

NetCache™ 5.4

DNFS Administration 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 210-00319
December 2002

Copyright and trademark information

Copyright information

Copyright © 1994–2002 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 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.

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

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

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

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

Trademark information

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

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

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	.ix
Chapter 1	The Distributed NFS Appliance	1
	The distributed NFS solution	2
	How the appliance works	4
	File handling	7
	Deployment preparation	10
	Export volumes	13
	Configuring export volumes	18
	Disconnected operations	23
Chapter 2	Administering Access	25
	Adding administrator accounts	26
	Access control lists	29
	Forming rules	31
Chapter 3	Network Interfaces	37
	Planning summary for network interface configuration	38
	About physical network interfaces	39
	Trunking	41
	Virtual interfaces	43
	Troubleshooting network interfaces	46
Chapter 4	Routing	47
	Planning summary for routing configuration.	48
	How the appliance handles routing.	50
	Binding types of traffic to specific interfaces	52
	Restricting interfaces for connecting to servers	55

	Using TCP large windows to improve performance	56
Chapter 5	Support for DNS.	59
	Basic DNS configuration	60
	Optional DNS configuration	61
	Handling DNS lookups for local servers	65
Chapter 6	Managing Disks	67
	Section A: RAID.	68
	Understanding RAID	69
	Understanding disk addresses	71
	Understanding usable space on each disk	73
	Using fibre channel disks of various sizes	74
	Handling disk failure	75
	Section B: RAID disk management	77
	Displaying the status and statistics of disks	78
	About adding and removing disks	80
	Troubleshooting a failed disk.	82
	Section C: Reinitializing all disks	83
Chapter 7	Support for Network Information Service.	85
	NIS database administration	86
Chapter 8	Bandwidth Allocation Management	89
	About bandwidth allocation	90
	Viewing existing bandwidth allocation rules.	91
Chapter 9	Using SNMP and the Appliance MIB	93
	About SNMP	94
	About the MIB	96

	Understanding SNMP traps.	98
	Creating a user-defined SNMP trap	99
Chapter 10	Automated Support with Autosupport	107
	Planning summary for Autosupport configuration	108
	About Autosupport	110
Chapter 11	Routine Administration.	117
	Halting, rebooting, and powering down the appliance	118
	Backing up and restoring the appliance configuration	119
	Messages log	121
	Index	123

Preface

About this guide

This guide describes how to administer, manage, and optimize appliances running Distributed NFS software. It provides information for the following supported NetCache platforms: C2100, C3100, and C6100.

For information about how to upgrade or service the system hardware, see the field service guide for your appliance model.

Audience

This guide is for system administrators who are familiar with computer networks. This guide does not cover basic system administration topics such as IP addressing, routing, and network topology. It emphasizes the characteristics of the appliance hardware and the Distributed NFS software.

Command and keyboard conventions

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

When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, “Ctrl-D” means pressing the “Control” and “D” keys simultaneously.

This guide uses the term “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 Enter. 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.

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. (The keys on your keyboard might not be labeled exactly as they are in this guide.)

Formatting conventions

The following table describes typographical conventions used in this guide.

Formatting convention	Type of information
<i>Italic type</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 hostname</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 it must be typed in uppercase letters to work properly.

About this chapter	This chapter discusses the distributed NFS appliance as a data services solution. The appliance extends file server capability by caching and serving content that originates from NFS file servers. Additionally, this chapter describes how the appliance facilitates seamless file sharing across a network.
Topics in this chapter	This chapter contains the following topics: ◆ “The distributed NFS solution” on page 2 ◆ “How the appliance works” on page 4 ◆ “File handling” on page 7 ◆ “Deployment preparation” on page 10 ◆ “Export volumes” on page 13 ◆ “Configuring export volumes” on page 18 ◆ “Disconnected operations” on page 23
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions.
Location of configuration options	Options for configuring distributed NFS are in the following location in the Appliance Manager: Setup tab > NFS
Using the CLI	To get information about NFS commands from the CLI, enter the following command: <pre>set config.nfs.?</pre> To get help for a specific appliance option, enter the following command: <pre>help config.nfs.optionname</pre>

The distributed NFS solution

Network latency and remote connectivity

Network attached storage (NAS) appliances are used to consolidate data and provide shared access to files over wide area networks (WANs). Large, geographically distributed organizations have high-speed local area networks (LANs) where data for each site is consolidated and administered from a central data repository.

Although data management at a central location is efficient, low bandwidth and high latency complicate long-distance connectivity for remote or branch offices. For example, remote offices experience delays in transferring large files for certain applications, and the WAN must be shared for all remote-office needs.

Maintaining network performance and connectivity to support distributed file systems is a fundamental challenge that administrators face every day. For example, remote offices often have no resident IT staff, and backup of remote office storage can be problematic.

What is a distributed NFS appliance?

A distributed NFS appliance is a device that supports remote data distribution within NFS client and NFS server configurations.

The appliance is an NFS proxy file server that

- ◆ Can interface with applications as if the applications were in direct communication with an NFS source filer
- ◆ Provides cached access to data originating from file servers in geographically distributed locations

Supported origin servers

A distributed NFS appliance supports the following NFS server operating systems:

- ◆ NetApp Filers
- ◆ Solaris 5.7/5.8 TCP/UDP V2/V3
- ◆ Red Hat Linux 7.1/7.2 UDP V2/V3. TCP is not supported when using Linux file servers. The reason is that these Linux versions do not have TCP server functionality built in.

Protocols and NFS versions supported

Currently, the distributed NFS appliance offers support for certain protocols and NFS versions. Specifically, the appliance

- ◆ Supports NFS in TCP and UDP mode
- ◆ Supports both NFS version 2 and NFS version 3
- ◆ Supports NLM
- ◆ Supports NIS
- ◆ Does not support HTTP, streaming, NNTP, or FTP
- ◆ Does not support the Solaris nfs_acl protocol

File sharing solutions

The appliance offers the following solutions for issues inherent to file sharing:

- ◆ Increases performance
The appliance consolidates data closer to the end user where cached access improves response time for users in remote locations.
- ◆ Reduces load on network infrastructure
The appliance assists NFS file servers by acting as a proxy to maintain and serve data.
- ◆ Reduces bandwidth consumption
The appliance eliminates the requirement to maintain and administer full dataset replication to service remote locations.
- ◆ Simplifies remote administration
Administrative functions for the appliance are performed from the central data location.
- ◆ Ensures data consistency
The appliance verifies object maturity at periodic intervals (set by an administrator) to ensure that data is consistent with the source storage device at the data center.

How the appliance works

Direct access to cached data

The NFS-enabled appliance caches and serves data to clients that are remote from an administrative data center. A file block is read through the network and cached. Subsequent requests for that file are then served from the appliance. Consequently, clients in remote locations are provided with direct access to cached data. This eliminates the requirement for client requests to travel repeatedly across the network to fetch the same file.


Cache hits and misses

When a client makes a read request, if the object or relevant block is located in the cache, a cache hit is served from the appliance. Cache hits are the result of a previous request. If a cache hit is served, the appliance generates a response and sends it to the client. If a locate fails (cache miss), the appliance uses the response from the origin server to load the object into the cache, then serves a response to the client.

A cache miss can be one of the following:

- ◆ Complete miss—The object is not cached.
- ◆ Sparse miss— Blocks that correspond to the client request are inadequate.
- ◆ Verify miss — Data in the cache is not valid because object content has changed on the server.

The following illustration maps a client request through a cache hit, cache miss sequence.



The following table describes the flow of the client request shown in the preceding illustration.

Stage	Description
1	The client makes a request for content to the appliance.
2	The appliance checks its cached objects for the requested content. a. If the request for content is a cache hit (that is, if the data exists in the cache), the cached object is served to the local client. b. If the request for content is a cache miss (that is, if the data does not exist in the cache), the appliance then acts as a client to fetch the object from the file server.
3	The file server satisfies the request for content and returns it to the appliance where the object is cached and served to the client.

Cache consistency

The appliance verifies object attributes prior to serving a request based on maximum age or schedules. Validation occurs at the file level, and cached data is validated when the appliance discovers that file attributes have changed. When

source file attributes change, the file is ejected from the cache. This means that if one block of a hundred-block file has been modified, the appliance ejects the entire file.

The maximum time for the appliance to serve a hit on a cached object without issuing a request to the origin server is a configurable option that is defined on a per- volume basis. Validation settings can range from zero (revalidate before each hit on a cached object is served) to x amount of time (revalidate every number of minutes or hours).

If a client is reading a cached file and a second client changes the state of the file on the origin server, the appliance ejects the cached object after verifying the object attributes. The requested blocks are then fetched from the origin server, the object is cached, and the appliance serves the client response.

File handling

Read-only cache support

The appliance supports caching of NFS requests and read-only operations. Read-only data is written to disk on the appliance and cached to serve future client requests. All other operations are passed through or proxied to the origin server. The appliance proxies any operation that modifies the state of the file system, for example, writing to a file or creating a directory.

File handles

The distributed NFS appliance caches objects by NFS file handle as opposed to file name. An NFS file handle is a reference to a file or a directory that is independent of a file name. All NFS operations use a file handle to identify the file or directory to which an operation applies.

The file handle, an opaque piece of data, is created by and returned from an NFS server in response to NFS requests. Only the server can interpret the data contained within that file handle. File handles provide the server flexibility in determining the most efficient way to reference a file.

A file handle is a persistent reference, which means that it remains valid until the file or directory is removed from the server.

Clients use file handles to refer to both files and directories on which they want to perform NFS operations. The operations most commonly used are `LOOKUP`, `GETATTR`, and `ACCESS`. Some common NFS operations include those shown in the following table.

NFS operation	Description
GETATTR <fh>	Returns the attributes of the file specified by the file handle
READ <fh, offset, count>	Reads the part of a file specified by the file handle
WRITE <fh, offset, count, data>	Writes data to a file.
READDIR <fh>	Returns a list of all of the files contained in a directory

NFS operation	Description
LOOKUP (fh, name)	Returns a file handle for a file or directory with a certain name found in a directory specified by a directory file handle in the request
ACCESS (fh, access)	Returns permission checking of future requests (NFS 3 only)

Caching granularity Cache space is expected to be significantly smaller than the size of the source file system. The appliance administers its own cache space by caching only parts of files.

Caching at a block level: The appliance caches at a block level. Caching is performed at a minimum of 4K block granularity. This means that, for example, if a client issues a read for three blocks of a file, the appliance is able to cache and serve hits on these three blocks of the file, even if the file itself is much larger than three blocks. Caching at this granularity results in efficient use of disk space and maximizes hits for “hot” objects, that is, objects that are frequently requested.

Data consistency The appliance supports varying levels of caching and varying levels of data consistency, depending on the data source.

The appliance acts as a sparse volume, which means that the appliance starts out empty and is populated as responses to client requests are cached. Only the data that has been requested by a user is cached. It is possible to have a file’s attributes cached but no data for the file.

Cached objects remain on the appliance as long as they remain consistent with the origin server. Server-side consistency is achieved in the following ways:

- ◆ Data is cached only when it is returned from the origin server.
- ◆ Data is invalidated on the appliance when any operation associated with a source file modification on the origin server is proxied through the appliance.

- ◆ Before the appliance serves a hit, based on consistency settings, the cached file attributes are verified with source file attributes. If the file has changed, it is discarded and the appliance forces a cache miss.

Note

NFS version 3 introduced a new subset of file attributes for verifying data consistency. Therefore, cache hit rates and overall performance for NFS version 3 is higher than that of NFS version 2.

Invalidation at a file level: Although caching granularity occurs at a block level, invalidation of cached data occurs at a file level. When the cache discovers that the attributes of the file have changed, it will eject the entire file. This means that if one block of a hundred block file is changed, the entire file will be ejected.

Deployment preparation

Forward- or reverse-proxy deployment

An NFS appliance can be deployed as a forward-proxy data cache or as a reverse-proxy accelerator.


Forward proxy: In a forward-proxy deployment, the appliance is remote from the data center where it caches popular data giving the end-user faster access to information.

Reverse proxy: In a reverse-proxy deployment or accelerator mode, the appliance is local to the administrative data center where it is used to offload work from busy file servers and free system resources.

Forward-proxy


In a forward-proxy deployment, the appliance is placed as close as possible to the remote office. Client requests are then explicitly directed to the appliance. This means that a client request for NFS is directed to the cache as opposed to the source file server. If valid data exists in the cache, that data or file is served directly to the client. If the data does not exist on the cache, it is retrieved from the source file server, written to disk on the appliance, and subsequently served to the client.

A forward-proxy deployment consolidates data at remote locations to free bandwidth and provide fast access to information. The following example illustrates such a forward-proxy deployment.



Reverse-proxy

In a reverse-proxy deployment, an appliance is used to offload busy data servers. Frequently accessed data, or “hot objects” are replicated and cached by the appliance. This supersedes full dataset replication by eliminating the requirement to push a full set of data to a remote location. Along with saving bandwidth, the appliance puts only the data that you need where you need it. The following example illustrates a reverse-proxy deployment.



Directing content to the appliance

A proxy configuration requires the administrator to explicitly provide the proxy IP address for NFS client, to direct requests to the proxy device. In this type of deployment, authenticated users know that the proxy exists and can simply remove the proxy information in order to bypass the cache.

Export volumes

Mapping volumes for export


Appliance export settings allow machines to mount disk partitions in the same way that an NFS server does. The appliance maps from a unique name to a volume on an NFS server. All deployments require you to configure at least one export on the appliance, but also allow you to configure multiple export volumes.

Export settings

Export settings are virtual and do not reflect actual volumes that exist on the cache. A volume on an origin server can be used by many exports. Once export settings are configured, MOUNT, NFS, and PORTMAP requests are then directed to the appliance, which either responds or proxies the request to the origin server, and then responds to the client.

How an NFS request is made

The following example illustrates the connect sequence or handshake that occurs the first time a client sends an NFS request to a server.



Stage	Description
1	Client asks the server (PORTMAP) what services are available.
2	Server (PORTMAP) responds that MOUNT v1 and v3, NFS v2 and v3, TCP, and UDP are available.
3	Client asks server (MOUNT) to mount /vol/vol0.
4	Server (MOUNT) responds with a file handle for /vol/vol0.

Stage	Description
5	Client requests server (NFS) to perform operation <i>x</i> on <file handle> with parameters <params>.
6	Server (NFS) responds with the <results of operation>.

Command options and syntax

The following table provides an in-depth description of each element in the export command syntax.

Command	Description
<virtual export name>	The name of the volume that a client will mount on the appliance itself.
<origin server name>	The name of the NFS server where this volume actually exists.
<origin server path>	The path to the volume on the NFS server.
<options>	There are several options you can specify per volume. Network Appliance supports the following mount options. Note_____ The following values are optional and can be configured from the Appliance Manager, or entered from the command line as a comma separated list. _____

Command	Description
<i>Maximum age</i> Time before reverification of objects CLI: max-age=#	This option enables you to set the amount of time the cache will answer a request without verifying the content with the origin server. If all of the traffic for this particular volume will be sent through the appliance, set the max-age value as high as possible. If you are modifying the volume on the origin server directly, the value should be set lower. If the value is set to 0, the cache revalidates every request. Standard client caches cache directories from 30 to 60 seconds and files from 10 to 30 seconds. Enter a number and an alphabetical character to indicate the unit (s = seconds, m = minutes, h = hours, d = days). The default setting, 30 seconds, is used when no unit is indicated.
<i>Disconnected operations</i> Type of behavior when servers are down CLI: hard, firm, or soft	This option allows you to determine the type of behavior the cache will exhibit when its connection to the origin server is lost. You can select from the list options: Hard: Cache blocks any operation that attempts to communicate with the origin server (default). Firm: Cache serves cached data only. All other attempts to communicate with the origin server are blocked. Soft: Cache serves cached data only, and returns error messages to the client for all other attempts to communicate with the origin server. Note_____ The consistency model for firm and soft settings is relaxed. This means that data is consistent, but content served is potentially out of date with respect to content on the origin server. _____
<i>Anon UID</i> Uid to use for requests with root uid CLI: anon=<uid>	This option enables you to specify the UID from which to translate requests from root. If you want root to be trusted on this volume, set anon=0. If you want root to be interpreted as the UNIX® user, set anon=65535 (default).

Command	Description
<i>Mount access</i> Hosts allowed to mount this volume CLI: access=<host list>	<i>Mount access</i> — This option enables you to set all mount requests to be rejected unless they come from the root user. Example: <hostname netgroup ip_address ip_address/bits> where host name is “company.com”, netgroup is “groupname”, IP address is “10.1.27.5”, and IP_address/bits is “10.1.27/24” The hosts listed are the only hosts allowed to mount the volume. The default is everyone.
<i>Root access</i> Hosts with root access to this volume CLI: root=<host list>	<i>Root access</i> — This option enables you to set a list of client IPs with root access to the specified volume. Example: <hostname netgroup ip_address ip_address/bits> where host name is “company.com”, netgroup is “groupname”, IP address is “10.1.27.5”, and IP_address/bits is “10.1.27/24” The hosts listed are the only clients that are trusted with user root.
<i>Read-write access</i> Hosts with read/write access to this volume CLI: rw=<host list>	This option enables you to set a list of client IPs with write access. Example: <hostname netgroup ip_address ip_address/bits> where host name is “company.com”, netgroup is “groupname”, IP address is “10.1.27.5”, and IP_address/bits is “10.1.27/24” The hosts listed are the only clients allowed to write to the file system. All other clients have only read-only access.
<i>Read-only</i> Restrict access for this volume to be read only CLI: ro	This option specifies to export a volume as read-only. The appliance rejects all operations that would modify the file system, such as writing a file or creating a directory.

Command	Description
<i>Bridging</i> Always use TCP to connect to origin server CLI: tcp	This option upgrades all connections between the appliance and server to TCP, regardless of the protocol that the client used to connect to the appliance. It is designed to be used over a WAN where UDP connections are unreliable. Note The appliance uses the same transport protocol to connect to the server that the client used to connect to the appliance. If you configure the export to use TCP to connect to the server, then the appliance uses TCP to connect to the server, regardless of what protocol the client used to connect to the appliance. If the server does not support TCP, the appliance detects this and uses UDP to connect to that server, regardless of what the client used.
<i>Verification schedule</i> Cron style entries for scheduling (reverification of objects) CLI: verify_sched= <cron entry>	This option enables you to set the schedule the appliance will use to determine when to reverify requested objects. A cron entry is a combination of the following: <minute><hour><day of month><month><year><day of week> or * for wild card. For example, 30 1 * * * 0 = every Sunday at 1:30 a.m. Use a larger max-age (the default max-age becomes infinite) if you can schedule times for objects on an export to become unverified.

Caution

The appliance access options supersede the origin server access options. On hits, only the appliance options root, anon, and access matter; on misses, the server and appliance options both matter. This means that unless the file server and the appliance have the same options, security holes could exist.

Configuring export volumes

Viewing export volumes from the GUI

This section describes configuring your appliance to export volumes. Before configuring the appliance to export volumes, you must find what, if any, volumes are currently being exported. The examples used in this section contain the following configurations:

- ◆ NFS server (filer.lan.company.com)
- ◆ Linux client (client.lan.company.com)
- ◆ Distributed NFS appliance (dnfs.appliance.company.com)

To check the server to see the volumes being exported on the client, complete the following steps.

Step	Action
1	From the Appliance Manager, select Setup tab > NFS > Exports.
2	View the existing configurations of exported and origin server volumes in the Configure Exported NFS Volumes table.

Viewing export volumes from a UNIX client

Many UNIX distributions include the `showmount` command. To view export volumes from a UNIX client, enter:

```
client [~] showmount -e filer.lan.company.com
```

The preceding command returns:

```
Export list for filer.lan.company.com:
/vol/vol0/home (everyone)
/vol/vol0      (everyone)
```

This shows that the filer (the origin server) is exporting two different volumes.

Again on the client, enter the following commands:

```
client [~] showmount -e dnfs.appliance
Export list for dnfs.appliance:
```

The preceding commands return:

Currently, the appliance is not exporting any volumes.

Configuring export volumes from the GUI

Continuing from the previous example, you can configure the appliance to export two volumes from the filer, the origin server.

To configure the appliance to export two volumes from the filer, complete the following steps.

Step	Action
1	Enable the NFS caching service by selecting Setup tab > NFS > General.
2	To configure a new volume, go to Setup tab > NFS > Exports and then click New Volume on the right side of the Configure Exported NFS volumes table. To edit an existing volume, locate the volume in the Configure Exported NFS Volumes table and click Edit.
3	Enter the basic options (these are mandatory). ◆ Enter the name of the volume to be exported on the appliance. Format: /vol/doom ◆ Enter the origin server and volume names. Format: /vol/build
4	Enter any advanced options that you want. These options include: Time before reverification of objects UID to use for requests with root UID Hosts allowed to mount this volume Hosts with read/write access to this volume Restrict access for this volume to read only Always use TCP to connect to origin server Cron style entries for scheduling a reverify
5	Click Commit Changes to have the changes take effect.

Configuring exports from the CLI

The `config.nfs.exports` option is used to create the mapping between volumes that you want the appliance to export and the actual volume on the origin NFS server.

To configure the export volume from the command line, enter

```
dnfs> set config.nfs.exports=\\
/vol/file filer.lan.company.com:/vol/vol0 anon=0,max-age=30
/vol/file/home filer.lan.company.com:/vol/vol0/home anon=0,max-age=100
\\
dnfs> set config.nfs.enable on
```

Configuring NFS mount points

To configure the export volume on the client, from the UNIX prompt

```
client [~] mount -o rw,bg,hard,intr,rsiz=8192,wsiz=8192
dnfs.appliance:/vol/file /t/dnfs.appliance/file
client [~] cd /t/dnfs.appliance/file
client [/t/dnfs.appliance/file]
```

To the client, this volume will look the same as `/vol/vol0` on `filer.lan.company.com`. An exception is if there are writes directly to the origin server that bypass the cache in the period between two reverifications.

Status and statistics commands

This section presents some useful commands for working with statistics and status information.

NFS status: Use the following commands to display NFS statistics.

```
dnfs> show status.nfs.bw_savings
status.nfs.bw_savings = 76%
or
```

```
dnfs> nfsstat
```

```
usage: nfsstat <interval>
```

```
usage: nfsstat -h [ <hostname or ip address> ]
```

```
usage: nfsstat -l
```

```
usage: nfsstat -z
```

```
usage: nfsstat
```

Client statistics: To display all of the NFS statistics for a specific client, enter

```
dnfs> nfsstat -h client.lan.company.com
```

Appliance statistics: Statistics can be useful to determine if traffic is sent through the appliance as expected. To display a running tally of current operations, enter

```
dnfs> sysstat 1
```

	CPU	.	ops/s	.	Net	kB/s	Disk	kB/s	BW savings/s
		tot	hit	miss	in	out	read	write	
	1%	0	0	0	0	0	84	132	0%
	0%	0	0	0	0	0	0	0	0%
	0%	0	0	0	0	0	0	0	0%

Option usage examples

A volume on an origin server can be used by many exports, as seen in the following example.

```
dnfs> set config.nfs.exports\\
/vol/home-ro      filer1.lan.company.com:/vol/home ro
/vol/home-tcp     filer1.lan.company.com:/vol/home tcp
/vol/special      filer1.lan.company.com:/thunder ro,max-age=15
/vol/hail         hail.lan.company.com/hail ro,max-age=4h
/vol/lightening   lightning.lan.company.com:/ max-age=0,anon=0
/vol/fireball     fireball.lan.company.com:/ verify-sched=0 0 * * * *
/vol/ice          ice.lan.company.com:/ verify-sched=30 1 * * * 0
\\
```

The following table shows the results of the preceding options.

Option	Result
/vol/home-ro	Exports /vol/home on filer1 read-only.
/vol/home-tcp	Exports /vol/home on filer1 as read/write, but upgrades all server requests to TCP.
/vol/special	Exports on filer1 as read-only, and the appliance serves hits on cached objects for 15 seconds before revalidating any of their attributes.
/vol/hail	Exports on hail as read-only, and the appliance serves hits on cached objects for four hours before revalidating any of their attributes.
/vol/lightening	Exports on lightening as trusting root on all clients, and the appliance serves hits on cached objects for 0 seconds before revalidating any of their attributes.
/vol/fireball	Objects are reverified on the first access after midnight daily.
/vol/ice	Objects are reverified after 1:30 a.m. on Sunday.

Disconnected operations

About disconnected operations

The distributed NFS appliance normally communicates with the origin server to

- ◆ Serve initial MOUNT requests
- ◆ Ensure that the cached content is current and consistent with that on the origin server (configured using the maximum-age and verification schedule options)
- ◆ Retrieve data that is cached (request misses)
- ◆ Write data that is not cacheable by definition

Disconnected operations is a mode that enables the distributed NFS appliance to serve as much data to clients as possible when a connection to the origin server is unavailable. Server unavailability might be due to any of the following conditions:

- ◆ The origin server is inoperable.
- ◆ The origin server is unreachable.
- ◆ Services on the origin server, such as NFS or exportfs permissions, are disabled.

If the previous conditions persist for more than a second or two, the distributed NFS cache switches to disconnected operations mode automatically.

Caching behavior during disconnected operations mode

Your distributed NFS appliance has limited file-serving capabilities when it operates in disconnected operations mode.

- ◆ Some functions, such as write and write-like operations and read operations of missing files or missing blocks of files, can be blocked or can return an error message to the client.

All volume export operations are treated as read-only (all file modification requests are rejected by the appliance).

- ◆ The appliance cannot verify object attributes prior to serving a request based on maximum age of content or verification schedules previously configured.

To monitor server health and connectivity, the appliance pings the origin server with a NULL NFS operation.

Troubleshooting inaccessible cached files

In addition to an unavailable or inoperable origin server, there are several reasons why files that had been cached might no longer exist on the appliance.

- ◆ Because of a client request, a file might be evicted from the cache to be replaced by the updated content on the origin server. During this window, the file becomes inaccessible.
- ◆ A file might be evicted from the cache because of low disk space.
- ◆ Write-type operations might make files obsolete and, therefore, inaccessible.
- ◆ Directories that include files that are deleted and created might also be evicted from the cache.

To reduce client opportunities that can flush content from a cache, you can make content available with a read-only export setting. The read-only export setting directs the appliance to reject all operations that would modify the file system, such as writing a file or creating directory. You can assign this export setting in the NFS - Exports page (Setup tab > NFS > Exports).

About this chapter

This chapter provides information about how to regulate administrative access to the appliance. It describes how to add and create groups of administrators. It also describes the authentication process, and how to control administrative access to the appliance using access control lists (ACLs). The appliance is permissive by default; that is, ACLs are not enabled and no requests are denied.

Topics in this chapter

This chapter provides information about the following topics:

- ◆ [“Adding administrator accounts”](#) on page 26
- ◆ [“Access control lists”](#) on page 29
- ◆ [“Forming rules”](#) on page 31

Procedures and option descriptions

See the Appliance Manager online Help for configuration procedures and option descriptions.

Location of configuration options

Options for adding users and configuring administrative access are in the following locations in the Appliance Manager:

Feature	Location
Administering users	Setup tab > Authentication
Controlling access	Setup tab > Access Control Lists

Adding administrator accounts

Accessing the Appliance Manager

The Appliance Manager is a browser-based administration tool that enables you to customize and maintain your configuration. To launch the Appliance Manager, open a Web browser and point it to

`http://hostname:port`

where *hostname* is the name of the appliance, and *port* is the Admin Port setting. The default Admin Port setting is 3132.

For example:

`http://cache1:3132/`

Default administrative user name and password: When you select an option from the Appliance Manager, you are prompted for a user name and password. By default, the user name is *admin*, and the default password is *NetCache*.

Administering access

User names, passwords, and group memberships are assigned using the authentication user and group options in the Appliance Manager. Use these options to do the following:

- ◆ Add users individually
Authenticated users who are not members of any groups are assigned default access permissions. There is only one set of default access permissions. Permissions cannot be modified on a per-user basis.
- ◆ Manage groups of users
Create groups, specify group memberships for users, set access permissions for the group, and edit user and group information for the local user database.

Note

A user can be a member of one group or multiple groups. Group access permissions can be set and modified on a per-group basis.

About authentication

Authentication is the process of determining a user's identity so that the user becomes known to the system. In addition to verifying a user's identity, authentication can be used in conjunction with access controls to enforce local security, allocate resources, and keep records of traffic.

The appliance authenticates a user for the first time when ACLs specifying authentication are applied, forcing authentication.

System administrators cannot access the Appliance Manager unless the user database contains at least one user name and password. The appliance provides a default administrator account and a default group called `NetCache`. The account user name is `admin` and the password is `NetCache`. All permissions assigned to the `NetCache` group apply to the default administrator account.

The default admin user account is configured with permissions to use all protocol services, as well as permissions to access the Appliance Manager and change the appliance configuration.

Note

The administrator user account user name and password are case-sensitive.

Using the Admin ACL option, you can create profiles with different levels of access for administrators. For example, you can apply different ACLs to different administrator accounts and allow some administrators to monitor the status of the appliance, but prevent them from making configuration changes.

When does the appliance authenticate a user?

The appliance authenticates a user for the first time when ACLs specifying authentication are applied. A user is authenticated against the local database. When authentication succeeds (that is, a match is found), any groups that a user belongs to are determined.

About the local user database

The user database contains administrator user account information that provides instructions for handling requests. Because it can contain group information, it can be used to assign access permissions based on group membership.

Groups

Grouping administrators is common practice and makes administering access less complicated than administering access by individual user. For example, the appliance can be configured to authorize access to administrative features for users in a senior administrative group, while prohibiting access to a junior group of administrators.

**Where
authentication
server information
is logged**

The appliance logs general information pertaining to authentication in the Messages log file. For example, if repeated authentication failures occur, the appliance logs a simple, explicit message noting that there were many authentication failures. You can use information such as this to assess if there might be a communication problem between the appliance and the authentication server rather than a true failed attempt to authenticate the user. The appliance logs this event as a connectivity error. Notification of this event in the log may indicate that a link is down between the appliance and the server.

Access control lists

What access controls do

Access control settings prevent unauthorized access to the appliance and restrict authorized access to services. The appliance supports different types of rules that impose a variety of restrictions. Sets of rules, or access control lists (ACLs), are used to impose restrictions inside the intranet domain.

Internet setting: In an Internet setting, access controls can be used to set rules that limit access for authorized users.

Intranet setting: In an intranet setting, access controls can help enforce authentication and authorize exactly what administrators have access to.

Authorized access

When a user makes an authenticated request, the appliance checks the user against all groups in the local user database, and attempts to match ACLs that are associated with those groups. If the request is not matched by user or group ACLs, then the request is checked against all other access control rules. If access is denied, a system message displays the rule and the name of the ACL that prevented access to services.

Controlling administrative access

ACLs can be applied explicitly to authorized users, but they are also effective for administering access control to individuals as part of a group, such as administrators. ACLs can be applied to members of administrative groups that are assigned “Admin, Full” and “Admin, Read-Only” access permissions.

ACL rule checking

The appliance sequentially reads ACLs and checks for the first matching entry. When a match occurs, the action is executed. If no match occurs, no action is executed and rule checking continues.

Rule checking sequence is shown in the following table.

Stage	Description
1	Check global settings that control authentication for protocol requests.

Stage	Description
2	Check ACL permission settings for the applicable group in the local user database.
3	Check request-type (protocol specific) ACL.
4	Check Web accelerator ACL.
5	Check global ACL.
6	Check group permissions settings.
7	Allow the request.

Order of precedence

Types of rules and order of precedence are described in the following table.

Order	Type	Description
1	User and group	User and group rules explicitly allow, deny, or redirect requests. User and group ACLs have precedence over request-type, global, and group access permissions.
2	Request-type (protocol specific)	Request-type rules control specific types of traffic. Access control rules are processed only for requests of a particular type, or protocol, such as NFS. Request-type access controls have precedence over Global access controls. If a request-type rule explicitly allows, denies, or redirects a request, the appliance follows this rule regardless of any rules configured in the Global Access Control option and any content category restrictions set in the SmartFilter options.
3	Accelerator	If a request originates from a Web-accelerator (reverse-proxy), this ACL is checked.
4	Global	Global rules impose restrictions on all traffic by allowing you to specify variables, such as a user name and a URL.
5	Group access permissions	Group access permissions specify default access to services for groups, and control rewrites or redirects for blocked requests. Default rules apply to requests that have no rule matches.

Forming rules

ACLs ACL entries allow or deny requests based on a number of variables, such as client IP addresses. ACLs are associated with specific users, and groups of users.

When ACL rules are formed, actions are combined with expressions that contain a variable-value pair. Variables and associated values form what is commonly referred to as a pattern-matching expression.

Syntax When applying actions (for example, allow, deny, redirect, rewrite, and auth) to matching expressions, use the `action expression` syntax, as shown in the following table.

Action	Expression
allow	client-ip 128.126.10.0
auth	group admin2
deny	server-domain net
allow	time 9am - 5pm
deny	group novice

Actions Actions and expressions form rules that control requests for access to Web content and supported protocols. The following table lists supported actions and describes how they govern client requests.

This action...	Specifies...
allow	Allow the request to proceed and disregard all the following rules.
auth	Require authentication and continue the rule checking sequence if already authenticated.
deny	Deny the request immediately.

This action...	Specifies...
<code>drop <ttl></code> where <i>ttl</i> is length of time in seconds. Default is 60 seconds.	Block packets from the source specified. For example, when used with the <i>client-ip</i> variable, packets from a matched client IP address are blocked at the IP level.
<code>no-log</code>	Selectively log and filter log file information by field. If applied as a global ACL, all logs are filtered. Entries in log files are replaced with a “-”. Note _____ Global authentication must be enabled. _____
<code>proxy</code>	Proxy session. Do not serve from cache.
<code>re-auth</code>	Force reauthentication.
<code>redirect <url></code>	Redirect the request immediately.
<code>rewrite</code>	Rewrite the URL, but continue the rule checking sequence.
<code>try-auth</code>	Require authentication and continue the rule checking sequence even if authentication fails.

Variables

Variables are assigned values that form expressions, for example, *time 9am - 5pm*. The following table lists variables and provides a description and example of each.

Variable/Description	Example
<code>accel</code> Web Accelerator requests	<code>deny accel</code>
<code>any</code> Any requests	<code>deny any</code>
<code>auth</code> Authenticated request	<code>auth any</code>

Variable/Description	Example
cache-ip IP address of this NetCache	allow cache-ip 206.80.5.0/24
cache-port Port number on this NetCache	deny cache-port 8080
client-domain Client domain name	deny client-domain .wisc.edu
client-ip Client IP address	deny client-ip 206.79.5.51
client-port Client port number	deny client-port 25
contains Substring search operator Note _____This variable supports user, group, req-headers, server-name, client-domain, and server-domain variables.	deny url contains bomb
contains-case Case-sensitive substring search operator Note _____This variable supports user, group, server-name, client-domain, and server-domain variables.	deny url contains-case VIOlence
date Date and time	deny group novice and date "sun nov 6 08:49:37 2002"
day Day	allow day mon - fri

Variable/Description	Example
group Multidomain support for groups. Note_____ A group name with no domain name matches any group with that name, regardless of the domain.	<pre>auth any allow group admins deny any</pre> This ACL allows access for members of group “admins” in all trusted domains
group\ Multidomain support for groups. Note_____ Using the “\” delimiter between the domain name and group name enables domain comparison.	<pre>auth any allow group "ADMINS\globalgroup" deny group "NOVICE\globalgroup" deny any</pre> This ACL allows access for members of ADMINS, and denies access for members of NOVICE. Globalgroup is a trusted domain.
http HTTP requests	<pre>deny url "http://www.corp.com"</pre>
server-domain Server domain name	<pre>allow server-domain .mycorp.com</pre>
server-ip Server IP address	<pre>deny server-ip 206.79.5.0/24</pre>
server-name Server name in URL	<pre>allow server-name www.mine.com</pre>
server-port Server port number in URL	<pre>deny server-port 1 - 1024</pre>

Variable/Description	Example
time Time Note ————— The time variable syntax is restrictive, and is supported only when expressed as a range or as part of a regex matching expression.	allow time 9am - 5pm
url URL string	deny url http://www.mine.com/
url-prefix Prefix of URL string	deny url-prefix http://www.mine
user User name of authenticated user	deny user bob

Expressions

Expressions are single predicates, variable-value pairs, range matches, or regular expressions (regex). The following table provides an example of each expression type.

Type	Definition	Example
Single predicate	<code><variable></code>	server-ip
Variable-value pair	<code><variable> <value></code>	client-ip 128.148.37.17
Range matching	<code><variable> <value> - <value></code>	server-port 1- 1024
Regular expression matching (regex)	<code><variable> matches <regex></code>	“^http://[0-9.]+/redcodeworm.asp\?X+”

Note—————

When you create rules, you must enclose values that contain white space or nonalphanumeric characters in double quotation marks.

Operators

Complex expressions are built using parentheses and boolean operators, such as `and` or `&&`, `or` or `||`, and `not`. For example:

- ◆ `<expr> and <expr>`
- ◆ `<expr> or <expr>`
- ◆ `(<expr>)`
- ◆ `not <expr>`

Caution

Using the operator `not` in combination with the action `allow` (for example, `allow not user user1`) does not result in an action to deny requests when the rule is applied.

Use caution when setting rules

The type of rule and the number of rules that the appliance must check before servicing each request can affect performance. For example, rules containing the `server-ip` variable prompt a cached DNS lookup that might be slower than using the `server-name` variable to specify server names.

In general, order your rules so that specific restrictions are listed before broad restrictions. Carefully plan the access restrictions you want to impose. Configure only the rules that are required for your environment, and where possible, collapse several rules into a single rule.

About this chapter	This chapter discusses the types of interfaces supported on the appliance, how the interfaces are named, and the concept of virtual interfaces.
Topics in this chapter	This chapter provides information about the following topics: ◆ “Planning summary for network interface configuration” on page 38 ◆ “About physical network interfaces” on page 39 ◆ “Trunking” on page 41 ◆ “Virtual interfaces” on page 43 ◆ “Troubleshooting network interfaces” on page 46
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions.
Location of configuration options	Options for configuring network interfaces are in the following location in the Appliance Manager: Setup tab > Network > Interfaces

Planning summary for network interface configuration

Determine network interface configuration

Configuring network interfaces will be faster if you plan ahead. Use the information in the following table to determine the network interface configuration that is necessary for your organization.

Decisions to make	Your options	Considerations
Whether you need to configure additional physical network interfaces	During the setup program, configuration of at least one physical interface is required.	You can configure additional physical interfaces after setup if they are needed.
Whether to configure one or more virtual interfaces	You can configure up to eight virtual interfaces.	Advantages of using virtual interfaces are as follows: ◆ Higher throughput because of multiple interfaces working as one interface. ◆ Fault tolerance. If one interface in a trunk goes down, the appliance can still stay connected to the network using the other interfaces. ◆ Protection against a switch becoming a single point of failure.
Whether you need to change the state of a physical interface	The state of any interface configured during setup is up by default.	◆ You must bring down any physical interface that is to be configured as part of a virtual interface. ◆ Before performing maintenance on the link, bring the interface down. When your work is complete, bring the interface back up.

About physical network interfaces

Supported interfaces

The appliance supports the following network interfaces:

- ◆ Physical interfaces
 - ❖ FDDI
 - ❖ Ethernet
 - ❖ Gigabit Ethernet
- ◆ Virtual

Note

For detailed information on virtual interfaces, see “[Virtual interfaces](#)” on page 43.

Naming conventions for network interfaces

The appliance uses the following interface naming conventions to refer to each network interface:

- ◆ Ethernet and Gigabit Ethernet interface names start with the letter *e* followed by the slot number (for example, *e0*) and the port number (if the adapter is a quad-port Ethernet card, for example, *e1a* and *e1b*).

Note

Port numbering on quad-port cards starts from the top port.

- ◆ FDDI interface names start with the letter *f* followed by the slot number (for example, *f1* and *f2*).

The appliance automatically creates host names for network interfaces by appending the slot number of the interface to the appliance’s host name. For example, if the appliance is named *toaster*, the host name for *e0* is *toaster-e0*, and the host name for *e1a* is *toaster-e1a*. If the interface is an FDDI interface, the letter *f* precedes the slot number. For example, the host name for *f1* is *toaster-f1*.

When to configure a network interface

During the setup program, at least one network interface must be configured. After setup, you can do the following:

- ◆ As necessary, configure additional network interfaces, for example, if you are bringing up an additional interface on your appliance
- ◆ Edit the options previously configured for a physical interface

- ◆ Change the state of a physical interface

Note

You do not need to configure a physical interface that you want to include in a virtual interface. The only requirement is that the state of that physical interface is down. See “[Virtual interfaces](#)” on page 43 for more information.

Trunking

About trunking

Trunking is a mechanism that groups together multiple network interfaces (links) into one logical interface (trunk). Different vendors refer to trunks by the following terms:

- ◆ Virtual aggregations
- ◆ Link aggregations
- ◆ Virtual interfaces (the term Network Appliance uses)
- ◆ EtherChannel

The appliance implementation closely follows the IEEE 802-3ad specification.


Advantages of trunking

Trunks provide several advantages over individual interfaces, such as the following:

- ◆ Higher throughput because of multiple interfaces working as one interface.
- ◆ Fault tolerance. If one interface in a trunk goes down, the appliance can still stay connected to the network using the other interfaces.
- ◆ Protection against a switch becoming a single point of failure.


Appliance interfaces before trunking

The following illustration shows four separate appliance interfaces, *e3a*, *e3b*, *e3c*, and *e3d*, before trunking.



Appliance interfaces after trunking

The following illustration shows the four appliance interfaces joined after trunking into a single trunk called *Trunk1*.



Virtual interfaces

About virtual interfaces

The feature that supports trunking on the appliance is known as virtual interface. After being created, a virtual interface is indistinguishable from a physical network interface.

Kinds of virtual interfaces

There are two kinds of virtual interfaces:

- ◆ Single-mode
- ◆ Multi-mode

Single-mode virtual interface operation


In a single-mode virtual interface, all interfaces that constitute the virtual interface share an IP address and a media access control (MAC) address. However, only one interface in the virtual interface is active. The other interfaces are on standby, ready to take over if the active interface fails. Failure means that the link status of the interface is down, which signals that the interface has lost connection with the switch.

Note

When you create a single-mode virtual interface, an interface is selected randomly to be the active interface. However, if you want to specify a certain interface to be active, you can override the random selection and specify that interface as *preferred*. For example, if you want to make a standby or a new interface active, you must make that interface the preferred interface. There can only be one preferred interface in a single-mode virtual interface.

A single-mode virtual interface can recover from a failure of up to $(n-1)$ interfaces, where n is the total number of interfaces that form the virtual interface.

Example: In the following illustration, *e0* and *e1* are part of the SingleVIF1 single-mode virtual interface. The active interface, *e0*, fails. The *e1* interface takes over and maintains the connection to the switch.



Multi-mode virtual interface operation

In a multi-mode virtual interface, all the interfaces in the virtual interface are active at the same time. This logical aggregation of links provides higher throughput than possible if the interfaces were configured as a single-mode virtual interface.

With a multi-mode virtual interface, the appliance load balances outgoing traffic over the outgoing interfaces that are trunked. The switch automatically balances the load among the interfaces of the virtual interface if the following conditions are met:

- ◆ The switch supports trunking over multiple ports
- ◆ The switch is configured so that all ports to which links of a virtual interface are connected are part of a single logical port and share a common MAC address.


For information about configuring the switch, see your switch vendor's documentation.

Note

Some switches might not support trunking of ports configured for jumbo frames. For more information, see your switch vendor's documentation.

A multi-mode virtual interface can recover from a failure of up to $(n-1)$ interfaces, where n is the total number of interfaces that form the virtual interface.

Example: In the following illustration, *e0*, *e1*, *e2*, and *e3* are part of the MultiVIF1 multi-mode virtual interface. All four interfaces in the MultiVIF1 multi-mode virtual interface are active.



In the preceding illustration, if any three of the interfaces were to fail one-by-one or simultaneously, the appliance still stays connected to the network.

Routing on an appliance with a multi-mode virtual interface

If this physical interface goes down or is removed from the multi-mode virtual interface, load balancing occurs over the remaining interfaces.

Prerequisites to creating a virtual interface

Before creating a virtual interface, you need to

- ◆ Identify the interfaces you want to aggregate into the virtual interface on your appliance.
- ◆ If you are creating a multi-mode virtual interface, configure a switch that supports trunking over multiple port connections according to your switch vendor's instructions.

Troubleshooting network interfaces

Use the `ifconfig` command for troubleshooting

When you are troubleshooting an appliance, you can use the `ifconfig` command to perform common interface configuration. See the *na_ifconfig* man page for more information.

Caution

Configuration changes made using the `ifconfig` command do not persist after a system reboot. To make changes persistent across a system reboot, use the `ifconfig -set` command.

Recovering from loss of connectivity to the appliance

If you bring down the administrative interface or all network interfaces, you can recover in one of the following ways:

If you have a backup configuration file: From the command line at the console, use the `config -rl` command to revert to a backup copy of the appliance configuration.

If you do not have a backup configuration file: Use the `ifconfig` command from the command line at the console to bring up an interface. Then use the Appliance Manager to bring up and, if necessary, reconfigure the other network interfaces.

About this chapter

This chapter discusses how the appliance uses default and static routes for incoming and outgoing requests, binds traffic to interfaces, and communicates with hosts and subnets using TCP large windows. This chapter also describes how to configure routing options.

Topics in this chapter

This chapter contains the following information:

- ◆ [“Planning summary for routing configuration”](#) on page 48
- ◆ [“How the appliance handles routing”](#) on page 50
- ◆ [“Binding types of traffic to specific interfaces”](#) on page 52
- ◆ [“Restricting interfaces for connecting to servers”](#) on page 55
- ◆ [“Using TCP large windows to improve performance”](#) on page 56

Procedures and option descriptions

See the Appliance Manager online Help for configuration procedures and option descriptions.

Location of configuration options

Options for configuring network interfaces are in the following location in the Appliance Manager:

Setup tab > Network > Routing

Planning summary for routing configuration

Determine routing configuration

Configuring routing will be faster if you plan ahead. Use the information in the following table and the remainder of the chapter to determine the routing configuration that is necessary for your organization.

Decisions to make	Considerations	Option to use in the Appliance Manager
Configuring basic routing support		
Whether you need to add static routes	You need to enter static routes in the appliance's routing table only if your default gateway is not aware of some addresses and you are not able to route packets to those addresses through RIP or IRDP.	Setup tab > Network > Routing > Static Routes tab
Whether you want the appliance to listen to Router Discovery messages	If your routers are configured to support IRDP, Network Appliance recommends that Route Discovery is enabled. Enabling Route Discovery eliminates the need to manually configure the default gateway.	Setup tab > Network > Routing > General tab > Route Discovery Enable option This option is disabled by default.
Controlling interfaces over which the appliance sends out and accepts traffic		
Whether you need to restrict interfaces through which the appliance accepts client requests	Typically, the appliance accepts client requests through any interface and IP address. You can add a level of security by using the Listen Addresses option to restrict the IP addresses, and therefore the interfaces, on which the appliance accepts client requests.	Setup tab > Network > General > Listen Addresses option Do not use this option if the appliance is configured for takeover.

Decisions to make	Considerations	Option to use in the Appliance Manager
Whether you need to balance traffic from servers to the appliances over appliance interfaces	Organizations that need balanced responses from servers across multiple lines might want to use the Load Balancing Interfaces option to balance traffic across interfaces.	Setup tab > Network > General > Load Balancing Interfaces option
Whether you want to restrict the interfaces from which the appliance can communicate with servers	Disable outgoing interfaces in the following circumstances: ◆ If you want to devote an interface for administrative purposes instead of using it for regular traffic. ◆ If one interface of an appliance is outside a firewall and another interface is inside the firewall, typically, you do not want to use the interface inside the firewall for traffic to outside servers.	Setup tab > Network > General Disabled Outgoing Interfaces option
Considerations for high-bandwidth, long-haul connections		
If you have a high-bandwidth, long-haul connection, whether TCP large windows configuration would be beneficial	TCP large windows improves performance over high-bandwidth, long-haul connections by using the capacity of the pipe more effectively.	Setup tab > Network > General > TCP Large Windows Enable and TCP Large Windows List

How the appliance handles routing

The appliance does not function as a router

Although the appliance can have multiple network interfaces, it does not function as a router; that is, the appliance does not route packets between its interfaces on behalf of other network hosts.

How the appliance routes its own packets

For routing its own packets, as well as dynamic routes, the appliance relies on the default route and static routes. Normally, the appliance learns explicit routes through Internet Control Message Protocol (ICMP) redirect messages received from the default router; you do not need to enter static routes in the appliance's routing table.

Configuration considerations: You need to enter static routes in the appliance's routing table only if your default gateway is not aware of some addresses and you are not able to route packets to those addresses through RIP or IRDP.

Management of network routing tables

The appliance uses the routed daemon to manage the network routing tables. Routed is a simple routing daemon that is enabled at boot time. Routed uses the Routing Information Protocol (RIP), which listens for packets being exchanged between routers on the network and solicits multicast router discovery Internet Control Message Protocol (ICMP) messages. The appliance uses RIP response packets to update its routing tables. The routing tables are updated with gateways that are currently recognized on the network, or with gateways that advertise a better metric than at least one of the gateways currently listed in the routing tables.

Configuration considerations: RIP is enabled by default. Disabling RIP stops the appliance from listening for RIP response packets, and information in the routing table remains static.

About Route Discovery

The Internet Router Discovery Protocol (IRDP), which is an extension of ICMP, enables hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers. When the Route Discovery feature is enabled, the appliance listens to Router Discovery messages, thereby eliminating the need to manually configure the default gateway. If the default router goes down, IRDP

informs the appliance about alternative routers, ensuring that the appliance maintains its connection to the Internet. Only routers that support IRDP can be discovered through the Route Discovery feature.

Configuration considerations: Route Discovery is disabled by default. If your routers are not configured to support IRDP, do not enable this option.

Binding types of traffic to specific interfaces

How the appliance routes requests on specific interfaces

By default, the appliance binds all types of traffic to all interfaces. You can, however, bind specific types of traffic to specific interfaces to balance the load across interfaces or to provide an additional level of security. Any time you bind traffic to a particular interface, the appliance sends or accepts that type of traffic only on that interface.

How the appliance typically accepts requests

Typically, the appliance accepts client requests through any interface and IP address, including any interface and IP address for an Internet link. You can add a level of security by restricting the IP addresses and, therefore, the interfaces on which the appliance accepts client requests.

Note

Interfaces can have more than one IP address.

Changing appliance behavior

Using the Listen Addresses option, you can limit the appliance IP addresses on which the appliance listens for client requests. If a TCP or UDP IP address is specified in this option, the appliance listens only on that IP address. If the Listen Addresses box is left blank, the appliance listens on all of its TCP or UDP interfaces.

Option location: The Listen Addresses option is on the Network - General page (Setup tab > Network > General).

Determining whether you should use this option

Consider the following when determining whether to use this option:

- ◆ Only the IP addresses associated with the interfaces of this appliance can be listed in this option.
- ◆ You must enter a full IP address, not a network address or a subnet address.
- ◆ Each interface for an IP address listed in this option must be configured as up; otherwise, the appliance will not listen on it.
- ◆ You cannot use this option if this appliance is part of a takeover pair.

Example of how Listen Addresses can be used

You can use the Listen Addresses option to allow client traffic from your internal network and deny traffic from the Internet

If your appliance is inside your firewall and your firewall does not block access from all clients external to your organization, you can specify in the Listen Addresses option only the IP addresses of the interfaces to your organization's internal networks. Because the IP address for the interface to the Internet is not listed in the Listen Addresses option, the appliance will not listen on that interface. Therefore, clients on the Internet cannot access your appliance.

Implication of source IP address on incoming traffic

Typically, all requests from the appliance are sent out using a single-source IP address, even if multiple interfaces are configured. Because the appliance has only one default gateway, outbound requests usually are sent through a single interface. The routing table determines the interface on which the request is sent *out*.

The source address of the outbound request determines which interface the origin server *replies* to and, therefore, is the IP address to which the server sends the requested object.

If the appliance has multiple interfaces, you can configure their IP addresses as outbound IP addresses so that traffic can be shared across the interfaces. Because the outbound traffic is spread over multiple IP addresses, incoming packets in reply to the request are directed to the appropriate interface based on the source IP address and, therefore, balanced over the interfaces.

Load balancing server-to-appliance traffic

You can use the Load Balancing Interfaces option to specify all interfaces through which the appliance communicates when initiating outgoing connections. Each interface listed participates in the load balancing of outbound traffic by splitting traffic across outbound links. The end result is that the incoming traffic is split over appliance interfaces.

If no IP addresses are listed in the Load Balancing Interfaces option, the appliance originates all outgoing requests from its primary interface.

Note

You cannot use the Load Balancing Interfaces option if your appliance is part of a takeover pair.

Option location: The Load Balancing Interfaces option is on the Network - General page (Setup tab > Network > General).

Who should use this option

This technique is useful for organizations that need to balance responses from servers across multiple lines. It is not useful for organizations that need to balance requests to servers across multiple lines. If you need to balance requests to servers across multiple lines, you must use another method, such as policy routing on the default gateway.

Restricting interfaces for connecting to servers

About the Disabled Outgoing Interfaces option

The Disabled Outgoing Interfaces option specifies interfaces that the appliance never uses when connecting to servers. Do not list the interfaces in this option that are listed in the Load Balancing Interfaces option.

Disable outgoing interfaces in the following circumstances:

- ◆ If you want to devote an interface for administrative purposes instead of using it for regular traffic.
- ◆ If one interface of an appliance is outside a firewall and another interface is inside the firewall, typically, you do not want to use the interface inside the firewall for traffic to outside servers.

Note

You cannot use the Disabled Outgoing Interfaces option if your appliance is part of a takeover pair.

Option location: The Disabled Outgoing Interfaces option is on the Network - General page (Setup tab > Network > General).

Using TCP large windows to improve performance

When using TCP large windows is beneficial

Long-haul networks have more capacity than can be used effectively with the standard TCP window size. Enabling TCP large windows dramatically improves performance over high-bandwidth, long-haul connections by using the capacity of the pipe more effectively. Using TCP large windows on local area networks, however, has no benefit.

Examples of networks that can benefit from enabling TCP large windows include the following:

- ◆ Two appliances communicating with each other over a satellite link
- ◆ Clients connecting to an appliance using DSL but with a high-capacity hop between them and the appliance

Determining which hosts or subnets to enable TCP large windows for

When the TCP Large Windows feature is configured, the appliance can communicate with specified hosts and subnets using TCP large windows. You can specify the systems or subnets with which the appliance should communicate using large windows, and also to specify the window size that the appliance should use with each.

Exercise caution when deciding which hosts or subnets the appliance should communicate with using large windows. Using large windows consumes more memory than using the standard window size; therefore, the appliance limits the number of simultaneous connections using large windows. The bigger the window size, the fewer simultaneous large windows connections the appliance can support.

Network Appliance recommends that you specify only a few hosts or subnets with which the appliance should communicate using large windows.

Option location: Setup tab > Network > General > TCP Large Windows Enable and TCP Large Windows List.

Determining the window size to use

To use the capacity of the link effectively, the window size should be greater than or equal to the *bandwidth-delay product*. Calculate the bandwidth delay product by multiplying the capacity of the link (bandwidth) by the round-trip time (delay), as follows:

$$\text{window size} \geq \text{bandwidth} \times \text{delay}$$

Setting the window size precisely to optimize the usage of the link can require some experimentation, however.

Guidelines for common connection types: If you are using a T1 or T3 connection, you can use the window sizes in the following table.

Connection type	Recommended window size
T1	11264 bytes
T3	345088 bytes

About this chapter	This chapter discusses setting up an appliance for host name resolution.
Topics in this chapter	This chapter contains the following information: ◆ “Basic DNS configuration” on page 60 ◆ “Optional DNS configuration” on page 61 ◆ “Handling DNS lookups for local servers” on page 65
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions relating to DNS.
Location of configuration options	To access the DNS options in the Appliance Manager, select: Setup tab > DNS

Basic DNS configuration

Host name resolution

The appliance uses DNS to resolve host names. The following information must be provided to ensure support of DNS:

- ◆ Host name
- ◆ Domain name
- ◆ DNS nameservers to which the appliance can send requests for name resolution

The first time the appliance boots, it broadcasts a request to the network to determine whether a DHCP server is available to provide specific information about the network. If a DHCP server replies, the appliance accepts the following information from the DHCP server, if the DHCP server has been configured to provide it:

- ◆ IP address
- ◆ Default gateway
- ◆ Domain name
- ◆ DNS nameservers

The appliance provides this information as default settings during the setup program.

If a DHCP server is not available, you are prompted for this DNS information during the setup program. You can change any of these settings later, if necessary, on the DNS - General page in the Appliance Manager. To access this page, select Setup tab > DNS > General.

Failover of nameservers

Entering multiple nameservers in the DNS configuration provides failover capability if a nameserver is unavailable. For example, if the first nameserver in the list goes down, the appliance tries the second nameserver, and so on. The appliance does not try the second nameserver if the first nameserver reports a DNS lookup failure. This behavior follows DNS standards.

Optional DNS configuration

Whether you need to handle domain names that are not fully qualified

If you expect that users will enter domain names that are not fully qualified, you should configure the appliance to handle such names. Users in enterprises usually know that they can enter partially qualified domain names in URLs. Therefore, system administrators of enterprises should consider configuring the appliance to handle domain names that are not fully qualified.

ISPs, as a rule, do not need to configure support for domain names that are not fully qualified because their users typically enter fully qualified domain names in URLs.

Handling domain names that are not fully qualified

The following options on the DNS - General page in the Appliance Manager can be configured to handle domain names that are not fully qualified:

- ◆ **DNS Tree-Search Enable**

Enables and disables a domain tree search for DNS lookups. When this option is enabled, the appliance uses the paths specified in the DNS Search Paths option, and a DNS lookup occurs even if the domain name is fully qualified. Keep in mind that a domain-tree search can generate a large amount of network traffic if the servers for the listed domains are not local.

- ◆ **DNS Search Paths**

Identifies the paths that the appliance should follow when trying to resolve a host name. The appliance follows the order in which the search paths are listed in this option, from the top to the bottom of the list.

Extra DNS lookups can occur with DNS tree searching

By default, when the appliance looks up a domain name in its DNS cache, it first attempts to look up the domain name with the local domain name appended, even if the requested domain name is fully qualified. If the first lookup fails, the DNS resolver then performs a lookup without the local domain name appended. The result is that the resolver performs an extra lookup when the URL contains a fully qualified domain name.

Note

The concept of fully qualified differs between DNS and most application-level protocols. For DNS, a domain name is fully qualified only if a dot appears at the end of the domain name. For example, although `www.books.com` appears to be a fully qualified URL, it is not fully qualified according to the DNS specification because no dot appears at the end. For this reason, the appliance appends the local domain name to the domain name.

If you expect that your users will always provide fully qualified domain names, Network Appliance recommends that you do not enable the DNS Tree-Search Enable option. The reason is that when the DNS Tree-Search Enable option is enabled, a DNS lookup occurs even if the domain name is fully qualified.

Avoiding an extra DNS lookup can be beneficial for the following reasons:

- ◆ A domain tree search can be slow.
- ◆ A domain tree search generates a large amount of network traffic if the servers for the listed domains are not local.
- ◆ Queries time out if no server is available for one of the domains.

What occurs when the DNS Tree-Search Enable option is enabled

The following table describes what occurs for domain names when the DNS Tree-Search Enable option is enabled. Assume for this example that the local domain is `books.com` and a valid host name is `web.books.com`.

If the domain name in the URL is...	Then...
Unqualified (that is, it does not contain dots). For example: <code>web</code>	If the unqualified domain name is <code>web</code> , for example, and the DNS Tree-Search Enable option is enabled, the first lookup is for <code>web.books.com</code> , which succeeds.
Partially qualified. For example: <code>web.books</code>	If the partially qualified domain name is <code>web.engr</code>, for example, the following occurs: 1. <code>books.com</code> is appended to the domain name so that the lookup is on <code>web.books.books.com</code>. This lookup fails. 2. The appliance then tries <code>web.books.com</code>. This lookup succeeds.

If the domain name in the URL is...	Then...
Fully qualified. For example: web.books.com	If a URL contains the fully qualified domain name web.books.com, for example, the following occurs when the DNS Tree-Search Enable option is enabled: 1. books.com is appended to the domain name so that the first lookup is for web.books.com.books.com. This lookup fails. 2. The second lookup is for web.books.com.com. This lookup fails. 3. The third lookup is for web.books.com. This lookup succeeds.

What occurs when the DNS Tree-Search Enable option is disabled

The following table describes what occurs for domain names when the DNS Tree-Search Enable option is disabled. Assume for this example that the local domain is books.com and a valid host name is web.books.com.

If the domain name in the URL is...	Then...
Unqualified (that is, it does not contain dots). For example: web	The local domain name is not appended to the domain name in the URL. This lookup fails.
Partially qualified. For example: web.engr	The local domain name is not appended to the domain name in the URL. The lookup of a partially qualified domain name fails right away. For example, web.engr is looked up as web.engr, and the lookup fails.
Fully qualified. For example: web.books.com	The DNS lookup succeeds the first time.

About identifying DNS search paths

By configuring DNS search paths, you increase the probability that the appliance can successfully resolve a host name if the domain name in the URL is not fully qualified. The reason is that the appliance can expand its search beyond the local domain. When the DNS Search Paths option is configured, the appliance tries to resolve the host name by searching in the following order:

1. Domain name
2. Search paths in this option, from the top to the bottom of the list

Identify DNS search paths if you have enabled the DNS Tree-Search Enable option.

Handling DNS lookups for local servers

About reducing DNS lookups on remote DNS nameservers

You might be able to reduce the number of DNS lookups performed by remote DNS nameservers by entering local host names and their associated IP addresses in the Hosts File option on the DNS - General page. Configuring the Hosts File option enables your appliance to handle DNS lookups for local servers. As a result, you do not need to configure host names and IP addresses for your local servers on remote DNS nameservers.

When to add entries to the Hosts File option

Typically, you do not need to add entries to the Hosts File option. Add entries to this option only if you have a special requirement to bypass the DNS nameserver.

For example, if you do not want to export DNS information to a DNS nameserver, you can add the host name and IP address to the Hosts File option. If you are setting up this appliance as an accelerator, you could, for example, add entries to this option for the servers that the appliance is accelerating. The DNS lookup would then be performed on the appliance, eliminating the need for a DNS lookup on a remote DNS nameserver.

What is checked first for DNS host name resolution

When the appliance tries to resolve a host name, it checks the local DNS cache first, then checks the entries in the Hosts File option. The appliance sends DNS queries to the remote DNS nameserver only if it cannot resolve the host name from the DNS lookup of the entries in the Hosts File option.

About this chapter

This chapter describes disk management on appliances. The information in this chapter covers disk management on appliances that support RAID (Redundant Array of Independent Disks) Level 4.

Sections in this chapter

This chapter contains the following sections:

- ◆ Section A, “[RAID](#),” on page 68
- ◆ Section B, “[RAID disk management](#),” on page 77
- ◆ Section C, “[Reinitializing all disks](#),” on page 83

Procedures and option descriptions

See the Appliance Manager online Help for configuration procedures and option descriptions related to the following disk maintenance topics:

- ◆ Adding disks to your system
- ◆ Removing disks from your system
- ◆ Replacing a failed disk on your system

Location of configuration options

Options for configuring the features discussed in this chapter are in the following locations in the Appliance Manager.

Feature	Location
Disk management	Maintenance tab > Disk > Management
Firmware management	Maintenance tab > Disk > Firmware
Qual Devices management	Maintenance tab > Disk > Qual Devices

Section A: RAID

About this section

This section contains information about RAID Level 4, different types of disks used in RAID, disk addressing, usable space on disks in the RAID group, and handling disk failures on your appliance.

Topics in this section

This section contains the following topics:

- ◆ [“Understanding RAID”](#) on page 69
- ◆ [“Understanding disk addresses”](#) on page 71
- ◆ [“Understanding usable space on each disk”](#) on page 73
- ◆ [“Using fibre channel disks of various sizes”](#) on page 74
- ◆ [“Handling disk failure”](#) on page 75

Understanding RAID

About RAID

Appliances (with the exception of the C1100 and C1105 systems) use RAID Level 4 to ensure data integrity even when one of the disks fails. The file system design, together with the support for RAID, optimizes appliance performance and enables you to expand the appliance's disk storage capacity incrementally.

The RAID disk structure on your appliance is factory-installed; you do not need to configure the storage system before using the appliance. However, when planning to optimize disk space or when replacing or adding disks, you must be familiar with how the appliance uses its disk array.

In the RAID file system, different disks have different functions. Most of the disks are *data disks*. One disk in the array is the *parity disk*, which enables the appliance to recover the data on a data disk if one fails. You can also install one or more disks to be used as *hot spares* in the event of disk failure.

Data disks

Data disks are those on which cached objects, logs, and other data are stored.

Parity disks

After a data disk failure, the appliance uses the *parity disk* in conjunction with the other data disks to reconstruct the failed disk's data and write it to a hot spare disk. The parity disk must be at least as large as the largest data disk. For each 28 disks added to an appliance, one drive is used for parity.

For more information about how the appliance recovers from disk failures by using the parity disk or hot spare disk, see [“Handling disk failure”](#) on page 75.

Hot spare disks

In addition to data and parity disks, you can install one or more *hot spare disks* in the appliance. A hot spare disk is not part of the RAID disk array and does not contain cached objects. After a disk failure, the appliance automatically rebuilds data onto a hot spare disk, which then replaces the failed disk in the RAID array. Using a hot spare disk avoids a system shutdown and returns the system to full performance.

A hot spare disk cannot replace a failed disk that is larger than itself. If you use only one hot spare disk in an appliance, the hot spare disk must be as large as the largest file system disk. If you have multiple hot spare disks installed, the system uses the smallest hot spare disk needed to replace the failed disk.

Hot swapping disks

Some models of the appliance allow you to *hot swap* disks. Hot swapping a disk means installing or removing it from the disk shelf while the appliance is running with minimal interruption to the service. For example, you might want to hot swap a disk into an appliance to replace a disk or to add a hot spare disk.

Understanding disk addresses

About disk addresses

You identify a disk by its address, which is listed in the Disk ID column of the Disk - Management page (Maintenance tab > Disk > Management) of the Appliance Manager. The Disk ID is represented in *x.y* format, where *x* refers to the host adapter and *y* refers to the disk ID number.

You use the disk address to do the following:

- ◆ Interpret system messages (for example, command output or error messages) that you see on your screen
- ◆ Locate a disk that a system message refers to

Locating a disk by its identifier

Use the Disk - Management page (Maintenance tab > Disk > Management) of the Appliance Manager to display disk information about all the disks on your appliance. This page lists the disk ID, the type (data, parity, or spare), the size, and the current status of every disk on your appliance.

This information is useful when you need to locate a disk referenced by the RAID disk number in a system message.

The following is an example of the Disk - Management page output.

Disk - Management

Use this page to view the type, size, and status of the disks installed on this NetCache Appliance. Additionally, use this page to perform disk management tasks, such as adding, removing, and replacing disks on this NetCache Appliance.

Insert or Remove Disk Drive: Swap Drive

To manually insert a disk drive into or remove a disk drive from the NetCache Appliance, first click Swap Drive. Only disk drives whose status is "Bad" can be manually removed from the NetCache Appliance.

Disk Status

Host Adapter: v0 Shelf: 0 Type: Fibre Channel

Bay	Disk ID	Size (MB)	Disk Type	Disk Operations	
0	v0.0	70	spare	Fail	Promote
1	v0.1	70	data	Fail	
2	v0.2	70	data	Fail	
3	v0.3	70	data	Fail	
4	v0.4	70	data	Fail	
5	v0.5	70	data	Fail	
6	v0.6	70	parity	Fail	

SCSI disk addresses

For SCSI disks, the Disk ID is a combination of the disk's SCSI adapter number and the disk's SCSI ID number. For example, disks attached to disk adapter 0 are number 0.0 through 0.6.

Fibre channel disk addresses

For fibre channel disks, the Disk ID is a combination of the disk's adapter number and disk's fibre channel loop ID. To create the fibre channel loop ID, multiply shelf ID by eight and add it to the bay number. For example, host adapter 8, shelf 1, disk 2 has disk ID of 8.10.

Understanding usable space on each disk

What usable disk space is

A disk’s *usable* space can be different from its *physical* space. Disks from different manufacturers might differ slightly in size even though they belong to the same size category. To ensure that you can replace disk from one manufacturer with a disk from another manufacturer, the appliance software automatically “right-sizes” its disks.

The appliance manages disk space so that objects never entirely fill the usable space. The following table shows the approximate usable disk space for each class of disk supported by the appliance.

Disk class	Usable disk space
1-GB disks	1,000 MB (2,048,000 blocks)
2-GB disks	2,000 MB (4,096,000 blocks)
4-GB disks	4,000 MB (8,192,000 blocks)
9-GB disks	8,600 MB (17,612,800 blocks)
18-GB disks	17,000 MB (36,864,000 blocks)
36-GB disks	34,500 MB (70,656,000 blocks)
72-GB disks	68,000 MB (139,264,000 blocks)

Understanding the overwriting-objects policy

The appliance uses a first-in, first-out (FIFO) replacement policy for writing cache objects. The FIFO policy operates by tracking inodes and using object ejection (truncation) to manage space. This process takes place asynchronously during the serving of requests. When an appliance disk reaches its upper capacity, objects are ejected until there is sufficient space to record new objects.

Using fibre channel disks of various sizes

Guideline for using fibre channel disks of various sizes

You can use fibre channel disks of various sizes with the same appliance, up to and including 18-GB, 36-GB, and 72-GB disks. However, unlike filers, the appliance does not support mixing disks of different sizes on the same disk shelf; they must be in separate shelves.

The following table shows a sample configuration.

Shelf	Data	Data	Parity	Hot Spare
Shelf 1	Disk 0 (9 GB)	Disk 1 (9 GB)	Disk 2 (9 GB)	Disk 3 (9 GB)
Shelf 2	Disk 0 (18 GB)	Disk 1 (18 GB)	Disk 2 (18 GB)	Disk 3 (18 GB)

Handling disk failure

About disk failure

The effects of a disk failure on the appliance operation depend on whether the appliance is equipped with a hot spare.

Disk failure with a hot spare

If the appliance has a hot spare installed, the system enters a state called “degraded mode” after a single disk failure. In this state, the RAID feature allows the appliance to continue to run without losing data (although the appliance’s performance is affected). In addition, while in degraded mode, parity calculations are performed for READs, and the appliance rebuilds the missing data on the hot spare in the background with minimal interruption to service.

The appliance logs this activity in the Messages log and does not automatically shut down. If you power off the appliance while it is in degraded mode, it stops data reconstruction. After you power on the appliance again, it restarts the data reconstruction process from the beginning.

Except for a slight loss in performance while data is rebuilt on the spare disk, the failure of a single disk is transparent to the user. The appliance exits degraded mode and returns to normal operation after it finishes reconstructing data.

Network Appliance recommends that you replace the failed disk with a new hot spare after the appliance has finished reconstructing data. This way, the appliance continues to have a hot spare that can be used in case another disk fails.

The Disk - Management page (Maintenance tab > Disk > Management) in the Appliance Manager displays which disk is reserved as the hot spare.

In addition to logging disk failure and hot spare replacement activity, the Messages log on the Log Files - Messages page (Data tab > Log Files > Messages) logs any failure in the periodic check of the hot spare.

Disk failures without a hot spare

If the appliance is not equipped with a hot spare, when a disk fails, the appliance enters a degraded mode. In this state, the RAID feature allows the appliance to continue to run without losing data, although the appliance’s performance is affected.

The system shuts down after 24 hours to encourage you to replace the disk.

Caution

Replace the failed disk as soon as possible. A second disk failure will cause the entire cache of objects to be lost, and you will have to reinstall the file system.

When a disk fails, the appliance logs one of the following warning messages in the Messages log and to the system console every hour:

Parity disk in x , RAID group y , is broken. Halting in m hours.

Data disk in volume x , RAID group y , is broken. Halting in m hours.

(In the actual message, x is the volume number, y is the RAID group number, and m is the number of hours before the system halts.)

In addition, immediately before the system halts, a message similar to the following is sent to the Messages log and the system console:

```
Sat Oct 29 13:26:42 PDT [statd]: When a disk is broken, the
system shuts down automatically every 24 hours to encourage you
to replace the disk. If you reboot the system it will run for
another 24 hours before shutting down.
```

The shutdown ensures that you notice the disk failure. You can restart the appliance without fixing the disk, but the appliance continues to shut itself off every 24 hours until the disk is replaced.

What happens to the data on the failed disk: If no hot spare disk is available, the failed disk data will not be reconstructed but will be inaccessible until you replace the failed disk. The failed disk data is reconstructed on the new disk.

What happens with a double disk failure: If a second disk fails while the appliance is operating in degraded mode, or if two disks fail at once, the file system is damaged and cannot be repaired. You must then reinstall the appliance software and restore the configuration from a backup. For more information, see [“Backing up and restoring the appliance configuration”](#) on page 119.

What happens when the parity disk fails: When the parity disk fails, the appliance operates in degraded mode until you replace the failed disk. After you have installed a replacement disk, the appliance uses the information on the data disks to rebuild the parity information onto the replacement disk.

Section B: RAID disk management

About this section

This section contains information required to maintain disks on RAID-enabled appliances.

Topics in this section

This section contains the following topics:

- ◆ [“Displaying the status and statistics of disks”](#) on page 78
- ◆ [“About adding and removing disks”](#) on page 80
- ◆ [“Troubleshooting a failed disk”](#) on page 82

Displaying the status and statistics of disks

Understanding the disk management display

You can view the status of disks on your appliance using the Disk - Management page (Maintenance tab > Disk > Management) of the Appliance Manager.

The following table describes the different columns in the Disk - Management page output.

Output heading	Description
Host Adapter	Indicates the host adapter number, which is the number of the slot in which the disk adapter is installed.
Shelf	<i>Shelf</i> is the number of the disk shelf that is connected to the disk adapter interface on the host adapter. Multiple disk shelves can be connected to a disk adapter interface if the disk adapter is a differential SCSI adapter or a fibre channel adapter. The disk shelf number starts at 0.
Type	Indicates the type of disk adapter, for example, Fibre Channel and SCSI.
Bay	<i>Bay</i> is the number of the disk slot in the disk shelf, which starts at 0.
Disk ID	Disk ID is represented in the <i>x.y</i> format, where <i>x</i> is the host adapter and <i>y</i> is the disk number. For more information, see “ Understanding disk addresses ” on page 71.
Size	Size of the disk in megabytes. The disk sizes are “right-sized,” as explained in “ Understanding usable space on each disk ” on page 73.
Disk Type	Indicates the current function of the disk. A disk can be a data, spare, or parity disk.

Output heading	Description
Disk Operations	The operations allowed on the disks in the appliance. The Fail operation changes the state of a disk to “Bad.” The Promote operation promotes a spare disk to a file system disk.

Understanding disk statistics

You can view the disk statistics of your appliance using the System Status - General page (Data tab > System Status > General) in the Appliance Manager. The options that display disk statistics on the System Status - General page are as follows:

- ◆ Disk Space in Use—Displays the amount of disk space consumed by objects stored on the appliance.
- ◆ Total Disk Space—Displays the total amount of data disk space available on an appliance. The disk space indicated by this option does not include the disk space of the parity disk.

About adding and removing disks

Why you might add or remove disks

You might want to add disks to the appliance to increase available disk space, to add a hot spare, or to replace a failed disk. You might want to remove a disk from your appliance because you want to use it in another appliance, because it has failed, or because it is producing excessive error messages.

Adding new and previously used disks

In most cases, the disk you add will be a new disk; that is, the disk has never been used.

If the disk you want to add to the appliance was previously used by another appliance or a Network Appliance filer, make sure that you *unfail* the disk before you begin. To unfail the disk, boot the appliance from a system boot diskette into Maintenance mode, then use the `disk unfail` command.

If you want to add a new disk to an existing system, first add the disk as a hot spare. Then, if you want the disk to be used as a data disk, promote the disk to be part of the file system.

When to add or remove disks in external disk platforms

The Appliance Manager enables you to add or remove a single disk from the system without interrupting service to your users. However, during the procedure, the appliance operates in proxy-only mode. That is, the appliance proxies requests on behalf of users by retrieving the objects from the source server. No new objects are saved to the cache, and no objects in the cache are served to users. Because users might notice slower performance during this time, it is a good idea to plan to add or remove disks when you expect the load on the appliance to be light.

If you are adding more than one disk or a new disk shelf to your system: If you are adding several disks or a new disk shelf to your system, the process might be faster if you shut down the system. If you shut down the system, you can add all the disks at once, then reboot the system. Otherwise, you must add the disks one at a time using the Appliance Manager.

If you must remove more than one disk from your system: If you must remove more than one disk from the appliance, Network Appliance recommends that you remove and replace that disk before proceeding to remove another disk.

If you remove more than one disk at a time from the appliance, the appliance will not be able to rebuild the data from the remaining disks. You will have to reinstall both the file system and the appliance software in that case.

If you are performing disk operations online: If you are performing disk operations online, you can perform only one operation at a time. The appliance does not allow you to begin a new disk operation if a previous disk operation is incomplete.

Using hot spares in external disk platforms

If you want to prevent loss of service time, you can add one or more hot spare disks. When a data disk on the appliance fails, a hot spare automatically replaces it without a system shutdown.

To determine whether your appliance contains a hot spare, see the Disk - Management page (Maintenance tab > Disk > Management) of the Appliance Manager. This page lists all appliance disks and indicates whether a disk is a parity, data, or spare disk.

You can add more hot spare disks to the system at any time, as long as disk slots are available.

Troubleshooting a failed disk

Recovery from disk replacement failure

If a SCSI disk fails during a disk hot swap operation, the appliance displays the following warning and then halts:

*** WARNING ***

A disk has failed during a disk swap operation. The system is halting to allow recovery. Please power off disk storage shelves for 15 seconds to completely spin down the disks. Then reapply power to shelves and reboot the system.

A system reboot can usually clear transient disk errors that occur during the disk hot swap. However, to recover from the error if it persists after the system reboot, complete the following steps.

Step	Action
1	Disconnect the disk shelf used during the hot swap from the power source and wait 15 seconds. This allows all drives to spin down.
2	Reconnect the power source to the disk shelf.
3	At the system monitor prompt, enter b to reboot the system. After the reboot, your change to the disk array takes effect.
4	Contact Network Appliance Technical Support if you encounter any further disk problems during the hot swap. When you call, be prepared, if possible, to provide information about your appliance, such as the output from the <code>sysconfig -v</code> command, messages in the Messages log, and a precise description of how the hot swap was performed.

Section C: Reinitializing all disks

When you need to reinitialize disks

You might need to reinitialize all disks in the following circumstances:

- ◆ You decide to redeploy an existing appliance and you need to completely reconfigure it.
- ◆ Network Appliance Technical Support advises you that the only way to recover from an error is to reinitialize all disks.

Caution—————
Reinitializing all disks causes all existing data to be lost.

To reinitialize all disks, complete the following steps.

Step	Action	
1	If the console is...	Then...
	Displaying the appliance prompt, for example, <code>netcache1></code>	Place the system boot diskette into the diskette drive of the appliance and enter reboot .
	Not displaying the appliance prompt	Reboot from diskette. ◆ If the console is showing the ok prompt, enter boot. ◆ If the console is unresponsive, enter boot floppy. If the console is still unresponsive, reset the appliance by turning it off, leaving it off for 30 seconds, and then turning it back on.
2	When the boot menu appears, enter 4 to choose Initialize all disks. After the initialization process finishes, the setup script runs. You will be prompted for the appliance IP address and other information.	

Step	Action
3	Upload the new software image through the Appliance Manager. See the Appliance Manager online Help for instructions. Until you upload the new software image, no operating system exists on the disks. Therefore, do not reboot the appliance until after you have uploaded the new operating system and have been prompted to reboot.

About this chapter	This chapter provides information about the Network Information Service (NIS) feature and provides information about how to configure NIS servers for centralized administration.
Topic in this chapter	This chapter provides information about the following topic: ◆ “NIS database administration” on page 86
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions.
Location of configuration options	Options for adding users and configuring administrative access are in the following locations in the Appliance Manager: Setup tab > NIS

NIS database administration

About NIS

NIS is a distributed database system that is used to centralize administration for a wide variety of network configuration information. NIS can be configured on the appliance to provide host name to IP address resolution, IP address to host name resolution, and netgroup authentication for access control of NFS clients.

A distributed NIS database consists of a single host or master server, and client servers with replicate databases. All administrative and database changes are performed on the master server and then propagated to NIS client servers, giving each host on the network knowledge about the entire system.

Netgroups

Netgroups are a NIS component used to associate a group of hosts with a group name. Netgroups are configured on the master NIS server and processed into netgroup maps that are propagated to client NIS servers.

Configuring NIS from the GUI

To configure NIS on the appliance, complete the following steps.

Step	Action
1	In the Appliance Manager, select Setup tab > NIS.
2	On the NIS - General page, ensure that the NIS Enable check box is selected. This enables resolution of hosts and netgroups through NIS.
3	Enter the name of the NIS domain to which the appliance belongs. The name must match the domain served by the specified NIS domain servers. For example: companyname.com

Step	Action
4	List IP addresses of NIS servers that serve the NIS domain specified in the NIS Domain Name option. All servers specified are monitored by health checks. If a connection is lost, failover to any of the listed NIS servers will occur. If no servers are specified, the appliance broadcasts on all interfaces and listens for NIS server responses. The appliance then binds to the NIS servers that respond.

Configuring NIS from the command line

To configure NIS from the command line, enter

```
set config.system.nis.enable <on|off>
set config.system.nis.domainname [<nis-server-ip>]
set config.system.nis.servers [<nis-server-ip> ...]
```

Administrative CLI commands

After you have configured on the appliance, you can use the following CLI commands to display status and configuration information:

Command	Description
<code>nis info</code>	Displays the status of NIS client services. This command displays the NIS domain name, IP address, binding type, state, binding status, last poll time, and number of pending client calls.
<code>ypmatch</code>	Prints NIS map key values.
<code>ypwhich</code>	Prints the name of the active NIS server.
<code>ypcat</code>	Prints all values in the NIS map.

About this chapter

This chapter describes rules that allocate various-size bandwidth “pipes” on your appliance and assign specific categories of connections to use these bandwidth pipes. Doing so enables you to restrict the share of your appliance’s total bandwidth capacity that can be used by any one category of connections.

You can also use the rules described in this chapter to completely deny or allow categories of connections through the appliance, without bandwidth restriction.

Topics in this chapter

This chapter provides information about the following topics:

- ◆ [“About bandwidth allocation”](#) on page 90
- ◆ [“Viewing existing bandwidth allocation rules”](#) on page 91

Procedures and option descriptions

See the Appliance Manager online Help for configuration procedures and option descriptions.

Location of configuration options

Options for configuring network interfaces are in the following location in the Appliance Manager:

Setup tab > Network > Bandwidth Allocation

About bandwidth allocation

What is bandwidth allocation?

The bandwidth allocation feature enables you to control the total bandwidth capacity that specified categories of connections can collectively consume on the appliance.

Using the Appliance Manager fields associated with bandwidth allocation, you can:

- ◆ Set aside portions of your total bandwidth capacity into various-size capacity *bandwidth pipes*.
- ◆ Assign connections matching a specified set of properties, such as protocol, source or destination IP addresses, and inbound (server side) or outbound (client side) communication to a bandwidth pipe.

All connections assigned to a specific bandwidth pipe share the total bandwidth space allocated for that bandwidth pipe. The appliance enforces the bandwidth limit, if necessary, by using buffers to space out both incoming and outgoing packets.

As a result, TCP connections, as well as protocol connections based on UDP, gracefully use the specified bandwidth pipe with a minimum of burstiness, allowing the appliance to enforce the bandwidth allocation rules efficiently.

Denying or allowing connections

You can also use bandwidth allocation rules to deny service to a category of connections completely, or allow service to that same category of connections with no bandwidth restriction whatsoever.

Practical uses for the bandwidth allocation feature

Practical uses for the bandwidth allocation management feature include the ability to

- ◆ Restrict aggregate bandwidth by single protocol (IP, TCP, UDP, and NFS)
- ◆ Limit aggregate bandwidth to or from a client or origin server
- ◆ Limit the bandwidth usage on any given interface

Viewing existing bandwidth allocation rules

Viewing the bandwidth allocation rules

If you have existing bandwidth allocation rules defined, you can view these rules in the Current Bandwidth Allocation Rules table.

Description of the Current Bandwidth Allocation Rules table

The following table provides descriptions of the fields in the Current Bandwidth Allocation Rules table.

Field	Description
Delete	Click this field to delete the listed aggregate bandwidth management rule.
Edit	Click this field to edit the listed bandwidth allocation rule. Clicking this field displays the Edit Bandwidth Allocation Rule page.
Rule #	Shows the number of the listed bandwidth allocation rule. Earlier-numbered rules are applied first and have precedence over later-numbered rules in the event of bandwidth pipe assignment conflict.
Protocol	Shows the protocol of the packets to which the listed rule applies.
Action	Shows the action associated with this rule. ◆ allow — Indicates that a particular flow is permitted without bandwidth pipe restriction. ◆ deny — Indicates that a particular flow is completely denied service. ◆ other — Indicates that the packet flow associated with this rule is assigned to a bandwidth pipe. The bandwidth pipe ID number and the bandwidth space of the pipe (indicated in bits or bytes per second) are displayed.

Field	Description
Filtering Options	Shows additional parameters that restrict the set of connections to which this bandwidth allocation rule applies. This field might show the following parameters. ◆ <code>from addr [port]</code>—Indicates that this rule is restricted to connections from a specified source IP address and optional port number. ◆ <code>to addr [port]</code>—Indicates that this rule is restricted to connections directed at a specified destination IP address and optional port number. ◆ <code>user username [, username]</code>—Indicates that this bandwidth allocation rule applies to the specified user or users. ◆ <code>not user username [, username]</code>—Indicates that this bandwidth allocation rule does not apply to the specified user or users. ◆ <code>group groupname [, groupname]</code> —Indicates that this bandwidth allocation rule applies to the specified group or groups of users. ◆ <code>not group groupname [, groupname]</code> — Indicates that this bandwidth allocation rule does not apply to the specified group or groups of users. ◆ <code>server traffic</code>—Indicates that this rule is restricted to traffic between this appliance and upstream servers from which this appliance retrieves requested data. ◆ <code>client traffic</code>—Indicates that this rule is restricted to traffic between this appliance and its clients or other downstream nodes. ◆ <code>in</code>—Indicates that this rule is restricted to incoming packets. ◆ <code>out</code>—Indicates that this rule is restricted to outgoing packets. ◆ <code>time - time</code> —Indicates that this rule is restricted to connections that are active during the specified time period. ◆ <code>via interface</code>—Indicates that this rule is restricted to traffic transmitted from or received at the specified interface.

First-numbered rules take precedence

If multiple bandwidth allocation rules are configured for one appliance, the rules are applied in the order in which they are numbered in the Current Bandwidth Allocation Rules table. If a category of connections is included in two bandwidth allocation rules that assign it to two different bandwidth pipes, the earliest numbered rule is the one that determines the bandwidth pipe to which this category of connections is assigned.

About this chapter	This chapter describes how SNMP (Simple Network Management Protocol) works on the appliance.
Topics in this chapter	This chapter contains the following information: ◆ “About SNMP” on page 94 ◆ “About the MIB” on page 96 ◆ “Understanding SNMP traps” on page 98 ◆ “Creating a user-defined SNMP trap” on page 99
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions relating to SNMP.
Location of configuration options	To access the SNMP options in the Appliance Manager, select: Setup tab > Administration

About SNMP

What the SNMP process does

If SNMP is enabled, SNMP managers might query the SNMP agent for information (specified in the Network Appliance custom MIB or the MIB-II specification) from time to time. In response, the SNMP agent gathers information and forwards it to the SNMP managers using the SNMP protocol. The SNMP agent also generates traps, whenever specific events occur, and sends these traps to the SNMP managers. The SNMP managers might then carry out actions based on information received in the trap notifications.

SNMP terms

To understand the SNMP information in this section, review the following terms and definitions:

Network management stations: Client workstations on a network that use third-party applications to process information.

Agent: The process that gathers status and diagnostic information and forwards it to network management stations.

MIB: ASCII files that describe the information that the agent sends to network management stations.

Trap: User-defined process that monitors a specific occurrence on the appliance, for example, a trap that monitors appliance disks.

Communities: Name that is used as a password by the SNMP manager to communicate with the SNMP agent.

SNMP agent and MIB groups supported

For diagnostic and other network management services, the appliance supports an SNMP version 1 compatible agent. This agent provides support for the MIB-II specification and a Network Appliance custom MIB. The following MIB-II groups are supported:

- ◆ System
- ◆ Interfaces
- ◆ Address translation
- ◆ IP
- ◆ ICMP

- ◆ TCP
- ◆ UDP
- ◆ SNMP

Note

Transmission and EGP MIB-II groups are not supported.

The Network Appliance custom MIB is discussed in “[Groups and objects in the custom MIB](#)” on page 96.

SNMP limits

The appliance supports the following SNMP limits:

- ◆ Specification of up to eight communities
- ◆ Trap notifications for up to eight management stations

About the MIB

Information contained in the MIB

The NetApp custom MIB provides detailed information about many aspects of the appliance operation.

Where to obtain the MIB

You can obtain the custom MIB from the software CD-ROM or from the NOW Web page at <http://now.netapp.com/> (perform a search for *MIB*). To locate the MIB on the CD-ROM, read the *index.htm* file in the top-level directory of the CD-ROM. You will need to install the MIB file on your network management station.

However, to obtain trap variable (OID) information from the MIB on the appliance, you can use the Appliance Manager MIB browser.

Installing the MIB

Install the MIB file on your network management workstation according to the installation procedure for your workstation to ensure that your workstation can obtain information from the appliance about the objects that are part of the MIB.

Groups and objects in the custom MIB

The following table describes the local groups and object types contained in the MIB.

Group name	Description
product	Product information, such as software version and system ID
sysStat	System-level statistics
Appliance options	Appliance objects you can directly affect: ◆ Whether the appliance is enabled on this system ◆ Whether the appliance is licensed on this system ◆ Whether the Domain Name System (DNS) is enabled on this system ◆ Whether NFS is licensed on this system

Group name	Description
Appliance info	Appliance information that requires rebooting of the device in order to change the object value: ◆ The appliance version string ◆ Appliance Admin Port number, the first in the list of admin ports configured for the appliance ◆ The number of entries in the acceleration monitor table ◆ Variable for user-defined SNMP traps to monitor whether there has been a status change of any of the acceleration servers that the appliance is probing ◆ Variable used in conjunction with amMonitor; it returns a string in the format of “IP:portstatus”, which reports the status of all the servers the appliance is accelerating since the last status change. ◆ A unique value for each entry in the acceleration monitor table ◆ The IP address of the acceleration server that the appliance is monitoring ◆ The corresponding port on the server that the appliance is monitoring ◆ The status of the service running on the server: active (in queue, but no probes performed yet); up (the server is reachable and responding to probes); down (the server is unreachable); inactive (monitoring is disabled) ◆ Whether the local configuration file on the appliance was changed without using DataFabric™ Manager configuration management ◆ The current cache configuration's version

Understanding SNMP traps

About traps

You set traps to inspect the value of MIB variables periodically. Whenever the value of a MIB variable meets the conditions you specify, a trap is sent to the network management stations on the trap host list. The trap host list specifies the network management stations that receive the trap information.

You can set traps on any numeric variable in the MIB. For example, you can set a trap to monitor the disks on an appliance and have the SNMP application on your network management station put a flashing message on your console that tells you that a disk has failed.

Traps are persistent. After you set a trap, it exists across reboots until you specifically remove it or modify it.

Types of traps

There are two types of traps:

- ◆ Built-in—Built-in traps are predefined in the appliance and are automatically sent to the network management stations on the trap host list if an event occurs. These traps are based on
 - ❖ RFC 1213, such as coldStart, linkDown, linkUp, and authenticationFailure
 - ❖ Specific traps defined in the Network Appliance custom MIB, such as diskFailedShutdown, cpuTooBusy, and overTempShutdown.
- ◆ User-defined—User-defined traps exist only after you create them using the Appliance Manager or the command-line interface.

Where to get more information

See the following RFCs (Requests For Comments) for more information:

- ◆ RFC 1157—Defines and describes SNMP.
- ◆ RFC 1213—Defines and describes the SNMP MIB-II specification.

Creating a user-defined SNMP trap

About user-defined traps

You can create custom SNMP traps to send specific types of information to your management station.

Viewing existing SNMP traps

To view existing user-defined traps, in the Appliance Manager, select Setup tab > Administration > SNMP, then click the Traps tab. Current user-defined traps are listed under the heading Current SNMP Traps. (The heading is displayed as “No Current SNMP Traps” if no traps have been defined.)

Identifying a MIB variable to be used with a new trap

To create a trap, you need to determine the information you want the trap to report on, and then determine the MIB variable that tracks that information. The NetApp MIB lists available variables with short descriptions.

To select a variable and its object identifier (OID) to be used for the trap, you can use the Appliance Manager MIB browser by clicking the Show MIB Browser button on the SNMP page.

You can also obtain the custom MIB from the software CD-ROM or from the NOW Web page at <http://now.netapp.com/> (perform a search for *MIB*). To locate the MIB on the CD-ROM, read the index.htm file in the top-level directory of the CD-ROM.

Obtaining the SNMP object identifier (OID) for a MIB variable

Each variable is associated with an OID. You can obtain the OID from the MIB browser by directly clicking on the OID; this imports it into the OID box on the New SNMP Trap page.

Alternatively, you can obtain the OID from the traps.dat file in the same directory as the MIB after you have determined which variable name you want the new trap to use and create the trap manually.

Parameters for user-defined traps

Use the descriptions in the following table to help you construct entries for trap parameters.

Parameter	Description
Trap Name	Specify the name of the new trap being defined. Valid characters include alphanumeric characters (a through z, A through Z, 0 through 9), the underscore (_), and the hyphen (-). The first character must be alphabetic. Internal spaces are not accepted.
Trap Enable	Enables the trap instance being defined. You do not have to enable the trap at this time. You can define the trap on this page and then enable it later on the Current SNMP Traps option on the Traps tab page.
Variable	Specify the object identifier (OID) for the MIB variable this trap should report on. Use the MIB browser to import the correct OID or enter the OID manually. A list of available variables is located in the NetApp MIB. Example: The OID for the MIB variable sysUptime is as follows: <pre>snmp.1.3.6.1.2.1.1.3.0</pre>
Trap ID	Optionally, specifies a trap ID value within the valid range of values reserved for the traps you define. You can use the value 0 (default) or a value within the range 1000 - 1999 inclusive. Specify a short name, or alias, for the trap you are defining. Remaining values in the ID space are reserved for Network Appliance and are not available for your use. If you specify no trap ID or if you specify 0 (the default value), triggering this trap will cause the appliance to issue a specific code based on the priority of this trap. The Trap ID parameter assigns a trap identifier (trap-id) to the information being sent to the management station. This enables the management station to signal alerts to administrators and indicate where the alert originated. Each trap defined in the Network Appliance custom MIB has its own identifier. For each user-defined trap, the management station uses the trap-id associated with the trap's priority.

Parameter	Description
Trigger	Specify the type of trigger that signals this trap to fire and send information to the management station. When a trap is triggered, data regarding the event that caused the trigger is sent to the network management stations. You can specify the following values: ◆ single-edge-trigger—Fires once when the value of the trap's MIB variable meets or crosses the threshold that you specify. You specify the threshold in the Edge-1 field. If the MIB variable value remains equal to the specified value or if it continues to exceed the threshold in the same direction, the single-edge-trigger will not fire again until the value crosses the threshold in the opposite direction and then crosses over again (once again exceeding the threshold specified in the Edge-1 field). The trigger fires each time the variable's value crosses the specified threshold in the specified direction. For a single-edge-trigger, only Edge-1 direction applies. You can set the direction to Up or Down, specifying the crossover direction in which the trigger occurs. ◆ double-edge-trigger—Fires when the value of the trap's MIB variable exceeds or drops below a threshold you specify. A double-edge trigger enables you to set two edges, each with its own direction (Up or Down). You specify these threshold values in the Edge-1 and Edge-2 options. For a double-edge-trigger, Edge-1 direction (either Up or Down) specifies the crossover direction that fires the trigger for one edge, and Edge-2 direction (either Up or Down) specifies the crossover direction that fires the trigger for the other edge. (Down is useful, for example, if you monitor for underutilization.) ◆ Level-trigger—Fires whenever the trap is evaluated if the value of the trap's MIB meets or exceeds the level that you specify, or remains at the level that you specify after meeting it. Specify the level in the Edge-1 option. ◆ Change-trigger—Keeps track of the last value received from the trap. If the current value differs from the previously received value, the trap is triggered. For example, the trap fires when the value changes from 0 to 1. ◆ Always-trigger—Instructs a trap to trigger every time the trap is evaluated.

Parameter	Description
Edge 1 and Edge 2	The Edge 1 and Edge 2 parameters are used in combination with the Trigger parameter values of single-edge trigger and double-edge trigger. Specify the threshold value or values that are compared during trap evaluation to determine whether to fire the trap and send data to the management station. The Edge 1 parameter specifies the value for the edge in a single-edge triggered trap or the first edge in a double-edge triggered trap. The default value for the Edge 1 parameter is MAXINT, or 2147483647. The Edge 2 parameter specifies the value for the second edge in a double-edge triggered trap. The default value for the Edge 2 parameter is 0.
Edge-1 Direction and Edge-2 Direction	Edge-triggered traps fire only when the edge threshold value is crossed in one direction. The Edge Direction parameters enable you to set or change the direction that is used to evaluate a trap. The edge-direction value specifies the direction (Up or Down) in which the crossover trigger is expected to occur. The default value for Edge-1 Direction is Up, and the default value for Edge-2 Direction is Down.
Interval	The Interval parameter is the time, in seconds, between evaluations of the variable associated with a trap. A trap can send data only as often as its associated variable is evaluated. The default value for the Interval parameter is 3600. Note The value of the Interval parameter for the Network Appliance predefined traps is one minute.
Interval Offset	The Interval Offset parameter is the time, in seconds, after the appliance starts up to delay the initial evaluation of this trap. This prevents too many traps from being evaluated at once immediately after the appliance starts up. The default value for the Interval Offset parameter is 0.

Parameter	Description
Rate-Interval	Enables you to specify the interval for monitoring the rate of change of a trap's MIB variable (specific object) over a period of time relative to a threshold value for that object. When the threshold is exceeded, the trigger fires the trap to notify you. Note The Interval option causes SNMP to evaluate the value of the object steadily every <i>x</i> number of seconds without relation to overall duration or threshold, but the Rate-Interval option causes SNMP to evaluate the object relative to a period of time that you set bounded by a threshold value. Example: You might use Rate-Interval to ensure that the amount of HTTP requests received by your appliance in one hour (360 seconds) does not exceed 90000 (requests). If the HTTP traffic surpasses that number of requests, you want to be notified (through triggering of a trap). The Rate-Interval option is used in conjunction with the Trigger option, the Edge-1 option, and the Edge-1 Direction option. These additional options enable you to define the threshold and direction for the trap variable value (object), which, if exceeded, causes a trap to fire that notifies you. (For more complex cases, you can use both directions, Edge-1 and Edge-2.) To measure the rate of HTTP requests per hour to cause a trigger to signal a trap to fire and send you a message when the threshold is exceeded, you would set the following cooperating options: ◆ Rate-Interval = 360 (seconds = 1 hour) ◆ Trigger = Single-edge trigger ◆ Edge-1 Direction = Up ◆ Edge-1 = 90000 Using the Rate-Interval option, you can scale the amount of time and the threshold values accordingly to the interval that you want to measure. Rate-Interval can be used also with Double-edge triggers, with either edge using either direction for its threshold boundary crossing. For example, you might want to use a Down direction to test for underutilization of an appliance. The default is 0.

Parameter	Description
Backoff Calculator	The Backoff Calculator parameter enables you to specify how often you want the trap evaluated after it fires the first time. After a trap fires and sends data, you might want the trap to be evaluated less frequently. For example, you might want to know when a file system is full within one minute, but want to be notified only once an hour thereafter that the file system is still full. The possible values for the Backoff Calculator parameter are as follows: ◆ Backoff Step ◆ Backoff Multiplier ◆ No Backoff The default value for the Backoff Calculator parameter is No Backoff.
Backoff Step	The Backoff Step parameter specifies the number of seconds by which the trap evaluation interval is increased. If a trap interval is 10 and its Backoff Step is 3590, the trap is evaluated every 10 seconds until it fires the first time and sends data, and is evaluated once an hour thereafter. When the trap is evaluated at the backoff step interval, the trap is fired only if the variable's value still exceeds the threshold. If the value no longer exceeds the threshold, the trap does not fire, and the trap returns to evaluating the variable with the frequency set in the Interval parameter. The default value for the Backoff Step parameter is 0.
Backoff Multiplier	The Backoff Multiplier parameter specifies the value by which to multiply a trap's evaluation interval each time the trap fires. If you set the Backoff Multiplier to Exponential Backoff and the Backoff Multiplier to 2, the evaluation interval doubles each time the trap fires. The default value for the Backoff Multiplier parameter is 1.

Parameter	Description
Priority	The Priority parameter specifies the priority of a trap. Also, if several traps are scheduled to fire at the same time, you can use the Priority parameter to decide which trap is serviced first. The possible values for the Priority parameter, in decreasing priority, are as follows: ◆ alert ◆ critical ◆ error ◆ warning ◆ notification ◆ informational ◆ debug The default value for the Priority parameter is notification.
Message	The Message parameter specifies a message that is sent out with a trap. The message could be a string of text or of the form <code>snmp.oid</code>. If an OID is specified, the result of evaluating that OID is sent. If the message parameter is not specified, the default value sent when a trap fires is a string that shows the OID value that triggered the trap. For example, the following string is sent to the network management stations for the total number of HTTP requests since last reboot: “Number of HTTP requests = <i><value></i>.” Note _____ If the message is a space-separated string, you must enclose the string in quotation marks (“”). _____

About this chapter	This chapter describes the Autosupport feature, which you can use to prevent potential problems from becoming problems.
Topics in this chapter	This chapter provides information about the following topics: ◆ “Planning summary for Autosupport configuration” on page 108 ◆ “About Autosupport” on page 110
Procedures and option descriptions	See the Appliance Manager online Help for configuration procedures and option descriptions.
Location of configuration options	Options for configuring Autosupport are in the following location in the Appliance Manager: Setup tab > Administration > Autosupport

Planning summary for Autosupport configuration

Determine Autosupport configuration

Configuring Autosupport will be faster if you plan ahead. Use the information in the following table to determine the Autosupport features that are appropriate for your organization.

Decisions to make	Your options	Considerations
Use Autosupport	By default, Autosupport sends messages to Network Appliance over HTTPS, with minimal detail (no private data). You can change the default settings.	◆ You can be alerted to potential problems before they become actual problems. ◆ Support for Autosupport is included as part of your basic warranty service.
Recipients of Autosupport messages	One or both of the following: ◆ Network Appliance Technical Support. ◆ Non-Network Appliance recipients. You can specify up to five e-mail addresses for event alerts and up to five addresses to receive both event messages and weekly reports.	Network Appliance looks for potential problems and has the experience to help you solve problems quickly.
Message types you want Autosupport to send	◆ If Network Appliance Technical Support is a recipient, event messages and weekly reports are sent. ◆ For non-Network Appliance recipients, you can separately identify recipients to receive event alerts and recipients to receive both event messages and weekly reports.	Sending event alerts to employees with pagers enables them to respond more quickly to problems when they are out of the office.

Decisions to make	Your options	Considerations
Level of detail in Autosupport messages (for messages to Network Appliance only)	◆ Include private information, such as IP addresses, in Autosupport messages. ◆ Do not include private information in Autosupport messages.	If you are concerned about revealing data that you consider to be sensitive external to your company, specify that the messages include a minimal level of detail.
How Autosupport messages can be sent	◆ For messages to Network Appliance, the transport can be HTTP Post, HTTPS Post, or SMTP e-mail. ◆ For non-Network Appliance recipients, the transport must be SMTP e-mail.	◆ HTTPS Post provides the highest level of security. Firewalls might block transport, however. ◆ Typically, messages sent by HTTP Post can bypass firewalls. ◆ SMTP e-mail requires a mail server. If a firewall blocks the mail server, use HTTP Post.
If SMTP e-mail is the transport for Autosupport messages	◆ You can choose the transport for messages to Network Appliance. ◆ The transport for non-Network Appliance recipients is always SMTP e-mail.	You must have at least one host at your site that listens on the SMTP port (25) to send mail—an SMTP server or a mail forwarder, such as the sendmail program or Microsoft Exchange server. This host must be reachable by the appliance. You can specify up to five e-mail hosts.

About Autosupport

Autosupport described

The Autosupport feature sends messages automatically to Network Appliance™ Technical Support, members of your organization, or both, depending on how you configure Autosupport. Autosupport messages alert recipients about problems and potential problems with the appliance. Network Appliance strongly encourages you to use the Autosupport feature.

Reasons for using Autosupport: Network Appliance recommends that you include Network Appliance Technical Support as a recipient of your Autosupport messages for the following reasons:

- ◆ Network Appliance has programmed scripts to automatically look for particular data in Autosupport weekly reports that might indicate a potential problem. Some Network Appliance customers have been surprised when they were alerted by Network Appliance Technical Support about potential system problems—before the customers were aware that their appliances were not functioning optimally.
- ◆ Network Appliance will help you solve problems that Autosupport detects. Network Appliance provides a range of support, including automated parts replacement, e-mail contact, and contact by a Technical Support engineer, depending on the type of problem. For example, if Network Appliance receives a message that a disk failure occurred on your appliance, a replacement disk is automatically sent to you.

Protection of private data: Revealing private data outside your company does not need to be a concern. You can set up Autosupport so that private data is not included in the messages sent to Network Appliance. Additionally, you have a choice of transports, including HTTPS, over which Autosupport messages to Network Appliance can be sent.

Types of Autosupport messages

Autosupport generates the types of messages shown in the following table. Details about the contents of Autosupport messages appear later in the chapter.

Type	Description
Event alert	A brief one-line message that Autosupport sends to recipients you have identified when an event tracked by Autosupport occurs. Event alerts are commonly sent to personnel with devices such as pagers who want to be notified of problems when they are in the field. Event alerts and event messages are sent at the same time.
Event message	A message that Autosupport sends to recipients you have identified when an event tracked by Autosupport occurs. The subject line is the same one-line description that is sent as an event alert. The message contains information to help you diagnose and troubleshoot the problem. Event alerts and event messages are sent at the same time.
Weekly report	A message that Autosupport sends automatically each week to recipients you have identified. Unlike event alerts and event messages, weekly reports are not triggered by a particular event.

Events that trigger Autosupport messages

If Autosupport is enabled, Autosupport generates an event alert and an event message when an event in the following table occurs.

Event	Subject line of the message
Low NVRAM lithium battery	BATTERY_LOW!!!
Disk failure	DISK_FAIL!!!
Disk failed reconstruction	DISK_RECONSTRUCTION FAILED!!!
Out-of-date disk firmware detected	DISK_FIRMWARE_NEEDED_UPDATE!!!

Event	Subject line of the message
Fan failure	FAN_FAIL!!!
Shutdown because of overheating	OVER_TEMPERATURE_SHUTDOWN!!!
Partial RPS failure (One of the two power supplies on an appliance using a redundant power supply (RPS) has ceased functioning.)	POWER_SUPPLY_DEGRADED!!!
System reboot	REBOOT
System reboot using the <code>reboot</code> command	REBOOT (reboot command)
System reboot using the <code>halt</code> command	REBOOT (halt command)
System reboot caused by an internal halt	REBOOT (internal halt)
System reboot caused by an internal reboot	REBOOT (internal reboot)
System reboot caused by a watchdog timer reset	REBOOT (watchdog reset)
System reboot caused by turning on power	REBOOT (power on)
System reboot following a panic	REBOOT (panic)
System reboot after loading firmware	REBOOT (after entering firmware)
System reboot after running a WAFL™ check	REBOOT (after WAFL_check)
System reboot for an unknown reason	REBOOT (unknown reason)

Event	Subject line of the message
System reboot caused by a power glitch, for example, a power interruption	REBOOT (power glitch)
wack command was run	WAFL_check!!!
Fibre channel arbitrated loop is unreliable	FC-AL UNRELIABLE!!!
Fibre channel arbitrated loop has link failures	FC-AL LINK_FAILURE!!!
Storage health monitor detected degraded disk I/O	DISK_IO_DEGRADED!!!
Storage health monitor detected a SCSI sense error	DISK_RECOVERED_ERRORS!!!
Storage health monitor detected low disk throughput	DISK_LOW_THRUPUT!!!
Storage health monitor detected an impending disk failure	DISK_FAILURE_PREDICTED!!!
Incorrect system hardware configuration detected (PCI card only)	SYSTEM_CONFIGURATION_ERROR
Fibre channel disk shelf configuration error	CONFIGURATION_ERROR!!!
Disk shelf error	SHELF_FAULT!!!
Spare disk failure	SPARE FAIL!!!
Weekly backup of /etc/messages	WEEKLY_LOG
option autosupport.doit command	USER_TRIGGERED (user supplied subject)
Log file push failure	LOG_PUSHING_FAILED

Event	Subject line of the message
Minicore sent to Network Appliance	Subject determined by Autosupport

Content of Autosupport event messages and weekly reports

The following table shows the type of content that can be included in event messages and weekly reports. The content that is actually included depends on whether you specify Complete or Minimal as the detail level for event messages and weekly reports.

“Filtered” in the Minimal column means that sensitive data is removed. For example, only nonsensitive content from the Messages log is sent if the level of detail is specified as Minimal. Regardless of content type, IP addresses, MAC addresses, and host names are always removed if the level of detail is specified as Minimal.

Content	Complete detail	Minimal detail
Date and time stamp of the message	X	X
Software licenses enabled on the appliance	X	X
Service usage (availability statistics)	X	X
System serial number (if available)	X	X
Appliance software version	X	X
Appliance system ID	X	X
Appliance host name	X	—
SNMP contact name and location (if specified in the SNMP page)	X	—
Messages log contents	X	Filtered
ECC memory scrubber statistics	X	X
Registry statistics	X	X
Appliance statistics	X	—

Content	Complete detail	Minimal detail
Appliance configuration changes from default	X	—
Volume status	X	X
Checksum status	X	X
Parity Inconsistency (PI) information	X	X
WAFL check logs	X	Filtered
Hit rates	X	—
Utilization statistics	X	Filtered
WAFL statistics	X	X
NFS statistics	X	X
Output from the following console commands:		
sysconfig -a sysconfig -c sysconfig -d sysconfig -r	X	X (-a filtered)
options	X	Filtered
ifconfig -a	X	—
ifstat -a	X	Filtered
df df -i	X	X
vol status vol status -v vol status -c	X	X

Content	Complete detail	Minimal detail
fcstat fcal_stats fcstat device_map fcstat link_stats	X	X
/etc/messages	X	Filtered
perf report -t	X	X

About this chapter

This chapter describes the appliance management utilities and common administrative tasks.

Topics in this chapter

This chapter provides information about the following topics:

- ◆ [“Halting, rebooting, and powering down the appliance”](#) on page 118
- ◆ [“Backing up and restoring the appliance configuration”](#) on page 119
- ◆ [“Messages log”](#) on page 121

Procedures and option descriptions

See the Appliance Manager online Help for configuration procedures and option descriptions relating to these topics.

Location of configuration options

Features discussed in this chapter are in the following locations in the Appliance Manager:

Feature	Location
Halt/Reboot	Maintenance tab > System Control
Back Up/Restore	Maintenance tab > Configuration Files
Messages log	Data tab > Log Files

Halting, rebooting, and powering down the appliance

When to use halt and reboot

The Appliance Manager provides a way to halt and reboot the appliance cleanly.

When to use halt: Use the Halt Appliance option if you want to stop the appliance from serving requests and do not want it to restart itself. This method is useful if you want to power down the appliance or if you want to boot the system from the system boot diskette.

When to use reboot: Use the Reboot Appliance option if you want to halt the appliance and reboot it in one step.

Powering down the appliance

Always halt the appliance before powering down the system. Perform the Halt Appliance operation on the Halt / Reboot page (Maintenance tab > System Control > Halt / Reboot) or run the `halt` command from the CLI.

Caution

Failure to properly shut down the appliance can result in data loss.

Unplugging the power cable from the appliance

If you need to unplug the appliance power cable for any reason, for example, you need to move the appliance, make sure to halt the system and ensure that the LCD displays `Halted` before you unplug the cable.

Which disk the system boots from

If the appliance has a boot diskette in its diskette drive, it boots from the diskette; otherwise, it boots from the hard disk.

Backing up and restoring the appliance configuration

Configuration backups help during disaster recovery

Although data saved in the cache is transitory and does not need to be backed up, the Appliance Manager enables you to save one or more backup copies of the system configuration to be restored in case of a configuration error or disaster. You provide the name of the backup image where the configuration is saved.

You can also download a backup configuration file to a host machine, then upload the backup configuration file to the appliance later so it can be restored. This enables you to save the backup off site for disaster recovery.

About system configuration files

Two appliance system configuration files exist by default:

- ◆ `default.cfg`—Contains all the default appliance configuration settings. This file is automatically generated by the system when the appliance is installed. This file must never be edited.
- ◆ `local.cfg`—Contains all the configuration settings that differ from the default settings listed in the `default.cfg` file. This file records all local changes from default that are made to the appliance during installation and configuration (through both the Appliance Manager and CLI).
- ◆ You can also create one or more arbitrarily named global configuration file containing appliance settings to be uniformly applied to other appliances. The administrator can then configure those appliances to periodically read the settings in the global file and pulled into the local configuration.

The appliance backs up the `local.cfg` and global configuration files into a single backup file. The `default.cfg` file is never modified and does not need to be backed up or restored.

Backup configuration file names

By default, the backup configuration files are named `backup.date`, where *date* takes the format *mmddyy*. For example, `backup.100200` is the appliance configuration file saved on October 2, 2000.

Using configuration backup files to restore the configuration

You can use the Appliance Manager to restore the system configuration files in the following ways:

- ◆ Restore the configuration locally

Selecting this option restores the local configuration and any global configuration files to the local appliance. This option does not restore global configurations to their original source.

- ◆ Restore the configuration independently

Selecting this option restores the `local.cfg` file locally and writes all the settings contained in any global configuration file into the `local.cfg` file. This makes all global settings part of the local configuration. These settings override any global configuration file changes.

Messages log

About the Messages log

The appliance provides a Messages log file to help you monitor events, such as reboots, device failures, and errors. Critical system events are also reported to the console. Use this log when you are looking for appliance configuration or runtime errors.

Messages log file management

The Messages log is always enabled and cannot be disabled. The Messages log normally uses only a small amount of memory and generating it should not affect normal appliance operations. The Messages log is not configurable. The appliance always rotates the Messages log. A new Messages log is started every Sunday at midnight. At this time, the recently replaced *messages* file is copied to *messages.0*, the *messages.0* file is copied to *messages.1*, and so on. The oldest Messages log is deleted.

Location of log file management options

Select Data tab > Log Files > *logfilename* in the Appliance Manager to access the log maintenance pages. From these pages, you can

- ◆ View an entire log file or the last lines of a log file
- ◆ Copy log files to your computer
- ◆ Delete one or more log files
- ◆ Send log files to a remote FTP or HTTP Web server

Index

A

- access to the appliance
 - access control lists 25
 - adding administrative accounts 25
 - administering 25
- ACLs
 - actions 31
 - described 31
 - expressions 35
 - operators 36
 - syntax 31
 - variables 32
- admin account, default 26
- Appliance Manager, accessing 26
- authentication, defined 27
- Autosupport
 - complete content 114
 - defined 110
 - event alerts 111
 - features 108
 - messages 110
 - minimal content 114

B

- backing up, configuration 119
- bandwidth allocation
 - defined 90
 - pipes 90
 - rules 91
- boot diskette
 - creating 83
 - reinitializing 83
- booting
 - reinitializing diskettes 83

C

- cache hit 4
- cache hit, illustrated 4
- cache miss
 - complete miss 4

- illustrated 4
 - sparse miss 4
 - verify miss 4
- caching
 - data consistency 8
 - granularity 8
 - invalidation 9
- CLI commands 1
- commands
 - CLI 1
 - statistics 20
 - status 20
 - syntax 14
 - sysconfig -v 82
- configuration files, system 119
- creating a boot diskette 83

D

- data disks 69
- database, local user 27
- deployment
 - forward proxy 10
 - reverse proxy 10
- DHCP 60
- disk replacement, errors during 82
- disks
 - adding, removing 80
 - data 69
 - failure 75
 - fibre channel 72, 74
 - hot spare 69
 - hot swap 70
 - hung 82
 - identifiers 71
 - maintaining 77
 - parity 69
 - removing 80
 - SCSI 72
 - status 79
 - usable and physical space 73
- distributed NFS, defined 2

DNS

- failover 60
- lookups 61, 65
- nameservers 60
- search paths 64
- support for 60
- tree-search 62

domain name

- not fully qualified 61
- qualified 61

Domain Name System (DNS) 60

Dynamic Host Configuration Protocol (DHCP) 60

E

error messages

- during disk replacement 82

export

- configuring 13
- configuring, example 19
- mapping volumes 13
- settings 13
- volume 13

F

file handles 7

forward proxy deployment 10

G

grouping administrators 27

groups 27

H

halting the appliance 118

hot spare disk 69

hot swap disks 70

hung disks during disk replacement 82

I

ICMP 50

interfaces

- configuring 38

load balancing 53

multi-mode virtual 44

naming 39

outgoing, disabling 55

physical 39

single-mode virtual 43

supported 39

troubleshooting 44

virtual 39, 43

Internet Control Message Protocol (ICMP) 50

Internet Router Discovery Protocol (IRDP) 50

L

large windows, TCP 56

load balancing interfaces 53

logging 28

M

maximum age, objects 5

Messages log file, defined 121

N

Network Information Service. *See* NIS

network interfaces

configuring 38

multi-mode virtual 44

naming 39

physical 39

single-mode virtual 43

supported 39

troubleshooting 44

virtual 39, 43

NFS operations

ACCESS 8

GETATTR 7

LOOKUP 8

READ 7

REaddir 7

WRITE 7

NIS 86

configuring 86, 87

defined 86

P

- parity disk 69
- powering down the system 118

R

- RAID, defined 69
- rebooting
 - from diskette 83
- rebooting the appliance 118
- Redundant Array of Independent Disks (RAID) 67
- reinitializing disks 83
- restoring, configuration 119
- reverse proxy, deployment 10
- RIP 50
- route discovery 50
- Routing Information Protocol (RIP) 50
- rules, types of 30

S

- SNMP
 - agent and supported MIBs 94
 - defined 94

- limits 95

MIB

- defined 96
 - installing 96
 - OID 99

- terms 94

traps

- defined 98
 - parameters 99
 - user-defined 99

- sysconfig -v command 82

- system configuration files 119

T

- TCP large windows 56

- trunking 41

V

- validation 5

virtual interfaces

- multi-mode 43
 - single-mode 43

