on a single filer.

Additionally, the SSP uses either physical interfaces or VLAN tags to securely forward traffic to appropriate vfilers. This approach

- ◆ Is cost effective
- ◆ Scales well
- ◆ Provides easy maintenance and monitoring

How an enterprise uses IPspaces

In an enterprise, the IT organization might have to maintain separate filers for each organization of the enterprise. As in the SSP approach, dedicating filers on an organizational basis can not only be expensive but also difficult to maintain. Setting up vfilers on a per-organization basis solves the problem. However, distinct IPspaces for these vfilers are usually not required. Because there is usually a level of trust between different organizations of the enterprise, network access to the various vfilers does not need to be controlled. Therefore, all vfilers can belong to the default IPspace.

IPspaces can be used in an enterprise setup as well if there is a need to route traffic from different vfilers using different default routes.

Managing IPspaces on a filer

Command for managing IPspaces

You manage IPspaces on a filer using the `ipspace` command. This command enables you to create, assign interfaces to, destroy, and list IPspaces on a filer.

For detailed information

For detailed information about how to perform specific tasks using the `ipspace` command, see the following topics:

- ◆ [“Creating an IPspace”](#) on page 55
- ◆ [“Listing IPspaces on a filer”](#) on page 56
- ◆ [“Assigning an interface to an IPspace”](#) on page 57
- ◆ [“Destroying an IPspace”](#) on page 59

Managing IPspaces on a filer

Creating an IPspace

Guidelines for creating an IPspace

- The following guidelines apply to creating an IPspace on a filer:
- ◆ You can create up to 10 IPspaces per filer.
 - ◆ You can use an alphanumeric string, 1 to 31 characters long, as the IPspace name.
 - ◆ All IPspace names you create on a filer must be unique. However, the IPspace names on cluster partners must be the same.

Creating an IPspace

To create an IPspace on a filer, complete the following step.

Step	Action
1	Enter the following command: <pre>ipspace create <i>ipspace</i>name</pre> <i>ipspace</i>name is the IPspace name that you want to create. Example: You can create an ipspace1 IPspace on a filer using the following command: <pre>ipspace create ipspace1</pre>

Managing IPspaces on a filer

Listing IPspaces on a filer

Listing IPspaces To list all IPspaces and the interfaces that are assigned to each IPspace, complete the following step.

Step	Action
1	Enter the following command: ipspace list Example: If you enter the <code>ipspace list</code> command on a filer that has three nondefault IPspaces, you see the following output: <pre>Number of ipspaces configured: 4 default-ipspace (e3 lo) ipspace1 (e2d) ipspace2 (e2c) ipspace3 (e10 e2b sf_vif)</pre> Note_____ In the preceding output, although loopback interface (lo) is shown in the default IPspace, all IPspaces support loopback traffic. The loopback traffic of each IPspace is completely isolated from the other IPspaces. _____

Assigning an interface to an IPspace

Requirement for assigning an interface

The interface you assign from one IPspace to another must not have an IP address configured on it.

Removing an IP address from an interface

To remove an IP address from an interface, complete the following steps.

Step	Action	
1	If...	Then...
	The interface is currently configured with an IP address that belongs to a vfiler in a nondefault IPspace	Remove the IP address configured for the interface from the vfiler. For more information about removing an IP address from a vfiler, see “ Changing resources for vfilers ” on page 22. You are done.
	The interface is in the default IPspace	Go to Step 2 .

Step	Action	
	If...	Then...
	The interface is configured with an IP alias	1. Enter the following command to remove the IP alias: ifconfig interface -alias address <i>interface</i> is the name of the interface. <i>address</i> is the IP address configured for the alias. 2. Go to Step 2.
	The interface is not configured with an IP alias	Go to Step 2 .
2	Enter the following command to remove the IP address: ifconfig interface 0.0.0.0 <i>interface</i> is the name of the interface from which you want to remove the IP address.	

Assigning an interface to an IPspace

To assign an interface to an IPspace, complete the following step.

Step	Action
1	Enter the following command: ipspace assign ipspace-name interface-name <i>ipspace-name</i> is the IPspace name to which the interface will be assigned. <i>interface-name</i> is the name of the interface to be assigned. You can assign only one interface at a time.

Managing IPspaces on a filer

Destroying an IPspace

Destroying
IPspaces

To destroy an IPspace, complete the following steps.

Step	Action
1	Make sure that there are no network interfaces or vfilers associated with the IPspace you want to destroy. Read “Listing IPspaces on a filer” on page 56 to identify any interfaces associated with the IPspace you want to destroy. Read “Managing IPspaces on a filer” on page 57 to assign network interfaces to another IPspace. Read “Displaying vfiler status” on page 32 and “Destroying a vfiler” on page 29 for procedures to identify any vfilers associated with the IPspace and to destroy those vfilers.
2	Enter the following command: ipspace destroy <i>ipspace</i> <i>ipspace</i> is the IPspace name to be destroyed.

Creating a vfiler in a nondefault IPspace

Guidelines for creating a vfiler in a nondefault IPspace

When creating a vfiler in a nondefault IPspace, you need to meet the same prerequisites and follow the same guidelines as you would when creating a vfiler in the default IPspace. For more information about the prerequisites and guidelines, see “[Planning for vfiler creation](#)” on page 12.

Creating a vfiler in a nondefault IPspace

To create a vfiler in a nondefault IPspace, complete the following steps.

Note—
You must follow these steps to configure IP addresses to be used by vfilers in nondefault IPspaces. Make sure you use the `-n` option to 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 `vfiler create` unless you use the `-n` option) does not allow duplicate IP addresses even if they are for interfaces in different IPspaces.

Step	Action
1	To create an IPspace, complete the steps in “ Creating an IPspace ” on page 55.
2	To assign an interface to be used by the vfiler in the newly created IPspace, complete the steps in “ Assigning an interface to an IPspace ” on page 57.
3	To verify that each interface to be used by the vfiler is ready to be configured, complete the steps in “ Ensuring that the network interface is ready ” on page 15.

Step	Action	
4	To create the vfiler, enter the following command: <pre>vfiler create vfiler_name -n -s ipspace -i ip_address [-i ip_address] ... path [path] ...</pre> <i>ipspace</i> is the IPspace in which the vfiler IP addresses reside. <i>ip_address</i> is an IP address. <i>path</i> 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. Example: The following command creates a vfiler named vfiler1 in the IPspace named ipspace1, using the IP address 123.123.123.123 and the /vol/vol1 volume as resources: <pre>vfiler create vfiler1 -n -s ipspace1 -i 123.123.123.123 /vol/vol1</pre>	
5	For each IP address used by the vfiler that is...	Go to...
	An IP alias	Step 6.
	A base IP address	Step 7.
6	Enter the following command to create the IP alias: <pre>ifconfig interface alias address</pre> If the interface to which you are assigning the alias is currently down, go to Step 7. Otherwise go to Step 8.	
7	Use the <code>ifconfig</code> command to configure the interface as Up with the IP address you specified in Step 4. Continue with Step 8.	

Step	Action
8	To modify the routing table for the vfiler, enter the following command: <pre>vfiler run vfiler_name route add destination</pre> Example: The following example adds a route to the routing table used by the vfiler named vfiler1: <pre>vfiler run vfiler1 route add 123.123.123.124</pre>
9	For each interface used by the vfiler, add the following to the /etc/rc file on the hosting filer: ◆ An <code>ifconfig</code> command to configure the interface as Up, with the address specified in Step 4 ◆ The <code>route</code> command you entered in Step 8 This configures the interfaces as Up, and enforces the <code>route</code> commands, across reboots.
10	If the hosting filer is part of a cluster, edit the /etc/rc file on each cluster partner to ensure that each interface the vfiler uses has a partner interface defined for it.

About this chapter

This chapter provides information you need for preparing the vfiler for NFS and CIFS.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“How to specify path names for NFS exports or CIFS shares”](#) on page 64
- ◆ [“Preparing the vfiler for NFS”](#) on page 65
- ◆ [“Preparing the vfiler for CIFS”](#) on page 66

How to specify path names for NFS exports or CIFS shares

Path names for NFS exports and CIFS shares

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

Example of a path for an NFS export

Suppose a vfiler 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, use /vol/vol1/home in the /etc/exports file or in the `exportfs` command.

Example of a path for a CIFS share

Suppose a vfiler named vfiler1 uses the /vol/vol1 volume for storage. To offer the home directory at the root of this volume as the HOME share, specify /vol/vol1/home as the path name for the HOME share.

Preparing the vfiler for NFS

When to use this section

If you use the default form of the `vfiler create` command, you are prompted for NFS setup information as soon as the vfiler has been created. At the end of the setup NFS is running (if NFS is licensed on the physical filer). Use this section only if you need to start NFS manually.

Starting the NFS protocol

To start NFS, complete the following step.

Step	Action
1	Do one of the following: ◆ From the vfiler, enter the following command: <code>nfs on</code> ◆ From the hosting filer, enter the following command: <code>vfiler run vfiler_name nfs on</code> Result: The NFS protocol server starts running on the vfiler.

Exporting all file systems in /etc/exports

To export all paths in the `/etc/exports` file, complete the following step.

Step	Action
1	Do one of the following: ◆ From the vfiler, enter the following command: <code>exportfs -a</code> ◆ From the hosting filer, enter the following command: <code>vfiler run vfiler_name exportfs -a</code> ◆ From a vfiler client that is allowed to connect to the vfiler through rsh, enter the following command: <code>rsh vfiler_name exportfs -a</code>

Preparing the vfiler for CIFS

Different methods of CIFS management

From the hosting filer, you can use the `vfiler run` command format to issue `cifs` commands for vfilers. From the vfilers, you use User Manager or Server Manager to manage user accounts and shares. The `cifs` command is not available through `rsh`.

Tasks that can be performed only on the hosting filer

Server Manager does not perform all the functions of the following `cifs shares` `-add` and `cifs shares -change` commands. You can execute these commands only from the vfiler command line (by means of the `vfiler context` command; see “[Executing commands and setting options for vfilers](#)” on page 34 for more information) or from the hosting filer (by means of the `vfiler run` command):

- ◆ `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 }`

No per-vfiler limit with CIFS

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

About local user accounts for vfilers

From the hosting filer, you can use the `useradmin` command to create local accounts for CIFS users of each vfiler. Each vfiler supports up to 96 local user accounts.

About this chapter

This chapter describes how to restrict and track the disk space and number of files used by a user, group, or qtree on a vfiler.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Understanding quotas”](#) on page 68
- ◆ [“Turning quotas on and off”](#) on page 69
- ◆ [“How to resize quotas”](#) on page 71
- ◆ [“Understanding the /etc/quotas file”](#) on page 72
- ◆ [“Displaying quota status”](#) on page 73
- ◆ [“Understanding quota reports”](#) on page 74
- ◆ [“Messages from exceeded quota thresholds and soft quotas”](#) on page 76

Understanding quotas

Types of quotas supported on vfilers

You can apply the same types of quotas—user, group, and tree quotas—on vfilers as you can on a filer.

Who manages quota specifications?

The vfiler administrator specifies the size of each quota in the vfiler's `/etc/quotas` file. That is, the vfiler administrator tracks and limits the amount of disk space and the number of files each user, group, or qtree uses.

Exception: If a qtree owned by the vfiler resides on a volume owned by the hosting filer, you, the hosting filer administrator, can also specify a quota for the qtree in the hosting filer's `/etc/quotas` file. The following example shows how the tree quota on the hosting filer affects a vfiler qtree:

Suppose the `/vol/vol1/qtree1` qtree is a storage unit of the vfiler, and the `/vol/vol1` volume is owned by the hosting filer.

In the `/etc/quotas` file of the vfiler, the vfiler administrator specifies that this qtree is limited to 20 GB of disk space. In the `/etc/quotas` file of the hosting filer, you can specify the disk space limit for the qtree to be 10 GB. This means that if quotas are turned on from the hosting filer for the `/vol/vol1` volume, the qtree cannot exceed the limit in either `/etc/quotas` file, whichever is lower. In this example, the qtree cannot exceed 10 GB.

Turning quotas on and off

How turning quotas on and off affects vfilers

You can turn quotas on and off on a per-volume basis for a vfiler. After you turn quotas on for a particular volume, Data ONTAP initializes quotas for the storage units residing on the volume, which are owned by the vfiler.

Who can turn quotas on and off?

Hosting filer administrators and vfiler administrators can turn quotas on and off for a volume as long as they own some storage space in the volume. For example, from the hosting filer, you can turn quotas on for the /vol/vol1 volume if vfiler0, which is the hosting filer, owns some storage space in the volume.

You can turn quotas on and off for a volume from a vfiler regardless of how the quota status for that volume is set from other vfilers, including vfiler0. For example, if vfiler1 owns the /vol/vol1/qtree1 qtree and the hosting filer owns the rest of the /vol/vol1 volume, you can turn quotas on and off from vfiler1 for the volume regardless of whether the hosting filer has quotas on or off for the same volume.

Prerequisite for turning quotas on and off

On a filer licensed for vfilers, you must enable quotas for a volume before you can turn quotas on and off for the volume. By default, quotas are enabled on all volumes.

Only hosting filer administrators can enable or disable quotas for a volume.

Effects of disabling quotas

Disabling quotas on a volume has these effects on all vfilers that have storage units in the volume:

- ◆ If quotas are currently off, you or the vfiler administrator cannot turn quotas on for that volume.
- ◆ If quotas are currently on, they are turned off immediately and they cannot be turned on.

Enabling or disabling quotas for a volume

To enable or disable quotas for a volume, complete the following step.

Step	Action
1	Enter the following command: quota allow disallow volume Result: After you enter the <code>quota allow</code> command, you can turn on quotas for the specified volume from a vfiler. After you enter the <code>quota disallow</code> command, vfilers are prevented from turning quotas on for the specified volume. If currently there are vfilers that have quotas on for the volume, quotas are turned off immediately for those vfilers.

Turning quotas on or off from a vfiler

To turn quotas on or off for a volume from a vfiler, complete the following step.

Step	Action	
1	If you manage the vfiler from...	Then...
	The hosting filer	Enter the following command: vfiler run vfilertemplate quota on off volume
	The vfiler	Enter the following command through an rsh connection to the vfiler: quota on off volume

How to resize quotas

Method to resize quotas

You resize quotas on the vfiler in the same way as you resize quotas on a filer.

Which quota records are adjusted?

When you enter the `quota resize` command from a vfiler to resize quotas for a volume, only the quota records for the vfiler are adjusted.

If the vfiler consists of qtrees, you can resize the quotas for this volume without turning quotas off and on for the entire volume.

Understanding the /etc/quotas file

Format of the /etc/quotas file

The format of the /etc/quotas file for a vfiler is the same format as for a filer. All path names in the file are complete path names.

Where user and group quotas take effect

For each user or group quota, you can specify in the Type field the qtree or volume in which the quota takes effect. You can specify only qtrees and volumes that are owned by the vfiler.

If you omit the path name in the Type field, the quota takes effect in the filer's root volume when quotas are turned on for the root volume.

Recommendation: A vfiler administrator should specify a path name in the Type field even though the quota is for a user or group in the filer's root volume. This ensures that the quota remains applicable even after you change the root volume to a different volume.

Displaying quota status

Volumes for which you can display quota status

You can display quota status for any volume on which your vfiler owns storage space.

Displaying quota status

To display quota status, complete the following steps.

Step	Action	
1	If you manage the vfiler from...	Then...
	The hosting filer	Enter the following command: vfiler run vfilertemplate quota
	The vfiler	Enter the following command through an rsh connection to the vfiler: quota
	Result: The command displays quota status information about all volumes in which the vfiler owns storage space. The following messages are sample messages in the command output: vol0: quotas are on. vol1: quotas are off. vol2: quotas are disabled.	

Understanding quota reports

About quota reports for vfilers

A quota report contains information for the vfiler from which you enter the `quota report` command.

Example: The `/vol/vol1` volume contains two qtrees, one owned by vfiler1 and the other owned by vfiler2. When you display the quota report for vfiler1, only the quota information pertaining to vfiler1 is shown. No information about the qtrees owned by vfiler2 or the hosting filer is included in the quota report.

About quota reports for the hosting filer

When you display a quota report for a hosting filer, it shows

- ◆ The disk space and the number of files owned by the quota targets in volumes and qtrees that do not belong to vfilers other than vfiler0.
- ◆ The disk space and the number of files in qtrees owned by other vfilers but that are restricted by tree quotas specified in the filer's `/etc/quotas` file. Information about users and groups in these qtrees, however, is not shown.

How the filer's tree quota affects the vfiler's quota report

The K-Bytes and Files fields include the limits for disk space and number of files in a qtree. However, if a qtree owned by a vfiler is restricted by the tree quota specified in the filer's `/etc/quotas` file, the K-Bytes and Files fields in the vfiler's quota report might not reflect the limits that are in effect for these reasons:

- ◆ If the vfiler does not impose a tree quota on the qtree, the quota report does not include an entry for the qtree, although the qtree is restricted by the tree quota imposed by the filer.
- ◆ If the vfiler imposes a tree quota on the qtree, the K-Bytes and Files fields in the quota report reflect the limits specified in the vfiler's `/etc/quotas` file, although the qtree is restricted by the tree quota imposed by the filer.

Example: The `/vol/vol1/qtree1` qtree is owned by vfiler1. Suppose the filer's `/etc/quotas` file specifies that this qtree can use up to 100 GB of disk space, and the vfiler's `/etc/quotas` file specifies that this qtree can use up to 150 GB of disk space. When you display a quota report for vfiler1, the entry for qtree1 shows that the maximum amount of disk space allowed is 150 GB, even though qtree1 is actually limited by the more restrictive quota, which is 100 GB.

**Command for
displaying vfiler
names in the quota
report**

The `quota report` command syntax for a vfiler is the same as the syntax for a filer, except that it also takes the `-v` argument. If you use the `quota report -v` command, the report displays the vfiler name before the Quota Specifier field.

Messages from exceeded quota thresholds and soft quotas

Where quota warning messages are sent

If a quota threshold or soft quota defined on a vfiler is exceeded, a warning message is logged to the filer console.

Example of the warning message

The following example is a warning message generated by when a quota threshold is exceeded:

```
[vfiler1@quota.softlimit.exceeded:notice]: Threshold exceeded for tree 3 on volume vol1 for vfiler "vfiler1"
```

About this chapter

This chapter covers two distinct but similar sets of tasks: preparing a disaster-recovery, or backup, vfiler (and activating it if a disaster occurs); and moving a vfiler from one physical filer to another (for example if you are replacing an old filer with a new one). Moving a vfiler is sometimes called “migration.”

Topics in this chapter

This chapter covers the following topics:

- ◆ [“Understanding disaster recovery and data migration”](#) on page 78
- ◆ [“How to create and activate a disaster-recovery vfiler”](#) on page 79
- ◆ [“How to move a vfiler”](#) on page 80
- ◆ [“Preparing the destination filer”](#) on page 81
- ◆ [“Storage and network checklists”](#) on page 92
- ◆ [“Creating the disaster-recovery vfiler”](#) on page 95
- ◆ [“Activating the disaster-recovery vfiler”](#) on page 98
- ◆ [“Reactivating the original vfiler”](#) on page 100
- ◆ [“Moving the vfiler”](#) on page 107

Understanding disaster recovery and data migration

Similarities between disaster recovery and data migration

Although the reasons for preparing a disaster-recovery vfiler are quite different from the reasons for moving or migrating a vfiler, many of the tasks involved are similar. In both cases you need to take inventory of the source vfiler and make sure the destination filer has the storage capacity and characteristics, and the network connectivity, to host what will be as nearly as possible an exact copy of the original vfiler. Then in both cases you use a small set of vfiler commands to actually create the new vfiler, populate it with the original vfiler's data, and prepare its connections to the network.

What the procedures do and do not cover: In both cases, the procedures that follow help you not only to move or copy the data (via SnapMirror), but also to configure all of the networking needed to serve the data. But they do not cover adjustments on the client side (for example, mounts onto a Solaris or Windows host). You will need to work closely with the administrators of those systems and with your network administrator to make those adjustments.

Note

A SnapMirror license is required on both the source and destination filers.

Differences between disaster-recovery and data migration

After the new vfiler has been created, the states of the source and destination vfilers depend on how you will use the new vfiler.

- ◆ If you are preparing a disaster-recovery vfiler:
 - ❖ The original (source) vfiler remains intact and continues to serve data to its clients.
 - ❖ The new (destination) vfiler is inactive, and its IP addresses remain unconfigured, until you activate it in the case of a disaster that destroys or incapacitates the original vfiler.
- ◆ If you are moving (migrating) a vfiler:
 - ❖ The original (source) vfiler is destroyed after the migration is complete.
 - ❖ The new (destination) vfiler becomes active and begins serving data as soon as the migration is complete.

How to create and activate a disaster-recovery vfiler

When to use this section

Use the procedures listed in this section if you want to create a backup copy of a vfiler (both the data and the vfiler configuration) for disaster-recovery purposes.

The procedures assume that the original vfiler will remain intact and running unless a disaster puts it out of action. If you want to destroy the original vfiler after the new vfiler is up and running, follow the procedures listed under “[How to move a vfiler](#)” on page 80 instead.

What to do

Before a disaster: To create a disaster-recovery vfiler, follow this process.

Stage	Description
1	Qualify and prepare the destination filer. Perform the “ Storage checks and actions ” on page 81 and the “ Network checks and actions ” on page 84.
2	Create the disaster-recovery vfiler. See “ Creating the disaster-recovery vfiler ” on page 95.

After a disaster: After a disaster has occurred, and the original vfiler is destroyed, incapacitated, or unreachable, follow this process.

Stage	Description
1	Activate the disaster-recovery vfiler. See “ Activating the disaster-recovery vfiler ” on page 98.
2	To reactivate the original vfiler after the damage caused by the disaster has been repaired, follow directions under “ Reactivating the original vfiler ” on page 100.

How to move a vfiler

When to use this section

Use this section if you want to move a vfiler (that is, the data and the vfiler configuration) from one NetApp appliance to another, for example from a filer that is overloaded to one that has spare capacity, or from an old filer to a new one. This is sometimes called “migrating” a vfiler.

The procedures assume that you want to destroy the original vfiler once the new vfiler is up and running. If you want to leave the original vfiler intact and create a backup copy of it for disaster-recovery purposes, use the procedures listed under “[How to create and activate a disaster-recovery vfiler](#)” on page 79 instead.

What to do

To move a vfiler, follow this process.

Stage	Description
1	Qualify and prepare the destination filer. Perform the “ Storage checks and actions ” on page 81 and the “ Network checks and actions ” on page 84.
2	Move the vfiler. See “ Moving the vfiler ” on page 107.

Preparing the destination filer

When to use this section Perform the “[Storage checks and actions](#)” that follow and the “[Network checks and actions](#)” on page 84 whenever you intend to replicate a vfiler, whether you are moving the vfiler or preparing a disaster-recovery vfiler.

Storage checks and actions Perform the following checks, and the accompanying actions if necessary, before moving (migrating) a vfiler or setting up a disaster-recovery vfiler. Photocopy the worksheet that follows this section and fill out the storage items as you go along.

Step	Action
1	Check that the destination filer has enough storage space to hold the source vfiler’s volumes. On the source filer, use the <code>vfiler status -r</code> command to see what volumes the vfiler is using. Then use the <code>df</code> command on each of those volumes to check how much space is being used. The destination volumes must have at least that much free space; run <code>df</code> on the destination filer to check. Fill in the <code>df</code> results on the worksheet. Note_____ If the source and destination filers use different-sized disks and have different block sizes, you will need to adjust the <code>df</code> numbers accordingly. _____

Step	Action	
2	If...	Then...
	There is enough space on the destination volumes	Continue with Step 3 .
	There is not enough space	a. If necessary, install new disk shelves. b. Use the <code>vol add</code> command to add new disks to the destination volumes.
3	Check to see if the destination filer has the same volume structure as the source. The volumes to be used by the destination vfiler must have the same path names as those used by the source vfiler.	
4	If...	Then...
	The destination filer has volumes whose path names match those used by the source vfiler	Continue with Step 5 .
	The destination filer does not have matching volume path names	Use the <code>vol create</code> command on the destination filer to create volumes whose names match those being used by the source vfiler, or <code>vol rename</code> to rename a volume.
5	Check to see if there are qtrees in the destination volumes whose names match those of qtrees in the source volumes. The migration process uses SnapMirror, which will replicate qtree names from the source to the destination volume, so these names should not already exist on the destination.	

Step	Action	
6	If...	Then...
	There are qtrees in the destination volumes whose names match those of qtrees in the source volumes	Delete the matching qtrees in the destination volumes. To remove a qtree, delete it from a client as you would a directory or folder.
	There are no matching qtree names in the destination volumes	Continue with Step 7 .
7	Check to see if quotas are being enforced from the hosting filer. Quotas enforced from the vfiler will be copied to the new vfiler, but quotas enforced from the hosting filer will not. To see where quotas are being enforced from, enter the command quota report from the hosting filer.	
8	If...	Then...
	The output of the <code>quota report</code> command shows that quotas for qtrees used by the vfiler are being enforced from the hosting filer	a. Print out a copy of the filer's quota file: <code>/vol/vol0/etc/quotas</code>. b. Copy the relevant entries into the destination filer's <code>/vol/vol0/etc/quotas</code> file.
	No quotas for qtrees used by the vfiler are being enforced from the hosting filer	Continue with Step 9 .
9	Continue with “ Network checks and actions ” on page 84.	

Network checks and actions

Perform the following checks, and the accompanying actions if necessary, before moving (migrating) a vfiler or setting up a disaster-recovery vfiler. Photocopy the worksheet that follows this section and fill out the network items as you go along.

Step	Action
1	Check to see if the destination vfiler can take over the source vfiler's IP addresses. Use the command <code>ifconfig -a</code> on the source vfiler and on the destination vfiler to display information about all their network interfaces. You can reuse the source IP addresses and aliases on the destination vfiler if the destination vfiler will be on the same subnet as the source vfiler.

Step	Action	
2	If...	Then...
	The source and destination filers are on the same subnet Note_____ This is the default case assumed by the <code>vfiler migrate</code> command.	Continue with Step 3 .
	They are on different subnets	a. Obtain as many new IP addresses for the destination vfiler as are in use on the source vfiler. Note_____ You may need to replicate subnet-separation arrangements that exist on the source vfiler. For example the source vfiler may use one IP address for a service network and another for an administration network. b. Make a note of the new IP addresses on the worksheet. The <code>vfiler dr configure</code> command will prompt you for these addresses; see “Creating the disaster-recovery vfiler” on page 95. The <code>vfiler migrate</code> command will also prompt you for these addresses, but you may then need to run <code>setup</code> on the destination vfiler. See “If you have moved the vfiler to a different subnet or Windows domain” on page 109.

Step	Action	
3	Check whether the source vfiler uses the default IPspace. Use the command <code>ipSPACE list</code> on the source vfiler to display information about IPspaces and the interfaces assigned to them. For more information about IPspaces, see Chapter 3 on page 45.	
4	If...	Then...
	<code>ipSPACE list</code> reports default-ipSPACE	Continue with Step 5 .
	<code>ipSPACE list</code> reports something other than default-ipSPACE	a. Use the command <code>ipSPACE create ipSPACENAME</code> to create a corresponding IPspace with the same name on the destination filer. b. Use the command <code>ipSPACE assign ipSPACENAME interfacename</code> to assign physical interfaces to the IPspace. These interfaces must all be attached to the same physical network.
5	Check to see if the destination vfiler will have access to the same NIS servers as the source. Note _____ You can skip this check if the source and destination vfilers are on the same subnet. Use the command <code>nis info</code> on the source vfiler to see what NIS servers are available to it. (The <code>ypwhich</code> command shows which server the filer is currently bound to.)	

Step	Action	
6	If...	Then...
	The destination vfiler can use the same NIS servers as the source vfiler Note This is the default case assumed by the <code>vfiler migrate</code> command.	Continue with Step 7 .
	The destination vfiler cannot use same NIS servers	a. Find out which NIS servers are available to the destination filer. b. Make a note of the IP addresses of those servers on the worksheet. The <code>vfiler dr configure</code> command will prompt you for these addresses; see “Creating the disaster-recovery vfiler” on page 95. The <code>vfiler migrate</code> command will <i>not</i> prompt you for these addresses. If you move a vfiler to a different subnet, you may need to run <code>setup</code> on the destination vfiler. See “If you have moved the vfiler to a different subnet or Windows domain” on page 109.

Step	Action
7	Check to see if the destination vfiler will have access to the same DNS servers as the source. Note_____ You can skip this check if the source and destination vfilers are on the same subnet. Use the command <code>dns info</code> on the source vfiler to see what DNS servers are available to it.

Step	Action	
8	If...	Then...
	The destination vfiler can use the same DNS servers as the source vfiler Note This is the default case assumed by the <code>vfiler migrate</code> command.	Continue with Step 9 .
	The destination vfiler cannot use the same DNS servers	a. Find out which DNS servers are available to the destination filer. b. Make a note of the IP addresses of those servers on the worksheet. The <code>vfiler dr configure</code> command will prompt you for these addresses; see “Creating the disaster-recovery vfiler” on page 95. The <code>vfiler migrate</code> command will <i>not</i> prompt you for these addresses. If you move a vfiler to a different subnet, you may need to run <code>setup</code> on the destination vfiler. See “If you have moved the vfiler to a different subnet or Windows domain” on page 109.
9	Check to see if the destination vfiler will have access to the same WINS servers and the same Windows security network as the source.	

Step	Action	
10	If...	Then...
	The destination vfiler can use the same WINS servers and Windows security network as the source vfiler	Continue with Step 11 .
	The destination vfiler cannot use the same WINS servers and Windows security network	a. Find out the name and type (Windows NT 4 or Windows 2000) of the domain the destination vfiler will be in. b. Make a note of this information on the worksheet. When you activate the disaster-recovery vfiler, you will need to configure it into the new domain. See “Activating the disaster-recovery vfiler” on page 98. If you move a vfiler into a different domain, you will need to configure it into the new domain. See “If you have moved the vfiler to a different subnet or Windows domain” on page 109.
11	Check to see if the destination vfiler can use the same trusted host for vfiler administration as the source vfiler.	

Step	Action	
12	If...	Then...
	The destination vfiler can use the same trusted host as the source vfiler	You are done.
	The destination vfiler cannot use the same trusted host	a. Find out the name of the new trusted host. b. Make a note of this information on the worksheet. You will need to configure the new trusted-host information after configuring the disaster-recovery vfiler, or after moving the vfiler. See “Creating the disaster-recovery vfiler” on page 95 or “If you have moved the vfiler to a different subnet or Windows domain” on page 109.

Storage and network checklists

How to use these checklists

It's a good idea to photocopy the checklists on this page and the next page and fill them out as you go through the checks earlier in this section. (By photocopying them you will keep the originals clean for next time.)

STORAGE CHECKLIST:

How much disk space is used and available on the volumes? (See [Step 1](#) of the “[Storage checks and actions](#)” on page 81.)

df of source filer's volumes: _____

df of destination filer's volumes: _____

Have you added a sufficient number of disks to the destination volumes if necessary? (See [Step 2](#) of the “[Storage checks and actions](#)”.) _____

Do the pathnames of the source and destination volumes match? (See [Step 4](#) of the “[Storage checks and actions](#)”.) _____

If you are managing qtree-based vfilers, is each destination volume free of qtree names that duplicate source qtree names? (See [Step 6](#) of the “[Storage checks and actions](#)”.) _____

Have you copied filer-based quota information from the source to the destination filer's /etc/quotas file if necessary? (See [Step 8](#) of the “[Storage checks and actions](#)”.) _____

NETWORKING CHECKLIST:

Are there enough IP addresses available for the vfiler on the destination network?
(See [Step 2](#) of the “[Network checks and actions](#)” on page 84.)

old interface:_____ new interface:_____

old interface:_____ new interface:_____

old interface:_____ new interface:_____

old interface:_____ new interface:_____

Check syntax carefully; interface names are case-sensitive.

Have you gathered all the authority servers? (See [Step 5](#) through [Step 10](#) of the
“[Network checks and actions](#)”).)_____

old NIS domain:_____ new NIS domain:_____

old NIS IP addr:_____ new NIS IP addr:_____

old NIS IP addr:_____ new NIS IP addr:_____

old DNS domain:_____ new DNS domain:_____

old DNS IP addr:_____ new DNS IP addr:_____

old DNS IP addr:_____ new DNS IP addr:_____

old DNS IP addr:_____ new DNS IP addr:_____

old WINS IP addr:_____ new WINS IP addr:_____

old WINS IP addr:_____ new WINS IP addr:_____

old NT domain type: NT4 W2K

old domain name (FQDN and NetBIOS):_____

new NT domain type: NT4 W2K

new domain name (FQDN and NetBIOS):_____

Can you use the same trusted host for vfiler administration? (See [Step 12](#) of the
“[Network checks and actions](#)”).)_____

old trusted host name:_____ new trusted host name:_____

What to do next

When you have finished the preparation, you are ready to create the destination vfiler.

- ◆ If you are preparing a disaster-recovery vfiler, follow directions under [“Creating the disaster-recovery vfiler”](#) on page 95.

In this case the original vfiler remains intact and running.

- ◆ If you are moving a vfiler from one physical filer to another, follow directions under [“Moving the vfiler”](#) on page 107.

In this case, the original vfiler is destroyed after it has been replicated on the destination filer.

Creating the disaster-recovery vfiler

Before you start

Before you begin creating the disaster-recovery vfiler, make sure the following are true:

- ◆ You have performed all the preparation steps listed under “[Storage checks and actions](#)” on page 81 and “[Network checks and actions](#)” on page 84.
- ◆ SnapMirror is licensed and enabled on both the source and the destination filer.
- ◆ The source and destination filers can communicate with each other over the network (for example by means of DNS lookup or entries in /etc/hosts).
- ◆ The source filer can communicate with the destination filer via the `rsh` command.
- ◆ If any of the source qtrees or volumes contain virtual disks, the destination filer is running a version of Data ONTAP that supports virtual disks.
- ◆ You know the source filer’s administrative ID and password.

Creating the vfiler

To create the disaster-recovery vfiler, perform the following steps.

Step	Action
1	<i>On the destination filer</i> , enter the command <code>vfiler dr configure source_vfiler @ source_filer</code>
2	Respond to the login prompt with a valid administrative ID and password for the source filer.
3	Respond to the IP address and binding prompts, using the addresses you entered on the “new interface” lines on the worksheet.
4	Respond to the NIS and DNS server prompts, using the addresses you entered under “New NIS addr” and “New DNS addr” on the worksheet.

Step	Action
5	Monitor progress by means of the following command: <pre>vfiler dr status source_vfiler @ source_filer</pre> Note_____ The <code>vfiler dr configure</code> command may take some time to complete, especially if a source qtree has many millions of inodes. When the <code>vfiler dr status</code> command reports all the storage units of the source vfiler as <code>snapmirrored</code>, proceed to the next step. Note_____ The disaster-recovery vfiler now exists, but has not been started. See “What the vfiler dr configure command does” on page 97 for more information.
6	If you copied quota information into the destination filer’s <code>/etc/quotas</code> (see Step 8 of “Preparing the destination filer” earlier in this chapter) activate the quotas on that filer. For each volume you are activating quotas on, use the following command: <pre>quota on volume</pre>
7	Edit the disaster-recovery vfiler’s <code>/etc/hosts.equiv</code> file, adding the name of the trusted host for administering the disaster-recovery vfiler. (This is the host name you entered on the worksheet as “new trusted host.”) Note_____ If the trusted host is a Windows system, or if it is a UNIX system and the trusted user is not the root user, you need to add the user name as well; for example, <code>adminhost joe_smith</code>
8	Add the path to the root volume and the name of the trusted host to the disaster-recover vfiler’s <code>/etc/exports</code> file. Example: <code>/vol/vol0 access=adminhost, root=adminhost</code>

What the `vfiler dr configure` command does

The `vfiler dr configure` command does the following:

- ◆ Checks that the destination filer is capable of receiving the source data.
- ◆ Configures and runs SnapMirror to copy the data from the source to the destination vfiler.
- ◆ Saves the IP configuration information you supply (see [Step 3](#) under “[Creating the vfiler](#)” on page 95).
- ◆ Saves the NIS and DNS server information you supply (see [Step 4](#) under “[Creating the vfiler](#)” on page 95).
- ◆ Saves the quota information from the source vfiler’s `/etc/quotas` file.
- ◆ Creates a special type of vfiler, known as a DR backup vfiler, on the destination filer.

This vfiler is stopped and cannot be started except as described under “[Activating the disaster-recovery vfiler](#)” on page 98. Before activation, the vfiler will respond only to the `vfiler status`, `vfiler dr delete`, and `vfiler dr configure` commands. You cannot use `ifconfig` to configure its addresses.

Activating the disaster-recovery vfiler

To activate the vfiler in the event of a disaster, perform the following steps.

Step	Action
1	On the destination filer, enter the following command: <code>vfiler dr activate source_vfiler @ source_filer</code>
2	To change the name of the Windows domain controller, use the <code>cifs prefdc</code> command.
3	To change the Windows WINS server, run <code>cifs setup</code> . Note _____ If the Windows domain has changed, you may need to change the permissions on the Windows data files to allow your users the same access they had in the old domain. _____
4	Make adjustments on the clients. These include remounting volumes and qtrees.

What activation does

The `vfiler dr activate` command does the following:

- ◆ Breaks the SnapMirror relationships between the source and destination filers.
- ◆ Converts the vfiler into an ordinary vfiler that can be started and that responds to all the commands vfilers support.
- ◆ Configures IP bindings on the destination vfiler according to the information you provided to 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 is removed from the destination `/etc/rc`.)
- ◆ 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.

Deleting the disaster-recovery vfiler

If you need to delete the disaster-recovery vfiler, you can do so, at any time after you have verified that the configuration process is complete, by running the following command on the destination filer:

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

This will also remove the SnapMirror relationships, and any other configuration information related to the disaster-recovery vfiler, from the source vfiler.

Reactivating the original vfiler

When to use this section

Use this section when you want to bring the original vfiler back online after repairing or replacing the original filer following a disaster.

How you do this depends on whether you have recovered the original filer, or are bringing up a new filer as a replacement.

- ◆ If you have recovered the original filer, follow instructions under [“Ways to reactivate the original vfiler”](#) immediately following.
- ◆ If you are bringing up a replacement filer, follow directions under [“Recreating the vfiler on a replacement filer”](#) on page 104.

Ways to reactivate the original vfiler

There are two ways to reactivate the original vfiler:

- ◆ If the filer was not damaged but suffered only a temporary failure, and you are comfortable using SnapMirror commands, follow the procedure [“Reactivating the vfiler using SnapMirror commands”](#).
This method is a little more complicated but involves less data movement.
- ◆ If the filer was severely damaged, or you are not comfortable using SnapMirror commands, follow the procedure [“Reactivating the vfiler using vfiler dr commands”](#) on page 103.
This method is less complicated but involves more data movement.

Reactivating the vfiler using SnapMirror commands: To reactivate the original vfiler on the original filer, perform the following steps.

Step	Action
1	Boot the original filer and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress. Note_____ If you don’t press the Del or Esc key in time, you can press Ctrl-C when prompted later during the boot; then choose option 5 (maintenance mode); then enter halt.

Step	Action
2	At the OK prompt, set the <code>no-vfiler-ips?</code> variable as follows: <pre>setenv no-vfiler-ips? true</pre> This ensures that the filer does not try to bind IP addresses already being used by the disaster-recovery vfiler. When the filer boots, the original vfiler will come up and report as running, but will not accept any read or write requests because all its interfaces will be unconfigured.
3	Resynchronize the mirrored volumes and qtrees. For each volume and qtree owned by the original vfiler, enter the following command on the <i>original filer</i>: <pre>snapmirror resync -S disaster_recovery_filer:/pathname production_filer:/pathname</pre> Example: <pre>snapmirror resync -S drfiler:/vol/vfiler1/mtree1 prfiler:/vol/vfiler1/qtree</pre> Troubleshooting: If the <code>snapmirror resync</code> command fails, telling you that there are no matching snapshots, you have probably accidentally deleted the snapshots that SnapMirror depends on. In this case you need to initialize SnapMirror using the <code>snapmirror initialize</code> command (see the <code>na_snapmirror</code> man page for details). Then continue with the next step of this procedure, but skip Step 7 when you get there.
4	Find out if the <code>snapmirror.access</code> option on the <i>disaster-recovery filer</i> is set to <code>legacy</code>. Enter the following command on the disaster-recovery filer: <pre>options snapmirror</pre>

Step	Action	
5	If...	Then...
	The options <code>snapmirror</code> command returns <code>snapmirror.access legacy</code>	Edit the file <code>/etc/snapmirror.allow</code> and add the host name of the <i>original filer</i> if it is not already there.
	The options <code>snapmirror</code> command returns a list of host names, and the name of the original filer is not in the list	Use the options <code>snapmirror</code> command to add the host name of the <i>original filer</i> . Example: <code>options snapmirror.access host=fridge,toaster,prfiler</code>
6	Run <i>setup on the disaster-recovery vfiler</i> and unconfigure its IP addresses.	
7	Update the data on the original vfiler. For each volume and qtree owned by the original vfiler, enter the following command <i>on the original filer</i>: <code>snapmirror update -S disaster_recovery_filer:/pathname production_filer:/pathname</code> Example: <code>snapmirror update -S drfiler:/vol/vfiler1/mtree1 prfiler:/vol/vfiler1/mtree1</code>	
8	Stop SnapMirror transfers to the disaster-recovery vfiler. For each volume and qtree owned by the original vfiler, enter the following command <i>on the original filer</i>: <code>snapmirror quiesce pathname</code> Example: <code>snapmirror quiesce /vol/vfiler1/mtree1</code> Note_____ This operation can take a long time. Use Ctrl-C to interrupt it if you need to. _____	

Step	Action
9	Check that all the paths are quiesced: <pre>snapmirror status</pre> The status column in the output should show each path as <code>Quiesced</code>.
10	Break the SnapMirror relationship. For each volume and qtrees owned by the original vfiler, enter the following command <i>on the original filer</i>: <pre>snapmirror break pathname</pre> Example: <pre>snapmirror break /vol/vfiler1/qtrees1</pre>
11	<i>On the original vfiler</i>, run setup to configure the vfiler's IP addresses and configure the NIS and DNS servers.
12	Destroy the disaster-recovery vfiler. <i>On the disaster-recovery filer</i>, enter the following command: <pre>vfiler destroy vfilername</pre> <i>vfilername</i> is the name of the disaster-recovery vfiler.
13	Recreate the disaster-recovery vfiler. From the <i>disaster-recovery filer</i>, perform the procedure “Creating the disaster-recovery vfiler” on page 95.

Reactivating the vfiler using vfiler dr commands: To reactivate the original vfiler on the original filer, perform the following steps.

Step	Action
1	Boot the original filer and interrupt the boot process by pressing the Del or Esc key while the memory self-test is in progress. Note_____ If you don't press the Del or Esc key in time, you can press Ctrl-C when prompted later during the boot; then choose option 5 (maintenance mode); then enter halt. _____

Step	Action
2	At the OK prompt, set the <code>no-vfiler-ips?</code> variable as follows: setenv no-vfiler-ips? true This ensures that the filer does not try to bind IP addresses already being used by the disaster-recovery vfiler.
3	At the OK prompt, enter boot.
4	Destroy the original vfiler. Enter the following command <i>on the original filer</i> : vfiler destroy vfilername <i>vfilername</i> is the name of the original vfiler.
5	Follow the procedure for “ Creating a replacement vfiler on the original filer or its replacement ” on page 104.

Recreating the vfiler on a replacement filer

To recreate the original vfiler on a replacement filer, perform the following steps.

Step	Action
1	Boot the replacement filer.
2	Follow the procedure for “ Creating a replacement vfiler on the original filer or its replacement ” on page 104.

Creating a replacement vfiler on the original filer or its replacement

Follow these steps to create a replacement vfiler on the original (production) filer or its replacement. Perform this procedure only after completing the preparatory steps under “[Recreating the vfiler on a replacement filer](#)” on page 104, or under “[Reactivating the vfiler using vfiler dr commands](#)” on page 103.

Step	Action
1	Prepare the new vfiler: Perform the preparation steps under “Preparing the destination filer” on page 81. Note Remember that you are performing these checks to inventory the resources of the <i>original filer or its replacement</i>.
2	Create the new vfiler, following directions under “Creating the vfiler” on page 95. Note Perform those steps on the <i>original filer or its replacement</i>. After performing those steps, you have a production vfiler and a backup vfiler, as you did before the disaster, but their roles are reversed: the vfiler you have just created is now the backup, and what was originally the disaster-recovery vfiler is still managing the live data. This may or may not suit your needs. ◆ If you want to put things back the way they were before the disaster, continue with the next step. ◆ If you don’t care about the roles of the two vfilers, you are done.
3	Run <i>setup on the disaster-recovery vfiler</i> and unconfigure its IP addresses.
4	Update the data on the new vfiler. For each volume and qtree owned by the new vfiler, enter the following command <i>on the original filer or its replacement</i>: <pre>snapmirror update -S <i>disaster_recovery_filer:/pathname</i> <i>production_filer:/pathname</i></pre> Example: <pre>snapmirror update -S drfiler:/vol/vfiler1/mtree1 prfiler:/vol/vfiler1/mtree1</pre>

5	Activate the new vfiler, following directions under “Activating the disaster-recovery vfiler” on page 98. Note_____ Perform these steps on the <i>original filer or its replacement</i>. _____
6	Destroy the disaster-recovery vfiler by entering the following command <i>on the disaster-recovery hosting filer</i>: vfiler destroy vfilername <i>vfilername</i> is the name of the original vfiler
7	Recreate the disaster-recovery vfiler. From the <i>original filer or its replacement</i>, perform the procedure “Creating the disaster-recovery vfiler” on page 95.

Moving the vfiler

Before you start

Before you begin the migration, make sure the following are true:

- ◆ You have performed all the preparation steps listed under “[Storage checks and actions](#)” on page 81 and “[Network checks and actions](#)” on page 84.
- ◆ SnapMirror is licensed and enabled on both the source and the destination filer.
- ◆ The source and destination filers are on the same subnet.

Note

If the filers are not on the same subnet, you will need to run `setup` (and `cifs setup` if necessary) on the new vfiler after the migration, to configure authority servers as specified on your worksheet. See “[If you have moved the vfiler to a different subnet or Windows domain](#)” on page 109 for details.

- ◆ The source and destination filers can communicate with each other over the network (for example by means of DNS lookup or entries in `/etc/hosts`).
- ◆ The source filer can communicate with the destination filer via the `rsh` command.
- ◆ If any of the source qtrees or volumes contain virtual disks, the destination filer is running a version of Data ONTAP that supports virtual disks.
- ◆ You know the source filer’s administrative ID and password.

How moving a vfiler affects clients

If you move a vfiler to another filer on the same subnet, you will see the following effects on clients.

- ◆ NFS volume-level mounts move transparently.
- ◆ NFS qtree-level mounts will need to be remounted.
- ◆ CIFS clients will experience the equivalent of a reboot.

For information on moving a vfiler to a filer on a different subnet, see “[If you have moved the vfiler to a different subnet or Windows domain](#)” on page 109.

Caution

The procedure that follows destroys the original vfiler after replicating it on the destination filer. If you want to leave the original vfiler intact and create a backup copy of it for disaster-recovery purposes, use the procedures listed under “[How to create and activate a disaster-recovery vfiler](#)” on page 79 instead.

Moving the vfiler

To move the vfiler, perform the following steps.

Step	Action
1	<i>On the destination filer, enter the following command:</i> <pre>vfiler migrate start source_vfiler @ source_filer</pre> Note_____ This procedure locks up the console for the minimum amount of time. If this is not a concern for you (if you want to let the job run overnight, for example) you can perform the migration with a single command: <pre>vfiler migrate source_vfiler @ source_filer</pre> In that case, skip Step 4 and Step 5. _____
2	Respond to the login prompt with a valid administrative ID and password for the source filer.
3	Respond to the IP address and binding prompts, using the addresses you entered on the “new interface” lines on the worksheet.
4	Monitor the progress of the migration by means of the following command: <pre>vfiler migrate status source_vfiler @ source_filer</pre> Note_____ The <code>vfiler dr migrate</code> command may take some time to complete, especially if a source qtree has many millions of inodes. _____

Step	Action
5	When the <code>status</code> command reports that SnapMirror has replicated all the storage units of the source vfiler, complete the migration by entering the following command: <pre>vfiler migrate complete source_vfiler @ source_filer</pre> Note_____ If instead you want to cancel the migration, you can do so now by issuing the following command on the destination filer: <pre>vfiler migrate cancel source_vfiler @ source_filer</pre> This will destroy the destination vfiler and remove the SnapMirror and other migration-related configuration information from the source vfiler. Note_____ If the <code>vfiler migrate complete</code> command reports that the process cannot complete, wait a few minutes and then enter the same <code>vfiler migrate complete</code> command again.
6	If you copied quota information into the destination filer's <code>/etc/quotas</code> file (see Step 8 of “Preparing the destination filer” earlier in this chapter) activate the quotas on that filer. For each volume you are activating quotas on, use the following command: <pre>quota on volume</pre>
7	Make adjustments on the clients. These will include remounting qtrees, but volume-level mounts should be intact.

If you have moved the vfiler to a different subnet or Windows

domain: If you have moved the vfiler to a different subnet, you may need to configure the network servers and make adjustments on the clients. If you have moved the vfiler to a different Windows domain, you may also need to modify data-file security attributes.

Perform the following steps *on the destination vfiler*.

Step	Action
1	To configure NIS and DNS servers, run <code>setup</code> .
2	To change the name of the Windows domain controller, use the <code>cifs prefdc</code> command.
3	To change the Windows WINS server, run <code>cifs setup</code>. Note If the Windows domain has changed, you may need 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, do the following: a. Edit the vfiler's <code>/etc/hosts.equiv</code> file, adding the name of the trusted host for administering the vfiler. (This is the hostname you entered on the worksheet as "new trusted host".) Note If the trusted host is a Windows system, or if it is a UNIX system and the trusted user is not the root user, you need to add the user name as well; for example <code>adminhost joe_smith</code> b. Add the path to the root volume and the name of the trusted host to the vfiler's <code>/etc/exports</code> file. Example: <code>/vol/vol0 access=adminhost, root=adminhost</code>
5	Make adjustments on the clients. These will include remounting volumes and qtrees.

What the vfiler migrate commands do

The `vfiler migrate start` command does the following:

- ◆ Checks that the destination filer is capable of receiving the source data.
- ◆ Configures and runs SnapMirror to copy the data from the source to the destination vfiler.
- ◆ Saves the IP configuration and binding information you supply (see [Step 3](#) under "[Creating the vfiler](#)" on page 95).

- ◆ Saves the quota information from the source vfiler's `/etc/quotas` file.

The `vfiler migrate complete` command does the following:

- ◆ Stops the source vfiler.
- ◆ Updates the data on the destination vfiler.
- ◆ Breaks the SnapMirror relationships.
- ◆ Configures IP bindings on the destination vfiler according to the information you provided to the `vfiler migrate start` command, adding the destination IP information to the destination `/etc/rc` file. (Any IP information that pertains only to the source vfiler is removed from the destination `/etc/rc`.)
- ◆ Configures any quota information saved by the `vfiler migrate start` command.
- ◆ Destroys the source vfiler.

About this chapter This chapter discusses how to perform virus scanning for vfilers that run the CIFS protocol.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Understanding virus scanning for vfilers](#)” on page 114
- ◆ “[Configuring virus scanning for vfilers](#)” on page 115

Understanding virus scanning for vfilers

Who can configure virus scanning

You can specify how virus scanning works on files owned by the hosting filer and files owned by vfilers. Vfiler administrators, however, cannot configure virus scanning for vfilers.

Filers with which virus scanners can be registered

Virus scanners can be registered with the hosting filer or any vfiler. You can determine whether files on a vfiler are scanned by virus scanners registered with the hosting filer or the vfiler.

Virus scanning for vfiler0

Virus scanning works for vfiler0 in the same way as it does for a filer not licensed for vfilers. See the *Data Protection Guide* for how to configure virus scanning for a filer.

Requirements for virus scanning for a vfiler other than vfiler0

These requirements must be met before you can scan files on a vfiler other than vfiler0:

- ◆ Virus scanning must be enabled for the vfiler.
- ◆ A virus scanner must be available for the vfiler because one of the following requirements is met:
 - ❖ A virus scanner must be registered with the vfiler.
 - ❖ A virus scanner must be registered with the hosting filer and the vfiler must be allowed to use it.

Effect of virus scanner availability on CIFS access

CIFS clients of the vfiler are not allowed to open any files on the vfiler if all the following conditions apply:

- ◆ Virus scanning is enabled.
- ◆ No virus scanner is available for the vfiler.
- ◆ For the vfiler, the `mandatory_scan` option for `vscan` is set to On.

Configuring virus scanning for vfilers

Configuring virus scanning

To specify the virus scanner for scanning files on a vfiler, complete the following steps.

Step	Action
1	Enter the following command to specify the virus scanner for scanning files on a vfiler: <pre>vfiler run vfilertemplate vscan options use_host_scanners { on off }</pre> Set the <code>use_host_scanners</code> option to On if you allow the vfiler to use the virus scanner registered with the hosting filer. The vfiler uses the virus scanner registered with the hosting filer only when the virus scanner registered with the vfiler is unavailable. Set the <code>use_host_scanners</code> option to Off if you do not want to allow the vfiler to use the virus scanner registered with the hosting filer. Note_____ The <code>use_host_scanners</code> option is applicable only to a vfiler you created. You cannot set it on vfiler0 or a filer. _____
2	Enter the following command to enable or disable virus scanning: <pre>vfiler run vfilertemplate vscan { on off }</pre>

Index

Symbols

- /etc directory
 - configuration files in 14
 - effects of vfiler creation on 12
- /etc/dgateways file 40
- /etc/exports file
 - changing path names in after renaming volume 38
 - exporting all file systems in 65
- /etc/messages file 33
- /etc/quotas file
 - format of 72
 - language for encoding 13
- /etc/snapmirror.conf file 41
- /etc/usermap.cfg file, language for encoding 13

A

- activating disaster-recovery vfiler 98
- adding and removing vfiler resources, effects of 22
- administering resources from the hosting filer 21
- after a disaster
 - summary of tasks 79
- allowing or disallowing
 - CIFS or NFS, effects of 30
 - protocols, steps for 31
 - rsh, effects of 30
- assigning interfaces to IPspaces 57
- autosupport messages for vfilers 33

B

- broadcast packets in an IPspace 49

C

- checklists for migration, disaster recovery
 - networking 93
 - storage 92
- CIFS
 - allowing or disallowing 30
 - local user accounts 66

- path names for shares 64
- prefdc command (domain controller) 98
- statistics 43
- virus scanning 114
- cifs
 - setup command, syntax for 20
 - setup command, when to run 19
 - stat command 43
- clients
 - effect of vfiler move on 107
- clusters
 - considerations when used with IPspaces 51
 - licensing vfilers for 13
 - naming IPspaces for 51
- commands for vfilers
 - add 23
 - allow 31
 - context 34
 - create 17
 - destroy 29
 - disallow 31
 - dr activate 98
 - dr configure 95
 - dr delete 99
 - executable from hosting filer 34
 - migrate 108
 - migrate complete 109
 - migrate start 108
 - move 24
 - prerequisites for entering through rsh 35
 - remove 24
 - run 34
 - setup 19, 65
 - start 27
 - status 81
 - stop 27
 - ways to execute 34
- context of vfiler
 - switching to 34
- creating disaster-recovery vfiler 95
- creating vfilers
 - in nondefault IPspaces 60

- options with default settings 14
- prerequisites for 12
- required status of network interface 15
- results of 14

D

daemon

- disabling the routed 11
- effects of disabling the routed 40
- enabling the routed 11

data migration 77

- explanation 78
- using vfilers for 3

decreasing the limit for vfilers per filer 26

default IPspace, interfaces in 48

default limits for vfilers per filer 25

default vfiler, definition of 4

deleting disaster-recovery vfiler 99

destination filer

- network checks for migration 84
- preparing for migration or disaster 81
- storage checks for migration 81

destroying IPspaces 59

destroying vfilers

- effects on quotas 29
- meaning of 29
- prerequisites for 29
- steps for 29

disabling or enabling quotas 69

disabling vfiler license

- effects of 11
- prerequisites for 11

disallowing or allowing

- CIFS or NFS, effects of 30
- protocols, steps for 31
- rsh, effects of 30

disaster recovery 77

- explanation 78
- network checks 84
- quotas 83
- reactivating original vfiler 100
- storage checks 81
- summary of tasks 79

disaster-recovery vfiler

activating 98

creating 95

deleting 99

preparing to create 81

disks for vfilers, moving 12

displaying

- quota reports 74
- quota status 73, 83

DNS servers

- and migration, disaster recovery 89

E

enabling or disabling quotas 69

enabling vfiler license, effects of 11

entering commands from the vfiler 34

establishing a vfiler, major steps for 10

F

filer reboot, effects on vfilers 37

H

hosting filer

- access to data on vfilers 4
- definition of 4
- format of quota reports on 74
- inventorying for vfiler replication 81
- managing vfiler storage from 21
- managing vfilers from 6
- quotas not copied to new vfiler 83

I

incoming packets for an IPspace 49

increasing the limit for vfilers per filer 25

interfaces, assigning to IPspaces 57

IP addresses

- and migration, disaster recovery 84
- assigned to vfilers 15
- configuring 15
- in different IPspaces 60
- planning for when creating vfilers 12
- removing from a vfiler 22
- removing from an interface 57

- unconfiguring 105
- IP alias, removing 15
- ipSPACE command 54
- IPspaces
 - and migration, disaster recovery 86
 - assigning interfaces to 57
 - case studies of 53
 - considerations when used with clusters 51
 - creating vfilers in nondefault 60
 - destroying 59
 - destroying vfilers in different 29
 - guidelines for using 46
 - incoming and outgoing packets in 49
 - interfaces in 48
 - listing 56
 - managing using the ipSPACE command 54
 - names for in a cluster 51
 - restricting broadcast packets in 49
 - routing tables for 40, 46, 48
 - typical applications of 46
 - used with VLAN tagging 50

L

- language, guidelines for 13
- license for MultiStore, enabling or disabling 11
- limit (vfilers per filer)
 - adjusting 25
- listing IPspaces 56
- local user accounts 66

M

- managing vfilers from vfilers 7
- mandatory_scan option 114
- maximum number of vfilers 4
- messages, autosupport for vfilers 33
- MIBs
 - NetApp custom 42
 - standard 42
- migrating (moving) a vfiler 107
 - summary of tasks 80
- migrating data, using vfilers for 3
- moving a vfiler 77, 107
 - effect on clients 107
 - network checks 84

- preparation 81
- quotas 83
- storage checks 81
- summary of tasks 80
 - to a different subnet 109
- moving or removing resources, requirements for 22
- moving resources, meaning of 22
- MultiStore, understanding 2

N

- naming storage resources in vfilers 12
- naming vfilers, guidelines for 13
- NetApp custom MIBs 42
- network checks
 - for migration, disaster recovery 84
- network interfaces
 - configuring down 16
 - for vfilers 4
 - removing an IP alias from 15
- network resources
 - adding 23
 - base IP address 15
 - changing 22
 - considerations for moving between vfilers 23
 - IP alias 15
 - planning for when creating vfilers 12
 - removing 23
 - requirements for moving and removing 22
 - steps for moving between vfilers 24
 - types of 2

NFS

- allowing or disallowing 30
- path names for exporting 64
- starting the protocol 65
- statistics 43
- nfsstat command 43
- NIS servers
 - and migration, disaster recovery 87
- no-vfiler-ips? variable
 - setting 104

O

- oplock settings for qtrees on vfilers, changing 38
- options

- after vfiler creation 14
- settable on vfilers 36
- specifying path names for 36
- options snapmirror command 101
- options snapmirror.access command 102
- original vfiler
 - reactivating after disaster 100
- outgoing packets for an IPspace 49

P

- packets
 - broadcast in an IPspace 49
 - incoming for an IPspace 49
- path names, for NFS exports and CIFS shares 64
- primary storage unit, definition of 12
- production (original) vfiler
 - reactivating after disaster 100
- protocols
 - allowing or disallowing 30
 - supported for vfilers 30

Q

- qtree command output, vfilers and hosting filer 39
- qtrees
 - assigning to vfilers 12
 - creating on vfiler0 38
 - creating on vfilers 38
 - oplock settings for 38
 - requirement for replication on new vfiler 82
 - security styles for 38
 - who can create on vfilers 38
- quota report command 83
- quotas
 - displaying status of 73
 - effects of destroying vfilers on 29
 - effects of disabling 69
 - guideline for 13
 - information in standard MIB 42
 - messages from exceeding thresholds 76
 - prerequisite for turning on and off 69
 - reports on 74
 - resizing 71
 - seeing where enforced from 83
 - tree 74

- turning on and off 69
- types supported for vfilers 68
- user and group 72
- vfiler names in report 74, 75
- where user and group quotas take effect 72
- which are copied to new vfiler 83
- who specifies 68

R

- reactivating original vfiler after disaster 100
- rebooting filer, effects on vfilers 37
- removing IP addresses from an interface 57
- removing IP aliases from an interface 15
- replacement vfiler
 - creating after disaster 104
- resizing quotas 71
- resources
 - administering from the hosting filer 21
 - changing 22
 - guidelines for assigning 12
 - network 2
 - storage 2
- routed daemon
 - disabling 11
 - effects of disabling 40
 - enabling 11
- routing table
 - for the filer 40
 - for vfilers in default IPspace 40
- rsh
 - access to vfiler from clients 35
 - allowing or disallowing on vfilers 30
 - enabling 35
 - entering commands for vfilers through 34
- rsh.access option 35
- rsh.enable option 30, 35
- running state of vfilers 14

S

- sample vfiler 4
- scanners, virus 114
- server consolidation, using vfilers for 2
- Server Manager 66
- service providers, using vfilers for data hosting 3

- setting up vfilers
 - introduction to 19, 65
 - steps for 19
- setup command for vfilers
 - purpose of 19, 65
 - running to unconfigure IP addresses 105
 - syntax for 20
- SnapMirror
 - license requirement 78
 - options snapmirror.access command 102
 - using to reactivate vfiler 100
- snapmirror break command 103
- snapmirror command
 - considerations for 41
 - where to enter 41
- snapmirror quiesce command 102
- snapmirror resync command 101
- snapmirror update command 102, 105
- snapmirror.access option 101
- snapmirror.allow file 102
- SNMP
 - on the hosting filer 42
 - using with vfilers 42
- standard MIB 42
- starting vfilers
 - meaning of 27
 - step for 27
- state of vfiler
 - after vfiler creation 14
 - depends on role after replication 78
 - effect of filer reboot on 37
- stopping vfilers
 - meaning of 27
 - step for 27
- storage checks
 - for migration, disaster recovery 81
- storage resources
 - adding 23
 - administering from the hosting filer 21
 - assigning qtrees 12
 - assigning volumes 12
 - changing 22
 - considerations for moving between vfilers 23
 - moving between vfilers 24
 - naming 12

- planning for when creating vfilers 12
 - primary unit 12
 - removing 23
 - requirements for moving and removing 22
 - types of 2
- syslogd messages 33
- sysstat command 43

T

- TCP statistics in standard MIB 42
- thresholds, quota 76
- timezone 19
- trusted host
 - and migration, disaster recovery 90
 - changing after migration 110
 - changing on disaster-recovery vfiler 96

U

- UDP statistics in standard MIB 42
- uptime command 43
- User Manager 66
- UUID for vfiler 32

V

- vfiler
 - activating backup for disaster-recovery 98
 - add command, syntax for 23
 - administering resources from hosting filer 21
 - allow command, syntax for 31
 - command family, path name in 18
 - command family, purpose of 18
 - commands, entering from hosting filer 34, 35
 - commands, entering from vfiler context 34
 - commands, entering through rsh 35, 36
 - create command, syntax for 17
 - creating backup for disaster-recovery 77
 - creating disaster-recovery 95
 - destroy command, syntax for 29
 - disallow command, syntax for 31
 - displaying status 32
 - group in NetApp custom MIB 42
 - limit per hosting filer 25
 - move command, syntax for 24

- moving 108
- moving (migrating) 77, 107
- moving (summary of tasks) 80
- moving to different subnet 109
- moving to different Windows domain 109
- preparing backup for disaster recovery 79
- reactivating original after disaster 100
- reactivating using vfiler dr commands 103
- reactivating via SnapMirror commands 100
- remove command, syntax for 24
- replacing on original filer 104
- run command, purpose of 34
- setup 19, 65
- start command, syntax for 27
- states, effect of filer reboot on 37
- states, stopped or running 27
- stop command, syntax for 27
- vfiler context command 34
- vfiler dr activate command 98
 - what it does 98
- vfiler dr configure command 95
 - what it does 97
- vfiler dr delete command 99
- vfiler migrate cancel command 109
- vfiler migrate command 108
- vfiler migrate complete command 109
 - what it does 111
- vfiler migrate start command 108
 - what it does 110
- vfiler status -r command 81
- vfiler0

- definition of 4
- included in vfiler limit 25
- virus scanning for 114
- vfiler template option, meaning of 18
- virus scanning
 - configuring 114
 - for vfiler0 114
 - registering scanners for 114
 - requirements for 114
- VLAN tagging, used with IPspaces 50
- volumes
 - requirement for replication on new vfiler 82
- volumes in vfilers
 - assigning 12
 - effects of renaming 38
 - taking offline 38

W

- Windows domain
 - and migration, disaster recovery 90
 - changing 98
- Windows domain controller
 - changing 98
- WINS server
 - and migration, disaster recovery 90
 - changing 98
- worksheet for migration, disaster recovery
 - networking 93
 - storage 92