

gFiler™ Gateway Series

Software Setup, Installation, and 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 210-00661
May 2004

Copyright and trademark information

Copyright information

Copyright © 1994–2004 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Portions of the software were created by Netscape Communications Corp.

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

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

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.
2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [Sdate-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de

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

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

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

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

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

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

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

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

Trademark information

NetApp, the Network Appliance design, NetApp—the Network Appliance Company, FAServer, NearStore, NetCache, and WAFL are registered trademarks of Network Appliance, Inc. in the United States and other countries. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. DataFabric, FilerView, SecureShare, SnapManager, SnapMirror, and SnapRestore are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP and Snapshot are trademarks of Network Appliance, Inc. in the United States and other countries. ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, gFiler, MultiStore, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapMover, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and SpinStor are registered trademarks of Spinnaker Networks, LLC in the United States, other countries, or both. SpinAV, SpinManager, SpinMirror, and SpinShot are trademarks of Spinnaker Networks, LLC in the United States, other countries, or both.

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

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Getting Started with Software Setup.	1
	Overview of setup and installation	3
	Getting started with terms	5
	Tools for administering your gFiler	8
Chapter 2	Planning for Software Setup of a gFiler	11
	Planning overview	12
	Requirement for enabling gFilers to read from and write to LUNs	14
	Rules for the disk (LUN) ownership software scheme	16
	Planning for disk (LUN) ownership assignment	19
	Planning for RAID mirroring for data protection	21
	Gathering data needed for the setup program	29
	Where to go from here	38
Chapter 3	Providing Setup Information	39
	Providing disk (LUN) ownership and system identification data	40
Chapter 4	Installing Data ONTAP Software	47
	Installing licenses on the gateway	48
	Installing the software from a UNIX client.	50
	Installing the software from a Windows client	53
	Downloading Data ONTAP	55
	Additional CIFS setup	57
Chapter 5	Completing the Initial Setup	59
	Setup completion process.	60

	Setting up host name resolution	63
	Verifying network connections	64
	Making gateway data available to NFS hosts	66
	Making gateway data available to CIFS hosts	69
	Verifying configuration for clustered gateways	71
	Configuring networking interfaces for clustered gateways.	75
	Checking and testing the configuration for clustered gateways	79
	Verifying the RMC connections and setup.	84
	Setup if your deployment includes a shared switch	86
	Where to go from here	88
Chapter 6	Managing Volumes, RAID Groups, and LUNs	89
	Volume configuration and management	90
	Management of RAID groups associated with NetApp volumes	96
	Managing LUNs and LUN ownership	100
Chapter 7	Managing Workload Distribution with SnapMover	109
	Planning for using SnapMover to change workload distribution.	110
	Changing workload distribution with SnapMover	115
Appendix A	Configuration Worksheet.	117
Appendix B	Time Zones.	119
Appendix C	Supported Languages	125
	Index	127

Preface

About this guide

This guide describes how to install and set up Data ONTAP™ software on NetApp® gFiler™ gateways. The information in this guide pertains to all the supported gFiler platforms: GF270c, GF825/GF825c, GF920/GF920c, GF940/GF940c, and GF960/GF960c.

Follow this guide, chapter by chapter, to complete the gateway setup. Appendixes in this guide provide worksheets to help you gather information you need when setting up the gateway on a network.

Audience

This guide is for users who are familiar with operating systems, such as UNIX®, Windows® 95, Windows NT®, and Windows 2000, and who will be installing gFiler gateways. This guide does not discuss basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the NetApp gateway.

Terminology

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 commands on either the system console or from any host 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.
Monospaced font	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note_____

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

Caution_____

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



Do not proceed with gFiler software upgrade or installation until you have read the *Getting Started Guide* for your vendor and the *Release Notes*. These documents contain information about storage system software and parameter settings that are essential for a successful gFiler software installation.

The following *Getting Started Guides* are currently available in the gFiler Gateway Series library:

- ◆ *Getting Started for Hitachi Data Systems*
- ◆ *Getting Started for IBM*
- ◆ *Getting Started for Hitachi, Limited*

Before you start the set up and installation process of Data ONTAP software on a gFiler, you must do the following:

- ◆ On the storage subsystem, create LUNs and make them available to the gFiler, as described in the following documentation:
 - ❖ The *gFiler Gateway Series Getting Started Guide* for your storage subsystem. (The process of making LUNs available varies among vendors.)
 - ❖ Your storage subsystem documentation
- ◆ Connect your gFiler and storage subsystem
 - ❖ See the *gFiler Gateway Series Hardware Installation Guide*.

Features described in this guide

This guide describes features that can only be used with gFilers. Data ONTAP provides a variety of other features that can be used both with filers and gFilers. Use the Data ONTAP documentation suite for information about setting up those features. The *gFiler Series Getting Starting Guide* lists Data ONTAP features that are not applicable to gFilers.

About this chapter

This chapter discusses the prerequisites to software setup, provides an overview of the setup process, and defines some gFiler terms used in this guide.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Overview of setup and installation”](#) on page 3
- ◆ [“Getting started with terms”](#) on page 5
- ◆ [“Tools for administering your gFiler”](#) on page 8

Overview of setup and installation

Overview of setup and installation tasks

The following table shows the major tasks in setting up a gFiler.

Step	Action	For information see...
1	Configuring specific gFilers to be able to read data from and write data to LUNs on the storage subsystem. This process is referred to as assigning <i>disk (LUN) ownership</i> .	Chapter 2, “ Planning for Software Setup of a gFiler ,” on page 11 and Chapter 3, “ Providing Setup Information ,” on page 39
2	Providing answers to prompts about the system using the data that you gathered in “ Gathering data needed for the setup program ” on page 29.	Chapter 3, “ Providing Setup Information ,” on page 39
3	Installing the gFiler storage system license and an NFS or a CIFS license. Before you can install the Data ONTAP software on a gFiler, you must install the gateway storage system license and an NFS license (when installing from a UNIX client) or a CIFS license (when installing from a Windows client). Licenses for other Data ONTAP features can be installed before or after the Data ONTAP software is installed.	Chapter 4, “ Installing Data ONTAP Software ,” on page 47
4	Installing and downloading Data ONTAP software.	Chapter 4, “ Installing Data ONTAP Software ,” on page 47
5	Verifying that your setup and installation is working properly.	Chapter 5, “ Completing the Initial Setup ,” on page 59

Step	Action	For information see...
6	Completing any additional setup that is dependent on Data ONTAP software installation first; for example:	
	◆ Configuration if you have switches	“Setup if your deployment includes a shared switch” on page 86
	◆ (Optional) Creating volumes to facilitate data management	Chapter 6, “Managing Volumes, RAID Groups, and LUNs,” on page 89
	◆ (Optional) Configuring SnapMover to enable you to change workload distribution	Chapter 7, “Managing Workload Distribution with SnapMover,” on page 109
	◆ (Optional) Configuring other Data ONTAP features that simplify administration, management, and data protection	Data ONTAP documentation suite

Getting started with terms

Cluster

A cluster is a pair of NetApp gFiler heads (or nodes) that are connected through an interconnect adapter and cables and are configured so that both gFiler heads share access to a set of Fibre Channel disks, subnets, and tape drives to provide fault tolerance. In most cases, the clustered nodes are two physically separate nodes. However, with the GF270c, a cluster consists of two internal gFiler heads that reside in the same appliance.

When one gFiler in a cluster undergoes a system failure and cannot reboot, the partner gFiler in the cluster can take over the failed gFiler functions and serve data. This process is known as a takeover.


In addition to enabling automatic takeover, you can also force a takeover at any time. You can cause a clustered gFiler to immediately halt and be taken over by its partner, allowing scheduled maintenance to be performed on the gFiler. Alternatively, you can cause a gFiler to halt while preventing its partner from taking over.

LUN

This guide uses the term *LUN* (Logical Unit Number) to describe an area on the storage subsystem that is available for a gFiler or non-gFiler host to read data from or write data to. You might be accustomed to hearing a different term to describe this area; the term varies among vendors and sometimes among platforms for the same vendor. See the *Gateway Series Getting Started Guide* for your storage subsystem for the specific term used for your platforms.

RAID protocols supported

gFilers support storage subsystems that use the RAID-1, RAID-5, and RAID-10 protocols, although the gFilers do not themselves use RAID-1, RAID-5, or RAID-10. The following illustration provides a conceptual view of RAID-0 used on the gFiler and the RAID-1, RAID-5, and RAID-10 storage subsystems.



The gFiler stripes across LUNs in RAID-0, thereby maximizing performance because it is using more spindles. Another benefit of the gFiler using RAID-0 for striping across LUNs is that you can add LUNs to a NetApp volume.

There is no need for the gFiler to provide disk protection. The LUNs are already protected with the redundancy provided by RAID-1, RAID-5, and RAID-10 on the storage systems.

Note

The filer uses the RAID-4 protocol for disk protection. RAID-4 protocol is not needed for gFilers because the storage subsystems provide disk protection. Most of the disk information in the Data ONTAP documentation does not apply to the gFiler because it is about RAID-4.

Volume

In Data ONTAP, *volumes* are the primary means of organizing and managing data. From the perspective of Data ONTAP, a volume is an identified collection of disks comprising a filesystem. From the perspective of a storage subsystem, a volume is an identified collection of LUNs on the storage subsystem. You manage data storage by managing volumes.

You can define more than one volume on a gFiler. You could, for example, create additional volumes to accommodate your users if they speak more than one language. For example, you could create one volume for files in French and another for files in German. See Chapter 6, “[Managing Volumes, RAID Groups, and LUNs](#),” on page 89 for more information about creating and managing volumes.

Tools for administering your gFiler

Where to find information about administration

Use the Data ONTAP documentation in conjunction with the gFiler series of documentation. The gFiler software documents contain the information that you need to set up the Data ONTAP software on your gFiler so that storage subsystem and the gFiler can operate together. The gFiler software documents also contain information about features that are not supported in Data ONTAP or work differently for the gFiler and the filer. See the *gFiler Gateway Series Documentation Roadmap* for information about the documents in the gFiler gateway series documentation library.

Of the guides in the Data ONTAP documentation set, NetApp recommends that you read the *Storage Management Guide* first.

Terminology differences in Data ONTAP documents: In some cases, different terminology is used between filer users and gFiler users (that is, users from a storage subsystem background). In most cases, the Data ONTAP documents use the terminology relevant to filer users. Therefore, you need to “translate” some of the words in the Data ONTAP guides to the context applicable to the gFiler. For example:

- ◆ Filers store data on disks in disk shelves. gFilers access data stored in LUNS on the storage subsystems.
 - ❖ Translate the term *filer* to *gFiler* or *gateway*.
 - ❖ Translate the term *disk* to *LUN*.
- ◆ Generally, you can disregard discussions about disk shelves.

Any specific platform information stated for the filer is true for the gFiler series counterpart.

Administration from FilerView

The FilerView® interface is a Web-based administration tool for Data ONTAP that enables you to perform most administration tasks from any host on your network that has a compatible Web browser. Through FilerView, you can perform tasks that otherwise require you to enter commands at the gateway prompt or edit configuration files.

You are not required to purchase an HTTP license to use FilerView.

Accessing FilerView

To access FilerView, complete the following steps.

Step	Action
1	Launch a Web browser, either Netscape Navigator 4.7 or later or Microsoft Internet Explorer 5 or later.
2	Point your browser to the gateway home page (http://hostname/na_admin), substituting the gateway name or IP address for <i>hostname</i>. Example: If your host's name is toaster, enter the following URL: http://toaster/na_admin Result: When your browser connects to the gateway, you see the gateway home page.

The initial page contains the choices shown in the following table.

Choice	Description
Install Documentation	Initiates a browser to select zip files to install the HTML and PDF versions of product documentation.
Manual Pages	Links to UNIX-style manual pages for the Data ONTAP console commands.
FilerView	Displays the FilerView program in a separate window.
Filer At-A-Glance	Displays a tool that monitors the performance of the gateway.
Initiate a Support Call	Initiates a support case to Network Appliance Customer Support.

How FilerView works

The FilerView interface consists of two frames: a left frame and a right frame.

Left frame: The left frame contains an expandable list of categories:

- ◆ The Real Time Status category contains choices that launch separate tools that monitor gateway performance.
- ◆ The other categories represent management functions.

You can view the real-time displays or the functions in a category by clicking the Real Time Status or a category name. You can select a real-time display or a function by clicking the display or function name.

Right frame: The right frame contains a logo initially.

As you select functions from the left frame, the right frame changes to display forms that provide information about the gateway configuration and enable you to change the gateway's configuration. You can alter the gateway configuration by changing information that is displayed in forms and then saving the information. The top of the window contains a Help button that you can click to view Help for the form that is currently displayed in the right frame.

Autosupport feature

Autosupport is automatically enabled when you run setup on a gateway for the first time. Depending on how Autosupport is configured, it sends event messages and weekly reports to your specified recipients, NetApp Technical Support, or both. Autosupport can significantly decrease the time it takes to assess and resolve problems that might occur on your system. If you do not want to send any Autosupport messages, you can disable Autosupport by entering `options autosupport.support.enable off` within 24 hours after you complete the setup process. You can also disable Autosupport any time after the first 24 hours by using the same command.

See your *Getting Started Guide* to determine whether your organization will be using Autosupport.

NOW Web site

The NOW site at <http://now.netapp.com> is available to some customers. It provides a variety of information, including a knowledge base. If you do not have access to NOW, contact your authorized reseller.

About this chapter

This chapter helps you understand what you need to plan for before you start setting up software on your gFilers.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Planning overview”](#) on page 12
- ◆ [“Requirement for enabling gFilers to read from and write to LUNs”](#) on page 14
- ◆ [“Planning for disk \(LUN\) ownership assignment”](#) on page 19
- ◆ [“Planning for RAID mirroring for data protection”](#) on page 21
- ◆ [“Gathering data needed for the setup program”](#) on page 29
- ◆ [“Where to go from here”](#) on page 38

Planning overview

Planning tasks

The following table summarizes the planning tasks for a gFiler after LUNs have been made available for gFilers on the storage subsystem. To make LUNs available for gFilers, a storage system administrator must create the LUNs on the storage subsystem and map the LUNs to ports on the storage subsystem to which the gFiler is connected.

Planning task	Summary
Plan which gFiler has read and write privileges to a LUN on a storage subsystem	Data ONTAP imposes a software control to allow only one gFiler to write to or read from a single LUN. You must determine which gFiler is to have read and write privileges to each LUN that has been made available on the storage subsystem to gFilers. See “Requirement for enabling gFilers to read from and write to LUNs” on page 14 for more information.
If you are using the SyncMirror™ product, plan how to configure each gFiler so that each of your storage subsystems has an exact copy of the data	A specific parameter is required when you configure the gFilers to support SyncMirror. See “Planning for RAID mirroring for data protection” on page 21.
If you are using a switch in your deployment, plan for limiting commands (read and write requests) on the storage subsystem port	Performance can be affected by too many commands on the storage subsystem port being sent by the gFiler and non-gFiler hosts. See “Setup if your deployment includes a shared switch” on page 86.
Determine your strategy for volumes and LUNs	Volumes function as containers to help you manage data on your storage subsystem. You determine how to allocate LUNs into volumes based on your requirements for performance and capacity. See Chapter 6, “Managing Volumes, RAID Groups, and LUNs,” on page 89 for more information.

Planning task	Summary
Plan for changing workload distribution by moving vFiler™ logical units	The volume of requests that a specific gFiler is required to handle might change after initial deployment. If you want the capability of easily changing the workload distribution by moving vFilers, you need to plan and configure Data ONTAP before you need to change the distribution. See Chapter 7, “ Managing Workload Distribution with SnapMover ,” on page 109.
Complete a worksheet of information that the setup program prompts you for	The setup program requires basic system information. See “ Gathering data needed for the setup program ” on page 29.


Requirement for enabling gFilers to read from and write to LUNs

Storage subsystems must make LUNs available on ports

Before a gFiler can see LUNs on a storage subsystem, one or more LUNs must be created on the storage subsystem and then configured to be available (visible) to the gFiler on the ports to which the gFiler is connected. The process varies among vendors for how LUNs are created and made available to gFilers. The process also varies as to how access controls on LUNs are imposed. See the *Gateway Series Getting Started Guide* for your storage subsystem for details.

Logical relationship between gFilers and LUNs

gFilers use a software-based scheme to control access of gFilers to LUNs that a storage subsystem has made available. This scheme enables you to make changes at the software level that in typical systems can only be made in a fixed LUN-ownership scheme. The following illustration shows the relationship between the LUNs that are made available for gFilers on the storage subsystem and the gFiler configuration scheme.



With the gFiler software scheme, for a gFiler to read from or write to a LUN that the storage subsystem has made available, you must establish a logical relationship between the gFiler and the LUN through a gFiler command (`disk assign`). From the point of view of the Data ONTAP software, this logical relationship is referred to as *disk ownership* (because a gFiler considers a LUN to be a virtual disk). From the point of view of the storage subsystem, you can think about disk ownership as the gFiler owning the LUN. Only one gFiler can own a specific LUN; Data ONTAP writes data to the LUN to identify the gFiler that is the “owner.”

Caution

The gFiler software-based scheme provides ownership control only for gFilers; it does not prevent a non-gFiler host from overwriting data in a LUN that a gFiler owns. Therefore, if a non-gFiler host is sharing an storage subsystem port with one or more gFilers, use LUN security on your storage subsystem to prevent gFilers and non-gFilers from overwriting LUNs.

Benefit of the disk ownership scheme

The primary benefits of the disk (LUN) ownership scheme are

- ◆ You can leave LUNs unused until you need them. When you need them, you can use a command at the gFiler console to add them to the storage that the gFiler is using.
- ◆ If you need to move data service from one gFiler to another because one gFiler is handling too much load, you can do so at the software level through a process that is transparent to users. You do not need to reconfigure hardware or the LUN access controls on the storage subsystem to change which gFiler is servicing data requests. Chapter 7, “[Managing Workload Distribution with SnapMover](#),” on page 109 provides more information about the gFiler scheme of moving vFilers from one gFiler to another, which is based on the disk ownership scheme.

Rules for the disk (LUN) ownership software scheme

Facts about disk (LUN) ownership


The following rules apply to the gFiler disk (LUN) ownership scheme:

- ◆ You can configure only one gFiler to own a specific disk (LUN).
- ◆ During initial setup of the gFiler, you must identify at least *one* disk (LUN). Otherwise, you cannot set up the gFiler.
- ◆ After you perform the initial setup of your gFiler, you can add additional LUNs that the gFiler is to own. See “[Managing LUNs and LUN ownership](#)” on page 100 for more information.
- ◆ A LUN must be part of a volume for it to be used for data storage.
You might want the LUN to be part of the root volume or part of a different volume that you create later. See “[Planning for disk \(LUN\) ownership assignment](#)” on page 19 for information about how the system handles LUNs if you enter multiple LUNs during the initial setup.
- ◆ Until the LUN that a gFiler owns is configured to be part of a volume, it is a *spare* LUN.
- ◆ You can remove all ownership status of a LUN by a gFiler. If you do so, the content of the LUN is not affected because disk (LUN) ownership is a logical relationship. The gFiler can no longer access the LUN after ownership status is removed, however. See “[Managing LUNs and LUN ownership](#)” on page 100 for information about removing ownership status of a LUN.
- ◆ In most circumstances, if you want to move data service from one gFiler to another, you need to set up your gFiler to support the SnapMover™ and MultiStore™ vFiler products. Through SnapMover, you can move the volume in which the LUN resides. See Chapter 7, “[Planning for using SnapMover to change workload distribution](#),” on page 110 for more information.
- ◆ Not all LUNs that the storage subsystems make available for the gFilers must be owned. From the point of view of Data ONTAP, a LUN that is unowned is not a spare LUN; a spare is owned by the gFiler.


The following scenarios summarize the rules in the previous list from a configuration point of view.

Viewing disk (LUN) ownership from a configuration point of view

Scenario 1—disk ownership is not configured yet: In this scenario, the LUNs have been made available (visible) on the storage subsystem. However, gFiler 1 has not yet been configured for disk (LUN) ownership. Therefore, the gFiler cannot read data from or write data to any LUNs on the storage subsystem.




Scenario—only some LUNs are owned: In this scenario, gFiler 1 has been configured to own LUNs 1 and 2, but not LUNs 3 and 4, because LUNs 3 and 4 are not yet needed for storage. The smallest of the two LUNs was automatically assigned to the root volume during initial setup of the gFiler. The other LUN was automatically allocated as a spare and remains a spare until you add it to a volume after Data ONTAP software installation; gFiler 1 can read data from and write data only to LUN 1 until LUN 2 is configured to be part of a volume.



After you perform initial setup of the gFiler, you could configure gFiler 1 to own LUN 3, LUN 4, or both. See “[Managing LUNs and LUN ownership](#)” on page 100.

Scenario 3—configuration with a cluster: In this scenario, two gFilers are configured in a cluster and the two gFilers share the same storage subsystem.



Assume that during initial setup gFiler 1 was assigned ownership of LUN 1 and LUN 2. LUN 1 was automatically added to the root volume, so the gFiler can access it for storage. LUN 2 remains a spare until you explicitly add it to a volume after Data ONTAP installation.

Similarly, assume that during initial setup gFiler 2 was assigned ownership of LUN 3 and LUN 4. LUN 4 remains a spare until you explicitly add it to a volume after Data ONTAP installation.

The key points of this scenario are as follows:

- ◆ Only one gFiler can own a specific LUN.
- ◆ By deploying the gFilers in a cluster, one gFiler can take over services for the other if a gFiler becomes unavailable.
- ◆ By deploying two switches, if one switch fails, the other switch provides the alternate path to the storage subsystem.

Command for adding disk (LUN) ownership

The command used for assigning disk (LUN) ownership is `disk assign`. See Chapter 4, “[Installing Data ONTAP Software](#),” on page 47 for its use during initial setup of the gFiler and Chapter 6, “[Managing Volumes, RAID Groups, and LUNs](#),” on page 89 for its use after initial setup.

Planning for disk (LUN) ownership assignment

Create a deployment level plan

Before you start configuration, develop a deployment-level plan for software. For example, determine

- ◆ Which gFilers will own which LUNs on the storage subsystems
- ◆ Whether you currently need all the storage provided by all the LUNs available to the gFilers or can you leave some LUNs unowned for the present
- ◆ What your LUN security plan is
- ◆ What your plan for volumes on the gFiler is

Preventing gFilers and non-gFilers from overwriting LUNs

Use LUN security or switch zoning so that LUNs for gFiler hosts are not visible to non-gFiler hosts, and the reverse. (Whether switch zoning is an effective method depends on the configuration.) The reason is that although the gFiler disk ownership scheme prevents one gFiler from overwriting a LUN owned by another gFiler, it does not prevent a gFiler from overwriting a LUN accessible by a non-gFiler host. Likewise, without a method of preventing overwriting, a non-gFiler host could overwrite a LUN used by a gFiler.

Planning for gFiler volumes

In Data ONTAP, *volumes* are the primary means of organizing and managing data. You manage data storage by managing volumes. You can think about a volume as a container for the LUNs that provide the capability for data storage.

Each gFiler has a special volume known as a *root volume*. The root volume is created automatically during the initial setup of the gFiler, after you assign disk (LUN) ownership in the boot menu Maintenance option. The root volume is comprised of one LUN at that time.

If you entered multiple LUNs during initial setup, one LUN, typically the smallest, is automatically selected to be part of the root volume. If the LUNs that you specify are within a plus or minus 20 percent range of each other, the system considers the LUNs to be the same size and randomly picks one of them as the smallest. The other LUNs are allocated as spares. If your organization wants to ensure that the actual smallest LUN is assigned to the root volume, enter only the actual smallest LUN in the boot menu Maintenance option.

Note

You cannot remove the LUN to be selected to be part of the root volume or move it to another volume later. However, after Data ONTAP is installed, you can create another volume, make it the root volume, and destroy the original root volume. See the *Data ONTAP Storage Management Guide* for more information.

The root volume size is at least 72 GB because 72 GB is required to hold core files and log files that might be generated. You can expand the root volume to 72 GB after initial setup.

After initial setup and Data ONTAP software installation, you can add LUNs to the root volume, or you can create additional volumes and add LUNs to those volumes. Creating multiple volumes provides you the ability to partition your data for management, which you might or might not need.

Note

After you add a LUN to a volume, you cannot remove it.

See “[Volume configuration and management](#)” on page 90 for volume metrics, considerations for creating multiple volumes, and other information, such as the performance impact of expanding volumes.

**Requirement for
spare LUNs for a
GF270c**

Each head of a GF270c must have one spare LUN of 1.5 GB or more assigned to it at all times, in case a problem occurs that results in a panic. The system can then dump core to the spare LUN and NetApp can analyze the core dump to determine the cause of the problem.

**Disk (LUN)
ownership
configuration if you
are using local
SyncMirror**

If you want to deploy the local SyncMirror product to provide synchronous RAID mirroring, you must plan ahead. The reason is that when you configure disk (LUN) ownership, you must correctly configure the pools associated with each LUN so that RAID mirroring works correctly. See “[Planning for RAID mirroring for data protection](#)” on page 21 for more information.

Planning for RAID mirroring for data protection

Reason to use SyncMirror

This section discusses how you can use the optional NetApp SyncMirror product to implement RAID mirroring for redundancy. Customers who choose to use SyncMirror typically have two storage subsystems and cannot tolerate any down time. Depending on your requirements, you might use SyncMirror alone or in conjunction with MetroCluster (see “[SyncMirror with MetroCluster](#)” on page 27).

SyncMirror cannot be used with the GF270c.

Note

The focus of this section is to help you plan for SyncMirror setup specific to a gFiler. See the *Data ONTAP 6.5 Data Protection Online Backup and Recovery Guide* for details about SyncMirror and the *gFiler Gateway Series Getting Started Guide* for information about the hardware configuration for a MetroCluster deployment.

Licenses required


A license is required to use SyncMirror. If you are using MetroCluster with SyncMirror, you must also have a license for MetroCluster.

Configuration requirements if you use SyncMirror

If you are using SyncMirror, when you configure disk (LUN) ownership information on a gFiler, you must correctly configure the pool parameter so that RAID mirroring works correctly.

Synchronous mirroring and updating of data with SyncMirror

The SyncMirror software creates two physically separated copies of the volume, called *plexes*. The plexes are simultaneously updated; therefore, the two copies of the data are always identical. The following illustration shows two different plexes of a volume. (Data ONTAP automatically names the plexes.)



If a failure occurs, the unaffected plex continues to serve data while the cause of the failure is fixed. After the problem that caused a failure is fixed, the two plexes can be resynchronized and the mirrored relationship reestablished.

The physical separation of the plexes protects against data loss if

- ◆ A LUN fails, for example, because of a double disk failure on the storage subsystem
- ◆ A storage subsystem becomes unavailable

Plexes are not created during the disk (LUN) ownership assignment process. They are created after you complete the disk (LUN) ownership process and then create a mirrored volume, as described in [“What else you need to do to use SyncMirror with a gFiler”](#) on page 26.


Note

SyncMirror requires twice as much storage space on a mirrored volume as you would need on an unmirrored volume. Each of the two plexes requires enough space to accommodate the data to be stored. For example, you need 200 GB of disk space to mirror a 100 GB volume, 100 GB for each plex of the mirrored volume.


Relationship of plexes to pools

One plex of a volume must reside on each of the storage subsystems so that the data can be mirrored. The scheme that Data ONTAP uses to specify the location of the plexes is through the assignment of pools. In Data ONTAP, pools are used to split a set of disks (LUNs) into two groupings so that the two groupings do not share common single points of failure. When you configure the disk assign command on a gFiler on which you are going to use SyncMirror, you must specify the pool parameter correctly. Your goal when you specify the pool parameters is to group the LUNs belonging to the two subsystems into two pools, so when you later create a mirrored volume, the LUNs for one plex are derived from the same pool.

The following illustration shows the relationships of plexes and pools—one plex is associated with pool 0 and one plex is associated with pool1. The number 0 is typically associated with the local pool and the number 1 is typically associated with the remote pool. For performance reasons, you want reading to and from the local pool—pool 0—thereby avoiding reads over long distances.



The following illustration shows a deployment with a gFiler cluster and two storage subsystems. For simplicity, assume that each gFiler has only one volume and one switch. Your goal is to mirror each gFiler's volume by the way that you configure the pools on the gFiler.



As the preceding illustration shows, for gFiler 1, one plex of the volume, pool 0, the local pool, is be on the Storage Subsystem A. The other plex of the volume, plex 1, pool 1, is be on StorageSubsystem B. Similarly, for gFiler 2, one plex of its volume, pool 0, is on Storage Subsystem B and the other plex of the volume, pool 1, is on Storage Subsystem A.

gFiler-specific configuration requirements for SyncMirror

The goal when you use SyncMirror is for the data to be mirrored exactly on the two storage subsystems. Then, if one storage subsystem becomes unavailable, all data can continue to be served.


If you use SyncMirror with gFilers, you must configure the pool parameters on all the gFilers when you configure the `disk assign` commands, and ensure that they are configured correctly. The `disk assign` command includes a `-p` parameter (where `-p` means pool). For LUNs for which you assign gFiler ownership during initial setup, you use the `disk assign` command with the `-p` parameter through the gFiler boot menu Maintenance option. After initial setup, you use the `disk assign` command with the `-p` parameter from the command line interface (CLI).

Caution

If you incorrectly configure the `-p` parameter when you configure disk ownership, you might not be able to create a mirrored volume. The mirror creation will not allow you to have mixed pools within a plex.

If you incorrectly configure the `-p` parameter, you can lose data if a disaster occurs that makes one storage system unavailable.

The following illustration provides a conceptual view of the configuration of the pools on the gFilers. The full command used on the gFiler for assigning disk (LUN) ownership is discussed in “[Actual commands](#)” on page 26. The following illustration shows the actual parameter for pool, `-p`, however.



1 Available LUNs: Notice that in the previous illustration the two storage subsystems, StoreA and StoreB, each made four LUNs available to the gFilers.

2 gFiler pool configuration: Now notice the way disk ownership is configured on each of the gFilers in relation to the exported LUNs on each storage subsystem. In this scenario, gFiler1 and gFiler2 are each configured to own two LUNs on each storage subsystem. In reality, each gFiler could be assigned to own fewer or more LUNs. The only restriction is that only one gFiler can own a specific LUN.

Look at how the pool (-p) parameters were configured on each gFiler. Pool 0 is typically assigned to the local pool and pool 1 is assigned to the remote pool.

3 Resulting pool configuration: As a result of the configuration of the -p parameter, each storage subsystem has a copy of each volume.

Actual commands: Let us look at this same scenario as in the previous illustration showing the actual commands that you would enter on each of the gFilers. The command for disk (LUN) ownership is named `disk assign`. When configuring the gFiler for SyncMirror, you must include the pool parameter (-p) with the `disk assign` command. (The -p parameter is not required when SyncMirror is not deployed.) You do not need to enter the `disk assign` command separately for each LUN ownership assignment, as was implied in the illustration. The actual commands for gFiler1 in the previous illustration are as follows:

```
disk assign storeAlun1 storeAlun2 -p 0
disk assign storeBlun1 storeBlun2 -p 1
```

The result of this configuration is that gFiler1 now owns four disks (LUNs), two on each storage subsystem.

The actual commands for gFiler2 in the previous illustration are as follows:

```
disk assign storeAlun3 storeAlun4 -p 1
disk assign storeBlun3 storeBlun4 -p 0
```

The result of this configuration is that gFiler2 now owns four LUNs, two on each storage subsystem.

What else you need to do to use SyncMirror with a gFiler

Assigning the pools correctly during the disk ownership configuration in the setup program is only one part of preparing your gFilers for using SyncMirror. To use the SyncMirror feature, you must do the following on each gFiler participating in mirroring:

- ◆ Install the SyncMirror license
- ◆ Install and download the Data ONTAP software
- ◆ Create a mirrored volume

Let us use the previous scenario to illustrate how to create a mirrored volume. In the previously-described scenario, the following was configured on gFiler1 for disk ownership, with the -p parameter configured to support SyncMirror:

```
disk assign storeAlun1 storeAlun2 -p 0
disk assign storeBlun1 storeBlun2 -p 1
```

The result of this configuration was that gFiler1 owns four LUNs. To create a mirrored volume (after installing the license), you would enter

```
vol create mirrorvol -m 4
```

- ◆ *mirrorvol* is the name of mirrored volume name. (You could enter another name for the volume.)
- ◆ -m specifies a mirrored volume to Data ONTAP.
- ◆ 4 is four disks (LUNs)

The *mirrorvol* command creates a mirrored volume with two plexes of two disks (LUNs) each. One plex contains the LUNs on storeA and the other plex contains the LUNs on storeB. The creation of the plexes does not allow mixing pools within a given plex.

After creating a mirrored volume on the first gFiler, you would create a mirrored volume on the other gFiler, using the same command. Creating mirrored volumes provides the maximum amount of fault isolation for handling the failures of the gFiler or either storage subsystem.

The SyncMirror chapter in the *Data ONTAP Data Protection, Online Backup, and Recovery Guide* provides detailed information about various tasks associated with using SyncMirror.


SyncMirror with MetroCluster

With a gFiler, SyncMirror is typically used in a MetroCluster deployment. The primary goal of a MetroCluster is to provide redundant storage services to mission-critical applications in case of site-specific disasters (for example, in the case of fire or long term power loss). The MetroCluster feature requires SyncMirror and a gFiler cluster configured for mirrored volumes to build a system that provides continuous updating of data and data protection over a distance.

Note

MetroCluster and SyncMirror cannot be used with the GF270c.

The following illustration shows an example of a simple MetroCluster deployment. The pools are configured for the MetroCluster in the same way as described in “[gFiler-specific configuration requirements for SyncMirror](#)” on page 24.



In a MetroCluster deployment with gFilers

- ◆ There can be up to six nodes in a neighborhood, four of which can be in MetroClusters.
A neighborhood is the number of nodes that can see each others' LUNs. A neighborhood is created on the storage subsystem by putting these nodes in each other's fabric. Because of the gFiler disk (LUN) ownership scheme, gFilers cannot write over each others' LUNs.
- ◆ In a given neighborhood, you can configure a maximum of two MetroClusters. If needed, you can have two additional gFilers outside of the MetroClusters.
- ◆ You can have multiple neighborhoods in a switch fabric, that is, multiple zones in the fabric.
- ◆ Volume migration through SnapMover is supported between the six nodes in a neighborhood.

Gathering data needed for the setup program

Information you need to gather

After you finish entering disk ownership information into the setup program, you are prompted for gFiler system information to set up your gFiler on the network. Gathering information ahead of time makes setup faster for you. This section describes the information that the setup program prompts you for. You can use the Configuration Worksheet in [Appendix A](#) on page 117 to record your values.

Note

If you are configuring a gFiler as part of a cluster, some information types must be unique for each gateway in the cluster and some information types must be identical on both gateways. In the tables in this section, the information type description indicates whether the information type must be unique for each gateway or identical on both gateways.

Setup information

The following information is needed to set up the gFiler or a gFiler cluster. Some information in the following sections applies only if your gateway uses the feature described in that section. For example, if you use virtual network interfaces, Data ONTAP prompts you for that information. The following sections are listed in the order in which the gateway prompts you for information when setup runs.

Host name: The name by which the gateway is known on the network.

If the gateway is licensed for the NFS protocol, the name can be no longer than 32 characters.

If the gateway is licensed for the CIFS protocol, the name can be no longer than 15 characters.

Caution

If you do not enter the name of an administration host, the gateway and its file systems are available to all hosts on the network. Users from any host can gain access to the gateway's root file system and access the configuration files in the /etc directory. This could cause a security breach.

Virtual network interfaces: If you want to create virtual network interfaces (vifs), gather the following information:

- ◆ Name of each vif
- ◆ Number of links for each vif
- ◆ Link names for each link

Note

The virtual network interface information must be identical on both gateways in a cluster.

See the *Network Management Guide* for details about vifs.

Network interfaces: If your network uses standard Ethernet or Gigabit Ethernet (GbE) interfaces, gather the following information for each interface. See the *Network Management Guide* for details about any of the following information types.

Information type	Description
Network interface name	The name of the Ethernet (or GbE) interface, depending on what port the Ethernet card is installed in. Examples include e0 (for Ethernet single); e1 (for GbE); and e3a, e3b, e3c, and e3d (for Ethernet quad-port). Network interface names are automatically assigned by Data ONTAP as it discovers them.
IP address	A unique address for each network interface.
Subnet mask	The subnet mask for the network to which each network interface is attached. Example: 255.255.255.0

Information type	Description
Partner IP address	If your gateway is licensed for cluster takeover, record the interface name or IP address belonging to the takeover partner that this interface should take over during a cluster takeover. Examples: e0 or 10.10.10.2 Note_____ For vifs, you must specify the interface name rather than the interface IP address. _____
Media type	The type of interface, one of the following: ◆ 100tx-fd (100Base-TX, full-duplex) ◆ tp-fd (10Base-T, full-duplex) ◆ 100tx (100Base-TX, half-duplex) ◆ tp (10Base-T, half-duplex) ◆ auto (autonegotiate speed and duplex)
Maximum Transmission Unit (MTU) size	If the interface supports jumbo frames, the default frame size defined for a jumbo frame on the gateway is 9,018 bytes, and the maximum supported size is 9,210 bytes. These sizes correspond to MTU sizes of 9,000 and 9,192, respectively.

Default gateway: If your network uses routers, gather the following information:

- ◆ Gateway name
- ◆ IP address of the router

Administration host: A host computer that is allowed to access the gateway through a Telnet host or through the Remote Shell (rsh) protocol.

Time zone: The time zone in which the gateway resides. See Appendix B, “[Time Zones](#),” on page 119 for a list of valid time zones.

The time zone must be identical on both gateways in a cluster.

Gateway location: A description of the physical location of the gateway. This information sets the SNMP location information.

Language: The language used for multiprotocol gateways if both the CIFS and NFS protocols are licensed. See Appendix C, “[Supported Languages](#),” on page 125 for a list of supported languages and their abbreviations.

The language must be identical on both gateways in a cluster.

HTTP directory: If your gateway is licensed for HTTP, determine where you are storing Web files and directories. The default directory is /home/http in the gateway’s root volume.

DNS services: If your network uses DNS services, gather the following information.

Information type	Description
DNS domain	The name of your network’s Domain Name Service (DNS). The DNS domain name must be identical on both gateways in a cluster.

Information type	Description	
DNS servers	The IP addresses of your DNS servers.	
	If...	Then...
	If your site uses DNS	Gather the IP address for one or two DNS servers that the gateway should use for host name lookups.
	If your site has a Windows 2000 domain	Gather the IP address of the DNS server that supports your Windows 2000 domain in the Configuration Worksheet. Note ————— When you run the <code>cifs</code> setup command (after setup is complete), Data ONTAP creates the Windows 2000 account through kerberos and LDAP. _____

NIS services: If your network uses NIS services, gather the following information.

Information type	Description
NIS domain name	The name of your NIS domain. The gateway can use an NIS domain to authenticate users and host computers. The NIS domain name must be identical on both gateways in a cluster.

Information type	Description	
NIS servers	The host names of your preferred NIS servers.	
	If...	Then...
	Your site uses NIS	Record the host names of your NIS servers.
	You want NIS to broadcast to find a server	Enter an asterisk (*).

Remote management controller: If your gateway has a remote management controller (RMC) for enhanced Autosupport, gather the following information.

Information type	Description
MAC address	If you are using DHCP service, record the MAC address for the RMC. Locate the MAC address by entering <code>sysconfig -v</code> at the gateway prompt and looking at the output for the RMC in slot 1.
IP address	If you are not using DHCP service, record the IP address for the RMC.
Network mask	If you are not using DHCP service, record the network mask for the RMC.
Gateway	If you are not using DHCP service, record the IP address for the RMC gateway.

Information type	Description
Media type	The type of interface for the RMC: ◆ 100tx-fd (100Base-TX, full-duplex) ◆ tp-fd (10Base-T, full-duplex) ◆ 100tx (100Base-TX, half-duplex) ◆ tp (10Base-T, half-duplex) ◆ auto (autonegotiate speed, duplex, and flow control)
Mail host	The name or IP address of the preferred mail host. The mail host delivers RMC alerts and local Autosupport e-mail.

CIFS protocol

If your gateway is licensed for the CIFS protocol, gather the following information. This information is not necessary for the initial gateway setup because that prepares the gateway for network operation. A subsequent CIFS setup is required to configure the gateway for CIFS operations.

Information type	Description
Windows domain	The name of your Windows domain. If your site uses Windows servers, it has one or more Windows domains. Record the name of the Windows domain to which the gateway should belong. Note_____ The Windows domain name value does not need to be identical on both gateways in a cluster. You can have each gateway in a cluster exist in a different domain and workgroup from its partner. However, if you have a multiprotocol environment and use UID to SID mapping, the UNIX security information must be compatible between the two domains. _____

Information type	Description
WINS servers	The servers that handle Windows Internet Name Service (WINS) name registrations, queries, and releases. If you choose to make the gateway visible through WINS, you need to record up to four WINS IP addresses. Note_____ The WINS server value does not need to be identical on both gateways in a cluster. You can have each gateway in a cluster exist in a different domain and workgroup from its partner. _____
Windows 2000 user	The name of the administrative Windows 2000 domain user who has sufficient privilege to add this gateway to the Windows 2000 domain. Note_____ This is only required if you are using a Windows 2000 domain. _____
Windows 2000 user password	The Windows 2000 domain password to be used for the Windows 2000 user or administrative user. If your domain is a Windows 2000 domain, you must use the password for the user having sufficient privilege to add the gateway to the domain. Note_____ Before you enter the password, be sure to create a secure connection (HTTPS); otherwise, the password is sent to the gateway unencrypted. _____

Information type	Description
Active Directory	The Windows 2000 Active Directory container in which to place gateway accounts. This can be either the default of “Computers” or a previously created organizational unit (OU) specified by you. The path for the OU must be specified in reverse order, separated by commas. Example: If the path is eng\dev\mgmt, the active directory distinguished name is as follows: ou=mgmt, ou=dev, ou=eng Note_____ A user in the Windows 2000 domain can pre-create the account. Data ONTAP updates that account. _____

Where to go from here

What to read next

Chapters in this guide are organized in the order of the tasks you would perform to complete setting up your gFiler. Next go to Chapter 3, “[Providing Setup Information](#),” on page 39.

About this chapter

Before you can install the Data ONTAP software, you must provide information through the boot menu and the setup program—specifically, assigning disk (LUN) ownership and providing information to identify your gFiler on the network. This chapter provides the procedures for how to set up your gFiler before installing the Data ONTAP software.

Note

If you are upgrading your gFiler, see the *gFiler Gateway Series Data ONTAP Upgrade Guide*.

Topic in this chapter

This chapter discusses “[Providing disk \(LUN\) ownership and system identification data](#)” on page 40.

Providing disk (LUN) ownership and system identification data

Information you provide

In the first part of the instructions in this section, you provide information about disk (LUN) ownership. In the second part of the instructions, you provide information that the setup program requires to identify the gFiler to the network.

Disk (LUN) ownership is a logical relationship between a specific gFiler and a specific LUN on a storage subsystem. This logical relationship makes the LUN visible to the gFiler that “owns” it, enabling that gFiler to read data from and write data to the LUN. [“Requirement for enabling gFilers to read from and write to LUNs”](#) on page 14 describes disk ownership in detail.

Providing LUN ownership and system ID data

To provide information for setup, complete the following steps.

Step	Action
1	Power on the gFiler to load the kernel from the CompactFlash™ card. Result: The system boot process begins.
2	Interrupt the boot process by pressing Ctrl-C when you see the following message on the console: Starting Press Ctrl-C for floppy boot menu Result: The following boot options menu appears: (1) Normal boot. (2) Boot without /etc/rc. (3) Change password. (4) No disks assigned (use 'disk assign' from the Maintenance Mode). (5) Maintenance mode boot.
3	Enter 5 for Maintenance mode.

Step	Action
4	Determine the disks (LUNs) on the storage subsystem that you want this gFiler to be configured to own. Enter disk show -v A list of all LUNs available on the storage system port to which the gFiler is physically connected is displayed. Caution LUN security can be configured on a storage subsystem to restrict the LUNs that are available to specific hosts. NetApp recommends that you use LUN security to prevent gFilers from accessing LUNs owned by non-gFiler hosts, and the reverse. If LUN security is not configured, all LUNs mapped to the port on the storage subsystem are visible to all gFilers and non-gFilers connected to the port. In this case you must pay close attention to how you assign disk (LUN) ownership in the next steps so that gFilers do not overwrite LUNs owned by non-gFilers.

Step	Action
5	Assign the LUNs (disks) on the storage subsystem that you want this gFiler to own. You must assign one LUN. You can assign more LUNs at this stage or wait to assign additional LUNs after Data ONTAP is installed. In most cases, you will use the <code>disk assign</code> command to identify specific disks (LUNs) that this gFiler is to own. If this gFiler is to own all the LUNs, you can use the <code>disk assign all</code> command. Caution _____ This step is a destructive operation. Ensure that you are not performing this step on a LUN that is to be used by a host such as a Hewlett-Packard® or Sun® host. Although the gFiler provides controls to prevent one gFiler from owning a LUN that is assigned to another gFiler, there are no controls to prevent you from assigning a gFiler ownership of a LUN owned by a non-gFiler host. Because the gFiler writes ownership information to the LUN when you use the <code>disk assign</code> command, and the non-gFiler host has already configured its version of ownership to the LUN, the LUN becomes corrupted. It is not possible to reverse the corruption. Using LUN security on the storage subsystem will enable you to restrict the hosts that are visible and, therefore, eliminate the possibility of a gFiler overwriting a non-gFiler owned LUN.
	If you are <i>not</i> going to be deploying SyncMirror, enter the command in the form <pre>disk assign fullLUNname1 fullLUNname2 ...</pre> You do not need to enter the pool parameter (-p). The pool parameter defaults to 0 (zero). You can enter all LUNs for a particular storage subsystem on the same command line.

Step	Action
	If you <i>are</i> going to be deploying SyncMirror, you must use the pool parameter with the disk assign command, and you enter the pool parameter correctly. See “Planning for RAID mirroring for data protection” on page 21 for details so that you can plan your configuration correctly. Command example: <pre>disk assign storeAlun1 storeAlun2 -p 0 disk assign storeBlun1 storeBlun2 -p 1</pre> Typically, pool 0 is assigned to the local pool and pool 1 is assigned to the remote pool.
6	Enter the following command: <pre>disk show -v</pre> The Local System ID is shown and the LUNs (disks) exported on the port are shown.
7	Review the list of disks (LUNs) to confirm that all disks (LUNs) that you expected to be assigned to this gFiler (the local system) are shown with the ID of this gFiler.
8	Enter halt, then type boot at the OK prompt.
9	Press Ctrl-C to interrupt the boot process. Result: The following boot options menu appears: <pre>(1) Normal boot. (2) Boot without /etc/rc. (3) Change password. (4) Initialize owned disks (n are owned by this filer). (5) Maintenance mode boot.</pre> <i>n</i> in Option 4 is the number of disks (LUNs) that were assigned in Option 5 for this gFiler to own.
10	Select Option 4. This step creates a root volume with one of the LUNs that you entered in Step 5. The LUN is automatically selected; it is typically the smallest LUN.

Step	Action
11	Enter y when prompted as to whether you want to install a new file system.
12	Enter y to the following message: This will erase all the data on the disks, are you sure? Result: The gFiler creates a root volume using one of the assigned LUNs and begins prompting for system information.
13	When you see the question “Would you like to continue setup through the Web interface?” enter n and continue using the setup command. Answer the prompts in the setup program, using the information that you gathered. See “ Gathering data needed for the setup program ” on page 29 for help with answering the setup prompts.
14	If you want remote console access, answer y when prompted as to whether you want to configure the RMC LAN interface. This prompt is displayed only for the gFiler 900 series models.
15	Install licenses and install and download Data ONTAP software, as described in Chapter 4, “ Installing Data ONTAP Software ,” on page 47.
16	If you have a gateway cluster, repeat the same procedure on the other gateway in the cluster.

Requirement for spare LUNs for a GF270c

Each head of a GF270c must have one spare LUN of 1.5 GB or more assigned to it at all times, in case a problem occurs that results in a panic. The system can then dump core to the spare LUN and NetApp can analyze the core dump to determine the cause of the problem. See “[Managing LUNs and LUN ownership](#)” on page 100 for details about assigning spare LUNs and adding them to NetApp volumes.

Disk show options

The `disk show` commands in the following table might be helpful to you.

disk show option	Information displayed
<code>disk show -v</code>	Shows all disks (LUNs) that are mapped to the physical port on the storage subsystem to which the gFiler is connected. The LUNs might be owned or unowned.
<code>disk show -a</code>	Shows all disks (LUNs) that have an owner.
<code>disk show -n</code>	Shows all disks (LUNs) that do not have an owner.

About this chapter

This chapter provides you with the information you need to install Data ONTAP software on the gateway.

Installation overview

Installing Data ONTAP software is a three-step process:

1. Installing licenses on the gateway.

Note

You must install the gateway storage system license and an NFS or a CIFS license before you can install the Data ONTAP software on a gFiler.

2. Installing system files from the Data ONTAP CD-ROM on a UNIX or Windows host.

The Data ONTAP CD-ROM is part of the Data ONTAP media kit.

3. Downloading the system files to the gateway.

Prerequisite to installing Data ONTAP

Before you start the installation process described in this chapter, ensure that you completed the procedures in Chapter 3, “[Providing Setup Information](#),” on page 39.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Installing licenses on the gateway](#)” on page 48
- ◆ “[Installing the software from a UNIX client](#)” on page 50
- ◆ “[Installing the software from a Windows client](#)” on page 53
- ◆ “[Downloading Data ONTAP](#)” on page 55
- ◆ “[Additional CIFS setup](#)” on page 57

Installing licenses on the gateway

Entering codes for new licenses

Before you can install Data ONTAP, you must install an NFS or a CIFS license. You can install other licenses, such as the license for SyncMirror or SnapMover, before or after you install Data ONTAP.

The procedure you use to install licenses differs slightly depending on whether you are installing licenses on a single gateway or a gateway cluster.

Enabling licenses on a single gateway

To enter new license codes, complete the following steps.

Step	Action
1	Enter the following command on the console: <pre>license add code ...</pre> <i>code</i> is the license code that you were given when you purchased the license for a service or protocol. Example: In this example, the license code for CIFS is ABCDEFG and the license code for gateway is TUVWXYZ. <pre>license add ABCDEFG TUVWXYZ</pre> Result: The gFiler acknowledges each license, similar to that shown in the following example: <pre>A cifs license has been installed. Use "cifs setup" to configure cifs. Run cifs setup to enable cifs. A gateway license has been installed. gateway storage will be enabled after reboot.</pre>
2	Reboot the gateway using the following command: <pre>reboot</pre>

Enabling licenses on a gateway cluster

To enter new license codes on a cluster, complete the following steps.

Step	Action
1	Enter the cluster license code using the following command on both gateway consoles: license add <i>code</i> <i>code</i> is the license code.
2	Reboot both gateways using the following command: reboot
3	Enable clustering on each gateway by entering the following command on each gateway console: cf enable
4	Ensure that clustering is enabled using the following command on each gateway console: cf status Result: Each gateway responds with a message similar to the following: Cluster enabled, gateway 2 is up
5	Repeat Step 1 for other licenses using the <code>license</code> command for each licensed product installed on the cluster.
6	Reboot both gateways of the cluster.

Installing the software from a UNIX client

When to perform UNIX client steps

Perform the steps described in this section only if you are using a UNIX workstation to install Data ONTAP on the gateway.

Procedure overview

This part of the installation involves performing the following procedures in order:

1. Mounting the gateway root volume on your host
2. Copying the system files to your client from the CD-ROM
3. Installing the system files

Note

To perform the installation, you need access to both the gateway’s system console and a client system.

Mounting the gateway on your client

To mount the gateway on your UNIX client, complete the following steps.

Step	Action
1	Make files and directories available for mounting by entering the following command on the console: exportfs -a
2	Create a directory on the client where you want to mount the gateway root file system: mkdir /mnt/gfiler <i>/mnt/gfiler</i> is the directory on the client where you want to mount the gateway root file system.
3	Mount the gateway’s root file system to the clients’s /mnt/gfiler directory: mount gfilername:/ /mnt/gfiler <i>gfilername</i> is the name of the gateway.

50

Installing the software from a UNIX client

Step	Action
4	Change the host directory to <i>/mnt/gFiler</i>. <pre>cd /mnt/gFiler</pre> <i>/mnt/gfiler</i> is the directory on the client where you mounted the gateway's root file system.

Copying the system files from the CD-ROM

To access the system files from the CD-ROM and copy them to your host, complete the following steps.

Step	Action
1	Mount the Data ONTAP CD-ROM onto your host.
2	From the host, extract the files by entering the following command: <pre>tar xvf /CD-ROM_MNT_PT/ONTAP/release/x86/SYSFILES.TAR</pre> <i>/CD-ROM_MNT_PT</i> is the name of your CD-ROM mount point. <i>release</i> is the Data ONTAP release number on the CD-ROM. Result: The <code>tar</code> command extracts the files and reports the names; they are <code>tar_image.Z</code> and <code>install_netapp</code>.
3	Complete the procedure described in “ Installing the system files ” on page 52.

Installing the system files

To install the system files, complete the following steps.

Step	Action	
1	If you want to...	Then...
	Delete the tar files after the files are extracted	From the host, enter the following command: <code>./install_netapp /mnt/gfiler</code> <i>/mnt/gfiler</i> is the directory on the host where you mounted the gateway's root file system.
	Keep the tar files after the files are extracted	From the host, enter the following command: <code>./install_netapp -k /mnt/gfiler</code> <i>/mnt/gfiler</i> is the directory on the host where you mounted the gateway's root file system.
	Result: The install script decompresses and extracts the files.	
2	From the client, unmount the gateway's root file system (/) by entering the following commands: <code>cd /</code> <code>umount /mnt/gfiler</code> <i>/mnt/gfiler</i> is the directory on the client where you mounted the gateway's root file system.	
3	Complete the procedure in “ Downloading Data ONTAP ” on page 55 to download the boot data on the gFiler.	

Installing the software from a Windows client

When to perform Windows client steps

Perform the steps described in this section only if you are using a Windows PC to install Data ONTAP on the gateway.

Procedure overview

This part of the installation involves performing the following procedures in order:

1. Mapping the gateway to a drive on your host
2. Copying system files to your host from the CD-ROM and installing them

Mapping the gateway to a drive

To map the root directory of your gateway to a drive on your host, complete the following steps.

Step	Action
1	Log in to your host as Administrator or log in using an account that has full control on the gateway C\$ share.
2	Map a drive to the C\$ directory of your gateway according to the procedure for your version of Windows.

Copying and installing the system files from the CD-ROM

To access the system files from the CD-ROM and copy them to your host, complete the following steps.

Step	Action
1	Insert the Data ONTAP CD-ROM into your CD-ROM drive.
2	Go to the following folder on the Data ONTAP CD-ROM: ONTAP\release\X86 <i>release</i> is the release number on the CD-ROM.

Step	Action
3	Double-click setup.exe. Result: The setup program displays a dialog box.
4	Make sure that you followed the steps described in the dialog box, then click OK. Result: The WinZip® dialog box appears.
5	Follow the WinZip prompts to unzip the file that you downloaded onto the drive to which you mapped the gateway. When appropriate, ensure that the following check boxes are selected: ◆ Overwrite Files Without Prompting ◆ When Done Unzipping Run... Leave the other options as they are. Result: In the dialog box's lower panel, confirmation messages are displayed as files are decompressed.
6	Complete the procedure in "Downloading Data ONTAP" on page 55 to download the boot data on the gFiler.

Downloading Data ONTAP

Downloading Data ONTAP

To download Data ONTAP, complete the following steps.

Step	Action
1	At the gateway console, enter the following command to copy the executables to the boot blocks of the gateway: download Result: The download command provides an acknowledgment similar to the following: <pre>download: You can cancel this operation by hitting Ctrl-C in the next 6 seconds. download: Depending on system load, it may take many minutes download: to complete this operation. Until it finishes, download: you will not be able to use the console. Fri Mar 26 15:04:06 PST [raid-u20: download.request:notice]: Operator requested download initiated download: Downloading boot device download: Downloading boot device (Service Area) download: Downloading disks: opening netapp_6.5.1 done download: boot blocks successfully copied to disk(s) raid-u20> Fri Mar 26 15:09:25 PST [raid-u20: download.requestDone:notice]: Operator requested download completed</pre>
2	At the gateway console, enter the following command: reboot
3	As needed, perform the procedures described in “Additional CIFS setup” on page 57.

Step	Action
4	If this is a gateway cluster, repeat this procedure on the partner gateway.

Additional CIFS setup

About concluding the installation

If you installed a CIFS license, you need to perform an additional step from either a UNIX or a Windows host. Otherwise, you are finished with the installation of the Data ONTAP software.

Setting up the gateway to use the CIFS protocol

To set up the gateway to use the CIFS protocol, complete the following step.

Step	Action
1	Enter the cifs setup command and use the CIFS information you gathered to answer questions posed by the prompts.

About this chapter This chapter discusses the verification and configuration processes you do to complete the setup.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Setup completion process”](#) on page 60
- ◆ [“Setting up host name resolution”](#) on page 63
- ◆ [“Verifying network connections”](#) on page 64
- ◆ [“Making gateway data available to NFS hosts”](#) on page 66
- ◆ [“Making gateway data available to CIFS hosts”](#) on page 69
- ◆ [“Verifying configuration for clustered gateways”](#) on page 71
- ◆ [“Configuring networking interfaces for clustered gateways”](#) on page 75
- ◆ [“Checking and testing the configuration for clustered gateways”](#) on page 79
- ◆ [“Verifying the RMC connections and setup”](#) on page 84
- ◆ [“Setup if your deployment includes a shared switch”](#) on page 86

Setup completion process

Verification and configuration processes

The following table describes the complete process for verifying that your gateway is communicating on the network, that it is online, and that it is ready to operate.

To make the gateway ready to use, complete the following steps.

Stage	What you do	For details
1	Set up host name resolution.	See “Setting up host name resolution” on page 63.
2	Verify the gateway’s network connection.	See “Verifying network connections” on page 64.
3	Verify that the gateway is online and has the appropriate volumes and licenses.	After initial setup, the gateway is online, and the following entities exist on the gateway: /vol/vol0 (a virtual root path) /vol/vol0/home (a directory) Note _____ /vol is not a directory—it is a special virtual root path under which the gateway mounts its volumes. You cannot mount /vol to view all the volumes on the gateway; you must mount each gateway volume separately. _____ For NFS: /vol/vol0 is exported to the administration host for root access; /vol0/home is exported to the administration host for root access and to all hosts for general access. For CIFS: /vol/vol0 is shared as C\$; only the system administrator with the root password has read and write access to the C\$ share. The /vol/vol0/home directory is shared as HOME without access granted to anyone.
4	For NFS, configure hosts to access the gateway.	See “Making gateway data available to NFS hosts” on page 66.
5	For CIFS, configure hosts to access the gateway.	See “Making gateway data available to CIFS hosts” on page 69.

Stage	What you do	For details
6	For HTTP, configure HTTP service.	If you purchased a license for the HTTP protocol for your gateway, turn on HTTP service by using the FilerView interface or the console commands described in the <i>File Access Management Guide</i> .
7	For CIFS, set up password and group files.	If you purchased a license for the CIFS protocol and you are not using a Windows NT or 2000 domain controller to authenticate users, configure password and group files according to the instructions in the <i>Storage Management Guide</i> .
8	Configure CIFS shares.	If you purchased a license for the CIFS protocol, configure CIFS shares so that users can access directories on the gateway according to the instructions in the <i>File Access Management Guide</i> .
9	Assign access rights to CIFS users and groups.	If you purchased a license for the CIFS protocol, assign access rights by using the instructions in the <i>File Access Management Guide</i> .
10	Check the configuration of clustered gateways.	If you configured clustered gateways, check the cluster configuration information according to “ Verifying configuration for clustered gateways ” on page 71.
11	Configure network interfaces for clustered gateways	If you configured clustered gateways, see “ Configuring networking interfaces for clustered gateways ” on page 75.
12	Check the RMC connections.	See “ Verifying the RMC connections and setup ” on page 84.
13	Configure the gFiler to control the number of initiators accessing a port on a storage subsystem.	See “ Setup if your deployment includes a shared switch ” on page 86.

Stage	What you do	For details
14	Install electronic documentation.	1. Obtain the documentation in zip file format from the CD-ROM: <i>drive:\doc\docs.zip</i>. 2. Use your browser to open the page at <i>http://filer_name/na_admin</i>. 3. Select the Install Documentation link. 4. Follow the prompts to locate and install the documentation zip file.

Setting up host name resolution

Interfaces host names require resolution

When you ran the setup program, the gateway generated a host name for each interface by appending the number of the interface to the gateway host name. You need to make sure that these host names are resolved into IP addresses.

For example, the interface name for the first interface on a gateway named “toaster” might be toaster-e0; the second interface might be toaster-e1.

Host name resolution

To ensure that host names are resolved into IP addresses, complete the following step.

Step	Action	
1	If you use...	Then...
	DNS or NIS for name resolution	Add an entry in your DNS or NIS databases for each of the gateway interfaces. Example: This shows how the entries might look for a gateway with four interfaces: 192.16.3.145 toaster-e0 192.16.3.145 toaster-e1 192.16.3.145 toaster-f0 192.16.3.145 toaster-a5
	/etc/hosts files for name resolution	Add an entry in each host’s /etc/hosts file for each of the gateway interfaces.

Verifying network connections

Verifying gateway connections to the network

To verify gateway connections to the network, complete the following steps.

Step	Action	
1	If you want to use the ping command to verify the...	Then...
	Host name interface	Complete Steps 2 and 3.
	IP address of the interface	Complete Steps 4 and 5.
2	At a host attached to the network, enter the following command: ping <i>hostname-interface</i> <i>hostname</i> is the host name that you assigned to the gateway when you ran the setup program. <i>interface</i> is one of the interface names that the gateway assigned when you ran the setup program. Example: This command tests the network connections for a gateway that has the host name “toaster” with an interface named e0 installed: ping toaster-e0	
3	Repeat the test once for each interface that is installed in the gateway.	

Step	Action
4	At a host attached to the network, enter the following command: ping <i>hostname-IP_address</i> <i>hostname</i> is the host name that you assigned to the gateway when you ran the setup program. <i>IP_address</i> is the IP address that the gateway assigned to that interface during the setup program. Example: This command tests the network connections for a gateway that has the host name “toaster” with an interface named e0 installed at 172.25.50.10” ping 172.25.50.10
5	Repeat the test once for each interface that is installed in the gateway.

Troubleshooting problems

To troubleshoot a problem when an interface does not respond to a ping command, complete the following steps.

Step	Action
1	Check to make sure that the interface cable is securely attached to the gateway and the switch or router, and that the connection and activity LEDs are lit.
2	Check to make sure that the media type is set up correctly if the interface is using a multiport Ethernet card with different port speeds.
3	Check to make sure that the routers function properly with correct routing information if the ping command is issued from a network not directly attached to the interface.

Making gateway data available to NFS hosts

Tasks to make gateway data available

To make gateway data available to NFS hosts, you need to do the following tasks:

- ◆ Export the gateway’s file system
- ◆ Mount the gateway’s file system to the NFS host

Exporting NFS hosts

You must export the gateway’s file systems before NFS hosts can mount them. To export the gateway’s file systems, complete the following steps.

Step	Action
1	Determine the directories that you want to export on various volumes. For example, valid path names for directories entered at the gateway prompt could be: /vol2/eng/hardware /vol2/eng/software
2	Use a text editor from an NFS host to open the gateway /etc/exports file.
3	Add the gateway directories to the /etc/exports file to export them. For information about specifying entries and access restrictions, see the chapter about file access using NFS in the <i>Data ONTAP System Administrator File Access Management Guide</i> . Example: /vol/vol0/home -access=10.1.1.20,root=10.1.1.20 /vol2/eng/hardware /vol2/eng/software
4	Save the file and exit the text editor.

Mounting NFS host
file systems

You can mount the gateway’s file system by using one of the following methods:

- ◆ Add entries for the directories exported from the gateway in each host’s /etc/vfstab file (for Solaris® hosts) or /etc/fstab file (for non-Solaris hosts).
- ◆ Use the UNIX mount command.
- ◆ Use the UNIX automounter tool.

Note_____

See your UNIX host’s man pages for information about the automounter tool.

Adding entries to the /etc/vfstab file: To mount the gateway’s file systems by adding entries to the /etc/vfstab file, complete the following steps.

Step	Action
1	Use a text editor from an NFS host to open the /etc/vfstab (or /etc/fstab) file.
2	Add the entries for the directories exported from the gateway in each host’s /etc/vfstab file (or /etc/fstab) file. Example: toaster:/vol/vol0/home nfs - yes - rw,bg,hard toaster:/vol2/eng/hardware toaster:/vol2/eng/software Note_____ The /vol/vol0/home directory is used as a sample gateway directory in this procedure. You should use a path that exists on your gateway. _____
3	Save the file and exit the text editor.
4	Reboot the hosts. Result: The hosts mount the file systems automatically when they boot.

Using the UNIX mount command: To mount the gateway's directory on the NFS host, complete the following steps.

Step	Action
1	At the NFS host, create a directory to act as a mount point, for example: mkdir /n/toaster/home
2	To mount the gateway's /vol/vol0/home directory to /n/toaster/home, add the following line to the NFS host's /etc/vfstab file (or /etc/fstab file) if the host is on the same network as the gateway's e0 Ethernet interface: toaster-e0:/vol/vol0/home /n/toaster/home nfs rw 0 0 Result: The file system is mounted when the host reboots.
3	To mount the /vol/vol0/home directory without a host reboot, enter the following command at the NFS host: mount toaster-e0:/vol/vol0/home /n/toaster/home Result: After the /vol/vol0/home directory on the gateway is mounted, the user on the NFS host can create, modify, or remove files in /n/toaster/home.
4	Repeat Step 1 through Step 3 for each NFS host that requires access.

Making gateway data available to CIFS hosts

How gateway data is visible to CIFS hosts

After setup is complete, the gateway becomes visible to CIFS hosts by automatically registering with the master browser on its local network. If cross-subnet browsing is configured correctly, the gateway is now visible to all CIFS hosts. For more information about cross-subnet browsing, see Microsoft networking documentation.

CIFS host connection methods

The following table describes the methods by which CIFS hosts can connect to the gateway. For more information about commands and windows mentioned in the table, see the operating system manual for your computer.

Note—
Most CIFS hosts can access the gateway using more than one method. For example, a Windows 2000 host can access the gateway by configuring a network drive or by clicking the gateway icon in the My Network Places window.

Method	Host operating system	Command or window
1. Connecting to the gateway as a network drive	Windows 95, Windows 98, Windows Me	Map Network Drive window—limited by Workgroup or domain membership
	Windows NT 4.0	Map Network Drive window (from NT Explorer)
	Windows 2000, Windows XP host	Map Network Drive
	Windows NT 3.51, Windows for Workgroups	Connect Network Drive window (from the File Manager)
2. Clicking the gateway icon	Windows 95, Windows 98, Windows Me, Windows NT 4.0	Network Neighborhood
	Windows Me, Windows 2000, Windows XP	My Network Places

Method	Host operating system	Command or window
3. Mapping a share to a drive letter on the command line	Operating systems that support the DOS command-line interface and the net use command, including Windows 2000 and XP	net use command Example: net use f:\\toaster\\home

Informing users how to access the gateway

To inform users about how to access the gateway, complete the following steps.

Step	Action	
1	Determine the connection methods that work for your CIFS hosts.	
2	If your CIFS hosts can...	Then...
	Use methods 1 and 3 to connect to the gateway	Inform your users of the gateway's UNC path, which consists of a computer name and a share name in the following format: <i>\\computer_name\share_name</i> Example: If the gateway is named toaster and the directory that contains users' directories is shared with the name "home," the UNC path is \\toaster\home. Result: Users can connect to the gateway, view directories, and read and write files.
	Use method 2 to connect to the gateway	Inform your users of the gateway name so that they can locate the gateway in Network Neighborhood.
3	Have users execute the appropriate procedure connection method, as described in " CIFS host connection methods " on page 69.	

Verifying configuration for clustered gateways

Values that need to match

When you configure your clustered gateways, the following configuration information needs to match on both nodes of a cluster.

- ◆ Configuration files
- ◆ Licenses
- ◆ Options
- ◆ Parameters

You can use the following tables to ensure that both gateways have the same values in their configuration after you configure them.

Configuration file matching

Use the following table to indicate you verified that the contents of the configuration files in the root volume of the gateway (not the UNIX administration host) are identical for both gateways.

File	Gateway 1	Gateway 2
/etc/resolv.conf		
/etc/httpd.mimetypes		
/etc/dgateways		
/etc/nsswitch.conf		

License matching

Use the following table to record the codes for each of the licenses for both gateways in the cluster to ensure that both gateways have their own unique licenses.

Note

Some licenses are not supported by the gateways. Such licenses do not appear in the following table.

License	Gateway 1 license code	Gateway 2 license code
nfs		
cifs		
cluster		
cluster_remote		
http		
gateway		
gateway_hitachi		
snapmanagerexchange		
snapmirror		
snapmover		
snaprestore		
syncmirror_local		
sv_linux_pri		
sv_ontap_pri		
sv_ontap_sec		
sv_unix_pri		
sv_windows_pri		
vfiler		
vld		

Option matching

Use the following table to compare the output of the `options` command for both gateways in the cluster to ensure that the options have identical values in both gateways. If an option is not identical, see [“Manually checking and correcting options and licenses”](#) on page 81 for information about correcting the option.

Option	Gateway 1	Gateway 2
<code>cifs.show_snapshot</code>		
<code>dns.domainname</code>		
<code>dns.enable</code>		
<code>http.admin.enable</code>		
<code>httpd.timeout</code>		
<code>httpd.timewait.enable</code>		
<code>ip.path_mtu_discovery.enable</code>		
<code>nfs.per_client_stats.enable</code>		
<code>nfs.v2.df.2gb.lim</code>		
<code>nfs.v3.enable</code>		
<code>nfs.v4.enable</code>		
<code>nis.domainname</code>		
<code>nis.enable</code>		
<code>pcnfsd.enable</code>		
<code>raid.reconstruct_speed</code>	Not applicable	Not applicable
<code>raid.timeout</code>	Not applicable	Not applicable
<code>snmp.enable</code>		
<code>telnet.enable</code>		
<code>timed.enable</code>		
<code>timed.log</code>		
<code>timed.max_skew</code>		
<code>time.min_skew</code>		

Option	Gateway 1	Gateway 2
timed.proto		
timed.sched		
timed.servers		
trusted.hosts		
vol.copy.throttle		
wafl.maxdirsize		
wafl.root_only_chown		

Parameter matching Use the following table to compare the parameters for both gateways in the cluster to ensure that the parameters have identical values in both gateways.

Parameter	Gateway 1	Gateway 2
ARP table		
date		
ndmp		
nfs		
route table		
routed		
timezone		

Configuring networking interfaces for clustered gateways

What the networking interfaces do

When a node in a cluster fails, the surviving node must be able to assume the identity of the failed node on the network. Networking interfaces allow individual nodes in the cluster to maintain communication with the network if the partner fails.

Network connection requirements

Both nodes in the cluster must meet the following network interface connection characteristics:

- ◆ Both nodes must have interfaces that access the same collection of networks and subnetworks.
- ◆ Both nodes must have interfaces that can take over the IP addresses of each other.
- ◆ Both nodes must have interfaces that can take over the MAC (media access control) addresses of each other.

Types of takeover configurations

The following table lists the three types of interface configurations that can be enabled in a cluster.

Interface type	Description
Shared	This type of interface supports both the local and partner nodes. It contains both the local node and partner node's IP addresses. During takeover, it supports the identity of both nodes.
Dedicated	This type of interface only supports the node in which it is installed. It contains the local node's IP address only and does not participate in network communication beyond local node support during takeover.
Standby	This type of interface is on the local node, but only contains the IP address of the partner node. During takeover, this interface supports a virtual filer identity. See the <i>Data ONTAP System Administrator's MultiStore Management Guide</i> for more information about virtual filers.

Information needed for configuration

You must gather the following information before configuring the interfaces:

- ◆ The IP address for both the local node and partner node
- ◆ The netmask for both the local node and partner node
- ◆ The Ethernet addresses for both the local node and partner node

Configuring networking

To configure the networking interfaces on a cluster, complete the following steps.

Note

If you configured a gateway cluster when you set up the gateway, you already performed this procedure.

Step	Action
1	Enter the following command at the local node console: setup
2	Answer the questions displayed on the screen. Please enter IP address for Network Interface ns0 [n.n.n.n]: Please enter netmask [255.255.255.0]: Should interface ns0 take over a partner IP address during failover? [y]: Please enter the IP address to take over: [m.m.m.m]:
3	Enter the following command to reboot the node: reboot
4	Repeat Steps 1 through 4 for the partner node.

Configuring interfaces

To configure an interface as a dedicated, standby, or shared interface, complete the following steps.

Step	Action
1	Check the current settings of the local node by entering the following command on the console: ifconfig -a Result: You should see output similar to the following: <pre>e0: Flags=40043<UP,BROADCAST,RUNNING,UP_1ARY> mtu 1500 inet 11.22.33.44 netmask ffffffff00 broadcast 11.22.33.45 secondary host inet 11.22.33.70 (not in use) ether 00:01:00:00:00:00:80 (Data Link Provider Interface)</pre>

Step	Action	
2	If you are setting the interface as a...	Then...
	Dedicated interface	Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig e0 local_node_ip_address Repeat Step 1 for the partner node.
	Shared interface	Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig NIC_name ip_address partner \ ip_address Example: ifconfig e0 199.9.204.255 \ partner 199.9.204.254 Next enter the following command at the prompt, or directly into the /etc/rc file of the partner node: ifconfig NIC_name ip_address partner \ ip_address Note _____ You do not have to enter the MAC address. _____
	Standby interface	Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig e1 local_node_ip_address Repeat Step 1 for the partner node.

Checking and testing the configuration for clustered gateways

Methods of checking cluster configuration

- There are two ways you can check your cluster configuration prior to placing your cluster online:
- ◆ Running the Cluster Configuration Checker script—The script runs under a UNIX shell or as a Common Gateway Interface (CGI) script from a UNIX Web server. It detects the following configuration problems:
 - ❖ Mismatched licenses
 - ❖ Mismatched option settings
 - ❖ Incorrectly configured network interfaces
 - ❖ Different versions of Data ONTAP on the local and partner gateways
 - ◆ Manually, through the command line

The Cluster Configuration Checker is available on the Tools & Utilities page on NOW.

Information you need

The following tables shows the requirements that you need to use the cluster configuration checker

UNIX shell	CGI
Names or addresses of each partner of the NetApp cluster	
Login names and passwords on both gFilers	
UNIX computer that can access both gFilers (you can use the ping command to confirm this)	◆ Web server that can access both gFilers ◆ A place to store the Perl script where the Web server can find it

Using the Cluster Configuration Checker—UNIX shell

To run the Cluster Configuration Checker script from a UNIX shell prompt, complete the following steps.

Step	Action
1	Execute the cf-config-check.cgi script from a Unix shell prompt, with the name or IP address of both gFilers on the command line.
2	Add the word verbose to the end of the command line to view more detail.
3	When prompted, enter the rsh login names and passwords of both gFilers.
4	Examine the script output and correct any errors or differences identified in the output.

Using the Cluster Configuration Checker—CGI

To run the Cluster Configuration Checker as a CGI script, complete the following steps.

Step	Action
1	Point your browser to the URL corresponding to the location of the Perl script.
2	Enter the names or IP addresses of the two gFilers.
3	Enter the rsh login names and passwords.
4	Click Check.
5	If no results are displayed, click your browser's Stop button.
6	Examine the script output and correct any errors or differences identified in the output.

Manually checking and correcting options and licenses

To check options manually on the local and partner gateways, complete the following steps.

Step	Action
1	View the options set on the local gateway by entering the following command at the gateway's console: options Result: The current option settings on the gateway are listed. You should see output similar to the following: <pre> autosupport.doit DONT autosupport.enable on . . dns.domainname gfiler1 (might be overwritten in takeover) dns.enable off (might be overwritten in takeover) ftpd.enable off . . </pre>
2	Note the values of these options for the local gateway, and repeat Step 1 on the partner gateway.
3	Verify that the options that have comments in parentheses are set to the same value for both gateways. The comments you see are as follows: ◆ Value might be overwritten in takeover ◆ Same value required in local+partner ◆ Same value in local+partner recommended
4	On either gateway console, correct any mismatched options using the following command: options <i>option_name</i> <i>option_value</i>
5	Check the license settings on both gateways using the following command: license
6	Correct mismatched license settings using the following command: license <i>code</i>

Testing takeover and giveback

To test takeover and giveback, complete the following steps.

Step	Action	
1	Check the cabling on the interconnect cables to make sure that they are secure.	
2	Verify that you can create and retrieve files on both nodes for each licensed protocol.	
3	From the local node, enter the following command and observe the output: cf enable Result: Takeover of the partner node is enabled on the local node.	
4	Enter the following command from the local node console: cf takeover Result: The local node takes over the partner node and gives the following output: takeover completed	
5	Make sure that the local node can access both the local and partner disks.	
6	Give back the partner node by entering the following command: cf giveback Result: The local node releases the partner node, which reboots and resumes normal operation. The following message is displayed on the console when the process is completed: giveback completed	
7	If takeover and giveback...	Then...
	Is completed successfully	Repeat Steps 2 through 6 on the partner node.
	Fails	Go to “Correcting takeover or giveback failure” on page 83.

Correcting takeover or giveback failure

The following table describes tests you can run to determine why takeover or giveback failed and also provides information to correct the problem.

Note

You must reboot the cluster after completing each test.

To correct a takeover or giveback failure, complete the following steps.

Step	Action	
1	Check communication between the local and partner nodes by entering the following command and observing the messages: cf status Results: If there is a problem with cabling or the adapter, the appropriate message is displayed.	
2	If the error message indicates...	Then...
	A cluster adapter error	Check the cluster adapter cabling. Make sure that the cabling is correct and properly seated.
	A cabling error	Check the cabling and reseal and tighten any loose cables.
	A networking error	Check for network connectivity. See the <i>Data ONTAP System Administrator's Guide</i> for more information.
3	Run the cluster configuration checker script or check the cluster configuration manually. Correct any errors or differences displayed in the output.	
4	Reboot the cluster and rerun the takeover and giveback tests.	
5	If clustering is still not enabled, call NetApp Technical Support or your authorized reseller.	

Verifying the RMC connections and setup

Verifying the RMC connections and setup

To verify the remote management controller connections and setup, complete the following steps.

Step	Action
1	Enter the following command: <pre>rmc setup</pre> Note _____ It might take a few minutes for the new network settings for the RMC to take effect.
2	Enter the following command to verify the connection from the gateway to the RMC interface: <pre>ping rmc interface</pre> <i>rmc interface</i> is the IP address of the RMC.
3	If Autosupport is to be used in your organization, verify that Autosupport is enabled and Autosupport options are valid by entering the following command: <pre>options autosupport</pre> Result: The Autosupport options are listed similarly to the following: <pre>autosupport.enable on autosupport.support.enable on autosupport.mailhost <i>name or IP address of mailhost</i></pre>
4	To test RMC mail delivery, enter the following command: <pre>rmc test email TESTMSG</pre> Result: The RMC should send e-mail within a few minutes. If this fails, verify gateway connectivity and whether the mail host is valid.

About the RMC network interface

The RMC network interface is not used for serving data, so it does not show up in the output for the `ifconfig` command.

Setup if your deployment includes a shared switch

Requirement to limit the target queue length limitation

If your deployment includes a shared switch, you need to plan for and configure the gFilers to limit the number of commands (read and write requests) that the storage subsystem port must handle as a result of gFilers and non-gFiler hosts accessing the port through the switch; that is, you need to plan for limiting the target queue length. When multiple initiators are accessing a target port on the storage subsystem, you do not want the outstanding commands in the queue buffer, from all initiators together, to exceed 256. Otherwise, the performance of gFiler and non-gFiler hosts can suffer.

On gFilers, you use the `disk.target_port.cmd_queue_depth` option to set the number of commands that the gFiler can send to the target port on the data subsystem. You do not set this option until after you set up and install Data ONTAP on your gFiler.

Non-gFiler hosts also provide a means for limiting the target queue length. See the documentation for the host for information about limiting target queue length.

Note

Target queue length might also be referred to as Q-Depth or Max Throttle.

Guidelines for determining the appropriate target queue length

You need to consider the impact of all the initiators on the port when you are planning the configuration for a specific gFiler or a specific non-gFiler host. If your deployment includes more than one initiator, you need to configure each initiator so that the total commands by all initiators does not exceed 256. The default for target queue depth value is 256.

Some guidelines for determining the appropriate target queue length are as follows:

- ◆ 0 (zero) means no limit on the outstanding commands. Therefore, do not configure a value of 0.
- ◆ You could divide 256 by the number of gFiler and non-gFiler hosts that will be acting as initiators on the target port on the storage subsystem, configuring each gFiler and non-gFiler host with the value that resulted.

- ◆ You could consider the volume of commands that specific initiators would be likely to send to the target port. You could then configure higher values for initiators likely to send a greater number of requests and a lower value for initiators likely to send a lesser number of requests.

Configure non-gFiler hosts according to the guidelines provided for those hosts.

Setting the target queue depth length

To set the `disk.target_port.cmd_queue_depth` option, complete the following step.

Step	Action
1	Enter the following command at the Data ONTAP command line: <code>disk.target_port.cmd_queue_depth <value></code> Do not use a value of 0 (zero) because it means no limit on the outstanding commands.

Where to go from here

What to read next

After completing this chapter, you have finished initial setup and installation of software on your gFiler. Next read

- ◆ Chapter 6, “[Managing Volumes, RAID Groups, and LUNs](#),” on page 89 to learn about managing volumes, RAID groups, and LUNs on your gFiler on an ongoing basis.
- ◆ Chapter 7, “[Managing Workload Distribution with SnapMover](#),” on page 109 to learn about using SnapMover to move the workload of processing requests from one gFiler to another. You must set up your system ahead of time to support SnapMover if you want to use this feature.

About this chapter

This chapter provides information to help you perform gFiler-specific tasks after initial setup that are associated with volumes, RAID groups, and LUNs.

Use this chapter in conjunction with the Data ONTAP documentation set.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Volume configuration and management”](#) on page 90
- ◆ [“Management of RAID groups associated with NetApp volumes”](#) on page 96
- ◆ [“Managing LUNs and LUN ownership”](#) on page 100

Volume configuration and management

About volumes

In Data ONTAP, *volumes* are the primary means of organizing and managing data. You can think about a NetApp Data ONTAP volume as a container for storing data. You identify the collection of LUNs that are to be in the volume (where collection could mean only one LUN); those LUNs provide the capability for data storage. You manage data storage by managing volumes through the Data ONTAP system administration interface.

During the initial setup process, a root volume is automatically created (see “[Planning for disk \(LUN\) ownership assignment](#)” on page 19). However, you can create multiple volumes to facilitate data management. For example, you could create additional volumes to accommodate your users if they speak more than one language, such as one volume for data in French and another for data in German.

The *Data ONTAP System Administration Storage Management Guide* provides more details about procedures for managing volumes than are provided in this section. This section provides specific information about volumes with the gFiler.

Volume metrics

The following table summarizes volume metrics for gFilers.

Metric	Information
Maximum number of volumes on a gFiler	200 volumes for both clustered and single gFilers
Minimum root volume size	72 GB (72 GB is required to hold core files and log files that might be generated.)
Minimum size for a non-root volume	1 GB (the same as the minimum size for a LUN)
Maximum volume size	This is dependent on the gFiler model. See the <i>gFiler Gateway Series Getting Started Guide</i> .
Number of LUNs in a volume	See “ Guidelines for determining volume size ” on page 91.

Metric	Information
Maximum number of RAID groups in a volume	150 RAID groups. See “Management of RAID groups associated with NetApp volumes” on page 96 for information about the role of RAID groups with volumes. Caution————— If you exceed the maximum volume size or maximum capacity for a gateway model, the maximum RAID groups in a volume is less than 150.

Guidelines for determining volume size

To determine the volume size, you must first determine the amount of space that you need to store data. Then specify the number of LUNs that you need for the capacity and performance that you require. Factors to consider are as follows:

- ◆ Capacity
 - How many LUNs you need to specify to support the storage capacity of the volume:
 - ❖ Consider the size of LUNs that were created on the storage subsystem.
 - ❖ Consider that the more LUNs you add to a volume, the more space is available.
 - ❖ Consider the maximum volume size for your gFiler model. See the *gFiler Gateway Series Getting Started Guide* for more information.
 - ❖ Consider the total number of LUNs for your gFiler model. See the *gFiler Gateway Series Getting Started Guide* for more information.
- ◆ Performance
 - How many LUNs are required for adequate I/O performance

Moving or deleting LUNs

After a LUN becomes part of a volume, the LUN cannot be moved to another volume or deleted. You can only move volumes, destroy volumes, or expand volumes. If you destroy a volume, all data in the volume is destroyed and is no longer accessible. See Chapter 7, “[Managing Workload Distribution with SnapMover](#),” on page 109 for more information about moving volumes by using SnapMover.

Should you add LUNs to a volume or create additional volumes?

If a volume has not reached its maximum number of LUNs, you can add LUNs to a volume.

If the gFiler volume was fully allocated with user data before you added LUNs to the volume, most of the free space for write actions is obtained from the new LUNs after the new LUNs are added. Therefore, add enough LUNs at the same time so that the system does not become spindle bound for the write actions. In general, adding two or more LUNs to a volume at the same time should satisfy this requirement.

Other than the performance issue just mentioned, there are no performance penalties that result from expanding an existing volume.

There is no performance penalty for expanding a volume with a different-sized LUN, as long as the new RAID group on the gFiler is created with the newly added LUN. See [“What to do if you are expanding a volume with a different-sized LUN”](#) on page 97.

Importance of volume names for migration

Consider whether you might need to move a volume in the future. If so, take care that names of volumes on the different gFilers are different. If you need to move volumes with vFilers later through SnapMover, the names of the vFilers and volumes being moved from the source to the destination must be unique. See [“Planning for using SnapMover to change workload distribution”](#) on page 110 for more information

Creating a volume

To create a volume, complete the following steps.

Note

See the *Data ONTAP Storage Management Guide* for details about creating mirrored volumes.

Step	Action
1	View a list of the spare LUNs on your NetApp appliance. At the gFiler console, enter the following command: <pre>vol status -s</pre> Result: The output of <code>vol status -s</code> lists all the spare LUNs that you can select for inclusion in the volume and their capacities.

Step	Action
2	At the system prompt, enter the following command: <pre>vol create vol-name [-r raid-group-size] [-l language_code] disk-list</pre> <i>vol-name</i> is the name for the new volume (without the /vol/ prefix). <i>raid-group-size</i> is the size of the RAID group you want. <i>language_code</i> specifies a language other than that of the root volume. See the <i>Data ONTAP Storage Management Guide</i> for a list of language codes. <i>disk-list</i> is one of the following: ◆ <i>ndisks[@disk-size]</i> <i>ndisks</i> is the number of LUNs in the volume. The value of <i>ndisks</i> must be at least 1 for unmirrored volumes and 2 for mirrored volumes. <i>disk-size</i> is the LUN size to use, in gigabytes. You must have at least <i>ndisk</i> available LUNs of the size you specify. Data ONTAP will select LUNs within a plus or minus 20 percent range of the size you enter. ◆ -d disk_name1 disk_name2... disk_nameN <i>disk_name1</i>, <i>disk_name2</i>, and <i>disk_nameN</i> are IDs of one or more LUNs that the gFiler owns; use a space to separate multiple LUNs.
3	Enter the following command to verify that the volume exists as you specified: <pre>vol status vol_name -r</pre> <i>vol_name</i> is the name of the volume whose existence you want to verify. Result: The system displays the RAID groups and LUNs of the specified volume on your gFiler.

Step	Action	
4	If you have...	Then...
	An NFS-only gFiler	Go to Step 6.
	A CIFS-only or multiprotocol gFiler	Create shares that refer to the new volume to enable clients to access it, then go to Step 5.
5	If you have...	Then...
	A CIFS-only gFiler	You are done.
	A multiprotocol gFiler	Go to Step 6.
6	Update the system /etc/exports file with the new volume.	
7	Run <code>export tfs</code> .	
8	Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the gFiler.	

Volume creation examples

Example 1: The following command creates a volume called `newvol`, with no more than eight LUNs in a RAID group, using the French character set, consisting of the LUNs with IDs `8a.0L4`, `8a.0L2`, and `8a.0L5`.

```
vol create newvol -r 8 -l fr -d 8a.0L4 8a.0L2 8a.0L5
```

Example 2: The following command creates the volume `newvol`. The RAID group size is 14 LUNs. Because the volume consists of 28 LUNs, those LUNs form two RAID groups, `rg0` and `rg1`:

```
vol create newvol -r 14 28@36
```

Parameters to accept or change after volume creation

After you create a volume, you can accept the defaults for CIFS oplocks and security style settings or change the values. You should decide what to do as soon as possible after creating the volume. If you change the parameters after files are in the volume, they might become inaccessible to users because of conflicts between the old and new values. For example, UNIX files available under mixed security might not be available after you change to NTFS security.

CIFS oplocks setting: The CIFS oplocks setting determines whether the volume uses CIFS oplocks. The default is to use CIFS oplocks.

Security style: The security style determines whether the files in a volume use NTFS security, UNIX security, or both.

With a new gFiler, the default depends on what protocols you licensed, as shown in the following table.

Protocol licenses	Default volume security style
CIFS only	NTFS
NFS only	UNIX
CIFS and NFS	UNIX

See the *Data ONTAP Storage Management Guide* for information about the security style.

Management of RAID groups associated with NetApp volumes

RAID group described

A NetApp volume is composed of one or more NetApp RAID groups. From the perspective of gFilers, Data ONTAP uses RAID groups to determine where to allocate data to the LUNs. A gFiler RAID group is composed of one or more LUNs.

The default size of a RAID group is 8 LUNs. Typically, you would not change the default RAID group size. However, you need to change the default RAID group size if you want to add LUNs to the volume that are of a different size than the LUNs in the current RAID group; that is, the RAID group to which Data ONTAP is automatically adding LUNs to (see “[How Data ONTAP creates RAID groups automatically](#)” on page 97). This section provides details about RAID groups and provides information about how to change the RAID group size if you need to.

Note_____

Current in this context means the RAID group to which LUNs are being added.

RAID group metrics

The following table shows metrics for RAID groups.

Metric	Information
Maximum RAID group size	14 LUNs
Minimum RAID group size	1 LUN
Default RAID group size	8 LUNs
Maximum number of RAID groups in a volume	150 RAID groups
Maximum number of RAID groups on a gFiler	400 RAID groups (applies to both clustered gFilers and a single gFiler)

Caution_____

Exceeding the maximum volume size or maximum capacity for a gFiler model limits the maximum RAID groups in a volume.

How Data ONTAP creates RAID groups automatically

When you create a volume or add LUNs to a volume, Data ONTAP creates new RAID groups as each RAID group is filled with its maximum number of LUNs, according to the *raid-group-size* parameter in the `vol` command. Within each volume, RAID groups are named `rg0`, `rg1`, `rg2`, and so on, in order of their creation. The last RAID group formed might contain fewer LUNs than are specified for the volume's RAID group size. In that case, any LUNs added to the volume are also added to the last RAID group until the specified RAID group size is reached.

Note

When creating a volume, you can use the `-r raid-group-size` command to specify the number of LUNs per RAID group. If no RAID group size is specified, the default value for your gFiler model is applied.

What to do if you are expanding a volume with a different-sized LUN

To avoid a performance penalty, you want all LUNs in a particular RAID group to be the same size. Therefore, when you want to add a LUN of a different size to a volume, you need to stop the current RAID group from growing beyond its current size (as would automatically occur if there is still room) and force a new RAID group to be created that would contain the LUNs of the different size. You can use the volume option named `raidsize` to limit the size the current RAID group to the number of LUNs already in it and to force a new RAID group to be created for newly-added LUNs. See [“Changing the size of the current RAID group”](#) on page 98.

What to do after changing the LUN size to force a new RAID group:

Your goal should be to keep the default RAID group size of 8. After forcing a new RAID group to be created for the LUNs of the different size, stop the new RAID group from expanding beyond the LUNs of the different size, then set the RAID group size back to 8. Future RAID groups then contain 8 LUNs. Follow the procedure in [“Changing the size of the current RAID group”](#) on page 98 for both stopping the RAID group size for the new, different-sized LUNs and returning to the default size of 8.

Changing the size of the current RAID group

You can use the `raidsiz` volume option to set the maximum RAID group size that can be reached when LUNs are added to a volume.

To...	Modify the <code>raidsiz</code> setting ...
Decrease the size of the current RAID group setting. This setting applies to future RAID groups until you change the <code>raidsiz</code> setting again.	Downward To stop a RAID group from expanding, set the value to the number of LUNs already in the current RAID group. You can use the <code>sysconfig</code> command and the <code>vol status -v</code> command to determine the settings for the current RAID group size setting and LUN size.
Increase the allowable size of the current RAID group and future RAID groups This setting applies to future RAID groups until you change the <code>raidsiz</code> setting again.	Upward

RAID groups that became full at a specific RAID group setting remain whatever size they were before the `raidsiz` setting was changed.

Changing the `raidsiz` setting: To change the `raidsiz` setting for a RAID group in an existing volume, complete the following step.

Step	Action
1	Enter the following command: <pre>vol options vol-name raidsiz size</pre> <i>vol-name</i> is the volume whose <code>raidsiz</code> setting you are changing. <i>size</i> is the number of LUNs you want in the most recently created RAID group and all future RAID groups in this volume. Example: The following command changes the <code>raidsiz</code> setting of the volume <code>yervol</code> to 16 LUNs: <pre>vol options yervol raidsiz 16</pre>

Verifying the raidsize setting: To verify the size of raidsize setting in a volume, enter the `vol options vol_name` command.

Modifying the size of an existing RAID group

Another way to increase the size of an existing RAID group is to use the `-g raidgroup` option in the `vol add` command

Adding new LUNs to a RAID group

If a volume has more than one RAID group, you can specify the RAID group to which you are adding LUNs. Typically you don't need to use this procedure, except if your RAID groups contain different size LUNs and the RAID groups are not yet full.

To add new LUNs to a specific RAID group of a volume, complete the following step.

Step	Action
1	Enter the following command: <pre>vol add vol-name ndisks[@disk-size] -d disk-name... -g raidgroup</pre> <i>raidgroup</i> is a RAID group in the volume specified by the <i>vol-name</i>. Example: The following command adds two disks to RAID group 0 of the vol0 volume: <pre>vol add vol0 -g rg0 2</pre>

Note_____

The number of LUNs you can add to a specific RAID group is limited by the raidsize setting of the volume to which that group belongs.

Managing LUNs and LUN ownership

One LUN is assigned to the root volume during setup

During the initial setup of a gFiler, you are required to specify one LUN that the gFiler is to own, but you can specify more. During the setup process, the system automatically creates a root volume and chooses one LUN (if you entered multiple LUNs) to be in the root volume. (Typically, the system chooses the smallest LUN.) Because that LUN was assigned to a volume through this process, the gFiler can read from and write to the LUN.

About spare LUNs

Until the LUN that a gFiler owns is configured to be part of a volume, it is a *spare* LUN.

During initial setup, the system adds one LUN that you specified in the disk assign command (in item 5 Maintenance on the boot menu) to the root volume. The system specifies any additional LUNs in item 5 as spares. After initial setup, you can also add spare LUNs to the gFiler, as described in “[About adding LUN ownership after initial setup](#)” on page 100.

How to enable a LUN to be used for storage

For a LUN to be used for storage, you must add it to a NetApp volume. You can include the spare in the vol create command for a new volume (see “[Creating a volume](#)” on page 92.) Alternatively you can include the spare in the vol add command for an existing volume. See “[Adding LUNs \(disks\) to a volume by using the disk add command](#)” on page 103.

Requirement for spare LUNs for a GF270c

Each head of a GF270c must have one spare LUN of 1.5 GB or more assigned to it at all times, in case a problem occurs that results in a panic. The system can then dump core to the spare LUN and NetApp can analyze the core dump to determine the cause of the problem.

About adding LUN ownership after initial setup

You can specify additional LUNs that a specific gFiler is to own anytime after Data ONTAP is installed. Circumstances in which you would add gFiler ownership of a LUN on a storage subsystem include the following:

- ◆ You left some LUNs that the storage subsystem made available for gFilers unowned and are now ready to designate the gFiler as the owner because you now need to use the storage.
- ◆ Other gFilers might have released ownership and you want this gFiler to own the read and write capability to the LUN.
- ◆ The LUNs were not available on the storage subsystem at the time you initially configured your gFiler.

LUN name format

LUN name format is shown in the following list:

- ◆ For a direct-attached configuration:
`<adapter>.<id>L<lun-id>`
 where
`<adapter>` is the adapter on the gFiler.
`<id>` is the channel adapter port on the storage subsystem.
`<lun-id>` is the full LUN ID

Example: 8a.0L4

- ◆ For a switch-attached configuration
`<switch-name>:<port>.<id>L<lun-id>`
 where
`<switch-name>` is the name of the switch.
`<port>` is the switch port.
`<id>` is the device ID.
`<lun-id>` is the LUN ID.

Example: s2:4.126l4 or s2:4.127l4 (depending on the type of switch)

Assigning disk (LUN) ownership

You use the `disk assign` command for assigning ownership of LUNs on a storage subsystem to a gFiler. To assign LUN ownership, complete the following steps.

Step	Action
1	If you are <i>not</i> going to be deploying SyncMirror, from the CLI prompt enter the <code>disk assign</code> command in the form <pre>disk assign fullLUNname1 fullLUNname2 ...</pre> You do not need to enter the pool parameter (<code>-p</code>) with the <code>disk assign</code> command. The pool parameter defaults to 0 (zero). You can enter multiple LUNs on the same command line.
2	If you <i>are</i> going to be deploying SyncMirror, enter the <code>disk assign</code> command with the pool parameter. You must use the pool parameter with the <code>disk assign</code> command, and you must enter the pool parameter correctly. See “Planning for RAID mirroring for data protection” on page 21 for details so that you can plan your configuration correctly. For example: <pre>disk assign storeAlun1 storeAlun2 -p 0 disk assign storeBlun1 storeBlun2 -p 1</pre> Typically, pool 0 is assigned to the local pool and pool 1 is assigned to the remote pool.
3	Enter the following command: <pre>disk show -v</pre> Result: The Local System ID is shown and the LUNs (disks) exported on the port are shown.
4	Review the list of LUNs (disks) to confirm that all disks (LUNs) that you expected to be assigned to this gFiler (the local system) are shown with the ID of this gFiler.
5	Enter halt, then type boot at the OK prompt.

Note

A LUN cannot be used for storage until it becomes part of a volume, as described in [“How to enable a LUN to be used for storage”](#) on page 100.

**Adding LUNs
(disks) to a volume
by using the disk
add command**

To add new disks to a volume by using the `disk add` command, complete the following steps.

Note_____

You can also add LUNs while creating a volume, as described in “[Creating a volume](#)” on page 92.

Step	Action
1	Verify that hot spare disks are available for you to add by entering the following command: vol status -s
2	Add the disks to a volume by entering the following command: vol add vol-name nluns[@lun-size] -d lun-name... <i>vol-name</i> is the name of the volume to which you are adding the disks. <i>nluns</i> is the number of LUNs to use. <i>lun-size</i> is the LUN size, in gigabytes, to use. You must have at least <i>nluns</i> available of the size you specify. <i>lun-name</i> is the disk (LUN) ID of one more spare LUN IDs; use a space to separate LUN IDs. The disk ID is under the HA.ID column in the <code>vol status -s</code> display. Example: The following command adds the LUNs 7.17 and 7.26 to the <code>thisvol</code> volume: <code>vol add thisvol -d 7.17 7.26</code>

Note_____

When adding LUNs to a volume by using automatic selection (that is, without using the `-d` option), Data ONTAP attempts to distribute I/O loads, such as alternating host adapters, to improve reconstruction performance. You can override this behavior by using the `-d` option to specify the LUNs you want.

Changing ownership of a spare LUN

If you want to change the ownership of a spare LUN from one gFiler to another within the same neighborhood, use the procedure in this section instead of using SnapMover. To change the ownership of a spare LUN, complete the following steps.

Step	Action
1	At the console of the gFiler that you expect owns the spare LUN, enter the following command to see a list of spare LUNs on the gFiler: <pre>vol status -s</pre> Result: The spare disks (LUNs) are listed.
2	Reassign ownership of the spare LUN to the other gFiler by entering the following command: <pre>disk assign <LUN-name> -o <name-of-the-receiving-system> -f</pre> <i>LUN-name</i> is the name of the LUN. <i>sysid-of-the-receiving-system</i> is the NVRAM of the gFiler to which the vfiler is being moved. You can find the ID of the gFiler to which the vFiler is being moved by using the sysconfig command on the destination gFiler. Result: A message that the status of the ownership changed is displayed.
3	Verify that the ownership of the spare LUN moved to the other gFiler by entering the following command: <pre>vol status -s</pre> Result: The spare LUN that you moved should no longer appear in the list of spare disks (LUNs).
4	On the destination gFiler, enter the following command to ensure that the spare LUN for which you changed ownership is listed in the spare LUNs listed: <pre>vol status -s</pre>

About making a LUN unused

A gFiler must cease all the input and output operations to a LUN for you to do the following on the storage subsystem:

- ◆ Unmap the LUN and use it for a non-gFiler host. See “[Removal of a LUN from a gFiler to move to a non-gFiler](#)” on page 106.
- ◆ Make LUN size or composition changes on the storage subsystem.

You can use the `disk remove -w` command to change the “state” of a LUN to unused. After the state is unused, the LUN can no longer be used by any neighborhood participants. (In contrast, if you want to remove disk ownership of a LUN but you want to preserve the ability for neighborhood participants to own it, you use the `disk assign -s unowned` command.)

The `disk remove -w` command does the following:

- ◆ Removes the label for RAID configuration information
- ◆ Makes the LUN unowned
- ◆ Moves the LUN into the broken pool; thereafter, the LUN cannot be used by the gFilers

After you use the `disk remove -w` command on a LUN (from the gFiler), you can make the LUN unavailable to the gFilers (from the storage subsystem).

Removing LUN ownership

To prepare the LUN to be physically removed, complete the following step.

Step	Action
1	From the CLI prompt of the gFiler, enter the following command: disk remove -w <LUNfullname> where <LUNfullname> is the full name of the LUN.

Changing LUN size and composition

The gFiler supports LUN size or composition changes only for *unused* LUNs. If a LUN is assigned as a spare LUN or an owned LUN (through the gFiler ownership scheme), you cannot change the size or composition of the LUN from the storage system management tool. Therefore, if you want to change the LUN size or composition, you must first make the LUN unused on the gFiler.

To change the size or composition of an unused LUN, complete the following steps.

Step	Action
1	On the gFiler, make a LUN unused by specifying the full name of the LUN in the <code>disk remove -w</code> command from the CLI command prompt (as described in “About making a LUN unused” on page 105). For example: <code>disk remove -w <LUNfullname></code>
2	On the storage subsystem, make the LUN invisible to the gFilers.
3	Change the size of the LUN.
4	Make the LUN available to the gFiler again. Result: The LUN will be seen to the gFiler as unused.
5	Assign the LUN ownership by a gFiler (by using the <code>disk assign</code> command).

Removal of a LUN from a gFiler to move to a non-gFiler

If a LUN is mapped to a gFiler and you want to unallocate the LUN from the gFiler and map it to another host, use a LUN that is a spare. Make the LUN unused by using the `disk remove -w` command, as described in [“About making a LUN unused”](#) on page 105.

Checksum space on each LUN

The gFiler formats LUNs in a special way to store checksum information that is used for data integrity checking on reads. One eighth of the space of the LUN is used for this purpose and is not available to the user.

Core dump space per LUN

For gFiler models other than the GF270c, one percent of the space on each LUN is reserved for core dumps. A minimum capacity of 272 GB results in 2.5 GB of core dump space (accounting for loss of space due to checksum formatting). Typically, this 1 percent allocation should be enough for core dumps.

Requirement for spare LUNs for a GF270c: Each head of a GF270c must have one spare LUN of 1.5 GB or more assigned to it at all times, in case a problem occurs that results in a panic. The system can then dump core to the spare LUN and NetApp can analyze the core dump to determine the cause of the problem.

About this chapter

This chapter describes using SnapMover to move the workload of processing of requests from one gFiler to another.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Planning for using SnapMover to change workload distribution](#)” on page 110
- ◆ “[Changing workload distribution with SnapMover](#)” on page 115

Planning for using SnapMover to change workload distribution

About this section	If the quantity of requests sent to a gFiler becomes high enough to affect performance, you can use SnapMover to move request servicing from one gFiler to another. This section provides information to help you plan for using SnapMover.
SnapMover described	SnapMover works in conjunction with the MultiStore feature. Its sole purpose is to enable you to move MultiStore <i>vFilers</i> quickly and efficiently between gFilers that are <i>sharing the same storage subsystem</i> . SnapMover uses the disk (LUN) ownership scheme to migrate (move) a vFiler instead of copying it, which makes migration quite fast. In contrast, the standard Data ONTAP vFiler migrate facility, which uses SnapMirror, copies data. Copying the data is time-consuming.
Licenses required	Two licenses are required to use SnapMover—one for SnapMover and one for MultiStore.
MultiStore vFiler technology described	With MultiStore, you can partition a single gFiler (or filer) into multiple logical units referred to as <i>vFilers</i>. Each vFiler owns its own exclusive storage and networking resources; no sharing of resources occurs even if the resources are on the same system. vFilers allow a lot of flexibility in how data can be partitioned and managed. For example, you could ◆ Create a vFiler to logically separate data belonging to a specific service provider from the data belonging to another service provider. ◆ Create a vFiler for data that you want to process on days 1 to 10 of the month and another vFiler for data on that you want to process on days 11 to 20 of the month. You could then move the vFilers, as necessary, to free up CPU processing. The <i>Data ONTAP MultiStore Guide</i> provides more examples of how you can use vFilers to partition and manage your data.

Requirement to associate vFilers with volumes: Before you can create a vFiler, you must define (create) one or more units of data storage called NetApp *volumes* on your gFilers. Then, when you create the vFiler, you identify the volume associated with the vFiler.


Note

The purpose of SnapMover is to move vFilers. You cannot use SnapMover to move a spare LUN, which is not part of a volume.


What the volume-level vFilers used with SnapMover are

SnapMover supports *volume-level vFilers* only; that is, vFilers for which all storage elements are composed of complete volumes. In practice, this means a vFiler with paths of the form `/vol/<volname>`. The data structure `/vol/<volname>/<subvolname>` is not supported for SnapMover.

The following illustration shows the configuration that is supported for SnapMover.



When SnapMover moves a vFiler, SnapMover moves all volumes associated with the vFiler, as the following illustration shows.



“[Information that is moved](#)” on page 114 provides details about what data SnapMover moves with the vFiler.

Requirements

The following requirements exist for using SnapMover and MultiStore to change the workload distribution:

- ◆ The following licenses must be installed on each gFiler after the disk (LUN) assignment process is complete and Data ONTAP is installed:
 - ❖ MultiStore
 - ❖ SnapMover

Other licenses used by the vFiler must match. For example, if CIFS is licensed on the source gFiler, it must be licensed on the destination gFiler. Otherwise, moving the vFiler causes CIFS to be unavailable after the vFiler is moved to the destination.
- ◆ Both the source gFiler and the destination gFiler (the gFilers between which you are moving a vFiler) must be connected to the same storage subsystem. The LUNs must be visible to both the source and destination gFilers.
- ◆ (Highly recommended) To ensure that disk (LUN) ownership changes are transparent to NFS users, the destination gFiler must have an Ethernet connection to the same subnet that the source gFiler uses.
- ◆ The volume that the vFiler will be using must be created before the vFiler is created. See “[Creating a volume](#)” on page 92.
- ◆ A vFiler must be created on the source.

- ◆ The vFiler's storage units must all be composed of complete volumes; that is, a vFiler with paths of the form `/vol/<volname>`.
 - ❖ The volumes for vFiler configuration information must be writable.
 - ❖ The IP address used by the vFiler must be unconfigured when you create the vFiler.
- ◆ At least one of the following protocols must be available: rsh, NFS, CIFS, or iSCSI

Guidelines for setting up volumes to support volume-vFilers for SnapMover

When SnapMover moves a vFiler, all volumes associated with that vFiler are moved. Therefore, when you create the vFilers, consider the following:

- ◆ SnapMover cannot move part of a volume when it moves a vFiler.
- ◆ The destination gFiler must be able to accommodate the size of the vFiler with all its associated volumes.
- ◆ There is a limit on the number of volumes and vFilers on a gFiler, depending on the memory capacity of the hosting gFiler. See "[Volume metrics](#)" on page 90 for the limit on the number of volumes. You can create up to 32 vFilers on a gFiler (64 in a cluster), depending on the memory capacity of the hosting gFiler. (A gFiler containing the maximum number of vFilers actually has 33 vFilers: vFiler0 plus the 32 vFilers that you, the system administrator, can create.)
- ◆ The names of the vFilers and volumes being moved from the source to the destination must be unique.

Although you can rename a volume, it is best not to do so if you have NFS clients because the renaming of the volume is not transparent to NFS clients. When the gFiler uses NFS to export a file system, the volume name is part of the exported path name. NFS clients are accustomed to trying to mount to the old path name; in order to access the data after the vFiler has been moved, clients will need to remount using the new path name.

Note

The *Data ONTAP MultiStore Guide* provides information about planning for, creating, and managing vFilers. Most of the information provided is applicable to volume-level vFilers. However, keep in mind when reading the instructions for creating a vFiler that the standard vFilers do not need to be a complete volume as is required for SnapMover.

qtree migration is not supported

You cannot migrate qtrees through SnapMover.

Information that is moved

When SnapMover moves the vFiler, it automatically updates the data associated with volumes, such as IP address, permissions, vFiler ID, and disk ownership (such as IP addresses). NFS clients are unaware that the vFiler is being moved. To CIFS client workstations, it appears as if the server was rebooted.

SnapMover command

The SnapMover command is `vfiler migrate -n nocopy`.

What the vfiler migrate -m nocopy command does

The `vfiler migrate -m nocopy` command does the following:

- ◆ Ensures that no vFiler with the same name exists on the destination
- ◆ Ensures that the SnapMover license is installed on both the source and destination gFilers
- ◆ Ensures that licenses and WAFL® versions are the same on both the source and destination gFilers
- ◆ Rewrites the disk ownership record so that the ownership of the volumes passes from the source vFiler to the destination vFiler
- ◆ Saves the IP configuration and binding information that you supplied when you created the vFiler
- ◆ Saves the quota information from the source vFiler's `/etc/quotas` file (how much disk space individual users can use)
- ◆ Stops the source vFiler
- ◆ Destroys the source vFiler
- ◆ Re-creates the vFiler on the destination gFiler

Reverting if you are using SnapMover

If you have migrated a volume vFiler and you want to move it back to the source, you cannot do so unless both the source and destination gFiler are running Data ONTAP 6.5 or later. You either need to revert the source gFiler to the same pre-6.5 Data ONTAP version as the destination gFiler or upgrade the destination gFiler to 6.5 or later.

Changing workload distribution with SnapMover

About this section

This section provides instructions for how to use SnapMover with the vFiler technology provided by the MultiStore software to move the workload from one gFiler to another.

Preparing for the use of SnapMover

Before you can move a vFiler you must prepare for using SnapMover by completing the following steps.

Step	Action
1	Install a SnapMover and MultiStore license on each of the two gFilers.
2	On the source gFiler, create a volume by following the instructions in “Creating a volume” on page 92. Example command entry: vol create vfl 2 <i>volume create</i> is the command for creating the volume. <i>vfl</i> is the name of the volume. 2 is the number of LUNs to put in the volume. Ensure that the volume name is not the same as the name of any volume on the destination gFiler.
3	Enter the following command to create a vFiler on the source gFiler: vfiler create <vfilername> -i <ipaddress> <pathname> Answer the prompts. In most cases, you can press the Enter key to accept the defaults. You need to specify the interface to use, however.

Moving a volume-level vFiler

The operation that is described in the following procedure is a pull operation; that is, you perform this operation on the gFiler to which you want the volume-level vFiler transferred (the *destination* gFiler). You do not have to shut down the destination gFiler to perform this operation on it.

To move a volume-level vFiler to the destination gFiler, complete the following steps.

Step	Action
1	On the <i>destination</i> gFiler, enter the following command: vfiler migrate -m nocopy <source_vfiler>@<source_gfiler> <i>source_vfiler</i> is the name of the vfiler that you are moving <i>source_gfiler</i> is the gFiler from which you are moving the vfiler.
2	Answer the prompts, including providing the following: a. A valid administrative login ID and password b. The IP address and binding information for the destination vfiler. Result: The vFiler is migrated from the source gFiler to the destination gFiler.
3	Verify that the vFiler was moved by entering the following command on the destination gFiler: vfiler status -r

Configuration Worksheet

A

How to use this worksheet

Use this worksheet to record the information you need to set up and configure your gateway. Gather the information before you begin setting up and configuring your system.

Parameters		Your values
Gateway	MAC address for the gateway's built-in Ethernet interface	
	Host name	
	Password	
	Time zone	
	Filer location	
	Language used for multiprotocol filers	
Administration host	Host name	
	IP address	
Ethernet	Interface name	
	IP address	
	Subnet mask	
	Partner IP address	
	Are jumbo frames supported?	
	MTU size for jumbo frames	
Virtual interface	Interface name	
	Number of links	
	Link names	
Router	Gateway name	
	IP address	

Parameters		Your values
DNS	Domain name	
	Server address 1	
	Server address 2	
	Server address 3	
NIS	Domain name	
	Server name 1	
	Server name 2	
	Server name 3	
Windows domain		
WINS servers	1	
	2	
	3	
Windows 2000 user		
Windows 2000 user password		
Active Directory (command-line setup only)		
Location of HTTP directory		
RMC	MAC address	
	IP address	
	Subnet mask	
	Gateway	
	Mailhost	

About this section

Enter one of the following time zones when the setup command prompts for a time zone. If you need to change this after setup is complete, see the *Storage Management Guide* for information about the `timezone` command.

Africa

Africa/Abidjan	Africa/Djibouti	Africa/Maputo
Africa/Accra	Africa/Douala	Africa/Maseru
Africa/Addis_Ababa	Africa/Freetown	Africa/Mbabane
Africa/Algiers	Africa/Gaborone	Africa/Mogadishu
Africa/Asmera	Africa/Harare	Africa/Monrovia
Africa/Bamako	Africa/Johannesburg	Africa/Nairobi
Africa/Bangui	Africa/Kampala	Africa/Ndjamena
Africa/Banjul	Africa/Khartoum	Africa/Niamey
Africa/Bissau	Africa/Kigali	Africa/Nouakchott
Africa/Blantyre	Africa/Kinshasa	Africa/Ouagadougou
Africa/Brazzaville	Africa/Lagos	Africa/Porto-Novo
Africa/Bujumbura	Africa/Libreville	Africa/Sao_Tome
Africa/Cairo	Africa/Lome	Africa/Timbuktu
Africa/Casablanca	Africa/Luanda	Africa/Tripoli
Africa/Conakry	Africa/Lumumbashi	Africa/Tunis
Africa/Dakar	Africa/Lusaka	Africa/Windhoek
Africa/Dar_es_Salaam	Africa/Malabo	

Aliases

GMT=Greenwich, UCT, UTC, Universal, Zulu

CET=MET (Middle European Time)

US/Eastern=Jamaica

US/Mountain=Navajo

America

America/Adak	America/Grenada	America/Noronha
America/Anchorage	America/Guadeloupe	America/Panama
America/Anguilla	America/Guatemala	America/Pangnirtung
America/Antigua	America/Guayaquil	America/Paramaribo
America/Aruba	America/Guyana	America/Phoenix
America/Asuncion	America/Halifax	America/Port_of_Spain

America/Atka	America/Havana	America/Port-au-Prince
America/Barbados	America/Indiana	America/Porto_Acre
America/Belize	America/Indianapolis	America/Puerto_Rico
America/Bogota	America/Inuvik	America/Rainy_River
America/Boise	America/Iqaluit	America/Rankin_Inlet
America/Buenos_Aires	America/Jamaica	America/Regina
America/Caracas	America/Jujuy	America/Rosario
America/Catamarca	America/Juneau	America/Santiago
America/Cayenne	America/Knox_IN	America/Santo_Domingo
America/Cayman	America/La_Paz	America/Sao_Paulo
America/Chicago	America/Lima	America/Scoresbysund
America/Cordoba	America/Los_Angeles	America/Shiprock
America/Costa_Rica	America/Louisville	America/St_Johns
America/Cuiaba	America/Maceio	America/St_Kitts
America/Curacao	America/Managua	America/St_Lucia
America/Dawson	America/Manaus	America/St_Thomas
America/Dawson_Creek	America/Martinique	America/St_Vincent
America/Denver	America/Mazatlan	America/Swift_Current
America/Detroit	America/Mendoza	America/Tegucigalpa
America/Dominica	America/Menominee	America/Thule
America/Edmonton	America/Mexico_City	America/Thunder_Bay
America/El_Salvador	America/Miquelon	America/Tijuana
America/Ensenada	America/Montevideo	America/Tortola
America/Fort_Wayne	America/Montreal	America/Vancouver
America/Fortaleza	America/Montserrat	America/Virgin
America/Glace_Bay	America/Nassau	America/Whitehorse
America/Godthab	America/New_York	America/Winnipeg
America/Goose_Bay	America/Nipigon	America/Yakutat
America/Grand_Turk	America/Nome	America/Yellowknife

Antarctica

Antarctica/Casey	Antarctica/Mawson	Antarctica/Palmer
Antarctica/DumontDUrville	Antarctica/McMurdo	Antarctica/South_Pole

Asia

Asia/Aden	Asia/Irkutsk	Asia/Qatar
Asia/Alma-Ata	Asia/Ishigaki	Asia/Rangoon
Asia/Amman	Asia/Istanbul	Asia/Riyadh
Asia/Anadyr	Asia/Jakarta	Asia/Saigon
Asia/Aqtau	Asia/Jayapura	Asia/Seoul
Asia/Aqtobe	Asia/Jerusalem	Asia/Shanghai
Asia/Ashkhabad	Asia/Kabul	Asia/Singapore
Asia/Baghdad	Asia/Kamchatka	Asia/Taipei

	Asia/Bahrain Asia/Baku Asia/Bangkok Asia/Beirut Asia/Bishkek Asia/Brunei Asia/Calcutta Asia/Chungking Asia/Colombo Asia/Dacca Asia/Damascus Asia/Dubai Asia/Dushanbe Asia/Gaza Asia/Harbin Asia/Hong_Kong	Asia/Karachi Asia/Kashgar Asia/Katmandu Asia/Krasnoyarsk Asia/Kuala_Lumpur Asia/Kuching Asia/Kuwait Asia/Macao Asia/Magadan Asia/Manila Asia/Muscat Asia/Nicosia Asia/Novosibirsk Asia/Omsk Asia/Phnom_Penh Asia/Pyongyang	Asia/Tashkent Asia/Tbilisi Asia/Tehran Asia/Tel_Aviv Asia/Thimbu Asia/Tokyo Asia/Ujung_Pandang Asia/Ulan_Bator Asia/Urumqi Asia/Vientiane Asia/Vladivostok Asia/Yakutsk Asia/Yekaterinburg Asia/Yerevan
Atlantic	Atlantic/Azores Atlantic/Bermuda Atlantic/Canary Atlantic/Cape_Verde	Atlantic/Faeroe Atlantic/Jan_Mayen Atlantic/Madeira Atlantic/Reykjavik	Atlantic/South_Georgia Atlantic/St_Helena Atlantic/Stanley
Australia	Australia/ACT Australia/Adelaide Australia/Brisbane Australia/Broken_Hill Australia/Canberra Australia/Darwin Australia/Hobart	Australia/LHI Australia/Lindeman Australia/Lord_Howe Australia/Melbourne Australia/NSW Australia/North Australia/Perth	Australia/Queensland Australia/South Australia/Sydney Australia/Tasmania Australia/Victoria Australia/West Australia/Yancowinna
Brazil	Brazil/Acre Brazil/DeNoronha	Brazil/East	Brazil/West
Canada	Canada/Atlantic Canada/Central Canada/East-Saskatchewan	Canada/Eastern Canada/Mountain Canada/Newfoundland	Canada/Pacific Canada/Saskatchewan Canada/Yukon

Chile	Chile/Continental	Chile/EasterIsland	
Etc	Etc/GMT	Etc/GMT+11	Etc/GMT-9
	Etc/GMT+0	Etc/GMT+12	Etc/GMT-10
	Etc/GMT+1	Etc/GMT0	Etc/GMT-11
	Etc/GMT+2	Etc/GMT-0	Etc/GMT-12
	Etc/GMT+3	Etc/GMT-1	Etc/GMT-13
	Etc/GMT+4	Etc/GMT-2	Etc/GMT-14
	Etc/GMT+5	Etc/GMT-3	Etc/Greenwich
	Etc/GMT+6	Etc/GMT-4	Etc/UCT
	Etc/GMT+7	Etc/GMT-5	Etc/Universal
	Etc/GMT+8	Etc/GMT-6	Etc/UTC
	Etc/GMT+9	Etc/GMT-7	Etc/Zulu
	Etc/GMT+10	Etc/GMT-8	
Europe	Europe/Amsterdam	Europe/Kiev	Europe/San_Marino
	Europe/Andorra	Europe/Kuybyshev	Europe/Sarajevo
	Europe/Athens	Europe/Lisbon	Europe/Simferopol
	Europe/Belfast	Europe/Ljubljana	Europe/Skopje
	Europe/Belgrade	Europe/London	Europe/Sofia
	Europe/Berlin	Europe/Luxembourg	Europe/Stockholm
	Europe/Bratislava	Europe/Madrid	Europe/Tallinn
	Europe/Brussels	Europe/Malta	Europe/Tirane
	Europe/Bucharest	Europe/Minsk	Europe/Vaduz
	Europe/Budapest	Europe/Monaco	Europe/Vatican
	Europe/Chisinau	Europe/Moscow	Europe/Vienna
	Europe/Copenhagen	Europe/Oslo	Europe/Vilnius
	Europe/Dublin	Europe/Paris	Europe/Warsaw
	Europe/Gibraltar	Europe/Prague	Europe/Zagreb
	Europe/Helsinki	Europe/Riga	Europe/Zurich
	Europe/Istanbul	Europe/Rome	
GMT	GMT	GMT+9	GMT-5
	GMT+1	GMT+10	GMT-6
	GMT+2	GMT+11	GMT-7
	GMT+3	GMT+12	GMT-8
	GMT+4	GMT+13	GMT-9
	GMT+5	GMT-1	GMT-10
	GMT+6	GMT-2	GMT-11

	GMT+7 GMT+8	GMT-3 GMT-4	GMT-12
Indian (Indian Ocean)	Indian/Antananarivo Indian/Chagos Indian/Christmas Indian/Cocos	Indian/Comoro Indian/Kerguelen Indian/Mahe Indian/Maldives	Indian/Mauritius Indian/Mayotte Indian/Reunion
Mexico	Mexico/BajaNorte	Mexico/BajaSur	Mexico/General
Miscellaneous	Arctic/Longyearbyen CET CST6CDT Cuba EET Egypt Eire EST EST5EDT Factory GB GB-Eire Greenwich Hongkong	HST Iceland Iran Israel Japan Kwajalein Libya MET MST MST7MDT Navajo NZ NZ-CHAT Poland	Portugal PRC PST8PDT ROC ROK Singapore Turkey UCT Universal UTC WET W-SU Zulu
Pacific	Pacific/Apia Pacific/Auckland Pacific/Chatham Pacific/Easter Pacific/Efate Pacific/Enderbury Pacific/Fakaofu Pacific/Fiji Pacific/Funafuti Pacific/Galapagos Pacific/Gambier Pacific/Guadalcanal	Pacific/Johnston Pacific/Kiritimati Pacific/Kosrae Pacific/Kwajalein Pacific/Majuro Pacific/Marquesas Pacific/Midway Pacific/Nauru Pacific/Niue Pacific/Norfolk Pacific/Noumea Pacific/Pago_Pago	Pacific/Ponape Pacific/Port_Moresby Pacific/Rarotonga Pacific/Saipan Pacific/Samoa Pacific/Tahiti Pacific/Tarawa Pacific/Tongatapu Pacific/Truk Pacific/Wake Pacific/Wallis Pacific/Yap

Pacific/Guam
Pacific/Honolulu

Pacific/Palau
Pacific/Pitcairn

System V

SystemV/AST4
SystemV/AST4ADT
SystemV/CST6
SystemV/CST6CDT
SystemV/EST5

SystemV/EST5EDT
SystemV/HST10
SystemV/MST7
SystemV/MST7MDT
SystemV/PST8

SystemV/PST8PDT
SystemV/YST9
SystemV/YST9YDT

Language choices The following table lists the languages supported by Data ONTAP on multiprotocol filers.

You can also view this list by entering `vol lang` at the gFiler prompt.

Language	Abbreviation	Language	Abbreviation
Danish	da	Korean	ko
Dutch	nl	Norwegian	no
English	en	Portugese	pt
English (U.S.)	en_US	POSIX	C
Finnish	fi	Simplified Chinese	zh
French	fr	Simplified Chinese (GBK)	zh.GBK
German	de	Spanish	es
Hebrew	he	Swedish	sv
Italian	it	Traditional Chinese euc-tw	zh_TW
Japanese euc-j	ja	Traditional Chinese Big 5	zh_TW.BIG5
Japanese PCK(sjis)	ja_JP.PCK		

Choosing the language code

To choose the language appropriate to your organization, complete the following step.

Step	Action
1	When prompted during setup, enter the code that corresponds to the appropriate language. Note_____ To use UTF-8 as the NFS character set, append <code>.UTF-8</code> to the abbreviation. _____ Example: <code>ko.UTF8</code>

Index

A

- administration tools 8
- Africa, time zone names for 119
- Australia, time zone names for 121
- Autosupport described 10

C

- Canada, time zone names for 121
- checksum space on LUN 106
- CIFS
 - making gateway data available to 69
 - oplocks, setting for volumes 95
 - requirement for using 57
- cluster
 - defined 5
 - with GF270c 5
- Cluster Configuration Checker 80
- clustered gFilers
 - checking configuration 79
 - defined 5
 - networking interfaces for 75
 - verifying configuration for 71
- configuration worksheet 117
- core dump space on LUNs 106

D

- dedicated interfaces 78
- disk ownership
 - assigning 101
 - benefits of 15
 - changing ownership of spares 104
 - command 18
 - control for gFilers only 15
 - defined 14
 - examples 17
 - planning for 19
 - rules for 16
 - scheme described 14
- disk show command options 45
- disk.target_port.cmd_queue_depth option 86

E

- Europe, time zone names for 122

F

- FilerView
 - accessing 9
 - described 8
 - how it works 10

G

- gFiler planning overview 12
- gFilers
 - administration tools 8
 - clusters, verifying configuration for 71
- GF270c
 - cluster with 5
 - MetroCluster and 27
 - spare LUN requirement 20, 44, 100, 107
 - SyncMirror and 21
- GMT, time zone names for 122

H

- host name resolution, setting up 63

I

- installation overview 3
- installing Data ONTAP software 47
- interfaces
 - dedicated description 75
 - shared description 75
 - standby description 75

L

- languages supported 125
- licenses
 - codes for 72
 - installing 48
 - MultiStore 110

- SnapMover 110
- SyncMirror 21
- Logical Unit Number, defined 5
- LUN ownership
 - assigning 101
 - benefits of 15
 - changing ownership of spares 104
 - command 18
 - control for gFilers only 15
 - defined 14
 - examples 17
 - planning for 19
 - rules for 16
 - scheme described 14
- LUN security, preventing LUN overwriting 19
- LUNs
 - adding to RAID groups 99
 - adding to volumes 92, 100
 - changing composition 105
 - changing size 105
 - checksum space on 106
 - core dump space on 106
 - defined 5
 - how gFilers access 14
 - making available to gFilers 14
 - making unused 105
 - mapping to a non-gFiler host 106
 - moving 91
 - name format 101
 - number in volume 90
 - removing ownership 105
 - spare defined 100
 - spare requirement for GF270c 20, 44, 100, 107

M

- MetroCluster
 - deployment example 28
 - GF270c and 27
 - maximum in a neighborhood 28
 - use with SyncMirror 27
- MultiStore vFilers with SnapMover 110

N

- neighborhood

- defined 28
- in switch fabric 28
- maximum MetroClusters in 28
- SnapMover with 28
- network connections, verifying 64
- New Zealand, time zone names for 121
- NFS hosts, making gateway data available to 66
- NOW Web site 10

O

- oplocks, setting for volumes 95
- overview, setup and installation 3

P

- planning overview 12
- plexes
 - described 21
 - relationship to pools 22
 - when created 22
- pools
 - relationship to plexes 22
 - with SyncMirror 22
- prerequisites to setup 1

R

- RAID groups
 - adding LUNs to 99
 - default size 96
 - when to change 96
 - definition 96
 - forcing a new one 97, 98
 - how created 97
 - maximum number 96
 - metrics 96
 - modifying size of 99
 - number in volume 91
 - performance issues 97
 - sizes 96
- RAID mirroring. *See* SyncMirror
- raidsize volume option 97
 - changing size of 98
 - forcing a new RAID group 98
 - verifying setting 99

- RAID-0 6
- RAID-1 5
- RAID-10 5
- RAID-5 5
- Remote Management Controller
 - connections and setup, verifying 84
- reverting SnapMover 114
- RMC connections and setup 84
- root volume
 - defined 19
 - minimum size 90

S

- security styles, setting for volumes 95
- setup data, providing 39
- setup overview 3
- setup program, gathering system data for 29
- shared switch configuration requirement 86
- SnapMirror, contrasted to SnapMover 110
- SnapMover
 - command 114
 - what it does 114
 - contrasted to SnapMirror 110
 - how it works 110
 - information moved 114
 - licenses for 110
 - migration in neighborhood 28
 - moving the vFiler 116
 - preparing for use of 115
 - purpose of 110
 - requirements 112
 - restriction 111
 - qtree movement 113
 - reverting 114
 - setting up volumes 113
 - vFilers with 110
 - volume-level vFilers
 - defined 111
- South America, time zones for 121, 122
- spare LUN
 - changing ownership of 104
 - defined 100
- spare LUNs
 - requirement for GF270c 20, 44, 100, 107

- storage subsystems
 - gFiler access to LUNs on 14
- switch fabric, neighborhoods in 28
- switch zoning, preventing LUN overwriting 19
- switches, configuration requirement with 86
- SyncMirror
 - configuration requirement 21
 - GF270c and 21
 - how it works 22
 - license for 21
 - MetroCluster
 - deployment example 28
 - planning for 20
 - plexes
 - described 21
 - pools with 22
 - requirements 24, 26
 - example configuration 25
 - storage space required 22
 - why use 21
 - with MetroCluster 27
- system data for setup program 29

T

- target queue depth length, setting 87
- time zones 119
- troubleshooting
 - giveback failure 83
 - takeover failure 83

V

- vFilers
 - defined 110
 - use with SnapMover 110
 - volume association 111
- volumes
 - adding LUNs to 20, 92
 - creating 92
 - examples 94
 - parameters after 94
 - creating additional 92
 - defined 6, 19, 90
 - guidelines to determine 91
 - LUNs in, number 90

- maximum number 90
- metrics 90
- moving LUNs in 91
- naming 92
- non-root
 - size 90
- plexes in 21
- RAID groups in 91
 - maximum number in 96
- root volume 19

- size 20, 90
- security style 95
- setting up for SnapMover 113
- size 90
 - guidelines to determine 91
- vFilers with 111

W

workload distribution. *See* SnapMover